ARCHITECTURE-AWARE NETWORK PRUNING FOR VISION QUALITY APPLICATIONS

Wei-Ting Wang, Han-Lin Li, Wei-Shiang Lin, Cheng-Ming Chiang, Yi-Min Tsai

MediaTek Inc.

ABSTRACT

Convolutional neural network (CNN) delivers impressive achievements in computer vision and machine learning field. However, CNN incurs high computational complexity, especially for vision quality applications because of large image resolution. In this paper, we propose an iterative architecture-aware pruning algorithm with adaptive magnitude threshold while cooperating with quality-metric measurement simultaneously. We show the performance improvement applied on vision quality applications and provide comprehensive analysis with flexible pruning configuration. With the proposed method, the Multiply-Accumulate (MAC) of state-of-the-art low-light imaging (SID) and super-resolution (EDSR) are reduced by 58% and 37% without quality drop, respectively. The memory bandwidth (BW) requirements of convolutional layer can be also reduced by 20% to 40%.

Index Terms— Pruning, Vision Quality, Network Architecture

1. INTRODUCTION

CNN is adopted as an essential ingredients in computer vision and machine learning areas [1, 2, 3]. Vision perception tasks including image classification, object detection and semantic segmentation are comprehensively investigated and associated with CNN. Even in image processing field such as super resolution, high dynamic range imaging and de-noising, CNN has progressive and promising improvement on image quality in recent years [4, 5].

However, compared to perception tasks, it requires higher computational complexity and BW requirements for vision quality tasks. MobileNetV1 [6] is designed with 569M MAC for ImageNet classification. On the other hand, in low-light photography SID [4] and super-resolution EDSR [5], it takes 560G MAC and 1.4T MAC per inference, respectively. It is

more challenge to deploy CNN models on mobile devices for vision quality applications.

Network pruning [7] is an effective methodology toward performance optimization. Sparsity is defined as the ratio of the number of zero weights divided by the number of total weights. Better pruning algorithm delivers higher sparsity and reduces more MAC and BW correspondingly. However, quality drop is one of the major challenges in network pruning. Fig. 1 shows visible defect on SID even with only 0.1 PSNR degradation.

In this paper, we propose architecture-aware pruning to maximize sparsity and MAC reduction without quality-metric (PSNR or SSIM) drops. We also analyze the effects of MAC and BW reduction with different configurations associating with pruned structures. The proposed method focus on algorithms including but not limited to SID and EDSR.

2. RELATED WORKS

Network pruning has been widely explored in existing literatures. To answer which weight should be pruned, some works add evaluation functions to loss function, such as group lasso [8] and MAC regularization [9]. However, it is difficult to find a proper ratio between additional pruning-related loss and original loss. Others works create evaluation functions, including sensitivity [10, 11] and weight magnitude [12]. The sensitivity method computes the impact of weights on the training loss and removes low-impact weights. Weight magnitude method simply prune weight if its absolute value is less than the threshold, which is easier to be applied on large-scale CNNs. In this work, we use weight magnitude method to prune the network.

Pruning granularity. There are two granularity of pruning, fine-grained pruning [12, 13] and coarse-grained pruning [14, 15, 16, 17]. Fined-grained method prunes individual weights (i.e., within a filter kernel), whereas coarse-grained method extensively considers network structures (i.e., along the output and the input channels). According to [18], fine-grained pruning needs additional dedicated hardware to handle irregular sparsity. Coarse-grained method may obtain higher compression ratio without the need of compression header [19]. Therefore, we focus on coarse-grained output-channel-wise pruning.

Fig. 1. Slightly quality-metric drop (PSNR -0.09) may incur visible defects (SID). (a) PSNR: 25.41. (b) PSNR: 25.32.

Iterative pruning. To prevent catastrophic accuracy degradation, iterative pruning is viewed as an effective retraining procedure [20, 21, 22]. For vision quality applications, quality metrics are required as a reference judgement for termination of pruning procedure.

