

# AN INTER-LAYER WEIGHT PREDICTION AND QUANTIZATION FOR DEEP NEURAL NETWORKS BASED ON SMOOTHLY VARYING WEIGHT HYPOTHESIS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

Due to a resource-constrained environment, network compression has become an important part of deep neural networks research. In this paper, we propose a new compression method, *Inter-Layer Weight Prediction* (ILWP) and quantization method which quantize the predicted residuals between the weights in all convolution layers based on an inter-frame prediction method in conventional video coding schemes. Furthermore, we found a phenomenon *Smoothly Varying Weight Hypothesis* (SVWH) which is that the weights in adjacent convolution layers share strong similarity in shapes and values, i.e., the weights tend to vary smoothly along with the layers. Based on SVWH, we propose a second ILWP and quantization method which quantize the predicted residuals between the weights in adjacent convolution layers. Since the predicted weight residuals tend to follow Laplace distributions with very low variance, the weight quantization can more effectively be applied, thus producing more zero weights and enhancing the weight compression ratio. In addition, we propose a new *inter-layer loss* for eliminating non-texture bits, which enabled us to more effectively store only texture bits. That is, the proposed loss regularizes the weights such that the collocated weights between the adjacent two layers have the same values. Finally, we propose an ILWP with an inter-layer loss and quantization method. Our comprehensive experiments show that the proposed method achieves a much higher weight compression rate at the same accuracy level compared with the previous quantization-based compression methods in deep neural networks.

## 1 INTRODUCTION

Deep neural networks have demonstrated great performance for various tasks in many fields, such as image classification (LeCun et al. 1990a; Krizhevsky et al. 2012; He et al. 2016), object detection (Ren et al. 2015; He et al. 2017; Redmon & Farhadi 2018), image captioning (Jia et al., 2015), and speech recognition (Hinton et al. 2012; Xiong et al. 2018). Wide and deep neural networks achieved great accuracy with the aid of the enormous number of weight parameters and high computational cost (Simonyan & Zisserman 2014; He et al. 2016; Huang et al. 2017).

However, as demands toward constructing the neural networks in the resource-constrained environments have been increasing, making the resource-efficient neural network while maintaining its accuracy becomes an important research area of deep neural networks. Several studies have aimed to solve this problem.

In LeCun et al. (1990b), Hassibi & Stork (1993), Han et al. (2015b) and Li et al. (2016), network pruning methods were proposed for memory-efficient architecture, where unimportant weights were forced to have zero values in terms of accuracy. In Fiesler et al. (1990), Gong et al. (2014) and Han et al. (2015a), weights were quantized and stored in memory, enabling less memory usage of deep neural networks. On the other hand, some literature decomposed convolution operations into sub operations (e.g., depth-wise separable convolution) requiring less computation costs at similar accuracy levels (Howard et al. 2017; Zhang et al. 2018; Sandler et al. 2018; Ma et al. 2018).

In this paper, we show that the weights between the adjacent two convolution layers tend to share high similarity in shapes and values. We call this phenomenon *Smoothly Varying Weight Hypothesis*

(SVWH). This paper explores an efficient neural network method that fully takes the advantage of SVWH.

Specifically, inspired by the prediction techniques widely used in video compression field (Wiegand et al. 2003; Sullivan et al. 2012), we propose a new weight compression scheme based on an inter-layer weight prediction technique, which can be successfully incorporated into the depth-wise separable convolutional blocks (Howard et al. 2017; Zhang et al. 2018; Sandler et al. 2018; Ma et al. 2018).

**Contributions:** Main contributions of this paper are listed below:

- From comprehensive experiments, we find out that the weights between the adjacent layers tend to share strong similarities, which lead us to establishing SVWH.
- Based on SVWH, we propose a simple and effective *Inter-Layer Weight Prediction* (ILWP) and quantization framework enabling a more compressed neural networks than only applying quantization on the weights of the neural networks.
- To further enhance the effectiveness of the proposed ILWP, we devise a new regularization function, denoted as *inter-layer loss*, that aims to minimize the difference between collocated weight values in the adjacent layers, resulting in significant bit saving for non-texture bits (i.e., bits for indices of prediction).
- Our comprehensive experiments demonstrate that, the proposed scheme achieves about 53% compression ratio on average in 8-bit quantization at the same accuracy level compared to the traditional quantization method (without prediction) in both MobileNetV1 (Howard et al., 2017) and MobileNetV2 (Sandler et al., 2018).

