

NIPQ: Noise proxy-based Integrated Pseudo-Quantization

Juncheol Shin^{*2}, Junhyuk So^{*1}, Sein Park², Seungyeop Kang³, Sungjoo Yoo³, and Eunhyeok Park^{1,2}

¹Department of Computer Science and Engineering, POSTECH

²Graduate School of Artificial Intelligence, POSTECH

³Department of Computer Science and Engineering, Seoul National University

{jchshin, junhyukso, seinpark, eh.park}@postech.ac.kr, {ksy8653, yeonbin}@snu.ac.kr

Abstract

Straight-through estimator (STE), which enables the gradient flow over the non-differentiable function via approximation, has been favored in studies related to quantization-aware training (QAT). However, STE incurs unstable convergence during QAT, resulting in notable quality degradation in low precision. Recently, pseudo-quantization training has been proposed as an alternative approach to updating the learnable parameters using the pseudo-quantization noise instead of STE. In this study, we propose a novel noise proxy-based integrated pseudo-quantization (NIPQ) that enables unified support of pseudo-quantization for both activation and weight by integrating the idea of truncation on the pseudo-quantization framework. NIPQ updates all of the quantization parameters (e.g., bit-width and truncation boundary) as well as the network parameters via gradient descent without STE instability. According to our extensive experiments, NIPQ outperforms existing quantization algorithms in various vision and language applications by a large margin.

1. Introduction

Neural network quantization is a representative optimization technique that reduces the memory footprint by storing the activation and weight in a low-precision domain. In addition, when hardware acceleration is available (e.g., low-precision arithmetics [30, 41, 43, 48] or bit-serial operations [14, 18, 42]), it also brings a substantial performance boost. These advantages make network inference affordable in large-scale servers as well as embedded devices [2, 32], which has helped popularize it in various applications. However, quantization has a critical disadvantage, quality degradation due to limited degrees of freedom.

^{*}Equal contribution.

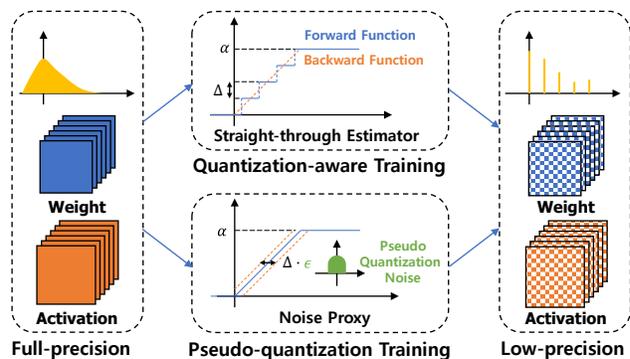


Figure 1. The proposed NIPQ as an alternative to QAT with STE.

The way to train the networks accurately within limited precision is critical and receiving much attention these days.

To mitigate the accuracy degradation, quantization-aware training (QAT) has emerged that trains a neural network with quantization operators to adapt to low precision. While the quantization operator is not differentiable, the straight-through estimator (STE) [5] allows the backpropagation of the quantized data based on linear approximation [19, 51]. This approximation works well in redundant networks with moderate precision (>4-bit). Thus, not only early studies [9, 34, 51] but also advanced ones [8, 16, 46] have proposed diverse QAT schemes based on STE and shown that popular neural networks (i.e., ResNet-18 [20]) can be quantized into 4-bit without accuracy loss.

However, STE-based QAT bypasses the approximated gradient, not the true gradient, and many studies have pointed out that it can incur instability and bias during training [19, 28, 29, 31, 35, 39]. For instance, PROFIT [31] points out that the instability is a major source of accuracy drop for the optimized networks (e.g., MobileNet-v2/v3 [22, 36]). More recently, an alternative training scheme, pseudo-quantization training (PQT) based on pseudo-quantization

noise (PQN), has been proposed [3, 10, 37] to address the instability of STE-based QAT. During PQT, the behavior of the quantization operator is simulated via PQN, and the learnable parameters are updated based on the proxy of quantization. While those studies are applied only to the weight, they can stabilize the training process significantly compared to QAT with STE and show the potential of PQT.

Nevertheless, the existing PQT algorithms have room for improvement. Various STE-based studies [8, 16, 51] have shown that truncation contributes significantly to reducing quantization errors, but even the advanced PQT studies [10, 37] use a naive min-max quantization. Integrating the truncation on the PQT framework can greatly reduce quantization error and enable us to exploit a PQT scheme for activation, which has an input-dependent distribution and requires a static quantization range. In addition, there is no theoretical support for whether PQT guarantees the optimal convergence of the quantization parameters. Intuitive interpretation exists, but proof of whether quantization parameters are optimized after PQT has yet to be provided.

In this paper, we propose a novel PQT method (Figure 1) called Noise proxy-based Integrated Pseudo-Quantization (NIPQ) that quantizes all activation and weight based on PQN. We present a novel idea, called Noise proxy, that shares the same quantization hyper-parameters (e.g., truncation boundary and bit-width) with the existing STE-based algorithm LSQ(+) [6, 16]. However, noise proxy allows to update the quantization parameters, as well as the network parameters, via gradient descent with PQN instead of STE. Subsequently, an arbitrary network can be optimized in a mixed-precision representation without STE instability. Our key contributions are summarized as follows:

- NIPQ is the first PQT that integrates truncation in addition to discretization. This extension not only further reduces the weight quantization error but also enables PQT for activation quantization.
- NIPQ optimizes an arbitrary network into the mixed-precision with awareness of the given resource constraint without human intervention.
- We provide theoretical analysis showing that NIPQ updates the quantization hyperparameters toward minimizing the quantization error.
- We provide extensive experimental results to validate the utility of NIPQ. It outperforms all existing mixed-precision quantization schemes by a large margin.

2. Related Work

Due to its practical usefulness, various studies have been proposed for multi-bit linear quantization [8, 16, 24, 51]. Most of these studies are based on STE, and the proposed

schemes have evolved by designing the quantization function to enable the optimization of the quantization parameters via gradient descent. However, in optimized networks such as MobileNet-v2, accuracy loss induced by STE instability has been reported for both activation [19] and weight [31]. They addressed the instability via a newly designed pipeline and non-linear approximation but suffered from the increased complexity and cost of QAT. In this work, NIPQ updates the quantization parameters without STE approximation, enabling stable convergence without additional cost or complexity and, most importantly, outperforming the existing methods by a large margin.

