



Citation for published version:

Singh, P, Verma, VK, Rai, P & Namboodiri, VP 2019, Play and Prune: Adaptive filter pruning for deep model compression. in S Kraus (ed.), *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019*. IJCAI International Joint Conference on Artificial Intelligence, vol. 2019-August, International Joint Conferences on Artificial Intelligence, pp. 3460-3466, 28th International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, 10/08/19.

Publication date:
2019

Document Version
Peer reviewed version

[Link to publication](#)

University of Bath

Alternative formats

If you require this document in an alternative format, please contact:
openaccess@bath.ac.uk

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Play and Prune: Adaptive Filter Pruning for Deep Model Compression

Pravendra Singh, Vinay Kumar Verma, Piyush Rai and Vinay P. Namboodiri

Department of Computer Science and Engineering, IIT Kanpur, India

{psingh, vkverma, piyush, vinaypn}@cse.iitk.ac.in

Abstract

While convolutional neural networks (CNN) have achieved impressive performance on various classification/recognition tasks, they typically consist of a massive number of parameters. This results in significant memory requirement as well as computational overheads. Consequently, there is a growing need for filter-level pruning approaches for compressing CNN based models that not only reduce the total number of parameters but reduce the overall computation as well. We present a new min-max framework for filter-level pruning of CNNs. Our framework, called Play and Prune (PP), jointly prunes and fine-tunes CNN model parameters, with an adaptive pruning rate, while maintaining the model’s predictive performance. Our framework consists of two modules: (1) An adaptive filter pruning (AFP) module, which minimizes the number of filters in the model; and (2) A pruning rate controller (PRC) module, which maximizes the accuracy during pruning. Moreover, unlike most previous approaches, our approach allows directly specifying the desired error tolerance instead of pruning level. Our compressed models can be deployed at run-time, without requiring any special libraries or hardware. Our approach reduces the number of parameters of VGG-16 by an impressive factor of 17.5X, and number of FLOPS by 6.43X, with no loss of accuracy, significantly outperforming other state-of-the-art filter pruning methods.

1 Introduction

Deep convolutional neural networks (CNN) have been used widely for object recognition and various other computer vision tasks. After the early works based on standard forms of deep convolutional neural networks [LeCun *et al.*, 1998; Krizhevsky *et al.*, 2012], recent works have proposed and investigated various architectural changes [Simonyan and Zisserman, 2015; Chollet, 2017; He *et al.*, 2016] to improve the performance of CNNs. Although these changes, such as adding more layers to the CNN, have led to impressive performance gains, and they have also resulted in a substantial

increase in the number of parameters, as well as the computational cost. The increase in model size and computations have made it impractical to deploy these models on embedded and mobile devices for real-world applications. To address this, recent efforts have focused on several approaches for compressing CNNs, such as using binary or quantized [Rastegari *et al.*, 2016] weights. However, these require special hardware or pruning of unimportant/redundant weights [Han *et al.*, 2015a; Han *et al.*, 2015b; Alvarez and Salzmann, 2016; Wen *et al.*, 2016], and give rather limited speedups.

As most of the CNN parameters reside in the fully connected layers, a high compression rate with respect to the number of network parameters can be achieved by simply pruning redundant neurons from the fully connected layers. However, this does not typically result in any significant reduction in computations (FLOPs based speedup), as most of the computations are performed in convolutional layers. For example, in case of VGG-16, the fully connected layers contain 90% of total parameters but account for only 1% of computations, which means that convolutional layers, despite having only about 10% of the total parameters, are responsible for 99% of computations. This has led to considerable recent interest in convolutional layer filter pruning approaches. However, most existing pruning approaches [Han *et al.*, 2015a; Han *et al.*, 2015b] result in irregular sparsity in the convolutional filters, which requires software specifically designed for sparse tensors to achieve speed-ups in practice [Han *et al.*, 2015a]. In contrast, some other filter pruning approaches [Luo *et al.*, 2017; Li *et al.*, 2017; He *et al.*, 2017] are designed to directly reduce the feature map width by removing specific convolutional filters via ℓ_2 or ℓ_1 regularization on the filters, and effectively reducing the computation, memory, and the number of model parameters. These methods result in models that can be directly used without requiring any sparse libraries or special hardware.

