

A one-step pruning-recovery framework for acceleration of convolutional neural networks

Author

Wang, D, Bai, X, Zhou, L, Zhou, J

Published

2019

Conference Title

Proceedings - International Conference on Tools with Artificial Intelligence, ICTAI

Version

Accepted Manuscript (AM)

DOI

[10.1109/ICTAI.2019.00111](https://doi.org/10.1109/ICTAI.2019.00111)

Rights statement

© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Downloaded from

<http://hdl.handle.net/10072/392920>

Griffith Research Online

<https://research-repository.griffith.edu.au>

A One-step Pruning-recovery Framework for Acceleration of Convolutional Neural Networks

Dong Wang, Lei Zhou, Xiao Bai
Beihang University, Beijing, China

Jun Zhou
Griffith University, Queensland, Australia

Abstract

Acceleration of convolutional neural network has received increasing attention during the past several years. Among various acceleration techniques, filter pruning has its inherent merit by effectively reducing the number of convolution filters. However, most filter pruning methods resort to tedious and time-consuming layer-by-layer pruning-recovery strategy to avoid a significant drop of accuracy. In this paper, we present an efficient filter pruning framework to solve this problem. Our method accelerates the network in one-step pruning-recovery manner with a novel optimization objective function, which achieves higher accuracy with much less cost compared with existing pruning methods. Furthermore, our method allows network compression with global filter pruning. Given a global pruning rate, it can adaptively determine the pruning rate for each single convolutional layer, while these rates are often set as hyper-parameters in previous approaches. Evaluated on VGG-16 and ResNet-50 using ImageNet, our approach outperforms several state-of-the-art methods with less accuracy drop under the same and even much fewer floating-point operations (FLOPs).

1. Introduction

Since AlexNet [14] demonstrated a great success in ILSVRC2012 [29], deep convolutional neural networks (CNNs) have led to rapid innovation in computer vision. Novel CNNs have been proposed to improve tasks ranging from image classification [31, 7], object detection [27, 26], to semantic segmentation [4]. One of the key reasons for this breakthrough is huge amount of parameters and depth of CNNs. Nevertheless, they also bring a side-effect that it is difficult to deploy complex CNNs in resource-constrained platforms like mobile phones.

Neural network compression and acceleration is an effective solution to this problem. Several neural network compression techniques have been proposed during the past years, for example, knowledge distillation [16, 39], tensor decomposition [11, 38], quantization [10, 6], and low-bit

arithmetic [35, 34]. Among these techniques, pruning is an important approach. Attempts in this direction include fine-grained and group-level pruning [32, 6] and structured sparsity learning [36]. The former discards connections between neurons according to their importance, while the latter regularizes filter, channel, filter shape and depth structures simultaneously. Although these methods have obtained relatively high compression ratios, their inference time is not significantly reduced due to irregularly pruned networks without co-designed hardware. To alleviate the problem, filter level pruning has been proposed to remove less important filters. The importance of the filters can be estimated using $L1$ norm of each filter [17], average percentage of zeros in the filter [9], or learning a sparse vector along the channel axis of its corresponding feature maps [8].

To avoid heavy accuracy drop after pruning filters, most approaches adopted an iterative pruning-recovery pipeline. After one layer is pruned with certain evaluation criterion, the pruned model has to be retrained to regain accuracy. Then the pruning-recovery cycle runs layer-by-layer iteratively to go through the whole network. This pipeline is very time consuming and can be even impractical as the CNNs go deeper, for which the training time increases linearly with the growth of the number of layers.

In this paper, we tackle this deficiency by proposing a one-step pruning-recovery framework. We adopt a layer selection method to select the layers to be pruned, while keeping dimensions of some important intermediate outputs constant. This can be implemented by keeping the numbers of filters in several important convolutional layers unchanged, and only prune those less important ones. To do so, we first learn an importance vector for each convolutional layer by adding a sparsity constraint to the normal CNN loss function. Then, this importance score is used to guide the whole network pruning process. After generating a complete pruned network, we can lightly regain the accuracy by reconstructing these crucial outputs simultaneously with our proposed optimization objective. In addition, our layer selection method also enables filter selection since it assigns a saliency score to each filter. In this way, given a global pruning rate, our method can automatically deter-

mine how many and which filters shall be eliminated from each single layer, and prune them in a one step operation rather than iteratively. This method solves the open problem on how to set the pruning hyper-parameters in previous works.

The main contributions of our method are summarized as follows:

1. **Concise pruning procedure.** We convert the conventional layer-by-layer pruning-recovery framework into a one-step framework, which significantly simplifies the pruning procedure.
2. **Low-cost optimization.** In the greedy pruning-recovery scenario, the cost of optimization is proportional to the number of layers in the network. Our one-step method significantly reduces this cost.
3. **Less hyper-parameters.** Since our method prunes filters globally and determines the pruning rate for each layer adaptively, it is not necessary to set hyper-parameters as in previous methods.
4. **Higher accuracy.** Our method is very effective. It outperforms the state-of-the-art CNN filter pruning methods by a large margin in accuracy under the same or even much fewer FLOPs.

