

# Modular Quantization-Aware Training for 6D Object Pose Estimation

Saqib Javed<sup>1†</sup>, Chengkun Li<sup>1</sup>, Andrew Price<sup>1</sup>, Yinlin Hu<sup>2</sup>, Mathieu Salzmann<sup>1</sup>  
<sup>1</sup>CVLab, EPFL <sup>2</sup>Magic Leap

Reviewed on OpenReview: <https://openreview.net/forum?id=1Iy0TEUou7>

## Abstract

Edge applications, such as collaborative robotics and spacecraft rendezvous, demand efficient 6D object pose estimation on resource-constrained embedded platforms. Existing 6D object pose estimation networks are often too large for such deployments, necessitating compression while maintaining reliable performance. To address this challenge, we introduce Modular Quantization-Aware Training (MQAT), an adaptive and mixed-precision quantization-aware training strategy that exploits the modular structure of modern 6D object pose estimation architectures. MQAT guides a systematic graduated modular quantization sequence and determines module-specific bit precisions, leading to quantized models that outperform those produced by state-of-the-art uniform and mixed-precision quantization techniques. Our experiments showcase the generality of MQAT across datasets, architectures, and quantization algorithms. Additionally, we observe that MQAT quantized models can achieve an accuracy boost ( $> 7\%$  ADI-0.1d) over the baseline full-precision network while reducing model size by a factor of  $4\times$  or more. [https://saqibjaved1.github.io/MQAT\\_](https://saqibjaved1.github.io/MQAT_)

## 1 Introduction

6D pose estimation is the process of estimating the 6 degrees of freedom (attitude and position) of a rigid body and has emerged as a crucial component in numerous situations, particularly in robotics applications such as automated manufacturing (Pérez et al., 2016), vision-based control (Singh et al., 2022), collaborative robotics (Vicentini, 2020) and spacecraft rendezvous (Song et al., 2022). However, such applications typically must run on embedded platforms with limited hardware resources.

These resource constraints often disqualify current state-of-the-art methods, such as ZebraPose (Su et al., 2022), SO-Pose (Di et al., 2021), and GDR-Net (Wang et al., 2021), which employ a two stage approach (detection followed by pose estimation) and thus entail a large memory footprint. By contrast, single-stage methods (Thalhammer et al., 2021; Peng et al., 2019; Song et al., 2020; Hu et al., 2021b; Wang et al., 2022; Chen et al., 2019; Rad & Lepetit, 2017; Hodañ et al., 2020) offer a more pragmatic alternative, yielding models with a good accuracy-footprint tradeoff. Nevertheless, they remain too large for deployment on edge devices.

To address this challenge, CA-SpaceNet (Wang et al., 2022) applies a *uniform* quantization approach to reduce the network memory footprint at the expense of a large accuracy loss; all network layers are quantized to the same bit width, except for the first and last layer.

In principle, *mixed-precision* quantization methods (Cai & Vasconcelos, 2020; Dong et al., 2020; Tang et al., 2022) could demonstrate similar compression with better performance, but they tend to require significant effort and GPU hours to determine the optimal bit precision for each layer. Furthermore, neither mixed-precision nor uniform quantization methods consider the importance of the order in which the network weights or layers are quantized, as searching for the optimal order is combinatorial in the number of, e.g., network layers.

<sup>†</sup>Corresponding author: [saqib.javed@epfl.ch](mailto:saqib.javed@epfl.ch)

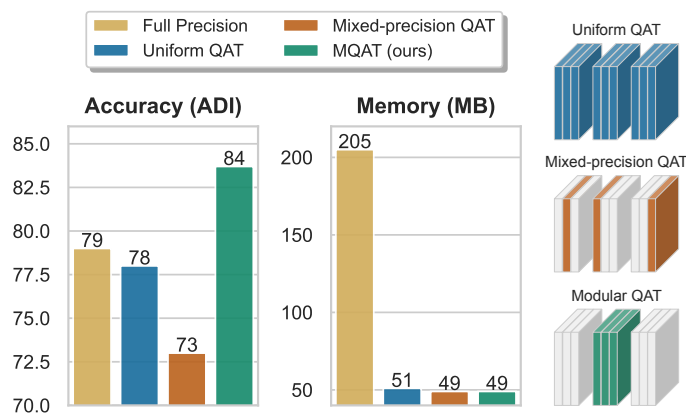


Figure 1: **Summary of this Work.** In contrast to uniform and mixed-precision quantization, MQAT accounts for the modularity of typical 6D object pose estimation frameworks. Uniform QAT quantizes an entire network simultaneously and uniformly; in mixed-precision QAT, layers are quantized to varying bit precisions regardless of their position; in contrast, MQAT applies quantization to network modules in a proposed order, with each module assigned the optimal bit precision. MQAT not only reduces the memory footprint of the network but can result in an accuracy boost that neither uniform nor mixed-precision quantization have demonstrated. The results shown in the figure represent a comparison of different quantization methods applied to the WDR network, evaluated on the SwissCube dataset.

In this work, we depart from such conventional Quantization-Aware Training (QAT) approaches and leverage the inherent modular structure of modern 6D object pose estimation architectures, which typically encompass components such as a backbone, a feature aggregation module, and prediction heads. Specifically, we introduce a Modular Quantization-Aware Training (MQAT) paradigm that relies on a graded quantization strategy, where the modules are quantized and fine-tuned in a mixed-precision way, following a recommended sequential order (i.e., quantization flow) based on their sensitivity to quantization. As the number of modules in an architecture is much lower than that of layers, it is possible to confirm a quantization flow is optimal.

Our experiments evidence that MQAT is particularly well-suited to 6D object pose estimation, consistently outperforming the state-of-the-art quantization techniques in terms of accuracy for a given memory consumption budget, as shown in fig. 1. We demonstrate the generality of our approach by applying it to different single-stage architectures, WDR (Hu et al., 2021b), CA-SpaceNet (Wang et al., 2022); different datasets, SwissCube (Hu et al., 2021b), Linemod (Hinterstoisser et al., 2012), and Occlusion Linemod (Brachmann et al., 2014); and using different quantization strategies, INQ (Zhou et al., 2017) and LSQ (Esser et al., 2020), within our modular strategy. We further show that our method also applies to two-stage networks, such as ZebraPose (Su et al., 2022), on which it outperforms both uniform and mixed-precision quantization. Furthermore, we extend MQAT to the task of object detection, further validating its efficacy in the computer vision domain. The results of applying MQAT to object detection are presented in appendix A.2, demonstrating its potential for broader applicability.

To summarize, our main contributions are as follows:

- For neural network compression, we are the first to identify and exploit the current trend of modular architectures. We propose Modular Quantization-Aware Training (MQAT), a novel mixed-precision Quantization-Aware Training (QAT) algorithm.
- We further demonstrate that to achieve the best results, the sequence of quantizing modules is significant and thus propose a recommended quantization order. We show that this order is optimal for the quantization of a 3 module 6D pose estimation network (Hu et al., 2021b).
- With MQAT, we show substantial accuracy gains over competitive QAT methods. Furthermore, unlike other QATs, MQAT significantly surpasses full-precision accuracy in the specific case of single-stage 6D pose estimation networks (Hu et al., 2021b; Wang et al., 2022).
- We validate MQAT across different datasets, neural network architectures, underlying quantization algorithms, and tasks, showcasing its adaptability and effectiveness in different settings.

## 2 Related Work

In this section, we survey recent advances in RGB-based 6D object pose estimation, outlining key architectures and their contributions to the field. We then explore the developments in quantization-aware training (QAT), particularly in relation to modular neural network designs. This review sets the groundwork for our proposed Modular Quantization-Aware Training framework, which is inspired by these advances and addresses their limitations.

### 2.1 6D Object Pose Estimation

**Single-Stage Direct Estimation.** PoseCNN (Xiang et al., 2018) was one of the first methods to estimate 6D object pose using a deep neural network. SSD6D (Liu et al., 2016; Kehl et al., 2017) and URSONet (Proença & Gao, 2020) used a discretization of the rotation space to form a classification problem instead of a regression one.