In this work, we propose a novel filter pruning formulation. Our formulation is based on a simple min-max game between two modules to achieve an adaptive maximum pruning with minimal accuracy drop. We show that our approach, dubbed as Play-and-Prune (PP), results in substantially improved performance as compared to other recently proposed filter pruning strategies while being highly stable and efficient to train. We refer to the two modules of our framework as an Adaptive Filter Pruning (AFP) and Pruning Rate Controller

(PRC). The AFP is responsible for pruning the convolutional filter, while the PRC is responsible for maintaining accuracy.

Unlike most previous approaches, our approach does not require any *external* fine-tuning. In each epoch, it performs an adaptive fine-tuning to recover from the accuracy loss caused by the previous epoch’s pruning. Moreover, while previous approaches need to pre-specify a pruning level for each layer, our approach is more flexible in the sense that it directly specifies an error tolerance level and, based on that, decides which filters to prune, and from which layer(s).

Through an extensive set of experiments and ablation studies on several benchmarks, we show that Play-and-Prune provides state-of-the-art filter pruning, and significantly outperforms existing methods.

2 Related Work

Among one of the earliest efforts on compressing CNNs by pruning unimportant/redundant weights, [Rastegari *et al.*, 2016; Chen *et al.*, 2015] includes binarizing/quantizing the network weights. These approaches [Zhou *et al.*, 2016; Han *et al.*, 2015b] learn sparse network weights, where most of the weights are zero, and consequently can lead to very small model sizes. However, despite sparsity of the weights, the pruned models are still not computationally efficient at run-time. Moreover, these models require special library/hardware for sparse matrix multiplication because activation/feature maps are still dense, which hinders practical utility.

In [Li *et al.*, 2017], the authors proposed filter pruning by ranking filters based on their sum of absolute weights. They assumed that if the sum of absolute weights is sufficiently small, the corresponding activation map will be weak. Similarly, [Luo and Wu, 2017] use a different approach to rank the filter importance, based on the entropy measures. The assumption is that high entropy filters are more important. Alternatively, [Hu *et al.*, 2016] use a data-driven approach to calculate filter importance, that is based on the average percentage of zeros in the corresponding activation map. Less important filters have more number of zeros in their activation map. Recently, [Molchanov *et al.*, 2017] proposed improving run time by using a Taylor approximation. This approach estimates the change in cost by pruning the filters. Another work [Luo *et al.*, 2017] uses pruning of filters based on the next layer statistics. Their approach is based on checking the activation map of the next layer to prune the convolution filters from the current layer. In a recent work [He *et al.*, 2017] used a similar approach as in [Luo *et al.*, 2017] but used lasso regression.

3 Proposed Approach

We assume we have a CNN model \mathcal{M} with K convolutional layer. Layer i is denoted as \mathcal{L}_i and consists of n_i filters denoted as $F_{\mathcal{L}_i} = \{f_1, f_2, \dots, f_{n_i}\}$. We assume that the unpruned model \mathcal{M} has the accuracy of \mathcal{E} and, post-pruning, the error tolerance limit is ϵ .

