

# Solving Oscillation Problem in Post-Training Quantization Through a Theoretical Perspective

Yuexiao Ma<sup>1\*</sup>, Huixia Li<sup>2</sup>, Xiawu Zheng<sup>3</sup>, Xuefeng Xiao<sup>2</sup>, Rui Wang<sup>2</sup>,  
Shilei Wen<sup>2</sup>, Xin Pan<sup>2</sup>, Fei Chao<sup>1</sup>, Rongrong Ji<sup>1,4†</sup>

<sup>1</sup> Key Laboratory of Multimedia Trusted Perception and Efficient Computing, Ministry of Education of China, School of Informatics, Xiamen University, 361005, P.R. China.

<sup>2</sup> ByteDance Inc. <sup>3</sup> Peng Cheng Laboratory, Shenzhen, China.

<sup>4</sup> Shenzhen Research Institute of Xiamen University.

bobma@stu.xmu.edu.cn, zhengxw01@pcl.ac.cn, {fchao, rrji}@xmu.edu.cn,  
{lihuixia, xiaoxuefeng.ailab, ruiwang.rw, zhengmin.666, panxin.321}@bytedance.com.

## Abstract

Post-training quantization (PTQ) is widely regarded as one of the most efficient compression methods practically, benefiting from its data privacy and low computation costs. We argue that an overlooked problem of oscillation is in the PTQ methods. In this paper, we take the initiative to explore and present a theoretical proof to explain why such a problem is essential in PTQ. And then, we try to solve this problem by introducing a principled and generalized framework theoretically. In particular, we first formulate the oscillation in PTQ and prove the problem is caused by the difference in module capacity. To this end, we define the module capacity (ModCap) under data-dependent and data-free scenarios, where the differentials between adjacent modules are used to measure the degree of oscillation. The problem is then solved by selecting top-k differentials, in which the corresponding modules are jointly optimized and quantized. Extensive experiments demonstrate that our method successfully reduces the performance drop and is generalized to different neural networks and PTQ methods. For example, with 2/4 bit ResNet-50 quantization, our method surpasses the previous state-of-the-art method by 1.9%. It becomes more significant on small model quantization, e.g. surpasses BRECQ method by 6.61% on MobileNetV2  $\times 0.5$ .

## 1. Introduction

Deep Neural Networks (DNNs) have rapidly become a research hotspot in recent years, being applied to various

\*This work was done when Yuexiao Ma was intern at ByteDance Inc. Code is available at: <https://github.com/bytedance/MRECG>

†Corresponding Author: rrji@xmu.edu.cn

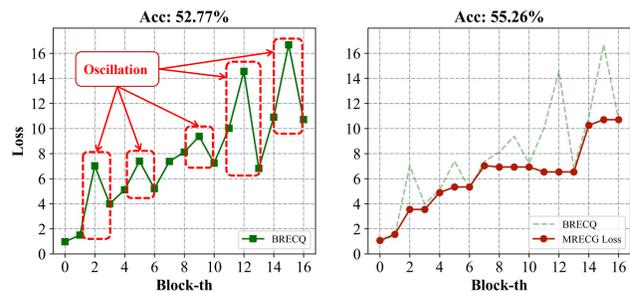


Figure 1. **Left:** Reconstruction loss distribution of BRECQ [17] on 0.5 scaled MobileNetV2 quantized to 4/4 bit. Loss oscillation in BRECQ during reconstruction see red dashed box. **Right:** Mixed reconstruction granularity (MRECG) smoothing loss oscillation and achieving higher accuracy.

scenarios in practice. However, as DNNs evolve, better model performance is usually associated with huge resource consumption from deeper and wider networks [8, 14, 28]. Meanwhile, the research field of neural network compression and acceleration, which aims to deploy models in resource-constrained scenarios, is gradually gaining more and more attention, including but not limited to Neural Architecture Search [18, 19, 33, 35–40, 42, 43], network pruning [4, 7, 16, 27, 32, 41], and quantization [3, 5, 6, 15, 17, 21, 22, 29, 31]. Among these methods, quantization proposed to transform float network activations and weights to low-bit fixed points, which is capable of accelerating inference [13] or training [44] speed with little performance degradation.

In general, network quantization methods are divided into quantization-aware training (QAT) [3, 5, 6] and post-training quantization (PTQ) [11, 17, 22, 31]. The former reduces the quantization error by quantization fine-tuning. Despite the remarkable results, the massive data require-

ments and high computational costs hinder the pervasive deployment of DNNs, especially on resource-constrained devices. Therefore, PTQ is proposed to solve the aforementioned problem, which requires only minor or zero calibration data for model reconstruction. Since there is no iterative process of quantization training, PTQ algorithms are extremely efficient, usually obtaining and deploying quantized models in a few minutes. However, this efficiency often comes at the partial sacrifice of accuracy. PTQ typically performs worse than full precision models without quantization training, especially in low-bit compact model quantization. Some recent algorithms [17, 22, 31] try to address this problem. For example, Nagel *et al.* [22] constructs new optimization functions by second-order Taylor expansions of the loss functions before and after quantization, which introduces soft quantization with learnable parameters to achieve adaptive weight rounding. Li *et al.* [17] changes layer-by-layer to block-by-block reconstruction and uses diagonal Fisher matrices to approximate the Hessian matrix to retain more information. Wei *et al.* [31] discovers that randomly disabling some elements of the activation quantization can smooth the loss surface of the quantization weights.