**Single-Stage with PnP.** In general, a better performing strategy for convolutional architectures consists of training a network to predict 2D-to-3D correspondences instead of the pose. The pose is then obtained via a RANdom SAMple Consensus (RANSAC) / Perspective-n-Point (PnP) 2D-to-3D correspondence fitting process. These methods typically employ a backbone, a feature aggregation module, and one or multiple heads. (Rad & Lepetit, 2017; Tekin et al., 2018) estimate these correspondences in a single global fashion, whereas (Peng et al., 2019; Jafari et al., 2018; Hu et al., 2019; Zakharov et al., 2019; Markus Oberweger, 2018; Chen et al., 2019) aggregate multiple local predictions to improve robustness. To improve performance in the presence of large depth variations, a number of works (Thalhammer et al., 2021; Hu et al., 2021b; Wang et al., 2022) use an FPN (Lin et al., 2016) to exploit features at different scales.

**Multi-Stage** The current state-of-the-art pose estimation frameworks incorporate a pipeline of networks that perform different tasks. In the first stage network, the target is localized and a Region of Interest (RoI) is cropped and forwarded to the second stage network. This isolates the position estimation task from the orientation estimation one and further provides the orientation estimation network with an RoI containing only object features. The second stage orientation estimation network can then more easily fit to the target object.

Multi-stage with PnP approaches have generated good results (Su et al., 2022; Di et al., 2021; Wang et al., 2021; Li et al., 2019b; Labbé et al., 2020), but some recent works avoid traditional PnP and RANSAC methods. PViT-6D (Stapf et al., 2023) uses an off-the-shelf object detector to crop the target and then reframes pose estimation as a direct regression task, using Vision Transformers to enhance accuracy. CheckerPose (Lian & Ling, 2023) improves robustness by progressively refining dense 3D keypoint correspondences with a graph neural network. MRC-Net (Li et al., 2024) introduces a two-stage process, combining pose classification and rendering with a multi-scale residual correlation (MRC) layer to capture fine-grained pose details. GigaPose (Nguyen et al., 2024) achieves fast and robust pose estimation through CAD-based template matching and patch correspondences. FoundationPose (Wen et al., 2024) unifies model-based and model-free setups using neural implicit representations and contrastive learning, achieving strong generalization without fine-tuning.

Multi-stage frameworks tend to yield more accurate results. However, they also have much larger memory footprints as they may include one object classifier network; one object position/RoI network; and  $N$  object pose networks. For hardware-restricted scenarios, a multi-stage framework may thus not be practical. Even for single-stage networks, additional compression is required (Blalock et al., 2020).

### 2.2 Quantization-Aware Training

Neural network quantization reduces the precision of parameters such as weights and activations. Existing techniques fall into two broad categories. **Post-Training Quantization (PTQ)** quantizes a pre-trained network using a small calibration dataset and is thus relatively simple to implement (Nagel et al., 2020; Li et al., 2021; Frantar & Alistarh, 2022; Zhao et al., 2019; Cai et al., 2020; Nagel et al., 2019; Shao et al.,

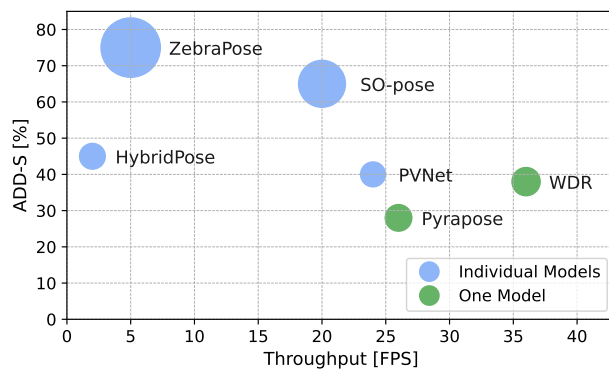


Figure 2: **Performance Comparison on Occluded-LINEMOD.** The marker size is proportional to the memory footprint. *Individual Models* refers to methods training one model for each object. *One Model* refers to methods training a single model for all objects.

2024; Lin et al., 2024; Chee et al., 2023). **Quantization-Aware Training (QAT)** retrains the network during the quantization process and thus better preserves the model’s full-precision accuracy. QAT methods include uniform QAT (Esser et al., 2020; Zhou et al., 2017; Bhalgat et al., 2020; Yamamoto, 2021) and mixed-precision QAT (Cai & Vasconcelos, 2020; Dong et al., 2020; Tang et al., 2022; Dong et al., 2019; Chen et al., 2021; Yao et al., 2020). As accuracy can be critical in robotics applications relying on 6D object pose estimation, we focus on QAT.

**Uniform QAT** methods quantize every layer of the network to the same precision. In Incremental Network Quantization (INQ) (Zhou et al., 2017), this is achieved by quantizing a uniform fraction of each layers’ weights at a time and continuing training until the next quantization step. Quantization can be achieved in a structured manner, where entire kernels are quantized at once, or in an unstructured manner. In contrast to INQ, Learned Step-size Quantization (LSQ) (Esser et al., 2020) quantizes the entire network in a single action. To this end, LSQ treats the quantization step-size as a learnable parameter. The method then alternates between updating the weights and the step-size parameter. Group-Wise Quantization (GWQ) (Yang et al., 2023) shows comparable accuracy while claiming reduced computational costs at inference time by assigning the same scale factor to different layers in the network.

Conversely, **Mixed-precision QAT** methods treat each network layer uniquely, aiming to determine the appropriate bit precision for each one. In HAWQ (Dong et al., 2019; 2020; Yao et al., 2020), the network weights’ Hessian is leveraged to assign bit precisions proportional to the gradients. In (Cai & Vasconcelos, 2020), the mixed precision is taken even further by applying a different precision to different kernels within a single channel. Finally, NIPQ (Shin et al., 2023) replace the standard straight-through estimator ubiquitous in quantization with a proxy, allowing the quantization truncation and scaling hyper-parameters to be included in the gradient descent. Mixed-precision QAT is a challenging task; existing methods remain computationally expensive for modern deep network architectures.

### 2.3 Quantization and Modular Deep Learning

In recent years, deep network architectures have increasingly followed a modular paradigm, owing to its advantages in model design and training efficiency (Jacobs et al., 1991; Pfeiffer et al., 2023; Ansell et al., 2021; Hu et al., 2021a; Pfeiffer et al., 2020). This approach leverages reusable modules, amplifying the flexibility and adaptability of neural networks and fostering parameter-efficient fine-tuning. In the quantization domain, several studies have also underscored the importance of selecting the appropriate granularity to bolster model generalization (Li et al., 2021) and enhance training stability (Zhang et al., 2023).

Furthermore, in the 6D object pose estimation domain, the demonstration of applying quantization remains limited, with none of the aforementioned quantization techniques addressing the specific task or underlying network architecture. CA-SpaceNet (Wang et al., 2022) applied LSQ quantization to their proposed architecture using three different quantization configurations. However they do so homogeneously, not accounting for the varying sensitivity to quantization throughout the network.

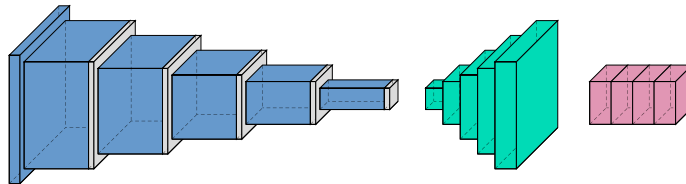


Figure 3: **Representative 6D Object Pose Estimation Network** with  $K = 3$  modules. From left to right, we denote them as *backbone*, *feature aggregation*, and *heads*.

To our knowledge, no existing research has advocated a systematic methodology for executing modular quantization-aware training, let alone studied the impact of quantization order on modular architectures. Thus, in this work, we aim to bridge this gap by introducing a comprehensive methodology for modular QAT, tailored but not limited to 6D object pose estimation.

### 3 Method

In this section, we first describe the general type of network architecture we consider for compact 6D object pose estimation and then introduce our Modular Quantization-Aware Training (MQAT) method.

