

SEAM: Searching Transferable Mixed-Precision Quantization Policy through Large Margin Regularization

Chen Tang^{1,2}, Kai Ouyang¹, Zenghao Chai^{1,3}, Yunpeng Bai¹,

Yuan Meng^{2†}, Zhi Wang^{1,4†}, Wenwu Zhu^{2†}

¹Shenzhen International Graduate School, Tsinghua University

²Department of Computer Science and Technology, Tsinghua University

³School of Computing, National University of Singapore ⁴Peng Cheng Laboratory

{tc20, oyk20, czh20, byp20}@mails.tsinghua.edu.cn

yuanmeng@mail.tsinghua.edu.cn, wangzhi@sz.tsinghua.edu.cn, wwzhu@tsinghua.edu.cn

ABSTRACT

Mixed-precision quantization (MPQ) suffers from the time-consuming process of searching the optimal bit-width allocation (*i.e.*, the policy) for each layer, especially when using large-scale datasets such as ISLVR-2012. This limits the practicality of MPQ in real-world deployment scenarios. To address this issue, this paper proposes a novel method for efficiently searching for effective MPQ policies using a small proxy dataset instead of the large-scale dataset used for training the model. Deviating from the established norm of employing a consistent dataset for both model training and MPQ policy search phases, our approach, therefore, yields a substantial enhancement in the efficiency of MPQ exploration. Nonetheless, using discrepant datasets poses challenges in searching for a transferable MPQ policy. Driven by the observation that quantization noise of sub-optimal policy exerts a detrimental influence on the discriminability of feature representations – manifesting as diminished class margins and ambiguous decision boundaries – our method aims to identify policies that uphold the discriminative nature of feature representations, *i.e.*, intra-class compactness and inter-class separation. This general and dataset-independent property makes us search for the MPQ policy over a rather small-scale proxy dataset and then the policy can be directly used to quantize the model trained on a large-scale dataset. Our method offers several advantages, including high proxy data utilization, no excessive hyper-parameter tuning, and high searching efficiency. We search high-quality MPQ policies with the proxy dataset that has only 4% of the data scale compared to the large-scale target dataset, achieving the same accuracy as searching directly on the latter, improving MPQ searching efficiency by up to 300×.

CCS CONCEPTS

• **Computing methodologies** → **Machine learning**; • **Computer systems organization** → **Neural networks**.

KEYWORDS

Model quantization, Efficient deep learning, Model compression

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

MM '23, October 29–November 3, 2023, Ottawa, ON, Canada

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0108-5/23/10.

<https://doi.org/10.1145/3581783.3611975>

ACM Reference Format:

Chen Tang^{1,2}, Kai Ouyang¹, Zenghao Chai^{1,3}, Yunpeng Bai¹, and Yuan Meng^{2†}, Zhi Wang^{1,4†}, Wenwu Zhu^{2†}. 2023. SEAM: Searching Transferable Mixed-Precision Quantization Policy through Large Margin Regularization. In *Proceedings of the 31st ACM International Conference on Multimedia (MM '23)*, October 29–November 3, 2023, Ottawa, ON, Canada. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3581783.3611975>

1 INTRODUCTION

With the success of deep learning, deep neural networks (DNNs) have been adopted for many artificial intelligence tasks such as image classification [18, 20, 35], object detection [33, 34], and meanwhile, become the indispensable part of modern multimedia applications [37]. However, the large computational resource requirements of DNNs remain one of the most giant stumbling blocks for deploying deep learning models. There are several compression techniques to reduce the redundancy in a deep model, such as pruning [29], knowledge distillation [19] and quantization [6, 26, 41, 53]. Quantization is a promising technique to reduce both the storage and computational resources overhead remarkably, by leveraging the fact that the inference precision is not strictly as high as training time. Therefore, quantization enables large models to run directly on the edge and mobile devices without redesigning a new model architecture, which empowers edge intelligence significantly.

Quantization can be divided into two categories: fixed-precision quantization and mixed-precision quantization (MPQ). Fixed-precision quantization [12, 28, 53], where an identical bit-width is designated for all layers in a deep model. While such a paradigm is proven to make the quantized model achieve sufficiently good performance at high bit-width (*e.g.*, ≥ 8 bits), a uniform bit-width is challenging for quantization in an ultra-low bit-width (*i.e.*, ≤ 4 bits) scenario. For example, BRQ [17] reports that there is more than 20% top-1 accuracy degradation in a 2 bit quantization for the MobilNetV2 model as compared to its full-precision (FP) counterpart.

Mixed-precision quantization (MPQ) [11, 14, 21, 46, 51] offers a flexible and efficient way to quantize deep models by allocating varying bit-widths to individual layers based on their diverse redundancy levels. Unlike fixed-precision quantization, MPQ assigns specific precisions to different layers, with higher redundancy layers receiving less bit-width than lower redundancy ones, thereby achieving an optimal accuracy-efficiency trade-off. The MPQ process typically involves two stages: firstly, a full-precision (FP) model \mathcal{M}_{FP} is trained on a training dataset \mathcal{D}_{train} ; subsequently, the FP

[†]Corresponding authors

model is served as a weight initialization to be quantized, while simultaneously searching for the optimal MPQ policy for determining the quantization precision to each layer, over a searching dataset \mathcal{D}_{search} . The searching process also performs quantization-aware training, therefore all search-based approaches [3, 21, 46] use the same dataset during both two processes, namely, $\mathcal{D}_{train} = \mathcal{D}_{search}$. Although using consistent datasets surely bring an accurate policy for the model to quantize, two problematic issues arise: **(a)** when $\mathcal{D}_{train} = \mathcal{D}_{search}$ and \mathcal{D}_{train} is large-scale, the combinatorial nature [40, 46] of the MPQ problem poses severe difficulties in search efficiency (e.g., BP-NAS consumes 35.6 GPU-hours to search for the ResNet-50 [47]). **(b)** in some sensitive user-data application scenarios, the training dataset is inaccessible.

Nevertheless, few research has explored to decouple the dataset used in model training and MPQ search stages. This is promising to improve the search efficiency since the searching process can be done on a small-scale proxy dataset, but inevitably encounters intractable challenges due to shifted data distributions and reduced data volume caused by the disparate datasets. For example, when using CIFAR-10 to search for MPQ policies for ResNet50 trained on ISLVR-2012, EdMIPS [3] lost nearly 7% of Top-1 accuracy [47]. Recently, GMPQ [47] indicates that, for an input image, preserving the attribution rank between the FP and quantized model can search a generalizable MPQ policy. They resort the feature visualization technique Grad-cam [36] to maintain the consistency of image attribution rank between the quantized and FP model. GMPQ can be regarded as an instance-level regularization over the proxy dataset, by enforcing a consistent relationship between FP and quantized model of each input instance. However, it does not exploit the information outside of an instance, i.e., in a class-level, let alone the fussy hyper-parameters tuning for aligning the attribution rank.