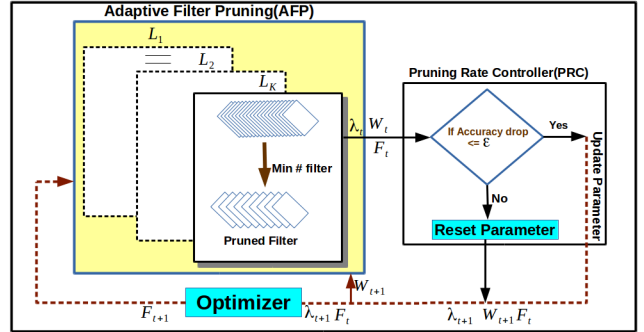


Figure 1: The figure shows the complete architecture. Here AFP minimizes the number of filter in model while PRC maximizes the accuracy during pruning. Here λ_t , \mathbf{W}_t and \mathbf{F}_t are the regularization, weight-threshold and remaining filters in the model respectively at t^{th} pruning iteration.

3.1 Overview

Our deep model compression framework is modeled as a min-max game between two modules, Adaptive Filter Pruning (AFP) and Pruning Rate Controller (PRC). The objective of the AFP is to iteratively minimize the number of filters in the model, while PRC iteratively tries to maximize the accuracy with the set of filters retained by AFP. The AFP will prune the filter only when the accuracy *drop* is within the tolerance limit (ϵ). If accuracy drop is more than ϵ then pruning stops and the PRC tries to recover the accuracy by fine-tuning the model. If PRC’s fine-tuning is unable to bring the error within the tolerance level ϵ , the AFP will not prune the filter from the model and game converges.

Let us denote the AFP by \mathcal{P} and the PRC by \mathcal{C} . Our objective function can be defined as follows:

$$\max_{\#w} \mathcal{C} \left(\min_{\#w = \sum_{i=1}^K n_i} \mathcal{P}(F_{\mathcal{L}_1}, F_{\mathcal{L}_2}, \dots, F_{\mathcal{L}_K}) \right) \quad (1)$$

As shown in the above objective, the AFP (\mathcal{P}) minimizes the number of filters in the network, and the PRC (\mathcal{C}) optimizes the accuracy given that the number of filters. Here $\#w$ is the number of remaining filters after pruning by AFP.

An especially appealing aspect of our approach is that the pruning rates in each iteration are decided *adaptively* based on the performance of the model. After each pruning step, the controller \mathcal{C} checks the accuracy drop (see Fig. 1). If the accuracy drop is more than ϵ , then the pruning rate is reset to zero, and the controller \mathcal{C} tries to recover the system performance (further details of this part are provided in the section on PRC). Eq. 1 converges when $\mathcal{C}(\#w)$ performance drop is more than the tolerance limit, and it is unable to recover it. In such a case, we rollback the current pruning and restore the previous model. At this point, we conclude that this is an optimal model that has the maximal filter pruning within ϵ accuracy drop.

3.2 Convolutional Filter Partitioning

The pruning module \mathcal{P} first needs to identify a candidate set of filters to be pruned. For this, we use a filter partitioning scheme in each epoch. Suppose the entire set of filters of the

model \mathcal{M} is partitioned into two sets, one of which contains the important filters while the other contains the unimportant filters. Let \mathcal{U} and \mathcal{I} be the set of unimportant and important filters, respectively, where

$$\mathcal{M} = \mathcal{U} \cup \mathcal{I} \quad \text{and} \quad \mathcal{U} \cap \mathcal{I} = \emptyset \quad (2)$$

$$\mathcal{U} = \{U_{\mathcal{L}_1}, U_{\mathcal{L}_2}, \dots, U_{\mathcal{L}_K}\} \quad \text{and} \quad \mathcal{I} = \{I_{\mathcal{L}_1}, I_{\mathcal{L}_2}, \dots, I_{\mathcal{L}_K}\}$$

Here $U_{\mathcal{L}_i}$ and $I_{\mathcal{L}_i}$ are set of unimportant and important filters, respectively, in layer \mathcal{L}_i . $U_{\mathcal{L}_i}$, selected as follows:

$$U_{\mathcal{L}_i} = \sigma_{\text{top } \alpha\%}(\text{sort}(\{|f_1|, |f_2|, \dots, |f_{n_i}|\})) \quad (3)$$

