

Task-Specific Zero-shot Quantization-Aware Training for Object Detection

Changhao Li^{1*}
Atlanta, USA
cli911@gatech.edu

Xinrui Chen^{2*}
Shenzhen, China
cxr22@tsinghua.org.cn

Ji Wang^{3*}
Beijing, China
wangji20@tsinghua.org.cn

Kang Zhao⁴
Beijing, China
zhaok14@tsinghua.org.cn

Jianfei Chen^{4†}
Beijing, China
jianfeic@tsinghua.edu.cn

¹ School of Computational Science and Engineering, Georgia Institute of Technology

² Shenzhen International Graduate School, Tsinghua University

³ School of Software, Tsinghua University

⁴ Dept. of Comp. Sci. and Tech., Institute for AI, Tsinghua-Bosch Joint ML Center, Tsinghua University

Abstract

Quantization is a key technique to reduce network size and computational complexity by representing the network parameters with a lower precision. Traditional quantization methods rely on access to original training data, which is often restricted due to privacy concerns or security challenges. Zero-shot Quantization (ZSQ) addresses this by using synthetic data generated from pre-trained models, eliminating the need for real training data. Recently, ZSQ has been extended to object detection. However, existing methods use unlabeled task-agnostic synthetic images that lack the specific information required for object detection, leading to suboptimal performance. In this paper, we propose a novel task-specific ZSQ framework for object detection networks, which consists of two main stages. First, we introduce a bounding box and category sampling strategy to synthesize a task-specific calibration set from the pre-trained network, reconstructing object locations, sizes, and category distributions without any prior knowledge. Second, we integrate task-specific training into the knowledge distillation process to restore the performance of quantized detection networks. Extensive experiments conducted on the MS-COCO and Pascal VOC datasets demonstrate the efficiency and state-of-the-art performance of our method. Our project is publicly accessible¹.

1. Introduction

Object detection neural networks are integral to a wide array of computer vision applications, ranging from autonomous driving to surveillance systems [1, 38, 39, 45]. As the demand for deploying deep neural networks on resource-constrained devices continues to grow, quantization has emerged as a critical technique to reduce network size and computational complexity while maintaining performance [7, 11, 20, 52]. However, traditional quantization-aware training methods often require access to the entire original training data, which may become inaccessible if the data are very huge or subjected to privacy protection. [26, 43]. In these scenarios, Zero-shot Quantization (ZSQ) [3, 34, 40, 54, 57] presents a promising solution, enabling the quantization of neural networks without relying on original training data. These methods primarily train their model on synthetic data generated by inverting the network with randomly sampled labels.

Most prior work in this area has focused on classification networks. Specifically, they use model inversion to synthesize data by optimizing the gauss noise image through gradient descent. Their loss functions are often designated to classification tasks, which generate classification task-specific images. For example, GDFQ [50] introduced a knowledge-matching generator to synthesize label-oriented data using cross-entropy loss and batch normalization statistics (BNS) alignment. TexQ [6] emphasized the detailed texture feature distribution in real samples and devised texture calibration for labeled image generation. PSAQ-ViT [31] introduced a patch similarity aware

*Equal contribution

†Corresponding Author

¹<https://dfq-dojogithub.io/dfq-toolkit-webpage>

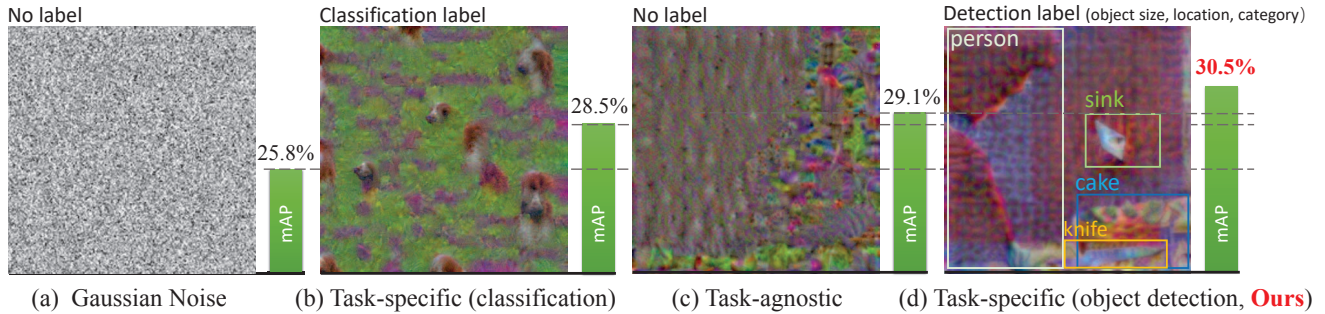


Figure 1. Comparative analysis of different synthetic images and their impact on zero-shot quantization-aware training with object detection network Mask R-CNN. It is observed that task-specific calibration set matters and results in better performance on the MS-COCO dataset. Inspired by this, we integrate novel task-specific manners into data synthesis and quantized detection network finetuning processes.

strategy to invert labeled images from Vision Transformers for quantization. These methods leverage synthetic data generated from the full-precision network to calibrate (*i.e.*, post-training quantization) or finetune the quantized network (*i.e.*, quantization-aware training) for accurate quantization. A more comprehensive discussion of data-driven quantization and ZSQ are presented in Appendix A.

Recently, ZSQ has been extended to downstream tasks such as object detection, but its application is limited due to inherent complexity. Classification tasks require only a randomly assigned category as the target label, but object detection demands that the target label comprise both the bounding box location and the classification label, making it difficult to determine. Consequently, existing approaches for detection networks drop the detection training loss and instead adopt a task-agnostic strategy for data generation and model quantization. For instance, PSAQ-ViT V2 [32] introduced an adaptive teacher–student strategy to generate task-agnostic images for finetuning the quantized model via knowledge distillation. Similarly, MimiQ [10] proposed inter-head attention similarity and applied head-wise structural attention distillation to align the attention maps of the quantized network with those of the full-precision teacher across downstream tasks. CLAMP-ViT [48] employed a two-stage approach, cyclically adapting between data generation and quantization. Although task-agnostic strategy enhances generalizability across different downstream tasks, it leads to a lack of task-specific bounding box size and location information in the object detection network, potentially resulting in suboptimal performance.