In this paper, we present a novel method to search the effective transferable MPQ policy by exploiting the class-level information on the proxy datasets, considering the class-level information is more luxuriant than instance-level [4]. Our idea is motivated by the observation that quantization poses side effects to the quantized model in the feature space compared to the FP model. Our significant finding has shed light on a common drawback of quantization: the quantization noise remarkably narrows the margin between classes and blurs the decision boundary (see Figure 2). On the other hand, maximizing inter-class separation while enhancing the intra-class compactness is highly favorable for classification, as there is a consensus that a large classification margin enhances the generalizability from statistical machine learning (e.g., SVM) to recent deep learning research [32, 45]. We hence propose to search the MPQ policy that can properly gather the features of the same classes and separate the features of different classes, making the features more robust to quantization noise.

Experimental results validate that a large margin between classes of proxy data helps search for a transferable MPQ policy for quantizing the model trained on challenging large-scale datasets. Our approach achieves competitive performance when searching on very small proxy datasets versus directly on large-scale datasets, in which the size of the former is only 4% of the latter. Consequently, we improve the MPQ policy search efficiency impressively. For ResNet18 and MobileNetV1, by using StanfordCars [23] as the proxy

dataset, our method achieves 375 \times and 300 \times speedup compared to the state-of-the-art MPQ approach FracBits [50], respectively.

2 RELATED WORK

2.1 Fixed-Precision Quantization

Fixed-precision quantization assigns a uniform bit-width for all layers. In this paper, we only consider quantization-aware training, as it can achieve higher compression ratio than post-training quantization [22, 30] and zero-shot quantization [25, 49, 52].

Dorefa [53] and PACT [6] uses a low-precision representation for weights and activations during forward propagation, and utilizes the Straight-Through Estimation (STE) [2] to estimate the gradient of piece-wise quantization function for backward propagation. To relieve the bias gradient of the STE, DSQ [13] employs tangent functions to approximate the non-differentiable quantization function. LSQ [12] introduces the learnable step-size scale factors to scale the tensor-wise weight and activation distributions. BSQ [17] further applies a bin regularization to ensure the weights fall in the center of quantization bins. All these works focus on training a well-performing quantized network, but suffer from severe performance degradation when the bit-width is decreased significantly.

2.2 Mixed-Precision Quantization

The fundamental of Mixed-precision quantization (MPQ) is that the different layers in a model have different redundancy, in which the high redundancy layers can be allocated small bit-width to ensure low complexity without a severe performance drop. However, the bit-width choice is discrete, and the combination of bit-width and layer (i.e., the policy) grows exponentially. Therefore, the main challenge is how to determine the optimal bit-width for each layer.

Obviously, brute-force is rather ineffective for the purpose of searching, as an L layers model with n bit-widths for activations and weights has n^{2L} possible policies [46]. To solve this, several studies make efforts to apply the intelligent algorithms to search the optimal MPQ policy. HAQ [46] and ReleQ [11] use reinforcement learning (RL) to train a bit-width allocator. SPOS [14], EdMIPS [3] and BP-NAS [51] adopt neural architecture search (NAS) methods to learn the bit-width. In particular, GMPQ [47] develops an instance-level regularization to make searching MPQ policy on a small dataset possible. However, GMPQ suffers from a fussy hyper-parameters tuning, including the approximated attribution rank level, number of interested pixels, etc.

Unlike learning the optimal MPQ policy, HAWQ [8, 9] and MPQCO [5] use the Hessian information as the quantization sensitivity metrics to assist bit-width assignment. LIMPQ [40] proposes to learn the layer-wise importance during a once quantization-aware training process. In contrast to these methods that aim to define some metrics to estimate the quantization sensitivity of layers, we propose to directly learn the effective bit-width configurations on a small proxy dataset.

2.3 Discriminative Feature Learning

Learning discriminative feature is highly favorable since it greatly facilitates the generalization of deep models, its core is to clarify the decision boundaries between classes. For nearly two decades, there are several studies to make efforts to achieve this.

bit-width for weights and activations of layer l is determined by an *argmax* function acts on its learnable weights vector α^l and β^l .

This paradigm and its variants [3, 21] require the searching dataset to be consistent with the full-precision model training one, otherwise resulting in a serious accuracy degradation [47]. Inevitably, using a consistent dataset leads to inefficiencies, especially on large-scale datasets like ISLVR2012 [7] with over 1 million samples to search for.

However, when searching an MPQ policy on a proxy dataset (e.g., a small-scale dataset CIFAR-10 with only 50000 training samples) through Equation 1 and then directly applying it to the model trained on a large-scale dataset (e.g., ISLVR2012), while the accuracy and complexity are both met, the accuracy on the proxy dataset is not of direct interest to us, because high accuracy on proxy dataset does not imply equivalent high accuracy on challenging large-scale datasets. One may argue that we can abridge the size of the target dataset to improve the efficiency, such as using a subset of target datasets to conduct MPQ search, but this would also result in serious performance degradation, as shown in Sec. 4.4.

Accordingly, instead of optimizing the above improper objective on the proxy dataset, we aim to search an MPQ policy that guarantees a large-margin on the proxy dataset to handle the incoming classes of the large-scale dataset.

3.2 Exploiting the Class-level Information

From the perspective of class-level features in a well-preforming MPQ policy, they should be well separated if not in the same class, and tightly gathered if in the same class. This has the following benefits: **a)** It alleviates the side effect of quantization on classification boundary. As shown in Figure 2(a) and Figure 2(b), we observe quantization sharply narrows the class boundaries in the feature space compared to the full-precision model. Therefore, an MPQ policy with an explicit feature separation guarantee can effectively alleviate the side effect of quantization. **b)** This is a widely pursued and *dataset-independent* attribute, as from classical statistical machine learning to recent deep learning research [27, 32, 45] both recognize a large classification margin in feature space can help generalization.

Motivated by this, we aim to search the MPQ policy that guarantees the large class margin on the proxy data distribution as much as possible. As we discussed above, such a general property in searched MPQ policies can ensure usability across the data distributions. However, the cross-entropy cannot provide this property, as the class margin is not explicitly formulated. Therefore, the objective is not only to optimize accuracy and complexity, but also to find an MPQ policy that maximizes the class margin.

We regard our approach as a class-level proxy data utilization, as it discovers the effective MPQ policy by leveraging the inter-class and intra-class information on the proxy dataset. The 2D visualization of our approach is shown in Figure 2(c), we observe that the t-SNE pattern is quite similar to the full-precision model, indicating an MPQ policy that is able to separate the features is searched for the quantized model.