Mixed-precision studies focus on allocating layer-wise or group-wise bit-width in consideration of the precision sensitivity of each layer to minimize accuracy drop within the given resource constraints. Various methods have been proposed, e.g., RL-based [15, 27, 45], Hessian-based [12, 13, 49], and differentiable [44, 46] algorithms, but RL-based and Hessian-based methods are relatively complex, requiring a lot of parameter tuning, and differentiable algorithms still suffer from STE approximation error. In the present work, we reinterpret the differentiable bit-width tuning in terms of PQT framework, enabling the layer-wise bit-width tuning via gradient descent without STE instability.

Robust quantization [1, 7, 40] aims to guide the convergence of the network toward a smooth and flat loss surface based on additional regularization. The robustness of neural networks is highly beneficial for deploying noisy devices or low-precision ALUs. In the case of noise proxy, it inherently improves the robustness during PQT with PQN. As far as we know, we observe for the first time that the robustness of activation quantization can be enhanced (Section 7.2).

The benefit of the PQT has been demonstrated in diverse studies [10, 37] in various perspectives. However, previous studies utilize PQT with PQN in a limited domain only for weight and have yet to provide theoretical support for the convergence of PQT corresponding to the quantization. We extend the idea of PQT to both activation and weight by integrating the idea of truncation in addition to discretization and prove that the optimization of quantization parameters through noise proxy minimizes the actual quantization error. In addition, our work is the first to demonstrate that the PQT-based pipeline outperforms STE-based mixed-precision quantization algorithms by a large margin.

3. Simple Example Problem

Before introducing the details, we first propose a simple problem for straightforward explanation. The objective is to minimize the difference from the target data $t \in \mathbb{R}^N$ and the quantized value of the learnable parameters $x \in \mathbb{R}^N$.

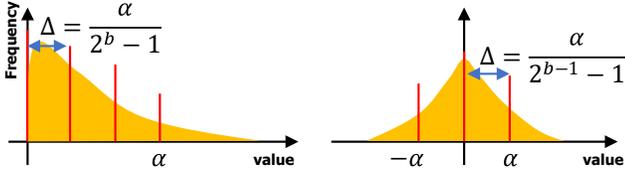


Figure 2. Visualization of the 2-bit quantization with hyper-parameters (e.g., α for truncation boundary and b for bit-width) for non-negative distribution (left) and symmetric distribution (right)

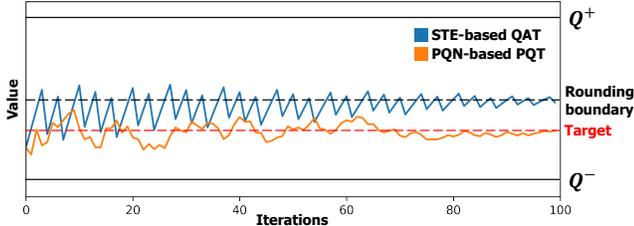


Figure 3. The trajectory of a scalar input during the minimization of \mathcal{L}^Q through SGD with STE-based QAT and PQN-based PQT. Q^+ represents the positive quantization level and Q^- the negative.

The example loss function is defined as follows:

$$\mathcal{L}_{exp}^Q(x, \alpha, b) = \sum_{i=0}^{N-1} \left(t_i - Q(x_i | \alpha, b) \right)^2, \quad (1)$$

where subscript i means the i -th element of the data, t is arbitrary distribution, and $Q(\cdot)$ is a given quantization function, parameterized α (truncation boundary) and b (bit-width). To minimize the loss, we need to update x , α , and b judiciously. We utilize this example for the rest of the paper.

4. Motivation

In this work, we focus on linear quantization, especially for a more specific case that has an edge on hardware acceleration (e.g., affine layer-wise quantization for activation and symmetry channel-wise quantization for weight) [43]. To get optimal quality within restricted resources, we need to tune the quantization parameters, shown in Figure 2. In this section, we analyze the limitations of existing STE-based and PQT-based algorithms and provide insight for next-step innovation.

4.1. Limitation of STE-based Quantization

STE [5] allows the gradient flow over the quantized data by approximating the discretization as linear identity mapping. STE-based QAT updates the quantization parameters via gradient descent, which enables the joint update of network and quantization parameters toward minimizing the target loss value. Many prior works show successful results relying on this approximation [8, 16, 24, 28, 31, 46, 51].

However, STE-based QAT has a critical limitation. As shown in Figure 3, the scalar value never converges to the

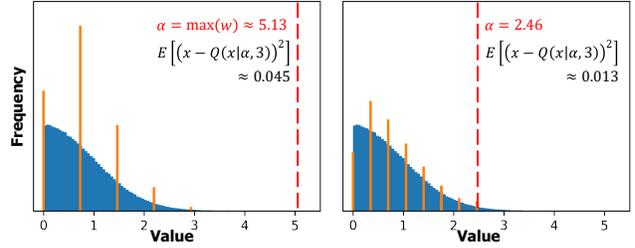


Figure 4. The comparison of 3-bit min-max quantization (left) and 3-bit quantization with truncation (right).

target value; instead, it keeps oscillating near the rounding boundary in the middle of the two nearest quantization levels. This is because when the scalar value is updated to cross the rounding boundary, it is mapped to a different quantization level. Then the sign of the gradient is reversed, updating the scalar value toward the opposite direction. The same process is repeated during training. Even at the later iterations, when the learning rate gradually decreases, the scalar value still oscillates near the rounding boundary.

The frequent flipping of the quantization value induces instability in network convergence. Moreover, the oscillation near the rounding boundary becomes the major source of large quantization errors. Even at the end of the training, $Q(x | \alpha, b)$ can be mapped to either Q^+ and Q^- , while Q^- has a minimal difference from the target value. Many studies have pointed out that this phenomenon degrades the quality of the quantized network, especially for the optimized networks [29, 31, 35].

4.2. Pros and Cons of Previous PQN-based PQT

To tackle the limitation of STE-based QAT, an alternative quantization pipeline, PQN-based PQT, has been proposed [10, 37]. During PQT, no quantization operator is used for forward and backward operation, but the behavior of the quantization operator is simulated via PQN. The learnable parameters are updated in a differentiable way without STE instability based on the proxy of quantization, becoming robust in the low-precision domain. As shown in Figure 3, the scalar value with PQT is updated stably toward the target value with small oscillation. Especially as the learning rate gets decreased, oscillation amplitude becomes reduced, eventually converging to the target value. After the PQT, the scalar value has been firmly mapped to Q^- , leading to minimal quantization error. Compared to STE-based QAT, PQN-based PQT is expected to have lower quantization error after the training.