We argue that incorporating task-specific information into ZSQ can significantly increase its effect. By augmenting training loss with object categories and bounding box information, our method can outperform previous task-agnostic methods and, in some settings, may even achieve comparable results to networks trained with full real-data.

The proposed task-specific framework consists of two stages. In the generation stage, we introduce a novel bound-

ing box and category sampling strategy to synthesize a calibration set from a pre-trained detection network, which reconstructs the location, size and category distribution of objects within the data without any prior knowledge. In the quantization stage, we integrate the detection training loss into the distillation process to further amplify the efficacy of quantized detection network finetuning.

Extensive experiments on MS-COCO and Pascal VOC confirm the state-of-the-art performance of our method. For example, when quantizing YOLOv5-1 to 6-bit, we achieve a 1.7% mAP improvement over LSQ trained with full real data. Furthermore, tests on YOLO11 and Swin Transformer models show our approach surpasses task-agnostic ZSQ by 2-3% in mAP across various quantization settings. Specifically, our contributions are threefold:

1. Drawback of task-agnostic calibration is revealed.

We emphasize task-specific synthetic images for zero-shot quantization of object detection networks. By developing a task-specific approach that optimizes both data synthesis and finetuning, we unlock the full performance potential of quantized object detection networks.

2. Task-specific object detection images synthesis.

We propose a bounding box sampling method tailored for object detection networks to reconstruct object categories, locations, and sizes in synthetic samples without any prior knowledge.

3. Task-specific quantized network distillation.

We integrate object detection task-specific finetuning into quantized network distillation, effectively restoring the performance of quantized object detection networks.

2. Motivation

2.1. Preliminary on Quantization

Network quantization has emerged as a critical technique for reducing network size and computational cost while maintaining performance [7, 11, 20, 52].

Given a floating-point tensor w_{fp} (weights or activa-

tions) and quantization bit width b , a commonly used per-tensor symmetric quantizer LSQ [15] for both weights and activations to quantize the data w_{fp} can be defined as:

$$w_{int} = clip(\lfloor \frac{w_{fp}}{s}, -2^{b-1}, 2^{b-1} - 1 \rfloor), \quad (1)$$

$$\hat{w}_{fp} = w_{int} \times s. \quad (2)$$

Here, w_{int} denotes the quantized integer representation of the data, $\lfloor input \rfloor$ rounds the input to its nearest integer, and the step size s is a quantization parameter that is obtained with calibration set and updated during quantization-aware training. Notably, the calculation of quantization parameters requires access to real training data as the calibration set, rendering traditional quantization methods inapplicable when training data is unavailable.

2.2. Revisiting ZSQs for Object Detection

Task-specific calibration set matters. ZSQs synthesize images as the calibration set for quantization. Despite the proven efficacy of task-agnostic images in enhancing various downstream tasks within the ZSQ framework, it is intuitively evident that these images lose a considerable amount of task-specific information, which potentially compromise the performance of specific tasks. As illustrated in Fig. 1, we visualize four types of synthetic images utilized for zero-shot quantization-aware training: Gaussian noise, classification task-specific, task-agnostic, and object detection task-specific images. The Gaussian noise image serves as a baseline for comparison. Task-agnostic images capture the general features extracted by the network’s backbone, whereas task-specific images are derived through model inversion with corresponding labels, such as classification categories and object detection bounding box localizations. The visualization analysis in Fig. 1 reveals that task-specific images extract a richer set of features compared to other types of images, including object location, label, and size. Subsequently, we conduct a comparative analysis of synthetic images to explore their efficacy in zero-shot quantization-aware training, indicating that task-mismatched images lead to performance degradation, as demonstrated by the classification task-specific image in Fig. 1(b), whereas task-specific images enhance performance, as shown by the object detection task-specific image in Fig. 1(d). Therefore, task-specific calibration is crucial in the zero-shot quantization for detection tasks.

Challenges on task-specific ZSQ for detection. While task-specific zero-shot quantization has achieved remarkable success in classification tasks, extending it to object detection faces significant challenges.

First, the label sampling methods in the classification task for synthesizing labeled images cannot be directly extended to object detection. For classification networks, data

synthesis typically requires only a randomly sampled category ID label [6, 9, 46, 47, 59]. In contrast, for object detection, the location and size of objects within the samples remain unknown and elusive in zero-shot scenarios, making artificial reconstruction without ground-truth information extremely challenging. Besides, object detection datasets typically exhibit significant category imbalance, which is not captured by random category sampling methods designed for classification tasks (e.g., category-balanced ImageNet or CIFAR-10/CIFAR-100). Therefore, randomly sampling object categories, locations and sizes often results in implausible category distribution, relative positions and sizes, leading to unrealistic synthetic data. Furthermore, the task-specific finetuning strategy for detection networks using synthetic calibration data remains underexplored. The currently used logits alignment methods, which are designed for classification networks, may not be sufficient for the more complex object detection networks. This insufficiency makes it difficult to fully leverage the limited synthetic data efficiently, thereby hindering the performance improvement of quantized detection models.

3. Methodology

In this section, we provide an overview of the proposed framework in Fig. 2. It contains two stages: generating a task-specific calibration set and performing quantization-aware training (QAT) with task-specific distillation.

3.1. Preliminaries on Task-Agnostic Data Synthesis

The calibration set utilized in model quantization needs to reflect the inherent distribution of model. Zero-shot quantization seeks to generate a synthetic calibration set that matches the model’s distribution [3]. The synthetic calibration set can be derived through noise optimization [3, 58, 59], which is usually instantiated by distribution approximation [3, 54]. Existing zero-shot methods in object detection typically require generating a calibration set the same size as the training set (120k images for MS-COCO) [4]. In contrast, we only generate a small amount of calibration set to extract features of the data.