3.2.1 Separation Regularization. The first term in Equation 1 is the soft-max cross-entropy loss [3, 47]. For simplicity, we revisit it here




Figure 2: The deep feature 2D visualization (t-SNE [42]) on a proxy dataset CIFAR-10 over (a) the full-precision ResNet18, (b) direct searched MPQ policy through EdMIPS [3] and (c) searched MPQ policy through proposed method. Colors represent different classes.

by considering a binary classification problem, which can be trivial generalized to multi-class classification,

$$\begin{aligned} \mathcal{L}_{task} &= -\log \frac{\exp(\mathbf{w}_1^\top \mathbf{g})}{\exp(\mathbf{w}_1^\top \mathbf{g}) + \exp(\mathbf{w}_2^\top \mathbf{g})} \\ &= -\log \frac{1}{1 + \underbrace{\exp(\mathbf{w}_2^\top \mathbf{g} - \mathbf{w}_1^\top \mathbf{g})}_{\text{equivalent optimized term}}}, \end{aligned} \quad (4)$$

where \mathbf{w}_1^\top and \mathbf{w}_2^\top are the weights for class 1 and class 2, respectively. \mathbf{g} is the deep feature of the model produced by several convolution layers (*i.e.*, layers that need to be quantized to mixed-precision).

Since the equivalent optimized term is not carried the margin objective during optimization, Equation 4 cannot explicitly guarantee any margin between classes. Some previous works even observe that the learned feature regions for some classes tend to be bigger than others. If this combines with the side effect of quantization on decision boundaries, it inevitably leads to the search for sub-optimal MPQ policies. In other words, the performance objective in Equation 1, the cross-entropy, is improper when the MPQ searching and full-precision model training datasets are inconsistent.

To this end, we introduce separation regularization to enforce a large margin guarantee in the searched policy. Firstly, a small intra-class variance should be achieved to compact the features,

$$\min_q \sum_{i=1}^N q_i, \quad \text{where } q_i = d(\mathbf{g}_i, \boldsymbol{\mu}_{y_i}), \quad (5)$$

where N is the number of samples, \mathbf{g}_i , y_i and $\boldsymbol{\mu}_{y_i}$ are the feature and label (ground truth) of sample i and the feature mean of class y_i , respectively. $d(\cdot, \cdot)$ is the metric for calculating the distance between the feature and its mean (*e.g.*, L2 distance).

Secondly, we consider the inter-class margin by minimizing a classification loss as

$$\begin{aligned} \min_o \mathcal{L}_{cls} &= \min_o \sum_{i=1}^N \sum_{j=1}^K o_{i,j}, \\ o_{i,j} &= \begin{cases} -\log \frac{\exp(h_j(\mathbf{g}_i; \mathbf{m}))}{\sum_{k=1}^K \mathbb{1}(k \neq y_i) \exp(h_k(\mathbf{g}_i; 0)) + \exp(h_j(\mathbf{g}_i; \mathbf{m}))}, & \text{if } j = y_i \\ 0, & \text{otherwise,} \end{cases} \end{aligned} \quad (6)$$

$h(\cdot; \cdot)$ is a map from feature space \mathbb{R}^D (i.e., \mathbf{g}) to class-wise prediction score and will be introduced in Equation 9. $\mathbb{1}(\cdot)$ is the indicator function and K represents the number of classes, respectively. m is a non-negative scalar that represents the margin of different classes to form an explicit classification margin between the label class of sample i and other classes in feature space, i.e., $h_j(\mathbf{g}_i; m) > h_k(\mathbf{g}_i; 0)$ ($k \neq j$, and $j = y_i$). One can see Equation 6 becomes the classic log-softmax cross-entropy loss when $h(\cdot; m)$ is a linear transformation and $m \equiv 0$, e.g., in classic softmax cross-entropy, a linear layer with weight $\mathbf{W} \in \mathbb{R}^{D \times K}$ and no biases is used to project the deep feature \mathbf{g}_i to \mathbb{R}^K —let us denote \mathbf{w}_j is the j -th column vector of \mathbf{W} , thus $h_j(\mathbf{g}_i; 0) = \mathbf{w}_j^\top \mathbf{g}_i$. Please note when $m \neq 0$, the classification margin requires the output sign of h should be always either positive or negative, which is not always satisfied in a classic softmax cross-entropy loss as the sign of linear projection is not certain.

We hence follow the previous work L-GM [45] that assumes the feature \mathbf{g}_i follows a Gaussian Mixture Distribution (GMD). Namely,

$$p(\mathbf{g}_i) = \sum_{k=1}^K p(k) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_k, \Sigma_k), \quad (7)$$

where $p(k)$ is the prior probability of class k , and $\boldsymbol{\mu}_k$ and Σ_k are the mean and covariance of class k . The posterior probability of feature \mathbf{g}_i is derived through the Bayes' rule,

$$\begin{aligned} p(y_i | \mathbf{g}_i) &= \frac{p(y_i) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_{y_i}, \Sigma_{y_i})}{p(\mathbf{g}_i)} \\ &= \frac{p(y_i) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_{y_i}, \Sigma_{y_i})}{\sum_{k=1}^K p(k) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_k, \Sigma_k)}. \end{aligned} \quad (8)$$

Under the GMD assumption, we can easily derive the additive inter-class margin according to

$$h_{y_i}(\mathbf{g}_i; m) = p(y_i) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_{y_i}, \Sigma_{y_i}, m) = p(y_i) |\Sigma_{y_i}|^{-\frac{1}{2}} \exp \left\{ - \underbrace{\frac{1}{2} (\mathbf{g}_i - \boldsymbol{\mu}_{y_i})^\top \Sigma_{y_i}^{-1} (\mathbf{g}_i - \boldsymbol{\mu}_{y_i}) + m}_{\text{non-negative}} \right\}, \quad (9)$$

where $h(\cdot)$ is formulated from a probability perspective, thus it is guaranteed to be non-negative. By replacing the subscript y_i of Equation 9 with k and setting $m = 0$, we can derive $h_k(\mathbf{g}_i; 0) = p(k) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_k, \Sigma_k, 0)$. Substitute it and Equation 9 into Equation 6, we can obtain the \mathcal{L}_{cls} accordingly. Finally, we apply a log-likelihood term [45] to restrict the feature \mathbf{g}_i centralization near its mean $\boldsymbol{\mu}_{y_i}$ to achieve intra-class compactness according to Equation 5 and Equation 7,

$$\begin{aligned} \mathcal{L}_{inc} &= \sum_{i=1}^N q_i = \sum_{i=1}^N d(\mathbf{g}_i, \boldsymbol{\mu}_{y_i}) \\ &= \sum_{i=1}^N -\log p(y_i) \mathcal{N}(\mathbf{g}_i; \boldsymbol{\mu}_{y_i}, \Sigma_{y_i}). \end{aligned} \quad (10)$$

We assume $p(y_i) = \frac{1}{K}$ and Σ_{y_i} is diagonal for both simplicity and considering its application in existing research [10, 44].