However, existing PQT-related studies still have critical limitations. The advanced methods [10, 37] utilize min-max quantization, which uses the minimum/maximum value of data for the quantization range. Many STE-based studies point out that truncation is essential to minimize the quanti-

zation error in the low-precision domain [8, 16, 24, 51]. As shown in Figure 4, applying quantization based on the minimum/maximum value of the data wastes invaluable quantization levels, resulting in higher quantization error. As a long tail or outlier pushes the boundary by a large margin, truncation reduces overall quantization error significantly, especially for near-zero values. In addition, quantization with truncation is crucial for extending PQT for activation quantization; because the activation is input-dependent, the static boundary needs to be trained over a training dataset for stable and fast inference. Unfortunately, the existing studies do not provide the theoretical integration of truncation on top of a PQN-based PQT framework.

5. NIPQ: Noise Proxy-based Integrated Pseudo-Quantization

In this work, we propose a novel PQT pipeline to compensate for the quality degradation after quantization without STE instability. Moreover, the proposed scheme includes the support of updating truncation, resulting in accurate quantization not only for weight but also for activation. Besides, we adopt the gradient-based bit-width optimization via continuous relaxation [10, 46]. This extension allows us to update all of the hyper-parameters for linear quantization via gradient descent, which greatly simplifies the layer-wise optimization process.

In this section, we define the proposed idea and present the analytical support for the optimal convergence of quantization parameters based on PQT. Our method is composed of a pair of algorithms: STE-based quantization $Q(x|\alpha, b)$ [6, 16] and the corresponding noise proxy $\tilde{Q}(x|\alpha, b)$. The algorithms are expressed as follows:

$$Q(x|\alpha, b) = \begin{cases} 0, & x \leq 0 \\ \lfloor x/\Delta \rfloor \cdot \Delta, & 0 < x < \alpha \\ \alpha, & x \geq \alpha, \end{cases} \quad (2)$$

$$\tilde{Q}(x|\alpha, b) = \begin{cases} 0, & x \leq 0 \\ x + \epsilon \cdot \Delta, & 0 < x < \alpha \\ \alpha, & x \geq \alpha, \end{cases} \quad (3)$$

where $\Delta = \frac{\alpha}{2^b - 1}$ represents the quantization step size or the difference between quantization levels, ϵ represents the PQN, and $\lfloor \cdot \rfloor$ means the rounding operation¹.

During PQT, $\tilde{Q}(x|\alpha, b)$ is applied on top of the quantization targets, activation and weight, and the learnable parameters are updated via gradient descent. Because α and b are included in the gradient path, the training with noise proxy updates all of the quantization parameters as well as the network's learnable parameters jointly. If we set the target loss

¹Please note that we present the quantization algorithm for non-negative distribution data, e.g., activation after ReLU, for simplicity. However, this approach holds the same for weight, except that quantization index is used in $[-2^{bit-1}, 2^{bit-1} - 1]$.

function, \mathcal{L}_{target} as the mixture of the task loss \mathcal{L}_{task} , and the cost loss \mathcal{L}_{cost} , for memory footprint or computation as follows:

$$\mathcal{L}_{target} = \mathcal{L}_{task} + \lambda \cdot \mathcal{L}_{cost}, \quad (4)$$

where λ is a hyper-parameter of balancing two loss functions. The layer-wise bit-width b is assigned within the available resource budget, and the corresponding α is tuned appropriately to achieve the best quality possible.

After PQT, $Q(x|\alpha, b)$ is used for the true quantization during inference. To achieve the high-quality output with the true quantization operator, the tuned α and b through PQT with noise proxy should also have the lowest loss with $Q(x|\alpha, b)$. In the following subsections, we provide theoretical support that the optimization of α and b through PQT guarantees the optimal convergence $Q(x|\alpha, b)$ when we sample PQN judiciously.

5.1. Input Activation Convergence

First of all, we evaluate the convergence property of the learnable input tensor x . In STE-based algorithm, the gradient of input data is bypassed through STE algorithm. Thereby the derivative is expressed as follows:

$$\frac{\partial Q(x|\alpha, b)}{\partial x} = \begin{cases} 0, & x \leq 0 \\ 1, & 0 < x < \alpha \\ 0, & x \geq \alpha. \end{cases} \quad (5)$$

In the case of the example in Eq. 1, the gradient of x_i within the quantization interval ($x_i \in [0, \alpha]$) is calculated as follows:

$$\frac{\partial \mathcal{L}_{exp}^Q(x, \alpha, b)}{\partial x_i} = 2(Q(x_i|\alpha, b) - t_i). \quad (6)$$

Therefore, x_i converges to t_i theoretically as the training progresses. However, as shown in Figure 3, x_i in practice fluctuates near the rounding boundary, except $Q(x_i|\alpha, b) = t_i$ exist, due to the instability of STE-based QAT.

On the other hand, the gradient of input data regarding noise proxy is expressed as follows:

$$\frac{\partial \tilde{Q}(x|\alpha, b)}{\partial x} = \begin{cases} 0, & x \leq 0 \\ 1, & 0 < x < \alpha \\ 0, & x \geq \alpha, \end{cases} \quad (7)$$

which is exactly the same as the $Q(\cdot)$ condition. When we update the example loss function via the proposed algorithm, the gradient of x_i within the quantization interval is evaluated as follows:

$$\frac{\partial \mathcal{L}_{exp}^{\tilde{Q}}(x, \alpha, b)}{\partial x_i} = 2(\tilde{Q}(x_i|\alpha, b) - t_i) = 2(x_i + \epsilon \cdot \Delta - t_i). \quad (8)$$

The gradient of x_i points toward the target, t_i , but with drift and noise from PQN. However, the drift can be amortized when the average of PQN is zero-centered ($E[\epsilon] = 0$).

Note that the distribution of actual quantization errors is empirically known to have zero-mean [47]. In addition, because the update step size is proportional to the learning rate, the oscillation amplitude becomes smaller when the learning rate becomes smaller as learning progresses; eventually, x_i converges to the target value, as shown in Figure 3. Unlike the STE-based QAT, noise proxy-based PQT precisely evaluates the effect of oscillation induced by PQN and discards the bias of PQN gradually during training.

Besides, the convergence property of oscillation near the target value is greatly beneficial to enhance the robustness of network, as pointed out in several previous studies [7]. When some noise, δ , exists, the objective function can be approximated via Taylor expansion as follows:

$$E[\mathcal{L}(x + \delta)] \quad (9)$$

$$\approx E[\mathcal{L}(x) + \delta \cdot \nabla_x \mathcal{L}(x) + \frac{1}{2} \delta^T \cdot \nabla_x^2 \mathcal{L}(x) \cdot \delta] \quad (10)$$

$$\approx \mathcal{L}(x) + \frac{E[\delta^2]}{2} \text{Tr}\{\nabla_x^2 \mathcal{L}(x)\}, \quad (11)$$

where the term related to the first derivative is removed since $E[\delta] = 0$, and the off-diagonal elements of the second derivative term become 0 because it relates to the expectation of multiplication of two i.i.d. samples. When the loss is converged to the minima point, the sum of eigenvalues, $\nabla_x^2 \mathcal{L}(x)$, should have a non-negative value. To minimize the average loss after training, the loss surface is guided to be converged to the minima having a lower Hessian trace value [7], resulting in enhanced network robustness.