Given a batch of N inputs $x \in R^{N \times 3 \times H \times W}$, where each pixel is initialized from random Gaussian noise $x_{i,c,h,w} \sim \mathcal{N}(0, 1)$, and a pre-trained full-precision detection network $\phi(\theta)$, synthetic calibration set are obtained through optimizing the inputs to match the batch normalization statistics (BNS) [55] in convolutional neural networks (e.g., YOLO [49], Mask R-CNN [21]):

$$\min_x \mathcal{L}_{prior}(x) = \sum_{l=1}^L (||\mu^l(\theta, x) - \mu^l(\theta)||_2 + ||\sigma^l(\theta, x) - \sigma^l(\theta)||_2), \quad (3)$$

where $\mu^l(\theta)/\sigma^l(\theta)$ are mean/variance parameters stored in the l -th BN layer of $\phi(\theta)$ and $\mu^l(\theta, x)/\sigma^l(\theta, x)$ are

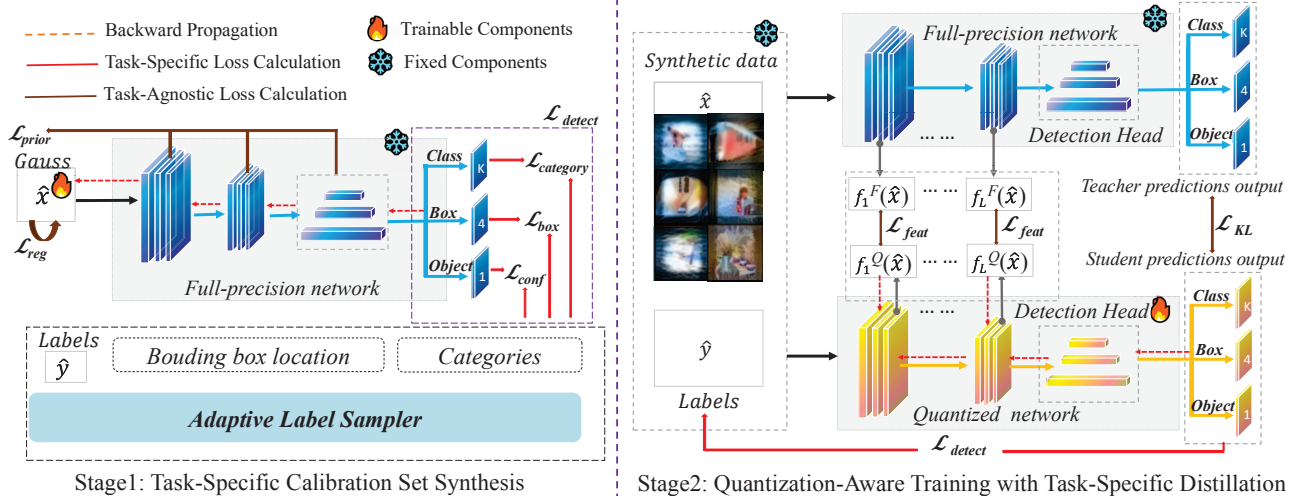


Figure 2. Overall architecture of our method. Our framework comprises: 1) constructing a task-specific condensed calibration set and 2) conducting quantization-aware training with task-specific distillation. See Section 3 for details.

mean/variance statistic calculated on inputs using $\phi(\theta)$. It enforces feature similarities at all levels by minimizing the distance between the feature map statistics for the synthesized image x and the real image. For vision transformer models (e.g., ViT [14], Swin Transformer [35]), since they only use layer normalization (LN) and do not store any runtime statistical information, we adopt the Patch Similarity Entropy loss (\mathcal{L}_{PSE}), as described in [31], as \mathcal{L}_{prior} to align inputs with the original data.

Besides the statistical alignment objective function, a regularization term consisting of the total variance and l_2 norm of the input image is always involved in the final loss function to steer images away from unrealistic images [37]:

$$\min_x \mathcal{L}_{reg}(x) = \alpha_{TV} \mathcal{L}_{TV}(x) + \alpha_{l_2} \|x\|_2^2, \quad (4)$$

where \mathcal{L}_{TV} promotes similarity between adjacent pixels by minimizing their Frobenius norm, consequently enhancing the smoothness, α_{TV} and α_{l_2} are hyper-parameters balancing the importance of two terms. Finally, we can regard task-agnostic data synthesis as a regularized minimization problem and optimize the following function:

$$\min_x \alpha_{prior} \mathcal{L}_{prior}(x) + \mathcal{L}_{reg}(x). \quad (5)$$

3.2. Stage I: Task-Specific Calibration Set Synthesis

Task-Specific Loss Calculation As discussed in Section 2.2, the object detection task requires labels that include the precise location and size of objects within an image, making artificial reconstruction extremely challenging. Consequently, previous zero-shot quantization approaches for object detection commonly utilize a task-agnostic loss, as described in Eq. 5, to generate images [10, 32, 48]. However, our experiments indicate that this approach leads to

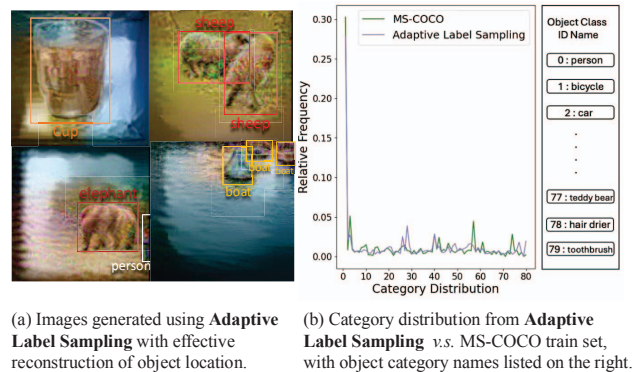


Figure 3. (a) Images generated by **Adaptive Label Sampling** on a YOLOv5 detector pre-trained on MS-COCO. (b) **Adaptive Label Sampling** can generate a category distribution frequency similar to MS-COCO in a zero-shot setting.

a significant loss of task-specific information, ultimately resulting in degraded performance on downstream tasks. Therefore, in this section, we introduce the training loss of the object detection network to optimize the sampled inputs, aiming to recover task-specific information.