#### 3.1 Modular Architecture for 6D Pose Estimation

As discussed in section 2, multi-stage networks (Su et al., 2022; Di et al., 2021; Wang et al., 2021; Labbé et al., 2020; Li et al., 2019b) tend to induce large memory footprints and large latencies, thus making poor candidates for hardware-restricted applications. This is further evidenced in fig. 2, where we compare a number of 6D object pose estimation architectures’ memory footprint, throughput and accuracy on the O-LINEMOD dataset. While demonstrating admirable accuracy, the size and latency of ZebraPose (Su et al., 2022) and SO-pose (Di et al., 2021) preclude their inclusion in hardware-restricted platforms. On this basis, we therefore focus on single-stage<sup>1</sup> 6D object pose estimation networks (Thalhammer et al., 2021; Wang et al., 2022; Tekin et al., 2018; Peng et al., 2019; Hu et al., 2021b).

In general, 6D pose estimation networks contain multiple modules, with each module performing a distinct task. A typical encoder-decoder-head 6D pose estimation architecture is illustrated in fig. 3. Explicitly:

1. A *backbone* or *encoder* module extracts features from an image at different depths. Common module examples include ResNets, DarkNets and MobileNets.
2. A *feature aggregator* or *decoder* aggregates or interprets the latent feature space. Common module examples include feature pyramid networks, attention layers, and simple upscaling CNNs (i.e., latter half of U-Nets).
3. A *head* module performs specific tasks or regresses specific quantities such as location vectors and correspondences. Commonly correspondence and location heads are fully-connected layers but may also be attention-based or CNNs for other tasks such as image segmentation.

Architectures for various tasks similarly consist of clearly identifiable modules. Therefore, we also demonstrate the effectiveness of our proposed MQAT method for object detection, as detailed in the Appendix.

#### 3.2 MQAT Algorithm

We introduce MQAT in algorithm 1 where we have also defined the notations. We begin by introducing the critical components for our method.

**Quantization Baseline** (*Lines 1:7*). MQAT starts with the aggressive quantization of individual modules to establish favourable bit precisions. For a modular network with  $K$  modules, we conduct  $K$  independent

<sup>1</sup>For completeness, we also demonstrate that our quantization method applies to two-stage networks (e.g., ZebraPose).

**Algorithm 1** MQAT Algorithm

**Input:** Training data, Quantization  $Q$ , Model  $M$  with modules  $\mathbf{B} = [B_1, B_2, \dots, B_K]$ , Bit-width search  $q = [2, 4, 8, \dots]$ , Accuracy metric  $ac(M)$

```

1: for  $k = 1, 2, \dots, K$  do
2:    $M_k^2 \leftarrow Q_k^2(M)$ 
3:   if  $ac(M_k^2) > ac(M) \wedge ac(M_k^2) > ac(M_{best})$  then
4:      $M_{best} \leftarrow M_k^2$ 
5:      $k_{best} \leftarrow k$ 
6:   end if
7: end for
8:  $flow = []$ 
9: if  $M_{best}$  then
10:   $M \leftarrow M_{best}$ 
11:   $flow \leftarrow index(SORT\_SIZE(\mathbf{B}))$ 
12:   $flow.pop(k_{best})$ 
13: else
14:   $flow \leftarrow index(SORT\_SIZE(\mathbf{B}))$ 
15: end if
16:  $\rho \leftarrow ILP(M, q)$ 
17:  $i \leftarrow 0$ 
18: while  $i < len(flow)$  do
19:   $M^Q \leftarrow Q_{flow(i)}^{\rho(i)}(M)$ 
20:   $retrain(M^Q)$ 
21: end while

```

**Output:** Modular quantized model ( $M^Q$ ).

**Variables:**

$flow$	Sequence of quantization order of modules.
$k_{best}$	Index of the quantized module which increased performance.
$K$	Number of modules.
$M_{best}$	Highest accuracy model containing a quantized module.
$M_k^q$	Model with only module $k$ quantized to $q$ bits.
$M^Q$	Model with modules quantized to different bit precisions.
$\rho$	List of bit precisions for each module.
$Q_c^j$	Quantization applied to module $B_c$ with $j$ -bits.
$SORT\_SIZE(\mathbf{B})$	Sort list of modules $\mathbf{B}$ based on their size.

2 bit quantizations for each module,  $B_k$ . The module is retained as quantized if it results in an improved accuracy for the entire model, (i.e.,  $M_k^q$  outperforms  $M$ ). This establishes a new baseline from which the remaining module bit precision optimization can be performed.

**Quantization Flow** (*Lines 8:15*). The sequence of module quantization—the *quantization flow*— is critical as quantizing the modules of a network is not commutative; we prioritize starting with modules that do not compromise accuracy, as quantization errors introduced early on are typically not mitigated by later steps. Moreover, we also observe that quantization-related noise can lead to weight instability (Défossez et al., 2022; Shin et al., 2023; Peters et al., 2023), hindering the performance of the quantized network. If no quantized module yielded an improved accuracy in the quantization baseline step, we proceed with quantizing the module with the lowest number of parameters first. Modules with a greater number of parameters will have more flexibility to adapt to aggressive quantization. The output of this algorithmic step is the recommended *quantization flow*, which is then passed to the next step.

**Optimal Bit Precision** (*Lines 16:21*). In MQAT, the optimal bit precision for each module (excepting a quantization baseline identified  $M_{best}$ ,  $k_{best}$ ) is ascertained through a process of constrained optimization, *Integer Linear Programming* (ILP), drawing inspiration from Yao et al. (2020). However, our methodology distinguishes itself by offering a lower degree of granularity. Central to our strategy is the uniform quantization of all layers within a given module to an identical bit-width. This design choice not only simplifies the computational complexity but also significantly enhances the hardware compatibility, an essential consideration for efficient real-world deployment.

To achieve this, we introduce the *importance metric* for modular quantization. This metric is conceptualized as the product of two factors: the sensitivity metric,  $\lambda$ , for each layer in a module which is computed by a similar approach to that in Dong et al. (2020), and the quantization weight error. The latter is calculated as the squared 2-norm difference between the quantized and full precision weights. Therefore, the importance metric is given by

$$\Omega_k = \frac{1}{L_k} \sum_{i=1}^{L_k} \frac{\lambda_i}{N_i} |Q(W_i) - W_i|_2^2, \quad (1)$$

with

$k$  the  $k$ -th module in the modular network;

- $i$  the  $i$ -th layer within module  $k$ ;
- $L_k$  the total number of layers in module  $k$ ;
- $N_i$  the number of parameters in the  $i$ -th layer of module  $k$ ;
- $\lambda_i$  the sensitivity of the  $i$ -th layer in module  $k$ ;
- $Q$  the quantization operation;
- $W_i$  the weights of the  $i$ -th layer in module  $k$ .

Finding the optimal bit precisions using ILP is formulated as

$$\begin{aligned} & \min_{\{\rho_k\}_{k=1}^K} \sum_{k=1}^K \Omega_k^{(\rho_k)}, \\ \text{subject to } & \sum_{k=1}^K S_k^{(\rho_k)} \leq \frac{\text{Full Precision Model Size}}{\text{Compression factor}}. \end{aligned} \quad (2)$$

In this formulation,  $\rho_k$  denotes the bit-width for the  $k^{\text{th}}$  module;  $S_k$  denotes the model size of module  $k$ ; and  $K$  represents the total number of layers.

## 4 Experiments

Historically, quantization and other compression methods have been used to exercise a trade-off between inference accuracy and deployment feasibility, particularly in resource-constrained circumstances. In the following sections, we will show that our approach may yield a significant inference accuracy improvement during compression.

We first introduce the datasets and metrics used for evaluation. Then, we present ablation studies to explore the properties of MQAT; this result is directly compared to uniform and mixed QAT methods. Finally, we demonstrate the generality of our method applied to different datasets, architectures, and QAT methods.

