

Quantization Aware Factorization for Deep Neural Network Compression

Daria Cherniuk

Stanislav Abukhovich

Anh-Huy Phan

Ivan Oseledets

Andrzej Cichocki

Skolkovo Institute of Science and Technology, Russia

DARIA.CHERNIUK@SKOLTECH.RU

STANISLAV.ABUKHOVICH@SKOLTECH.RU

A.PHAN@SKOLTECH.RU

I.OSELEDETS@SKOLTECH.RU

A.CICHOCKI@SKOLTECH.RU

Julia Gusak

Inria, University of Bordeaux, France

JULGUSAK@GMAIL.COM

Abstract

Tensor decomposition of convolutional and fully-connected layers is an effective way to reduce parameters and FLOP in neural networks. Due to memory and power consumption limitations of mobile or embedded devices, the quantization step is usually necessary when pre-trained models are deployed. A conventional post-training quantization approach applied to networks with decomposed weights yields a drop in accuracy. This motivated us to develop an algorithm that finds tensor approximation directly with quantized factors and thus benefit from both compression techniques while keeping the prediction quality of the model. Namely, we propose to use Alternating Direction Method of Multipliers (ADMM) for Canonical Polyadic (CP) decomposition with factors whose elements lie on a specified quantization grid. We compress neural network weights with a devised algorithm and evaluate its prediction quality and performance. We compare our approach to state-of-the-art post-training quantization methods and demonstrate competitive results and high flexibility in achieving a desirable quality-performance tradeoff.

1. Introduction

Recent years more and more applications based on deep learning algorithms appear on the market, and many of them are developed to work on mobile and edge devices. However, neural networks which are in the core of such algorithms typically cannot be utilized in their original form, as they often fail to meet the memory and energy constraints of devices or do not execute the inference phase rapidly enough to fulfill user requirements. Thus, many approaches to reduce and accelerate neural networks have been proposed in the literature.

Conventionally, these approaches include pruning (Dong, Li, Tang, Liu, Pan, Wang, & Chu, 2024), tensor/factorization methods, quantization, knowledge distillation (Gou, Yu, Maybank, & Tao, 2021), and architecture search (White, Safari, Sukthanker, Ru, Elsken, Zela, Dey, & Hutter, 2023). Also, a significant body of research has specifically focused on training neural networks with mobile devices, primarily through approaches such as federated learning combined with adaptive communication techniques (Tang, Shi, Li, & Chu, 2023a), (Tang, Shi, Wang, Li, & Chu, 2023b).

The availability of additional information, such as training data or loss/gradients values, yields additional opportunities to improve performance of compressed models. From this point of view, we

distinguish among techniques that do not use additional data, use data but do not perform training, and those that run several training epochs to tune parameters of the compressed model.

In our paper we propose a method that jointly performs factorization and quantization of neural networks. Namely, we replace a convolutional layer represented with parameters in float32 format with a decomposed layer, which is constructed as a sequence of convolutional layers with weights in lower precision format (e.g., int8, int6, int4). By applying our method we benefit from reduction in number of parameters and operations due to factorization and further model size reduction and acceleration due to lower bit representation of weights.

There are many papers related to neural network speed-up and compression. The following subsections review the most relevant to our proposed technique:

1.1 Factorization

Various tensor decompositions (Kolda & Bader, 2009), (Cichocki, Lee, Oseledets, Phan, Zhao, & Mandic, 2016) are used to factorize weights (Panagakis, Kossaifi, Chrysos, Oldfield, Nicolaou, Anandkumar, & Zafeiriou, 2021), (Kossaifi, Bulat, Tzimiropoulos, & Pantic, 2019), (Gusak, Kholiavchenko, Ponomarev, Markeeva, Blagoveschensky, Cichocki, & Oseledets, 2019) or activations (Cui, Zhang, Daulbaev, Gusak, Oseledets, & Zhang, 2020), (Gusak, Daulbaev, Ponomarev, Cichocki, & Oseledets, 2021) of deep neural networks. These approaches allow to store less parameters during training and inference and to speedup the computation by reducing the number of required elementary operations. For linear layers, prior works have utilized Tensor-Train (TT) format (Novikov, Podoprikin, Osokin, & Vetrov, 2015). For convolutional layers, whose weights are 4D tensors, authors in (Lebedev, Ganin, Rakhuba, Oseledets, & Lempitsky, 2015) used CPD. Authors in (Phan, Sobolev, Sozykin, Ermilov, Gusak, Tichavský, Glukhov, Oseledets, & Cichocki, 2020) also decomposed convolutional layers through CP-decomposition but proposed to flatten spatial dimension and factorize 3D tensors instead of the original 4D. Other tensor approximations techniques like Tucker (Kim, Park, Yoo, Choi, Yang, & Shin, 2016) or block term decompositions are used by researches as well. In our paper, we focus on CP decomposition for convolutional layers with kernel spatial dimensions greater than 1 and matrix factorization for linear layers or 1x1 convolutional layers. We prefer to study these type of decompositions to others as they show promising accuracy/speed-up trade-off in existing research papers.

