

MULTI-LOSS-AWARE CHANNEL PRUNING OF DEEP NETWORKS

Yiming Hu, Siyang Sun, Jianquan Li, Jiagang Zhu, Xingang Wang, Qingyi Gu

Research Center of Precision Sensing and Control, Institute of Automation, Chinese Academy of Sciences
School of Computer and Control Engineering, University of Chinese Academy of Sciences
{huyiming2016, sunsiyang2015, lijianquan2015, zhujiagang2015, xingang.wang, qingyi.gu}@ia.ac.cn

ABSTRACT

Channel pruning, which seeks to reduce the model size by removing redundant channels, is a popular solution for deep networks compression. Existing channel pruning methods usually conduct layer-wise channel selection by directly minimizing the reconstruction error of feature maps between the baseline model and the pruned one. However, they ignore the feature and semantic distributions within feature maps and real contribution of channels to the overall performance. In this paper, we propose a new channel pruning method by explicitly using both intermediate outputs of the baseline model and the classification loss of the pruned model to supervise layer-wise channel selection. Particularly, we introduce an additional loss to encode the differences in the feature and semantic distributions within feature maps between the baseline model and the pruned one. By considering the reconstruction error, the additional loss and the classification loss at the same time, our approach can significantly improve the performance of the pruned model. Comprehensive experiments on benchmark datasets demonstrate the effectiveness of the proposed method.

Index Terms— deep neural networks, object classification, model compression, channel pruning

1. INTRODUCTION

In recent years, deep convolutional neural networks (CNNs) have been widely applied to various computer vision tasks, e.g., image classification, object detection, action recognition, since AlexNet [1] won the ImageNet Challenge: ILSVRC 2012 [2]. However, it is very hard to deploy these deep models on resource-constrained devices including mobile robots, unmanned aerial vehicles, smartphones, etc. due to their high storage and computational cost.

In order to make CNNs available on resource-constrained devices, many model compression and acceleration methods have been presented including channel pruning [3, 4, 5], network quantization [6, 7, 8] and low-rank approximation [9, 10]. The method proposed in this paper falls into the category of channel pruning. Different from simply making sparse connections, channel pruning reduces the model size

by directly removing redundant channels. In contrast, pruning the entire channel can achieve fast inference without special software or hardware implementation. To prune redundant channels, existing reconstruction-based methods [3, 4, 5] usually minimize the reconstruction error of feature maps between the baseline model and the pruned one. DCP [11] seeks to conduct channel selection by introducing additional discrimination-aware losses. These methods suffer from two aspects of limitations. First, directly matching feature maps ignores the feature and semantic distributions within them. In other words, one should be encouraged to focus more on really important channels and spatial activations of feature maps instead of simply aligning every location. Second, the reconstruction error of a layer or additional losses do not always truly reflect the change of the classification loss, which may lead to mistakenly pruning some important channels.

In this paper, we address these problems from two aspects. First, the feature and semantic correlation loss is defined, through which the pruned model can preserve the feature and semantic distributions within feature maps of the baseline network. Second, we reconsider real contribution of channels to the overall performance. To this end, the classification loss of the pruned model is used to supervise channel selection. Then, we conduct layer-wise channel pruning by considering the reconstruction error, the additional loss and the classification loss at the same time.

Our main contributions are summarized as follows. First, we present a multi-loss-aware channel pruning method for deep networks compression by the introduction of the feature and semantic correlation loss and the classification loss. Second, extensive experiments on benchmark datasets show that the proposed method is theoretically reasonable and practically effective. On CIFAR-10, our approach outperforms the previous state-of-the-art method by 0.55% in accuracy.

2. RELATED WORK

The basic idea of weight pruning is to prune some redundant weights. Han et al. [12] proposed a conceptually simple pruning method: weights lower than a certain threshold are considered as low contribution ones that can be pruned, and then fine-tuning is taken for restoring the network accuracy.

