

# Fine-grained Data Distribution Alignment for Post-Training Quantization

Yunshan Zhong<sup>1,2</sup>, Mingbao Lin<sup>3</sup>, Mengzhao Chen<sup>2</sup>, Ke Li<sup>3</sup>,  
Yunhang Shen<sup>3</sup>, Fei Chao<sup>2</sup>, Yongjian Wu<sup>3</sup>, Rongrong Ji<sup>1,2\*</sup>

<sup>1</sup>Institute of Artificial Intelligence, Xiamen University.

<sup>2</sup>Media Analytics and Computing Lab, Department of Artificial Intelligence,  
School of Informatics, Xiamen University. <sup>3</sup>Tencent Youtu Lab.

zhongyunshan@stu.xmu.edu.cn, linmb001@outlook.com, cmzxmu@stu.xmu.edu.cn,  
{tristanli.sh, shenyunhang01}@gmail.com,  
fchao@xmu.edu.cn, littlekenwu@tencent.com, rrji@xmu.edu.cn

**Abstract.** While post-training quantization receives popularity mostly due to its evasion in accessing the original complete training dataset, its poor performance also stems from scarce images. To alleviate this limitation, in this paper, we leverage the synthetic data introduced by zero-shot quantization with calibration dataset and propose a fine-grained data distribution alignment (FDDA) method to boost the performance of post-training quantization. The method is based on two important properties of batch normalization statistics (BNS) we observed in deep layers of the trained network, *i.e.*, inter-class separation and intra-class incohesion. To preserve this fine-grained distribution information: 1) We calculate the per-class BNS of the calibration dataset as the BNS centers of each class and propose a BNS-centralized loss to force the synthetic data distributions of different classes to be close to their own centers. 2) We add Gaussian noise into the centers to imitate the incohesion and propose a BNS-distorted loss to force the synthetic data distribution of the same class to be close to the distorted centers. By utilizing these two fine-grained losses, our method manifests the state-of-the-art performance on ImageNet, especially when both the first and last layers are quantized to the low-bit. Code is at <https://github.com/zysxmu/FDDA>.

**Keywords:** Batch normalization statistics; Post-training quantization; Synthetic data

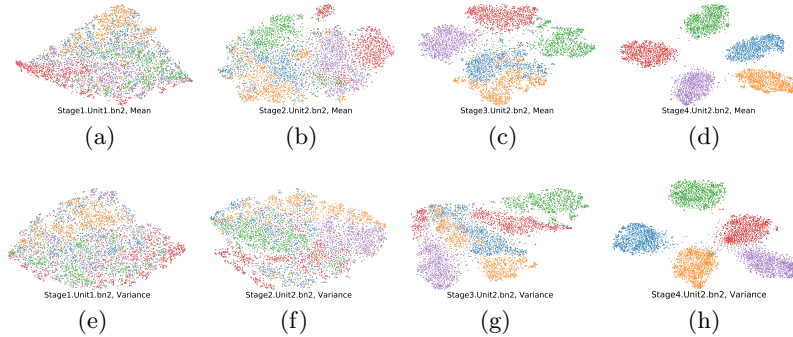
## 1 Introduction

Recent years have witnessed the rising of deep neural networks (DNNs) in computer vision. Nevertheless, the increasing model size barricades the deployment of DNNs on resource-limited platforms such as mobile phones, embedding devices, *etc.* To overcome this dilemma, varieties of methods [9,17] are explored to reduce the complexity of DNNs. Network quantization, which represents

---

\* Corresponding Author





**Fig. 2.** t-SNE visualization (five classes) of BNS in different layers of pre-trained ResNet-18 on ImageNet. BNS in shallow layers are overlapping. For deep layers, different classes have varying BNS and there exists small distortion within data from the same class. Similar observations can be found in other networks as well (See the supplementary materials). Best viewed in color.

to fit well the real data distribution. Many researches on zero-shot quantization are indicated to synthesizing fake images using data optimizer [38,2,39,10] or data generator [37,23,5]. The synthetic data is then used to train the quantized model. Though this manner partly alleviates data lacking, performance drops greatly if simply using synthetic data. For example, Li *et al.* [19] observed only 21.71% top-1 accuracy when quantizing ResNet-18 to 4-bit in a single zero-shot manner [2]. So far, combining real calibration dataset with synthetic data remains unexplored in post-training quantization, and we believe, might be a promise of boosting the low-bit performance.

Motivated by the above analyses, in this paper, we propose fine-grained data distribution alignment (termed FDDA) for post-training quantization, as illustrated in Fig. 1. FDDA is inspired by the fact that the representation of synthetic data has been demonstrated to be feasible in many other tasks such as image super-resolution [40], light-weight student network [3] and zero-shot quantization [37]. Thus, for the first time, we apply it to post-training quantization to tackle the insufficient data problem, providing a new perspective for post-training quantization. Following GDFQ [37], we take full use of the pre-trained model to guide the generator to synthesize fake data. Except for the distillation of output logits between the pre-trained model and its quantized version to improve the quantization performance, GDFQ also retains the distribution information of training data, modeled by the batch normalization statistics (BNS) in the pre-trained model. To this end, the mean and variance of the synthetic data distribution are constrained to be the same as those of the real data distribution. However, we realize that, based on our two insightful observations on BNS, this information retaining manner is very coarse for deep layers.