1.2 Quantization

Due to memory/power consumption limitations of edge devices, the quantization step is necessary, when pre-trained models are deployed. A comprehensive overview of papers on neural network quantization can be found in (Gholami, Kim, Dong, Yao, Mahoney, & Keutzer, 2022) and (Nagel, Fournarakis, Amjad, Bondarenko, van Baalen, & Blankevoort, 2021).

Quantization might greatly impede the quality of a neural network prediction. There are different techniques to restore the quality, some of them requiring additional data or training. Best accuracy restoration can be obtained with quantization-aware training: fine-tuning the quantized model while restricting model weights to be on the quantization grid. Another approach, taken in (van Baalen, Louizos, Nagel, Amjad, Wang, Blankevoort, & Welling, 2020), is to learn bit-width and quantization parameters simultaneously with model weights. However, both techniques require dataset and resources to train.

Because of training expenses, methods that allow for more accurate quantization but avoid training/fine-tuning have been a research interest in several papers: authors in (Eldad, Alexander, Uri, & Mark, 2019) scale the weights of consecutive layers which allows them achieve smaller quality drop on a wide variety of networks; AdaRound (Nagel, Amjad, Van Baalen, Louizos, & Blankevoort, 2020) proves that rounding to the nearest node of the quantization grid is not the best strategy and propose to choose rounding by optimization process which requires only a few thousand samples of data, researches in (Hubara, Nahshan, Hanani, Banner, & Soudry, 2020) further develop the idea of AdaRound, complementing it with integer programming to determine the best bit-width for each layer.

1.3 ADMM

Alternating Direction Method of Multipliers or ADMM (Boyd, Parikh, Chu, Peleato, & Eckstein, 2010) is an optimization algorithm for convex problems. The advantages of ADMM algorithm are that it can be efficiently parallelized (Shang, Liu, & Cheng, 2014), (Liavas & Sidiropoulos, 2015) and the convergence is guaranteed in convex case (Boyd et al., 2010). However, for non-convex problems the algorithm is not guaranteed to converge or might not converge to a global minimum. Authors in (Diamond, Takapoui, & Boyd, 2016) proposed using ADMM in conjunction with local improvement methods to find sufficient heuristic solutions.

In (Huang, Sidiropoulos, & Liavas, 2016) authors have applied ADMM to a task of finding a CP-decomposition that satisfied particular constraints like non-negativity, sparsity, etc.. They called this approach AO-ADMM, since one of the objectives was an Alternating Optimization (in particular, Alternating Least Squares) that is used to find CP factors.

Authors in (Yin, Sui, Liao, & Yuan, 2021) have embedded ADMM into training neural network to impose low TT-rank structure on its weights. After training, the weights are decomposed into a TT-format and fine-tuned.

(Leng, Dou, Li, Zhu, & Jin, 2018) aimed to apply ADMM to training neural network with weights that lie on a grid of scaled powers of two. The first objective function in their case is model prediction loss, the second - projection of weight tensors on the grid of scaled powers of two. The authors solve a discrete non-convex constraint problem by alternating between optimizing a scaling factor and projecting onto a fixed grid. The method converges, however, the projection even on the unscaled grid is a non-convex problem and is not guaranteed to converge to global minimum.

Contributions. Our main contributions are the following:

- We introduce an ADMM-based algorithm, which approximates a tensor with a CP-factorization, where factors are represented in lower precision format, e.g., int8, int6, int4.
- We propose a neural network compression method that simultaneously performs factorization and quantization of weight tensors based on introduced tensor approximation technique. As far as we are concerned, this is the first work that introduces weights approximation by CP-decomposition with lower precision factors.
- We show that in terms of compression/accuracy trade-off our approach performs better than conventional tensor approximation methods followed by quantization of factors.
- We demonstrate the effectiveness of our method in comparison with other quantization approaches that are applicable when only a small set of training data is available.

Our source code and experiments are available as an open github repository¹.

2. Methodology

2.1 Convolutional Layer Factorization

Similar to (Lebedev et al., 2015) and, later, (Phan et al., 2020), we represent convolution layer weights as a 3-way tensor by flattening spatial dimension and decompose it into 3 consecutive convolutional layers. Let $\mathbf{K} \in \mathbb{R}^{T \times S \times D \times D}$ be a kernel tensor, corresponding to convolutional layer with S input channels, T output channels, and $D \times D$ spatial convolution. Let $\overline{\mathbf{K}} \in \mathbb{R}^{T \times S \times D^2}$ denote the \mathbf{K} tensor after reshape operation. Using rank- R CP decomposition, one element of $\overline{\mathbf{K}}$ can be represented as:

$$\overline{\mathbf{K}}(t, s, dd) \cong \sum_{r=1}^R \overline{\mathbf{K}}^{dd}(dd, r) K^s(s, r) K^t(t, r) \quad (1)$$

Then, in order to reshape convolutional layer back to original shape, we reshape factor-matrix $\overline{\mathbf{K}}^{dd}$ to a tensor \mathbf{K}^{dd} of shape $D \times D \times R$.

$$\mathbf{K}(t, s, j, i) \cong \sum_{r=1}^R \mathbf{K}^{dd}(j, i, r) K^s(s, r) K^t(t, r) \quad (2)$$