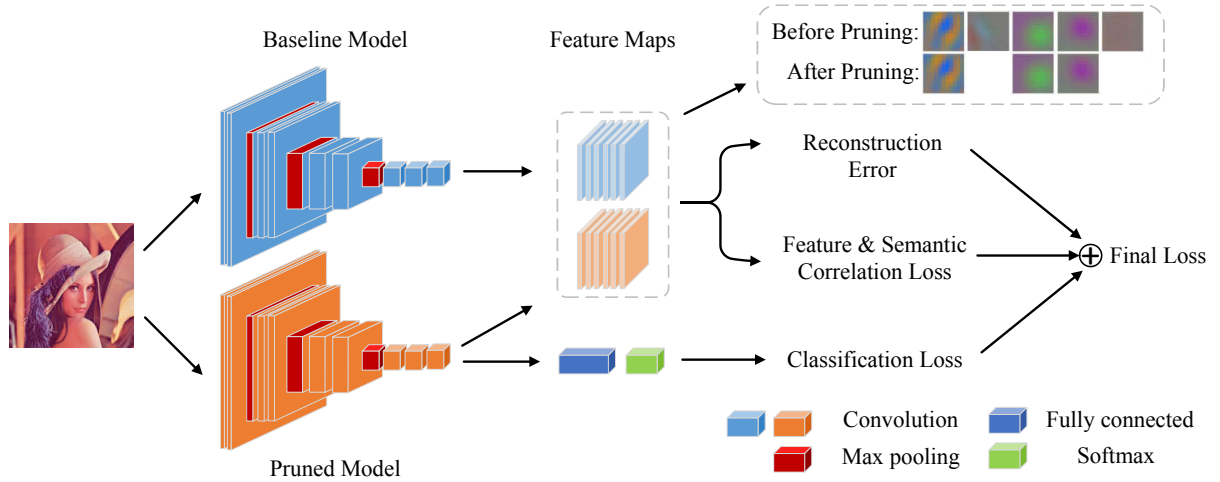


Fig. 1. Illustration of multi-loss-aware channel pruning. The final loss is used to supervise layer-wise channel selection.

However, the method needs special software or hardware implementation for fast inference. In view of this issue, Structured Sparsity Learning (SSL) [13] imposed regularizations on different levels of structures such as filters, channels or layers. Then weights belonging to the same level of structure would go to zero at the same time, which could be removed. This method obtains extraordinary performance but is computationally expensive and difficult to converge. Afterwards, some inference-based channel pruning methods [4, 14] were proposed, and they sought to do channel selection by defining different selection criteria. ThiNet [4] explicitly formulated channel selection as an optimization problem, and conducted channel pruning by minimizing the reconstruction error of feature maps between the baseline model and pruned one. DCP [11] aimed at selecting the most discriminative channels for each layer by considering both the reconstruction error and the discrimination-aware loss. However, these methods ignore the feature and semantic distributions within feature maps and real contribution of channels to the overall performance.

3. OUR METHOD

For better description of the proposed method, some notations are given first. Considering a layer of the CNN model, X represents the input tensor, W denotes the convolution filter and \otimes represents the convolution operation. We further use $F \in \mathbb{R}^{M \times H \times Z}$ to denote output feature maps of the layer. Here, M , H , Z represents the number of output channels, the height and the width of the feature maps respectively.

3.1. Motivation

Existing methods [3, 4, 11] conduct channel pruning by minimizing the reconstruction error of a layer, which can be de-

finied as the Euclidean distance of feature maps between the original model and the pruned one:

$$\mathcal{L}_r = \frac{1}{2T} \|X \otimes W - X \otimes W_{\mathcal{P}}\|_2^2 \quad (1)$$

where $T = M \times H \times Z$, \mathcal{P} denotes the index set of the selected channels, $W_{\mathcal{P}}$ represents the submatrix indexed by \mathcal{P} . Jiang et al. [5] explore the correlation between the layer-wise reconstruction error and the classification loss, and they demonstrate that these two variables show positive correlation. However, the correlation coefficient varies with different layers and some layers show low correlation. Consequently, simply minimizing the reconstruction error may mistakenly remove some important channels for some layers. Moreover, assuming each channel of F corresponds to a type of feature and each spatial activation encodes a type of semantic information, the feature and semantic distributions reflect which channels and spatial areas a CNN model focuses on. However, directly matching feature maps is not a good choice to preserve the feature and semantic distributions of the baseline model.

These limitations of existing methods inspire us to define a new loss: the feature and semantic correlation loss. The pruned model can preserve the feature and semantic distributions within feature maps of the baseline network using this loss. Meanwhile, we will explicitly use the classification loss to supervise layer-wise channel selection.

3.2. Formulation

To transfer rich semantic information of the baseline model, one can align the semantic distribution of two feature maps [15, 16]. In this paper, we consider aligning both the feature and semantic distributions by introducing the feature and semantic correlation loss. Here, the loss can be defined as the

Algorithm 1 The proposed method

Require: $\{W^l : 0 < l < L\}$: parameters of the pre-trained model, X : the training data, the pruning rate.