### 4.1 Datasets and Metrics

The LINEMOD (LM) and Occluded-LINEMOD (LM-O) datasets are standard BOP benchmark datasets for evaluating 6D object pose estimation methods, where the LM dataset contains 13 different sequences consisting of ground-truth poses for a single object. LM-O extends LM by including occlusions. Similar to GDR-Net (Wang et al., 2021), we utilize 15% of the images for training. For both datasets, additional rendered images are used during training (Wang et al., 2021; Peng et al., 2019). Similarly to previous works, we use the ADD and ADI/ADD-S error metrics<sup>2</sup> (Hodaň et al., 2016) expressed as

$$e_{ADD}(P^{est}, P^{gt}) = \frac{1}{V} \sum_{i=1}^V \|P_i^{est} - P_i^{gt}\|_2, \quad (3)$$

$$e_{ADI}(P^{est}, P^{gt}) = \frac{1}{V} \sum_{i=1}^V \min_{j \in [1, V]} \|P_i^{est} - P_j^{gt}\|_2, \quad (4)$$

where  $P_i^{est}$  and  $P_i^{gt}$  denote the  $i^{\text{th}}$  vertex of the 3D mesh after transformation with the predicted and ground-truth pose, respectively. We then report the accuracy using the *ADD-0.1d* and *ADD-0.5d* metrics, which encode the proportion of samples for which  $e_{ADD}$  is less than 10% and 50% of the object diameter. Similarly, we report *ADI-0.1d* and *ADI-0.5d* metrics for swisscube dataset to be consistent with (Hu et al., 2021b).

While LM and LM-O present their own set of challenges, their scope is restricted to household objects, consistently illuminated without significant depth variations. Conversely, the SwissCube dataset (Hu et al.,

<sup>2</sup>The ADI/ADD-S metric was proposed for image sets with *indistinguishable* poses. ADI is used to compare fairly with previous works. (Hu et al., 2021b)




Figure 4: **A heuristic search of quantization flow sequences** to demonstrate quantization flow optimality for  $K=3$  module WDR. This corresponds to lines 1-7 in Algo.1.

MQAT	ADI	
First Quantized Module	0.1d	0.5d
Full Precision	78.79	98.98
Backbone	69.08	96.79
Head	67.84	98.12
Feature Pyramid Network (FPN)	<b>83.8</b>	<b>99.4</b>

Table 1: Effect of Starting MQAT with Different Modules.

2021b) embodies a challenging scenario for 6D object pose estimation in space, incorporating large scale variations, diverse lighting conditions, and variable backgrounds. To remain consistent with previous works, we use the same training setup and metric as Wang et al. (2022).

## 4.2 Implementation Details

We use PyTorch to implement our method. For the retraining of our partially quantized pretrained network, we employ an SGD optimizer with a base learning rate of  $1e-2$ . It is common practice for quantization algorithms to start with a pre-trained model (Esser et al., 2020; Dong et al., 2019; Zhou et al., 2017); we similarly do so here. We employed 30 epochs to identify the starting module at 2bits (Alg. 1 Lines 1-7) and then trained 30 epochs per module (Alg. 1 Lines 18-21). Therefore MQAT training time will scale linearly with additional modules. Specifically, WDR with  $M=3$  results in 120 epochs and ZebraPose with  $M=2$  results in 90 epochs. We use the same training time respectively for LSQ, INQ, HAWQ for a fair comparison. For all experiments, we use a batch size of 8 and employ a hand-crafted learning scheduler which decreases the learning rate at regular intervals by a factor of 10 and increases it again when we quantize a module with INQ<sup>3</sup>. However, when we quantize our modules using LSQ, the learning rate factor is not increased, only decreased by factors of 10. We use a  $512 \times 512$  resolution input for the SwissCube dataset and  $640 \times 480$  for LM and LM-O as in Peng et al. (2019).

## 4.3 MQAT Paradigm Studies

In this section, we conduct comprehensive studies using SwissCube and demonstrate the superior performance of our approach over conventional QAT ones.

### 4.3.1 MQAT Quantization Flow

We first perform an ablation study to validate the recommended quantization flow’s sequence optimality. As discussed in section 3.2, the module quantization sequence is not commutative. Using WDR, we perform aggressive quantization to every combination of modules in the network. This is an  $O(2^K)$  search; this results in eight module quantization combinations for a network with  $K = 3$  modules. The results are visualized in fig. 4. The backbone and head modules exhibit greater sensitivity to aggressive quantization. Conversely, the accuracy of the network is enhanced when using 2 bit quantization on the Feature Pyramid Network (FPN) module only. No other combination of module quantizations yields an accuracy increase. This further emphasizes the importance of carefully selecting a module quantization flow.

We additionally perform ablation studies on the optimal order (i.e., flow) of module quantization. We begin by quantizing different modules first, instead of the FPN. Table 1 shows the results of both the head and backbone modules when they are the first module quantized. We observe that the inference accuracy decreases dramatically for both cases. No combination of module flow or bit precision schedule is able to

<sup>3</sup>The learning rate and quantization schedulers are provided in the appendix.



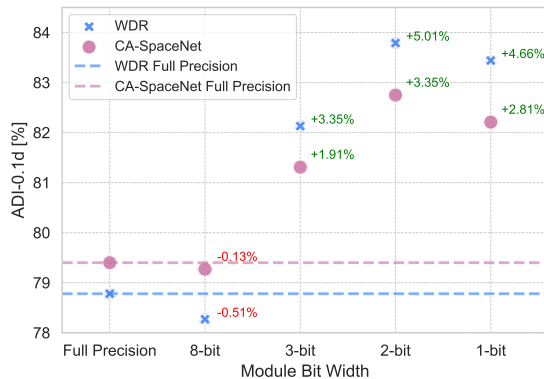


Figure 5: **Ablation study on the FPN bit-width.** We compare the performance by varying the bit-width of the feature aggregation module in each model.

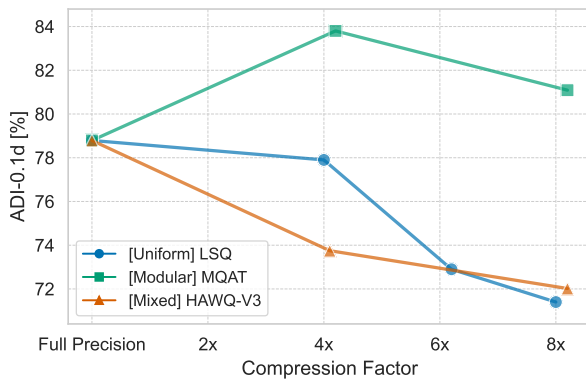


Figure 6: Comparison between our proposed paradigm MQAT, uniform QAT (LSQ), and layer-wise mixed-precision QAT (HAWQ-V3).

Table 2: **Comparison with the state-of-the-art on SwissCube.** We report ADI-0.1d scores for three different depth ranges. A \* indicates applying MQAT with 2-bit precision FPN to the model.

Network	Near	Medium	Far	All
SegDriven-Z (Hu et al., 2019)	41.1	22.9	7.1	21.8
DLR (Chen et al., 2019)	52.6	45.4	29.4	43.2
CA-SpaceNet	91.0	86.3	61.7	79.4
CA-SpaceNet*	95.5	90.7	66.2	82.7
WDR	92.4	84.2	61.3	78.8
WDR*	<b>96.1</b>	<b>91.5</b>	<b>68.2</b>	<b>83.8</b>

recover the inference accuracy after it was lost. The 2 bit aggressive FPN quantization yields improved accuracy only when the FPN is quantized first.

### 4.3.2 Quantized FPN Sensitivity Study.

To expand upon the 2 bit FPN accuracy enhancement, we perform a higher granularity bit-precision search on the FPN module. Again, the FPN module was quantized, but to five different bit-widths for comparison; the results are presented in fig. 5. The accuracy of two full-precision networks, WDR (Hu et al., 2021b) and CA-SpaceNet (Wang et al., 2022), are shown with dashed lines. The highest accuracy is achieved with a 2 bit or ternary  $\{-1, 0, 1\}$  FPN. Further pushing the FPN to binary weights  $\{-1, 1\}$  slightly reduces the accuracy, but maintains a significant improvement over both baselines<sup>4</sup>.

In the context of our study, as depicted in fig. 5, our findings support the premise that varying bit-precisions across different modules within QAT, such as the FPN, can significantly influence overall performance. This can be attributed to a redistribution of computes or inherent regularization effects on the modules. The improvements in performance reinforce the potential of MQAT in enhancing model generalizability.

### 4.3.3 MQAT Compared to Uniform and Mixed QAT.

For direct comparison, we apply three different quantization paradigms. Starting from a full precision WDR network, we apply a uniform QAT method, LSQ (Esser et al., 2020), a mixed-precision QAT method, HAWQ-V3 (Yao et al., 2020), and finally our proposed MQAT method with increasing compression factors. The results are provided in fig. 6. Again, MQAT demonstrates a significant accuracy improvement comparing to other methods while sustaining the requested compression factor; it is the only quantization approach to show an increase in inference accuracy during compression.