3. PROPOSED METHOD

3.1. Architecture-aware Pruning

An output channel is pruned if its maximum absolute weight value is less than magnitude threshold. For convolutional layer, the weight kernel has tensor shape $i \times o \times k \times k$, where i is the number of input channels, o is the number of output channels and k is kernel size. Output-channel-wise pruning removes the weights along output channels. The kernel shape becomes $i \times (o - o') \times k \times k$ if o' output channels are pruned. The output-channel pruned ratio is defined as o'/o.

Once output channels of a layer are pruned, the corresponding input channels of the following layer are also removed. We defined one layer's sparsity as $1 - ((i - i')(o - o'))/(i \times o)$, where i' is the number of pruned input channel in the layer. The network sparsity is defined as the ratio of the number of zero weights of a pruned network divided by the number of total weights of the original network.

3.1.1. Keep Layer Depth

Usually, in vision quality applications, each layer in network is semantically designed for quality-sensitive primitives, such as edge and chroma, with respect to different resolutions. Intensively removing a layer can severely degrade the quality. Therefore, we keep the network architecture by preserving minimum number of output channels.

3.1.2. Enhance MAC Efficiency

A pruned network with higher weight sparsity may not imply higher computation reduction. We define MAC/weight, R_l (Eq. 1), for each layer as an indicator of MAC efficiency.
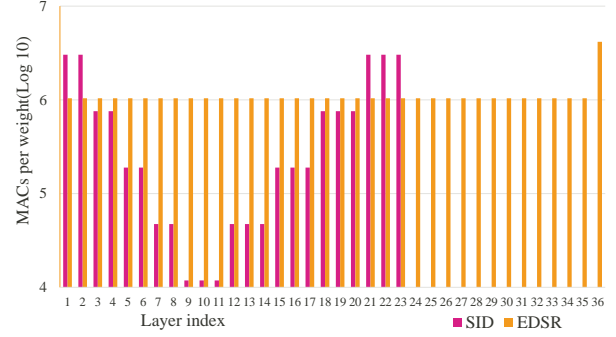
Fig. 2. MAC/weight are much larger on top and bottom layers in SID. However, MAC/weight for most layers are uniform in EDSR.

Fig. 3. An example of residual block. There are 4 layers with 6 output channels in each layer. Color parts represent removed output channels. One color denotes one group of channels in balance pruned output channel method.

M_l and W_l is number of MAC and weights of a layer, respectively. Fig. 2 shows that MAC/weight are much larger on both top and bottom layers on SID because of its U-Net [23] network topology. Therefore, to productively reduce computation, output channels of a layer with higher R_l are tend to be pruned more because of higher magnitude threshold.

$$R_l = log_{10}(M_l/W_l) \qquad (1)$$

3.1.3. Balance Pruned Output Channel

Residual block is universally used in network topology design such as in EDSR which is a variation of ResNet [24] with long shortcut. However, to prune output channels from residual block is arduous because of element-wise operations (i.e., element-wise ADD) or concatenation. Fig. 3 illustrates an example that a magnitude threshold is applied to 4-layer residual block. Because of element-wise ADD after layer Conv-B and layer Conv-D, the output channel of a given layer (Conv-D) and its preceding layer (Conv-B) with the same index (5) should be grouped and pruned at the same time. Conv-B layer

Algorithm 1 Architecture-aware and Quality Metric Guaranteed Pruning

---

1: **Input:** Target quality Q, Target Sparsity Increment S_i, Threshold Increment T_i, Total Step G
2: Target Sparsity $S = S_i +$ total-network-sparsity
3: Initial Threshold Base $T_b = T_i$
4: **repeat**
5:    **for** layer **in** network **do**
6:       S_l = pruned-output-channel-ratio(layer)
7:       M_l = MAC(layer)
8:       W_l = weight-size(layer)
9:       $R_l = log_{10}(M_l/W_l)$
10:      $T_l = T_b \times (1 - S_l) \times R_l$
11:      prune-output-channels-by-threshold T_l
12:    **end for**
13:    S_c = calculate-total-network-sparsity
14:    $T_b = T_b + T_i$
15: **until** $S_c > S$
16: **repeat**
17:    retrain-pruned-network
18:    Q_t = evaluate-quality-metric
19:    g = get-current-step
20: **until** ($Q_t > Q$ or $g >= G$)
21: **if** $g < G$ **then**
22:    jump to line 2
23: **end if**

---

and Conv-D layer have less pruned output channels compared to layer Conv-A and layer Conv-C.