Ensure: $\{W_{\mathcal{P}}^l : 0 < l < L\}$: weights of the pruned model.

- 1: Initialize $W_{\mathcal{P}}^l$ with W^l for $\forall 1 \leq l \leq L$
- 2: **for** $l = 1, 2, \dots, L$ **do**
- 3: Construct the final loss \mathcal{L} shown as in Fig. 1
- 4: Perform channel selection for layer l by Eq 5
- 5: Update $W_{\mathcal{P}}^l$ w.r.t. the selected channels by Eq 6
- 6: **end for**
- 7: Fine-tune the pruned model

square error of two Gram matrices:

$$\mathcal{L}_s = \frac{1}{4N^2M^2} (\|G^f - G_{\mathcal{P}}^f\|_2^2 + \|G^s - G_{\mathcal{P}}^s\|_2^2) \quad (2)$$

where $N = H \times Z$, the Gram matrix $G^f \in \mathbb{R}^{M \times M}$ and $G^s \in \mathbb{R}^{N \times N}$ encode the feature and semantic correlation for a layer of the baseline model respectively. Similarly, $G_{\mathcal{P}}^f$ and $G_{\mathcal{P}}^s$ are Gram matrices corresponding to the pruned model. $G^f \in \mathbb{R}^{M \times M}$ is computed by the inner product between different vectorized features. $G^s \in \mathbb{R}^{N \times N}$ is computed by the inner product between different vectorized spatial activations across channels:

$$G_{ij}^f = \sum_{k=1}^N F_{ik} F_{jk}, G_{i'j'}^s = \sum_{k'=1}^M F_{k'i'} F_{k'j'} \quad (3)$$

To ensure important channels being kept, we explicitly use the classification loss to supervise layer-wise channel selection. By considering the reconstruction error, the feature and semantic correlation loss and the classification loss at the same time, the problem of channel pruning can be formulated to minimize the following joint loss function:

$$\begin{aligned} \min_{W_{\mathcal{P}}} \quad & \mathcal{L}(W_{\mathcal{P}}) = \mathcal{L}_r(W_{\mathcal{P}}) + \alpha \mathcal{L}_s(W_{\mathcal{P}}) + \beta \mathcal{L}_c(W_{\mathcal{P}}) \\ \text{s.t.} \quad & \|\mathcal{P}\|_0 \leq \mathcal{K} \end{aligned} \quad (4)$$

where α and β are two positive coefficients balancing the regularization terms, \mathcal{L}_c is the classification loss and \mathcal{K} is the number of channels to be selected.

3.3. Optimization Algorithm

It is a combinatorial optimization problem to choose a subset of channels according to Eq. 4. Directly searching competitive network structures in the whole solution space is NP-hard. Existing methods [4, 11, 5] seek to do channel selection with greedy algorithm by only considering the weight magnitude or the gradient of Eq. 4 w.r.t. the channel. In this paper, we take both of them into consideration. To be specific, the

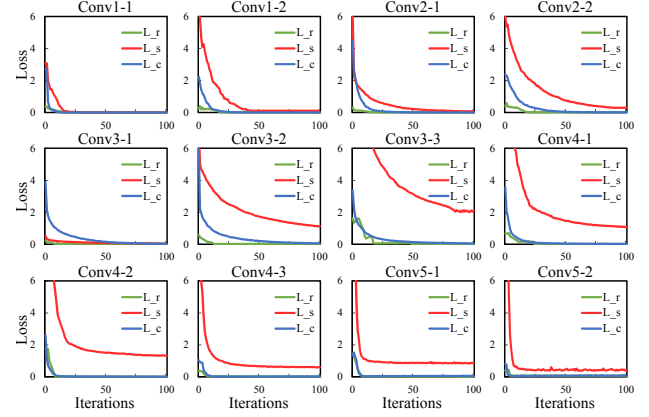


Fig. 2. The evolution curves of different losses over epochs on CIFAR-100.

sensitivity of the k -th channel is defined as:

$$\delta_k = \sum_{i=1}^H \sum_{j=1}^W \left\| \frac{\mathcal{L}}{W_{k,i,j}} W_{k,i,j} \right\|_2^2 \quad (5)$$

which is the square sum of weights of the k -th channel multiplying gradients of Eq. 4 w.r.t. the channel. Then we retain the channels with the i largest sensitivity and remove others. After this, Eq. 4 w.r.t. the selected channels is further optimized according to stochastic gradient descent (SGD). $W_{\mathcal{P}}$ is updated by:

$$W_{\mathcal{P}} = W_{\mathcal{P}} - \eta \frac{\partial \mathcal{L}}{\partial W_{\mathcal{P}}} \quad (6)$$

where η represents the learning rate. After updating $W_{\mathcal{P}}$, the channel pruning of a single layer is finished. The pruning process of the whole model is described in Algorithm 1.