<sup>4</sup>For the interested reader, the FPN layer-wise ADI-0.1d accuracies are provided in the appendix.

Table 3: **Quantized FPN in WDR network on different datasets.** We report ADI for Swisscube and ADD for LM/LM-O, with the quantization flow  $F \rightarrow H \rightarrow B$ .

MQAT Mode	Bit-Precisions	SwissCube		LM		LM-O	
		0.1d	0.5d	0.1d	0.5d	0.1d	0.5d
Full precision	Full precision	78.8	98.9	56.1	99.1	37.8	85.2
MQAT <sub>LSQ</sub>	8-2-8	83.4	99.3	63.5	99.2	39.8	86.4
MQAT <sub>INQ</sub>	8-2-8	<b>83.7</b>	<b>99.4</b>	<b>63.9</b>	<b>99.5</b>	<b>40.2</b>	<b>86.7</b>

#### 4.4 MQAT Generality

Finally, we demonstrate the generality of MQAT to different datasets, QAT methods, and network architectures. Additionally, we showcase its applicability to a different task: object detection.<sup>5</sup>

##### 4.4.1 Dataset and QAT Generality

As discussed in section 4.1, the image domains of LM, LM-O and SwissCube are vastly different. The full precision and MQAT quantized models results for all three datasets are shown in table 3. MQAT demonstrates an accuracy improvement in all datasets. We use the ADI metric for evaluation on the SwissCube dataset as in (Hu et al., 2021b; Wang et al., 2022), while we use the ADD metric for LM and LM-O as used by (Wang et al., 2021; Su et al., 2022; Thalhammer et al., 2021; Labbé et al., 2020; Peng et al., 2019).

Accuracy improvements of 5.0%, 7.8% and 2.4% are demonstrated on SwissCube, LM and LM-O, respectively, when MQAT with INQ is utilized. Replacing INQ with LSQ yields accuracy improvements of 4.6%, 7.5% and 2.0%, respectively. This evidences that the performance enhancement is independent of the dataset domain and the applied QAT method.

As discussed in section 3.2 and section 4.3.1, it is difficult to recover accuracy once it is lost during quantization. To this end, since INQ (Zhou et al., 2017) quantizes only a fraction of the network at once, it follows that the remaining unquantized portion of the network is left flexible to adapt to aggressive quantization. Conversely, LSQ (Esser et al., 2020) quantizes the entire network in a single step; no fraction of the network is left unperturbed. Consequently, INQ demonstrates superior results in table 3. While any QAT method may be used, we recommend partnering MQAT with INQ for optimal aggressive quantization results.

##### 4.4.2 Architecture Generality

In table 2, we compare several single-stage PnP architectures on the SwissCube dataset. To demonstrate the generality of our performance enhancement, we apply MQAT to aggressively quantize the FPN of both CA-SpaceNet (Wang et al., 2022) and WDR (Hu et al., 2021b). We demonstrate an accuracy improvement of 4.5%, 4.4% and 4.5% for Near, Medium and Far images, respectively, on CA-SpaceNet, resulting in a total testing set accuracy improvement of 3.3%. Recall the already presented total testing set accuracy improvement of 5.0% for WDR. Previously, the full precision CA-SpaceNet had shown a performance improvement over the full precision WDR, but WDR sees greater gains from the application of MQAT.

In addition, (Wang et al., 2022) published accuracy results for a uniform QAT quantized CA-Space network, shared in table 4. Specifically, CA-SpaceNet explored three quantization modes (B, BF and BFH). These correspond to quantizing the backbone, quantizing the backbone and FPN (paired), and quantizing the whole network (uniformly), respectively.

Finally, we evaluate MQAT on a multi-stage network architecture. Specifically, in Table 5, we demonstrate the performance of our method on the state-of-the-art 6D object pose estimation network, the two-stage ZebraPose network (Su et al., 2022). Note that, when quantized using MQAT, the model’s performance comes close to that of its full-precision counterpart while being more than four times smaller. Furthermore, MQAT outperforms the state-of-the-art HAWQ-V3 (Yao et al., 2020) by  $\sim 2.6\%$ , with the added advantage of further network compression. To further demonstrate the efficacy, we quantize another two stage network i.e

<sup>5</sup>The results for the object detection task are provided in the appendix.

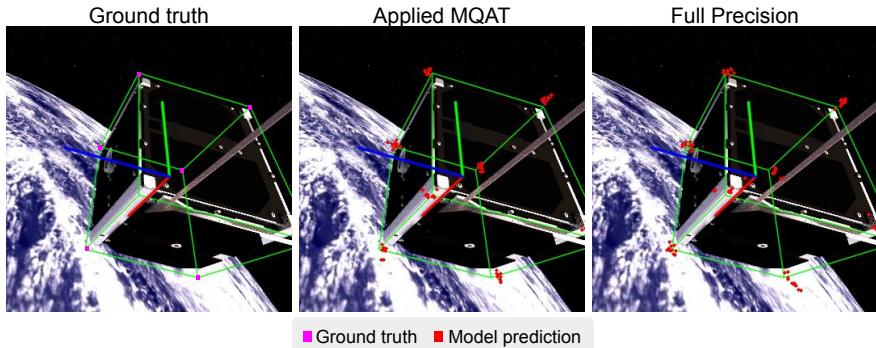


Figure 7: **Visualization of predictions.** Comparison between our proposed MQAT paradigm, and full-precision network. The model applying MQAT yields predictions that are on par with, or more concentrated than its full-precision counterpart.

Table 4: **CA-SpaceNet Published Quantization vs MQAT.** We report ADI scores on the SwissCube dataset sorted by the compression factor of the network, for MQAT methods, we use quantization flow  $F \rightarrow H \rightarrow B$ .

Quantization Method	ADI-0.1d	Compression	Bit-Precisions (B-F-H)
LSQ	79.4	1×	32-32-32
LSQ B	76.2	2.2×	8-32-32
LSQ BF	75.0	3.2×	8-8-32
LSQ BFH	74.7	4.0×	8-8-8
<b>MQAT (Ours)</b>	<b>82.7</b>	<b>4.7×</b>	<b>8-2-8</b>
LSQ B	75.1	2.9×	3-32-32
LSQ BF	74.5	5.9×	3-3-32
<b>MQAT (Ours)</b>	<b>80.2</b>	<b>8.2×</b>	<b>4-2-4</b>
LSQ BFH	68.7	10.6×	3-3-3

GDR-Net with existing uniform and mixed-precision methods, and MQAT. We employed ADD-0.1d metric for 6D object pose evaluation for O-Linmod dataset as Wang et al. (2021). It is evident from table 6 that MQAT outperforms both LSQ (Esser et al., 2020) and HAWQ-V3 (Yao et al., 2020) even with slightly more compressed network.

As we demonstrated in section 4.3.1, quantizing network modules all together greatly reduces inference accuracy as the smaller unquantized fraction of the network is not able to adapt to the quantization. Additionally, quantizing from backbone to head does not consider the sensitivity of the network modules to quantization. As a final note, CA-SpaceNet does not quantize the first and last layer in any quantization mode. In contrast, MQAT quantizes the entire network.

## 5 Limitations and Discussion

**Module Granularity.** As conclusively demonstrated in section 4.3.1, MQAT exploits the modular structure of a network. Therefore, if the network does not contain distinct modules, MQAT simply converges to a uniform QAT methodology. In principle, MQAT can apply to any architectures with  $K \geq 2$ .

**Latency.** Directly reporting latency measurements involves hardware deployment, which goes beyond the scope of this work. However, as shown in Yao et al. (2020), latency is directly related to the bit operations per second (BOPs). With lower-precision networks, both the model size and the BOPs are reduced by the same compression factor, which we provide in our experiments. Therefore, it is expected that MQAT would demonstrate a latency improvement proportional to the network compression factor.

Nonetheless, we can share some results running an MQAT quantized network on a CPU. Running WDR on our **Intel Core i7-9750H** CPU demonstrates a latency of **650ms** for full precision and **299ms** for our MQAT int8 quantized model. However this does not exploit the full benefits of our MQAT quantized model as all the CPU deployed parameters are stored at the highest common denominator of INT8 (whereas our model employs lower bit-widths for certain parameters).