Therefore, having an input, X , to the convolutional layer, the output tensor, Y , is calculated as:

$$\begin{aligned} \mathbf{Y}(t, h', w') = & \\ & \sum_{h=h'-\delta}^{h'+\delta} \sum_{w=w'-\delta}^{w'+\delta} \sum_{s=1}^S \mathbf{K}(t, s, h - h' + \delta, w - w' + \delta) \mathbf{X}(s, h, w) \end{aligned} \quad (3)$$

Substituting kernel expression into the formula above, performing rearrangements and grouping summands, we obtain the following consecutive expressions for the approximate evaluation of the convolution

$$\begin{aligned} \mathbf{Z}^1(r, h, w) &= \sum_{s=1}^S K^s(s, r) \mathbf{X}(s, h, w), \\ \mathbf{Z}^2(r, h', w') &= \sum_{h=h'-\delta}^{h'+\delta} \sum_{w=w'-\delta}^{w'+\delta} \mathbf{K}^{dd}(h - h' + \delta, w - w' + \delta, r) \mathbf{Z}^1(r, h, w), \\ \mathbf{Y}(t, h', w') &= \sum_{r=1}^R K^t(t, r) \mathbf{Z}^2(r, h', w'), \end{aligned}$$

where $\delta = D/2$.

Therefore, we substitute original convolution with a sequence of smaller convolutions: point-wise convolution that reduces the number of input channels from S to R , convolution with the

1. <https://github.com/KamikaziZen/admm-quantization>

same spatial dimension as the original but with R number of input and output channels, and another point-wise convolution that changes the number of channels from R to T . The last convolution in sequence adds the original bias.

1x1 kernel size. Convolutional layer with 1x1 kernel is equivalent to a Linear layer with weight matrix of shape $S \times T$ acting on the input with flattened spatial dimension and, thus, can be decomposed into two matrix factors and replaced with a sequence of two convolutional layers both with kernel size 1x1.

2.2 Quantization

Quantizing of neural network means transforming its weights and/or activations into low-bit fixed-point representations, e.g., int8. It saves memory occupied by model parameters, reduces the amount of data transfer, size and energy consumption of the MAC operation (Nagel et al., 2021).

In our work we use both signed symmetric and asymmetric uniform per-tensor quantization (Nagel et al., 2021) in which tensors are mapped to their int versions in the following way:

$$x_{\text{int}} = \text{clip} \left(\left\lfloor \frac{x}{\text{scale}} \right\rfloor + z; -2^{b-1}, 2^{b-1} - 1 \right), \quad (4)$$

where $\lfloor x \rfloor$ denotes rounding x to the nearest integer, scale is a step of the quantization grid, z is a zero-point and b is a number of bits in quantization. In case of symmetric quantization z is assumed to be 0. Real-valued approximation is obtained by extracting zero-point and multiplying by scale: $x \approx \text{scale} * (x_{\text{int}} - z)$. For per-tensor MinMax quantization (Nagel et al., 2021), scale is determined from minimum and maximum values of tensor X :

$$\text{scale} = \frac{\max\{X\} - \min\{X\}}{2^b - 1} \quad (5)$$

MinMax quantization, however, might suffer from large outliers. A way to alleviate this issue is to use MSE approach to set range of values (Nagel et al., 2021). In case of symmetric quantization (to reduce number of optimized parameters from 2 to 1 and provided that the value distribution is symmetric, as shown in Figure 2), scale calculation takes the following form:

$$\text{scale} = \frac{2q_{\max}}{2^b - 1} \quad (6)$$

and q_{\max} is determined through optimizing

$$\arg \min_{q_{\max}} \|X - \hat{X}(q_{\max})\|_F^2 \quad (7)$$

where \hat{X} is a quantized tensor X . In Section 3, we perform an ablation study to determine the best quantization scheme (Figure 2).

2.3 Quantization-aware Factorization

Finding a decomposition that produces less error after quantization can be formulated as a constrained tensor factorization problem. A method called AO-ADMM(a hybrid of the Alternating Optimization and the Alternating Direction Method of Multipliers) has been shown (Huang et al.,

2016) to successfully solve problems that involve non-negativity, sparsity or simplex constraints. In our work we extend the field of its application by introducing a constraint function that ensures low quantization error of the derived factors and construct a corresponding algorithm.

Searching for a factorization of $n \times m$ matrix X that satisfies quantization constraint (i.e. obtained factors are equal to their quantized versions) means minimizing the following function:

$$\underset{A,B}{\text{minimize}} \frac{1}{2} \|X - AB^T\|_F^2 + I_Q(A) + I_Q(B) \quad (8)$$

where A is a matrix of shape $n \times r$, B - of shape $r \times m$ ($r \leq n, m$) and I_Q is an indicator function such that $I_Q(Y) = 0$ when $Y \in Q$ and $I_Q(Y) = +\infty$ if $Y \notin Q$. In case of symmetric per-tensor quantization (Nagel et al., 2021), Q is a set of tensors whose elements belong to a discrete quantization grid $\{(-2^{b-1}) * \text{scale}, \dots, (2^{b-1} - 1) * \text{scale}\}$.