2. Related Work

To guarantee high capacity of models, modern CNNs are often over-parameterized [30, 5]. This leads to long training time and difficulty in deployment on environment with limited computation resources. Model reduction by pruning is a widely adopted solution to this problem. Tran et al. and He et al. [32, 7] proposed iterative training and pruning methods to guarantee the accuracy of model while shrinking the network aggressively. Connections with weights smaller than a predefined threshold are pruned in order to gain a high compression ratio. However, these methods do not lead to acceleration in inference since the irregular pruned outcome needs the support of co-designed hardware. To address the disadvantage of non-structured pruning, some structured sparsity learning algorithms have been proposed. Lebedev and Lempitsky [15] proposed a group-wise brain damage process to produce sparse convolution kernels, achieving one sparsity pattern per group (2D kernels) in the convolutional layers. Then the entire group with small weights can be removed. Structured sparsity learning method has also been proposed to regularize filter, channel, filter shape and depth structures [36].

Recently, filter-level pruning has attracted considerable interests from both academia and industry. It aim to remove trivial filters and then fine-tune the networks to recover accuracy. Filter pruning methods differ on how they measure

the importance of the filters. The importance can be calculated based on absolute sum of weights of filters [17] or the average percentage of zeros in a filter [9]. Some methods converted the filter elimination problem to a feature map channel selection problem. Early attempts require a large number of random trials to filter channels, making the task time consuming [1]. On the other hand, activation pruning can be formulated as an optimization problem. Filters to be pruned can be learned based on the statistics of neighboring layer [21], or along the channel axis of the activation by leveraging LASSO regression [8]. Importance of filters can be calculated based on channel scaling factor [20], or by propagating the importance scores of final responses backwards. Lin et al. [19] accelerated CNNs via global and dynamic filter pruning. Furthermore, reinforcement learning [18] has also been used to determine the filters to be pruned.

To avoid heavy accuracy drop, all above-mentioned methods sorted to the iterative pruning-recovery pipeline which is time-consuming in practice. In our work, we tackle this drawback by proposing a one-step pruning-recovery CNN acceleration framework. Much less training cost and lower accuracy drop can be achieved with our approach.

Apart from pruning, other techniques for CNN acceleration include quantization [10, 6], knowledge distillation [16, 39], tensor decomposition [11, 38] and low-bit arithmetic [35, 34]. These methods are complementary and perpendicular to our pruning-based method, so we do not cover these approaches in the experiments, as a common practice in other works [21, 20].

3. The Proposed Method

In this section, we firstly describe notions and the general filter pruning framework. Then details are given on how our method converts the traditional iterative layer-by-layer pruning-recovery compression procedure into a one-step pruning-recovery method. After generating the complete pruned network in one step, our recovery procedure regains the accuracy of the pruned model with a novel optimization objective in one step. Fig. 1 shows the proposed framework and its comparison with the conventional iterative framework. Finally, we discuss the difference between our method and major filter pruning and knowledge distillation methods.

3.1. Layer-by-layer Filter Pruning

Considering a typical CNN with L convolutional layers, the weight of the i -th ($1 \leq i \leq L$) convolutional layer is a 4-D tensor $W^{(i)} \in \mathbb{R}^{C_{i+1} \times C_i \times M_i \times K_i}$, where C_{i+1} , C_i , M_i and K_i are the dimensions of the weight tensor along the axes of filter, channel, spatial height and spatial width, respectively. The output of the layer is a 3-D tensor $X^{(i)} \in \mathbb{R}^{C_{i+1} \times H_i \times W_i}$