Thus, the optimization objective during MPQ searching is

$$\mathcal{L} = \mathcal{L}_{cls} + \lambda \mathcal{L}_{inc} + \gamma \mathcal{L}_{comp}, \quad (11)$$

where \mathcal{L}_{cls} is the classification loss, \mathcal{L}_{inc} is the intra-class compactness loss and \mathcal{L}_{comp} is the complexity loss. λ and γ are the hyper-parameters to weight the corresponding loss in the optimization process.

4 EXPERIMENT

4.1 Settings

4.1.1 Datasets. The proxy (MPQ policy searching) datasets \mathcal{D}_{search} are CIFAR-10 [24] and StanfordCars [23]. CIFAR-10 has 10 categories, and each category has 5000 training samples and 1000 test samples. StanfordCars has 196 categories of cars; and the training set has 8144 training samples, and the test set has 8041 test samples. The target (model training) dataset \mathcal{D}_{train} is ISLVR-2012 [7] with 1000 categories, containing about 1.28M training samples and 50000 validation samples.

We search the MPQ policy on the training set of proxy datasets. The training samples of proxy datasets are used to search MPQ policies. For CIFAR-10, we adopt 4 pixel zero-paddings on each side and perform 32×32 random croppings. For StanfordCars, we perform 224×224 random croppings only. After searching, we finetune (quantize) the model with the searched policies on the target dataset. We use the basic data augmentation methods during finetuning and evaluate the final performance on the ISLVR-2012 validation set.

4.1.2 Models. We conduct the experiments on three representative models including the ResNet-{18, 50} [18] and the MobileNet [20]. Particularly, we use the standard architecture for ResNet.

4.1.3 Hyper-parameters. For ResNet and MobileNet, the bit-width candidates of weights and activations are $\mathbf{B}^w = \mathbf{B}^a = \{2, 3, 4, 6\}$ and $\mathbf{B}^w = \mathbf{B}^a = \{2, 3, 4, 5, 6\}$, respectively. Following the previous arts [12, 40, 46], the first and last layers are fixed to 8 bits.

For searching, we adopt the SGD optimizer, and the initial learning rate is set to 0.01 for 15 epochs. Empirically, we find the intra-class compactness regularization is not sensitive to the hyper-parameter and set $\lambda = 0.1$ for all proxy datasets, more details for $\lambda = 0.1$ can be found in the ablation study. We set the class margin $m = 0.3$ and $m = 0.01$ for CIFAR-10 and StanfordCars respectively while multiplying by the non-negative term in Equation 9. We fine-tune the hyperparameter γ in line with prior works on differentiable MPQ [3, 47]. A higher γ value corresponds to a less computation complexity policy to search for.

For finetuning (quantizing), we follow the basic quantization-aware training settings in LSQ [12] and LIMPQ [40]. Specifically, we use the full-precision model (trained on \mathcal{D}_{train}) as the initialization and adopt the SDG optimizer with Nesterov momentum [39] and the initial learning rate and weight decay are set to 0.04 and 2.5×10^{-5} , respectively. For fair comparisons, we do not apply knowledge distillation in main results (Sec. 4.2), and moreover, the large-margin regularization is not applied during finetuning. We use the cosine learning rate scheduler and finetune the model 90 epochs and the first 5 epochs are used as warm-up.

4.2 Comparisons with the State-of-the-Art

We compare our method with the SOTA quantization works on the classification task.

Table 1: Accuracy and efficiency results for ResNet. “Top-1 Q/FP” represents the Top-1 accuracy of quantized model and full-precision model. “MP” means mixed-precision quantization. “Cost” denotes the MPQ policy search time that is measured by GPU-hours. “*”: reproduces through the vanilla ResNet architecture [18]. “#”: the result of shortening the search epochs to half. “Ours-C”: denotes the MPQ policies search on CIFAR-10. “Ours-S”: denotes the MPQ policies search on StanfordCars. The lowest accuracy degradation results are bolded in each metric.

Method	W-bits	A-bits	Top-1 Q/FP (%)	BitOPs (G)	Cost (h)
ResNet18					
PACT	3	3	68.1 / 70.4	23.09	-
LSQ*	3	3	69.4 / 70.5	23.09	-
EdMIPS	3MP	3MP	68.2 / 69.6	-	9.8
GMPQ*	3MP	3MP	68.6 / 70.5	22.8	0.6
DNAS	3MP	3MP	68.7 / 71.0	25.38	-
FracBits	3MP	3MP	69.4 / 70.2	22.93	150.1
LIMPQ	3MP	3MP	69.7 / 70.5	23.07	3.3
Ours-C	3MP	3MP	70.0 / 70.5	23.07	0.9
Ours-S	3MP	3MP	69.6 / 70.5	23.06	0.3
PACT	4	4	69.2 / 70.4	35.04	-
LSQ*	4	4	70.5 / 70.5	35.04	-
DNAS	4MP	4MP	70.6 / 71.0	-	-
FracBits	4MP	4MP	70.6 / 70.2	34.7	151.3
FracBits#	4MP	4MP	70.3 / 70.2	34.7	76.8
LIMPQ	4MP	4MP	70.8 / 70.5	35.04	3.3
Ours-C	4MP	4MP	70.8 / 70.5	34.7	0.9
Ours-S	4MP	4MP	70.5 / 70.5	34.7	0.4
ResNet50					
HAQ	4MP	8	76.1 / 76.2	136.5	-
BP-NAS	4MP	4MP	76.7 / 77.5	64.4	35.6
FracBits	4MP	4MP	76.5 / 77.5	71.17	630.6
Ours-C	4MP	4MP	76.8 / 77.5	70.43	1.73
Ours-S	4MP	4MP	76.2 / 77.5	71.6	1.3

For fixed-precision works, we compare our method with PACT [6], PROFIT [31] and LSQ [12]. For MPQ works, we compare our method with DNAS [48], HMQ [15], HAQ [46], BP-NAS [51], FracBits [50], GMPQ [47], SDQ [21] and LIMPQ [40].

Specifically, since original LSQ and GMPQ use the Pre-Activation ResNet architecture, we re-implement them for fair comparisons through the vanilla ResNet [18].

4.2.1 ResNet. We show the mixed-3bits and mixed-4bits results of ResNet-{18, 50}, as listed in Table 1. We provide the full-precision accuracy to compare the *absolute accuracy degradation* between the full-precision and quantized model.

For ResNet18, under 3-bits level BitOPs constraints, “Ours-C” causes only 0.5% Top-1 accuracy degradation compared to the full-precision model, which is the lowest one among recent works. Under 4bits level BitOPs constraints, “Ours-C” achieves the highest Top-1 accuracy. Meanwhile, it achieves about 160× policy search speedup compared with FracBits. Thanks to the small data amounts

Table 2: Accuracy and efficiency results for MobileNetv1. “Top-1/5” represents Top-1 and top-5 accuracy respectively.

