
Accurate Post Training Quantization With Small Calibration Sets

Itay Hubara^{*12} Yury Nahshan^{*1} Yair Hanani¹ Ron Banner¹ Daniel Soudry²

Abstract

Lately, post-training quantization methods have gained considerable attention, as they are simple to use, and require only a small unlabeled calibration set. This small dataset cannot be used to fine-tune the model without significant overfitting. Instead, these methods only use the calibration set to set the activations' dynamic ranges. However, such methods always resulted in significant accuracy degradation, when used below 8-bits (except on small datasets). Here we aim to break the 8-bit barrier. To this end, we minimize the quantization errors of each layer or block separately by optimizing its parameters over the calibration set. We empirically demonstrate that this approach is: (1) much less susceptible to overfitting than the standard fine-tuning approaches, and can be used even on a very small calibration set; and (2) more powerful than previous methods, which only set the activations' dynamic ranges. We suggest two flavors for our method, parallel and sequential aim for a fixed and flexible bit-width allocation. For the latter, we demonstrate how to optimally allocate the bit-widths for each layer, while constraining accuracy degradation or model compression by proposing a novel integer programming formulation. Finally, we suggest model global statistics tuning, to correct biases introduced during quantization. Together, these methods yield state-of-the-art results for both vision and text models. For instance, on ResNet50, we obtain less than 1% accuracy degradation — with 4-bit weights and activations in all layers, but first and last. The suggested methods are two orders of magnitude faster than the traditional Quantize Aware Training approach used for lower than 8-bit quantization. We open-sourced our code <https://github.com/papers-submission/CalibTIP>.

^{*}Equal contribution ¹Habana Labs – An Intel company, Caesarea, Israel ²Department of Electrical Engineering - Technion, Haifa, Israel. Correspondence to: Itay Hubara <itay-hubara@gmail.com>.

1. Introduction

The pursuit of advanced Deep Neural Networks (DNNs) causes researchers to construct deeper and wider networks, making them expensive to use in terms of power and time. This increases the need for efficient implementations of these networks. Efficient networks reduce cloud-vendor costs and make it possible to run them on low-power devices such as smartphones and wearable devices. The most common off-the-shelf approach to improving network efficiency is quantization, which reduces the numerical precision of the network and its complexity and memory footprint.

DNN quantization techniques can be classified as either post-training or quantization-aware training (QAT) techniques (Han et al., 2015; Courbariaux et al., 2015; Hubara et al., 2017; Zhou et al., 2016). Although QAT techniques, in general, achieve better results, there are important real-world scenarios in which they are not applicable. These are the cases where the training data is sensitive or simply unavailable at the time of deployment. For instance, when off-the-shelf or legacy models are being used, or when medical records are involved. Therefore, much attention has recently been dedicated to post-training quantization methods (Nagel et al., 2019; Banner et al., 2018; Zhao et al., 2019), which can be more easily applied in practice. These methods allow for network quantization to happen seamlessly when deployed, without requiring additional information from the user except a small unlabeled calibration set.

Unfortunately, post-training quantization below 8-bit usually incurs significant accuracy degradation, and in some cases even higher numerical precision is required. In this paper, our goal is to break this barrier by distilling all the information the pre-trained model and calibration set encode. Our goal is to find an optimal scheme for the current state of the art hardware which usually supports 16,8,4 bits data types with per-channel quantization of the weights. To that end, we suggest a three-stage pipeline that consists of methods applied solely on a small calibration set to reduce the local error introduced during the quantization process (e.g., round-off errors) followed by integer programming to determine the bit-width of different layers so that the overall accuracy degradation is minimized. Even without using mixed-precision, the suggested method is much less prone to overfitting than current methods and yields best in

class results for 8-bits Mobilenet-V2 and BERT-base trained on ImageNet and SQuAD1.1 datasets, respectively. Our paper suggests several contributions for mixed-precision post-training quantization:

1. **AdaQuant (section 3.1):** A layer-by-layer optimization method that minimizes the error between the quantized layer output and the full-precision layer output. This method can consume only a small calibration dataset from training data without overfitting. In a comprehensive study, we show that AdaQuant defines a new state-of-the-art for post-training quantization on several networks and tasks, including vision models (Resnet18, Resnet50, MobilenetV2) and language (BERT).
2. **Integer programming (section 3.2):** As some parts of the network may allow lower precision compared to other layers, we suggest an integer-linear programming based approach for determining the precision level of different layers. This method aims at maximizing either the expected speedup or savings in power consumption without violating a predefined constraint on network accuracy degradation or compression.
3. **Para-normalization (section 3.3):** Following quantization, we observe an inherent bias in the mean and the variance of batch norm statistics. We show that by employing the re-estimated statistics in batch normalization, much of the quantized network degradation can be recovered.
4. **Light and Advanced pipelines (section 4):** We analyze the advantages and disadvantages of each of the given methods and suggest two pipelines: (1) light pipeline that does not require a backward pass, thus can be invoked even on inference-only hardware; and (2) Advanced pipeline that includes also AdaQuant and bias tuning.

2. Related work

There has been a significant effort to accelerate inference via quantization (Courbariaux et al., 2015; Han et al., 2015; Rastegari et al., 2016; Zhou et al., 2017). These works involve re-training in order to compensate for the degradation due to the quantization process. Post-training quantization, on the other hand is applied to a model after it was trained. Thus, it avoids re-training and as such it is much simpler to use. However, naively quantizing a full-precision model to INT4 or lower to accelerate computation usually incurs significant accuracy degradation (Krishnamoorthi, 2018; Jacob et al., 2018).