Given a full-precision detection network $\phi(\theta)$ and a Gaussian-initialized input x , the standard form of the training loss in classification networks is expressed as $\mathcal{L}_{classify}(\phi(x), c)$, where the target label c is an integer generated through random sampling. However, in object detection tasks, label information is more complex, often encompassing the object's position and size. This can be formulated as $\mathcal{L}_{detect}(\phi(x), y)$, which typically consists of three components: a box category loss $\mathcal{L}_{category}$, a box dimension loss \mathcal{L}_{box} , and a grid location loss \mathcal{L}_{conf} . Through this formulation, we can reconstruct the category and coordinate information of bounding boxes in real images. Com-

binning with the task-agnostic loss, we derive the following task-specific data synthesis function:

$$\min_x \alpha_{prior} \mathcal{L}_{prior}(x) + \alpha_{detect} \mathcal{L}_{detect}(\phi(x), y) + \mathcal{L}_{reg}(x). \quad (6)$$

Adaptive Label Sampling To obtain task-specific information in a zero-shot scenario, we propose an adaptive label sampling method for label synthesis that effectively extracts essential bounding box coordinates and category information for object detection. Our approach relies solely on a pre-trained network and does not require any additional information (e.g., metadata, feature activations) or extra networks (e.g., pre-trained generative models).

Motivated by [56], which integrates soft labels into the data reconstruction process to better align synthetic data and labels, we propose an adaptive label sampling method, detailed as follows:

We start by randomly generating a label containing a single object, where the bounding box coordinates and category are uniformly sampled within a valid range (Details in Table 6). This label is then used as the ground truth target, and the input is optimized using Eq. 6. After a fixed number of iterations, we re-detect objects in the input using a pre-trained full-precision detection network. High-confidence regions from the teacher network’s feature map are added as new labels, while low-confidence regions are removed, ensuring that each image retains at least one label. The overall process is detailed in Algorithm 1.

The input image and target labels are updated alternately, progressively aligning with each other throughout the process. We also provide a visualization of this alignment in Fig. 5. With this sampling strategy, we eliminate the need for real detection labels. Besides, as presented in Fig. 3(b), the approach can eventually produce bounding box categories that closely resemble the actual distribution, while also reconstructing objects’ relative positions, sizes, and counts. This capability enables tackling the challenging zero-shot quantization-aware training for object detection.

Task-Specific Data Synthesis Through adaptive label sampling, we obtain labels that resemble the bounding box category and coordinate information in real images. We then fix the generated labels and optimize a newly Gaussian-initialized input toward these targets using the task-specific training loss introduced in Eq. 6. See Appendix B.1 for more details.

Compared to randomly sampling multiple bounding box coordinates and category labels, our adaptive label sampling method effectively utilizes the knowledge embedded in pre-trained object detection networks, resulting in higher-quality image synthesis. As illustrated in Fig. 3(a), images generated using our approach accurately capture object locations within the image, producing more realistic results.

Further studies, as discussed in Appendix E, demonstrate that our method yields clearer objects and more coherent layouts than those generated through random sampling.

3.3. Stage II: QAT with Task-Specific Distillation

In this section, we propose to minimize the knowledge gap between the full-precision pre-trained network (teacher) and the low-precision quantized network (student) through knowledge distillation.

Knowledge distillation [22] is a widely used technique for knowledge transfer. Previous studies [13, 29] have used it to enhance performance in classification tasks involving quantized CNNs and LLMs. However, the hidden states from backbone and prediction head of a full-precision pre-trained YOLO network contain much of the statistical information from real training data [55], which is challenging to fully utilize. To address this, we propose using feature-level distillation to align intermediate features and prediction-matching distillation to ensure consistency between the predictions of the quantized and pre-trained networks.

Prediction-matching Distillation As proposed in Section 3.2, our synthetic calibration set $\{(\hat{x}^i, \hat{y}^i)\}_{i=1}^N$ is the result of the network backpropagating through pre-defined labels, directly aligning predictions of the quantized network with the targets would lead to severe over-fitting issues. Therefore, we introduce the Kullback–Leibler (KL) divergence loss [28] between the predictions of the quantized network and the full-precision network as soft labels to align their outputs, thereby recovering the performance of the quantized network, which is formulated as:

$$\min_{\theta'} \mathcal{L}_{KD} = \frac{\tau^2}{N} \sum_{i=1}^N KL(z^F(\hat{x}_i; \theta), z^Q(\hat{x}_i; \theta')), \quad (7)$$

where $\{\hat{x}_i\}_{i=1}^N$ is a batch of the calibration set images, $z^F(\hat{x}_i; \theta)/z^Q(\hat{x}_i; \theta')$ are output predictions from full-precision / quantized network and τ is the distilling temperature. We denote parameters of full-precision / quantized network as θ/θ' .

Feature-level Distillation We extend the knowledge transfer approach to the feature level by introducing a feature distillation method that explicitly aligns intermediate features between the teacher and student. This significantly improves training stability in low-bit settings, where QAT at ultra-low bit widths often leads to rapid error accumulation. By ensuring feature consistency between the teacher and student through feature distillation, we effectively reduce error propagation throughout the training process.

In the quantization-aware training stage, given a batch of synthetic image $\{\hat{x}_i\}_{i=1}^N$, we impose the mean squared error constraints between the feature maps from teachers

Table 1. Comparison with real data QATs on YOLOv5/YOLO11 on MS-COCO validation set.