## 2 RELATED WORK

**Network pruning:** Network pruning methods prune the unimportant weight parameters, enabling to reduce the redundancy of weight parameters inherent in neural networks. LeCun et al. (1990b) and Hassibi & Stork (1993) reduced the number of weight connections implicitly through setting an proper objective function for training. Han et al. (2015b) successfully removed the unimportant weight connections through certain thresholds for the weight values, showing no harm of accuracy in the state-of-the-art convolutional neural network models. Recently, structured (filter/channel/layer-wise) pruning methods have been proposed in Liu et al. (2017) and Li et al. (2016), where a set of weights is pruned based on certain criteria (e.g., the sum of absolute values in the set of weights), demonstrating significantly reduced number of weight parameters and computational costs. Furthermore, He et al. (2018) use AutoML for channel pruning and their proposed method get accuracy 13.2% more than filter pruning method (Li et al., 2016). Our paper is linked to the pruning methods in perspective of assigning more zero weights for weight compression.

**Quantization:** Quantization reduces the representation bits of original weights in neural networks. Fiesler et al. (1990) proposed a weight quantization using weight discretization in neural networks. Han et al. (2015b) incorporated a vector quantization into the pruning, proving that quantization and pruning can jointly work for weight compression without accuracy degradation. This pruning-quantization framework, i.e., called Deep Compression, became a milestone in model compression research of deep neural networks. Lin et al. (2016) proposed a fixed-point quantization using a linear scale factor for weight values where bit-widths for quantization are adaptively found for each layer, thus enabling 20% reduction of the weight size in memory without any loss in accuracy compared to the baseline fixed-point quantization method. Furthermore, Sung et al. (2015), Zhuang et al. (2018) and Zhao et al. (2019) use clipping weights before applying linear quantization and those methods are improve accuracy than linear quantization without clipping.

For 1-bit quantization which is named as Binary Neural Networks, several studies (Courbariaux et al. 2015; Courbariaux et al. 2016; Hubara et al. 2017) binarized the weights and/or activations in the process of back-propagation, enjoying considerably reduced usage of memory space and computation overhead.

Compared to the aforementioned quantization techniques, our work applies quantization in the combination of the residuals for the weights in inter layers rather than the weight values themselves. Empirically, we found that the residuals tend to produce much larger portion of zero values when quan-

tized, since they often follow very narrow Laplace distributions. Therefore, our proposed method can significantly reduce the memory consumption for neural networks as shown in Section 4.

**Prediction in conventional video coding:** Prediction technique is considered one of the most crucial parts in video compression, aiming at minimizing the magnitudes of signals to be encoded by subtracting the input signals to the most similar encoded signals in a set of prediction candidates (Wiegand et al. 2003; Rijkse 1996; Sullivan et al. 2012). The prediction methods can produce the residuals of signals with low magnitudes as well as a large number of non/near zero signals. Therefore, they have effectively been incorporated into transforms and quantization for concentrating powers in low frequency regions and reducing the entropy, respectively. There are two prediction techniques: inter- and intra- predictions. The inter-prediction searches the best prediction signals from the encoded neighbor frames out of the current frame, while the intra-prediction generates a set of prediction signals from the input signals and determine the best prediction (Wiegand et al. 2003; Sullivan et al. 2012). This is because intra frames that use only intra-prediction for compression are used as reference frames for subsequent frames to be predicted.

Note that a few studies explored to apply the transform techniques of video and/or image coding to the weight compression problem in neural networks. Wang et al. (2016) and Ko et al. (2017) applied DCT (Discrete Cosine Transform) used in the JPEG (Joint Picture Encoding Group) algorithm to the model compression problem of deep neural networks such that the energy of weight values became concentrated in low frequency regions, thus producing more non/near-zero DCT coefficients for the weights. Compared to the aforementioned papers, our work does not adopt transform techniques to reduce model sizes, because the transforms introduce much computation in inference, decreasing the effectiveness of the weight compression in practical applications.

In this paper, we found out that the inter-prediction technique can play a crucial role for weight compression under the SVWH condition (i.e., the weights between adjacent layers tend to be similar). The proposed *inter-layer loss* reinforces SVWH for reducing the non-texture bits in the training process. As a result, the proposed inter-prediction and quantization framework for weight compression yields impressive compression performance enhancement at the similar accuracy level compared to the baseline models.