AdaQuant: A recent post-training quantization method (Nagel et al., 2020), termed AdaRound, suggested optimiz-

ing the rounding policy. Instead of using the predominant rounding-to-nearest approach, they suggest formulating a per-layer quadratic optimization problem to optimize the round-off error. Our proposed method, AdaQuant, takes another step and relaxes AdaRound’s implicit constraint which forces the quantized weights to be within ± 1 of their round-to-nearest value. This is done by optimizing the weights and quantization parameters of each layer separately, over the calibration set, to minimize the MSE between the layer’s original and quantized outputs. As oppose to AdaRound we apply AdaQuant to find optimal quantization not only to weights but also to activations. In addition we suggest two flavors for AdaQuant: (1) **parallel-AdaQuant** suited for mixed-precision setting; (b) **sequential-adaquant** which suited for fixed configuration.

Integer programming: Early work by Lin et al. (2016) used a convex optimization formulation which results in a simple greedy compression scheme. Aflalo et al. (2020) used a combinatorial optimization approach for network pruning. Their problem was formulated as a Knapsack problem that optimizes the trade-off between the channels importance and their associated computational cost. Cai et al. (2020) finds a mixed-precision configuration with a guaranteed Pareto efficient allocation with respect to model size and accuracy degradation. While this provides a ”best-effort” standard (e.g., the configuration cannot be further compressed without hurting accuracy), it does not suggest which of all possible outcomes is best. To the best of our knowledge, this work is the first to formalize a generic integer program, which can easily be adapted to various types of models and requirements with a clear objective and constraints.

Batch norm tuning: Finkelstein et al. (2019) were the first to recognize that a significant source of degradation is a shift in the mean activation value. They show a simple method to compensate for this bias by updating the bias terms. Nagel et al. (2019) suggest equalizing the weight ranges in the network and correct biases in the error that are introduced during quantization. Recently Sun et al. (2019) suggested batch norm tuning for FP8 models. Here we detail how to perform this procedure on a per-channel quantized (PCQ) model with fused batch-norm layers. The procedure is light as it only requires to invoke the quantized model a few times (on the calibration set) and adjust the quantization parameters. Moreover after re-tuning the BN layers can be reabsorbed which reduces the inference complexity. To the best of our knowledge, this work is the first to suggest it.

3. Optimizing The Quantization Pipeline

In most post-training quantization settings, a model and a small unlabeled calibration set are given. To avoid overfitting the calibration set, most studies utilize it only to extract

the network’s internal statistics, which is later used to set the quantization parameters.

Here we suggest using the calibration set much more extensively to tune the model while avoiding over-fitting the data. In the following subsections, we detail three different optimization methods over the calibration set: (1) AdaQuant, a layerwise optimization of weights and quantization parameters; (2) an integer programming formulation for a mixed-precision setting; and (3) Para Normalization (PN), for tuning the model’s internal statistics to match the numerical precision setting. We discuss the strengths and weaknesses of each method and suggest an optimization flow that exploits all the additive merits and leads to state-of-the-art results.

3.1. AdaQuant - Layer/Block - wise Optimization over the Calibration Set

A per-tensor optimization for reducing the quantization error have been suggested (Nahshan et al., 2019; Banner et al., 2018; Jacob et al., 2018). Their goal was to find an optimized quantization step size $\hat{\Delta}$ obtained by

$$\begin{aligned} \hat{\Delta} &= \arg \min_{\Delta} \|X - Q_{\Delta}(X)\|^2 \\ Q_{\Delta}(X) &= \Delta \cdot \left\lfloor \frac{X}{\Delta} \right\rfloor, \end{aligned} \quad (1)$$

where $Q(\cdot)$ is the quantization function. Although these methods are fast and easy to use, they often result in an inferior solution — the loss in eq. (1) is sub-optimal, as it penalizes all the quantization errors equally. Whereas it should penalize more quantization errors that affect the classification. Accordingly, researchers (Hubara et al., 2017; Zhou et al., 2016) suggested *Quantization-Aware-Training* (QAT) methods to fix this error by training the entire model at once. However, those methods have three limitations: (a) they require a large training set to avoid over-fitting, (b) they approximate the back-propagation gradients through discrete function (the quantizer), and (c) they have high computational and memory footprints. We suggest a modified objective for per-layer joint optimization of the weights and quantization parameters.

$$\begin{aligned} (\hat{\Delta}_w, \hat{\Delta}_x, \hat{V}) &= \\ \arg \min_{\Delta_w, \Delta_x, V} &\|W'X - Q_{\Delta_w}(W')Q_{\Delta_x}(X)\|^2, \end{aligned} \quad (2)$$

where $W' = W + V$ and V is a continuous variable added to W . The quantized network weights are defined as $W_q = Q_{\hat{\Delta}_w}(W + \hat{V})$. In this new objective, the quantized tensor is not required to be "close" to the original tensor, as in eq. (1), and thus benefits from the flexibility that QAT methods have. Yet, it can be executed in parallel over all layers and is much less prone to over-fitting. Moreover, under a fixed

configuration we can optimize the model globally and infer the error between layers. Thus, instead of running AdaQuant on all layers in parallel we can run it sequentially and fix the error induced by quantizing previous layers. Thus, we suggest a second objective:

$$\begin{aligned} &(\hat{\Delta}_{w_l}, \hat{\Delta}_{x_l}, \hat{V}_l) \\ &= \arg \min_{\Delta_{w_l}, \Delta_{x_l}, V_l} \|W_l X_l - Q_{\Delta_{w_l}}(W'_l) \cdot Q_{\Delta_{x_l}}(X_l^q)\|^2 \quad (3) \\ X_q &= \sigma(Q_{\Delta_{w_{l-1}}}(W_{l-1} + V_{l-1}) \cdot Q_{\Delta_{x_l}}(X_{l-1}^q)), \end{aligned}$$

where $\sigma(\cdot)$ is some activation function.