5.2. Alpha and Bit-width Convergence

In addition to the input data, the noise proxy should guarantee the convergence of the quantization parameters in the optimal point that minimizes the loss with $Q(x|\alpha, b)$. First, let's consider the gradient of α regarding quantization and noise proxy function.

$$\frac{\partial Q(x|\alpha, b)}{\partial \alpha} = \sum_{x_i \geq \alpha} 1 + \sum_{0 < x_i < \alpha} \frac{1}{N_{lv}} \lfloor \frac{N_{lv}}{\alpha} x \rfloor - \frac{x}{\alpha}, \quad (12)$$

$$\frac{\partial \tilde{Q}(x|\alpha, b)}{\partial \alpha} = \sum_{x_i \geq \alpha} 1 + \sum_{0 < x_i < \alpha} \epsilon \cdot \frac{1}{N_{lv}}, \quad (13)$$

where $N_{lv} = 2^b - 1$ denotes the number of levels.

Two gradients have a difference only for the case when $x \in (0, \alpha)$. However, note that Equations 12 and 13 become identical when PQN ϵ is sampled from the quantization noise distribution $\lfloor x/\Delta \rfloor - x/\Delta$. The same conclusion is observable for the number of quantization levels, which

is expressed as:

$$\frac{\partial Q(x|\alpha, b)}{\partial N_{lv}} = \sum_{0 < x_i < \alpha} \frac{x}{N_{lv}} - \frac{\alpha}{N_{lv}^2} \lfloor \frac{N_{lv}}{\alpha} x \rfloor, \quad (14)$$

$$\frac{\partial \tilde{Q}(x|\alpha, b)}{\partial N_{lv}} = \sum_{0 < x_i < \alpha} -\epsilon \frac{\alpha}{N_{lv}^2}. \quad (15)$$

Therefore, when we use the sampled noise following the quantization error distribution, the gradient of quantization parameters becomes identical to the true quantization operator while still enjoying the benefit of stochasticity of PQN.

In short, when we apply NIPQ whose noise is sampled from the quantization noise distribution, the gradient direction of α and b is identical in both $Q(x|\alpha, b)$ and $\tilde{Q}(x|\alpha, b)$. Therefore, when α converges to the optimal point via PQT where $\frac{\partial \mathcal{L}^{\tilde{Q}}}{\partial \alpha} = \sum_i \frac{\partial \mathcal{L}}{\partial \tilde{Q}_i(x_i|\alpha, b)} \cdot \frac{\partial \tilde{Q}_i(\cdot)}{\partial \alpha} = 0$, then we can easily validate that $\frac{\partial \mathcal{L}^Q}{\partial \alpha} = \sum_i \frac{\partial \mathcal{L}}{\partial Q_i(x_i|\alpha, b)} \cdot \frac{\partial Q_i(\cdot)}{\partial \alpha} = 0$. Likewise, when $\frac{\partial \mathcal{L}^{\tilde{Q}}}{\partial b} = 0$, then $\frac{\partial \mathcal{L}^Q}{\partial b} = 0$ and that is the point where the loss value with $Q(x|\alpha, b)$ is minimized. Therefore, optimizing the quantization parameters α and b based on NIPQ also minimizes the target loss with true quantization.

6. Additional Details

As explained above, NIPQ enables the compensation of quality degradation of the quantized network without STE approximation. However, in practice, we need to consider several details to maximize obtainable accuracy.

6.1. Stochastic Rounding for Bit-width Parameter

First, during PQT phase, we use the bit-width with stochastic rounding as follows:

$$b \leftarrow 2 + 14 \cdot \text{Sigmoid}(\beta), \quad (16)$$

$$b \leftarrow \lfloor b + U(-0.5, 0.5) \rfloor, \quad (17)$$

where $U(x, y)$ represents the uniform noise in $[x, y]$, and β is a trainable value in the continuous domain corresponding to the bit-width b . The bit-width is narrowed into a range of 2 to 16 bits, then mapped to the integer value via stochastic rounding. The idea of continuous relaxation of bit-width and updating it through gradient descent has been proposed in several previous studies [10, 46]. They have used the continuous value as it is, but we find that these settings yield suboptimal convergence because of the domain difference in that only discrete bit-width can be a candidate of precision during inference. To resolve the domain gap, we propose the bit-width update with stochastic rounding. It is an unbiased estimator but makes a decision in discrete space, whereby the domain difference could be mitigated. Our experiment shows that this update improves the final accuracy by a large margin, especially in the sub-4-bit domain. The experiments are included in the Supplementary Materials.

6.2. Stabilization in the Late Training Stage

In addition, we need to update the batch normalization layers after finishing NIPQ training (BN update). PQN in NIPQ induces instability in the running statistics of the normalization layers, so we update the batch normalization statistics with the true quantization operator after finishing training. We also observe additional performance improvement by switching from NIPQ to STE-based QAT at the later stage of training and updating the last few epochs with a very small learning rate (QAT finetune), while the amount of improvement is far smaller than BN update. We speculate that QAT finetune gives an opportunity to finetune the learnable parameters as well as normalization statistics of the network to mitigate the minor instability from noise. The details are included in the Supplementary Materials.

6.3. PQN Sampling

Finally, we provide the convergence condition in the last section in that PQN should be sampled following the quantization error distribution. However, we observe that NIPQ with a uniform noise $U(-\frac{\Delta}{2}, \frac{\Delta}{2})$ also shows the comparable result to the noise with the desired property². In this configuration, the relative frequency of the quantization error is ignored. This makes the quantization noise overestimated, and the low bit tends not to be selected, but a similar result is achievable by making λ of the cost function larger. In particular, the difference in final quality is negligible if the QAT finetune is applied. In practice, sampling of quantization error distribution slows down the training significantly. Uniform noise with QAT finetune can be considered as a practical approach of reducing PQN sampling overhead while exploiting the benefit of PQT.

7. Experimental Setup

In order to validate the performance of NIPQ, we conduct comprehensive studies to measure the output quality and observe diverse properties in various applications. The results in this paper are obtained based on PQN sampled from the quantization noise, unless otherwise specified explicitly. Likewise, QAT finetune is used at the last few epochs, otherwise specified explicitly. We implement our algorithm on PyTorch [33] library, and the code is available at our repository³. The details of target loss and training pipeline are available in the Supplementary Materials.