However, we observe that all the above methods show different degrees of oscillation with the deepening of the layer or block during the reconstruction process, as illustrated in the left sub-figure of Fig. 1. We argue that the problem is essential and has been overlooked in the previous PTQ methods. In this paper, through strict mathematical definitions and proofs, we answer 3 questions about the oscillation problem, which are listed as follows:

(i). **Why the oscillation happens in PTQ?** To answer this question, we first define module topological homogeneity, which relaxes the module equivalence restriction to a certain extent. And then, we give the definition of module capacity under the condition of module topological homogeneity. In this case, we can prove that when the capacity of the later module is large enough, the reconstruction loss will break through the effect of quantization error accumulation and decrease. On the contrary, if the capacity of the later module is smaller than that of the preceding module, the reconstruction loss increases sharply due to the amplified quantization error accumulation effect. Overall, we demonstrate that the oscillation of the loss during PTQ reconstruction is **caused by the difference in module capacity**;

(ii). **How the oscillation will influence the final performance?** We observe that the final reconstruction error is highly correlated with the largest reconstruction error in all the previous modules by randomly sampling a large number of mixed reconstruction granularity schemes. In other words, when oscillation occurs, the previous modules obviously have larger reconstruction errors, thus leading to worse accuracy in PTQ;

(iii). **How to solve the oscillation problem in PTQ?**

Since oscillation is caused by the different capacities of the front and rear modules, we propose the **Mixed REConstruction Granularity (MRECG)** method which jointly optimizes the modules where oscillation occurs. Besides, our method is applicable in data-free and data-dependent scenarios, which is also compatible with different PTQ methods. In general, our contributions are listed as follows:

- We reveal for the first time the oscillation problem in PTQ, which has been neglected in previous algorithms. However, we discover that smoothing out this oscillation is essential in the optimization of PTQ.
- We show theoretically that this oscillation is caused by the difference in the capability of adjacent modules. A small module capability exacerbates the cumulative effect of quantization errors making the loss increase rapidly, while a large module capability reduces the cumulative quantization errors making the loss decrease.
- To solve the oscillation problem, we propose a novel **Mixed REConstruction Granularity (MRECG)** method, which employs loss metric and module capacity to optimize mixed reconstruction granularity under data-dependency and data-free scenarios. The former finds the global optimum with moderately higher overhead and thus has the best performance. The latter is more effective with a minor performance drop.
- We validate the effectiveness of the proposed method on a wide range of compression tasks in ImageNet. In particular, we achieve a Top-1 accuracy of 58.49% in MobileNetV2 with 2/4 bit, which exceeds current SOTA methods by a large margin. Besides, we also confirm that our algorithm indeed eliminates the oscillation of reconstruction loss on different models and makes the reconstruction process more stable.

## 2. Related Work

**Quantization:** Quantization can be divided into two categories: **Quantization-Aware Training (QAT)** and **Post-Training Quantization (PTQ)**. QAT [3, 5, 45] uses the entire training dataset for quantization training and updates the gradients by back-propagation of the network to eliminate quantization errors. Although QAT integrates various training methods to achieve higher accuracy, this process is often resource intensive and limited in certain data privacy scenarios. Our research interest is not in this field.

PTQ has attracted increasing attention in recent years due to the advantages of efficient model deployment and low data dependence. Since quantization training is not included, PTQ algorithms usually use a small calibration dataset for reconstruction to obtain a better-quantized model

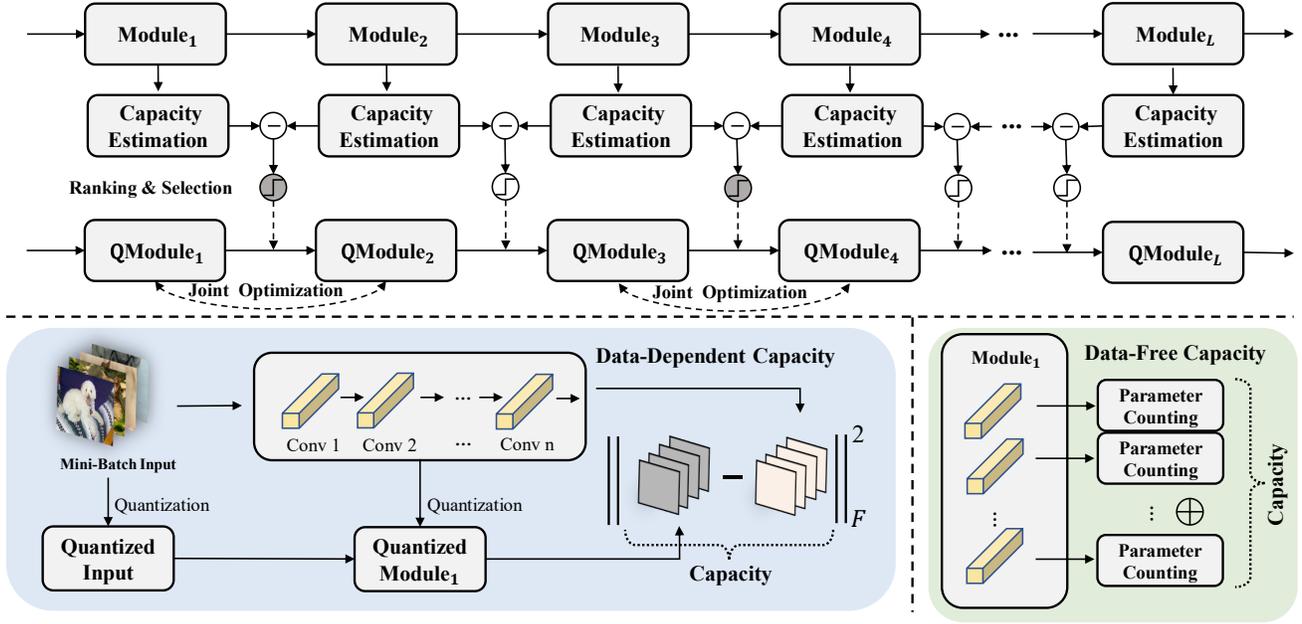


Figure 2. Overview of our method. Top: Optimization process for MRECG. We first estimate the capacity of each module, then we rank the capacity difference of adjacent modules, and finally we jointly optimize the adjacent modules corresponding to the top-k capacity difference. Bottom left: Module capacity estimation for data-dependent scenarios. We calculate the squared Frobenius norm of the feature map difference before and after quantization in the PTQ process. Bottom right: Module capacity estimation for the data-free scenario. We sum the convolutional layer capacities corresponding to the modules.

by optimizing the approximation of the second-order Taylor expansion term for task loss, which is introduced in Sec. 1.