### 3 METHOD

#### 3.1 BASIC INTER-LAYER WEIGHT PREDICTION

Figure 1-(a) shows an simple example of the proposed ILWP method with a full search strategy (ILWP-FSS). The full search strategy (FSS) indicates that, for the  $j$ -th weight in the kernel of the  $i$ -th convolution layer ( $K_{i,j}$ ), it searches the most similar weight kernel (i.e.,  $K_{u,v}$  in Figure 1-(a)) in the range of  $[1, i - 1]$ -th layers given a pre-trained neural network model. We then compute the residuals ( $R_{i,j}$ ) between the current weight and the best prediction, which is finally quantized in a certain bit representation and stored in memory space with the index of the best prediction (i.e.,  $(u, v)$  in Figure 1-(a)).

The proposed ILWP is performed from the second layer to the last layer in a sequential manner. It should be noted that, because of the large portion of non-texture bits (indices of the best prediction), ILWP-FSS tends to produce more bits (both texture (residual) and non-texture bits) than the standard weight values without ILWP, which makes ILWP worthless in compressing the weights. In the next subsection, we describe how SVWH solves this problem effectively.

#### 3.2 SMOOTHLY VARYING WEIGHT HYPOTHESIS

As shown in Figure 2, the dominant portions of the best predictions in the current ( $i$ -th layer) layer are obtained from its previous ( $(i - 1)$ -th layer) layers, being consistently observed in all the layers of the neural networks. From these observations, we claim that the adjacent layers tend to consists of the similar weight values, leading us to propose SVWH. The proposed SVWH can be mathematically expressed as

$$P[L(K_{i,j}, K_{i-1,v}) > L(K_{i,j}, K_{u,v})] > P[L(K_{i,j}, K_{i-1,v}) \leq L(K_{i,j}, K_{u,v})], \quad (1)$$

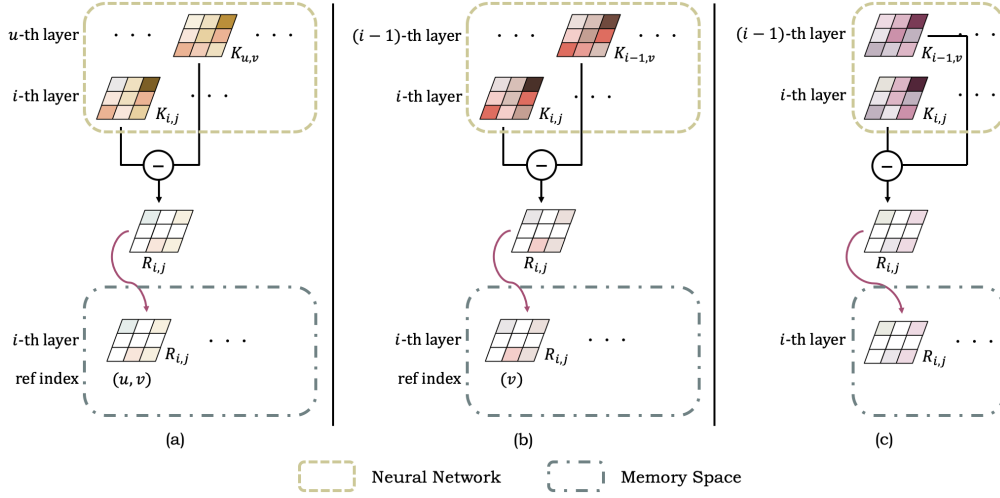


Figure 1: Our *Inter-Layer Weight Prediction* (ILWP) scheme. (a) ILWP with the full search strategy described in Section 3.1; (b) ILWP with the local search strategy described in Section 3.2; (c) ILWP trained with inter-layer loss and without searching described in Section 3.3. The ref index is the referenced kernel index for ILWP, i.e. the index of the most similar weight kernel on target layer. Other notation is described in each subsection.




Figure 2: Heatmaps that show the percentage in the number of best predictions with respect to minimizing  $L_1$  distance between the source layer ( $x$ -axis) and the target layer ( $y$ -axis) in the neural networks trained on CIFAR10 (top) and CIFAR100 (bottom). From left to right are MobileNet, MobileNetV2, ShuffleNet, and ShuffleNetV2

where  $P[\cdot]$  is a probability,  $L(\cdot, \cdot)$  is a distance between two kernels, and  $(u, v)$  are the indices of the layer and kernel, where  $1 \leq u < i - 1$  (See Figure 1).

Inspired by SVWH, we propose an enhanced version of ILWP, i.e. ILWP with a local search strategy (ILWP-LSS) that finds the best prediction only from the the previous layer (See Figure 1-(b)). The local search strategy (LSS) can effectively reduce the non-texture bits, because the non-texture bits for the layer index are not required. Moreover, the LSS allows the network to keep the weights only in the previous one layer in inference, thus reducing the memory buffer for keeping the weights. Our experimental results in Section 4 show that, the local search strategy enhances the compression performance compared to the FSS in the ILWP framework of deep neural networks. In this paper, to further increase the effectiveness of the proposed ILWP, we devise a new regularization loss, namely inter-layer loss.