7.1. Sensitivity-aware Mixed-precision

Mixed-precision quantization aims at improving accuracy while minimizing the overall storage footprint and

²Unlike the previous studies [3, 10], PQN from normal distribution shows inferior performance in our algorithm.

³<https://github.com/ECOLab-POSTECH/NIPQ>

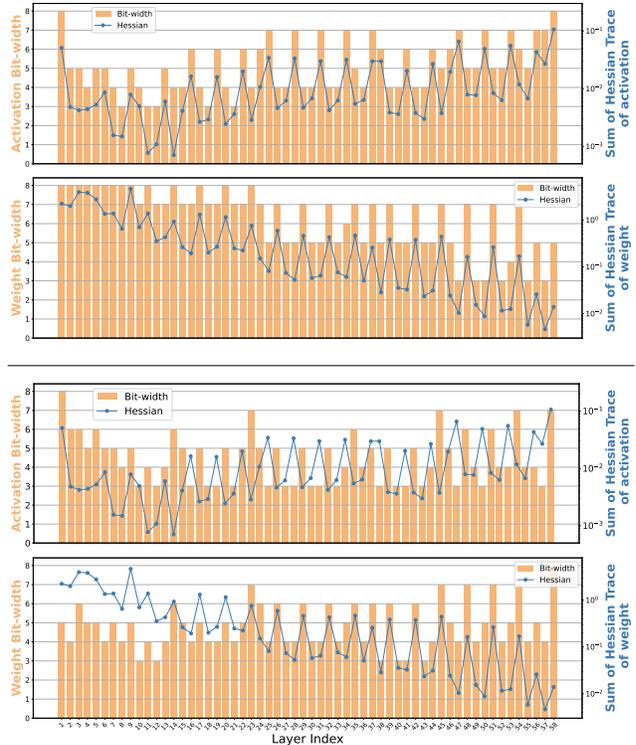


Figure 5. Layer-wise assigned bit-width of activation (top) and weight (bottom) and corresponding sensitivity measure (the sum of Hessian trace) of MobileNet-v2 on the CIFAR-100 dataset [25]. The average precision of activation weight is restricted to 4-bit (top) and the computation is restricted to 1.5 GBOPs (bottom).

computational cost. To maintain accuracy within the resource budget, more bit-width should be assigned to the sensitive layer to minimize overall quantization errors. Recently, the Hessian of the loss function has often been used as a metric of sensitivity against quantization [12, 13, 49]. The smaller the sum of the Hessian trace, the lower the sensitivity, and vice versa. Figure 5 (top) visualizes the assigned bit-width of activation and weight and the corresponding sum of the Hessian trace estimated via the Hutchinson algorithm [12] when applying NIPQ to MobileNet-v2 on the CIFAR-100 dataset. Each plot is obtained by optimizing the weights and activation independently. As shown in the figure, the more sensitive the layer is, the more bit-width is assigned. The redundant layers tend to have a fewer bit-width to match the target objective. Note that NIPQ does not have any additional stages that measure the sensitivity of the layer. Instead, we just train the network with an additional penalty term to restrict the average precision to 4-bit. The noise proxy algorithm allocates precision by itself, considering the sensitivity of the target layer, thereby quantizing the network with the highest accuracy as efficiently as possible within the constraints.

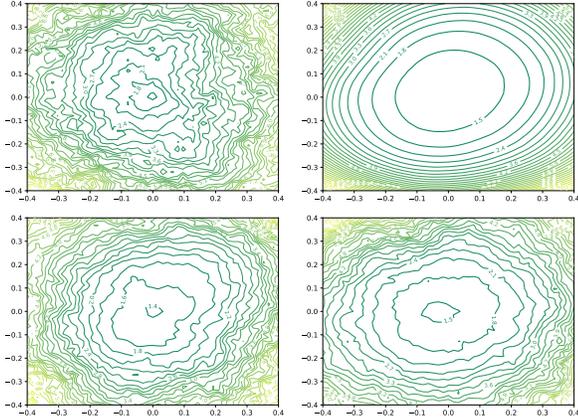


Figure 6. Loss landscapes [26] of 3-bit quantized MobileNet-V2 fine-tuned by STE-based LSQ [16] (top left), 3-bit NIPQ, BN update (top right), mixed-precision NIPQ (avg. 3-bit), BN update (bottom left), mixed-precision NIPQ (avg. 3-bit), QAT finetune (bottom right) on the CIFAR-100 dataset.

On the other hand, the automated bit-width allocation enables tuning of network considering complex metric that is practically improbable with the conventional Hessian-based methods [12, 13, 49]. Figure 5 (bottom) shows the assigned bit-width of activation and weight when restricting the computation cost (bit-operations, BOPs [4, 31]) as a 1.5G BOPs, which is equal to the computation cost of the quantized 4-bit model using PACT [8] or LSQ [16]. Unlike Figure 5 (top), the bit-width of activation is slightly misaligned with the sum of the hessian trace. In Figure 5 (bottom), we aim to optimize the average bit-width of activation and weight independently. Those two configurations are optimized via disjoint target losses thereby, each precision is assigned proportionally to the sensitivity of activation and weight separately. However, in the case of BOPs, the computation cost is proportional to the product of bit-width of activation and weight. Thereby, when we restrict the overall computation cost, the activation bit-width is allocated with awareness of the activation sensitivity as well as the corresponding weight sensitivity, and vice versa. This experimental result indicates that assigning the bit-width based on layer-wise sensitivity naively might not be an optimal policy for computation-aware quantization.

7.2. Robust Quantization for Practical Deployment

Another strength of NIPQ is the enhanced robustness of the network against unexpected noise. This brings diverse advantages to deploying the network in practice. For instance, the analog sum-based device [38] has significant energy efficiency but has inevitable noise induced by process variation or temperature drift, which causes instability of output. Even in the case of digital NPU, the low-precision implementation is fragmented because there are hundreds

	MNLI	CoLA	MRPC	QNLI	QQP	RTE	SST-2	STS-B
FP	84.46	56.53	90.62	90.66	87.76	71.19	92.32	88.64
DiffQ-MP-6 [10]	81.66	0.04	85.85	50.54	86.65	47.29	89.91	74.86
NIPQ-MP-6	83.87	59.60	89.68	90.92	87.48	68.95	91.97	88.11
DiffQ-MP-3 [10]	31.82	0.02	81.22	49.46	0	47.29	50.92	11.27
NIPQ-MP-3	81.88	53.03	88.62	90.23	86.68	69.31	91.51	88.42

Table 1. Results of weight quantization for BERT-base on GLUE

of hardware manufacturers [40]. When we deploy a neural network to the target device, the implementation difference can introduce unexpected noise on the data, resulting in accuracy degradation. The robustness of the network allows accuracy to be maintained in this environment, so securing this property is a significant advantage.