Method	Real Data	Num Data	Prec.	mAP / mAP50					
				YOLOv5-s	YOLOv5-m	YOLOv5-l	YOLO11-s	YOLO11-m	YOLO11-l
Pre-trained	✓	120k(full)	FP	37.4/56.8	45.4/64.1	49.0/67.3	47.0/65.0	51.5/70.0	53.4/72.5
LSQ	✓	120k(full)	W8A8	35.7/54.9	43.2/62.2	46.0/64.9	44.9/61.8	49.1/66.2	50.4/67.4
LSQ+	✓	120k(full)		35.4/54.6	43.3/62.4	46.3/64.9	45.1/61.8	49.6/66.7	50.9/67.7
LSQ	✓	2k		31.6/50.6	36.5/55.6	40.3/59.1	44.0/60.8	47.6/64.5	48.8/65.8
LSQ+	✓	2k		31.5/50.3	36.6/55.8	40.1/58.6	43.8/60.7	47.8/64.7	48.5/65.3
Ours	×	2k		35.8/55.0	43.6/62.3	47.3/65.6	45.6/62.3	50.0/66.5	51.8/68.4
LSQ	✓	120k(full)	W6A6	31.5/49.9	41.3/60.0	43.3/62.1	43.0/59.7	47.4/64.2	48.6/65.3
LSQ+	✓	120k(full)		32.3/50.9	41.3/60.3	43.4/62.3	43.2/59.8	47.6/64.3	48.9/65.8
LSQ	✓	2k		28.9/47.2	35.0/53.9	37.7/55.7	41.5/58.3	45.0/61.9	45.8/62.5
LSQ+	✓	2k		28.6/46.7	34.2/52.6	37.5/55.8	41.6/58.2	44.8/61.7	45.9/62.8
Ours	×	2k		32.7/51.4	41.0/59.7	45.1/63.3	43.0/59.3	47.1/63.2	48.4/64.6
LSQ	✓	120k(full)	W4A8	32.2/51.0	41.0/59.9	44.6/63.5	42.4/59.1	47.6/64.4	48.7/65.6
LSQ+	✓	120k(full)		32.3/51.1	41.2/60.1	44.4/63.2	42.7/59.3	47.8/64.8	49.4/66.3
LSQ	✓	2k		28.1/46.5	35.8/54.6	39.0/57.5	40.9/57.5	45.2/62.4	46.1/63.0
LSQ+	✓	2k		29.3/47.8	37.8/56.9	40.6/59.7	40.7/57.3	45.2/62.3	46.4/63.4
Ours	×	2k		33.0/52.5	42.6/61.7	46.2/64.7	42.6/58.9	47.7/64.1	49.4/65.7

and students. With L being the number of distilling network layers, the feature distillation loss \mathcal{L}_{feat} is expressed as:

$$\min_{\theta'} \mathcal{L}_{feat} = \frac{1}{NL} \sum_{i=1}^N \sum_{l=1}^L \|f_l^F(\hat{x}_i; \theta) - f_l^Q(\hat{x}_i; \theta')\|_2^2. \quad (8)$$

Task-Specific Quantization-Aware Training Previous works do not explicitly incorporate task-specific loss into their QAT training objectives [10, 32, 48]. While this allows their networks to remain flexible and adaptable to various downstream tasks (e.g. instance segmentation, object classification), we find that it often compromises performance on the target task — object detection in our case. To mitigate this, we introduce the task-specific training loss \mathcal{L}_{detect} during the QAT phase, enabling the quantized network to learn bounding box information directly from labels.

To this end, the total loss for quantization-aware training can be summarized as:

$$\min_{\theta'} \mathcal{L}^Q = \beta_{KL} \mathcal{L}_{KD} + \beta_{feat} \mathcal{L}_{feat} + \beta_{detect} \mathcal{L}_{detect}, \quad (9)$$

where β_{KL} , β_{feat} and β_{detect} are hyper-parameters to balance the three terms.

4. Experiments and Results

In this section, we validate the proposed task-specific scheme on MS-COCO 2017 [33] and Pascal VOC [16] datasets. Following LSQ [15], we apply symmetric quantization to both weights and activations. Through extensive experiments, we demonstrate that our method is effective




Figure 4. Visualization of images generated by YOLO11 and Transformer-backbone Mask R-CNN. More examples can be found in Appendix E.

on various architectures, including the YOLOv5 [51] series, YOLO11 [24] series, CNN-based Mask R-CNN [21], as well as Transformer-based Mask R-CNN [35] for object detection tasks. We further compare our approach with standard quantization-aware training baselines, including LSQ [15] and LSQ+ [2], using real data. Implementation details are provided in Appendix B, and ablation studies on different settings and components are presented in Appendix C. We also present visualizations of images generated by various models, including YOLO11 and Transformer-based Mask R-CNN, as shown in Fig. 4. More images and further analysis are provided in Appendix E.

4.1. Comparison with YOLO Networks

For YOLO object detection networks, we conduct experiments using the widely adopted YOLOv5-m/s/l models as well as the latest YOLO11-s/m/l. As baselines, we include competitive real-data QAT methods such as LSQ [15] and LSQ+ [2]. Extensive experiments show that our method outperforms both LSQ and LSQ+, which rely on full real-image training, by using only a small amount of ground truth label information. We compare our performance against LSQ and LSQ+, both trained on 120k real images from the MS-COCO dataset. In contrast, our approach utilizes only 2k ground truth labels for calibration set generation. The results are summarized in Table 1, with further details on the performance of the YOLOv5 series models at lower precision provided in Appendix C.5.

Bit-width Our method demonstrates strong performance across different bit-widths. Notably, when quantizing both weights and activation parameters to 8-bit precision, we find that our approach outperforms full-data LSQ/LSQ+ across all YOLOv5 and YOLO11 models by 0.3%–1.0%. Even when the quantization precision is further reduced to 6-bit or lower, our method still achieves results comparable to full-data LSQ/LSQ+, while significantly surpassing LSQ/LSQ+ with the same data amount by 2%–6%.

Network Size Larger networks tend to exhibit poor performance with existing quantization-aware training methods, particularly in low-bit-width cases. For instance, in the 6-bit case, LSQ+ applies to YOLOv5-s resulting in a 5.1% decrease in mAP compared to the pre-trained network, which achieves 5.6% with YOLOv5-l. In contrast, our approach yields only a 4.7% gap in mAP when quantizing YOLOv5-s to 6-bit precision, and the difference further reduces to 3.9% with YOLOv5-l.

Efficiency Our method performs quantization-aware training using a condensed synthetic detection calibration set, just 1/60 the size of the original, yet achieves superior performance compared to traditional QAT methods that require the full training dataset. This enables more efficient training, reducing computational costs and time while producing higher-quality low-precision models.

4.2. Comparison with Mask R-CNN Networks

CNN-Backbone Mask R-CNN We further conduct experiments on Mask R-CNN with CNN backbone. Results in Table 2 compare LSQ baselines trained on 120k/5k real samples from the MS-COCO/Pascal VOC datasets with our method, which uses only 2k/50 synthetic samples.

As presented, we achieve state-of-the-art results on both VOC and MS-COCO datasets, surpassing LSQ trained with full real data at 8-bit width. Specifically, on the smaller VOC dataset, our method surpasses LSQ trained with the entire dataset by 0.5% while using only 1/100 of the train-

Table 2. Comparison with real data QATs on CNN-based Mask R-CNN.