### 3.3 ILWP WITH INTER-LAYER LOSS

To further exploit SVWH, we propose a new inter-layer loss which makes the collocated filter weights between the adjacent layers have almost the same values. We find out that our ILWP can more effectively be applied to the depth-wise convolution ( $3 \times 3$  spatial convolution) layer in the depth-wise separable convolution block (Howard et al., 2017), compared to the traditional 3D convolutions ( $3 \times 3 \times C$  convolution). This is because high dimensionality of the weights in the traditional 3D convolution filter tends to hinder finding out the best prediction sharing strong similarity. This can introduce a longer tail and wider shape of the Laplace distribution for the residuals of the weights, thus decreasing the compression efficiency. Moreover, it is not possible to predict the weight kernels having different channel dimensions from the current weight kernel. This can limit the usability of the proposed ILWP.

On the other hand, the depth-wise convolution consists only of nine elements, and all the depth-wise convolutions have canonical kernel shapes with  $3 \times 3$  size in whole networks. Also, since the point-wise convolution ( $1 \times 1$  convolution) learns channel correlations for the output features of the depth-wise convolution layer, forcing the collocated weights to be similar does hardly affect the accuracy of neural networks. These characteristics of depth-wise convolution enhance the usability of the proposed ILWP in the weight compression of neural networks.

Moreover, it becomes more popular to use depth-wise separable convolutions in the recent neural network architectures, due to its high efficiency (Kaiser et al. 2017; (Sandler et al., 2018); Zhang et al. 2018; Ma et al. 2018). Therefore, we apply the proposed method into the spatial convolution in the depth-wise separable convolution block. Our proposed *inter-layer loss* can dramatically eliminate the non-texture bits and is defined as follows:

$$\mathcal{L}_{\text{inter-layer}} = \frac{1}{Z} \sum_{i=2}^N \sum_{j=1}^{c_{i-1}} |K_{i,j} - K_{i-1,v}|, \quad (2)$$

where  $Z$  is the number of weights to be predicted, and  $N$  is the number of depth-wise convolution layers in a whole neural network. In Eq. (2),  $v$  is the index in the previous layer which is equal to  $j \bmod c_{i-1}$  in the inter layer loss. The proposed loss function in Eq. (2) not only regularizes the weights but also allows us to eliminate all the non-texture bits for indices of the best predictions, since the network can always predict the weight values of the current layer from the collocated weights in its previous layer. For training the neural networks with our proposed loss, our total loss is formulated as follows:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{cls}} + \lambda \mathcal{L}_{\text{inter-layer}}, \quad (3)$$

where  $\mathcal{L}_{\text{cls}}$  is the conventional classification loss using the cross-entropy loss (He et al., 2016).  $\lambda$  in Eq. (3) is the control parameter for the *inter-layer loss* and is set to 1 over all experiments. From our comprehensive experiments, we found that setting  $\lambda$  in Eq. (3) to 1 is suitable to match the trade-off between the performance of neural networks and the parameter size of neural networks (See Figure 4 in Section 4.1).

Through our new *inter-layer loss*, SVWH is explicitly controlled in a more elaborate manner as shown in Figure 3 that shows the Heatmaps for the percentage in the number of best predictions with respect to minimizing  $L_1$  distance between the source layer ( $x$ -axis) and the target layer ( $y$ -axis) in MobileNet, MobileNetV2, ShuffleNet, and ShuffleNetV2 trained with the proposed inter-layer loss. Finally, the reconstruction of the current weight kernel is performed as follow:

$$\hat{K}_{i,j} = \tilde{R}_{i,j} + \hat{K}_{i-1,v}, \quad (4)$$

where  $\tilde{R}_{i,j}$  is the quantized residual at the  $i$ -th layer and  $j$ -th filter position.  $\hat{K}$  in Eq. (3) is the final reconstruction of  $K$ . Note that, to reconstruct the current weights, the weights only in the previous layer (i.e.,  $\hat{K}_{i-1,v}$  in Eq. (4)) are required (See Figure 1-(c)). The residuals of the weights are quantized through a linear quantization technique, and then saved using Huffman coding for the purpose of efficient model capacity (Han et al., 2015b). Thanks to the prediction scheme, our method




Figure 3: Heatmaps that show the percentage in the number of best predictions with respect to minimizing  $L_1$  distance between the source layer ( $x$ -axis) and the target layer ( $y$ -axis) in the neural networks trained on CIFAR10 (top) and CIFAR100 (bottom) with *inter-layer loss*. From left to right are MobileNet, MobileNetV2, ShuffleNet, and ShuffleNetV2

usually remains more high-valued non-zero weights than the traditional quantization methods. Since high weight values importantly contribute to the prediction performance (Han et al., 2015b), our method can retain both high accuracy and weight compression performance.