As explained in Section 5.1, NIPQ regularizes the loss surface. Figure 6 visualizes the sharpness of the loss surface by measuring the change of loss values after adding noise on top of the trained weight along the two random vectors [26]. As shown in the figure, NIPQ converges to a flat and smooth loss surface (top left vs. right). When layer-wise mixed precision is enabled, the loss surface becomes complex (top right vs. bottom left). We speculate that mixed-precision optimization exploits the sensitivity to trade robustness for better output quality. NIPQ enables such a trade-off in an automatic manner by forcing the network to adapt to the pseudo-noise that generalizes the quantization noise.

Besides, NIPQ enhances the robustness of the quantization parameters as well as the network parameters. Check the Supplementary Materials for the detailed results.

7.3. Importance of Truncation for Quantization

The last line of the related study is DiffQ [10], which is based on PQN-based PQT for weight quantization. The key difference from NIPQ is that DiffQ uses linear quantization based on the min-max value of weight instead of truncation. However, to minimize quantization errors with the limited bit-width, the presence of truncation is extremely helpful. In general, the data distribution of a natural network follows a bell-shaped distribution, where the majority of data is concentrated near zero. Because min-max quantization increases the quantization interval near zero, the quantization error greatly increases. Table 1 compares the accuracy after the quantization between DiffQ and NIPQ. Note that only the weight is quantized, and bit-width is allocated per layer. As shown in the table, NIPQ outperforms DiffQ by a large margin in the same bit-width. Integration of truncation on quantization minimizes the overall quantization error significantly, especially when the precision is limited.

7.4. Quantization of Large-scale Vision Tasks

To demonstrate the outstanding performance of the proposed method, we apply NIPQ for large-scale vision applications, including ImageNet [11] classification, multi-scale super-resolution, and VOC [17] object detection.

Model	Method	Bit(W/A)	BOPs(G)	Top-1
ResNet-18 [20]	FP	32/32	1857.6	70.54
	FP+KD	32/32	1857.6	72.17
	PACT [8]	4/4	34.7	69.2
	LSQ [16]	4/4	34.7	69.39
	DJPQ [46]	MP-BOPs	35.0	69.3
	HAQ [45]	MP-BOPs	34.4	69.2
	HAWQ [13]	MP-BOPs	34.0	68.5
	HAWQ-V3 [49]	MP-BOPs	34.0	68.5
	HAWQ-V3 [49]	MP-BOPs	72.0	70.2
	NIPQ-MP-4	3.89/3.98	45.5	69.84
	NIPQ-BOPs	4.53/3.67	34.1	69.47
	NIPQ-MP-4+KD	4.09/4.17	50.73	71.83
NIPQ-BOPs+KD	4.47/3.65	34.2	71.24	
MobileNet-v2 [36]	FP	32/32	306.8	72.6
	FP+KD	32/32	306.8	73.41
	DSQ [19]	4/4	15.8	64.8
	LSQ [16]	4/4	15.8	70.46
	DJPQ [46]	MP-BOPs	7.9	69.3
	HAQ [45]	MP-BOPs	8.3	69.5
	DuQ+KD [31]	4/4	5.3	69.86
	NIPQ-MP-4	3.79/4.00	8.67	71.52
	NIPQ-BOPs	6.19/5.31	8.25	72.26
	NIPQ-BOPs	5.19/4.24	5.37	70.92
	NIPQ-MP-4+KD	3.80/4.01	8.73	72.16
	NIPQ-BOPs+KD	6.21/5.35	8.31	72.94
NIPQ-BOPs+KD	5.24/4.24	5.34	71.58	
MobileNet-v3 [22]	FP	32/32	218.7	74.52
	FP+KD	32/32	218.7	75.8
	PACT [8]	4/4	3.46	67.98
	PACT+KD [8]	4/4	3.46	70.16
	DuQ [31]	4/4	3.46	69.50
	DuQ+KD [31]	4/4	3.46	71.01
	NIPQ-MP-4	3.97/3.99	7.77	72.41
	NIPQ-BOPs	5.02/3.71	3.26	70.96
	NIPQ-MP-4+KD	4.05/4.00	7.96	74.49
	NIPQ-BOPs+KD	4.89/3.69	3.29	72.41

Table 2. Top-1 accuracy (%) of quantized networks on ImageNet dataset. MP-BOPs represents the mixed-quantization with the bit-operations (BOPs) constraint while MP-N is that with N-bit average bit-width. ‘*’ denotes the first and last layers remaining 8-bit, and KD denotes the knowledge distillation [21].

	Target Bit-width (Weight / Activation)						
	FP/FP	5/8	4/8	3/8	5/5	4/4	3/3
DoReFa [51]	0.857	0.593	0.541	0.359	0.588	0.498	0.288
PACT [8]	0.857	0.845	0.835	0.806	0.838	0.811	0.708
LSQ [16]	0.857	0.85	0.843	0.815	0.843	0.823	0.761
NIPQ-MP	0.857	0.851	0.848	0.833	0.848	0.836	0.801

Table 3. mAP comparison of Yolov5-S [23] on the PASCAL VOC [17].

First, we apply quantization to diverse networks on the ImageNet [11] classification task, and performs comparisons with various existing studies. Existing studies report mixed results with/without using well-known teacher-student-based knowledge distortion, so we report the accuracy for both cases, with EfficientNet-B0 as a teacher when necessary. Only the input image has 8-bit precision, and every layer in the network, including the first and last layers, is quantized via noise proxy.

Table 2 shows top-1 accuracy of the quantized networks.

NIPQ shows outstanding results for the optimized but hard to quantize networks such as MobileNet-v2/v3, as well as redundant networks such as ResNet-18. As shown in the table, existing methods are inferior to the proposed method, which achieves high accuracy in the same bit-width or bit-operations. These outstanding results come from two factors: first, NIPQ allows the network to converge to a more robust space without STE-induced instability, and second, within the resource budget, the quantization hyperparameters could be automatically tuned without the intervention of any hand-crafted manipulations. The benefit of these properties is maximized in optimized networks such as MobileNet-v2/v3. Note that when the average precision is constrained, NIPQ tries to increase accuracy with additional operations and vice versa. The automated tuning allows us to quantize the network considering our target goal.