Dataset	Method	Real Data	Num Data	Precision	mAP
VOC	Pre-trained	✓	5k(full)	FP32	75.6
	LSQ	✓	5k(full)		72.4
	LSQ	✓	50	W8A8	70.9
	Ours	×	50		72.9
MS-COCO	Pre-trained	✓	120k(full)	FP32	38.1
	LSQ	✓	120k(full)		35.0
	LSQ	✓	2k	W8A8	32.9
	Ours	×	2k		35.2
	LSQ	✓	120k(full)		34.6
	LSQ	✓	2k	W4A8	32.3
	Ours	×	2k		34.6

Table 3. Comparison with real data QATs on Transformer-based Mask R-CNN on MS-COCO validation set.

Method	Real Data	Num Data	Prec.	mAP / mAP50	
				Swin-T	Swin-S
Pre-trained	✓	120k(full)	FP	46.0/68.1	48.5/70.2
LSQ	✓	120k(full)		45.9/68.0	48.1/69.7
LSQ	✓	2k	W8A8	44.4/65.9	47.0/68.6
Ours	×	2k		45.1/66.7	47.1/68.8
LSQ	✓	120k(full)		44.7/66.8	47.1/68.8
LSQ	✓	2k	W6A6	41.2/62.9	44.4/65.9
Ours	×	2k		42.0/63.0	45.1/65.8
LSQ	✓	120k(full)		45.5/64.7	47.8/69.4
LSQ	✓	2k	W4A8	43.3/65.2	45.9/67.3
Ours	×	2k		43.0/64.2	46.2/67.1

ing data, and exceeds LSQ trained on a similar dataset size by 2%. On the larger MS-COCO dataset, our approach outperforms LSQ trained with the full dataset by 0.2% using only 1/60 of the training data, and surpasses LSQ trained on a comparable dataset size by 2.3%. These findings highlight the robustness and strong generalization ability of our method across datasets of varying scales.

Transformer-Backbone Mask R-CNN In this section, we validate the proposed method on Transformer-based object detection networks. Specifically, we use Swin-T/S [35] as the backbone model, combined with the Mask R-CNN prediction head to form our model. The results on the MS-COCO dataset are presented in Table 3.

Our method can be seamlessly extended to Transformer-based object detection networks. Compared to LSQ trained with the same amount of data, our approach consistently outperforms it by 0.3%–0.8% across different bit-widths. While our method shows a slight performance drop of 0.8%–2.7% compared to LSQ trained on the full dataset, it significantly improves QAT efficiency by requiring substantially less data.

Table 4. Comparison with Task-Agnostic Methods on MS-COCO validation set. L_{detect} represents the task-specific training loss. All methods use 2k synthetic images for QAT.

Method	Prec.	mAP / mAP50			
		YOLO11-s	YOLO11-m	Swin-T	Swin-S
Pre-trained	FP	47.0/65.0	51.5/70.0	46.0/68.1	48.5/70.2
Ours	W8A8	45.6/62.3	50.0/66.5	45.1/66.7	47.1/68.8
	w/o L_{detect}	43.6/60.2	49.8/66.7	44.6/66.8	45.1/67.1
Ours	W4A8	42.6/58.9	47.7/64.1	43.0/64.2	46.2/67.1
	w/o L_{detect}	39.7/56.0	47.3/64.2	43.0/64.8	43.6/65.4
Ours	W6A6	43.0/59.3	47.1/63.2	42.0/63.0	45.1/65.8
	w/o L_{detect}	40.7/57.0	46.6/63.2	40.4/61.7	42.8/64.6

Table 5. Comparison with Data free Methods on MS-COCO validation set. All methods use 2k synthetic images for W6A6 QAT on YOLOv5-s. Info. denotes detailed information about labels including bounding box categories and coordinates. Distri. represents quantity distribution information about the labels per image. "-" indicates that the network diverges.

Method	Info.	Distri.	mAP	mAP50
Real Label	✓	✓	32.7	51.4
Gaussian noise	×	×	-	-
Tile (Out-of-distri.)	×	×	23.9	39.0
Tile (In-distri.)	×	✓	24.0	39.3
MultiSample (Out-of-distri.)	×	×	28.2	46.7
MultiSample (In-distri.)	×	✓	29.7	48.0
Ours (Adaptive Label Sampling)	×	×	32.0	50.0

4.3. Comparison with Task-Agnostic Methods

Previous zero-shot quantization works in object detection primarily employ task-agnostic training methods for both image generation and quantization-aware training [10, 32, 48]. We are the first to introduce task-specific loss at both stages. Through extensive experiments, as shown in the Table 4, we demonstrate that incorporating task-specific loss improves detection performance across various models, including YOLO11-s/m and Swin-T/S, as well as across different quantization levels, such as 6-bit and 8-bit precision.

This improvement is due to two key factors. First, task-specific training loss enriches the image generation process with detailed information, such as bounding box categories, size and coordinates, resulting in a distribution that more closely resembles real images. Second, during quantization-aware training, it enables the model to learn directly from labels, enhancing its ability to extract meaningful information from images.

4.4. Comparison with Data Free Methods

Furthermore, we explore a completely data-free scenario where no information about real images or labels is available and demonstrate the robustness of our adaptive label sampling method. The results are presented in Table 5.

First, we establish a weak baseline by using **Gaussian** noise as targets for QAT training and find that the quantized network fails to converge. This highlights the importance of synthetic data quality for QAT. For other proxy datasets, we primarily consider two types: *in-distribution* and *out-of-distribution*. *In-distribution* datasets assume that the number of bounding boxes per image is known, making the generated images statistically closer to real ones. In contrast, *out-of-distribution* datasets assume no prior knowledge about the original labels.

For baseline methods, we include **Tile**, which divides the image into uniform grids and randomly generates a unique bounding box in each grid, including its category and coordinate information, and **MultiSample**, which randomly samples multiple labels for each image. As shown in Table 5, QAT using images generated by our adaptive label sampling method surpasses the best *in-distribution* proxy dataset by 2.3% at 6-bit precision. Furthermore, when comparing our sampling method with images generated using real labels, we observe only a 0.7% performance gap. This demonstrates that our adaptive label sampling method can effectively extract real label information from the network, thus enabling the generation of high-quality synthetic data. Additional results across different precision levels are also provided in Appendix C.4.