Eq. 3 sorts the set in increasing order of $|f_j|$, σ is the select operator and selects the $\alpha\%$ filters with least importance. The remaining filters on \mathcal{L}_i belongs to set $I_{\mathcal{L}_i}$. Here $|f_j|$ is the sum of absolute values of weights in convolutional filter f_j and can be seen as the filter importance. A small sum of absolute values of filter coefficients implies less importance. Our approach to calculate filter importance uses their ℓ_1 norm (Eq. 3), which has been well-analyzed and used in prior works [Li *et al.*, 2017; Ding *et al.*, 2018; He *et al.*, 2018]. Our approach isn't however tied to this criterion, and other criteria can be used, too. We are using this criterion because of its simplicity.

3.3 Weight Threshold Initialization

After obtaining the two sets of filters \mathcal{U} and \mathcal{I} , directly removing \mathcal{U} may result in a sharp and potentially irrecoverable accuracy drop. Therefore we only treat \mathcal{U} as a candidate set of filters to be pruned, of which a *subset* will be pruned eventually. To this end, we optimize the original cost function for the CNN, subject to a group sparse penalty on the set of filters in \mathcal{U} , as shown in Eq. 4. Let $C(\Theta)$ be the original cost function, with Θ being original model parameters. The new objective function can be defined as:

$$\Theta = \arg \min_{\Theta} (C(\Theta) + \lambda_A \|\mathcal{U}\|_1) \quad (4)$$

Here λ_A is the ℓ_1 regularization constant. This optimization penalizes \mathcal{U} such that $|f_j|$ (sum of absolute weights of coefficients in each filter f_j) tends to zero, where $f_j \in U_{\mathcal{L}_i} \forall i \in \{1, 2, \dots, K\}$. This optimization also helps to transfer the information from \mathcal{U} to the rest of the model. If a filter f_j has approximately zero sum of absolute weights then it is deemed safe to be pruned. However, reaching a close-to-zero sum of absolute weights for the whole filter may require several epochs. We therefore choose an *adaptive* weight threshold (W_{γ_i}) for each layer \mathcal{L}_i , such that removing $\forall f_j \in U_{\mathcal{L}_i}$ s.t. $|f_j| \leq W_{\gamma_i}$ results in negligible (close to 0) accuracy drop.

We calculate the initial weight threshold (W_{γ_i}) for \mathcal{L}_i as follows: optimize Eq. 4 for one epoch with $\lambda_A = \lambda$, where λ is the initial regularization constant, which creates two clusters of filters (if we take the sum of the absolute value of filters) as shown in Fig 2. On left cluster (right plot) using the binary search find the maximum threshold W_{γ_i} for \mathcal{L}_i such that accuracy drop is nearly zero.




Figure 2: Histogram of the sum of absolute value of convolutional filters for CONV5.1 in VGG-16 on CIFAR-10. Where left plot is for original filters and the right plot shows the sum of the absolute value of filters after optimization.

3.4 Adaptive Filter Pruning (AFP)

The objective of the AFP module is to minimize the number of filters in the model. Initially, based on the sparsity hyperparameter λ , we calculate the weight thresholds \mathbf{W} . Now instead of using the constant pruning rate, we change the pruning rate adaptively given by the pruning rate controller (PRC; described in the next section) in each epoch. This adaptive strategy helps to discard the filter in a balanced way, such that we can recover from the accuracy drop. In each epoch, from the current model, we select $\alpha\%$ of the filter of lowest importance from each layer, partition them into \mathcal{U} and \mathcal{I} , and perform optimization using Eq. 4, where λ_A is given by PRC. The optimization in Eq. 4 transfers the knowledge of unimportant filters into the rest of the network. Therefore some filter from the \mathcal{U} can be safely discarded. This removal of the filter from the model is done based on the threshold (\mathbf{W}_A) given by the PRC module. Now, from each layer, the filters below the adaptive threshold \mathbf{W}_A are discarded. In each epoch, the weight thresholds and regularization constant is updated dynamically by the PRC module, and a subset of \mathcal{U} is pruned. Hence, in the same epoch, we can recover from the accuracy drop from the previous epoch's pruning, making the model ready to prune filters in the current epoch.