Method	W-bits	A-bits	Top-1/5 (%)	BitOPs (G)	Cost (h)
PACT	4	4	62.4 / 82.2	9.68	-
LSQ	3	3	68.3 / 88.1	5.8	-
HMQ	3MP	4MP	69.3 / -	-	-
FracBits	3MP	3MP	68.7 / 88.2	5.78	237.2
LIMPQ	3MP	3MP	69.5 / 89.1	5.78	3.4
Ours-C	3MP	3MP	69.9 / 89.3	6.28	1.0
Ours-S	3MP	3MP	69.6 / 89.2	6.13	0.8
PACT	6	4	67.5 / 87.8	14.13	-
PROFIT	4	4	69.1 / 88.4	9.68	-
LSQ	4	4	71.2 / 90.0	9.68	-
HAQ	4MP	4MP	67.5 / 87.9	-	35.6
HAQ	6MP	4MP	70.4 / 89.7	-	-
FracBits	4MP	4MP	71.4 / 90.0	9.63	250.2
LIMPQ	4MP	4MP	71.8 / 90.4	9.68	3.6
Ours-C	4MP	4MP	71.8 / 90.5	9.30	1.1
Ours-S	4MP	4MP	71.7 / 90.3	9.86	0.8

of StanfordCars, “Ours-S” uses only 8041 training samples to search a very competitive MPQ policy.

For ResNet50, we search 4bits level policies. One can see that our method achieves quite similar performance compared to gradient-based methods BP-NAS and FracBits while further reducing the search time significantly.

In summary, we observe that using CIFAR-10 as the proxy dataset causes the least accuracy degradation, probably due to the categories in CIFAR-10 are somewhat similar to ISLVR-2012. However, using the StanfordCars as the proxy dataset can further reduce the search time, since its samples are less than CIFAR-10. Overall, our method not only achieves a comparable accuracy as searching directly on ISLVR-2012, but also significantly improves the searching efficiency.

4.2.2 MobileNet. Table 2 summarizes the results of mixed-3bits and mixed-4bits on MobileNetv1.

For mixed-3bits searched on CIFAR-10, we observe our method both outperforms the existing SOTA mixed-precision work LIMPQ and fixed-precision work LSQ. In particular, our method arises a 1.8% absolute gain on Top-1 accuracy compared to LSQ, and 1.2% higher accuracy than FracBits. We further narrow the gap between the full-precision and quantized MobileNet. Please note that we are the first work to provide a 3-bits level MobileNet that almost achieves 70% Top-1 accuracy. For mixed-4bits searched on CIFAR-10, our method has up to 237× searching efficiency improvement compared to FracBits and up to 0.4% higher accuracy compared to the SOTA efficient MPQ approach LIMPQ.

For mixed-3bits and mixed-4bits searched on StanfordCars, they show 0.3% and 0.1% absolute Top-1 accuracy degradation compared to the CIFAR-10 but further save about 20% searching cost. This further proves that our method can still be very effective even if the proxy dataset (*i.e.*, all cars) has much lower class-similarity to the target dataset.




Figure 3: Complexity-accuracy trade-off for ResNet18 and MobileNet.

4.2.3 Discussion for Proxy Datasets. In this subsection, we observe that using CIFAR-10 as a proxy dataset can search for more well-performing MPQ policies better than StanfordCars. On the other hand, StanfordCars has higher search efficiency than CIFAR-10. We conjecture this is because the category of CIFAR-10 is more similar to the target dataset ISLVR-2012, and the data amounts of CIFAR-10 are more than that of StanfordCars. Meanwhile, we find that the performance loss of policies searched on StanfordCars is slightly larger than CIFAR-10 when the complexity constraint becomes tighter, *e.g.*, the mixed-3bits results for MobileNet.

Therefore, while it is feasible to search a well-performing MPQ policy by using an arbitrary proxy dataset, if the model requires more aggressive quantization, a proxy dataset with more class-level similarity compared to the target dataset could be considered to further improve the performance.

4.3 Complexity-Accuracy Trade-off

In Figure 3, we show the complexity-accuracy trade-off of LSQ [12], EdMIPS [3] and our method for ResNet18 and MobileNet. Unless otherwise specified, the proxy dataset used in our method is CIFAR-10.

For ResNet18, our method achieves significant performance gains compared to the mixed-precision approach EdMIPS. We even consistently have an absolute advantage of over 2% Top-1 accuracy.

For MobileNet, our method provides a very high accuracy improvement within the constraints of approximate complexity. Especially, our method improves 4.9% Top-1 accuracy compared to LSQ at 3G BitOPs constraint. Meanwhile, our method has a much fine-grained trade-off thanks to the mixed-precision quantization.

4.4 Ablation Study

In this subsection, we investigate: (a) the effectiveness of using a subset of \mathcal{D}_{train} as the searching dataset; (b) what happens when one adds SEAM to the baseline; (c) performance difference under various hyper-parameters settings.

4.4.1 Subset of ISLVR-2012. Although GMPQ has shown direct searching over the proxy dataset incurs severe performance degradation, there is no relevant literature to study the effect of using a subset of the target dataset (*e.g.*, ISLVR-2012) as the proxy dataset. To this end, we randomly sample 4% (roughly the same sample size as CIFAR-10) training data from ISLVR-2012 and use them to search a 3-bits level policy for ResNet18 without/with the proposed method.

Table 3: Results of ablation study. \mathcal{D}_{search} denotes the dataset used for MPQ policy searching. ISLVR-2012 (4%) indicates a subset of ISLVR-2012 with a sample of 4% of the full training set.

Proxy (\mathcal{D}_{search})	SEAM	Top-1 Quant (%)
ISLVR-2012 (4%)	✗	69.1
ISLVR-2012 (4%)	✓	69.8
CIFAR-10	✗	68.4
CIFAR-10	✓	70.0

As shown in Table 3, the subset of ISLVR-2012 without proposed method still has about 1% performance degradation compared to CIFAR-10 with proposed method. This is because the data distribution in the subset is significantly different from the full set. When the proposed method is enabled, this subset yields superior performance than StanfordCars. That further demonstrates the effectiveness of our method, and indicates that we can gain more performance by leveraging the class-similarity between proxy and target datasets.

4.4.2 Performance Improvement over Baseline. To show that MPQ benefits from discriminability of feature representations, we further add proposed method on EdMIPS [3] – a baseline MPQ approach. Specifically, EdMIPS is a conventional differentiable mixed-precision quantization approach, requiring consistency dataset of model training and policy searching. We directly apply proposed large-margin regularization term on it to search a MPQ policy. As shown in Figure 4, we observe that the proposed method can help the baseline to discover better MPQ policy.