4. EXPERIMENTS

In this section, our approach is evaluated on CIFAR-10 [19] and CIFAR-100 [19]. The proposed method is compared with several state-of-the-art methods including ThiNet [4], CP [3] and DCP [11]. Our approach is implemented using the PyTorch [20] framework. For the experiments on CIFAR-10 and CIFAR-100, we adopt the same experiments setting as in [11]. Besides, to balance the effects of different losses, α and β (Eq. 4) are set to 0.001 and 1 respectively.

4.1. Overall Comparison Results

To verify the effectiveness of the proposed approach, we conduct comparison experiments with several state-of-the-art methods on VGGNet and ResNet-56. Table 1 summarizes overall comparison results on CIFAR-10. It can be observed that the accuracy gap between our baseline model and the comparison reference model is very small. For VGGNet,

Table 1. Overall comparison results for VGGNet and ResNet-56 on CIFAR-10.

Models	Methods	Baseline (%)	Err. Gap (%)	#Param.	#Param. ↓	#FLOPs	#FLOPs ↓
VGG-16	CP [3]	6.01	+0.32	7.70M	1.92×	155.80M	2.00×
	ThiNet [4]		+0.14				
	DCP [11]		-0.17				
	Ours	6.08	-0.13	5.50M	2.26×	140.30M	2.23×
ResNet-56	CP [3]	6.20	+1.00	0.42M	1.97×	63.20M	1.99×
	ThiNet [4]		+0.82				
	WM [18]		+0.56				
	DCP [11]		+0.31				
	Ours	6.26	-0.24				

Table 2. The comparison results on ResNet-56 with different losses on CIFAR-10.

Methods	Training Err. (%)	Test Err. (%)
\mathcal{L}_r	3.75	9.74
\mathcal{L}_s	6.53	10.95
\mathcal{L}_c	0.85	8.35
$\mathcal{L}_r + \mathcal{L}_s$	1.31	8.27
$\mathcal{L}_r + \mathcal{L}_c$	0.73	8.14
$\mathcal{L}_s + \mathcal{L}_c$	1.03	8.20
$\mathcal{L}_r + \mathcal{L}_s + \mathcal{L}_c$	1.09	8.00

the proposed approach has a higher speedup and compression rate than other methods with the same accuracy drop. For ResNet-56, our approach achieves the best performance compared with previous state-of-the-art methods. With a compression rate of 1.97× and a speedup of 1.99×, the proposed method outperforms the DCP by 0.55% in accuracy drop. Moreover, our pruned ResNet-56 even outperforms the baseline model by 0.24% in accuracy.

4.2. Analysis of the Proposed Approach

To study the effects of different losses, we prune 30% channels of ResNet-56 on CIFAR-10 with different losses without fine-tuning. As shown in Table 2, the test error of the pruned model gets an immediate decline when introducing the feature and semantic correlation loss or the classification loss. It's worth noting that the pruned model obtains the best performance with three losses together. This demonstrates that multi-loss fusion is an effective strategy for channel pruning. Besides, the single \mathcal{L}_s may not be a better criterion of pruning compared with \mathcal{L}_r and \mathcal{L}_c , yet it can complement other losses in some aspects. This also explains why all mixed losses have better performance than a single loss.

Going one step further, we verify the convergence of our method. The loss function in Eq. 4 is optimized using SGD when pruning a single layer. Fig. 2 shows three losses (Eq. 4) change over epochs for the first twelve layers of VGGNet on CIFAR-100. It can be found that models with different

Table 3. The comparison results on ResNet-18 and ResNet-34 with different pruning rate on CIFAR-100.

Methods	Pruned rate (%)	Test Err. (%)
ResNet-18 (Baseline 21.89%)	30	24.39
	50	24.91
	70	25.15
ResNet-34 (Baseline 21.16%)	30	21.87
	50	22.41
	70	22.89

losses can converge to a minor error after 50 epochs. At Conv5-1 and Conv5-2, curves of three losses do not drop after 10 epochs. The above results demonstrate that the proposed method has a fast rate of convergence.