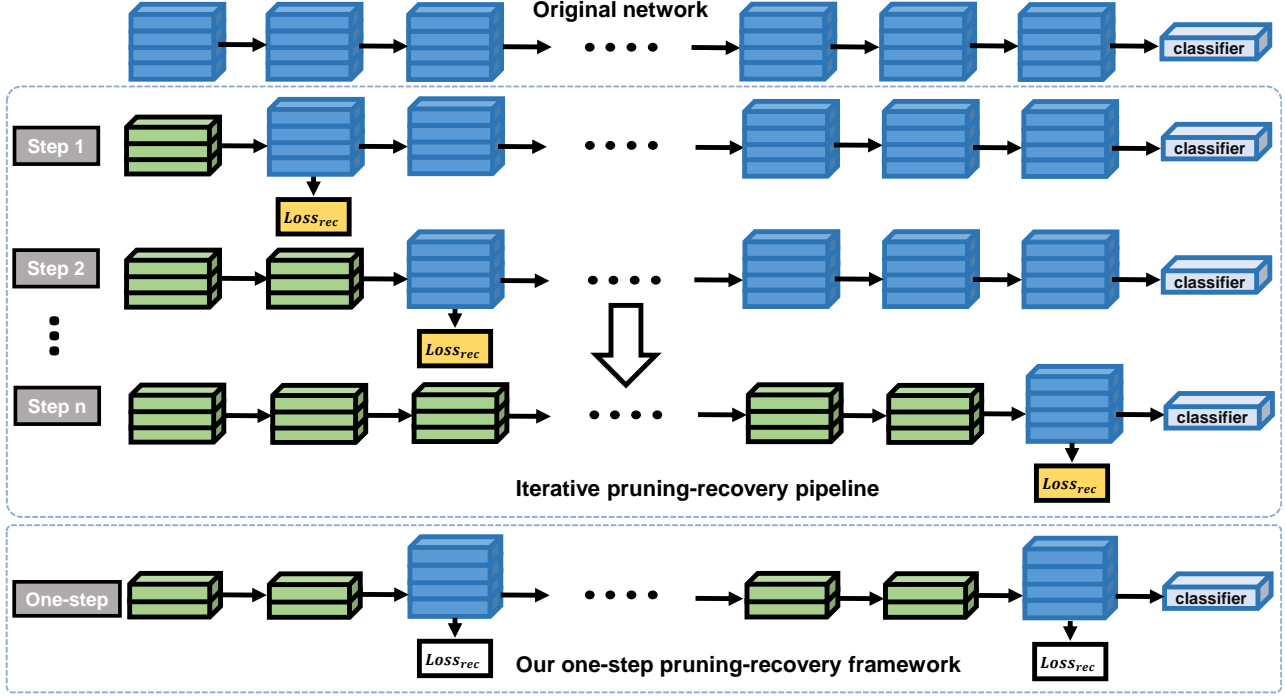


Figure 1. Illustration of two pruning framework. The blue and green cuboids are denoted as the original and the pruned convolutional layers, respectively. The traditional filter pruning methods prunes the whole network using a layer-by-layer pruning strategy. In each iteration, one layer is pruned and the reconstruction loss is minimized to regain accuracy of the network (represented by the yellow rectangle). In our framework, the whole network is only pruned in one step. Our method only prunes less important layers, leaving some crucial layers unchanged. Then the pruned model is restored by simultaneously reconstructing the output of these crucial layers with the proposed optimization objective (denoted by the white rectangle), so the model restoration only needs to be optimized once.

which has C_{i+1} feature maps with $H_i \times W_i$ in spatial size. In particular, we denote the input of the network as $X^{(0)}$.

Traditional filter pruning methods [8, 21, 20] work in such a scenario: given two consecutive convolutional layers, after filter pruning is applied to the former layer, the amount of input channels of the latter layer is reduced, but the number of feature maps from the latter layer remains unchanged. This property guarantees that the performance of the pruned network does not drop much by resorting to feature map reconstruction, which can be achieved by optimizing the following objective function:

$$\min_{W^{(i+1)}} \|X^{(i+1)} - X^{(i)} * W^{(i+1)}\|_F^2 \quad (1)$$

where $\|\cdot\|_F$ is the Frobenius norm and $*$ is the convolution operation. $X^{(i)} = X^{(i-1)} * W^{(i)} \in \mathbb{R}^{C'_{i+1} \times H_i \times W_i}$ is the feature map produced by the pruned i -th convolutional layer with the shrinking weight $W^{(i)} \in \mathbb{R}^{C'_{i+1} \times C_i \times M_i \times K_i}$. $W^{(i+1)} \in \mathbb{R}^{C_{i+2} \times C'_{i+1} \times M_{i+1} \times K_{i+1}}$ is the parameter tensor of the $(i+1)$ -th convolutional layer which is compressed along the channel axis. To compress the whole network, this single layer pruning strategy is applied to the model layer-by-layer, whose complexity and time cost is proportional to the number of convolutional layers.

In principle, iterative pruning-recovery algorithms could prune the whole network in one step with their proposed filter evaluation criteria and directly reconstruct the output of the final convolutional layer to restore the totally pruned model. However this will be slow to converge. Moreover, significant accuracy drop may happen and can not be compensated even by the subsequent time consuming fine-tuning. Therefore, iterative pruning-recovery becomes a preferred option in the literature.

3.2. One-step filter pruning

We aim to generate the complete pruned network that will be restored in our one-step recovery procedure. Intuitively if a score can be assigned to each filter in every layers in one step, filters with low scores can be removed simultaneously to obtain the pruned model. To achieve this goal, we calculate the rank of all filters by applying a sparsity constraint to the channel dimension. Specifically, for the i -th convolutional layer, we multiply an importance vector $\beta^{(i)} \in \mathbb{R}^{C_{i+1}}$ to the $X^{(i)}$ channel:

$$Z_{j,:}^{(i)} = \left| \beta_j^{(i)} \right| X_{j,:}^{(i)} \quad \forall j \in \{1, \dots, C_{i+1}\} \quad (2)$$

Then we continue the forward propagation by feeding the scaled intermediate output $Z^{(i)}$ to the rest of the network.

The importance vector for each convolutional layer can be learned by optimizing the normal training loss function of a CNN with a sparsity constraint that pushes values in β vector to 0:

$$Loss_{sel}(X^{(0)}, Y; \mathbf{W}, \beta) = Loss(X^{(0)}, Y; \mathbf{W}, \beta) + \lambda \|\beta\|_1 \quad (3)$$

where \mathbf{W} is the original parameters of the network, Y is the ground-truth, and λ is a hyper-parameter that trades off the loss and the sparsity constraint. During the training, \mathbf{W} is fixed and only β is updated. We initialize the values of entries in β with 1. After optimization, a score is obtained for each channel of each convolutional layer.

3.3. One-step recovery

Normally the last convolutional layer of a CNN would not be pruned after the whole network pruning since it is usually followed by a fully-connected layer. Base on this rule, a reasonable insight to recover a pruned model is directly reconstructing the output of the final layer. However, we find that this option is slow to converge and the accuracy decays by a large margin. Motivated by [24, 3], we leverage intermediate supervision technology to achieve quick convergence and much less drop of accuracy.

Our method reconstructs several significant intermediate outputs as well as the last feature maps simultaneously. Since the premise of activation reconstruction is that the shape of the activation is constant after pruning, the dimensionality of the feature maps produced by these significant layers should remain unchanged after pruning the whole network. Thus, a constraint should be applied to our one-step network pruning procedure, i.e., the number of filters in a crucial layer shall not change even if the layer contains trivial filters. Therefore, the network pruning only applies to less important layers.