5. Conclusions

For the first time, we revisit the current task-agnostic zero-shot quantization (ZSQ) methods for object detection tasks and identify the inherent limitations in their performance due to their task-agnostic nature. We propose a novel zero-shot quantization framework specifically tailored for object detection. The proposed framework consists of two key components: a novel task-specific synthesis process for generating the calibration set and a task-specific distillation process. The task-specific synthesis process leverages a bounding box and category sampling strategy to extract more relevant information from the original model, while the task-specific distillation process utilizes this information to fine-tune the quantized model, thereby significantly enhancing its performance. Extensive experiments demonstrate that our proposed method is both efficient and accurate, achieving performance comparable to traditional quantization-aware training (QAT) methods, such as Learned Step Size Quantization (LSQ), which rely on full real data, while significantly outperforming task-agnostic counterparts. This empowers zero-shot quantization with immense practical significance for object detection tasks.

Acknowledgments

This work was supported by the NSFC Project (No. 62376131).

References

- [1] Abhishek Balasubramaniam and Sudeep Pasricha. Object detection in autonomous vehicles: Status and open challenges. *arXiv preprint arXiv:2201.07706*, 2022. 1
- [2] Yash Bhalgat, Jinwon Lee, Markus Nagel, Tijmen Blankevoort, and Nojun Kwak. Lsq+: Improving low-bit quantization through learnable offsets and better initialization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 696–697, 2020. 6, 7, 12
- [3] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13169–13178, 2020. 1, 3, 14, 15
- [4] Akshay Chawla, Hongxu Yin, Pavlo Molchanov, and Jose Alvarez. Data-free knowledge distillation for object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3289–3298, 2021. 3
- [5] Xinrui Chen, Renao Yan, Junru Cheng, Yizhi Wang, Yuqiu Fu, Yi Chen, Tian Guan, and Yonghong He. Adeq: Adaptive diversity enhancement for zero-shot quantization. In *International Conference on Neural Information Processing*, pages 53–64. Springer, 2023. 12
- [6] Xinrui Chen, Yizhi Wang, Renao Yan, Yiqing Liu, Tian Guan, and Yonghong He. Texq: Zero-shot network quantization with texture feature distribution calibration. *Advances in Neural Information Processing Systems*, 36, 2024. 1, 3, 12
- [7] YH Chen, TJ Yang, J Emer, and V Sze Eyeriss. v2: A flexible accelerator for emerging deep neural networks on mobile devices., 2019, 9. DOI: <https://doi.org/10.1109/JETCAS>, pages 292–308, 2019. 1, 2
- [8] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018. 12
- [9] Kanghyun Choi, Deokki Hong, Noseong Park, Youngsok Kim, and Jinho Lee. Qimera: Data-free quantization with synthetic boundary supporting samples. *Advances in Neural Information Processing Systems*, 34:14835–14847, 2021. 3
- [10] Kanghyun Choi, Hye Yoon Lee, Dain Kwon, SunJong Park, Kyuyeun Kim, Noseong Park, and Jinho Lee. Mimiq: Low-bit data-free quantization of vision transformers with encouraging inter-head attention similarity. *arXiv preprint arXiv:2407.20021*, 2024. 2, 4, 6, 8, 12
- [11] Lei Deng, Guoqi Li, Song Han, Luping Shi, and Yuan Xie. Model compression and hardware acceleration for neural networks: A comprehensive survey. *Proceedings of the IEEE*, 108(4):485–532, 2020. 1, 2
- [12] Terrance DeVries and Graham W Taylor. Improved regularization of convolutional neural networks with cutout. *arXiv preprint arXiv:1708.04552*, 2017. 13
- [13] Xin Ding, Xiaoyu Liu, Yun Zhang, Zhijun Tu, Wei Li, Jie Hu, Hanting Chen, Yehui Tang, Zhiwei Xiong, Baoqun Yin, et al. Cbq: Cross-block quantization for large language models. *arXiv preprint arXiv:2312.07950*, 2023. 5
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. 4
- [15] Steven K Esser, Jeffrey L McKinstry, Deepika Bablani, Rathinakumar Appuswamy, and Dharmendra S Modha. Learned step size quantization. *arXiv preprint arXiv:1902.08153*, 2019. 3, 6, 7, 12, 13
- [16] Mark Everingham, Luc Van Gool, Christopher KI Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International journal of computer vision*, 88:303–338, 2010. 6, 12
- [17] Alex Finkelstein, Ella Fuchs, Idan Tal, Mark Grobman, Niv Vosco, and Eldad Meller. Qft: Post-training quantization via fast joint finetuning of all degrees of freedom. In *Computer Vision—ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 115–129. Springer, 2023. 12
- [18] Elias Frantar and Dan Alistarh. Optimal brain compression: A framework for accurate post-training quantization and pruning. *Advances in Neural Information Processing Systems*, 35:4475–4488, 2022. 12
- [19] Cong Guo, Yuxian Qiu, Jingwen Leng, Xiaotian Gao, Chen Zhang, Yunxin Liu, Fan Yang, Yuhao Zhu, and Minyi Guo. Squant: On-the-fly data-free quantization via diagonal hessian approximation. *arXiv preprint arXiv:2202.07471*, 2022. 12
- [20] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*, 2015. 1, 2
- [21] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 2961–2969, 2017. 3, 6
- [22] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 5
- [23] Yongkweon Jeon, Chungman Lee, and Ho-young Kim. Genie: show me the data for quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12064–12073, 2023. 14, 15
- [24] Glenn Jocher, Jing Qiu, and Ayush Chaurasia. Ultralytics YOLO, 2023. 6
- [25] Sangil Jung, Changyong Son, Seohyung Lee, Jinwoo Son, Jae-Joon Han, Youngjun Kwak, Sung Ju Hwang, and Changkyu Choi. Learning to quantize deep networks by optimizing quantization intervals with task loss. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 4350–4359, 2019. 12
- [26] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018. 1