Specifically, we calculate the mean and variance *w.r.t.* each image and the results are visualized in Fig. 2. As can be seen, the BNS of different classes are

overlapping in shallow layers. However, two properties of deep BNS including inter-class separation and intra-class incohesion are observed. The former indicates different classes possess varying BNS while the latter indicates a small distortion of BNS among data from the same class. The BNS captured in the pre-trained model only reflects the distribution of the whole dataset, which are applicable to shallow layers with mixed class-wise BNS. While these BNS in the pre-trained model are coarse for deep layers with separable class-wise BNS. Thus, the synthetic data needs a fine-grained BNS alignment for these deep layers.

To this end, we further introduce a BNS-centralized loss and a BNS-distorted loss, respectively to align the fine-grained BNS properties of inter-class separation and intra-class incohesion. In contrast to zero-shot quantization [2], an additional calibration dataset, usually comprising one image per class, can be available in post-training quantization. To fully utilize this bonus, we derive means and variances of each image in deep layers of the pre-trained model, and then define the computed means and variances as the BN centers of each class. To preserve inter-class separation, our BNS-centralized loss forces the synthetic data distributions of different classes closely to their own centers. To preserve intra-class incohesion, we add Gaussian noise into the centers to mimic the distortion and the BNS-distorted loss forces the synthetic data distribution of the same class closely to the distorted centers. Through a fine-grained BNS alignment, our proposed FDDA significantly improves the quantization performance over existing methods on ImageNet [19,36,14], particularly when the first and last layers are also quantized to low-bit. Our contributions are three-fold:

- To our best knowledge, we are the first to explore combining calibration dataset with synthetic data in post-training quantization, which might provide a new perspective for post-training quantization.
- We observe properties of inter-class separation and intra-class incohesion in deep BNS. Besides, we devise a BNS-centralized loss and a BNS-distorted loss to preserve these two properties in synthetic data.
- Extensive experiments demonstrate that our FDDA can well improve the performance on ImageNet. For example, our FDDA outperforms the current SOTA, BRECQ [19], by 6.64% in the top-1 accuracy when all layers of MobileNet-V1 are quantized to 4-bit.

## 2 Related Work

In this section, we briefly discuss the most related work to ours including post-training quantization and zero-shot quantization. A more comprehensive survey is referred to [7].

**Post-training Quantization.** Most existing post-training quantization methods attempt to alleviate the accuracy deterioration problem from two perspectives: designing more sophisticated quantization methods and introducing a new rounding function or loss function. From the first perspective, Liu *et al.* [22] proposed to close the gap between the full-precision weight vector and its low-bit version by using the linear combination of multiple low-bit vectors. Wang *et*

*al.* [36] finished quantization in a two-stage manner of bit-split and bit-stitching. In the bit-split stage, the  $K$ -bit constraint of integer is split into  $(K - 1)$  ternary learning problems, and each bit is then separately solved in an iterative optimization procedure. In the bit-stitching stage, the  $K$ -bit integer is recovered by the linear combination with a base of  $2^{k-1}$  for the  $k$ -th bit. In addition, they also use filter-wise quantizer for activations and integrate the scaling factor into corresponding 2D kernels to avoid extra storage. Piece-wise linear quantization [6] splits the whole weights into two non-overlapping areas, including one dense region comprising low-magnitude weights, and one sparse region comprising high-magnitude weights. On top of the splitting, both areas are respectively quantized into the same low-bit. To model the quantization parameters, a linear regressor is constructed to predict the  $\alpha$ -quantile of activations [18], which eliminates the involvement of complex sorting algorithm. From the second perspective, AdaRound [28] analyzes that it is not advisable to simply round full precision weight to its nearest fixed-point value. Alternatively, the rounding problem is formulated as a per-layer quadratic unconstrained binary optimization problem, based on which, a continuous relaxation is introduced to find an adaptive rounding. BRECQ [19], one of the state-of-the-art methods, builds a block-wise reconstruction between the outputs of the full precision network and quantized network to achieve a balance between cross-layer dependency and generalization error. Besides, trainable clipping [4] for activations is also considered by BRECQ. Similar motivation can also be found in earlier works [15,36,1].

**Zero-Shot Quantization.** Zero-shot quantization is not permitted to access the training dataset. Thus, synthetic samples become an alternative to calibrate and fine-tune the quantized models. According to the methodology of data synthesis, we categorize existing studies into two groups: data optimizer [38,39,2,10,20] and data generator [37,23,5]. Data optimizer based methods produce synthetic images from Gaussian noise [38]. ZeroQ [2] forwards propagated the Gaussian inputs to collect BNS. Then the optimization that minimizes the difference between the collected BNS and the BNS in the pre-trained model is constructed to update the Gaussian inputs to ensure that the synthetic data does not deviate from the real data distribution. Except for aligning the BNS, a Domain Prior loss and an Inception loss are introduced in [10]. The former encourages nearby pixels between the input image and its Gaussian-smoothed variant to be similar. The latter prevents the model from producing inputs that lead to exploding outputs. To break the data homogenization, DSG [39] slacks the alignment of BNS and introduces a layer-wise enhancement to enhance diverse data samples. As for data generator, this group is featured with a generator in Generative Adversarial Networks (GAN) [8] to synthesize images. GDFQ [37] exploits the classification boundary knowledge and distribution information in the pre-trained model, and then devises a knowledge matching generator to produce synthetic data for model quantization. To diversify generated data, ZAQ [23] trains the quantized model and generator in an adversarial fashion by adopting an elaborated two-level discrepancy. To capture the distribution of the original

data lies on the decision boundaries, Qimera [5] introduces superposed latent embeddings to produce boundary supporting samples.