**Module Capacity:** Some common parameters affect the module capacity, for example, the size of the filter, the bit width of the weight parameter, and the number of convolutional groups. In addition, some investigations show that the stride and residual shortcuts also affect the module capacity. Kong *et al.* [12] show that convolution with stride=2 can be equivalently replaced by convolution with stride=1. Meanwhile, the filter size of the replaced convolution is larger than that of the original convolution, implying an increase in module capacity. The MobileNetV2 proposed by [26] contains depth-wise convolution, which does not contain the exchange of information between channels, so it somehow compromises the model performance. Liu *et al.* [20] argue that the input of full precision residual shortcuts increases the representation capacity of the quantized module.

### 3. Methodology

In this section, we first prove that the oscillation problem of PTQ is highly correlated with the module capacity by a theorem and corollary. Secondly, we construct the capacity difference optimization problem and present two solutions in data-dependent and data-free scenarios, respectively. Finally, we analyze expanding the batch size of the calibration

data to reduce the expectation approximation error, which shows a trend of diminishing marginal utility.

#### 3.1. Oscillation problem of PTQ

Without loss of generality, we employ modules as the basic unit of our analysis. In particular, different from BRECG, module granularity in this paper is more flexible, which represents the layer, block, or even stage granularity. Formally, we set  $f_i^{(n_i)}(\cdot)$  be the  $i$ -th ( $i=1,2,\dots,L$ ) module of the neural network which contains  $n_i$  convolutional layers. We propose a more general reconstruction loss under the module granularity framework as follows,

$$\mathcal{L}(W_i, X_i) = \mathbb{E} \left[ \|f_i^{(n_i)}(W_i, X_i) - f_i^{(n_i)}(\widetilde{W}_i, \widetilde{X}_i)\|_F^2 \right], \quad (1)$$

where  $W_i, X_i$  are the weights and inputs of the  $i$ -th module, respectively.  $\widetilde{W}_i, \widetilde{X}_i$  are the corresponding quantized version. When  $f_i^{(n_i)}(\cdot)$  contains only one convolution layer, Eq. 1 degenerates to the optimization function  $\|W_i X_i - \widetilde{W}_i \widetilde{X}_i\|_F^2$  in AdaRound [22]. When  $f_i^{(n_i)}(\cdot)$  contains all convolutional layers within the  $i$ -th block, Eq. 1 degenerates to the optimization function  $\mathbb{E} \left[ \Delta z^{(i),T} H z^{(i)} \Delta z^{(i)} \right]$  in BRECG [17]. Note that we ignore the regularization term in AdaRound which facilitates

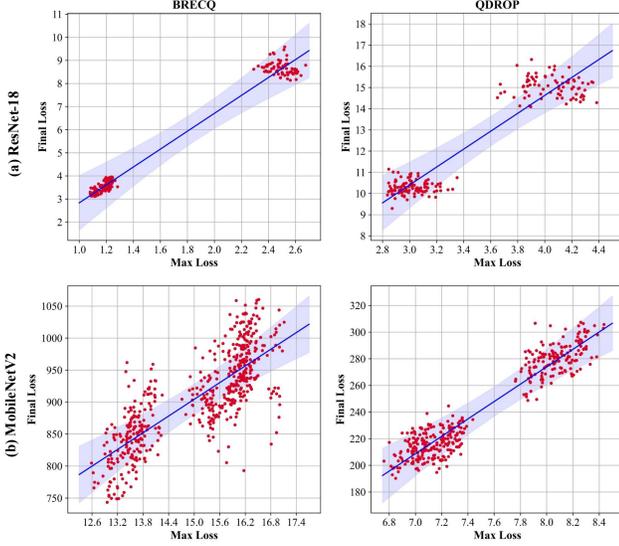


Figure 3. The relationship between the final reconstruction error and the maximum reconstruction error of the previous module. We randomly sample a number of mixed reconstruction granularity schemes and recover the accuracy using BRECQ or QDROP. We record the final reconstruction error and the maximum reconstruction error of the previous modules for each scheme. The first row of the figure represents the experiments on 4/4 bit ResNet-18. The second row of the figure represents the experiments on 4/4 bit 0.5 scaled MobileNetV2. The first and second columns represent using the BRECQ or QDROP algorithms to recover the accuracy.

convergence since it converges to 0 after each module is optimized. In addition, we omit the squared gradient scale in BRECQ for simplicity. If two modules have the same number of convolutional layers and all associated convolutional layers have the same hyper-parameters, they are called to be equivalent (including but not limited to additional residual input, kernel size, channels, groups, and stride).

There exists an accumulation effect of quantization error in quantization [34], which is manifested by the increasing tendency of quantization error in the network influenced by the quantization of previous layers. Since PTQ does not contain quantization training, this accumulation effect is more obvious in reconstruction loss. We propose the following theorem to demonstrate that the accumulation effect of quantization error in PTQ leads to an incremental loss.

**Theorem 1** *Given a pre-trained model and input data. If two adjacent modules are equivalent, we have,*

$$\mathcal{L}(W_i, X_i) \leq \mathcal{L}(W_{i+1}, X_{i+1}). \quad (2)$$

Detailed proof is provided in the supplementary material. Theorem 1 illustrates that under the condition that two adjacent modules are equivalent, the quantization error accumulation leads to the increment of loss

However, due to the difference of one or several convolution hyper-parameters (additional residual input, kernel size, channels, groups, stride, etc.), the condition of adjacent module equivalence in the above theorem is difficult to be satisfied in real scenarios. Non-equivalent modules will inevitably result in differences in module capacity. It is well-known that the number of parameters and the bit-width of the module affects the module capacity, in which the effect of these hyper-parameters on the module capacity is easy to quantify. In addition, we introduce in Sec. 2 that the residual input [20], the convolution type [26], and the convolution hyperparameters [12] all affect the module capacity, which the effects are difficult to quantify. Therefore, we define the concept of module topological homogeneity as follows, which relaxes the restrictions on module equivalence to some extent, while making it possible to compare capacities between different modules.