Table 5: **Quantization of ZebraPose (Su et al., 2022)**. We report ADD scores on the LM-O dataset and compare MQAT to mix-precision quantization (HAWQ-V3).

Quantization Method	ADD 0.1d	Compression	Bit-Precisions	Quantization Flow
Full precision	76.90	1×	Full precision	N/A
HAWQ-V3 (Yao et al., 2020)	71.11	4×	Mixed (layer-wise)	N/A
HAWQ-V3 (Yao et al., 2020)	69.87	4.60×	Mixed (layer-wise)	N/A
<b>MQAT <math>K = 2</math> (ours)</b>	<b>72.54</b>	<b>4.62×</b>	8-4 (B-D)	D → B

Table 6: **Quantization of GDR-Net (Wang et al., 2021)**. We report ADD scores on the LM-O dataset and compare MQAT to uniform (LSQ) and mix-precision quantization (HAWQ-V3). B, R and P indicates *Backbone*, *Rotation Head* and *PnP-Patch* modules.

Quantization Method	ADD 0.1d	Compression	Bit-Precisions	Quantization Flow
Full precision	56.1	1×	Full precision	N/A
LSQ (Esser et al., 2020)	50.7	4.57×	Uniform(7-bit)	N/A
HAWQ-V3 (Yao et al., 2020)	50.3	4.9×	Mixed (layer-wise)	N/A
<b>MQAT (ours)</b>	<b>51.8</b>	<b>4.97×</b>	8-4-4 (B-R-P)	R → P → B

**Recommended Quantization Flow Optimality.** A major contribution of this paper is the identification of the existence of an asynchronous optimal quantization sequence or flow. In section 3.2 we provide a recommended quantization flow and exhaustively demonstrate its optimality for  $K = 3$  in fig. 4. However, the recommended quantization flow’s optimality has yet to be analytically proven for all combinations of networks and number of modules,  $K$ .

## 6 Conclusion

We have introduced Modular Quantization-Aware Training (MQAT) for networks that exhibit a modular structure, such as 6D object pose estimation architectures. Our approach builds on the intuition that the individual modules of such networks are unique, and thus should be quantized uniquely while heeding an optimal quantization order. Our extensive experiments on different datasets and network architectures, and in conjunction with different quantization methods, conclusively demonstrate that MQAT outperforms uniform and mixed-precision quantization methods at various compression factors. Moreover, we have shown that it can even enhance network performance. In particular, aggressive quantization of the network FPN resulted in 7.8% and 2.4% test set accuracy improvements over the full-precision network on LINEMOD and Occluded-LINEMOD, respectively. In the future, we will investigate the applicability of MQAT to tasks other than 6D object pose estimation and another potential follow up would be to apply MQAT to architectures with even more modules or to instead classify large modules as two or more modules for the purposes of MQAT.

## 7 Acknowledgements

We thank [Ziqi Zhao](#) for assisting us with the experiments on Multi-stage architecture. This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No. 945363. Moreover, this work was funded in part by the Swiss National Science Foundation and the Swiss Innovation Agency (Innosuisse) via the BRIDGE Discovery grant No. 194729.

## References

Alan Ansell, E. Ponti, Anna Korhonen, and Ivan Vulic. Composable sparse fine-tuning for cross-lingual transfer. In *Annual Meeting of the Association for Computational Linguistics*, 2021.

- Yash Bhalgat, Jinwon Lee, Markus Nagel, Tijmen Blankevoort, and Nojun Kwak. LSQ+: Improving low-bit quantization through learnable offsets and better initialization. In *Computer Vision and Pattern Recognition*, 2020.
- David Blalock, Jose Javier, Gonzalex Ortiz, Jonathan Frankle, and John Gutta. What is the State of Neural Network Pruning? *Machine Learning and Systems*, 2020.
- Eric Brachmann, Alexander Krull, Frank Michel, Stefan Gumhold, Jamie Shotton, and Carsten Rother. Learning 6D Object Pose Estimation using 3D Object Coordinates. In *European Conference on Computer Vision*, 2014.
- Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. ZeroQ: A Novel Zero Shot Quantization Framework. In *Computer Vision and Pattern Recognition*, 2020.
- Zhaowei Cai and Nuno Vasconcelos. Rethinking Differentiable Search for Mixed-Precision Neural Networks. In *Computer Vision and Pattern Recognition*, 2020.
- Jerry Chee, Yaohui Cai, Volodymyr Kuleshov, and Christopher De Sa. QuIP: 2-Bit Quantization of Large Language Models With Guarantees. In *Neural Information Processing Systems*, 2023.
- Bo Chen, Jiewei Cao, Alvaro Parra, and Tat-Jun Chin. Satellite Pose Estimation with Deep Landmark Regression and Nonlinear Pose Refinement. In *International Conference on Computer Vision*, 2019.
- Peng Chen, Jing Liu, Bohan Zhuang, Mingkui Tan, and Chunhua Shen. AQD: Towards Accurate Quantized Object Detection. In *Computer Vision and Pattern Recognition*, 2021.
- Alexandre Défossez, Yossi Adi, and Gabriel Synnaeve. Differentiable Model Compression via Pseudo Quantization Noise. *Transactions on Machine Learning Research*, 2022.
- Yan Di, Fabian Manhardt, Gu Wang, Xiangyang Ji, Nassir Navab, and Federico Tombari. SO-Pose: Exploiting Self-Occulsion for Direct 6D Pose Estimation. In *International Conference on Computer Vision*, 2021.
- Zhen Dong, Zhewei Yao, Amir Gholami, Michael Mahoney, and Kurt Keutzer. HAWQ: Hessian AWARE Quantization of Neural Networks with Mixed-Precision. In *International Conference on Computer Vision*, 2019.
- Zhen Dong, Zhewei Yao, Yaohui Cai, Daiyaan Arfeen, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. HAWQ-V2: Hessian Aware trace-Weighted Quantization of Neural Networks. In *Neural Information Processing Systems*, 2020.
- Steven K. Esser, Jeffrey L. McKinstry, Deepika Bablani, Rathinakuma Appuswamy, and Dharmendra S. Modha. Learned Step Size Quantization. In *International Conference on Learning Representations*, 2020.
- Elias Frantar and Dan Alistarh. Optimal Brain Compression: A Framework for Accurate Post-Training Quantization and Pruning . In *Neural Information Processing Systems*, 2022.
- Stefan Hinterstoisser, Vincent Lepetit, Slobodan Ilic, Stefan Holzer, Gary R. Bradski, Kurt Konolige, and Nassir Navab. Model Based Training, Detection and Pose Estimation of Texture-Less 3D Objects in Heavily Cluttered Scenes. In *Asian Conference on Computer Vision*, 2012.
- Tomáš Hodaň, Dániel Baráth, and Jiří Matas. EPOS: Estimating 6D Pose of Objects with Symmetries. In *Computer Vision and Pattern Recognition*, 2020.
- Tomáš Hodaň, Jiří Matas, and Štěpán Obdržálek. On Evaluation of 6D Object Pose Estimation. In *European Conference on Computer Vision*, 2016.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models. In *International Conference on Learning Representations*, 2021a.