Table 4: Effectiveness of proposed method SEAM upon EdMIPS.

Method	SEAM	\mathcal{D}_{search}	Cost (h)	Top-1 (%)	W&A bits
EdMIPS	✗	ISLVR-2012	9.5	65.9	2MP
EdMIPS	✓	ISLVR-2012	11.5	66.4	2MP

4.4.3 Effectiveness of λ . The setting of fixed λ is inspired by several regularization-based quantization-aware training studies [1, 17, 38]. Specifically, these studies enable their regularization term until after tens of training epochs. This delay is intended to **avoid optimization interference by the different loss terms**, ensuring that cross-entropy (CE) term dominates early training to optimize the parameters properly. Once the CE loss becomes small, the regularization term is added and plays a major role in optimization. In this paper, we direct use a small value for the regularization term λ to simulate the above optimization idea. We ablate this hyper-parameter in Table 5.

One can see that λ does need a relatively small value, which conforms with our optimization principle.

4.5 Bit-width Assignment Behavior

In Figure 4, we visualize the searched MPQ policies for the mixed-3bit ResNet18, ResNet50 and MobileNet. For ResNet, we clearly see




Figure 4: Bit-width assignment for mixed-3bit ResNet18, ResNet50 and MobileNetv1.

Table 5: Performance of different λ values.

Network	λ	Top-1 (%)	Cost (h)	W & A bits
ResNet18	0.02	69.5	0.8	3MP
ResNet18	0.1	70.0	0.9	3MP
ResNet18	0.5	69.2	0.9	3MP
ResNet18	2.5	69.0	0.9	3MP

Table 6: Results of finetuning the ResNet18 with an external teacher model ResNet101 (*: result from Table 1).

Method	Teacher	Top-1 (%)	BitOPs (G)
Full-precision	-	70.5	FP
Ours - w/o KD*	-	70.0 (-0.5)	23.07
GMPQ - KD	ResNet101	69.5 (-1.0)	22.8
SDQ - KD	ResNet101	70.2 (-0.3)	23.5
Ours - KD	ResNet101	70.7 (+0.2)	23.07

that almost the highest bit-width is given for the residual convolution layers. That is because these layers are more important for bypassing signals from shallow to deep layers [43], as well as having fewer parameters. For MobileNet, we find that higher bit-width is assigned to the Depthwise-Convolution (DW) layers than the Pointwise-Convolution (PW) layers, as the DW layer is typical less redundant [40].

4.6 Effectiveness of Knowledge Distillation

As knowledge distillation (KD) is a widely-used technique to regain the performance of quantized model [21, 31], we study the effectiveness of KD for our method here. Follow SDQ [21], we use a ResNet101 as the full-precision distillation teacher during the fine-tuning time. The distillation temperature is set to 1. We compare our method with GMPQ and SDQ at the 3-bits levels (about 23G BitOPs) search policies.

As shown in Tab 6, our approach achieves the highest performance when knowledge distillation is applied. In particular, compared to the state-of-the-art work SDQ under approximate complexity, our method attains an absolute accuracy improvement of 0.5%, indicating our method can search the optimal MPQ policy properly on a small-scale proxy dataset for the purpose of knowledge distillation.

5 CONCLUSION

In this work, we propose to search the MPQ policy on a small-scale proxy dataset for a model trained on a large-scale one. To bridge the inconsistent data distributions, we not only focus on optimizing the accuracy on the proxy dataset, but also enforce a large-margin of the searched MPQ policy should be met. We regard this as a *class-level* data exploitation for the limited proxy data, which is more data efficient than the *instance-level* data exploitation [47]. Our class-level data exploitation renders the search policies can compact the features in the same classes and separate the feature into different classes, which is a favorable and dataset-independent property. The experiments validate our idea, and we use only 4% of data to search for the high quality MPQ policies, achieving the same accuracy as searching directly on the large-scale dataset, and speeding up the MPQ searching process by up to 300 \times .

6 ACKNOWLEDGMENT

This work is supported in part by Shenzhen Science and Technology Program (Grant No. RCYX20200714114523079 and JCYJ2022081810-1014030). The authors would like to thank the anonymous reviewers for their valuable comments.