**Definition 1 (Module Topological Homogeneity)** *Suppose two modules have the same number of convolutional layers. If the hyper-parameters of the corresponding convolutional layers of the two modules are the same except for the kernel size and channels, we claim that the two modules are topologically homogeneous.*

From Definition 1, we relax the restriction of module equivalence on the equality of kernel size and channels in module topology homogeneity, which only results in differences in the number of module parameters. In other words, we obviate the problem that the impact of hyper-parameters such as residual inputs and groups makes the module capacity hard to be quantified. Specifically, if the module weight containing  $n$  convolutional layers is  $W = [W_1, W_2, \dots, W_n]$  and  $W_i$  is the weight of the  $i$ -th convolution layer in the module, then the module capacity (ModCap) is defined as the following equation.

$$\text{ModCap} = \sum_{i=1}^n \text{params}(W_i) \times b_i \times \alpha_i, \quad (3)$$

where  $\text{params}(\cdot)$  is a function that counts the number of parameters,  $b_i$  is the bit-width of the  $i$ -th convolutional layer, and  $\alpha_i$  is used to make convolutional layers that have different strides comparable. Specifically, convolutional layers of different strides are equated by some transformation [12], accompanied by a change in the number of parameters. Therefore, we convert the stride = 2 layer in the network to implicit stride = 1 layer by multiplying the scaling parameter  $\alpha_i$ . Under this conversion, all layers of the network satisfy the definition of topological homogeneity and thus are comparable with each other. Theoretically, according to the analysis of [12],  $\alpha_i$  can be set to 1.6, which is generalized in different networks in our paper.

Then, we derive a corollary of Theorem 1 to explain why the oscillations occur in PTQ.

**Corollary 1** *Suppose two adjacent modules be topologically homogeneous. If the module capacity of the later module is large enough, the loss will decrease. Conversely, if the latter module capacity is smaller than the former, then the accumulation effect of quantization error is exacerbated.*

Detailed proofs are provided in the supplementary material. From Corollary 1, we can conclude that the oscillation problem of PTQ on a wide range of models is caused by the excessive difference in capacities with adjacent modules. The logical/correlation chain of our paper: Oscillation  $\propto$  The largest error  $\propto$  Final error  $\propto$  Accuracy. **Oscillation  $\propto$  The largest error:** From Fig. 6 and a similar figure in the supplementary material, the more severe the oscillation of the loss distribution corresponding to the different algorithms, the larger the peak of the loss. That is, the degree of oscillation is positively correlated with the largest error. **The largest error  $\propto$  Final error:** In addition, our observations in Fig. 3 show that the largest error is positively correlated with the final error on different algorithms for different models. **Final error  $\propto$  Accuracy: Theoretically,** by performing a Taylor expansion on the accuracy loss function according to BRECC [17] and Adaround [22], we can derive the final reconstruction error that is highly correlated to the performance. **Empirically,** extensive experiments conducted in the above papers also prove the statement. In conclusion, the degree of oscillation of the error is positively correlated with the accuracy. Fig. 1 in our paper also demonstrates that reducing the degree of oscillation is beneficial to accuracy. In the next section, we will introduce how to eliminate oscillations in PTQ by optimizing module capacity differences.

### 3.2. Mixed REConstruction Granularity

The oscillation problem analyzed in Sec. 3.1 indicates that the information loss due to the difference in module capacities will eventually affect the performance of PTQ. And since there is no quantization training process in PTQ, this information loss is something that cannot be recovered even by increasing the model capacity of subsequent modules. Therefore, we smoothen the loss oscillations by jointly optimizing modules with large capacity differences, thus reducing the final reconstruction loss.

From Theorem 1 and Corollary 1, it is clear that a small capacity of the later module will aggravate the cumulative effect of the quantization error and thus increase the reconstruction loss sharply. Conversely, a large capacity of the later module will reduce the loss and increase the probability of information loss in the subsequent module. Therefore, we hope that the capacities of two adjacent modules are as close as possible. We construct the capacity difference optimization problem for the model containing  $L$  modules based on the capacity metric (CM) as follows,

$$\operatorname{argmin}_{\mathbf{m}} \sum_{l=1}^{L-1} (CM_l - CM_{l+1})^2 m_l + \lambda (\mathbf{m} \cdot \mathbf{1} - k)^2, \quad (4)$$

where  $\mathbf{m} \in \mathbb{R}^{L-1}$  is a binary mask vector. When  $m_i=1$ , it means we perform joint optimization for the  $i$ -th and  $(i+1)$ -th modules.  $\mathbf{1} \in \mathbb{R}^{L-1}$  is a vector with all elements of 1.  $k$  is a hyper-parameter that controls the number of jointly optimized modules.  $\lambda$  controls the importance ratio of the regularization term and the squared capacity difference optimization objective.

We calculate the square of the difference in capacities of all  $L - 1$  kinds of adjacent modules for ranking and select the top  $k$  adjacent modules for joint optimization. We use ModCap and reconstruction loss as our capacity metrics in the data-free and data-dependent scenarios, respectively. Details are shown in Fig. 2. The mixed reconstruction granularity scheme obtained according to the ModCap metric is computationally efficient because it does not involve the reconstruction. However, once we combine a pair of adjacent modules, this combined module cannot compare ModCap for further combination because it is not topologically homogeneous with the adjacent module. Therefore, this optimization scheme can only obtain the local optimal solution for the mixed reconstruction granularity.

On the other hand, according to Corollary 1, the ModCap difference is positively correlated with the difference in reconstruction loss, we can consider the reconstruction loss itself as a capacity metric. This scheme can get the global optimal solution, but it requires a PTQ reconstruction to obtain the reconstruction loss, which is relatively inefficient.

**Batch size of calibration data.** PTQ requires a small portion of the dataset to do calibration of the quantization parameters and model reconstruction. We note that Eq. 1 contains the expectation of the squared Frobenius norm. The expectation represents the average of the random variables, and we take the average of a sampled batch to approximate the expectation in the reconstruction process. The law of large numbers [9] proves that the mean of samples is infinitely close to the expectation when the sample size  $N$  tends to infinity, *i.e.*,

$$\begin{aligned} & \lim_{N \rightarrow +\infty} \frac{1}{N} \sum_{m=1}^N \|f_i^{(n)}(W_i, X_i^{(m)}) - f_i^{(n)}(\widetilde{W}_i, \widetilde{X}_i^{(m)})\|_F^2 \\ &= \mathbb{E} \left[ \|f_i^{(n)}(W_i, X_i) - f_i^{(n)}(\widetilde{W}_i, \widetilde{X}_i)\|_F^2 \right], \end{aligned} \quad (5)$$

Our experiments show that expanding the batch size of calibration data can improve the accuracy of PTQ. And this trend shows the diminishing marginal utility. Specifically, the improvement of PTQ accuracy slows down as the batch size increases. Details are shown in Sec. 4.3.

