

## 505 7 Appendix

### 506 7.1 Details for Problem Setups

507 **Baselines.** While our method aims to solve the QAT problem, we extensively compare our *BiTAT*  
508 against various Post-training Quantization (PTQ)- or QAT methods: BRECQ [18] is a PTQ method  
509 that considers weight dependencies using the Hessian of the task loss. DBQ [8] is a QAT method based  
510 on continuous relaxation of the quantizer function. EBConv [3] conditionally selects appropriate  
511 binarized weights based on the task information. Bi-Real Net [20] adds residual connections to  
512 propagate full-precision values, preventing information loss to activation quantization. Real-to-  
513 Bin [22] constrains a loss term at the end of each convolution to minimize the output discrepancy  
514 between the full-precision and the quantized model. LCQ [35] devises a trainable quantization  
515 function in order to reduce the quantization error. MeliusNet [2] proposes a new architecture that  
516 better propagates full-precision values throughout the network. ReActNet [21] is the state-of-the-art  
517 binary quantization method, which additionally adopts residual connections, and element-wise shift  
518 operations before/after the activation and the sign operation. Note that DBQ, LCQ, and MeliusNet  
519 keep some crucial layers of MobileNet in full-precision, leading to inefficiency at evaluation time.

520 **Training.** Following the setup from ReActNet [21], we quantize all layers’ weights and activations  
521 except the initial and final layers. We use the Adam optimizer [16]. For the ImageNet experiment,  
522 a learning rate is 0.002 and 0.0002 for quantization training and the fine-tuning, respectively, with  
523 linear learning rate decay. We set batch size as 512 both the quantization phase and the fine-tuning  
524 phase is done for 5 epochs per layer. For the CIFAR-100 experiment, a learning rate is  $3 \times 10^{-4}$  for  
525 quantization training and the fine-tuning with linear learning rate decay. We set batch size as 800 and  
526 both the quantization and fine-tuning are done with 40 epochs per layer. For all experiments, we set  
527  $\lambda = 100$ , and  $\gamma = 10^{-5}$ , which notes that simple choice of the hyperparameters for our regularization  
528 terms is sufficient to show impressive performance. The number of input dimension groups is set  
529  $k = 256$ , applying the grouped weight correlation to layers with input dimensions smaller than  $k$ .

530 **Inference.** In deployment, the highly efficient XNOR-Bitcounting operations can be used for the  
531 convolutional layers, also used in existing neural network binarization works [5, 30, 21].

### 532 7.2 Extension to Convolutional Layers

533 Let us consider a convolutional layer of size  $n_{out} \times n_{in} \times k \times k$ , where  $n_{in}$  and  $n_{out}$  are the number  
534 of input and output channels, respectively, and  $k$  is the kernel size. We define the set  $\mathcal{P}_x$  as the set of  
535 all patches of size  $n_{in} \times k \times k$  extracted from the training image  $x$ . This patch-extracting operation  
536 is sometimes called `im2col` or `F.unfold` in PyTorch.

537 A convolutional layer applied to  $x$  can be thought of as a fully-connected layer individually applied  
538 to all patches in  $\mathcal{P}_x$  and then concatenated:

$$w * x = \{\text{Reshape}_{(n_{in}k^2) \times (n_{out})}(w)^\top \text{vec}(p)\}_{p \in \mathcal{P}_x}, \quad (14)$$

539 where  $*$  denotes the convolution operation,  $\text{Reshape}_{shape}(\cdot)$  denotes the reshaping of the tensor into  
540 the specified *shape*, and  $\text{vec}(\cdot)$  denotes the flattening operation. Each pixel of the output feature map  
541 corresponds to a matrix multiplication between a patch and the weight matrix. Therefore, we can  
542 analogically apply the same transformation as explained in Section 3 to convolutional layers.

### 543 7.3 Details on Cross-layer Dependency

544 In this section, we further explain the detailed experimental setting for Figure 3. We take the standard  
545 MobileNetV2 [31] model and train it to convergence on the CIFAR-100 dataset with standard SGD  
546 with a weight decay. Then, we add noise to the same pretrained model parameters before evaluating  
547 the test accuracy based on the following two different ways:

- 548 1. **Layer-dependent noise addition.** We first compute the covariance of the input to the first  
549 layer and perform PCA using obtained covariance values to compute  $\tilde{w}^{(1)}$  in Equation 4.  
550 Now, we add independent gaussian noise with varying scales to the top five rows of  $\tilde{w}^{(1)}$ .  
551 Next, we sequentially repeat the process to the consecutive layers, and after that, we evaluate