We use the learning results in our one-step pruning procedure to define the importance of a convolutional layer. Specifically, the mean of all absolute channel scores in a convolutional layer is taken as its importance score. For the i -th convolutional layer, its score is calculated as:

$$score^{(i)} = \sum_{j=1}^{C_{i+1}} |\beta_j^{(i)}| \quad (4)$$

Subsequently, our one-step recovery procedure reconstructs the outputs of the crucial nodes simultaneously. Formally, define $F^{(i)}$ and $F'^{(i)}$ as the sub-network of the original network and the pruned one starting from the input layer to the i -th convolutional layer, respectively. The learnable parameters of $F^{(i)}$ and $F'^{(i)}$ are denoted as $\mathbf{W}^{(i)}$ and $\mathbf{W}'^{(i)}$. Given the indices set $S = \{l_1, \dots, l_N\}$ of the convolutional nodes where reconstruction will occur, we enforce

the response of the l_i -th convolutional layer of the pruned network to mimic the corresponding output of the original model as follows:

$$Loss_{l_i}(X^{(0)}; \mathbf{W}^{(l_i)}, \mathbf{W}'^{(l_i)}) = \|X^{(l_i)} - X'^{(l_i)}\|_F^2 \quad (5)$$

$$X^{(l_i)} = F^{(l_i)}(X^{(0)}; \mathbf{W}^{(l_i)}) \quad (6)$$

$$X'^{(l_i)} = F'^{(l_i)}(X^{(0)}; \mathbf{W}'^{(l_i)}) \quad (7)$$

where $\mathbf{W}^{(l_i)}$ is fixed and only $\mathbf{W}'^{(l_i)}$ is updated during the training. To regain the accuracy after one-step filter pruning, we optimize the reconstruction error per crucial layer simultaneously:

$$Loss_{rec} = \frac{1}{N} \sum_{i=1}^N Loss_{l_i} \quad (8)$$

Feature map reconstruction can either be employed before or after non-linear activation. Li et al. [17] gathered all feature map statistics prior to non-linear activations or batch normalization, whilst others [9, 23] adopted post non-linear activation. However, when all responses are reconstructed concurrently, the feature maps must be restored after the non-linear activation, e.g., ReLU [22], otherwise exploding gradients will occur. Thus, for a convolutional layer followed by a non-linear activation, Eq. (5) can be re-written as:

$$Loss_{l_i} = \|\sigma(X^{(l_i)}) - \sigma(X'^{(l_i)})\|_F^2 \quad (9)$$

where σ is the non-linear activation.

Furthermore, besides the traditional mean square error (MSE), alternative metrics can be used to mimic the feature maps. In total, four mimicking functions are considered: 1) MSE, 2) LASSO, 3) KL-divergence, and 4) JS-divergence. In the experiments, KL and JS divergence measure the distribution similarity between two groups of feature maps, which achieve better performance compared with the element-wise norm reconstruction. The default option in our framework is the KL-divergence. Therefore, for the l_i -th convolutional layer, we have:

$$Loss_{l_i} = \frac{1}{H_{l_i} \times W_{l_i}} \sum_{j=1}^{H_{l_i}} \sum_{k=1}^{W_{l_i}} KL(P_{j,k}^{(l_i)}, P'_{j,k}{}^{(l_i)}) \quad (10)$$

$$P_{j,k}^{(l_i)} = softmax(\sigma(X_{:,j,k}^{(l_i)})) \quad (11)$$

$$P'_{j,k}{}^{(l_i)} = softmax(\sigma(X'_{:,j,k}{}^{(l_i)})) \quad (12)$$

where

$$KL(p, q) = \sum_i p_i \log \frac{p_i}{q_i} \quad (13)$$

$$softmax(z)_i = \frac{e^{z_i}}{\sum_{k=1}^K e^{z_k}} \quad (14)$$

for $i = 1, \dots, K$. It should be noted that reconstruction based on the KL or JS divergence will not work if only the response of the last convolutional layer is considered. This is because the constraint is too weak for the following fully connected layer for classification. Thus, two conditions are required: (1) $N \geq 2$ and (2) the final convolutional layer needs to be manually appended to S if it does not exist in the set.

3.4. Discussion

Filter selection and pruning. While He et al. [8] removed filters by learning sparse vectors along the channel dimension, the impact between layers was not taken into consideration due to the layer-by-layer pruning strategy. On the contrary, our method assesses the filters based on a global context. A similar optimization objective for filter selection focuses on online iterative filter pruning [20]. This is completely different from our method which aims to get the complete pruned network in one step. While some methods [23, 37] also measure significance of all filters in one-step via back propagation, their estimation is biased since these approaches only depend on a mini-batch of samples. In contrast, our one-step filter pruning strategy estimates filters with the whole training dataset, leading to more reliable assessment.

In summary, the core contribution of this paper is the one-step recovery strategy. Leveraging it, we can regain a complete pruned network with less accuracy drop as well as fewer FLOPs, when compared with the traditional iterative pruning-recovery methodology. In addition, the cost of training can be greatly reduced with our framework. Thus, filter selection strategy is auxiliary in our work and we will simplify it aggressively in our future work.

Knowledge distillation (KD). Our method can also be regarded as the application of knowledge distillation to filter pruning. The original and pruned network can be taken as a teacher and student model, respectively. However, current KD-based approaches compress CNNs by devising a new tiny model or with layer-level pruning [39]. Our method is fundamentally different by operating on a different data granularity and avoiding the conventional iterative pruning-recovery process.