The objective of the AFP module can be defined as:

$$\Theta' = \sigma_{\#w \in \Theta'} \left[\mathcal{P} \left(\arg \min_{\Theta'} \left(C(\Theta') + \lambda_A \sum_{i=1}^K \|\mathcal{U}\|_1 \right) \right) \right] \quad (5)$$

Here Θ' is the collection of remaining filters after pruning, and σ is the select operator. $\#w$ is the collection of all the filter from each layer \mathcal{L}_i that has a sum of absolute value greater than the W_{γ_i} . From Eq.-5, it is clear that it minimizes the number of the filters based on $W_{\gamma_i} \in \mathbf{W}_A, \forall i \in \{1, 2, \dots, K\}$.

3.5 Pruning Rate Controller (PRC)

Let $\mathbf{W} = [W_{\gamma_1}, W_{\gamma_2}, \dots, W_{\gamma_K}]$ denote the initial weight thresholds for the K layers (described in Weight Threshold Initialization section). Now the adaptive thresholds \mathbf{W}_A are calculated as follows:

$$\mathbf{W}_A = \delta_w \times T_r \times \mathbf{W} \quad (6)$$

$$T_r = \begin{cases} \mathcal{C}(\#w) - (\mathcal{E} - \epsilon) & : \mathcal{C}(\#w) - (\mathcal{E} - \epsilon) > 0 \\ 0 & : \text{Otherwise} \end{cases} \quad (7)$$

where $\mathcal{C}(\#w)$ is the accuracy with $\#w$ remaining filters, \mathcal{E} is the accuracy of the unpruned network, and the number $\mathcal{C}(\#w) - (\mathcal{E} - \epsilon)$ denotes how far we are from tolerance error level ϵ . Here, δ_w is a constant used to accelerate or decrease the pruning rate. The regularization constant λ_A in Eq. 4 also adapted based on the model performance after pruning and its updates are given as follows

$$\lambda_A = \begin{cases} (\mathcal{C}(\#w) - (\mathcal{E} - \epsilon)) \times \lambda & : \mathcal{C}(\#w) - (\mathcal{E} - \epsilon) > 0 \\ 0 & : \text{Otherwise} \end{cases} \quad (8)$$

Form Eq. 8 it is clear that we set the regularizer constant to zero if our pruned model performance is below the tolerance limit. Otherwise, it is proportional to the accuracy above the tolerance limit. λ is the initial regularization constant.

The PRC module essentially controls the rate at which the filters will get pruned. In our experiments, we found that if the pruning rate is high, there is a sharp drop in accuracy after pruning, which may or may not be recoverable. Therefore pruning saturates early, and we are unable to get the high pruning rate. Also if the pruning rate is too slow, the model may get pruned very rarely and spends most of its time in fine-tuning. We, therefore, use a pruning strategy that adapts the pruning rate *dynamically*. In the pruning process, if in some epoch, the system performance is below the tolerance limit, we reset the pruning rate to zero. Therefore the optimization will focus only on the accuracy gain until the accuracy is recovered to be again within the tolerance level ϵ . Note that the adaptive pruning rate depends on model performance. When the model performance is within ϵ , the pruning depends on how far we are from ϵ . From Eq 6, it is clear that the \mathbf{W}_A depends on the performance of the system over the $\#w$ filters in the model. In this way, by controlling the pruning rate, we maintain a balance between filter pruning and accuracy. This module tries to maximize accuracy by reducing pruning rate. The objective function of PRC can be defined as:

$$\max_{\Theta'} \mathcal{C}(\Theta', D) \quad (9)$$

Here \mathcal{C} calculates the performance, i.e., accuracy. It is the function of all the convolutional filters Θ' that remain after pruning, and D is the validation set used to compute the model accuracy.