We propose a guidance (Eq. 2) to prune output channels of residual block easier by increasing magnitude threshold on layers with lower ratio of pruned output channels. MAC efficiency mentioned in Sec. 3.1.2 is also applied. S_l is the ratio of pruned output channels of layer l. T_b is the original magnitude threshold base. Thus, output channels of a layer with lower S_l have higher tendency to be pruned.

$$T_l = T_b \times (1 - S_l) \times R_l \qquad (2)$$

3.2. Quality Metric Guarantee

To maintain quality metric (PSNR and SSIM) while maximizing pruned MAC, our algorithm prunes and retrains network iteratively. The iteration terminates when either the target quality-metric criteria or maximum training steps is reached. The proposed overall flow is shown in Algorithm 1.
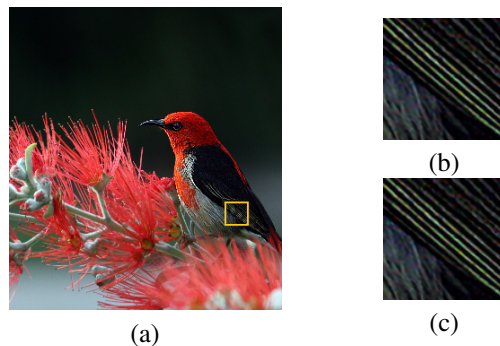
4. EXPERIMENTAL RESULT

4.1. Experiment Setup

We generally investigate both SID for low-light photography and EDSR (baseline network, ×2) for super resolution. SID uses its own dataset [4] and EDSR adopts DIV2K dataset

Fig. 4. SID results of our method compared to original (without pruning). (a)(c) Original (PSNR: 28.54, SSIM: 0.767). (b)(d) Pruned with Method-D (PSNR: 28.55, SSIM:0.768)

Fig. 5. EDSR results of our method compared to original (without pruning). (a) Image sampled from DIV2K. (b) Original (PSNR: 34.42, SSIM: 0.942). (c) Proposed Method-D (PSNR: 34.42, SSIM: 0.942)

[25]. The input size of the network is set to the maximum image resolution in the datasets, 1424×2128 for SID and 1020×1020 for DIV2K, to calculate MAC and BW.

In SID dataset, we use images captured by Sony α7SII camera as our training and validation data, which contains 280 pairs and 93 pairs, respectively. The pre-process stage is aligned with the setting in SID paper. In DIV2K dataset, we use the pre-process setting mentioned in [5] to generate 5,458,040 training patches from 800 training images and use 100 validation images as our validation data.

4.2. Result

The comprehensive analysis is elaborated in Table 1. We evaluate four distinct approaches. Method-A stands for magnitude threshold pruning without any structural hints.

**Table 1**. Detailed results. BW, considering only convolutional layers, consists of both weights and activations. Each weight and activation is represented with 4-byte floating-point numerical precision.