4. Experiment

We have completed comprehensive experiments to evaluate our method. The first part of the experiments is to evaluate each component of our algorithm. We aim to demonstrate (1) crucial layers can be chosen with our one-step pruning strategy; (2) how the amount of mimicked nodes

¹Symbols from res2a to res5c represent layers in ResNet-50 with increasing depth gradually. For convenience, we omit the suffix, `.relu`, for each layer.

Table 1. Mimicking location.

Model	Layers
Model_auto	res2c ¹ , res3d, res4f, res5c
Model_manual1	res2a, res3a, res4a, res5c
Model_manual2	res2b, res3b, res4b, res5c
Model_manual3	res2c, res3c, res4c, res5c
Model_manual4	res2c, res3d, res4d, res5c
Model_manual5	res2c, res3d, res4e, res5c
Model_mimic1	res5c
Model_mimic2	res4f, res5c
Model_mimic3	res3d, res4f, res5c

affect the performance of our framework; (3) the behavior of our method under different reconstruction functions; and (4) our framework can be aggressively boosted with our one-step filter pruning algorithm since it works as a powerful filter selector compared with other naive versions. These evaluations are done on ResNet-50 [7] using CIFAR-10 [13] which consists of 50k images for training and 10k for testing in 10 classes.

In the second part of the experiments, we evaluate our approach on VGG-16 [31] and ResNet-50 [7] using ImageNet [29]. All experiments were implemented using PyTorch [25] on two NVIDIA 1080Ti GPUs. The networks were optimized with Adam [12].

4.1. Identifying crucial layers with our one-step pruning strategy

We first show the effectiveness of our one-step pruning strategy for selecting crucial layers. By setting λ (in Eq. (3)) to 1.0 and optimizing its objective for 15 epoches with a learning rate $1.0e^{-5}$, the rank of all convolutional layers was obtained. The first 4 nodes were chosen as shown in the column *Layers* of the 1th row in Table 1. The column *Layers* indicates mimicked nodes and *Model* represents the model trained by reconstructing corresponding layers. We denoted the model trained by mimicking the four nodes as Model_auto. In comparison, we manually selected several groups of nodes with different depth. Models that were trained based on these manually selected layers were named from Model_manual1 to Model_manual5 respectively. Details can be seen in Table 1.

For training models with feature maps reconstruction, we first copied the classification layers from the original network and fixed them during the training. Then, for convenience, we shrank each convolutional layer to be pruned with a compression ratio 8.0 and randomly copied matching amount of filters from the original model. Finally, the models were trained with a batch size 128 for 15 epoches. Learning rate was initially set as $1.0e^{-3}$ and then divided by 10 after every 5 epoches. Eq. (8) with MSE mimicking function was employed as the learning objective. This

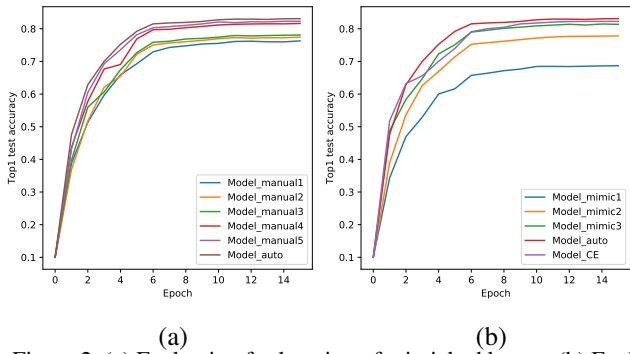


Figure 2. (a) Evaluation for location of mimicked layers. (b) Evaluation for amount of mimicked layers.

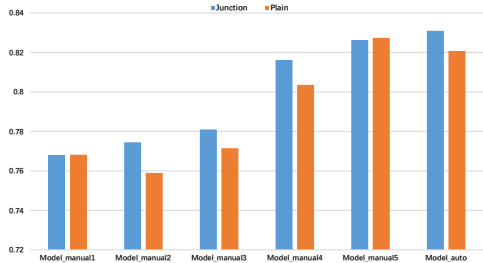


Figure 3. Comparison between mimicking junction node and the plain.

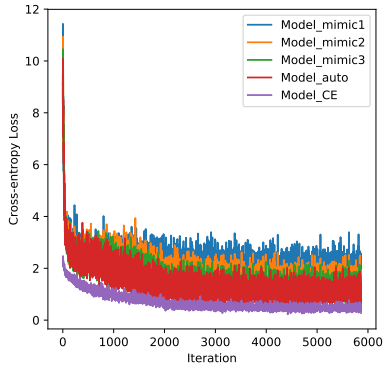


Figure 4. Variation of minibatch cross-entropy loss.

training setting was also adopted in other module evaluation, so we do not repeat it unless specific points need to be addressed.

The results are shown in Fig. 2(a). From this figure we can see that our layer identification strategy could overwhelm the handcrafted counterparts. One reason for it may be that our strategy regards the deeper layers as the crucial ones whilst those manually selected layers are the shallow ones, which is also consistent with the experiences and conclusions shared in several works [8, 21]. Another reason for it could be that better high-level semantic features

which dominate in performance of a model can be obtained by reconstructing deeper layers. This leads to the conclusion that our one-step pruning strategy could select the crucial layers for reconstruction and greatly improve our whole framework.