Methods	W/A	Res18	Res50	MBV2×1.0	MBV2×0.75	MBV2×0.5	MBV2×0.35
Full Prec.	32/32	71.01	76.63	72.20	69.95	64.60	60.08
ACIQ-Mix [1]	4/4	67.00	73.80	-	-	-	-
ZeroQ [2]	4/4	21.71	2.94	26.24	-	-	-
LAPQ [23]	4/4	60.30	70.00	49.70	-	-	-
AdaQuant [10]	4/4	69.60	75.90	47.16	-	-	-
Bit-Split [30]	4/4	67.56	73.71	-	-	-	-
AdaRound [22]	4/4	67.96	73.88	61.52	55.32	40.71	35.13
BRECQ* [17]	4/4	68.69	74.88	67.51	62.94	53.02	48.88
Ours+BRECQ	4/4	<b>69.06 (+0.37)</b>	74.84	<b>68.56 (+1.05)</b>	<b>64.55 (+1.61)</b>	<b>55.26 (+2.24)</b>	<b>50.67 (+1.79)</b>
QDROP [31]	4/4	69.10	75.03	67.89	63.26	54.19	49.79
Ours+QDROP	4/4	<b>69.46 (+0.36)</b>	<b>75.35 (+0.32)</b>	<b>68.84 (+0.95)</b>	<b>64.39 (+1.13)</b>	<b>55.64 (+1.45)</b>	<b>50.94 (+1.15)</b>
LAPQ [23]	2/4	0.18	0.14	0.13	-	-	-
AdaQuant [10]	2/4	0.11	0.12	0.15	-	-	-
AdaRound [22]	2/4	62.12	66.11	36.31	25.58	15.12	12.46
BRECQ* [17]	2/4	63.71	68.55	52.30	47.14	34.55	30.80
Ours+BRECQ	2/4	<b>65.61 (+1.9)</b>	<b>70.04 (+1.49)</b>	<b>58.49 (+6.19)</b>	<b>52.50 (+5.36)</b>	<b>41.16 (+6.61)</b>	<b>35.46 (+4.66)</b>
QDROP [31]	2/4	64.66	70.08	52.92	49.00	37.13	32.37
Ours+QDROP	2/4	<b>66.18 (+1.52)</b>	<b>70.53 (+0.45)</b>	<b>57.85 (+4.93)</b>	<b>53.71 (+4.71)</b>	<b>40.09 (+2.96)</b>	<b>35.85 (+3.48)</b>
AdaQuant [10]	3/3	60.09	67.46	2.23	-	-	-
AdaRound* [22]	3/3	63.91	64.85	34.55	18.16	8.13	4.45
BRECQ* [17]	3/3	64.83	70.06	52.03	45.54	29.79	25.52
Ours+BRECQ	3/3	<b>65.64 (+0.81)</b>	<b>70.68 (+0.62)</b>	<b>57.14 (+5.11)</b>	<b>50.21 (+4.67)</b>	<b>35.11 (+5.32)</b>	<b>30.26 (+4.74)</b>
QDROP [31]	3/3	65.56	71.07	54.27	49.26	35.14	29.40
Ours+QDROP	3/3	<b>66.30 (+0.74)</b>	<b>71.92 (+0.85)</b>	<b>58.40 (+4.13)</b>	<b>51.78 (+2.52)</b>	<b>38.43 (+3.29)</b>	<b>32.96 (+3.56)</b>
BRECQ* [17]	2/2	46.89	40.18	7.03	5.60	1.87	1.62
Ours+BRECQ	2/2	<b>52.02 (+5.13)</b>	<b>43.72 (+3.54)</b>	<b>13.84 (+6.81)</b>	<b>9.46 (+3.86)</b>	<b>3.43 (+1.56)</b>	<b>3.22 (+1.6)</b>
QDROP [31]	2/2	51.14	54.74	8.46	8.67	3.31	2.77
Ours+QDROP	2/2	<b>54.46 (+3.32)</b>	<b>56.82 (+2.08)</b>	<b>14.44 (+5.98)</b>	<b>11.40 (+2.73)</b>	<b>4.18 (+0.87)</b>	<b>3.09 (+0.32)</b>

Table 1. A comparison of our algorithm with the State-Of-The-Art method. We combine our algorithm with two latest PTQ algorithms, BRECQ and QDROP, and show significant improvements on a wide range of models. “\*” represents our reproduction of the algorithm based on an open-source codebase in a uniform experimental setup. “W/A” represents the bit width of the weights and activations, respectively. “Full Prec.” is the full precision pre-trained model. Under different bit configurations, we show in the table the comparison of our algorithm with a wide range of PTQ methods on ResNet-18, ResNet-50 and different scaled MobileNetV2.

## 4. Experiments

In this section, we perform a series of experiments to verify the superiority of our algorithm. First, we introduce the implementation details of our experiments. Second, we compare our algorithm quantized to different low bit-widths with the State-Of-The-Art on ImageNet datasets. Finally, we design a variety of ablation experiments to comprehensively analyze the properties of our algorithm, including the Pareto optimum subject to model size, the respective contributions of MRECG and expanded batch size, the phenomenon of diminishing marginal utility of expanded batch size, and loss distribution of different algorithms.

### 4.1. Implementation Details

We validate the performance of our algorithm on the ImageNet dataset [25], which consists of 1.2M training images and 50,000 validation images. We take 16 batches of data to perform the reconstruction of PTQ. We fix the batch size to 256 for ResNet-18 and MobileNetV2, and 128 for ResNet-50. In addition, our experimental analysis of the expanded batch size is shown in Sec. 4.3. Our data preprocessing follows [8] as everyone does. We use the Pytorch library [24] to complete our algorithm. We perform our experiments on an Nvidia Tesla A100 as well as an Intel(R) Xeon(R) Platinum 8336C CPU.