The weight kernels in the first layer are not quantized since it is as important as Intra-frames (i.e. reference frames) in a group of pictures of the video coding scheme (Wiegand et al., 2003). If the weight kernels in the first layer are quantized, the weight kernels in subsequent layers which are predicted based on the weight kernels in the first layer are negatively affected, leading to accuracy drop in neural networks.

## 4 EXPERIMENTS

In this section, we describe and prove the superiority of the proposed method by applying it on image classification tasks, specifically for CIFAR-10 and CIFAR-100 datasets. For securing generality of our proposed IWLP, we applied our method in four convolutional neural networks, MobileNet, MobileNetV2, ShuffleNet, and ShuffleNetV2. Four NVidia 2080-Ti GPU with the Intel i9-7900X CPU are used to perform the experiments. For the hyper-parameter setting in the training process, we set the initial learning rate as 0.1, which is multiplied by 0.98 every epochs. We used Stochastic Gradient Descent (SGD) optimizer with Nesterov momentum (Sutskever et al., 2013) factor 0.9. All the neural networks are trained for 200 epochs with a batch size of 256. The baseline model is each aforementioned four convolutional neural networks which is quantized by linear quantization on the weights of depthwise convolutional kernel. In all the experiments, the test accuracy and parameters in kilobyte (KB) are marked from the average of 5 runs.

### 4.1 OPTIMAL $\lambda$ SELECTION

First, we found the optimal  $\lambda$  in Eq. (3) that is suitable to match the trade-off between the accuracy performance and the parameter size of neural networks. Figure 4 shows the parameter sizes and top-1 test accuracy in MobileNet trained on the CIFAR-10 dataset for different  $\lambda$  in Eq. (3) after quantization and Huffman coded. We can see that the case of  $\lambda = 1$  has the smallest parameter size in neural networks, and slightly lower top-1 test accuracy in the CIFAR-10 dataset. Furthermore, not only MobileNet trained on CIFAR-100, but also other models, i.e., MobileNetV2, ShuffleNet, and ShuffleNetV2 trained on CIFAR-10 and CIFAR-100 datasets also has very similar results. So, we experimented by setting the  $\lambda$  in Eq. (3) to 1.




Figure 4: Comparison of parameter size in kilobytes (KB) and top-1 test accuracy in MobileNet trained on CIFAR-10 dataset for different  $\lambda$  in Eq. (3) after  $n$ -bit linear quantization. 32-bit means using full precision floating point.

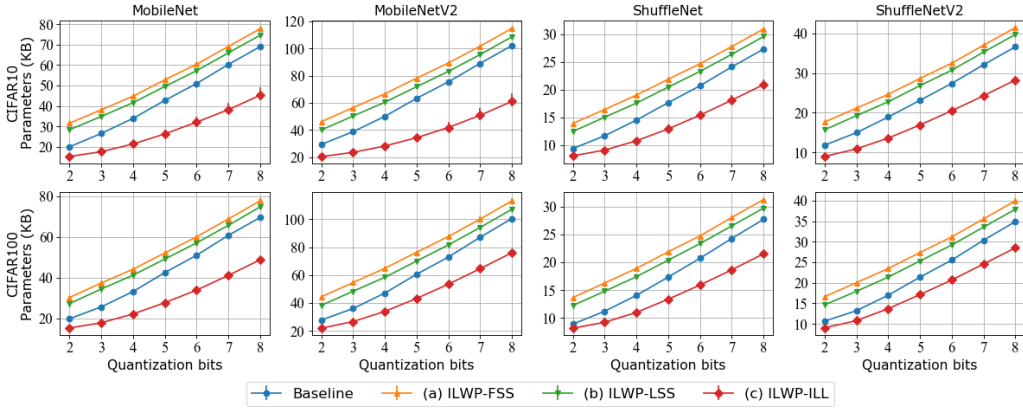


Figure 5: Comparison of Parameter size in kilobytes (KB) for different quantization bits in  $\{2, 3, 4, \dots, 8\}$  on CIFAR-10 (top) and CIFAR-100 (bottom) datasets. From left to right are MobileNet, MobileNetV2, ShuffleNet, and ShuffleNetV2. (a) ILWP with the full search strategy (ILWP-FSS); (b) ILWP with the local search strategy (ILWP-LSS); (c) ILWP trained with inter-layer loss and without searching (ILWP-ILL) described in Section 3.3. In all the figures, each dot represents the average test accuracy over 5 training times in each quantization bit in  $\{2, 3, 4, \dots, 8\}$ .