### 3 Methodology

#### 3.1 Preliminaries

**Quantizer.** Following [2, 37], we adopt asymmetric uniform quantization in this paper. Given the data  $\mathbf{x}$  (weights or activations), bit-width  $b$ , lower bound  $l$  and upper bound  $u$ , the quantizer is defined as:

$$\mathbf{q} = \text{round}\left(\frac{\text{clip}(\mathbf{x}, l, u)}{s}\right), \quad (1)$$

where  $\text{clip}(\mathbf{x}, l, u) = \min(\max(\mathbf{x}, l), u)$ ,  $\text{round}(\cdot)$  rounds its input to the nearest integer,  $s = \frac{u-l}{2^b-1}$  is the scaling factor that projects a floating-point number to a fixed-point integer, and  $\mathbf{q}$  is the quantized fixed-point number. The corresponding de-quantized item  $\bar{\mathbf{x}}$  can be obtained as:

$$\bar{\mathbf{x}} = \mathbf{q} \cdot s. \quad (2)$$


We use layer-wise quantizer and filter-wise quantizer for activations and weights, respectively. The lower bound  $l$  and upper bound  $u$  are set to the minimum and maximum of per-layer activations (per-channel weights).

**Data Synthesis.** Ideally, post-training quantization completes network compression with a small calibration dataset  $D = \{(\hat{\mathbf{x}}, \hat{y})\}$ , typically consisting of one image per class<sup>1</sup>. However, this small calibration dataset fails to retain performance when quantizing the network to very low precision, such as 4-bit. Inspired by zero-shot quantization, we resort to data synthesis. As shown in Fig. 1, we deploy a generator  $G$  to synthesize an image  $\tilde{\mathbf{x}}$  from a random Gaussian noise  $\tilde{\mathbf{z}}$  conditioned on the target label  $\tilde{y}$ , *i.e.*,  $\tilde{\mathbf{x}} = G(\tilde{\mathbf{z}}|\tilde{y})$ . We expect that the synthetic data to be similar to the real data. Despite the inaccessibility of the whole training data, we can turn to the data distribution information captured by the batch normalization statistics (BNS) in the pre-trained model  $F$ . Following [37], the BNS loss can be adopted to preserve the distribution:

$$\mathcal{L}_{\text{BNS}}^G = \sum_{l=1}^L \|\boldsymbol{\mu}'_l(\tilde{\mathbf{x}}) - \boldsymbol{\mu}_l^F\|^2 + \|\boldsymbol{\sigma}'_l(\tilde{\mathbf{x}}) - \boldsymbol{\sigma}_l^F\|_2^2, \quad (3)$$

where  $\boldsymbol{\mu}_l^F$  and  $\boldsymbol{\sigma}_l^F$  are the running mean and variance in the  $l$ -th layer of pre-trained  $F$ .  $\boldsymbol{\mu}'_l(\cdot)$  and  $\boldsymbol{\sigma}'_l(\cdot)$  return the mean and variance of input data in the  $l$ -th layer of  $F$ . This loss can be regarded as matching the first and second moment of real data and synthetic data.

<sup>1</sup> Occasionally, the label  $\hat{y}$  is not available. In this case, it can be predicted by the pre-trained full-precision model.



**Fig. 3.** The Average of Silhouette Coefficient values *w.r.t.* the BNS in different layers.

**Classification.** We also use cross-entropy loss to ensure synthetic data can be correctly classified by the pre-trained model  $F$ :

$$\mathcal{L}_{\text{CE}}^G = \mathbb{E}_{(\mathbf{x}, y) \sim \{(\tilde{\mathbf{x}}, \tilde{y})\}} [\text{CE}(F(\mathbf{x}), y)]. \quad (4)$$

Note that, we fix  $F$  during the whole training process, and the generator  $G$  is updated instead.

### 3.2 Our Insights

The BNS in pre-trained  $F$  are calculated by a series of averages of different batches of the full training set, *a.k.a.*, moving average. Consequently, they capture the distribution of the whole dataset. However, it is unclear whether they can be a representative of per-class BNS or even per-image BNS. To verify this, we feed the whole ImageNet to ResNet-18 [11] and calculate per-image mean vector  $\mu'_i$  and variance vector  $\sigma'_i$  in each layer. Fig. 2 visualizes several examples via t-SNE [25].

As can be seen, the BNS over different classes vary a lot across different layers. Specifically, the BNS of different classes are overlapping in shallow layers while there is regularity in deep layers, which, we outline into two properties: **inter-class separation** and **intra-class incohesion**. The former indicates that the BNS within the same class are formed into one cluster and different classes are separable. The latter refers to a small distortion of BNS among data from the same class. To quantitatively measure these two properties, we introduce Silhouette Coefficient (SC) [33], a value of which reflects how similar an object is to its own cluster (cohesion) in comparison with other clusters (separation). SC value of one sample  $\mathbf{v}$  is defined as:

$$\text{SC}(\mathbf{v}) = \frac{b(\mathbf{v}) - a(\mathbf{v})}{\max(a(\mathbf{v}), b(\mathbf{v}))}, \quad (5)$$

where  $a(\mathbf{v})$  denotes the average of intra-cluster distance for sample  $\mathbf{v}$ , and  $b(\mathbf{v})$  is the average of nearest-cluster distance for sample  $\mathbf{v}$ . Note that  $b(\mathbf{v})$  is the distance between  $\mathbf{v}$  and the nearest cluster that  $\mathbf{v}$  is not a part of. The value

of  $SC(\mathbf{v})$  ranging from -1 (separation) to +1 (cohesion). Values near 0 indicate overlapping clusters. Negative values generally indicate that a sample has been assigned to the wrong cluster, as a different cluster is more similar. With Eq. (5), we can obtain SC values for each mean vector  $\boldsymbol{\mu}'_l$  and variance vector  $\boldsymbol{\sigma}'_l$ . Fig. 3 displays the average of all input samples in each layer. As can be seen, the SC values are very small, even negative, in shallow layers, which indicates overlapping clusters. On the contrary, SC values increase in deep layers, which indicates inter-class separation. However, the increase of SC is very limited (no more than 0.25), which indicates a relatively large intra-cluster distance, *i.e.*, intra-class incohesion. These analyses are consistent with the observations in Fig. 2.