Note, that sequential AdaQuant should not be applied before the bit allocation was set as it optimizes over noisy inputs obtained from predecessor quantized layers. We evaluate both flavors of adaquant named, *AdaQuant* and *sequential-AdaQuant* and detail our finding in section 5.1. We note that AdaQuant also optimizes over biases and offsets and optimized fused conv-bn-relu layers when present; these were removed from the formulation in Equation 2 for simplicity. In the short time since we release the first version of the paper Yuhang et al. (2021) published their results. They cites our method and extends it to Block-wise optimization. In section 5.1 we apply block optimization as well and compare our results.

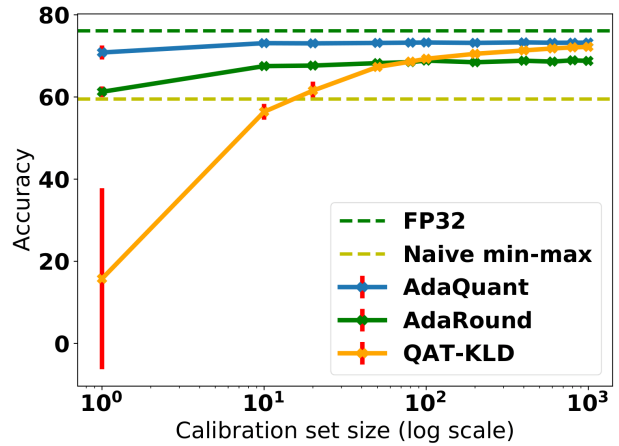


Figure 1. Comparison of different optimization methods over ResNet-50 quantized to 4 bit except the first and the last layers which were kept in 8bit. Even optimizing on a single image drastically improves the results but as expected have a high variance (red bar). The variance decreases rapidly as the calibration set size increases.

Size of calibration set Perhaps surprisingly, although we experiment with a very small calibration set, no over-fitting is observed.

In fig. 1 we compare AdaQuant to current state-of-the-art methods including QAT with knowledge distillation (QAT-KLD) (Kim et al., 2019) and AdaRound (Nagel et al., 2020). For each method, we measured the top-1 accuracy with respect to the number of samples in the calibration set over five runs and present the mean and standard deviation. As can be seen, AdaQuant is superior to previous methods and specifically excels on small calibration sets. Remarkably, AdaQuant does not overfit even when optimized on a single image. Additional details can be found in sections (1) and (4) of the supplementary material.

3.2. Per-layer bit allocations with integer programming

AdaQuant significantly enhances network accuracy at lower bit widths. However, it is often not sufficient by itself to attain acceptable accuracy. Therefore, in practical use cases, the user would like to balance between accuracy and performance (e.g., power and speed), by setting several layers to higher precision. Our high-level goal in this section would be to optimize the overall network performance while maintaining a predefined accuracy degradation or a model compression constraint.

In the following, we provide an integer-programming (IP) formulation for optimizing per-layer bit allocations. Depending on the needs, our performance metrics \mathbb{P} would be either the execution time of the network or its power consumption. Also, with every layer quantization, there is an associated quantization error that affects the training loss \mathcal{L} . We chose the latter to be our penalty metric. Integer programming is applied in those situations where a given problem can clearly be represented in the form of a linear relationship between different decision variables. Unlike other previous works on compression, it attains a **global optimum**. For example, (Lin et al., 2016) suggested a convex optimization problem, but the constraints and the objective are not linear. This typically has a drastic impact on convergence time, and the quality of the results since the Simplex method can no longer be applied (Van Doormaal & Raithby, 1984).

Basic formulation We are given a neural network with L layers. For each layer l , we have weights W_l that need to be multiplied with activations of the previous layer X_{l-1} . Such lower bit width multiplications can be executed by quantizing the weights and activations to achieve higher throughput and energy-efficient solutions. Let W_l^k and X_{l-1}^n represent a quantized version of W_l and X_{l-1} to k and n bits, respectively. For each layer l , a low-bit width multiplication $W_l^k \cdot X_{l-1}^n$ results in a loss degradation $\Delta\mathcal{L}_l^{k,n}$ and in performance improvement $\Delta\mathbb{P}_l^{k,n}$ with respect to the original product $W_l \cdot X_{l-1}$. This performance improvement measure needs to be additive and sum up to a total benefit in end-to-end network performance (e.g., power, model size, etc.).

Our goal would be to maximize the total performance improvement without exceeding the total network degradation $\Delta\mathcal{L}$.