## 4.2 EXPERIMENTAL RESULTS

Figure 5 shows the parameter size of our proposed methods for different quantization bits in  $\{2, 3, 4, \dots, 8\}$  on CIFAR-10 and CIFAR-100 datasets. Figure 5 shows that our proposed ILWP-ILL yields higher compression ratio in the weight parameters. On the other hand, ILWP-FSS and ILWP-LSS have larger parameter sizes than the baseline. This is because both ILWP-FSS and ILWP-LSS contain non-texture bits (bits for indices of the best predictions), resulting in bit overheads.

For further analysis on the texture and non-texture bits in the proposed ILWP, Table 1 shows the performance comparison of our proposed methods in terms of the parameter sizes for the depth-wise convolution kernels and top-1 test accuracy. Due to the inter-layer prediction scheme, all the proposed methods show less amount of texture bits compared to the baseline. However, the total amounts of bits of ILWP-FSS and ILWP-LSS are higher than the baseline model, which is due to the presence of the non-texture bits ( $u$  and  $v$  in Figure 1). However, ILWP-ILL demonstrated much reduced amount of total bits compared to the baseline, which is due to the property that this method does not require saving any indices of reference kernels while maintaining good accuracy under the SVWH condition.

Figure 6 shows the results of our proposed methods on CIFAR-10 and CIFAR-100 datasets for the parameter size in kilobytes (KB) and test accuracy. As shown in Figure 6, the ILWP-ILL outperforms the other compared ones in terms of the trade-off between the amounts of compressed bits and accuracy. This is because ILWP-ILL significantly saves the weight bits by eliminating the non-

Table 1: Parameter size (texture bits, non-texture bits, total bits) in kilobytes (KB) and top-1 test accuracy of depth-wise convolution kernels after 8 bit quantization and Huffman coding, in MobileNet trained on CIFAR-10. We compared the performance of the three variations of ILWP: ILWP-FSS, ILWP-LSS, and ILWP-ILL.

Method	texture bits	non-texture bits	total bits	Test accuracy
Baseline	70.192	-	70.192	92.125
ILWP-FSS	67.716	11.177	78.893	92.225
ILWP-LSS	67.415	8.540	75.955	<b>92.235</b>
ILWP-ILL	<b>33.373</b>	-	<b>33.373</b>	91.48

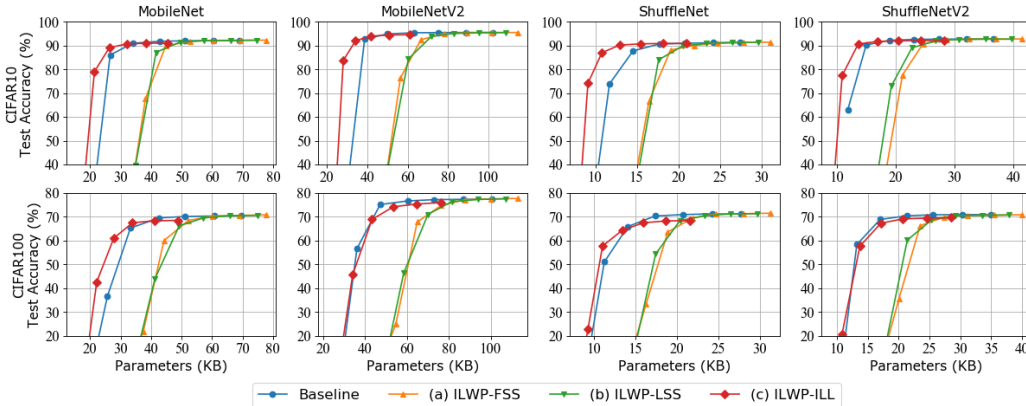


Figure 6: Performance results of our proposed methods on CIFAR-10 (top) and CIFAR-100 (bottom) datasets for the parameter size in kilobytes (KB) and test accuracy. From left to right are MobileNet, MobileNetV2, ShuffleNet, and ShuffleNetV2.  $x$ -axis: the total size of the quantized and Huffman coded depth-wise convolutional kernels in storage;  $y$ -axis: Top-1 test accuracy. (a) ILWP-FSS; (b) ILWP-LSS; (c) ILWP-ILL. In all the figures, each dot represents the average test accuracy over 5 training times in each quantization bit in  $\{2, 3, 4, \dots, 8\}$ .

texture bits as well as increasing the effectiveness of the quantization process as the residuals tend to follow much narrower Laplace distributions than the original weight values (See Section 5).