Since our selection results are the junction nodes, we manually chose several groups of the plain layers, e.g., convolutional layers in the main branch of the residual block for ResNet, as the comparative trials. We replaced each layer in *Layers* for each model in Table 1 with its corresponding parent node located in the main branch of the bottleneck block in ResNet-50. Fig. 3 shows the details on this comparison. In most cases, the mimicking junction nodes surpass the plain layer reconstruction by a large margin. Two exceptions are with Model_manual1 and Model_manual5, in which the slight difference could be negligible. In summary, our framework should mimic the important nodes, the junction or the deep layers, and our node identification strategy can give the right dependence for the subsequent one-step recovery procedure.

4.2. Influence of amount of mimicked nodes

In this experiment, we analyse the influence of amount of mimicked layers. For convenience, we selected Model_auto in Table 1 as the baseline. Three models for comparison, denoted as Model_mimic3, Model_mimic2 and Model_mimic1 respectively, were generated by iteratively deleting one layer from the layers of Model_auto. A model trained with a classic classification criterion of cross-entropy loss was denoted as Model_CE.

The comparison results are shown in Fig. 2(b). Three observations can be made from this figure. First, although solely reconstructing the final activation could work, which is also consistent with intuition, its testing accuracy is low. Second, the accuracy is proportional to the number of the mimicked layers while the effect would tend towards saturation with the number of layers increasing. Last, our reconstruction strategy surpasses the option of training with conventional classification objective. Variation of minibatch cross-entropy loss for each model is shown in Fig. 4.

Interestingly, while Model_CE is suboptimal, its cross-entropy loss is the minimal. This manifests that our method could strengthen the generalization capability to a certain extent, compared with directly optimizing the ordinary classification objective. In conclusion, whilst there is a positive correlation between model performance and the number of mimicked layers, the improvement becomes not significant with a large number of mimicked layers. In practice, a small quantity is sufficient for performance boosting. Thus, the constraint on the amount of unaltered nodes imposed on our approach would not obviously cripple its strength of pruning filters.

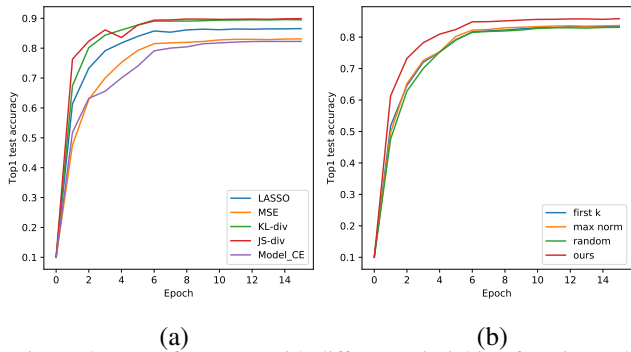


Figure 5. (a) Performance with different mimicking functions. (b) Performance with different filter selection strategies.

4.3. Comparison between diverse mimicking functions

The aim of this experiment is to study how different mimicking functions can affect the model performance. Four mimicking functions were taken into consideration: (1) MSE, (2) LASSO regression, (3) KL-divergence and (4) JS-divergence. Model_{auto} in Table 1 was selected as the pruned model with each mimicking function applied in turn. The results are displayed in Figure 5(a).

This figure shows that frequently used MSE is the worst in practice. In addition, it is worth noting that mimicking the probability distribution is more efficient than fitting the real-valued activation. This perhaps could be ascribed to two reasons: (1) real-valued mimicking may be unreliable for values near zero, and the error can be accumulated; (2) probability distribution fitting acts as a regularizer as revealed in [2, 28]. However, merely mimicking the response of the final convolutional layer with KL-divergence or JS-divergence does not work, which indirectly proves that reconstructing the distribution for several intermediate activations and the last one guarantees the content similarity between the final responses produced by the original network and the pruned model.

4.4. Boosting performance with our node identification strategy

Since our one-step filter pruning strategy works not only for layer selection but also as a filter selector, we explore how this function affects the performance of our method. To evaluate it, we considered several alternative selection strategies:

- *random* - shrinks a single convolutional layer randomly with a predefined compression ratio.
- *first k* - selects the first k filters.
- *max response* - selects filters that have high absolute weight sum [17].

From Fig. 5(b), we can see that whilst there is no obvious difference between three naive criteria, our method surpasses them by a large margin. An intuitive explanation on this advantage is that our filter selection strategy takes the coupling between layers into consideration, while others [8, 21, 17, 33] independently evaluate filters layer-by-layer. In addition, given a predefined pruning rate, our node identification strategy determines how many filters are to be remained for each layer according to the learned channel scores, solving an open problem for hyper-parameter setting in previous works [8, 21]. To sum up, besides picking right nodes, our node identification strategy can also work as a filter selector to decide how many filters shall be pruned for every selected convolutional layer. This also improved the accuracy of the pruned networks.

4.5. Experiments on ImageNet ILSVRC-12

In the second part of the experiments, we evaluated the performance of our method for VGG-16 [31] and ResNet-50 [7] on large-scale ImageNet classification task. The ILSVRC-12 dataset [29] consists of over 1.28 million training images drawn from 1000 categories. Images were resized such that the shorter side is 256. The training and testing were on random and center crop of 224×224 pixels, respectively, followed by mean-std normalization.