For PTQ reconstruction, our weight rounding scheme follows [22]. The rest of our reconstruction hyper-

Methods	Top-1 Acc(%)
Baseline	50.83
Baseline+MRECG	51.91
Baseline+ExpandBS	51.57
Baseline+both(Ours)	52.25

Table 2. Top-1 accuracy (%) for different component combinations. Baseline is 0.5 scaled MobileNetV2 quantized to 4/4 bit using the BRECQ algorithm. MRECG is mixed reconstruction granularity. ExpandBS represents expanded batch size.

parameters, such as reconstruction iterations, loss ratios, etc., are consistent with Adaround [22], BRECQ [17] and QDrop [31]. Also, as mentioned in QDrop, BRECQ additionally relaxes the bit-width of the first layer output to 8 bit, which somehow brings some accuracy gains. We re-executed BRECQ in an experimental setup uniform with all algorithms and obtained the accuracy under different bit configurations of models. Please refer to the supplementary material for other hyper-parameter settings.

## 4.2. ImageNet Classification

We validated our algorithm on a wide range of models, including ResNet-18, ResNet-50 and MobileNetV2 at different scales. As shown in Tab. 1, our algorithm outperforms other methods by a large margin on a wide range of models corresponding to different bit configurations. Specifically, we achieve 66.18% and 57.85% Top-1 accuracy on ResNet-18, MobileNetV2 with 2/4 bit, which are 1.52% and 4.93% higher than SOTA, respectively. Secondly, our method shows strong superiority at lower bits. For example, the accuracy gain of 2/4 bit PTQ quantization on MobileNetV2×0.75 is significantly higher than that of 4/4 bit (5.36% vs. 1.61%). The oscillation problem may be more significant because the low bit causes an increase in the magnitude of PTQ reconstruction error. So smoothing this oscillation at low bit quantization is crucial to the PTQ optimization process. Also, our method shows considerable improvement on MobileNetV2 at different scales, implying that the model size does not limit the performance of our algorithm. Besides, we notice that our algorithm is more effective for MobileNetV2. Through our observations, the oscillation problem in MobileNetV2 is more severe than that of the ResNet family of networks. The deep separable convolution increases the difference in module capacity and thus makes the magnitude of the loss oscillation larger. Therefore, our algorithm achieves a considerable performance improvement by solving the oscillation problem in MobileNetV2.

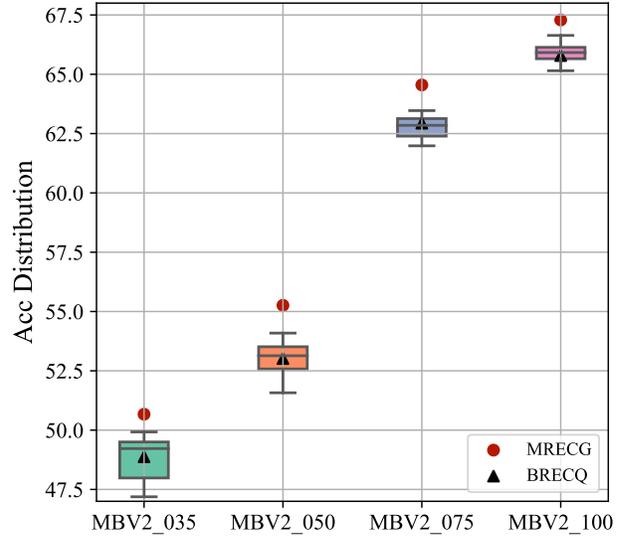


Figure 4. Pareto optimality of MRECG. We randomly sampled a large number of mixed reconstruction granularity schemes on four scales of MobileNetV2. We combine these schemes with the BRECQ algorithm to recover accuracy. The accuracy distribution of the sampled schemes on different models is given. Also, we mark the accuracy values of MRECG and BRECQ, respectively.

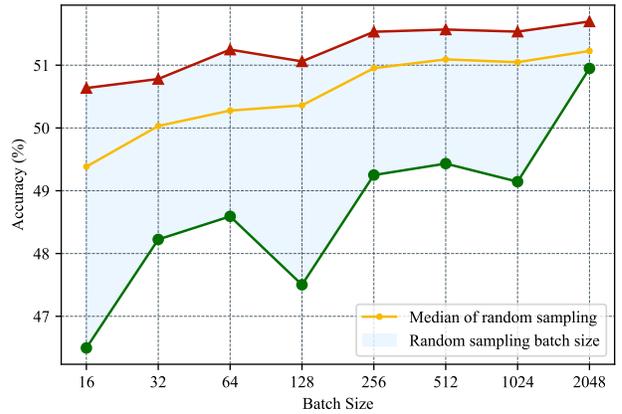


Figure 5. Impact of expanding batch size on PTQ accuracy. We sample some batch sizes on 4/4 bit MobileNetV2 and recover the accuracy. We demonstrate the median of sampled accuracy as the batch size increases to prevent the effect of outlier points.

## 4.3. Ablation Study

**Pareto optimal.** We randomly sample a large number of mixed reconstruction granularity schemes on different scaled MobileNetV2 and optimize these schemes with 4/4 bit PTQ using the BRECQ algorithm to obtain the sampling accuracy. We plot the accuracy distribution of these schemes in Fig. 4. Also we mark the accuracies of MRECG and BRECQ algorithms on MobileNetV2 at

different scales, respectively. From Fig. 4, we can observe that MRECG has stable accuracy improvement on MobileNetV2 at all scales. Overall, MRECG achieves the Pareto optimal state under the limitation of model size.

**Component Contribution.** Here we examine how much mixed reconstruction granularity and expanded batch size contribute to accuracy, respectively. We take the 0.5 scaled MobileNetV2 quantized to 4/4 bit by BRECG as the baseline. As shown in Tab. 2, the mixed reconstruction granularity and expanded batch size have 1.08% and 0.74% accuracy improvement, respectively. Meanwhile, the combination of the two methods further boosts model performance.