We now turn to solve the above problem using an integer program. We define a binary variable $I_l^{k,n}$, which is set to one if and only if the weights W_l^k are multiplied with the activations X_{l-1}^n at layer l ; otherwise we set the indicator to zero i.e., $I_l^{k,n} = 0$. Then, the basic bit allocation problem can be formulated as follows:

$$\text{Maximize } \sum_{l=0}^{L-1} \Delta\mathbb{P}_l \quad (4a)$$

$$\text{Subject to } \sum_l \Delta\mathcal{L}_l \leq \Delta\mathcal{L}, \quad (4b)$$

$$\begin{aligned} \forall l \in \{1, \dots, L\} : \Delta\mathbb{P}_l &= \sum_{k,n} I_l^{k,n} \cdot \Delta\mathbb{P}_l^{k,n}, \Delta\mathcal{L}_l \\ &= \sum_{k,n} I_l^{k,n} \cdot \Delta\mathcal{L}_l^{k,n} \end{aligned} \quad (4c)$$

$$\forall l \in \{1, \dots, L\} : \sum_{k,n} I_l^{k,n} = 1, I_l^{k,n} \in \{0, 1\} \quad (4d)$$

The objective function (3a) maximizes the total performance improvement. Constraints (3b) and (3c) ensure that the total degradation in loss and the total improvements in performance due to the quantization of layer l to k -bit-weights and n -bit-activations would be $\Delta\mathcal{L}_l$ and $\Delta\mathbb{P}_l$, respectively. Eq (3d) states that the restriction on total degradation of $\Delta\mathcal{L}$ is obeyed and ensures that only one configuration (of quantized weights and activation) per layer is selected.

3.3. Para-Normalization Tuning

A common practice is fusing BN layers into their predecessor weight layers before applying post-training quantization to reduce the amount of Multiply-Accumulate (MAC) operations. However, the reduction in bit-width after quantization can cause the model’s internal statistics to deviate further from those of the full precision model. To compensate for this deviation, we suggest new method, we called *Para-Normalization*¹ for updating the BN statistics.

This Para-Normalization (PN) method has a few steps. First, we need to reconstruct the BN layers then re-tune the BN layers’ statistics (by a few iterations of running-mean to re-collect the statistics). Finally, re-absorb (re-fuse) the BN layers into the weight layers (this is possible only in a per-channel weights quantization setting, which is the current standard). Next, we give more details on each phase.

Reconstructing BN layers Assume the original (pre-fusing) BN parameters γ_o, β_o and ϵ are known, as is usually

¹Name meaning: This ”paranormal” method ”resurrects” the absorbed BN.

the case. We would like to initialize μ, σ^2 , as well as the BN parameters γ_r and β_r (r for "reconstructed") so that the reconstructed BN

$$BN_r(x) = \gamma_r \frac{x - \mu}{\sqrt{\sigma^2 + \epsilon}} + \beta_r \approx x \quad (5)$$

will re-adjust the model statistics. To do so, first we initialize the reconstructed BN layers by setting the following parameters (denoted by r):

$$\mu = \beta_r = \beta_o; \quad \sigma^2 = \gamma_o^2; \quad \gamma_r = \sqrt{\gamma_o^2 + \epsilon} \quad (6)$$

so that $BN_r(x) = x$. Then, we update μ and σ^2 by collecting running mean and running variance on the calibration data. We stress that the BN parameters, γ_r, β_r , do not change while applying PN, as we only invoke forward propagation.

Re-fusing BN layers Due to the per-channel quantization setting we use, the collected statistics can be fused back into the current quantization scale as follows:

$$W'_i = W_i \frac{\gamma_r}{\sigma}; \quad b'_i = \frac{\gamma_r}{\sigma} (b_i - \mu) + \beta_r; \quad \Delta'_{w_i} = \frac{\gamma_r}{\sigma} \Delta_{w_i}$$

Thus, in addition to the regular BN fusion, the quantization step is adjusted by $\gamma_r \sigma^{-1}$. Additional details are given in section (2) of the supplementary material.

Bias tuning Much like Finkelstein et al. (2019), we suggest applying a global bias-tuning procedure on the final mixed-precision model by applying quantization-aware training to minimize the Knowledge Distillation (KD) loss (which does not require labels). Since we restrict the trainable variables to be the biases only, we can train only on the calibration set without experiencing overfitting.

4. Quantization Flow

Past years have seen the rapid development of efficient deployment techniques (Nagel et al., 2019; Haroush et al., 2019). Deployment flows can vary based on the user setting such as hardware constraints, deployment time and task/dataset availability. While some users are willing to pay at initialization the time and effort to gain another fraction of accuracy, others require a simple and fast solution. We address this by suggesting two novel pipelines, light and advanced. Our pipelines are designed to the current, most common setting: per-channel quantization with a small calibration set.

Our *light pipeline* requires three steps: (1) Fuse layers and define quantization parameters; (2) Find optimal mixed-precision configuration using IP; and (3) Use BN tuning to correct the internal statistics. We note that all steps

Table 1. Comparing the light and advanced pipelines building blocks

METHOD/PIPELINE	LIGHT	ADVANCED
ADAQUANT	×	✓
MIXED PRECISION (IP)	✓	✓
BN TUNING	✓	✓
BIAS TUNING	×	✓

do not require back-propagation and thus are very light and fast. In addition to the light setting, in the *advanced pipeline* we apply AdaQuant to reduce each layer’s output distortion from its full precision counterpart before invoking the IP algorithm. A detailed comparison between the two pipelines is given in table-1. Models that were optimized using AdaQuant to different bit-widths can be seamlessly stitched thus having the ability to create an optimized model in a mixed-precision setting. Subsequently, global methods such as tuning both BN statistics and the layers’ biases can be applied to reduce a Knowledge Distillation loss. Although there are additional post-training quantization techniques that could be potentially combined with our methods, such as bias correction (Banner et al., 2018), equalization (Meller et al., 2019), and outlier channel splitting (Zhao et al., 2019), we did not find it necessary: our results demonstrate that our relatively simple pipeline yields state of the art accuracy on both vision and text models, even without combining such methods. In the following sections, we show our findings and give an ablation study that highlights the importance of each method and their combination.