We introduce an auxiliary variable, \tilde{B} , and formulate a constrained Least Squares problem, which can be solved using alternating update scheme by fixing one variable and minimizing the objective function over the other. For fixed factor, A , the subproblem takes the following form:

$$\begin{aligned} & \underset{B, \tilde{B}}{\text{minimize}} \frac{1}{2} \|X - A\tilde{B}\|_F^2 + I_Q(B) \\ & \text{subject to } B = \tilde{B}^T \end{aligned} \quad (9)$$

The ADMM method introduces a dual variable, U , for the equality constraint, $B = \tilde{B}^T$, and alternates between optimizing each part of the objective function:

$$\tilde{B} = \arg \min_{\tilde{B}} \left(\frac{1}{2} \|X - A\tilde{B}\|_F^2 + \frac{\rho}{2} \|B - \tilde{B}^T + U\|_F^2 \right) \quad (10)$$

$$B = \arg \min_B \left(I_Q(B) + \frac{\rho}{2} \|B - \tilde{B}^T + U\|_F^2 \right) \quad (11)$$

$$U = U + B - \tilde{B}^T \quad (12)$$

The optimal solution for the auxiliary variable \tilde{B} is given in closed form:

$$\tilde{B} = (A^T A + \rho I)^{-1} (X^T A + \rho(B + U))^T \quad (13)$$

For efficiency, (13) is solved through Cholesky decomposition of $A^T A + \rho I$.

Solving (10) for the factor, B , gives the minimum distance between $(\tilde{B}^T - U)$ and set Q :

$$B = \underset{B \in Q}{\text{argmin}} \|B - \tilde{B}^T + U\|_F^2 = \text{proj}_Q(\tilde{B}^T - U) \quad (14)$$

Overall, the ADMM procedure for the factor, B , is defined in Algorithm 1. For $X \approx A\tilde{B}^T$ matrix factorization G is a Gram matrix of the factor that is fixed on current iteration of Alternating Least Squares and K is an original matrix multiplied by a fixed factor ($A^T A$ and $X^T A$ from (13)). We adopt expressions for r , s and ρ from (Huang et al., 2016). Similarly, the task is solved for the remaining factor through minimizing $\|X - \tilde{A}^T B^T\|_F^2$ over set Q . The procedure continues alternating between factors until convergence.

Algorithm 1 Solve 9 using ADMM

Input: B, U, K, G, rank

Output: B, U

- 1: $\rho \leftarrow \text{trace}(G)/\text{rank}$
 - 2: $L \leftarrow \text{Cholesky}(G + \rho I)$
 - 3: **repeat**
 - 4: $\tilde{B} \leftarrow L^{-T}L^{-1}(K + \rho(B + U))^T$
 - 5: $B_0 \leftarrow B$
 - 6: $B \leftarrow \text{proj}_Q(\tilde{B}^T - U)$
 - 7: $U \leftarrow U + B - \tilde{B}^T$
 - 8: $r \leftarrow \|B - \tilde{B}^T\|_F^2 / \|B\|_F^2$
 - 9: $s \leftarrow \|B - B_0\|_F^2 / \|U\|_F^2$
 - 10: **until** $r < \epsilon$ and $s < \epsilon$
-

For 3-way tensors, the matrix multiplication of two factors in (8) is replaced with the Canonical Polyadic decomposition (CP):

$$\underset{A,B,C}{\text{minimize}} \left(\frac{1}{2} \|\mathbf{X} - \sum_{r=1}^R A_{:,r} \otimes B_{:,r} \otimes C_{:,r}\|_F^2 + I_Q(A) + I_Q(B) + I_Q(C) \right) \quad (15)$$

where \mathbf{X} is a tensor of shape $n \times m \times k$, A is a matrix of shape $n \times r$, B - matrix of shape $m \times r$, C - matrix of shape $k \times r$, \otimes denotes the outer product. The constrained ALS subproblem for factor B takes the following form:

$$\begin{aligned} & \underset{B, \tilde{B}}{\text{minimize}} \frac{1}{2} \|\mathbf{X}_{(2)} - \tilde{B}^T (C \odot A)^T\|_F^2 + I_Q(B) \\ & \text{subject to } B = \tilde{B}^T \end{aligned} \quad (16)$$

where $\mathbf{X}_{(2)}$ is a matrix unfolding of tensor \mathbf{X} by mode 2, \odot denotes Khatri-Rao product. The ADMM procedure is depicted in Algorithm 2, with $G = (A^T A) * (C^T C)$ and $K = \mathbf{X}_{(2)}(C \odot A)$ (Smith, Beri, & Karypis, 2017). Optimization with respect to factors A and C is done in the same manner (Smith et al., 2017). Alternating between factors proceeds until convergence.

The set Q is non-convex, thus, it is not guaranteed that solution converges to global or even local minimum. Moreover, as shown in Figure 1, convergence depends on initialization. Authors in (Diamond et al., 2016) proposed to use Multistart, Polishing and Neighbor search to improve solution. However, in our case, search among neighbors (even sampling search) of current iteration factor would require number of computations comparable to the number of parameters in a layer. Polishing on Q reduces to simple recomputation of \tilde{B}_{k+1} with B_{k+1} instead of B_k since the only convex restriction (Diamond et al., 2016) that can be chosen in this case is the trivial one. Carefully choosing initialization, however, proved to be a plausible technique to obtain a better solution (Figure 1).