**In-depth Analysis.** The observed inter-class separation and intra-class incohesion can be explained by the fact that networks extract class-unrelated universal low-level features in shallow layers such as edges and curves. While in deep layers, networks are learned to extract class-related semantic features distinguishable from other classes in deep layers, leading to inter-class separation. The intra-class incohesion results from the varying image contents though these images are from the same class. For shallow layers, the overlapping clusters hardly model a subtle per-class distribution, thus the BNS in the pre-trained model can be an alternative. However, the BNS in the pre-trained model are very coarse and the constraint of Eq. (3) cannot model the properties of per-class separation and intra-class incohesion. Thus, in addition to the coarse-grained alignment, a fine-grained BNS alignment is also necessary for deep layers.

### 3.3 Fine-grained BNS Alignment

To preserve the properties of inter-class separation and intra-class incohesion, in this subsection, we respectively introduce a BNS-centralized loss and a BNS-distorted loss. Details are presented below.

**BNS-centralized Loss.** Since the BNS of each class are formed into one cluster, we can place a centroid as an explicit supervisory signal for each class and force the per-class distribution of synthetic images to be close to the assigned centroid. Recall that a real image  $(\hat{\mathbf{x}}, \hat{y})$  per class can be available from the calibration dataset  $D$  in post-training quantization. It is reasonable to use the BNS of image  $\hat{\mathbf{x}}$  as the corresponding centroid of class  $\hat{y}$  since its BNS already fall into the target cluster and are separable from BNS of other classes. To this end, given  $\hat{\mathbf{x}}$  with its label  $\hat{y} = c$ , we define the following BNS-centralized loss such that synthetic images can be further aligned to their corresponding centroids:

$$\begin{aligned} \mathcal{L}_{\text{C-BNS}}^G = & \sum_{l=K}^L \|\boldsymbol{\mu}'_l(\tilde{\mathbf{x}}|\tilde{y}=c) - \boldsymbol{\mu}'_l(\hat{\mathbf{x}}|\hat{y}=c)\|^2 \\ & + \|\boldsymbol{\sigma}'_l(\tilde{\mathbf{x}}|\tilde{y}=c) - \boldsymbol{\sigma}'_l(\hat{\mathbf{x}}|\hat{y}=c)\|^2, \end{aligned} \quad (6)$$

where  $K$  is a pre-given hyper-parameter denoting the start of deep layers. In all experiments, we set  $K = \text{ceil}(\frac{L}{2}) - 2$  where  $\text{ceil}(\cdot)$  is the rounding up function.



**BNS-distorted Loss.** Our BNS-centralized loss ensures the fine-grained separableness across different classes. However, how to retain the incohesion within the same class remains an issue. To solve this, we further propose to distort the centroid of per-class BNS by introducing Gaussian noise and define the following BNS-distorted loss:

$$\begin{aligned} \mathcal{L}_{\text{D-BNS}}^G = & \sum_{l=K}^L \left\| \boldsymbol{\mu}'_l(\tilde{\mathbf{x}}|\tilde{y}=c) - \mathcal{N}(\boldsymbol{\mu}'_l(\hat{\mathbf{x}}|\hat{y}=c), \mathbf{v}_\mu) \right\|^2 \\ & + \left\| \boldsymbol{\sigma}'_l(\tilde{\mathbf{x}}|\tilde{y}=c) - \mathcal{N}(\boldsymbol{\sigma}'_l(\hat{\mathbf{x}}|\hat{y}=c), \mathbf{v}_\sigma) \right\|^2, \end{aligned} \quad (7)$$

where  $\mathbf{v}_\mu = 0.5$  and  $\mathbf{v}_\sigma = 1.0$  are used to control the distortion degrees of mean and variance. For each synthetic data  $\tilde{\mathbf{x}}$ , its target is sampled from a Gaussian distribution centered on the class centroid of  $\tilde{\mathbf{x}}$ . Thus, our BNS-distorted loss provides diverse distorted centroids which prevent the BNS of per synthetic data from overfitting its centroid. This further retain the intra-class incohesion.

Our experimental results in Sec. 4.3 show that the BNS-distorted loss can retain the inter-class separation to some extent since the distorted centroid for synthetic data is centered on the corresponding class centroid. However, it is hard to manually model the Gaussian noise exactly such that the inter-class separableness and intra-class incohesion can be well preserved in the synthetic images simultaneously. Thus, both BNS-centralized loss and BNS-distorted loss are necessary as verified in the experiment.