VGG-16. For layer selection we optimized Eq.(3) with 10 training epochs by setting batch size, λ , and learning rate with 128, 0.1 and $1.0e^{-4}$ respectively. According to the learnt result we made the last three nodes, i.e., conv5_1, conv5_2 and conv5_3, mimicked. Then the model was trained with batch size 128 for 15 epoches to optimize Eq. (8). Learning rate was initially set to 0.001 and then divided by 10 after every 5 epoches. After reconstruction, we retrained the intermediate model for 20 epoches with batch size 128 and fixed the learning rate to $1.0e^{-5}$. We evaluated our method by reducing various amounts of FLOPs. The results are shown in Table 2.

As can be seen from this table, our method leads to less accuracy loss with similar and even much fewer FLOPs. For example, we can surpass [8] under the same FLOPs. When compared with Taylor ($2.7\times$) [23], the accuracy loss from our method is lower, but with reduced FLOPs by 80.03% against 62.86%. More importantly, our approach is efficient yet simple because the traditional layer-by-layer pruning-recovery optimization is not needed. After determining the structure of the pruned network, we can regain the accuracy by optimizing our novel reconstruction objective end-to-end for the whole network, with slight accuracy drop that can be aggressively compensated by fine-tuning. Therefore, the training cost can be greatly reduced with our method. In addition, we also reported the performance of our algorithm before fine-tuning in Table 4 to aggressively show the effectiveness of our proposed reconstruction objective,

Table 2. Comparison of several state-of-the-art filter pruning methods for VGG-16 on ImageNet. All accuracies were tested on the validation set using single view central patch crop. Note that methods are ordered based on their top-5 accuracy drop.

Method	Top-5 Acc. Baseline (%)	Top-5 Acc. (%)	Top-5 Acc. Drop (%)	Pruned FLOPs (%)
GDP [19]	89.42	87.95	1.47	75.48
Ours (4.4×)	90.38	88.84	1.54	77.28
CP (4.4×) ² [8]	89.90	88.10	1.70	77.28
Ours (5×)	90.38	88.38	2.00	80.03
Taylor (2.7×) [23]	89.30	87.00	2.30	62.86
RNP (3×) [18]	89.90	87.58	2.32	66.67
SSS [9]	90.84	88.20	2.64	75.24
Ours (6×)	90.38	87.33	3.05	83.48
RNP (4×) [18]	89.90	86.67	3.23	75.00
RNP (5×) [18]	89.90	86.32	3.58	80.00
Ours (7×)	90.38	85.89	4.49	85.78
Taylor (3.9×) [23]	89.30	84.50	4.80	74.16

Table 3. Comparison of several state-of-the-art filter pruning methods for ResNet-50 on ImageNet. All accuracies were tested on the validation set using single view central patch crop. Note that methods are ordered based on their top-5 accuracy drop.

Method	Top-5 Acc. Baseline (%)	Top-5 Acc. (%)	Top-5 Acc. Drop (%)	Pruned FLOPs (%)
ThiNet-50 [21]	91.14	90.02	1.12	55.83
Ours (2.8×)	92.87	91.64	1.23	64.64
CP (2.8×) [8]	92.20	90.80	1.40	64.64
Ours (3×)	92.87	91.43	1.44	66.84
SSS (ResNet-26) [9]	92.86	90.79	2.07	43.04
GDP [19]	92.30	90.14	2.16	59.33
ThiNet-30 [21]	91.14	88.30	2.84	71.50
Ours (4×)	92.87	89.75	3.12	75.00
Ours (5×)	92.87	87.74	5.13	80.00

i.e., Eq.(8). From the results, we can see that our method outperforms [8] by a large margin with the same FLOPs before fine-tuning and only suffers 3.48% top-5 accuracy drop. The superiority is ascribed to our novel reconstruction objective.

ResNet-50. For layer identification, we adopted the same settings as in our VGG-16 experiment. Based on the learnt results, *Layers* of Model.auto in Table 1 was mimicked. Subsequently the model was trained with batch size 64 for 15 epoches. The learning rate was initially set to 0.001 and then divided by 10 after every 5 epoches. After reconstruction, we retrained the intermediate model for 20 epoches with batch size 64 and fixed learning rate $1.0e^{-5}$. The results were shown in Table 3. Our method exhibits consistent outstanding performance, achieving higher accuracy under fewer FLOPs on ResNet. Performance before fine-tuning is also reported in Table 4. The effectiveness of our optimization objective is consistent.

Table 4. Performance of our method before fine-tuning for VGG-16 and ResNet-50 on ImageNet.

Method	Model	Top-5 Acc. Drop
Ours (4.4×)	VGG-16	3.48%
CP (4.4×) [8]	VGG-16	22.00%
Ours (2.8×)	ResNet-50	2.69%
CP (2.8×) [8]	ResNet-50	8.00%

5. Conclusions

To simplify conventional filter pruning procedure for CNNs, we have introduced an efficient framework which replaces the traditional layer-by-layer pruning-recovery framework with a one-step version by simultaneously mimicking several crucial nodes which are unaltered during whole network pruning. Our approach is simple yet efficient. Much lower optimization cost is required to regain the performance of the pruned network when compared with alternative approaches. Furthermore, our approach can achieve less accuracy drop under the same and even much fewer FLOPs compared with several state-of-the-art methods. The advantages of our method has been demonstrated by extensive experiments on benchmark CNNs and

²The precise speed-up ratios of [8] for VGG-16 (5×) and ResNet-50 (2×) are 4.4× and 2.8× respectively, according to the models released by the authors of [8].

datasets.