Network	Solution	# of MAC ($\times 10^9$)		# of Weights ($\times 10^3$)		# of Activations ($\times 10^6$)		BW (MByte/Inference)		Validation PSNR	Validation SSIM
SID	Original	560	(100%)	7757	(100%)	1915	(100%)	1922	(100%)	28.54	0.767
SID	Method-A	458	(82%)	6918	(89%)	1632	(85%)	1639	(85%)	28.54	0.768
SID	Method-B	354	(63%)	5275	(68%)	1485	(78%)	1491	(78%)	28.54	0.771
SID	Method-C	270	(48%)	5584	(72%)	1219	(64%)	1225	(64%)	28.54	0.769
SID	Method-D	236	(42%)	4241	(55%)	1169	(61%)	1173	(61%)	28.55	0.768
EDSR	Original	1428	(100%)	1367	(100%)	5076	(100%)	5077	(100%)	34.42	0.942
EDSR	Method-A	1085	(76%)	1037	(76%)	4481	(88%)	4481	(88%)	34.43	0.942
EDSR	Method-B	1085	(76%)	1037	(76%)	4481	(88%)	4481	(88%)	34.43	0.942
EDSR	Method-C	1085	(76%)	1037	(76%)	4481	(88%)	4481	(88%)	34.43	0.942
EDSR	Method-D	897	(63%)	857	(63%)	4083	(80%)	4083	(80%)	34.42	0.942

Fig. 6. Pruned output channel per layer on SID
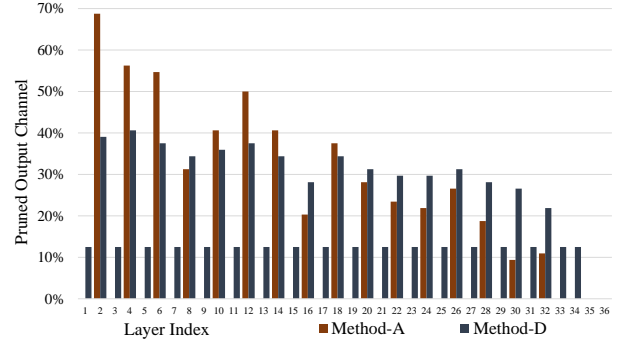
Fig. 7. Pruned output channel per layer on EDSR

Method-B keeps the depth of network (Sec. 3.1.1). Method-C further considers MAC/weight ratio (Sec. 3.1.2) on the basis of Method-B. Method-D integrates all proposed techniques. All methods are conducted in company with quality-metric constraints (Sec. 3.2). For both SID and EDSR, it shows no PSNR and SSIM drop on all methods. Fig. 4 and Fig. 5 reveal the indistinguishable quality difference.

Keep Layer Depth. In SID, Method-B reduces MAC from 82% to 63% compared to Method-A. As shown in Fig. 6, Method-A may remove all the output channels of a layer due to not keeping layer depth, which leads to severe quality metric drops that cannot be recovered in retraining steps.

Enhance MAC Efficiency. In SID, Fig. 6 shows that Method-C prunes more weights on both top and bottom layers that have larger MAC/weight (Eq. 1). Therefore, Method-C reduces MAC from 63% to 48% compared to Method-B.

Balance Pruned Output Channel. Method-D increases

17% weight sparsity but only reduces 6% MAC compared to Method-C in SID. Fig. 6 illustrates that Method-D prunes less on top and bottom layers which have larger MAC/weight.

In EDSR, there is no difference among Method-A, Method-B and Method-C because no layer is pruned by Method-A and MAC/weight are identical for all layers (last layer could not be pruned) as shown in Fig. 2. Fig. 7 shows that Method-D reduces MAC from 76% to 63%, which is more than 10%, in shortcut-connected layers.

In summary, our methodology has significant reduction on both MAC and BW, which implies reduction on inference latency. For BW, we have 39% and 20% reduction on SID and EDSR, respectively. Our methodology also works well on complex network architecture.

5. CONCLUSION

To minimize computation complexity without quality drop on vision quality applications, our architecture-aware pruning is optimized for pruning more for complexity metric (e.g., MAC) on SID and shortcut-connected layers on EDSR. The MAC of SID and EDSR are reduced by 58% and 37%, respectively. Memory bandwidth is also reduced without degradation of PSNR, SSIM and subjective quality. The reduction of computation complexity and memory bandwidth could benefit on general mobile devices without special hardware design.

6. REFERENCES

[1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in Neural Information Processing Systems 25*, F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, Eds., pp. 1097–1105. Curran Associates, Inc., 2012.

[2] Ross B. Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik, "Rich feature hierarchies for accurate object detection and semantic segmentation," *CoRR*, vol. abs/1311.2524, 2013.