### 3.4 Model Quantization

**Classification.** To take full use of the available data, both calibration images and synthetic images are used to fine-tune the quantized model  $Q$ , which can be realized through the cross-entropy loss:

$$\mathcal{L}_{\text{CE}}^Q = \mathbb{E}_{(\mathbf{x}, y) \sim D \cup \{(\tilde{\mathbf{x}}, \tilde{y})\}} [\text{CE}(Q(\mathbf{x}), y)]. \quad (8)$$

**Distillation.** It is possible that the synthetic image does not include corresponding class-specific features. As a result,  $\tilde{y}$  may be unreliable. Thus, we apply knowledge distillation (KD) [12] to transfer the outputs of full-precision model  $F$  to quantized model  $Q$  so that even though the synthetic image may have an inaccurate label,  $Q$  can still be correctly optimized by learning the soft target provided by  $F$ . Moreover, KD is also beneficial to the learning of calibration data. The KD loss is defined by the Kullback-Leibler distance  $\text{KL}(\cdot, \cdot)$  as:

$$\mathcal{L}_{\text{KD}}^Q = \mathbb{E}_{(\mathbf{x}, y) \sim D \cup \{(\tilde{\mathbf{x}}, \tilde{y})\}} [\text{KL}(Q(\mathbf{x}), F(\mathbf{x}))]. \quad (9)$$

### 3.5 Training Process

The training of our method consists of updating the generator  $G$  and the quantized model  $Q$ , where  $Q$  is obtained by quantizing the pre-trained full-precision model  $F$ .  $G$  produces a set of synthetic images while  $Q$  is trained with the aid of synthetic images and calibration images. We also emphasize that the  $F$  is fixed without any updating during the whole training process.

**Updating Generator  $G$ .** With a random Gaussian noise  $\tilde{z}$  conditional on label  $\tilde{y}$  as its input, the generator  $G$  synthesizes an image  $\tilde{x} = G(\tilde{z}|\tilde{y})$ , which is then used for classification of Eq. (4) and preserving distribution of training set including the coarse alignment of Eq. (3), inter-class separation of Eq. (6) and intra-class incohesion of Eq. (7). Thus, the overall loss for the generator  $G$  is derived as:

$$\mathcal{L}^G = \alpha_1 \cdot \mathcal{L}_{\text{CE}}^G + \alpha_2 \cdot \mathcal{L}_{\text{BNS}}^G + \alpha_3 \cdot \mathcal{L}_{\text{D-BNS}}^G + \alpha_4 \cdot \mathcal{L}_{\text{C-BNS}}^G, \quad (10)$$

where the  $\alpha_1, \alpha_2, \alpha_3$  and  $\alpha_4$  are the trade-off parameters.

**Updating Quantized Model  $Q$ .** The quantized model  $Q$  takes the synthetic data and calibration data as its inputs, and then the classification loss of Eq. (8) and distillation loss of Eq. (9) are constructed to retain the performance. Thus, the overall loss for the quantized model  $Q$  is derived as:

$$\mathcal{L}^Q = \mathcal{L}_{\text{CE}}^Q + \alpha_5 \cdot \mathcal{L}_{\text{KD}}^Q, \quad (11)$$

where  $\alpha_5$  is a trade-off parameter.

## 4 Experimentation

### 4.1 Implementation Details

We choose to quantize ResNet-18 [11], MobileNetV1 [13], MobileNetV2 [35] and RegNet-600MF [32]. All experiments are conducted on ImageNet with 1.2 million training images and 50,000 validation images from 1,000 classes [34]. The calibration dataset consists of 1,000 images including one image per class. We report the top-1 accuracy and the code is implemented using Pytorch [31]. Following the GDFQ [37], all full-precision models are downloaded from *pytorchcv*.

For ease of implementation, we directly import the generator from GDFQ [37] to produce synthetic images. The initial learning rates for the generator and quantized network are set to  $10^{-3}$  and  $10^{-6}$  respectively. For the generator, the optimizer is Adam [16] with 0.9 as the momentum and the learning rate are multiplied by 0.1 every 100 epochs. For the quantized network, the optimizer is SGD with Nesterov [30] with  $10^{-4}$  as the weight decay and we adjust the learning rate using the cosine annealing [24]. Before formal training, we set up a warm-up updating of the generator  $G$  for 50 epochs. Then, a total of 350 epochs are used to update the generator  $G$  and quantized model  $Q$ .

**Table 1.** Comparison with GDFQ [37] by quantizing all layers of ResNet-18 methods when all layers of ResNet-18 to 4-bit. TAQ denotes training-aware are quantized to 4-bit. “C” indicates quantization. “C” indicates the calibration dataset. “F” represents fine-grained data distribution alignment. Note that GDFQ + C + F = FDDA.

Method	Acc. (%)
Full precision	71.47
TAQ	68.24
GDFQ	60.60
GDFQ + C	65.64
GDFQ + C + F (FDDA)	<b>68.88</b>
FDDA + w/o label	<b>68.68</b>

**Table 2.** Comparison with zero-shot methods when all layers of ResNet-18 to 4-bit. TAQ denotes training-aware are quantized to 4-bit. “C” indicates quantization. “C” indicates the calibration dataset.