**Diminishing marginal utility of expanding batch size.**

By the law of large numbers [9], when the number of samples increases, the mean of samples converges to the expectation in Eq. 1 thus yielding a smaller approximation error. However, when the sample size is large enough, the accuracy gain from the reduction of the approximation error is negligible. In other words, expanding batch size presents a trend of diminishing marginal utility. In Fig. 5, we randomly sample some batch sizes and obtained PTQ accuracy of 4/4 bit by BRECG on 0.5 scaled MobileNetV2. We demonstrate the median of the sampling accuracy to prevent the effect of outliers. We notice that the sampling accuracy fluctuates more when the batch size is small, implying that the approximation error of a smaller batch size generates larger noise. This situation is moderated as the batch size is expanded. In addition, we observe an increase and then stabilization of the median accuracy, implying that expanding the batch size can bring accuracy gains to PTQ. Meanwhile, this gain is constrained by the diminishing marginal utility.

**Loss distribution.** As in Fig. 6, we present the loss distribution of different algorithms. From the distribution of Adaround we can see that it has the largest oscillation amplitude. Therefore, in the deeper layers of the model, the reconstruction loss of Adaround increases rapidly. In addition, BRECG uses block reconstruction for PTQ optimization, which somehow alleviates the loss oscillation problem and therefore brings performance improvement. Our mixed reconstruction granularity smoothes out the loss oscillations by jointly optimizing adjacent modules with large differences in capacity. It can be seen that the loss variation of MRECG is more stable compared to Adaround and BRECG. Moreover, when comparing the MRECG loss in different scenarios, we find that the MRECG obtained based on the data-free scenario still has a small amplitude loss oscillation. We believe that this modest oscillation is related to the fact that the ModCap MRECG can only jointly optimize two adjacent modules, which can only achieve local optimality. In contrast, the loss MRECG has the smoothest loss curve. However, this global optimum comes at the cost of prolonging the time to obtain the quantized model.

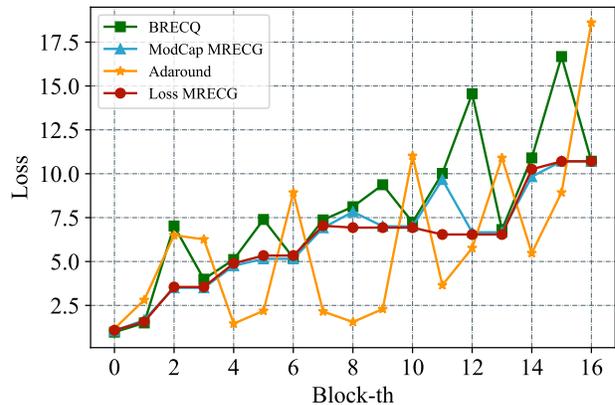


Figure 6. The reconstruction loss distributions of different algorithms for 4/4 bit configurations on 0.5 scaled MobileNetV2, including Adaround, BRECG, and MRECG constructed using ModCap and Loss metric for data-free or data-dependency scenarios. we omit the final fully connected layer loss for clarity.

**5. Conclusion**

In this paper, we discover an oscillation problem in the PTQ optimization process for the first time, which has been neglected in all previous PTQ algorithms. We then theoretically analyze that this oscillation is caused by the difference in the capacity of adjacent modules. Meanwhile, we observe that the final reconstruction error is positively correlated with the largest value of reconstruction error. Therefore, we construct a mixed reconstruction granularity optimization problem to smoothen the loss oscillations by jointly optimizing the adjacent modules with large capacity differences in data-dependent and data-free scenarios, which reduces the final reconstruction error. In addition, we observe that increasing the batch size of calibration data can reduce the expectation approximation error of the objective function. And this gain is of diminishing marginal utility. We validate the effectiveness of our method on extensive models with different bit configurations.

**Acknowledgments**

This work was supported by National Key R&D Program of China (No.2022ZD0118202), the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No. U21B2037, No. U22B2051, No. 62176222, No. 62176223, No. 62176226, No. 62072386, No. 62072387, No. 62072389, No. 62002305 and No. 62272401), and the Natural Science Foundation of Fujian Province of China (No.2021J01002, No.2022J06001), Guangdong Basic and Applied Basic Research Foundation (No. 2019B1515120049), CAAI-Huawei MindSpore Open Fund.

## References

- [1] Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapid-deployment. *Neural Information Processing Systems(NeurIPS)*, 32, 2019. 6
- [2] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Computer Vision and Pattern Recognition (CVPR)*, pages 13169–13178, 2020. 6
- [3] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018. 1, 2
- [4] Sara Elkerdawy, Mostafa Elhoushi, Hong Zhang, and Nilanjan Ray. Fire together wire together: A dynamic pruning approach with self-supervised mask prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12454–12463, June 2022. 1
- [5] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019. 1, 2
- [6] 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. In *International Conference on Computer Vision (ICCV)*, pages 4852–4861, 2019. 1
- [7] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*, 2015. 1
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. 1, 6
- [9] Pao-Lu Hsu and Herbert Robbins. Complete convergence and the law of large numbers. *Proceedings of the national academy of sciences*, 33(2):25–31, 1947. 5, 8
- [10] Itay Hubara, Yury Nahshan, Yair Hanani, Ron Banner, and Daniel Soudry. Improving post training neural quantization: Layer-wise calibration and integer programming. *arXiv preprint arXiv:2006.10518*, 2020. 6
- [11] Itay Hubara, Yury Nahshan, Yair Hanani, Ron Banner, and Daniel Soudry. Accurate post training quantization with small calibration sets. In *International Conference on Machine Learning*, pages 4466–4475. PMLR, 2021. 1
- [12] Chen Kong and Simon Lucey. Take it in your stride: Do we need striding in cnns? *arXiv preprint arXiv:1712.02502*, 2017. 3, 4
- [13] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018. 1
- [14] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60:84–90, 2017. 1
- [15] Huixia Li, Chenqian Yan, Shaohui Lin, Xiawu Zheng, Baochang Zhang, Fan Yang, and Rongrong Ji. Pams: Quantized super-resolution via parameterized max scale. In *European Conference on Computer Vision*, pages 564–580. Springer, 2020. 1
- [16] Yawei Li, Kamil Adamczewski, Wen Li, Shuhang Gu, Radu Timofte, and Luc Van Gool. Revisiting random channel pruning for neural network compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 191–201, June 2022. 1
- [17] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. {BRECQ}: Pushing the limit of post-training quantization by block reconstruction. In *International Conference on Learning Representations (ICLR)*, 2021. 1, 2, 3, 5, 6, 7
- [18] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European conference on computer vision (ECCV)*, pages 19–34, 2018. 1
- [19] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 1
- [20] Zechun Liu, Baoyuan Wu, Wenhao 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. In *European Conference on Computer Vision (ECCV)*, pages 722–737, 2018. 3, 4
- [21] Yuexiao Ma, Taisong Jin, Xiawu Zheng, Yan Wang, Huixia Li, Yongjian Wu, Guannan Jiang, Wei Zhang, and Rongrong Ji. Ompq: Orthogonal mixed precision quantization. *arXiv preprint arXiv:2109.07865*, 2021. 1
- [22] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pages 7197–7206. PMLR, 2020. 1, 2, 3, 5, 6, 7
- [23] Yury Nahshan, Brian Chmiel, Chaim Baskin, Evgenii Zheltonozhskii, Ron Banner, Alex M Bronstein, and Avi Mendelson. Loss aware post-training quantization. *Machine Learning*, 110(11):3245–3262, 2021. 6
- [24] 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
- [25] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 6
- [26] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Computer Vision and Pattern Recognition (CVPR)*, pages 4510–4520, 2018. 3, 4