5. Experiments

In this section, we demonstrate our methods and pipelines on several models and datasets. We first start by analyzing image recognition models such as ResNet18/50, MobileNet-V2, which were trained over the ImageNet dataset. Next, we demonstrate our method robustness by applying it on question answering task using the popular BERT model (Devlin et al., 2018), which was fine-tuned on the SQuAD1.1 dataset (Rajpurkar et al., 2016). In all our experiments, we used a small calibration set taken from the training dataset. Unless stated otherwise, we applied asymmetric per-channel quantization (i.e. GEMLOWP (Wu et al., 2016)) with quantized offset (i.e., zero point). Next, we analyze each method’s strengths and weaknesses separately and argue for its validity. Additional implementation details can be found in section and the code are given in sections (5) of the supplementary material.

5.1. AdaQuant

Recently several researchers suggested different types of MSE optimization. In most cases, the optimization was done per-tensor (i.e., for the weights and activations separately).

Here we argue that by optimizing both quantization parameters and the weights jointly we can reduce the MSE even further and hence improve the accuracy as demonstrated in fig. 2. In contrast to AdaRound (Nagel et al., 2020) which restricted the change of the weights to be within ± 1 we allow the weights to change as needed. As can be seen in fig. 3 the weights indeed change their quantized value by more than one. Since our pipeline is focused on the mixed-precision setting we optimize each layer separately to enable maximum flexibility when stitching the optimized models. Under that setting AdaQuant can be performed in parallel across all layers. However, since most recent papers do not show full compression-accuracy curves and only a few attempt 4-bit compression, we also compare our results to common fixed configurations using our *sequential-AdaQuant* flavor. While sequential AdaQuant cannot be parallelized or used for the mixed-precision setting it yields best-in-class results for per-layer optimization on all models tested as can be seen in table-2 and 3. For instance, on the extensively studied 8bit MobileNet-V2 (MobileNet-V2) topology we achieved 71.6% top-1 accuracy — less than 0.5% degradation compared to the full precision counterparts (71.9%). If the system at hand supports more extensive fine-tuning process and the calibration set is large enough, one can optimize per-block (Yuhang et al., 2021). A block is a set of sequential layers. For ResNet we used the "Bottleneck/Basic Residual Block" for BERT we used a "transformer layer". Block sequential-Adaquant minimize the Adaquant objective per-block output by optimizing all the weights and quantization parameters in the block. As expected this approach yields even better result and currently hold the state-of-art accuracy for ResNet-50. Note, that Block Seq AdaQuant as opposed to BRECQ (Yuhang et al., 2021) carries the errors from one layer to the next thus, can potentially fix the quantization errors introduced by previous layers. Using BRECQ baseline model we achieved 75.18% top-1 on ResNet50.

5.2. Integer Programming

Our Integer programming formulation requires us to have two quantities per-layer: (1) loss degradation and; (2) performance improvement. Obtaining those quantities requires invoking the model over a small calibration set L times (once per layer) and measure the loss degradation and the performance gain. In our experiments, we set the performance value to be the number of parameters, but this measure could be changed to any additive measure. In all experiments, we used 1000 samples from the training set as our calibration set. Our setting considers only a mixture of 8-bit and 4-bit layers; to further test IP capabilities, we investigate a mixture of 2-4-8 bits as well. Unfortunately, since 2-bits quantization in post-training setting results in high degradation, the IP algorithm chose only a mixture of 4-8 bits for compression ratio higher than 12.5%. Yet for 12.5% com-

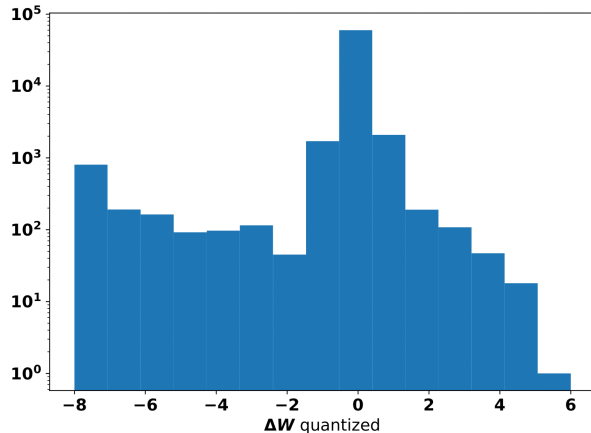


Figure 2. AdaQuant vs. AdaRound. (a) A histogram of ΔW distribution. AdaRound restricts this additive term to be $\Delta W = \pm 1$. Relaxing this constraint provides a more powerful optimization.