Method	Acc. (%)
Full precision	71.47
DI [38] + C	65.68
ADI [38] + C	66.30
ZeroQ [2] + C	66.89
ZAQ [23] + C	29.50
GDFQ [37] + C	65.64
Qimera [5] + C	66.19
DSG [39] + C	66.83
<b>FDDA</b>	<b>68.88</b>

## 4.2 Experimental Results

**Comparison with Zero-Shot Methods.** We first compare with the zero-shot GDFQ [37], since the data synthesis of our FDDA is built upon the framework of GDFQ. Table 1 displays our experimental results when all layers of ResNet-18 [11] are quantized to 4-bit. The accuracy of the full-precision model decreases from 71.47% to around 68.24% when the training-aware quantization, which requires all training data, is applied. Regarding GDFQ which only considers the synthetic data for fine-tuning the quantized model, performance severely degenerates to 60.60%. Such poor performance disables the application of GDFQ. Given the calibration dataset, GDFQ can increase to 65.48%, well demonstrating the correctness of our motive in combining real calibration dataset with synthetic data. Nevertheless, the coarse BNS alignment in GDFQ fails to model the fine-grained properties of inter-class separableness and intra-class incohesion. In contrast, our FDDA increases the performance to 68.88%, better than the training-aware quantization. This result shows the efficacy of our BNS-centralized loss and BNS-distorted loss in synthesizing better images.

In our settings, we assume to have access to the image labels. However, these labels are sometimes not available in real applications. Luckily, the pre-trained full-precision model can be used to predict these labels. In Table 1, we also report the performance of our FDDA, *i.e.*, 68.68%, using predicted labels. The slight drops are attributed to some of the misclassified labels. Nevertheless, our FDDA without real labels still maintains better performance than GDFQ with real images, well demonstrating the importance of preserving the fine-grained inter-class separableness and intra-class incohesion in learning to synthesize images.

Table 2 further shows the comparison between our FDDA with advanced zero-shot studies including data optimizer based methods [2,39,38] and data generator based methods [37,23,5]. For the former group, we use the calibration dataset as well as 10,000 synthetic images to train the quantized model, and the training process is the same with ours. For the latter group, we insert the

**Table 3.** Comparisons with existing post-training quantization methods. WBAB indicates the weights and activations are quantized to B-bit while FBLB indicates the first layers and last layers are quantized to B-bit.

	Methods	ResNet-18	MobileNetV1	MobileNetV2	RegNet-600MF
Settings	Full precision	71.47	73.39	72.49	73.71
W5A5, F8L8	ACIQ-Mix [1]	68.34	52.34	61.74	69.53
	AdaQuant [14]	68.56	-	65.19	-
	Bit-Split [36]	69.10	-	-	-
	BRECQ [19]	70.60	70.16	70.83	73.38
	<b>FDDA (Ours)</b>	<b>70.86</b>	<b>71.16</b>	<b>71.99</b>	<b>73.99</b>
W4A4, F8L8	ACIQ-Mix [1]	67.0	5.06	39.49	54.22
	AdaQuant [14]	67.50	-	34.95	-
	Bit-Split [36]	67.56	-	-	-
	BRECQ [19]	69.60	63.66	66.57	68.33
	<b>FDDA (Ours)</b>	<b>69.76</b>	<b>65.76</b>	<b>69.32</b>	<b>70.33</b>
W5A5, F5L5	ACIQ-Mix [1]	66.80	51.65	60.42	69.13
	AdaQuant [14]	68.19	-	63.61	-
	Bit-Split [36]	68.88	-	-	-
	BRECQ [19]	70.27	66.51	70.26	72.78
	<b>FDDA (Ours)</b>	<b>70.56</b>	<b>70.26</b>	<b>71.63</b>	<b>73.62</b>
W4A4, F4L4	ACIQ-Mix [1]	57.47	4.68	34.84	51.74
	AdaQuant [14]	63.45	-	34.64	-
	Bit-Split [36]	67.49	-	-	-
	BRECQ [19]	67.94	57.11	63.64	66.17
	<b>FDDA (Ours)</b>	<b>68.88</b>	<b>63.75</b>	<b>68.38</b>	<b>68.96</b>

calibration dataset into their training process. In Table 2, when the calibration dataset is applied to all methods, our FDDA still outperforms the advance of DSG by a large margin of 2.05%, which again demonstrates the efficacy of our BNS-centralized loss and BNS-distorted loss.

**Comparison with Competitors.** We compare with the recent studies on post-training quantization [1, 14, 36, 19]. The quantized networks include ResNet-18 [11], MobileNetV1 [13], MobileNetV2 [35] and RegNet-600MF [32]. All networks are quantized to the low precision 5-bit and 4-bit. Besides, to show the advantage of our FDDA, we quantize the first and last layers to 8-bit and lower precision (5-bit or 4-bit). Table 3 shows the experimental results.

If the first and last layers of full-precision models are quantized to 8-bit (F8L8), both our FDDA and recent SOTA BRECQ [19] can retain a high performance of the full-precision models regardless of 5-bit (W5A5) or 4-bit (W4A4) weights and activations in other layers. Comparing to BRECQ, FDDA obtains performance gains by 0.26%, 1.00%, 1.16% and 0.61% when quantizing ResNet-18, MobileNetV1, MobileNetV2 and RegNet-600MF to W5A5, while they are 0.16%, 2.10%, 2.75% and 2.00% when quantized to W4A4. Our FDDA retains better performance than BRECQ when quantizing light-weight models such as MobileNets, particularly when lower precision, such as 4-bit, is performed.




Fig. 4. Influence of the trade-off parameters.


When the first and last layers of full-precision models are quantized to lower precision (F5L5 or F4L4) as well, we notice that FDDA outperforms BRECQ by margins. Specifically, our FDDA increases the performance of BRECQ in W5A5 by 0.29%, 3.75%, 1.37% and 0.84% *w.r.t.* ResNet-18, MobileNetV2, MobileNetV2 and RegNet-600MF, and the performance gains are 0.94%, 6.64%, 4.74%, 2.79% in the case of W4A4. These results verify our statement in the introduction that the performance improvements of existing studies are on the premise that the first and last layers are quantized to higher precision, and also demonstrates the efficacy of our combining real calibration dataset with synthetic data.