Algorithm 2 Solve (15) using ADMM-EPC

Input: \mathbf{X} , rank
Output: A, B, C

- 1: Initialize: A, B, C with CPD-EPC
- 2: Initialize: U_A, U_B, U_C with 0
- 3: **repeat**
- 4: $G \leftarrow B^T B * C^T C$
- 5: $K \leftarrow \mathbf{X}_{(1)}(C \odot B)$
- 6: $A, U_A \leftarrow \text{ADMM}(A, U_A, K, G, \text{rank})$
- 7:
- 8: $G \leftarrow A^T A * C^T C$
- 9: $K \leftarrow \mathbf{X}_{(2)}(C \odot A)$
- 10: $B, U_B \leftarrow \text{ADMM}(B, U_B, K, G, \text{rank})$
- 11:
- 12: $G \leftarrow A^T A * B^T B$
- 13: $K \leftarrow \mathbf{X}_{(3)}(B \odot A)$
- 14: $C, U_C \leftarrow \text{ADMM}(C, U_C, K, G, \text{rank})$
- 15: **until** e_{quant} (17) ceases to improve

3. Results

We use pretrained ResNet18 model from torchvision model zoo² and ImageNet³ dataset to evaluate our method. First, we conduct several ablation studies to determine quantization scheme, initialization, convergence properties e.t.c.. Then, we compare variations of our method with state-of-the-art approaches from previous papers as well as a naive approach of directly quantizing factors. In all our experiments we utilize per-tensor quantization. For fair comparison, we only consider methods that do not perform model fine-tuning and require only a small sub-sample of the training dataset. **Metrics.** In some ablation studies we compute a quantized reconstruction error metric:

$$e_{\text{quant}} = \frac{\|\mathbf{X}_{(2)} - \hat{B}(\hat{C} \odot \hat{A})^T\|_F}{\|\mathbf{X}_{(2)}\|_F} \quad (17)$$

where $\hat{A}, \hat{B}, \hat{C}$ are quantized factors of tensor \mathbf{X} . $\mathbf{X}_{(2)}$ is a matrix unfolding by mode 2. We could use any other unfolding, the purpose of this notation is merely to address the fact that Frobenius norm is formulated for matrices.

Taking into consideration benefits of reducing both MAC operations (factorization) and operand's bit-width (quantization), we adopt BOP metric, introduced in (van Baalen et al., 2020) and use it to compare our method with other approaches. For layer l , BOP count is computed as follows:

$$\text{BOPs}(l) = \text{MACs}(l)b_w b_a \quad (18)$$

where b_w is a bit-width of weights and b_a is a bit-width of (input) activations. While computing this metric for the whole model we take into consideration all layers including BatchNorms.

2. <https://pytorch.org/vision/stable/models.html>

3. <https://www.image-net.org/>

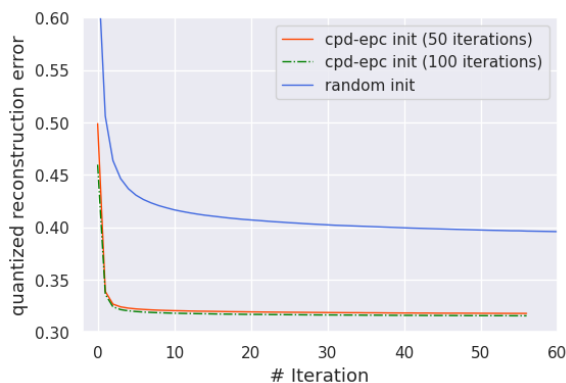


Figure 1: Convergence of ADMM algorithm for a single layer of ResNet18 model with random and CPD-EPC initialization. We have found that using CPD-EPC as initialization instead of random sampling leads to lower quantized reconstruction error, while only a few number of CPD-EPC iterations is sufficient. This result reproduces throughout all model layers.

Factorization ranks. To choose a rank for each layer factorization, we set up a parameter reduction rate (i.e. the ratio of the number of the factorized layer’s parameters over that of the original layer) and define the rank as $N/(n + m)/rate$ for fully-connected layers and $N/(n + m + k)/rate$ for convolutions. N is the number of parameters in the original layer, n , m and k - dimensions of reshaped convolution weights (Section 2.1), $rate$ denotes parameter reduction rate.

BatchNorm calibration. Since Resnet models have a lot of BatchNorm layers, and both factorization and quantization disturb the distribution of outputs, calibrating BatchNorm layers’ parameters (performing inference with no gradients or activations accumulation) can boost model prediction quality. The same procedure was used, for example, in AdaQuant(Hubara et al., 2020). In all our experiments we perform calibration on 2048 samples from training dataset. It is the same number of samples as used in AdaRound(Nagel et al., 2020) and AdaQuant.

3.1 Ablation Studies

CP factors quantization. First, we try to perform factorization and quantization successively. We obtain factors using CPD and CPD-EPC (Phan et al., 2020) and compare results after quantizing weights while leaving activations in full precision (Table 1). Although CPD factors provide close to unfactorized model accuracy, they degrade to very poor quality of predictions upon quantization. CPD-EPC factors provide better model accuracy after quantization, while having lower accuracy before quantization. This is explained by the fact that CPD-EPC minimizes Sensitivity (Tichavský, Phan, & Cichocki, 2019) of a decomposition that results in factors with lower value range than that of the classic CPD (Figure 2). It alleviates the problem of large outliers for uniform quantization.