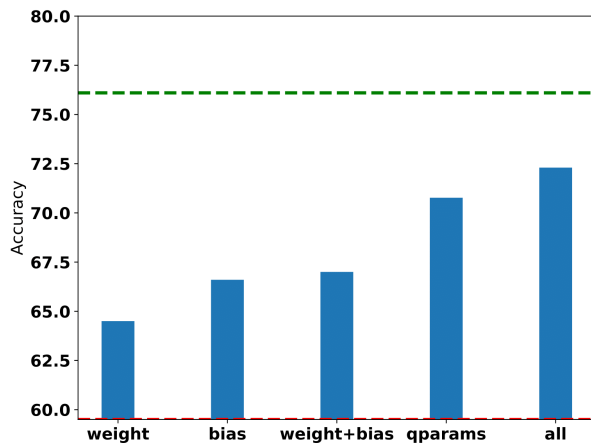


Figure 3. AdaQuant vs. AdaRound. Ablation study on parameters optimization for ResNet50 over ImageNet. AdaRound is based exclusively on weight optimization, while AdaQuant optimizes the weights, biases, and other quantization parameters jointly.

pression ratio, the IP method found that by setting one layer to 2-bits while setting 8 smaller layers to 8-bits accuracy gains over 5.5% with respect to uniform 4-bit quantization. Also, by allowing a less hardware friendly setting where numerical precision can have the form of any integer between 2-8, yields the highest compression-accuracy ratio (fig. 4 - *relaxed advanced pipeline*).

5.3. Para-Normalization

Para-Normalization (PN) has a significant advantage, as it does not require any weight optimization (section 3.3). Since PN is applied by invoking the entire model, we must apply it only after setting the mixed-precision bit-width

Accurate Post Training Quantization With Small Calibration Sets

	RN-18	RN-34	RN-50	RN-101	RNnext-50	Inc-V3
ACIQ* (Banner et al., 2018)	64.5%	69.1%	68.1%	68.1%	68.6%	60.4%
DFQ* (Nagel et al., 2019)	57.1%	59.4%	64.5%	64.6%	65.5%	58.2%
BRECQ [†] (Yuhang et al., 2021)	69.6%	N/A	75.05%	N/A	N/A	N/A
AdaQuant	67.4%	70.3%	73.7%	74.4%	74.0%	72.6%
Sequential-AdaQuant	69.4%	71.7%	75.1%	75.5%	75.6%	73.4%
Block-Sequential-AdaQuant	69.6%	73.2%	75.9%	76.4%	76.6%	74.5%
FP32	71.97%	73.3%	77.2%	77.3%	79.22%	77.4%

Table 2. INT-4 quantization of weights and activations. Top-1 score over ImageNet dataset for different post-training quantization methods. All layers were quantized to 4-bit except the first and last layers which were set to 8-bit. (*) These methods were implemented according to the paper. (†) the results of this method could not be reproduced and extended to other models (their open sourced code has a bug); thus we were only able to add their reported results. In all our experiments we apply per-channel quantization of the weights.

configuration. This is the case for all global optimization methods including bias-tuning. Notably, PN requires only a few (at most 10) forward passes over the calibration set and yield significant gains (fig. 4). In this study, we applied PN on models trained with BN layers only. However, it might be possible to extend this method to models without BN layers by reconstructing them from the statistics. We encourage the reader to investigate this path.

	MobileNet V2 (top-1)	BERT-Base SQuAD1.1 (F1)
min-max	70.9%	87.83%
DFQ (Nagel et al., 2019)	71.2%	N/A
ZeroQ (Cai et al., 2020)	72.91%	N/A
AdaQuant	73.03%	88.35%
Seq-AdaQuant	72.94%	88.45%
Block Seq-AdaQuant	72.96%	88.7%
FP32	73.03%	88.81%

Table 3. INT-8 quantization of weights and activations. For min-max quantization we use average min/max values based on the calibration set (Krishnamoorthi, 2018). A comparison with DFQ and naive quantization methods (which uses the channel’s full dynamic range). In all our experiments we apply per-channel quantization of the weights and quantized all layers to 8-bit.

Testing the strength of this method on both vision and text topologies resulted in state-of-the-art results. As can be seen in table 3, on BERT-base model over SQuAD1.1 dataset (BERT-Base-SQuAD1.1) we managed to obtain 88.45% F1 score using just AdaQuant — less than 0.5% of its full precision counterpart (81.3%). Throughout our experiments, we avoided using any augmentation technique and follow the standard (He et al., 2016) validation set preprocessing.

5.4. Full pipeline and ablation study

Several researchers suggested different methods for post-training mixed-precision quantization yet, none offer their code. Each paper focuses on a different quantization setting

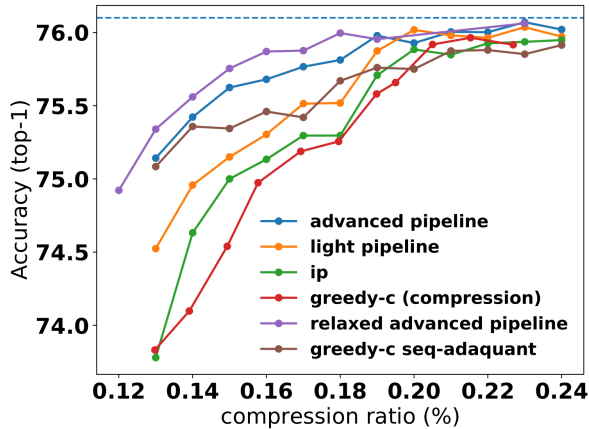
(e.g., quantizing only the weights, per-tensor quantization, etc.). Therefore, to demonstrate our pipeline strength, we created two different baselines based on common practices:

Greedy-accuracy: recent studies suggested measuring each layer sensitivity and, based on the compression target, reduce the precision for the most robust layers.