Compared to ILWP-FSS, ILWP-LSS shows very slight improvement in trade-off between the size of parameters and accuracy as ILWP-LSS does not store the layer index of the best prediction ( $u$  in Figure 1). This is due to the fact that the portion of the bits for the layer indices is much smaller than the portion of the bits for the kernel indices. Compared to the baseline models, it is worth noting that both ILWP-FSS and ILWP-LSS have often worse compression efficiency compared to the baseline. This is because they introduce non-texture bits consisting of a large portion in the total bits for weight parameters. Therefore, it can be concluded that our foundation of SVWH and the proposed inter-layer loss allow the network to make use of full advantages in the inter-layer prediction scheme of the conventional video coding frameworks.

## 5 ANALYSIS

Figure 7 compares the distributions for the weights and residuals in all the depth-wise convolution kernels in MobileNet trained on CIFAR-100. Figure 7-(a), -(b), and -(c) visualize the distribution of the weight kernels in baseline model, the distribution of residuals in ILWP-FSS and the distribution of the residuals in ILWP-ILL, respectively. As shown in Figure 7, ILWP-FSS produces a single modal Laplace distribution. This nice property contributes to high compression capacity for storing the weights in neural networks when these residuals are quantized and Huffman coded. However, this method requires saving the additional indices, leading to consuming more amount of bits.






Figure 7: Comparison of the distributions for the weights and residuals in all the depth-wise convolution kernels in MobileNet trained on CIFAR-100. (a) Distribution of the weight kernels in baseline model; (b) Distribution of residuals in ILWP-FSS; (c) Distribution of the residuals in ILWP-ILL.

As shown in Figure 7-(c), it can be observed that for ILWP-ILL, the residuals are located in near-zero positions (about 46% of the weights), following a very sharp Laplace distribution. This indicates that ILWP-ILL with quantization allows neural networks to achieve remarkable weight compression performance by generating a large amount of zero coefficients after quantization. Furthermore, in terms of information theory, the Shannon entropy  $H_x$  of Laplace distribution is derived as follows<sup>1</sup>:

$$\begin{aligned}
 H_x &= - \int_{-\infty}^{\infty} f_L(x) \log(f_L(x)) \\
 &= - \int_{-\infty}^{\infty} \left( \frac{1}{2b} e^{-\frac{|x-\mu|}{b}} \right) \log\left( \frac{1}{2b} e^{-\frac{|x-\mu|}{b}} \right) \\
 &= \log(2b) + 1,
 \end{aligned} \tag{5}$$

where,  $f_L(\cdot)$  is the probability density function of the Laplace distribution and  $b$  and  $\mu$  is scale and location factors of the Laplace distribution, respectively. So, The Shannon entropy is proportional to scale parameter  $b$ , which controls width of Laplace distribution, i.e., small  $b$  has a narrow Laplace distribution.

As shown in Figure 7, it is observed that distribution of residuals in ILWP-ILL has much narrower Laplace distribution than distribution of residuals in ILWP-FSS. Consequently, the information entropy of the distribution of the residuals in ILWP-ILL is lower than the information entropy of the distribution of the residuals in ILWP-FSS. Meanwhile, the entropy coding as Huffman coding is more compressed in small information entropy. As a result, the ILWP-ILL method is more compressed than the ILWP-FSS method after quantization and entropy coding.

## 6 CONCLUSION

We propose a new inter-layer weight prediction with inter-layer loss for efficient deep neural networks. Motivated by our observation that the weights in the adjacent layers tend to vary smoothly, we successfully build a new weight compression framework combining the inter-layer weight prediction scheme, the inter-layer loss, quantization and Huffman coding under SVWH condition. Intuitively, our prediction scheme significantly decreases the entropy of the weights by making them much narrower Laplace distributions, thus leading remarkable compression ratio of the weight parameters in neural networks. Also, the proposed inter-layer loss effectively eliminates the non-texture bits for the best predictions. To the best of our knowledge, this work is the first to report the phenomenon of the weight similarities between the neighbor layers and to build a prediction-based weight compression scheme in modern deep neural network architectures.