In addition, we conduct an experiment to quantize the object detection task, which is known to be difficult to quantize. The difficulty is rapidly increased because we apply quantization to the advanced optimized network, Yolov5-S [23]. Table 3 shows the comparison of existing quantization studies in the same average bit-width. Due to large accuracy loss, existing 4-bit solutions are difficult to use in reality, but NIPQ shows practical, reliable quality in 4-bit precision. The results validate the stability of the NIPQ regardless of the difficulty of the target task.

Finally, to validate the superiority of NIPQ on the regression application, we apply quantization to the super-resolution task. In the case of super-resolution task, that conducts image restoration, the advantage of dynamic quantization, whose quantization parameters are updated regarding input data, is well demonstrated. Even in this case, NIPQ outmatches the best dynamic method, DDTB [50]. More results are available in the Supplementary Materials.

8. Conclusion

To achieve reliable output in low-precision, we propose NIPQ, which enhances the benefit of PQT pipeline by integrating the idea of truncation and providing theoretical support for it. NIPQ automates the layer-wise low-precision optimization for an arbitrary network by updating the network parameters and quantization parameters jointly via gradient descent toward minimizing the target loss. Our extensive results show that the proposed scheme outperforms all of the existing studies by a large margin for various vision and language applications.

Acknowledgements.

This work was supported by IITP grant funded by the Korea government (MSIT, No.2019-0-01906, No.2021-0-00105, and No.2021-0-00310). We appreciate valuable comments from Myeonghwan Ahn at SNU.

References

- [1] Milad Alizadeh, Arash Behboodi, Mart van Baalen, Christos Louizos, Tijmen Blankevoort, and Max Welling. Gradient ℓ_1 regularization for quantization robustness. *International Conference on Learning Representations (ICLR)*, 2020. [2](#)
- [2] Mário Almeida, Stefanos Laskaridis, Abhinav Mehrotra, Lukasz Dudziak, Ilias Leontiadis, and Nicholas D. Lane. Smart at what cost?: characterising mobile deep neural networks in the wild. *Internet Measurement Conference (IMC)*, 2021. [1](#)
- [3] Chaim Baskin, Natan Liss, Yoav Chai, Evgenii Zheltonozhskii, Eli Schwartz, Raja Giryes, Avi Mendelson, and Alexander M. Bronstein. NICE: noise injection and clamping estimation for neural network quantization. *CoRR*, abs/1810.00162, 2018. [2](#), [6](#)
- [4] Chaim Baskin, Eli Schwartz, Evgenii Zheltonozhskii, Natan Liss, Raja Giryes, Alexander M. Bronstein, and Avi Mendelson. UNIQU: uniform noise injection for the quantization of neural networks. *arXiv:1804.10969*, 2018. [7](#)
- [5] Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *CoRR*, abs/1308.3432, 2013. [1](#), [3](#)
- [6] Yash Bhalgat, Jinwon Lee, Markus Nagel, Tijmen Blankevoort, and Nojun Kwak. LSQ+: improving low-bit quantization through learnable offsets and better initialization. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. [2](#), [4](#)
- [7] Xiangning Chen and Cho-Jui Hsieh. Stabilizing differentiable architecture search via perturbation-based regularization. *International Conference on Machine Learning (ICML)*, 2020. [2](#), [5](#)
- [8] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. PACT: parameterized clipping activation for quantized neural networks. *CoRR*, abs/1805.06085, 2018. [1](#), [2](#), [3](#), [4](#), [7](#), [8](#)
- [9] Matthieu Courbariaux and Yoshua Bengio. Binarynet: Training deep neural networks with weights and activations constrained to +1 or -1. *CoRR*, abs/1602.02830, 2016. [1](#)
- [10] Alexandre Défossez, Yossi Adi, and Gabriel Synnaeve. Differentiable model compression via pseudo quantization noise. *CoRR*, abs/2104.09987, 2021. [2](#), [3](#), [4](#), [5](#), [6](#), [7](#)
- [11] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. *Computer Society Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009. [7](#), [8](#)
- [12] Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. HAWQ-V2: hessian aware trace-weighted quantization of neural networks. *Neural Information Processing Systems (NeurIPS)*, 2020. [2](#), [6](#), [7](#)
- [13] Zhen Dong, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. HAWQ: hessian aware quantization of neural networks with mixed-precision. *International Conference on Computer Vision (ICCV)*, 2019. [2](#), [6](#), [7](#), [8](#)
- [14] Charles Eckert, Xiaowei Wang, Jingcheng Wang, Arun Subramanian, Ravi R. Iyer, Dennis Sylvester, David T. Blaauw, and Reetuparna Das. Neural cache: Bit-serial in-cache acceleration of deep neural networks. *International Symposium on Computer Architecture (ISCA)*, 2018. [1](#)
- [15] Ahmed T. Elthakeb, Prannoy Pilligundla, Fatemehsadat Mireshghallah, Amir Yazdanbakhsh, and Hadi Esmaeilzadeh. Releq : A reinforcement learning approach for automatic deep quantization of neural networks. *IEEE Micro*, 40(5):37–45, 2020. [2](#)
- [16] Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S. Modha. Learned step size quantization. *International Conference on Learning Representations (ICLR)*, 2020. [1](#), [2](#), [3](#), [4](#), [7](#), [8](#)
- [17] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John M. Winn, and Andrew Zisserman. The pascal visual object classes (VOC) challenge. *Int. J. Comput. Vis.*, 88(2):303–338, 2010. [7](#), [8](#)
- [18] Boyuan Feng, Yuke Wang, Tong Geng, Ang Li, and Yufei Ding. APNN-TC: accelerating arbitrary precision neural networks on ampere GPU tensor cores. *The International Conference on High Performance Computing, Networking, Storage and Analysis (SC)*, 2021. [1](#)
- [19] Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazhen Lin, Fengwei Yu, and Junjie Yan. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. *International Conference on Computer Vision (ICCV)*, 2019. [1](#), [2](#), [8](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016. [1](#), [8](#)
- [21] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Dark knowledge. *Presented as the keynote in BayLearn*, 2(2), 2014. [8](#)
- [22] Andrew Howard, Ruoming Pang, Hartwig Adam, Quoc V. Le, Mark Sandler, Bo Chen, Weijun Wang, Liang-Chieh Chen, Mingxing Tan, Grace Chu, Vijay Vasudevan, and Yukun Zhu. Searching for mobilenetv3. *International Conference on Computer Vision (ICCV)*, 2019. [1](#), [8](#)
- [23] Glenn Jocher, Alex Stoken, Jirka Borovec, Liu Changyu, Adam Hogan, L Diaconu, F Ingham, J Poznanski, J Fang, L Yu, et al. ultralytics/yolov5: v3. 1-bug fixes and performance improvements. *Zenodo*, 2020. [8](#)
- [24] Sangil Jung, Changyong Son, Seohyung Lee, JinWoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, and Changkyu Choi. Learning to quantize deep networks by optimizing quantization intervals with task loss. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. [2](#), [3](#), [4](#)
- [25] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. *Citeseer*, 2009. [6](#)
- [26] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. *Advances in Neural Information Processing Systems (NeurIPS)*, 2018. [7](#)
- [27] Jing Liu, Jianfei Cai, and Bohan Zhuang. Sharpness-aware quantization for deep neural networks. *CoRR*, abs/2111.12273, 2021. [2](#)