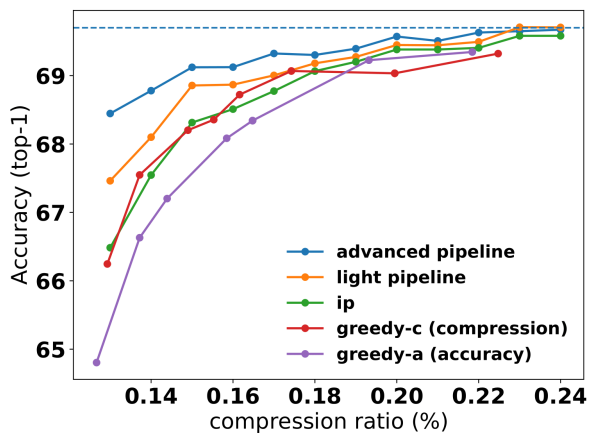
Greedy-compression: the complementary greedy approach (Lin et al., 2016) to sort the layers by their number of parameters and increase the precision of the layers from the smallest to the largest layer until the compression budget is reached.

In fig. 4 we present an ablation study results over ResNet50/18 and MobileNet. As detailed in table 1 our **advanced pipeline** consists of AdaQuant, IP-mixed-precision, para-normalization (PN) and bias-tuning while our **light pipeline** consists of only IP-mixed-precision and PN. For ResNet-50, fig. 4a, we removed the greedy accuracy plot as it yields far inferior results and added two additional plots instead, **greedy-c seq-adaquant** and **relaxed advanced pipeline**. The former decides the bit allocation using a greedy compression scheme and apply on top of it Sequential AdaQuant. The latter is similar to the advanced pipeline but allows the integer-programming to choose any bit-width between 2-8 and not just 4-bit or 8-bit.

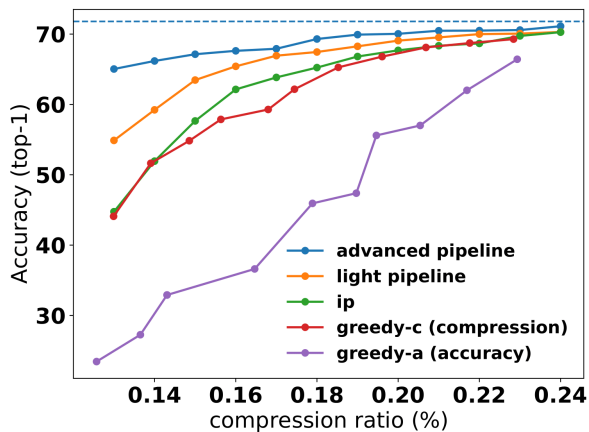
Surprisingly, although the size of the layer should correlate with its sensitivity to quantization, the two greedy methods yield entirely different configurations. Investigating the configuration *greedy-compression* found that sorting by compression correlates with the location of the layers in the model. In most vision models, the layers closer to the input have fewer parameters. This aligns with current common practice (Banner et al., 2018). For high and low compression rates IP has limited power as most of the layers are in 8bit or 4bit. Thus, the interesting part lies in the middle. Fig 4 demonstrates that for ResNet50 and MobileNet-V2: when compression rates are 0.15-0.18, IP results improves by 10over greedy-c which stress its importance.



(a) ResNet-50 ablation study.



(b) ResNet-18 ablation study.



(c) MobileNet V2 ablation study.

Figure 4. Ablation study over ResNet-50/18 and MobileNet-V2 -compression-accuracy curves. The compression ratio is measured as the ratio between the compressed model and the full-precision (32-bit) mode thus 0.25 compression rate indicate that the entire model uses 8-bit precision and respectively for 4-bit the compression rate is 0.125

Next, we turn to consider the light and advanced pipelines. Under challenging compression rates, our light-pipeline results highlight the importance of para-normalization. As can be seen in our experiment fig. 4, by merely invoking the model at inference mode for a few iterations and fixing the intermediate statistics, one can recover more than 1.5% of the accuracy (73.7% v.s 75.37%). As expected, by applying the advanced pipeline, one can obtain state-of-the-art accuracy. Arguably, our most impressive results are at 0.13% compression rate in which we managed to stay within 1% of the full precision accuracy while converting 96% of the model to 4-bit. For the challenging MobileNet-V2 we managed to switch 25% of the layers to 4bit (weights and activations) while maintaining less than 2% degradation; Additionally, we achieved, for the first time, reasonable top-1 accuracy of 65% when almost the entire model is in 4-bit. While one can explore many additional combinations of the suggested method, perhaps the most effortless yet promising one would harness the greedy compression simplicity and sequential-AdaQuant strength. We explore this setting and reported the results in fig. 4a as expected this method mostly benefits the highly compressed models as Sequential-AdaQuant optimizes each layer with respect to the error induced by quantizing former layers. When the model is less compressed the rule of sequential-AdaQuant is less important and the nonoptimal bit-allocation harms the results.