We have also found (Figure 1) that using a few iterations of CPD-EPC as initialization of ADMM algorithm decreases quantized reconstruction error (17) by keeping value distribution narrowly around zero (Figure 2) which ensures a better quality of model predictions after quantization, especially when using symmetric quantization schemes. We have observed that only a few iterations of cpd-epc is sufficient, after which accuracy of a quantized model ceases to improve.

#bits W/A	CPD	CPD-EPC
32/32	69.26	66.65
8/32	0.1	66.5
6*/32	0.1	65.47
4*/32	0.1	37.54

Table 1: Comparison (% top-1 test accuracy) between CPD and CPD-EPC decompositions of a ResNet18 model(with reduction rate 2) before and after per-tensor MinMax quantization of weights. *We quantize all factorized layers to a specified bit-width, but leave batchnorm, downsample, first and last layers in 8 bits.

Quantization scheme. We conducted an ablation study to choose between MinMax and MseMinMax quantization schemes to be used in a projection step of ADMM algorithm (14). Results are presented in Table 2. Although MSEMinMax slightly increases the overall computation time, it produces consistently better results due to scale adjustment. Hence, we use MSEMinMax scheme in all subsequent experiments.


#bits W/A	MinMax	MseMinMax
8/8	67.99	68.56
6*/8	67.51	68.41
4*/8	62.66	67.75

Table 2: Comparison (% top-1 test accuracy) between per-tensor MinMax and MSEMinMax quantization schemes used in projection step of ADMM iteration (the algorithm was run with reduction rate 2). Post-factorization quantization of weights is performed with per-tensor MSEMinMax scheme in both cases. Quantization of activations is performed by histogram observer evaluated on 500 samples. *We quantize all factorized layers to a specified bit-width, but leave batchnorm, downsample, first and last layers in 8 bits.


Bit-allocation. Determining the best trade-off between bit-allocation and reduction rate is not a trivial task. We performed several ablation studies (Figures 4a and 4b) to choose between different configurations and came to the conclusion that our ADMM-based approach benefits most from lower bit-width, while for conventional decomposition followed by quantization, the best results were achieved with 6-bit weight quantization.

3.2 Baselines


We use top-1 test accuracy and BOP count (18) to compare our method to other post-quantization approaches: AdaRound(Nagel et al., 2020), AdaQuant(Hubara et al., 2020), AdaRound applied to factors obtained with CPD-EPC decomposition and naive approach of successive CPD-EPC factorization and quantization.




(a) CPD



(b) CPD-EPC



(c) ADMM with CPD-EPC initialization



(d) ADMM with random initialization

Figure 2: Histograms of values of factor matrix obtained with different decomposition methods from a single layer (layer1.0.conv1) of the ResNet18 model. Our ADMM-based algorithm produces factors whose values lie on a specified quantization grid.

We used AdaRound implementation from the AIMET⁴ framework and AdaQuant from the author’s repository⁵. AdaRound uses per-tensor quantization, while AdaQuant performs per-channel

4. <https://github.com/quic/aimet>

5. <https://github.com/itayhubara/CalibTIP>




Figure 3: Comparison between post-training quantization methods on ResNet18 model. wXaY denotes quantization of weights with bit-width X and activations with bit-width Y. rZ labels indicate that factorization was performed with reduction rate Z. ADMM-EPC defines our ADMM-based approach with CPD-EPC initialization. *We quantize all factorized layers to a specified bit-width, but leave batchnorm, downsample, first and last layers in 8 bits. **AdaQuant quantizes all layers to 4 bits, except for the first and last layer which are quantized to 8 bits. AdaQuant uses per-channel quantization.

quantization of weights. For AdaQuant, we used the advanced pipeline (Hubara et al., 2020) with bias-tuning, even though it involves quantization-aware training step. We report the best accuracy we could achieve, though it is not equal to accuracy depicted in AdaQuant paper. But considering 71.97% FP32 benchmark reported for ResNet18 model (in comparison with 69.76% used as starting point in this paper), the quality drop remains the almost the same. We have also noticed that AdaQuant does not consider quantizing inputs to residual connection summation.

With ADMM, CPD-EPC and CPD-EPC+AdaRound methods we do not factorize the first, last and downsample layers of the model. These layers contain small number of parameters but are placed in crucial parts of the architecture. Thus, factorizing them will not bring sufficient reduction in mac operations but might significantly degrade accuracy of predictions. Due to the fact that we do not factorize these layers, their weights still lie in full range of FP32 values in contrast with ADMM-factorized layers whose values lie on a uniform quantization grid of a specific bit-width. We leave these layers to 8-bit quantization for ADMM and CPD-EPC benchmarks. This distribution of bit-widths is concordant with other works on quantizing resnet-like architectures, e.g., (Hubara et al., 2020), even though factorization of layers is not involved.

Results are presented in Figure 3, more details can be found in the supplementary materials. It can be seen that our ADMM-based method outperforms other approaches on smaller BOP count values and shows competitive results on larger ranges.