To further explore limits of the proposed approach, we prune ResNet-18 and ResNet-34 with different pruning rate on CIFAR-100. Experimental results in Table 3 show that the pruned model gets an immediate accuracy drop with increase of pruning rate. Besides, ResNet-18 is more sensitive to channel pruning than ResNet-34 due to its fewer parameters.

5. CONCLUSION

In this paper, a multi-loss-aware channel pruning method is presented for deep networks compression. The proposed approach achieves the state-of-the-art performance on benchmark datasets by considering the reconstruction error, the classification loss and the feature and semantic correlation at the same time. As for future works, we will combine existing pruning strategies with other network compression methods to explore more compact models with less accuracy drop.

6. ACKNOWLEDGMENT

This work has been supported by the National Key Research and Development Program of China No. 2018YFD0400902 and the National Natural Science Foundation of China (Grant No. 61673376 and 61573349).

7. REFERENCES

- [1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [2] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al., "Imagenet large scale visual recognition challenge," *International Journal of Computer Vision*, vol. 115, no. 3, pp. 211–252, 2015.
- [3] Yihui He, Xiangyu Zhang, and Jian Sun, "Channel pruning for accelerating very deep neural networks," in *International Conference on Computer Vision (ICCV)*, 2017.
- [4] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin, "Thinet: A filter level pruning method for deep neural network compression," *arXiv preprint arXiv:1707.06342*, 2017.
- [5] Chunhui Jiang, Guiying Li, Chao Qian, and Ke Tang, "Efficient dnn neuron pruning by minimizing layer-wise nonlinear reconstruction error.," in *IJCAI*, 2018, vol. 2018, pp. 2–2.
- [6] Peisong Wang, Qinghao Hu, Yifan Zhang, Chunjie Zhang, Yang Liu, Jian Cheng, et al., "Two-step quantization for low-bit neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 4376–4384.
- [7] Yuhui Xu, Yongzhuang Wang, Aojun Zhou, Weiyao Lin, and Hongkai Xiong, "Deep neural network compression with single and multiple level quantization," *arXiv preprint arXiv:1803.03289*, 2018.
- [8] Aojun Zhou, Anbang Yao, Kuan Wang, and Yurong Chen, "Explicit loss-error-aware quantization for low-bit deep neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 9426–9435.
- [9] Min Wang, Baoyuan Liu, and Hassan Foroosh, "Factorized convolutional neural networks.," in *ICCV Workshops*, 2017, pp. 545–553.
- [10] Wenqi Wang, Yifan Sun, Brian Eriksson, Wenlin Wang, and Vaneet Aggarwal, "Wide compression: Tensor ring nets," *learning*, vol. 14, no. 15, pp. 13–31, 2018.
- [11] Zhuangwei Zhuang, Minghui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu, "Discrimination-aware channel pruning for deep neural networks," in *Advances in Neural Information Processing Systems*, 2018, pp. 883–894.
- [12] 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*, 2015, pp. 1135–1143.
- [13] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li, "Learning structured sparsity in deep neural networks," in *Advances in Neural Information Processing Systems*, 2016, pp. 2074–2082.
- [14] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang, "Network trimming: A data-driven neuron pruning approach towards efficient deep architectures," *arXiv preprint arXiv:1607.03250*, 2016.
- [15] Zehao Huang and Naiyan Wang, "Like what you like: Knowledge distill via neuron selectivity transfer," *arXiv preprint arXiv:1707.01219*, 2017.
- [16] Sergey Zagoruyko and Nikos Komodakis, "Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer," *arXiv preprint arXiv:1612.03928*, 2016.
- [17] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang, "Learning efficient convolutional networks through network slimming," in *Computer Vision (ICCV), 2017 IEEE International Conference on*. IEEE, 2017, pp. 2755–2763.
- [18] 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.
- [19] Alex Krizhevsky and Geoffrey Hinton, "Learning multiple layers of features from tiny images," Tech. Rep., Citeseer, 2009.
- [20] Adam Paszke, Soumith Chintala, Ronan Collobert, Koray Kavukcuoglu, Clement Farabet, Samy Bengio, Iain Melvin, Jason Weston, and Johnny Mariethoz, "Pytorch: Tensors and dynamic neural networks in python with strong gpu acceleration, may 2017," .