### 4.3 Ablation Studies

We further study the influence of the trade-off parameters in Eq. (10) and Eq. (11), and the number of available images in the calibration dataset. All experiments are conducted by quantizing all layers of ResNet-18 to 4-bit.

**Trade-off Parameters.** We first display the influence of different trade-off parameters in Fig. 4. The  $\alpha_1$ ,  $\alpha_2$ ,  $\alpha_3$ , and  $\alpha_4$  from Eq. (10) balance different losses in updating the generator while  $\alpha_5$  from Eq. (11) balances the losses in updating the quantized model. Each  $\alpha_i$  is first empirically initialized. Then, for  $\alpha_i$ , we search its optimal value using the grid search with others fixed. From Fig. 4, we can see that the optimal configurations of these three parameters are  $\alpha_1 = 0.5$ ,  $\alpha_2 = 0.2$ ,  $\alpha_3 = 0.9$ ,  $\alpha_4 = 0.05$  and  $\alpha_5 = 20$ , which are also our settings for all the aforementioned experiments. Though they might not be the optimal for all networks, we find these configurations already bring better performance than the recent state-of-the-arts. Also, we observe that  $\alpha_4 \ll \alpha_3$ , in which  $\alpha_3$  and  $\alpha_4$  respectively balance the importance of the proposed BNS-distorted loss and BNS-centralized loss. This is due to the distorted centroid for synthetic data is centered on the corresponding class centroid, thus the BNS-distorted loss can retain the inter-class separation to some extent, which partly relieves the involvement of BNS-centralized loss and leads to a small  $\alpha_4$ .

**Effect of Available Classes.** The calibration dataset consists of 1,000 images including one image per class by default. However, some images might be missing in real-world applications. Consequently, the class information of corresponding image is not available for our fine-grained BNS alignment. In Fig. 5, we



**Fig. 5.** Effect of Available Classes.

further excavate the influence of available classes on our final performance. For unavailable classes, we omit their C-BNS and D-BNS loss when computing the Eq. (10). Note that in this case, the size of the calibration dataset is equal to the number of available classes. It can be seen from Fig. 5 that performance drops as the available classes decrease. Nevertheless, comparing to the recent advance, BRECCQ, which obtains only 67.94% top-1 accuracy (see Table 3), our FDDA still maintains a higher performance of 68.09% even when only 700 classes are available. The good performance can be attributed to two reasons. On one hand, synthetic data benefits post-training quantization even though some classes are missing. On the other hand, the fine-grained data alignment helps to synthesize better images for fine-tuning the quantized model.

## 5 Conclusion

In this paper, we proposed a fine-grained data distribution alignment (FDDA) to solve the insufficient data problem in post-training quantization. We observed two important BNS properties of the inter-class separation and intra-class incohesion in the deep layers of neural network. To retain these two fine-grained distribution information, we respectively proposed the BNS-centralized loss and BNS-distorted loss. Using a real image from the calibration dataset as the centroid of each class, the BNS-centralized loss constrains the BNS of synthetic data to be close to the BNS of its class centroid, while the BNS-distorted loss introduces Gaussian noise to distort the class centroid for the purpose of incohesion. By retaining these two properties in the synthetic data, our FDDA shows its superiority over the state-of-the-art competitors on ImageNet, particularly in the hardware-friendly case where the first and last layers of networks are also quantized to low precision.

**Acknowledgements.** This work was supported by the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No. U21B2037, No. 62176222, No. 62176223, No. 62176226, No. 62072386, No. 62072387, No. 62072389, and No. 62002305), Guangdong Basic and Applied Basic Research Foundation (No.2019B1515120049), and the Natural Science Foundation of Fujian Province of China (No.2021J01002).

## References

1. Banner, R., Nahshan, Y., Soudry, D., et al.: Post training 4-bit quantization of convolutional networks for rapid-deployment. In: *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*. pp. 7950–7958 (2019)
2. Cai, Y., Yao, Z., Dong, Z., Gholami, A., Mahoney, M.W., Keutzer, K.: Zeroq: A novel zero shot quantization framework. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 13169–13178 (2020)
3. Chen, H., Wang, Y., Xu, C., Yang, Z., Liu, C., Shi, B., Xu, C., Xu, C., Tian, Q.: Data-free learning of student networks. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 3514–3522 (2019)
4. Choi, J., Wang, Z., Venkataramani, S., Chuang, P.I.J., Srinivasan, V., Gopalakrishnan, K.: Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085* (2018)
5. Choi, K., Hong, D., Park, N., Kim, Y., Lee, J.: Qimera: Data-free quantization with synthetic boundary supporting samples. In: *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)* (2021)
6. Fang, J., Shafiee, A., Abdel-Aziz, H., Thorsley, D., Georgiadis, G., Hassoun, J.H.: Post-training piecewise linear quantization for deep neural networks. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 69–86 (2020)
7. Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M.W., Keutzer, K.: A survey of quantization methods for efficient neural network inference. *arXiv preprint arXiv:2103.13630* (2021)
8. Goodfellow, I.J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*. pp. 2672–2680 (2014)
9. Han, S., Pool, J., Tran, J., Dally, W.J., et al.: Learning both weights and connections for efficient neural network. In: *Proceedings of the Advances in Neural Information Processing Systems (NeurIPS)*. pp. 1135–1143 (2015)
10. Haroush, M., Hubara, I., Hoffer, E., Soudry, D.: The knowledge within: Methods for data-free model compression. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8494–8502 (2020)
11. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 770–778 (2016)
12. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531* (2015)
13. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
14. Hubara, I., Nahshan, Y., Hanani, Y., Banner, R., Soudry, D.: Improving post training neural quantization: Layer-wise calibration and integer programming. *arXiv preprint arXiv:2006.10518* (2020)
15. Hubara, I., Nahshan, Y., Hanani, Y., Banner, R., Soudry, D.: Accurate post training quantization with small calibration sets. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 4466–4475 (2021)
16. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: *Proceedings of the International Conference on Learning Representations (ICLR)* (2014)