REFERENCES

- [1] Milad Alizadeh, Arash Behboodi, Mart van Baalen, Christos Louizos, Tijmen Blankevoort, and Max Welling. 2020. Gradient L1 regularization for quantization robustness. *arXiv preprint arXiv:2002.07520* (2020).
- [2] Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. Estimating or Propagating Gradients Through Stochastic Neurons for Conditional Computation. *CoRR abs/1308.3432* (2013). arXiv:1308.3432 <http://arxiv.org/abs/1308.3432>
- [3] Zhaowei Cai and Nuno Vasconcelos. 2020. Rethinking Differentiable Search for Mixed-Precision Neural Networks. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13–19, 2020*. Computer Vision Foundation / IEEE, 2346–2355. <https://doi.org/10.1109/CVPR42600.2020.00242>
- [4] Can Chen, Shuhao Zheng, Xi Chen, Erqun Dong, Xue (Steve) Liu, Hao Liu, and Dejing Dou. 2021. Generalized DataWeighting via Class-Level Gradient Manipulation. In *Advances in Neural Information Processing Systems 34: Annual Conference on Neural Information Processing Systems 2021, NeurIPS 2021, December 6–14, 2021, virtual*, Marc’Aurelio Ranzato, Alina Beygelzimer, Yann N. Dauphin, Percy Liang, and Jennifer Wortman Vaughan (Eds.). 14097–14109. <https://proceedings.neurips.cc/paper/2021/hash/75eb02f92fc30a8040bbd625af99f1-Abstract.html>
- [5] Weiha Chen, Peisong Wang, and Jian Cheng. 2021. Towards Mixed-Precision Quantization of Neural Networks via Constrained Optimization. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021*. IEEE, 5330–5339. <https://doi.org/10.1109/ICCV48922.2021.00530>
- [6] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. 2018. PACT: Parameterized Clipping Activation for Quantized Neural Networks. *CoRR abs/1805.06085* (2018). arXiv:1805.06085 <http://arxiv.org/abs/1805.06085>
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009)*, 20–25 June 2009, Miami, Florida, USA. IEEE Computer Society, 248–255. <https://doi.org/10.1109/CVPR.2009.5206848>
- [8] Zhen Dong, Zhewei Yao, Daiyaan Arfeen, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2020. HAWQ-V2: Hessian Aware trace-Weighted Quantization of Neural Networks. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc’Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/d77c703536718b95308130ff2e5cf9e-Abstract.html>
- [9] Zhen Dong, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2019. HAWQ: Hessian AWARE Quantization of Neural Networks With Mixed-Precision. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 – November 2, 2019*. IEEE, 293–302. <https://doi.org/10.1109/ICCV.2019.00038>
- [10] Sandrine Dudoit, Jane Fridlyand, and Terence P Speed. 2002. Comparison of discrimination methods for the classification of tumors using gene expression data. *Journal of the American statistical association* 97, 457 (2002), 77–87.
- [11] Ahmed T. Elthakeb, Pranoy Pilligundla, Fatemehsadat Miresghallah, Amir Yazdanbakhsh, and Hadi Esmailzadeh. 2020. ReLeQ : A Reinforcement Learning Approach for Automatic Deep Quantization of Neural Networks. *IEEE Micro* 40, 5 (2020), 37–45. <https://doi.org/10.1109/MM.2020.3009475>
- [12] Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakumar Apuswamy, and Dharmendra S. Modha. 2020. Learned Step Size quantization. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26–30, 2020*. OpenReview.net. <https://openreview.net/forum?id=rkgO66VKDS>
- [13] Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazhen Lin, Fengwei Yu, and Junjie Yan. 2019. Differentiable Soft Quantization: Bridging Full-Precision and Low-Bit Neural Networks. In *2019 IEEE/CVF International Conference on Computer Vision, ICCV 2019, Seoul, Korea (South), October 27 – November 2, 2019*. IEEE, 4851–4860. <https://doi.org/10.1109/ICCV.2019.00495>
- [14] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. 2020. Single Path One-Shot Neural Architecture Search with Uniform Sampling. In *Computer Vision – ECCV 2020 – 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XVI (Lecture Notes in Computer Science, Vol. 12361)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 544–560. https://doi.org/10.1007/978-3-030-58517-4_32
- [15] Hai Victor Habi, Roy H. Jennings, and Arnon Netzer. 2020. HMQ: Hardware Friendly Mixed Precision Quantization Block for CNNs. In *Computer Vision – ECCV 2020 – 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI (Lecture Notes in Computer Science, Vol. 12371)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 448–463. https://doi.org/10.1007/978-3-030-58574-7_27
- [16] Raia Hadsell, Sumit Chopra, and Yann LeCun. 2006. Dimensionality Reduction by Learning an Invariant Mapping. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2006)*, 17–22 June 2006, New York, NY, USA. IEEE Computer Society, 1735–1742. <https://doi.org/10.1109/CVPR.2006.100>
- [17] Tiantian Han, Dong Li, Ji Liu, Lu Tian, and Yi Shan. 2021. Improving Low-Precision Network Quantization via Bin Regularization. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021*. IEEE, 5241–5250. <https://doi.org/10.1109/ICCV48922.2021.00521>
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep Residual Learning for Image Recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2016, Las Vegas, NV, USA, June 27–30, 2016*. IEEE Computer Society, 770–778. <https://doi.org/10.1109/CVPR.2016.90>
- [19] Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. Distilling the Knowledge in a Neural Network. *CoRR abs/1503.02531* (2015). arXiv:1503.02531 <http://arxiv.org/abs/1503.02531>
- [20] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications. *CoRR abs/1704.04861* (2017). arXiv:1704.04861 <http://arxiv.org/abs/1704.04861>
- [21] Xijie Huang, Zhiqiang Shen, Shichao Li, Zechun Liu, Xianghong Hu, Jeffery Wicaksana, Eric P. Xing, and Kwang-Ting Cheng. 2022. SDQ: Stochastic Differentiable Quantization with Mixed Precision. In *International Conference on Machine Learning, ICML 2022, 17–23 July 2022, Baltimore, Maryland, USA (Proceedings of Machine Learning Research, Vol. 162)*, Kamalika Chaudhuri, Stefanie Jegelka, Le Song, Csaba Szepesvári, Gang Niu, and Sivan Sabato (Eds.). PMLR, 9295–9309. <https://proceedings.mlr.press/v162/huang22h.html>
- [22] Itay Hubara, Yury Nahshan, Yair Hanani, Ron Banner, and Daniel Soudry. 2021. Accurate post training quantization with small calibration sets. In *International Conference on Machine Learning*. PMLR, 4466–4475.
- [23] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 2013. 3D Object Representations for Fine-Grained Categorization. In *2013 IEEE International Conference on Computer Vision Workshops, ICCV Workshops 2013, Sydney, Australia, December 1–8, 2013*. IEEE Computer Society, 554–561. <https://doi.org/10.1109/ICCVW.2013.77>
- [24] Alex Krizhevsky, Geoffrey Hinton, et al. 2009. Learning multiple layers of features from tiny images. (2009).
- [25] Huantong Li, Xiangmiao Wu, Fanbing Lv, Daihai Liao, Thomas H Li, Yonggang Zhang, Bo Han, and Mingkui Tan. 2023. Hard Sample Matters a Lot in Zero-Shot Quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 24417–24426.
- [26] Jing Liu, Bohan Zhuang, Peng Chen, Chunhua Shen, Jianfei Cai, and Mingkui Tan. 2023. Single-path bit sharing for automatic loss-aware model compression. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2023).
- [27] Weiyang Liu, Yandong Wen, Zhiding Yu, and Meng Yang. 2016. Large-Margin Softmax Loss for Convolutional Neural Networks. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19–24, 2016 (JMLR Workshop and Conference Proceedings, Vol. 48)*, Maria-Florina Balcan and Kilian Q. Weinberger (Eds.). JMLR.org, 507–516. <http://proceedings.mlr.press/v48/liud16.html>
- [28] Zechun Liu, Kwang-Ting Cheng, Dong Huang, Eric P Xing, and Zhiqiang Shen. 2022. Nonuniform-to-uniform quantization: Towards accurate quantization via generalized straight-through estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 4942–4952.
- [29] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. 2019. Rethinking the Value of Network Pruning. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6–9, 2019*. OpenReview.net. <https://openreview.net/forum?id=rJlnB3C5Ym>
- [30] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. 2020. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*. PMLR, 7197–7206.
- [31] Eunhyeok Park and Sungjoo Yoo. 2020. PROFIT: A Novel Training Method for sub-4-bit MobileNet Models. In *Computer Vision – ECCV 2020 – 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VI (Lecture Notes in Computer Science, Vol. 12351)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 430–446. https://doi.org/10.1007/978-3-030-58539-6_26
- [32] Kanchana Ranasinghe, Muzammal Naseer, Munawar Hayat, Salman H. Khan, and Fahad Shahbaz Khan. 2021. Orthogonal Projection Loss. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021*. IEEE, 12313–12323. <https://doi.org/10.1109/ICCV48922.2021.01211>
- [33] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. 2016. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 779–788.
- [34] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems* 28 (2015).
- [35] Mark Sandler, Andrew G. Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. Computer Vision Foundation / IEEE