[3] Jonathan Long, Evan Shelhamer, and Trevor Darrell, "Fully convolutional networks for semantic segmentation," *CoRR*, vol. abs/1411.4038, 2014.

[4] Chen Chen, Qifeng Chen, Jia Xu, and Vladlen Koltun, "Learning to see in the dark," *CoRR*, vol. abs/1805.01934, 2018.

[5] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee, "Enhanced deep residual networks for single image super-resolution," *CoRR*, vol. abs/1707.02921, 2017.

[6] 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," *CoRR*, vol. abs/1704.04861, 2017.

[7] R. Reed, "Pruning algorithms-a survey," *IEEE Transactions on Neural Networks*, vol. 4, no. 5, pp. 740–747, Sep. 1993.

[8] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li, "Learning structured sparsity in deep neural networks," *CoRR*, vol. abs/1608.03665, 2016.

[9] Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Tien-Ju Yang, and Edward Choi, "Morphnet: Fast & simple resource-constrained structure learning of deep networks," *CoRR*, vol. abs/1711.06798, 2017.

[10] Yann Le Cun, John S. Denker, and Sara A. Solla, "Optimal brain damage," in *Proceedings of the 2Nd International Conference on Neural Information Processing Systems*, Cambridge, MA, USA, 1989, NIPS'89, pp. 598–605, MIT Press.

[11] A. P. Engelbrecht, "A new pruning heuristic based on variance analysis of sensitivity information," *Trans. Neur. Netw.*, vol. 12, no. 6, pp. 1386–1399, Nov. 2001.

[12] Song Han, Jeff Pool, John Tran, and William J. Dally, "Learning both weights and connections for efficient neural networks," *CoRR*, vol. abs/1506.02626, 2015.

[13] Suraj Srinivas, Akshayvarun Subramanya, and R. Venkatesh Babu, "Training sparse neural networks," *CoRR*, vol. abs/1611.06694, 2016.

[14] Sajid Anwar and Wonyong Sung, "Compact deep convolutional neural networks with coarse pruning," *CoRR*, vol. abs/1610.09639, 2016.

[15] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz, "Pruning convolutional neural networks for resource efficient transfer learning," *CoRR*, vol. abs/1611.06440, 2016.

[16] A. Polyak and L. Wolf, "Channel-level acceleration of deep face representations," *IEEE Access*, vol. 3, pp. 2163–2175, 2015.

[17] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf, "Pruning filters for efficient convnets," *CoRR*, vol. abs/1608.08710, 2016.

[18] Yihui He and Song Han, "ADC: automated deep compression and acceleration with reinforcement learning," *CoRR*, vol. abs/1802.03494, 2018.

[19] Huizi Mao, Song Han, Jeff Pool, Wenshuo Li, Xingyu Liu, Yu Wang, and William J. Dally, "Exploring the regularity of sparse structure in convolutional neural networks," *CoRR*, vol. abs/1705.08922, 2017.

[20] Tien-Ju Yang, Yu-Hsin Chen, and Vivienne Sze, "Designing energy-efficient convolutional neural networks using energy-aware pruning," *CoRR*, vol. abs/1611.05128, 2016.

[21] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang, "Network trimming: A data-driven neuron pruning approach towards efficient deep architectures," *CoRR*, vol. abs/1607.03250, 2016.

[22] G. Castellano, A. M. Fanelli, and M. Pelillo, "An iterative pruning algorithm for feedforward neural networks," *IEEE Transactions on Neural Networks*, vol. 8, no. 3, pp. 519–531, May 1997.

[23] Olaf Ronneberger, Philipp Fischer, and Thomas Brox, "U-net: Convolutional networks for biomedical image segmentation," *CoRR*, vol. abs/1505.04597, 2015.

[24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," *CoRR*, vol. abs/1512.03385, 2015.

[25] Eirikur Agustsson and Radu Timofte, "Ntire 2017 challenge on single image super-resolution: Dataset and study," July 2017.