- [28] Zechun Liu, Baoyuan Wu, Wenhan Luo, Xin Yang, Wei Liu, and Kwang-Ting Cheng. Bi-real net: Enhancing the performance of 1-bit cnns with improved representational capability and advanced training algorithm. *European Conference on Computer Vision (ECCV)*, 2018. **1, 3**
- [29] Markus Nagel, Marios Fournarakis, Yelysei Bondarenko, and Tijmen Blankevoort. Overcoming oscillations in quantization-aware training. In *International Conference on Machine Learning*, pages 16318–16330. PMLR, 2022. **1, 3**
- [30] Int4 precision for ai inference. <https://devblogs.nvidia.com/int4-for-ai-inference/>, 2019. Accessed: 2022-11-12. **1**
- [31] Eunhyeok Park and Sungjoo Yoo. PROFIT: A novel training method for sub-4-bit mobilenet models. *European Conference on Computer Vision (ECCV)*, 2020. **1, 2, 3, 7, 8**
- [32] Jongsoo Park, Maxim Naumov, Protonu Basu, Summer Deng, Aravind Kalaiah, Daya Shanker Khudia, James Law, Parth Malani, Andrey Malevich, Nadathur Satish, Juan Miguel Pino, Martin Schatz, Alexander Sidorov, Viswanath Sivakumar, Andrew Tulloch, Xiaodong Wang, Yiming Wu, Hector Yuen, Utku Diril, Dmytro Dzhulgakov, Kim M. Hazelwood, Bill Jia, Yangqing Jia, Lin Qiao, Vijay Rao, Nadav Rotem, Sungjoo Yoo, and Mikhail Smelyanskiy. Deep learning inference in facebook data centers: Characterization, performance optimizations and hardware implications. *CoRR*, abs/1811.09886, 2018. **1**
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32, 2019. **6**
- [34] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. *European Conference on Computer Vision (ECCV)*, 2016. **1**
- [35] Charbel Sakr, Steve Dai, Rangha Venkatesan, Brian Zimmer, William Dally, and Brucec Khailany. Optimal clipping and magnitude-aware differentiation for improved quantization-aware training. *International Conference on Machine Learning (ICML)*, 2022. **1, 3**
- [36] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *IEEE Conference on Computer Vision and Pattern Recognition, CVPR*, 2018. **1, 8**
- [37] Pedro Savarese, Xin Yuan, Yanjing Li, and Michael Maire. Not all bits have equal value: Heterogeneous precisions via trainable noise. *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. **2, 3**
- [38] Ali Shafiee, Anirban Nag, Naveen Muralimanohar, Rajeev Balasubramonian, John Paul Strachan, Miao Hu, R. Stanley Williams, and Vivek Srikumar. ISAAC: A convolutional neural network accelerator with in-situ analog arithmetic in crossbars. *International Symposium on Computer Architecture (ISCA)*, 2016. **7**
- [39] Sungho Shin, Jinhwan Park, Yoonho Boo, and Wonyong Sung. Hihlp: Quantized neural networks training for reaching flat minima in loss surface. *The Thirty-Fourth AAAI Conference on Artificial Intelligence (AAAI)*, 2020. **1**
- [40] Moran Shkolnik, Brian Chmiel, Ron Banner, Gil Shomron, Yury Nahshan, Alex M. Bronstein, and Uri C. Weiser. Robust quantization: One model to rule them all. *Neural Information Processing Systems (NeurIPS)*, 2020. **2, 7**
- [41] Jinook Song, Yunkyo Cho, Jun-Seok Park, Jun-Woo Jang, Sehwan Lee, Joon-Ho Song, Jae-Gon Lee, and Inyup Kang. 7.1 an 11.5 tops/w 1024-mac butterfly structure dual-core sparsity-aware neural processing unit in 8nm flagship mobile soc. *International Solid-State Circuits Conference (ISSCC)*, 2019. **1**
- [42] Andrew Tulloch and Yangqing Jia. High performance ultra-low-precision convolutions on mobile devices. *arXiv:1712.02427*, 2017. **1**
- [43] Andrew Tulloch and Yangqing Jia. Quantization and training of neural networks for efficient integer-arithmetic-only inference. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. **1, 3**
- [44] Stefan Uhlich, Lukas Mauch, Fabien Cardinaux, Kazuki Yoshiyama, Javier Alonso García, Stephen Tiedemann, Thomas Kemp, and Akira Nakamura. Mixed precision dnns: All you need is a good parametrization. *8th International Conference on Learning Representations, ICLR*, 2020. **2**
- [45] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. HAQ: hardware-aware automated quantization with mixed precision. *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. **2, 8**
- [46] Ying Wang, Yadong Lu, and Tijmen Blankevoort. Differentiable joint pruning and quantization for hardware efficiency. *European Conference on Computer Vision (ECCV)*, 2020. **1, 2, 3, 4, 5, 8**
- [47] Bernard Widrow, Istvan Kollar, and Ming-Chang Liu. Statistical theory of quantization. *IEEE Transactions on instrumentation and measurement*, 45(2):353–361, 1996. **5**
- [48] Hao Wu. NVIDIA Low Precision Inference on GPU. *GPU Technology Conference*, 2019. **1**
- [49] Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Ghodami, Jiali Yu, Eric Tan, Leyuan Wang, Qijing Huang, Yida Wang, Michael W. Mahoney, and Kurt Keutzer. HAWQ-V3: dyadic neural network quantization. *Proceedings of the 38th International Conference on Machine Learning, ICML*, 2021. **2, 6, 7, 8**
- [50] Yunshan Zhong, Mingbao Lin, Xunchao Li, Ke Li, Yunhang Shen, Fei Chao, Yongjian Wu, and Rongrong Ji. Dynamic dual trainable bounds for ultra-low precision super-resolution networks. *arXiv preprint arXiv:2203.03844*, 2022. **8**
- [51] Shuchang Zhou, Zekun Ni, Xinyu Zhou, He Wen, Yuxin Wu, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *CoRR*, abs/1606.06160, 2016. **1, 2, 3, 4, 8**