- [27] Raghuraman Krishnamoorthi. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018. [12](#)
- [28] Solomon Kullback and Richard A Leibler. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86, 1951. [5](#)
- [29] Huantong Li, Xiangmiao Wu, Fanbing Lv, Daihai Liao, Thomas H Li, Yonggang Zhang, Bo Han, and Mingkui Tan. Hard sample matters a lot in zero-shot quantization. In *Proceedings of the IEEE/CVF conference on Computer Vision and Pattern Recognition*, pages 24417–24426, 2023. [5](#), [12](#)
- [30] Yuhang Li, Ruihao Gong, Xu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. Brcq: Pushing the limit of post-training quantization by block reconstruction. *arXiv preprint arXiv:2102.05426*, 2021. [12](#)
- [31] Zhikai Li, Liping Ma, Mengjuan Chen, Junrui Xiao, and Qingyi Gu. Patch similarity aware data-free quantization for vision transformers. In *European conference on computer vision*, pages 154–170. Springer, 2022. [1](#), [4](#)
- [32] Zhikai Li, Mengjuan Chen, Junrui Xiao, and Qingyi Gu. Psaq-vit v2: Toward accurate and general data-free quantization for vision transformers. *IEEE Transactions on Neural Networks and Learning Systems*, 2023. [2](#), [4](#), [6](#), [8](#), [12](#)
- [33] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part V 13*, pages 740–755. Springer, 2014. [6](#), [12](#)
- [34] Yuang Liu, Wei Zhang, and Jun Wang. Zero-shot adversarial quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1512–1521, 2021. [1](#)
- [35] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. [4](#), [6](#), [7](#)
- [36] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. [13](#)
- [37] Aravindh Mahendran and Andrea Vedaldi. Understanding deep image representations by inverting them. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5188–5196, 2015. [4](#)
- [38] Jiageng Mao, Shaoshuai Shi, Xiaogang Wang, and Hongsheng Li. 3d object detection for autonomous driving: A comprehensive survey. *International Journal of Computer Vision*, 131(8):1909–1963, 2023. [1](#)
- [39] Pawan Kumar Mishra and GP Saroha. A study on video surveillance system for object detection and tracking. In *2016 3rd international conference on computing for sustainable global development (INDIACom)*, pages 221–226. IEEE, 2016. [1](#)
- [40] Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1325–1334, 2019. [1](#)
- [41] Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1325–1334, 2019. [12](#)
- [42] Markus Nagel, Rana Ali Amjad, Mart Van Baalen, Christos Louizos, and Tijmen Blankevoort. Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pages 7197–7206. PMLR, 2020. [12](#)
- [43] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart Van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021. [1](#)
- [44] Markus Nagel, Marios Fournarakis, Rana Ali Amjad, Yelysei Bondarenko, Mart Van Baalen, and Tijmen Blankevoort. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295*, 2021. [12](#)
- [45] Kanyifeechukwu Jane Oguine, Ozioma Collins Oguine, and Hashim Ibrahim Bisallah. Yolo v3: Visual and real-time object detection model for smart surveillance systems (3s). In *2022 5th Information Technology for Education and Development (ITED)*, pages 1–8. IEEE, 2022. [1](#)
- [46] Biao Qian, Yang Wang, Richang Hong, and Meng Wang. Adaptive data-free quantization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7960–7968, 2023. [3](#)
- [47] Biao Qian, Yang Wang, Richang Hong, and Meng Wang. Rethinking data-free quantization as a zero-sum game. In *Proceedings of the AAAI conference on artificial intelligence*, pages 9489–9497, 2023. [3](#)
- [48] Akshat Ramachandran, Souvik Kundu, and Tushar Krishna. Clamp-vit: Contrastive data-free learning for adaptive post-training quantization of vits. In *European Conference on Computer Vision*, pages 307–325. Springer, 2024. [2](#), [4](#), [6](#), [8](#), [12](#)
- [49] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. [3](#)
- [50] Xu Shoukai, Li Haokun, Zhuang Bohan, Liu Jing, Cao Jiezhong, Liang Chuangrun, and Tan Mingkui. Generative low-bitwidth data free quantization. In *The European Conference on Computer Vision*, pages 1–17. Springer, 2020. [1](#), [12](#)
- [51] S Ultralytics. Yolov5: A state-of-the-art real-time object detection system. 2021. [6](#)
- [52] Huan Wang, Suhas Lohit, Michael N Jones, and Yun Fu. What makes a ” good ” data augmentation in knowledge distillation—a statistical perspective. *Advances in Neural Information Processing Systems*, 35:13456–13469, 2022. [1](#), [2](#)
- [53] Xiuying Wei, Ruihao Gong, Yuhang Li, Xianglong Liu, and Fengwei Yu. Qdrop: Randomly dropping quantization for extremely low-bit post-training quantization. *arXiv preprint arXiv:2203.05740*, 2022. [12](#)

- [54] Shoukai Xu, Haokun Li, Bohan Zhuang, Jing Liu, Jiezhong Cao, Chuangrun Liang, and Mingkui Tan. Generative low-bitwidth data free quantization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, pages 1–17. Springer, 2020. [1](#), [3](#)
- [55] Hongxu Yin, Pavlo Molchanov, Jose M Alvarez, Zhizhong Li, Arun Mallya, Derek Hoiem, Niraj K Jha, and Jan Kautz. Dreaming to distill: Data-free knowledge transfer via deep-inversion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8715–8724, 2020. [3](#), [5](#)
- [56] Zeyuan Yin, Eric Xing, and Zhiqiang Shen. Squeeze, recover and relabel: Dataset condensation at imagenet scale from a new perspective. *Advances in Neural Information Processing Systems*, 36, 2024. [5](#)
- [57] Edouard Yvinec, Arnaud Dapogny, Matthieu Cord, and Kevin Bailly. Spiq: Data-free per-channel static input quantization. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3869–3878, 2023. [1](#)
- [58] Xiangguo Zhang, Haotong Qin, Yifu Ding, Ruihao Gong, Qinghua Yan, Renshuai Tao, Yuhang Li, Fengwei Yu, and Xianglong Liu. Diversifying sample generation for accurate data-free quantization. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15658–15667, 2021. [3](#)
- [59] Yunshan Zhong, Mingbao Lin, Gongrui Nan, Jianzhuang Liu, Baochang Zhang, Yonghong Tian, and Rongrong Ji. 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*, pages 12339–12348, 2022. [3](#)