17. Krishnamoorthi, R.: Quantizing deep convolutional networks for efficient inference: A whitepaper. arXiv preprint arXiv:1806.08342 (2018)
18. Kryzhanovskiy, V., Balitskiy, G., Kozyrskiy, N., Zuruev, A.: Qpp: Real-time quantization parameter prediction for deep neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10684–10692 (2021)
19. Li, Y., Gong, R., Tan, X., Yang, Y., Hu, P., Zhang, Q., Yu, F., Wang, W., Gu, S.: Breq: Pushing the limit of post-training quantization by block reconstruction. In: Proceedings of the International Conference on Learning Representations (ICLR) (2021)
20. Li, Y., Zhu, F., Gong, R., Shen, M., Dong, X., Yu, F., Lu, S., Gu, S.: Mixmix: All you need for data-free compression are feature and data mixing. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4410–4419 (2021)
21. Lin, M., Ji, R., Xu, Z., Zhang, B., Wang, Y., Wu, Y., Huang, F., Lin, C.W.: Rotated binary neural network. In: Proceedings of the Advances in Neural Information Processing Systems (NeurIPS). pp. 7474–7485 (2020)
22. Liu, X., Ye, M., Zhou, D., Liu, Q.: Post-training quantization with multiple points: Mixed precision without mixed precision. In: Proceedings of the AAAI Conference on Artificial Intelligence (AAAI). pp. 8697–8705 (2021)
23. Liu, Y., Zhang, W., Wang, J.: Zero-shot adversarial quantization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1512–1521 (2021)
24. Loshchilov, I., Hutter, F.: Sgdr: Stochastic gradient descent with warm restarts. In: Proceedings of the International Conference on Learning Representations (ICLR) (2016)
25. van der Maaten, L., Hinton, G.: Visualizing data using t-sne. *Journal of Machine Learning Research (JMLR)* **9**, 2579–2605 (2008)
26. Martinez, J., Shewakramani, J., Liu, T.W., Bărsan, I.A., Zeng, W., Urtasun, R.: Permute, quantize, and fine-tune: Efficient compression of neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 15699–15708 (2021)
27. Martinez, J., Zakhmi, S., Hoos, H.H., Little, J.J.: Lsq++: Lower running time and higher recall in multi-codebook quantization. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 491–506 (2018)
28. Nagel, M., Amjad, R.A., Van Baalen, M., Louizos, C., Blankevoort, T.: Up or down? adaptive rounding for post-training quantization. In: Proceedings of the International Conference on Machine Learning (ICML). pp. 7197–7206 (2020)
29. Nahshan, Y., Chmiel, B., Baskin, C., Zheltonozhskii, E., Banner, R., Bronstein, A.M., Mendelson, A.: Loss aware post-training quantization. arXiv preprint arXiv:1911.07190 (2019)
30. Nesterov, Y.E.: A method of solving a convex programming problem with convergence rate  $o(k^{-2})$ . In: Proceedings of the Russian Academy of Sciences (RAS). pp. 543–547 (1983)
31. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Proceedings of the Advances in Neural Information Processing Systems (NeurIPS). pp. 8026–8037 (2019)
32. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 10428–10436 (2020)



33. Rousseeuw, P.J.: Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics (JCAM)* **20**, 53–65 (1987)
34. Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al.: Imagenet large scale visual recognition challenge. *International Journal of Computer Vision (IJCV)* **115**, 211–252 (2015)
35. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 4510–4520 (2018)
36. Wang, P., Chen, Q., He, X., Cheng, J.: Towards accurate post-training network quantization via bit-split and stitching. In: *Proceedings of the International Conference on Machine Learning (ICML)*. pp. 9847–9856 (2020)
37. Xu, S., Li, H., Zhuang, B., Liu, J., Cao, J., Liang, C., Tan, M.: Generative low-bitwidth data free quantization. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. pp. 1–17 (2020)
38. Yin, H., Molchanov, P., Alvarez, J.M., Li, Z., Mallya, A., Hoiem, D., Jha, N.K., Kautz, J.: Dreaming to distill: Data-free knowledge transfer via deepinversion. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 8715–8724 (2020)
39. Zhang, X., Qin, H., Ding, Y., Gong, R., Yan, Q., Tao, R., Li, Y., Yu, F., Liu, X.: Diversifying sample generation for accurate data-free quantization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 15658–15667 (2021)
40. Zhang, Y., Chen, H., Chen, X., Deng, Y., Xu, C., Wang, Y.: Data-free knowledge distillation for image super-resolution. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7852–7861 (2021)
41. Zhong, Y., Lin, M., Nan, G., Liu, J., Zhang, B., Tian, Y., Ji, R.: Intraq: Learning synthetic images with intra-class heterogeneity for zero-shot network quantization. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 12339–12348 (2022)