References

- [1] S. Anwar and W. Sung. Compact deep convolutional neural networks with coarse pruning. *CoRR*, abs/1610.09639, 2016.
- [2] L. J. Ba and R. Caruana. Do deep nets really need to be deep? In *NIPS*, pages 2654–2662, 2014.
- [3] Z. Cao, T. Simon, S.-E. Wei, and Y. Sheikh. Realtime multi-person 2d pose estimation using part affinity fields. In *CVPR*, 2017.
- [4] L. C. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected CRFs. In *ICLR*, 2015.
- [5] E. Denton, W. Zaremba, J. Bruna, Y. Lecun, and R. Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *NIPS*, pages 1269–1277, 2014.
- [6] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *ICLR*, 2016.
- [7] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [8] Y. He, X. Zhang, and J. Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, 2017.
- [9] H. Hu, R. Peng, Y. Tai, and C. Tang. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. *CoRR*, abs/1607.03250, 2016.
- [10] B. Jacob, S. Kligys, B. Chen, M. Zhu, M. Tang, A. Howard, H. Adam, and D. Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *CVPR*, pages 11–20, 2017.
- [11] M. Jaderberg, A. Vedaldi, and A. Zisserman. Speeding up convolutional neural networks with low rank expansions. In *BMVC*, 2014.
- [12] D. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [13] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.
- [14] A. Krizhevsky, I. Sutskever, and G. E. Hinton. ImageNet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [15] V. Lebedev and V. Lempitsky. Fast ConvNets using group-wise brain damage. In *CVPR*, pages 2554–2564, 2016.
- [16] J. B. Lei and R. Caruana. Do deep nets really need to be deep? In *NIPS*, pages 2654–2662, 2014.
- [17] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf. Pruning filters for efficient ConvNets. In *ICLR*, pages 1–13, 2017.
- [18] J. Lin, Y. Rao, J. Lu, and J. Zhou. Runtime neural pruning. In *NIPS*, pages 1097–1105, 2017.
- [19] S. Lin, R. Ji, Y. Li, Y. Wu, F. Huang, and B. Zhang. Accelerating convolutional networks via global & dynamic filter pruning. In *IJCAI*, 2018.
- [20] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang. Learning efficient convolutional networks through network slimming. In *ICCV*, pages 2755–2763, 2017.
- [21] J.-H. Luo, J. Wu, and W. Lin. ThiNet: A filter level pruning method for deep neural network compression. In *ICCV*, 2017.
- [22] A. L. Maas, A. Y. Hannun, and A. Y. Ng. Rectifier nonlinearities improve neural network acoustic models. In *ICML*, volume 30, page 3, 2013.
- [23] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz. Pruning convolutional neural networks for resource efficient inference. In *ICLR*, 2017.
- [24] A. Newell, K. Yang, and J. Deng. Stacked hourglass networks for human pose estimation. In *ECCV*, 2016.
- [25] A. Paszke, S. Gross, S. Chintala, G. Chanan, E. Yang, Z. DeVito, Z. Lin, A. Desmaison, L. Antiga, and A. Lerer. Automatic differentiation in PyTorch. In *Workshop on Automatic Differentiation, NIPS*, 2017.
- [26] J. Redmon, S. Divvala, R. Girshick, and A. Farhadi. You only look once: Unified, real-time object detection. In *CVPR*, pages 779–788, 2016.
- [27] S. Ren, R. Girshick, R. Girshick, and J. Sun. Faster R-CNN: Towards real-time object detection with region proposal networks. *IEEE TPAMI*, 39(6):1137–1149, 2015.
- [28] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio. FitNets: Hints for thin deep nets. In *ICLR*, 2015.
- [29] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, and M. Bernstein. ImageNet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [30] B. Shakibi, B. Shakibi, M. Ranzato, M. Ranzato, and N. D. Freitas. Predicting parameters in deep learning. In *NIPS*, pages 2148–2156, 2013.
- [31] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015.
- [32] J. Tran, J. Tran, J. Tran, and W. J. Dally. Learning both weights and connections for efficient neural networks. In *NIPS*, pages 1135–1143, 2015.
- [33] D. Wang, L. Zhou, X. Zhang, X. Bai, and J. Zhou. Exploring linear relationship in feature map subspace for ConvNets compression. *CoRR*, abs/1803.05729, 2018.
- [34] P. Wang and J. Cheng. Fixed-point factorized networks. In *CVPR*, 2017.
- [35] W. Wen, Y. He, S. Rajbhandari, M. Zhang, W. Wang, F. Liu, B. Hu, Y. Chen, and H. Li. Learning intrinsic sparse structures within long short-term memory. In *ICLR*, pages 1–13, 2018.
- [36] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li. Learning structured sparsity in deep neural networks. In *NIPS*, pages 2082–2090, 2016.
- [37] R. Yu, A. Li, C. F. Chen, J. H. Lai, V. I. Morariu, X. Han, M. Gao, C. Y. Lin, and L. S. Davis. NISP: Pruning networks using neuron importance score propagation. In *CVPR*, pages 2755–2763, 2018.
- [38] X. Yu, T. Liu, X. Wang, and D. Tao. On compressing deep models by low rank and sparse decomposition. In *CVPR*, pages 67–76, 2017.
- [39] S. Zagoruyko and N. Komodakis. Paying more attention to attention: Improving the performance of convolutional neural networks via attention transfer. In *ICLR*, 2016.