- Yinlin Hu, Joachim Hugonot, Pascal Fua, and Mathieu Salzmann. Segmentation-Driven 6D Object Pose Estimation. In *Computer Vision and Pattern Recognition*, 2019.
- Yinlin Hu, Sébastien Speierer, Wenzel Jakob, Pascal Fua, and Mathieu Salzmann. Wide-Depth-Range 6D Object Pose Estimation in Space. In *Computer Vision and Pattern Recognition*, 2021b.
- Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive Mixtures of Local Experts. In *Neural computation*, 1991.
- Omid Hosseini Jafari, Siva Karthik Mustikovela, Karl Pertsch, Eric Brachmann, and Carsten Rother. iPose: Instance-Aware 6D Pose Estimation of Partly Occluded Objects. In *Asian Conference on Computer Vision*, 2018.
- Wadim Kehl, Fabian Manhardt, Federico Tombari, Slobodan Ilic, and Nassir Navab. SSD-6D: Making RGB-Based 3D Detection and 6D Pose Estimation Great Again. In *International Conference on Computer Vision*, 2017.
- Yann Labbé, Justin Carpentier, Mathieu Aubry, and Josef Sivic. CosyPose: Consistent multi-view multi-object 6D pose estimation. In *European Conference on Computer Vision*, 2020.
- Rundong Li, Yan Wang, Feng Liang, Hongwei Qin, Junjie Yan, and Rui Fan. Fully Quantized Network for Object Detection. In *Computer Vision and Pattern Recognition*, 2019a.
- Yuelong Li, Yafei Mao, Raja Bala, and Sunil Hadap. Mrc-net: 6-dof pose estimation with multiscale residual correlation. In *2024 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2024.
- Yuhang Li, Ruihao Gong, Zu Tan, Yang Yang, Peng Hu, Qi Zhang, Fengwei Yu, Wei Wang, and Shi Gu. BRECCQ: Pushing the Limit of Post-Training Quantization by Block Reconstruction. In *International Conference on Learning Representations*, 2021.
- Zhigang Li, Gu Wang, and Xiangyang Ji. CDPN: Coordinates-Based Disentangled Pose Network for Real-Time RGB-Based 6-DoF Object Pose Estimation. In *International Conference on Computer Vision*, 2019b.
- Ruyi Lian and Haibin Ling. CheckerPose: Progressive Dense Keypoint Localization for Object Pose Estimation with Graph Neural Network. In *International Conference on Computer Vision*, 2023.
- Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. AWQ: Activation-aware Weight Quantization for LLM Compression and Acceleration. In *MLSys*, 2024.
- 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 *European Conference on Computer Vision*, 2014.
- Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature Pyramid Networks for Object Detection. In *Computer Vision and Pattern Recognition*, 2016.
- Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: Single Shot MultiBox Detector. In *European Conference on Computer Vision*, 2016.
- Zechun Liu, Kwang-Ting Cheng, Dong Huang, Eric Xing, and Zhiqiang Shen. Nonuniform-to-Uniform Quantization: Towards Accurate Quantization via Generalized Straight-Through Estimation. In *Computer Vision and Pattern Recognition*, 2022.
- Vincent Lepetit Markus Oberweger, Mahdi Rad. Making Deep Heatmaps Robust to Partial Occlusions for 3D Object Pose Estimation. In *European Conference on Computer Vision*, 2018.

- Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Data-Free Quantization Through Weight Equalization and Bias Correction. In *International Conference on Computer Vision*, 2019.
- Markus Nagel, Rana Ali Amjad, Mart van Baalen, Christos Louizos, and Tijmen Blanevoort. Up or Down? Adaptive Rounding for Post-Training Quantization. In *International Conference on Machine Learning*, 2020.
- Van Nguyen Nguyen, Thibault Groueix, Mathieu Salzmann, and Vincent Lepetit. GigaPose: Fast and Robust Novel Object Pose Estimation via One Correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.
- Tae Ha Park, Marcus Märtens, Gurvan Lecuyer, Dario Izzo, and Simone D’Amico. Speed+: Next-generation dataset for spacecraft pose estimation across domain gap. In *Aerospace Conference*, 2022.
- Sida Peng, Yuan Liu, Qixing Huang, Hujun Bao, and Xiaowei Zhou. PVNet: Pixel-wise Voting Network for 6DoF Pose Estimation. In *Computer Vision and Pattern Recognition*, 2019.
- Luis Pérez, Ínigo Rodríguez, Nuria Rodríguez, Rubén Usamentiaga, and Daniel F. García. Robot Guidance Using Machine Vision Techniques in Industrial Environmnets: A Comparative Review. *Sensors*, 2016.
- Jorn Peters, Marios Fournarakis, Markus Nagel, Mart van Baalen, and Tijmen Blankevoort. QBitOpt: Fast and Accurate Bitwidth Reallocation during Training. In *International Conference on Computer Vision*, 2023.
- Jonas Pfeiffer, Ivan Vulić, Iryna Gurevych, and Sebastian Ruder. MMAD-X: An Adapter-Based Framework for Multi-Task Cross-Lingual Transfer. *EMNLP*, 2020.
- Jonas Pfeiffer, Sebastian Ruder, Ivan Vulić, and Edoardo Maria Ponti. Modular deep learning. *Transactions on Machine Learning Research*, 2023.
- Pedro F. Proença and Yang Gao. Deep Learning for Spacecraft Pose Estimation from Photorealistic Rendering. In *International Conference on Robotics and Automation*, 2020.
- Mahdi Rad and Vincent Lepetit. BB8: A Scalable, Accurate, Robust to Partial Occlusion Method for Predicting the 3D Poses of Challenging Objects without Using Depth. In *International Conference on Computer Vision*, 2017.
- Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, and Ping Luo. OmniQuant: Omnidirectionally Calibrated Quantization for Large Language Models. In *International Conference on Learning Representations*, 2024.
- Juncheol Shin, Junhyuk So, Sein Park, Seungyeop Kang, Sungjoo Yoo, and Eunhyeok Park. NIPQ: Noise Proxy-Based Integrated Pseudo-Quantization. In *Computer Vision and Pattern Recognition*, 2023.
- Abhilasha Singh, V. Kalaichelvi, and R. Karthikeyan. A survey on vision guided robotic systems with intelligent control strategies for autonomous tasks. *Cogent Engineering*, 2022.
- Chen Song, Jiaru Song, and Qixing Huang. HybridPose: 6D Object Pose Estimation under Hybrid Representations. In *Computer Vision and Pattern Recognition*, 2020.
- Jianing Song, Duarte Rondao, and Nabil Aouf. Deep Learning-based Spacecraft Relative Navigation Methods: A Survey. *Acta Astronautica*, 2022.
- Sebastian Stapf, Tobias Bauernfeind, and Marco Riboldi. PViT-6D: Overclocking Vision Transformers for 6D Pose Estimation with Confidence-Level Prediction and Pose Tokens, 2023. URL <https://arxiv.org/abs/2311.17504>.
- Yongzhi Su, Mahdi Saleh, Torben Fetzer, Jason Rambach, Nassir Navab, Benjamin Busam, Didier Stricker, and Deferico Tombari. ZebraPose: Coarse to Fine Surface Encoding for 6DoF Object Pose Estimation. In *Computer Vision and Pattern Recognition*, 2022.

- Mingxing Tan, Ruoming Pang, and Quoc V. Le. EfficientDet: Scalable and Efficient Object Detection. In *Computer Vision and Pattern Recognition*, 2020.
- Chen Tang, Kai Ouyang, Zhi Wang, Yifei Zhu, Yaowei Wang, Wen Ji, and Wenwu Zhu. Mixed-Precision Neural Network Quantization via Learned Layer-wise Importance. In *European Conference on Computer Vision*, 2022.
- Bugra Tekin, Sudipta N. Sinha, and Pascal Fua. Real-Time Seamless Single Shot 6D Object Pose Prediction. In *Computer Vision and Pattern Recognition*, 2018.
- Stefan Thalhammer, Markus Leitner, Timothy Patten, and Markus Vincze. PyraPose: Feature Pyramids for Fast and Accurate Object Pose Estimation under Domain Shift. In *International Conference on Robotics and Automation*, 2021.
- Federico Vicentini. Collaborative Robotics: A Survey. *Journal of Mechanical Design*, 2020.
- Gu Wang, Favian Manhardt, Federico Tombari, and Xiangyang Ji. GDR-Net: Geometry-Guided Direct Regression Networks for Monocular 6D Pose Estimation. In *Computer Vision and Pattern Recognition*, 2021.
- Shunli Wang, Shuaibing Wang, Bo Jiao, Ding kang Yang, Liuzhen Su, Peng Zhai, Chixiao Chen, and Lihua Zhang. CA-SpaceNet: Counterfactual Analysis for 6D Pose Estimation in Space. In *2022 International Conference on Intelligent Robots and Systems*, 2022.
- Bowen Wen, Wei Yang, Jan Kautz, and Stan Birchfield. FoundationPose: Unified 6D Pose Estimation and Tracking of Novel Objects. In *Computer Vision and Pattern Recognition*, 2024.
- Yu Xiang, Tanner Schmidt, Venkatraman Narayanan, and Dieter Fox. PoseCNN: A Convolutional Neural Network for 6D Object Pose Estimation in Cluttered Scenes. *Robotics: Science and Systems*, 2018.
- Kohei Yamamoto. Learnable Companding Quantization for Accurate Low-bit Neural Networks. In *Computer Vision and Pattern Recognition*, 2021.
- Jiaming Yang, Chenwei Tang, Caiyang Yu, and Jiancheng Lv. GWQ: Group-Wise Quantization Framework for Neural Networks. In *Asian Conference on Machine Learning*, 2023.
- Zhewei Yao, Zhen Dong, Zhangcheng Zheng, Amir Gholami, Jiali Yu, Eric Tan, Leyuan Wang, Qijing Huang, Yida Wang, Michael W. Mahoney, and Kurt Keutzer. HAWQV3: Dyadic Neural Network Quantization. In *International Conference on Machine Learning*, 2020.
- Sergey Zakharov, Ivan Shugurov, and Slobodan Ilic. DPOD: 6D Pose Object Detector and Refiner. In *International Conference on Computer Vision*, 2019.
- Yifan Zhang, Zhen Dong, Huanrui Yang, Ming Lu, Cheng-Ching Tseng, Yuan Du, Kurt Keutzer, Li Du, and Shanghang Zhang. QD-BEV: Quantization-aware View-guided Distillation for Multi-view 3D Object Detection. In *International Conference on Computer Vision*, 2023.
- Ritchie Zhao, Yuwei Hu, Jordan Dotzel, Chris De Sa, and Zhiru Zhang. Improving Neural Network Quantization without Retraining using Outlier Channel Splitting. In *International Conference on Machine Learning*, 2019.
- Aojun Zhou, Anbang Yao, Yiwen Guo, Lin Xu, and Yurong Chen. Incremental Network Quantization: Towards Lossless CNNs with Low-Precision Weights. In *International Conference on Learning Representations*, 2017.