<sup>1</sup>See Section 2.1.3.5 (p.21) in Kotz et al. (2012)

## REFERENCES

- Matthieu Courbariaux, Yoshua Bengio, and Jean-Pierre David. Binaryconnect: Training deep neural networks with binary weights during propagations. In *Advances in neural information processing systems*, pp. 3123–3131, 2015.
- Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Binarized neural networks: Training deep neural networks with weights and activations constrained to+1 or-1. *arXiv preprint arXiv:1602.02830*, 2016.
- Emile Fiesler, Amar Choudry, and H John Caulfield. Weight discretization paradigm for optical neural networks. In *Optical interconnections and networks*, volume 1281, pp. 164–174. International Society for Optics and Photonics, 1990.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. Compressing deep convolutional networks using vector quantization. *arXiv preprint arXiv:1412.6115*, 2014.
- Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015a.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015b.
- Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in neural information processing systems*, pp. 164–171, 1993.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pp. 2961–2969, 2017.
- Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 784–800, 2018.
- Geoffrey Hinton, Li Deng, Dong Yu, George Dahl, Abdel-rahman Mohamed, Navdeep Jaitly, Andrew Senior, Vincent Vanhoucke, Patrick Nguyen, Brian Kingsbury, et al. Deep neural networks for acoustic modeling in speech recognition. *IEEE Signal processing magazine*, 29, 2012.
- Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017.
- Xu Jia, Efstratios Gavves, Basura Fernando, and Tinne Tuytelaars. Guiding the long-short term memory model for image caption generation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2407–2415, 2015.
- Lukasz Kaiser, Aidan N Gomez, and Francois Chollet. Depthwise separable convolutions for neural machine translation. *arXiv preprint arXiv:1706.03059*, 2017.
- Jong Hwan Ko, Duckhwan Kim, Taesik Na, Jaeha Kung, and Saibal Mukhopadhyay. Adaptive weight compression for memory-efficient neural networks. In *Proceedings of the Conference on Design, Automation & Test in Europe*, pp. 199–204. European Design and Automation Association, 2017.

- Samuel Kotz, Tomasz Kozubowski, and Krzysztof Podgorski. *The Laplace distribution and generalizations: a revisit with applications to communications, economics, engineering, and finance*. Springer Science & Business Media, 2012.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.
- Yann LeCun, Bernhard E Boser, John S Denker, Donnie Henderson, Richard E Howard, Wayne E Hubbard, and Lawrence D Jackel. Handwritten digit recognition with a back-propagation network. In *Advances in neural information processing systems*, pp. 396–404, 1990a.
- Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pp. 598–605, 1990b.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.
- Darryl Lin, Sachin Talathi, and Sreekanth Annapureddy. Fixed point quantization of deep convolutional networks. In *International Conference on Machine Learning*, pp. 2849–2858, 2016.
- Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2736–2744, 2017.
- Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 116–131, 2018.
- Joseph Redmon and Ali Farhadi. Yolov3: An incremental improvement. *arXiv preprint arXiv:1804.02767*, 2018.
- Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*, pp. 91–99, 2015.
- Karel Rijkse. H. 263: video coding for low-bit-rate communication. *IEEE Communications magazine*, 34(12):42–45, 1996.
- Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 4510–4520, 2018.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12):1649–1668, 2012.
- Wonyong Sung, Sungho Shin, and Kyuyeon Hwang. Resiliency of deep neural networks under quantization. *arXiv preprint arXiv:1511.06488*, 2015.
- Ilya Sutskever, James Martens, George Dahl, and Geoffrey Hinton. On the importance of initialization and momentum in deep learning. In *International conference on machine learning*, pp. 1139–1147, 2013.
- Yunhe Wang, Chang Xu, Shan You, Dacheng Tao, and Chao Xu. Cnnpack: Packing convolutional neural networks in the frequency domain. In *Advances in neural information processing systems*, pp. 253–261, 2016.
- Thomas Wiegand, Gary J Sullivan, Gisle Bjontegaard, and Ajay Luthra. Overview of the h. 264/avc video coding standard. *IEEE Transactions on circuits and systems for video technology*, 13(7): 560–576, 2003.

- Wayne Xiong, Lingfeng Wu, Fil Alleva, Jasha Droppo, Xuedong Huang, and Andreas Stolcke. The microsoft 2017 conversational speech recognition system. In *2018 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 5934–5938. IEEE, 2018.
- Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 6848–6856, 2018.
- Ritchie Zhao, Yuwei Hu, Jordan Dotzel, Christopher De Sa, and Zhiru Zhang. Improving neural network quantization using outlier channel splitting. *arXiv preprint arXiv:1901.09504*, 2019.
- Bohan Zhuang, Chunhua Shen, Minghui Tan, Lingqiao Liu, and Ian Reid. Towards effective low-bitwidth convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7920–7928, 2018.