6. Conclusion

Most techniques for obtaining quantized neural networks require the entire training set. This approach can be impractical due to a lack of data or a time-consuming process. In this paper, we attempt to open the door to enabling 4-bits post-training-quantization of Deep Neural Networks with a small, unlabeled calibration set. The performance gain on currently available hardware using 4-bits instead of 8-bits can be significant, even if applied only for a subset of the network’s layers. Hence, we believe this is a highly practical scenario. We present a novel method, named AdaQuant and its flavor sequential-AdaQuant both require only a small calibration set and enable 4-bit vision models to achieve near full-precision accuracy. In addition, we demonstrate how classical integer programming formulation leads to optimal bit allocation solution. Finally, we suggest two pipelines, light and advanced. The former does not require further training of the model and thus can be executed on a computationally bound hardware (light pipeline). The latter performs slight model optimization, but does not require high memory or compute capabilities and does not prone to over-fitting. All our methods are relatively light and require only few minutes (see sec 4.4 of the supplementary material). As an example, for ResNet50, even the most time-consuming version, seq-AdaQuant, takes less than 5

minutes on one device (GeForce 1080). Thus, our methods are two orders of magnitude faster than QAT. We believe further research in that path would ensure 4-bit data type to be a first-class-citizen for post training quantization.

7. Acknowledgements

The research of DS was supported by the Israel Science Foundation (grant No. 1308/18), and by the Israel Innovation Authority (the Avatar Consortium).

References

- Aflalo, Y., Noy, A., Lin, M., Friedman, I., and Zelnik, L. Knapsack pruning with inner distillation. *arXiv preprint arXiv:2002.08258*, 2020.
- Banner, R., Nahshan, Y., Hoffer, E., and Soudry, D. Acicq: Analytical clipping for integer quantization of neural networks. 2018.
- Cai, Y., Yao, Z., Dong, Z., Gholami, A., Mahoney, M. W., and Keutzer, K. Zeroq: A novel zero shot quantization framework. *arXiv preprint arXiv:2001.00281*, 2020.
- Courbariaux, M., Bengio, Y., and David, J.-P. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pp. 3123–3131, 2015.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Finkelstein, A., Almog, U., and Grobman, M. Fighting quantization bias with bias. *arXiv preprint arXiv:1906.03193*, 2019.
- Han, S., Mao, H., and Dally, W. J. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- Haroush, M., Hubara, I., Hoffer, E., and Soudry, D. The knowledge within: Methods for data-free model compression. *arXiv preprint arXiv:1912.01274*, 2019.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hubara, I., Courbariaux, M., Soudry, D., El-Yaniv, R., and Bengio, Y. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- Jacob, B., Kligys, S., Chen, B., Zhu, M., Tang, M., Howard, A., Adam, H., and Kalenichenko, D. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2704–2713, 2018.
- Kim, J., Bhargat, Y., Lee, J., Patel, C., and Kwak, N. Qkd: Quantization-aware knowledge distillation. *arXiv preprint arXiv:1911.12491*, 2019.
- Krishnamoorthi, R. Quantizing deep convolutional networks for efficient inference: A whitepaper. *arXiv preprint arXiv:1806.08342*, 2018.
- Lin, D., Talathi, S., and Annapureddy, S. Fixed point quantization of deep convolutional networks. In *International conference on machine learning*, pp. 2849–2858, 2016.
- Meller, E., Finkelstein, A., Almog, U., and Grobman, M. Same, same but different-recovering neural network quantization error through weight factorization. *arXiv preprint arXiv:1902.01917*, 2019.
- Nagel, M., Baalen, M. v., Blankevoort, T., and Welling, M. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 1325–1334, 2019.
- Nagel, M., Amjad, R. A., van Baalen, M., Louizos, C., and Blankevoort, T. Up or down? adaptive rounding for post-training quantization. *arXiv preprint arXiv:2004.10568*, 2020.
- Nahshan, Y., Chmiel, B., Baskin, C., Zheltonozhskii, E., Banner, R., Bronstein, A. M., and Mendelson, A. Loss aware post-training quantization. *arXiv preprint arXiv:1911.07190*, 2019.
- Rajpurkar, P., Zhang, J., Lopyrev, K., and Liang, P. Squad: 100,000+ questions for machine comprehension of text. *arXiv preprint arXiv:1606.05250*, 2016.
- Rastegari, M., Ordonez, V., Redmon, J., and Farhadi, A. Xnor-net: Imagenet classification using binary convolutional neural networks. In *European conference on computer vision*, pp. 525–542. Springer, 2016.
- Sun, X., Choi, J., Chen, C.-Y., Wang, N., Venkataramani, S., Srinivasan, V. V., Cui, X., Zhang, W., and Gopalakrishnan, K. Hybrid 8-bit floating point (hfp8) training and inference for deep neural networks. In *Advances in Neural Information Processing Systems*, pp. 4901–4910, 2019.
- Van Doormaal, J. P. and Raithby, G. D. Enhancements of the simple method for predicting incompressible fluid flows. *Numerical heat transfer*, 7(2):147–163, 1984.

Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., et al. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*, 2016.

Yuhang, L., Ruihao, G., Xu, T., Yang, Y., Peng, H., Qi, Z., Fengwei, Y., Wei, W., and Shi, G. Brecq: Pushing the limit of post-training quantization by block reconstruction. In *ICLR*, 2021. URL <https://openreview.net/forum?id=POWv6hDd9XH>.

Zhao, R., Hu, Y., Dotzel, J., De Sa, C., and Zhang, Z. Improving neural network quantization without retraining using outlier channel splitting. In *International Conference on Machine Learning*, pp. 7543–7552, 2019.

Zhou, A., Yao, A., Guo, Y., Xu, L., and Chen, Y. Incremental network quantization: Towards lossless cnns with low-precision weights. *arXiv preprint arXiv:1702.03044*, 2017.

Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., and Zou, Y. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.