Another benefit of a joint factorization-quantization approach is that it is more adaptive in achieving a desirable trade-off between accuracy and efficiency: varying reduction rates gives more flexibility than changing bit-widths configurations.

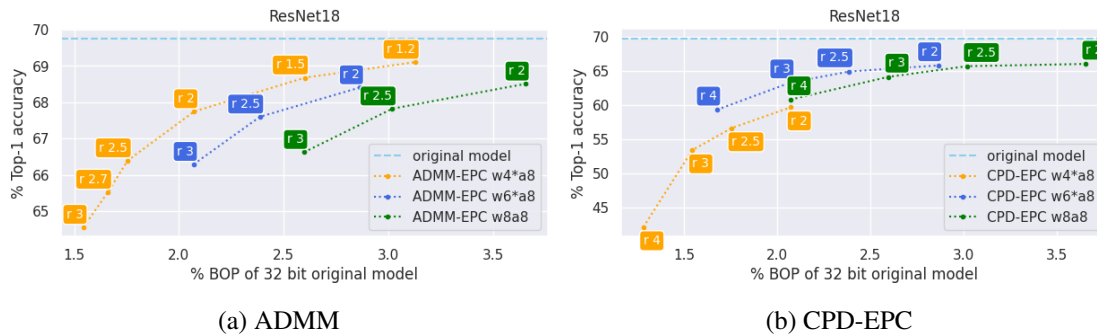


Figure 4: Ablation studies for best bit-width configuration. wXaY denotes quantization of weights with bit-width X and activations with bit-width Y. rZ labels indicate that factorization was performed with reduction rate Z. ADMM-EPC defines our ADMM-based approach with CPD-EPC initialization. *We quantize all factorized layers to a specified bit-width, but leave batchnorm, downsample, first and last layers in 8 bits.

4. Conclusion and Discussion

In this paper we proposed a technique for compressing neural networks that performs joint factorization and quantization. We introduced ADMM-based approximation algorithm to replace weights with float32 elements with their factorized representations, whose factors contain elements of lower precision (int8/int6/int4).

Our method allows to benefit from both factorization and quantization techniques, resulting in decrease of model size and acceleration in model inference. Our experiments have shown that ADMM-based joint quantization and factorization shows superior or on par results in comparison to other recent approaches to post-training quantization that do not involve model training and only require a small subset of training data. Moreover, our method allows for a flexible trade-off between accuracy and acceleration not restricted solely by bit-width configurations.

Our experiments indicate that the proposed ADMM-based optimization method is particularly advantageous when applied to tensor decompositions involving at least three factors, specifically in the context of Canonical Polyadic Decomposition (CPD) as utilized in this study. This method is likely to enhance the performance of neural networks with substantial convolutional layers, where tensor decompositions can be effectively employed, such as in CNNs designed for MRI image analysis or 3D shape recognition.

Through multiple ablation studies (Figures 1 and 2, Table 1) we determined that CPD-EPC initialization is the favorable choice of initialization for our optimization algorithm. It mitigates the

challenges associated with random initialization by reducing the need to conduct multiple optimization runs to achieve improved results.

There are several potential avenues for enhancing our approach. For example, incorporating other tensor decompositions and more sophisticated quantization techniques. Another possible direction would be to devise an algorithm of jointly choosing best factorization rank and bit-width similar Bayesian optimization in PARS (Sobolev, Ermilov, Phan, & Cichocki, 2022) or integer programming approach in (Hubara et al., 2020).

5. Acknowledgments

This work was partially supported by the Russian Science Foundation under grant number 25-41-00091. Julia Gusak contributed to the paper during her time at Skolkovo Institute of Science and Technology prior to joining Inria.

References

- Boyd, S., Parikh, N., Chu, E., Peleato, B., & Eckstein, J. (2010). Distributed optimization and statistical learning via the alternating direction method of multipliers..
- Cichocki, A., Lee, N., Oseledets, I., Phan, A. H., Zhao, Q., & Mandic, D. P. (2016). Tensor networks for dimensionality reduction and large-scale optimization part 1 low-rank tensor decompositions. *Foundations and Trends in Machine Learning*, 9.
- Cui, C., Zhang, K., Daulbaev, T., Gusak, J., Oseledets, I., & Zhang, Z. (2020). Active Subspace of Neural Networks: Structural Analysis and Universal Attacks. *SIAM Journal on Mathematics of Data Science*, 2(4), 1096–1122. blue <https://epubs.siam.org/doi/pdf/10.1137/19M1296070paper>.
- Diamond, S., Takapoui, R., & Boyd, S. (2016). A general system for heuristic solution of convex problems over nonconvex sets..
- Dong, P., Li, L., Tang, Z., Liu, X., Pan, X., Wang, Q., & Chu, X. (2024). Pruner-zero: Evolving symbolic pruning metric from scratch for large language models..
- Eldad, M., Alexander, F., Uri, A., & Mark, G. (2019). Same, same but different: Recovering neural network quantization error through weight factorization.. Vol. 2019-June.
- Gholami, A., Kim, S., Dong, Z., Yao, Z., Mahoney, M. W., & Keutzer, K. (2022). A survey of quantization methods for efficient neural network inference..
- Gou, J., Yu, B., Maybank, S. J., & Tao, D. (2021). Knowledge distillation: A survey. *Int. J. Comput. Vision*, 129(6), 1789–1819.
- Gusak, J., Daulbaev, T., Ponomarev, E., Cichocki, A., & Oseledets, I. (2021). Reduced-Order Modeling of Deep Neural Networks. *Computational Mathematics and Mathematical Physics Journal*. blue <https://arxiv.org/abs/1910.06995paper>.
- Gusak, J., Kholiavchenko, M., Ponomarev, E., Markeeva, L., Blagoveschensky, P., Cichocki, A., & Oseledets, I. (2019). Automated Multi-Stage Compression of Neural Networks. In *Proceedings of the IEEE International Conference on Computer Vision Workshops (ICCVW)*, pp. 0–0. blue