METHOD	Accuracy (%)	$Q_{orig}$	$Q_{ours}$
REACTNET [21]	$65.51 \pm 0.74$	<b>13.35</b>	475.94
BiTAT (Ours)	<b><math>68.36 \pm 0.45</math></b>	39.77	<b>434.92</b>

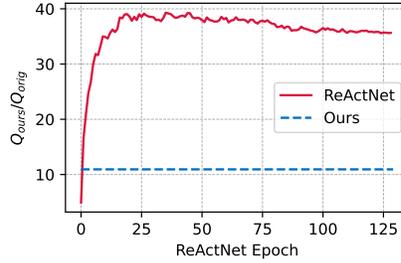


Figure 8: **Additional ablation studies.** **Left:** The comparison of the final  $Q_{orig}$  (Equation 2) and  $Q_{ours}$  (Equation 6) values in ReActNet and BiTAT. **Right:** The evolution of the ratio of  $Q_{ours}$  to  $Q_{orig}$  in ReActNet, in comparison to the final ratio in our BiTAT. Cross-layer dependencies not considered in both computations.

552 the model performance, which is shown in solid red lines. The same process is done but  
 553 with the bottom five rows of each layer instead of the top five, shown in solid blue lines.

554 2. **Layer-independent noise addition.** Before adding noise to model parameters, we compute  
 555 the covariance of the input values for all layers. Next, we perform PCA and compute  $\tilde{w}^{(l)}$   
 556 with Equation 4 per layer using these initial covariance values. That is, a layer cannot  
 557 reflect the weight change through noise addition in others layers, as different from the first  
 558 approach. Independent gaussian noise with varying scales is added to the top five rows of  
 559  $\tilde{w}^{(l)}$  for each layer, and then the performance of the model is evaluated, shown in red dashed  
 560 lines. The same is done with the bottom five rows of  $\tilde{w}^{(l)}$ , shown in blue dashed lines.

## 561 7.4 Additional Analysis

562 This paper suggests that the proposed quantization loss on disentangled weights is a better indicator  
 563 for prediction accuracy than the general quantization loss (Equation 2), which is evident in multiple  
 564 validation analyses and the superior model performance in our BiTAT as described in the main  
 565 text. Here, we provide the quantitative analysis to show that ReActNet [21] fails to minimize the  
 566 quantization loss on disentangled weights while our BiTAT successfully does. In Figure 8 Left, we  
 567 show the  $Q_{orig}$  (Equation 2) and  $Q_{ours}$  (Equation 6 w/o  $\ell_1$  norm) between the initial full-precision  
 568 weights of the pre-trained model and the obtained binarized weights from ReActNet and BiTAT.  
 569  $Q_{orig}$  represents the naive MSE between the full-precision weights and the binarized weights.  $Q_{ours}$   
 570 represents the dependency-weighted MSE between the full-precision weights and the binarized  
 571 weights. Note that, in this analysis, we obtain  $s$  and  $V$  for each layer from the initial pre-trained  
 572 model by Equation 10 to compute  $Q_{ours}$  and neglect the weight dependency across different layers,  
 573 which is hard to be computed analytically.

574 We observe that while the value of  $Q_{orig}$  is lower in ReActNet than in BiTAT,  $Q_{ours}$  is higher in  
 575 ReActNet than in BiTAT. As shown in Figure 8 Right, the ratio  $r = Q_{ours}/Q_{orig}$  in ReActNet (Red)  
 576 drastically increases at the beginning stage and is maintained in a high degree until the completion  
 577 of the quantization-aware training, compared to the  $r$  value of the model obtained by BiTAT (Blue  
 578 dashed). While disregarding the first few epochs of ReActNet training, where the accuracy is very low,  
 579 ReActNet’s  $r$  value dominates that of BiTAT. The value slowly decreases as the ReActNet training  
 580 proceeds, but never reaches the level of BiTAT, demonstrating the inefficiency of the ReActNet  
 581 training procedure compared to ours.

## 582 7.5 Limitations

583 We consider two limitations of our work in this section. First, our BiTAT framework is built based on  
 584 a sequential quantization strategy, which progressively quantizes the layers from the bottom to the  
 585 top. Therefore, the training time of our algorithm depends on the number of layers in the backbone  
 586 network architecture. While we have already validated the cost-efficiency of our proposed method  
 587 against ReActNet using MobileNet (26 stacked layers) in Figure 7 Left, we might spend more training  
 588 time quantizing all layer weights for the backbones, composed of much more layers like ResNet-1001  
 589 (1001 stacked layers). Next, our method focuses on the cross-layer weight dependency within each  
 590 neural block, including a few consecutive layers. We thereby avoid the excessive computational  
 591 cost of obtaining the relationship across all layers in the backbone architecture, yet we consider it a  
 592 tradeoff between accurate dependency and computation budgets.