- Computer Society, 4510–4520. <https://doi.org/10.1109/CVPR.2018.00474>
- [36] Ramprasaath R. Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2020. Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization. *Int. J. Comput. Vis.* 128, 2 (2020), 336–359. <https://doi.org/10.1007/s11263-019-01228-7>
- [37] Xiangxi Shi, Jianfei Cai, Shafiq Joty, and Jiuxiang Gu. 2019. Watch it twice: Video captioning with a refocused video encoder. In *Proceedings of the 27th ACM International Conference on Multimedia*. 818–826.
- [38] Moran Shkolnik, Brian Chmiel, Ron Banner, Gil Shomron, Yury Nahshan, Alex M. Bronstein, and Uri C. Weiser. 2020. Robust Quantization: One Model to Rule Them All. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6–12, 2020, virtual*, Hugo Larochelle, Marc'Aurelio Ranzato, Raia Hadsell, Maria-Florina Balcan, and Hsuan-Tien Lin (Eds.). <https://proceedings.neurips.cc/paper/2020/hash/3948ead63a9f2944218de038d8934305-Abstract.html>
- [39] Ilya Sutskever, James Martens, George E. Dahl, and Geoffrey E. Hinton. 2013. On the importance of initialization and momentum in deep learning. In *Proceedings of the 30th International Conference on Machine Learning, ICML 2013, Atlanta, GA, USA, 16–21 June 2013 (JMLR Workshop and Conference Proceedings, Vol. 28)*. JMLR.org, 1139–1147. <http://proceedings.mlr.press/v28/sutskever13.html>
- [40] Chen Tang, Kai Ouyang, Zhi Wang, Yifei Zhu, Wen Ji, Yaowei Wang, and Wenwu Zhu. 2022. Mixed-Precision Neural Network Quantization via Learned Layer-Wise Importance. In *Computer Vision - ECCV 2022 - 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XI (Lecture Notes in Computer Science, Vol. 13671)*, Shai Avidan, Gabriel J. Brostow, Moustapha Cissé, Giovanni Maria Farinella, and Tal Hassner (Eds.). Springer, 259–275. https://doi.org/10.1007/978-3-031-20083-0_16
- [41] Chen Tang, Haoyu Zhai, Kai Ouyang, Zhi Wang, Yifei Zhu, and Wenwu Zhu. 2022. Arbitrary Bit-width Network: A Joint Layer-Wise Quantization and Adaptive Inference Approach. In *Proceedings of the 30th ACM International Conference on Multimedia*. 2899–2908.
- [42] Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-SNE. *Journal of machine learning research* 9, 11 (2008).
- [43] Andreas Veit, Michael J. Wilber, and Serge J. Belongie. 2016. Residual Networks Behave Like Ensembles of Relatively Shallow Networks. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5–10, 2016, Barcelona, Spain*, Daniel D. Lee, Masashi Sugiyama, Ulrike von Luxburg, Isabelle Guyon, and Roman Garnett (Eds.). 550–558. <https://proceedings.neurips.cc/paper/2016/hash/37bc2f75bf1bcfe8450a1a41c200364c-Abstract.html>
- [44] Weitao Wan, Cheng Yu, Jiansheng Chen, Tong Wu, Yuanyi Zhong, and Ming-Hsuan Yang. 2022. Shaping deep feature space towards gaussian mixture for visual classification. *IEEE transactions on pattern analysis and machine intelligence* 45, 2 (2022), 2430–2444.
- [45] Weitao Wan, Yuanyi Zhong, Tianpeng Li, and Jiansheng Chen. 2018. Rethinking Feature Distribution for Loss Functions in Image Classification. In *2018 IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2018, Salt Lake City, UT, USA, June 18–22, 2018*. Computer Vision Foundation / IEEE Computer Society, 9117–9126. <https://doi.org/10.1109/CVPR.2018.00950>
- [46] Kuan Wang, Zhijian Liu, Yujun Lin, Ji Lin, and Song Han. 2019. HAQ: Hardware-Aware Automated Quantization With Mixed Precision. In *IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2019, Long Beach, CA, USA, June 16–20, 2019*. Computer Vision Foundation / IEEE, 8612–8620. <https://doi.org/10.1109/CVPR.2019.00881>
- [47] Ziwei Wang, Han Xiao, Jiwen Lu, and Jie Zhou. 2021. Generalizable Mixed-Precision Quantization via Attribution Rank Preservation. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10–17, 2021*. IEEE, 5271–5280. <https://doi.org/10.1109/ICCV48922.2021.00524>
- [48] Bichen Wu, Yanghan Wang, Peizhao Zhang, Yuandong Tian, Peter Vajda, and Kurt Keutzer. 2018. Mixed Precision Quantization of ConvNets via Differentiable Neural Architecture Search. *CoRR abs/1812.00090* (2018). [arXiv:1812.00090](http://arxiv.org/abs/1812.00090) <http://arxiv.org/abs/1812.00090>
- [49] Shoukai Xu, Shuhai Zhang, Jing Liu, Bohan Zhuang, Yaowei Wang, and Minghui Tan. 2023. Generative Data Free Model Quantization with Knowledge Matching for Classification. *IEEE Transactions on Circuits and Systems for Video Technology* (2023).
- [50] Linjie Yang and Qing Jin. 2021. FracBits: Mixed Precision Quantization via Fractional Bit-Widths. In *Thirty-Fifth AAAI Conference on Artificial Intelligence, AAAI 2021, Thirty-Third Conference on Innovative Applications of Artificial Intelligence, IAAI 2021, The Eleventh Symposium on Educational Advances in Artificial Intelligence, EAAI 2021, Virtual Event, February 2–9, 2021*. AAAI Press, 10612–10620. <https://ojs.aaai.org/index.php/AAAI/article/view/17269>
- [51] Haibao Yu, Qi Han, Jianbo Li, Jianping Shi, Guangliang Cheng, and Bin Fan. 2020. Search What You Want: Barrier Panelty NAS for Mixed Precision Quantization. In *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part IX (Lecture Notes in Computer Science, Vol. 12354)*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm (Eds.). Springer, 1–16. https://doi.org/10.1007/978-3-030-58545-7_1
- [52] Edouard Yvinec, Arnaud Dapogny, Matthieu Cord, and Kevin Bailly. 2023. Powerquant: Automorphism search for non-uniform quantization. *arXiv preprint arXiv:2301.09858* (2023).
- [53] Shuchang Zhou, Zekun Ni, Xinyu Zhou, He Wen, Yuxin Wu, and Yuheng Zou. 2016. DoReFa-Net: Training Low Bitwidth Convolutional Neural Networks with Low Bitwidth Gradients. *CoRR abs/1606.06160* (2016). [arXiv:1606.06160](http://arxiv.org/abs/1606.06160) <http://arxiv.org/abs/1606.06160>