## A Appendix

### A.1 Multi-scale Fusion Analysis of MQAT Applied to WDR (Hu et al., 2021b)

In this section we study the inference performance of our MQAT when applied to WDR, our primary focus is to understand the impact of MQAT on the multi-scale fusion inference of WDR, particularly examining changes in performance across individual layers of FPN module in WDR. To this end, we applied MQAT to the WDR architecture and conducted a layer-by-layer performance analysis of FPN module similar to the original WDR paper. The results are presented in Table 7. A notable observation from this analysis is the overall enhancement in performance across most layers and the improvement is consistent across various scales and depth ranges. Particularly, Layer 1 exhibits a significant performance boost, especially for objects classified as *Far*. This layer-specific insight underscores the effectiveness of MQAT in optimizing the WDR network at a granular level.

Layer	Near	Medium	Far	All
1	13.6 (+0.8)	83.8 (+1.0)	55.1 (+6.2)	52.9(+3.0)
2	13.6 (+2.6)	83.8 (+1.3)	55.1 (+0.3)	54.1 (+1.3)
3	16.3 (-0.5)	77.8 (+1.6)	17.1 (-1.2)	37.2 (-0.1)
4	13.0 (+0.7)	0.41 (+0.7)	0 (0)	3.8 (+0.4)

Table 7: **Layer-wise 6D Pose Validation Results with MQAT on WDR’s FPN.** ADI-0.1d are reported for each layer. Notable performance enhancements, particularly in Layer 1 for *Far* objects, illustrate the effective quantization of WDR by MQAT.

### A.2 Generality of MQAT in Object Detection Problem

In our study, while the primary focus is on 6D pose estimation, where our method’s efficacy is already demonstrated, we further extend our evaluation to object detection tasks. This extension is aimed at underscoring the generality of our approach.

To this end, we applied our quantization technique to the Faster R-CNN network, which utilizes a ResNet-50 backbone, a widely recognized model in object detection tasks. Our evaluation was conducted on the comprehensive COCO dataset, a benchmark for object detection.

Network	QAT Method	mAP	Compression
FasterRCNN	Full Precision	37.9	1x
	FQN	32.4	8x
	INQ	33.4	8x
	LSQ	33.8	8x
	<b>MQAT (ours)</b>	<b>35.1</b>	<b>8x</b>
EfficientDet-D0	Full Precision	33.16	1x
	N2UQ	20.11	10x
	<b>MQAT (ours)</b>	<b>21.67</b>	<b>10x</b>

Table 8: **Quantization for Object Detection.** We evaluate the given networks on COCO dataset and report mAP.

The results of this experiment are summarized in Table 8. Here, our method, denoted as MQAT, is compared against other quantization approaches: INQ (Zhou et al., 2017), LSQ (Esser et al., 2020), and the FQN (Li et al., 2019a) method, which has been specifically tailored for object detection tasks. The comparison reveals



Figure 8: **Visualization of the Difference in Object Detection Performance** on MSCOCO (Lin et al., 2014) between N2UQ and MQAT at the same compression ratio.

that MQAT not only adapts well to a different task domain but also achieves superior performance over these established Quantization-Aware Training techniques. This underlines the adaptability and robustness of our approach, extending its potential applications beyond 6D pose estimation to broader areas within computer vision.

Moreover, we also created a baseline for the network, EfficientDet (Tan et al., 2020) by quantizing it with a recent quantization method: Non-Uniform to Uniform Quantization (N2UQ) (Liu et al., 2022)). This comparison is crucial to validate the effectiveness of MQAT across various modular architectures and quantization methods. Remarkably, our MQAT approach demonstrated a performance improvement of approximately 1.6% over N2UQ as shown in both Table 8. and fig. 8. This enhancement was observed under comparable compression ratios and identical training durations, further substantiating the superiority of MQAT in terms of efficiency and effectiveness.

### A.3 More Results on Speed+ Dataset (Park et al., 2022)

#### A.3.1 Overview of Speed+ Dataset

The Next Generation Spacecraft Pose Estimation Dataset (Speed+) addresses the the domain gap challenge in spacecraft pose estimation. it encompasses 60,000 synthetic images, divided into an 80:20 train-to-validation ratio. The test set comprises of 9,531 Hardware-In-the-Loop images of the half-scale mockup model of the Tango spacecraft.

### A.3.2 Results on Speed+ Dataset

Our method was further tested on the Speed+ dataset. The WDR network was quantized with our proposed MQAT and the results are shown in Table 9, where we exclusively assess the model on the validation dataset.

Network	ADI-0.1d
Full precision	96.2
MQAT <sub>(8-2-8)</sub>	<b>99.1</b>

Table 9: **Evaluation of MQAT on Speed+ dataset.**

## A.4 More Implementation Details

### A.4.1 LR Schedule for INQ (Zhou et al., 2017)

As mentioned in our experiments section, we employ a SGD optimizer with a base learning rate ( $lr_b$ ) of  $1e^{-2}$ . We trained for 30 epochs with a batch size of 8. We created a hand-crafted learning scheduler which decreases the learning rate at regular intervals by a factor of 10 and increases it again when we quantize a module. The gamma ( $\gamma$ ) and quantization fraction scheduler are shown in Table 10. The quantization fraction corresponds to the percentage of weights quantized at each epoch. At each epoch, the learning rate ( $lr$ ) is computed as:

$$lr = lr_b * \gamma$$

Epoch	Schedule	
	$\gamma$	fraction
0	1	0.2
3	0.1	-
5	1	0.4
7	0.1	-
9	1	0.6
11	0.1	-
13	1	0.8
15	0.1	-
17	1	0.9
19	0.1	-
21	1	0.95
23	0.1	-
25	1	0.975
27	0.1	-
29	1	1
30	0.1	-

Table 10: **Learning Rate Schedule.**

### A.4.2 LR Schedule for LSQ (Esser et al., 2020) and HAWQ-V3 (Yao et al., 2020)

As mentioned in our experiments section, we employ a SGD optimizer with a base learning rate ( $lr_b$ ) of  $1e-2$ . However, we used a multi-step decay scheduler here. Learning rate was decreased by factors of 10 at epochs  $\{10, 20, 25\}$  of the training. We trained for 30 epochs with a batch size of 8.

## **A.5 Reproducibility**

The source code will be provided publicly along with the details of the environments and dependencies. Moreover, we will provide instructions to reproduce the main results in the manuscript.