https://openaccess.thecvf.com/content_ICCVW2019/papers/LPCV/Gusak_Automated_Multi-Stage_Compression_of_Neural_Networks_ICCVW2019_paper.pdf

- Huang, K., Sidiropoulos, N. D., & Liavas, A. P. (2016). A flexible and efficient algorithmic framework for constrained matrix and tensor factorization.. Vol. 64.
- Hubara, I., Nahshan, Y., Hanani, Y., Banner, R., & Soudry, D. (2020). Improving post training neural quantization: Layer-wise calibration and integer programming..
- Kim, Y. D., Park, E., Yoo, S., Choi, T., Yang, L., & Shin, D. (2016). Compression of deep convolutional neural networks for fast and low power mobile applications..
- Kolda, T. G., & Bader, B. W. (2009). Tensor decompositions and applications..
- Kossaifi, J., Bulat, A., Tzimiropoulos, G., & Pantic, M. (2019). T-net: Parametrizing fully convolutional nets with a single high-order tensor. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7822–7831.
- Lebedev, V., Ganin, Y., Rakhuba, M., Oseledets, I., & Lempitsky, V. (2015). Speeding-up convolutional neural networks using fine-tuned cp-decomposition..
- Leng, C., Dou, Z., Li, H., Zhu, S., & Jin, R. (2018). Extremely low bit neural network: Squeeze the last bit out with admm..
- Liavas, A. P., & Sidiropoulos, N. D. (2015). Parallel algorithms for constrained tensor factorization via alternating direction method of multipliers. *IEEE Transactions on Signal Processing*, 63.
- Nagel, M., Amjad, R. A., Van Baalen, M., Louizos, C., & Blankevoort, T. (2020). Up or down? adaptive rounding for post-training quantization. In *International Conference on Machine Learning*, pp. 7197–7206. PMLR.
- Nagel, M., Fournarakis, M., Amjad, R. A., Bondarenko, Y., van Baalen, M., & Blankevoort, T. (2021). A white paper on neural network quantization. *ArXiv, abs/2106.08295*.
- Novikov, A., Podoprikin, D., Osokin, A., & Vetrov, D. (2015). Tensorizing neural networks.. Vol. 2015-January.
- Panagakis, Y., Kossaifi, J., Chrysos, G. G., Oldfield, J., Nicolaou, M. A., Anandkumar, A., & Zafeiriou, S. (2021). Tensor methods in computer vision and deep learning. *Proceedings of the IEEE*, 109(5), 863–890.
- Phan, A. H., Sobolev, K., Sozykin, K., Ermilov, D., Gusak, J., Tichavský, P., Glukhov, V., Oseledets, I., & Cichocki, A. (2020). Stable low-rank tensor decomposition for compression of convolutional neural network.. Vol. 12374 LNCS, pp. 522–539. Springer Science and Business Media Deutschland GmbH.
- Shang, F., Liu, Y., & Cheng, J. (2014). Generalized higher-order tensor decomposition via parallel admm.. Vol. 2.
- Smith, S., Beri, A., & Karypis, G. (2017). Constrained tensor factorization with accelerated AO-ADMM..
- Sobolev, K., Ermilov, D., Phan, A.-H., & Cichocki, A. (2022). Pars: Proxy-based automatic rank selection for neural network compression via low-rank weight approximation. *Mathematics*, 10.

- Tang, Z., Shi, S., Li, B., & Chu, X. (2023a). Gossipfl: A decentralized federated learning framework with sparsified and adaptive communication. *IEEE Transactions on Parallel and Distributed Systems*, 34(3), 909–922.
- Tang, Z., Shi, S., Wang, W., Li, B., & Chu, X. (2023b). Communication-efficient distributed deep learning: A comprehensive survey..
- Tichavský, P., Phan, A. H., & Cichocki, A. (2019). Sensitivity in tensor decomposition. *IEEE Signal Processing Letters*, 26.
- van Baalen, M., Louizos, C., Nagel, M., Amjad, R. A., Wang, Y., Blankevoort, T., & Welling, M. (2020). Bayesian bits: Unifying quantization and pruning.. Vol. 2020-December.
- White, C., Safari, M., Sukthanker, R., Ru, B., Elsken, T., Zela, A., Dey, D., & Hutter, F. (2023). Neural architecture search: Insights from 1000 papers..
- Yin, M., Sui, Y., Liao, S., & Yuan, B. (2021). Towards efficient tensor decomposition-based dnn model compression with optimization framework..