- [27] Maying Shen, Pavlo Molchanov, Hongxu Yin, and Jose M. Alvarez. When to prune? a policy towards early structural pruning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12247–12256, June 2022. 1
- [28] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alexander A Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2017. 1
- [29] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. Haq: Hardware-aware automated quantization with mixed precision. In *Computer Vision and Pattern Recognition (CVPR)*, pages 8612–8620, 2019. 1
- [30] Peisong Wang, Qiang Chen, Xiangyu He, and Jian Cheng. Towards accurate post-training network quantization via bit-split and stitching. In *International Conference on Machine Learning*, pages 9847–9856. PMLR, 2020. 6
- [31] Xiuying Wei, Ruihao Gong, Yuhang Li, Xianglong Liu, and Fengwei Yu. Qdrop: Randomly dropping quantization for extremely low-bit post-training quantization. *arXiv preprint arXiv:2203.05740*, 2022. 1, 2, 6, 7
- [32] Paul Wimmer, Jens Mehnert, and Alexandru Condurache. Interspace pruning: Using adaptive filter representations to improve training of sparse cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12527–12537, June 2022. 1
- [33] Xin Xia, Xuefeng Xiao, Xing Wang, and Min Zheng. Progressive automatic design of search space for one-shot neural architecture search. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2455–2464, 2022. 1
- [34] Stone Yun and Alexander Wong. Do all mobilenets quantize poorly? gaining insights into the effect of quantization on depthwise separable convolutional networks through the eyes of multi-scale distributional dynamics. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2447–2456, 2021. 4
- [35] Shaokun Zhang, Feiran Jia, Chi Wang, and Qingyun Wu. Targeted hyperparameter optimization with lexicographic preferences over multiple objectives. In *The Eleventh International Conference on Learning Representations*, 2023. 1
- [36] Shaokun Zhang, Xiawu Zheng, Chenyi Yang, Yuchao Li, Yan Wang, Fei Chao, Mengdi Wang, Shen Li, Jun Yang, and Rongrong Ji. You only compress once: Towards effective and elastic bert compression via exploit-explore stochastic nature gradient. *arXiv preprint arXiv:2106.02435*, 2021. 1
- [37] Xiawu Zheng, Xiang Fei, Lei Zhang, Chenglin Wu, Fei Chao, Jianzhuang Liu, Wei Zeng, Yonghong Tian, and Rongrong Ji. Neural architecture search with representation mutual information. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 11912–11921, June 2022. 1
- [38] Xiawu Zheng, Rongrong Ji, Yuhang Chen, Qiang Wang, Baochang Zhang, Jie Chen, Qixiang Ye, Feiyue Huang, and Yonghong Tian. Migo-nas: Towards fast and generalizable neural architecture search. *IEEE TPAMI*, 43(9):2936–2952, 2021. 1
- [39] Xiawu Zheng, Rongrong Ji, Lang Tang, Baochang Zhang, Jianzhuang Liu, and Qi Tian. Multinomial distribution learning for effective neural architecture search. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1304–1313, 2019. 1
- [40] Xiawu Zheng, Rongrong Ji, Qiang Wang, Qixiang Ye, Zhen-guo Li, Yonghong Tian, and Qi Tian. Rethinking performance estimation in neural architecture search. In *CVPR*, 2020. 1
- [41] Xiawu Zheng, Yuexiao Ma, Teng Xi, Gang Zhang, Errui Ding, Yuchao Li, Jie Chen, Yonghong Tian, and Rongrong Ji. An information theory-inspired strategy for automatic network pruning. *arXiv preprint arXiv:2108.08532*, 2021. 1
- [42] Xiawu Zheng, Chenyi Yang, Shaokun Zhang, Yan Wang, Baochang Zhang, Yongjian Wu, Yunsheng Wu, Ling Shao, and Rongrong Ji. Ddpnas: Efficient neural architecture search via dynamic distribution pruning. *International Journal of Computer Vision*, pages 1–16, 2023. 1
- [43] Xiawu Zheng, Yang Zhang, Sirui Hong, Huixia Li, Lang Tang, Youcheng Xiong, Jin Zhou, Yan Wang, Xiaoshuai Sun, Pengfei Zhu, Chenglin Wu, and Rongrong Ji. Evolving fully automated machine learning via life-long knowledge anchors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):3091–3107, 2021. 1
- [44] Feng Zhu, Ruihao Gong, Fengwei Yu, Xianglong Liu, Yan-fei Wang, Zhelong Li, Xiuqi Yang, and Junjie Yan. Towards unified int8 training for convolutional neural network. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1969–1979, 2020. 1
- [45] Bohan Zhuang, Chunhua Shen, Mingkui Tan, Lingqiao Liu, and Ian Reid. Towards effective low-bitwidth convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7920–7928, 2018. 2