In addition to dynamically controlling the pruning rate, the PRC offers several other benefits, discussed next.

Iterative Pruning Bounds Accuracy Drop

Eq. 7 and Eq. 8 ensure that compressed model will not go beyond error tolerance limit, which is controlled by the PRC. Experimentally we found that, in a non-iterative one round pruning, if model suffers from a high accuracy drop during pruning then pruning saturates early and fine tuning will not recover accuracy drop properly. We have shown an ablation study which shows the effectiveness of iterative pruning over the single round pruning to justify this fact.

Pruning Cost

The cost/effort involved in pruning is mostly neglected in most of the existing filter pruning methods. Moreover, most of the methods perform pruning and fine-tuning separately. In contrast, we jointly prune and fine-tune the CNN model

		Baseline	PP-1	PP-2
Input Size		32×32×3	32×32×3	32×32×3
Layers	CONV1.1	64	18	18
	CONV1.2	64	48	48
	CONV2.1	128	65	65
	CONV2.2	128	65	65
	CONV3.1	256	104	96
	CONV3.2	256	112	112
	CONV3.3	256	114	110
	CONV4.1	512	207	186
	CONV4.2	512	163	79
	CONV4.3	512	79	79
	CONV5.1	512	74	74
	CONV5.2	512	48	48
	CONV5.3	512	60	60
	FC6	512	512	512
FC7	10	10	10	
Total parameters		15.0M	1.13M (13.3×)	0.86M (17.5×)
Model Size		60.0 MB	4.6 MB (13.0×)	3.5 MB (17.1×)
Accuracy		93.49	93.46	93.35
FLOPs		313.7M	54.0M (5.8×)	48.8M (6.43×)

Table 1: Layer-wise pruning results and pruned models (PP-1, and PP-2) statistics for VGG-16 on CIFAR-10.

parameters, with an adaptive pruning rate, while maintaining the model’s predictive performance. Therefore, in a given epoch, we can recover from the accuracy drop suffered due to the previous epoch’s pruning and making the model ready to prune filters in the current epoch.

Layer Importance

Most previous methods [Ding *et al.*, 2018; He *et al.*, 2017; Li *et al.*, 2017; He *et al.*, 2018; Yu *et al.*, 2018] use user-specified desired model compression rate but finding optimal compression rate is not so easy and involves many trials. In CNNs, some layers are relatively less important, and therefore we can prune many more filters from such layers. In contrast, if we prune a large number of filters from important layers, then this might result in an irrecoverable loss in accuracy. Our approach is more flexible since it directly specifies an error tolerance level ϵ and, based on that, adaptively decides which filters to prune, and from which layer(s), using Eq. 6 to determine layer-specific pruning rates.

4 Experiments and Results

To show the effectiveness of the proposed approach, we have conducted extensive experiments on small as well as large datasets, CIFAR-10 [Krizhevsky and Hinton, 2009] and ILSVRC-2012 [Russakovsky *et al.*, 2015]. Our approach yields state-of-art results on VGG-16 [Simonyan and Zisserman, 2015], and RESNET-50 [He *et al.*, 2016] respectively.

For all our experiments, we set $\lambda = 0.0005$ (set initially but later adapted), $\delta_w = 1$ and $\alpha = 10\%$. We follow the same parameter settings and training schedule as [Li *et al.*, 2017; He *et al.*, 2017; He *et al.*, 2018]. We also report an ablation study for various values of α .

4.1 VGG-16 on CIFAR-10

We experimented with the VGG-16 on the CIFAR-10 dataset. We follow the same parameter settings and training schedule as [Li *et al.*, 2017]. Table 1 shows the layer-wise pruning statistics for PP-1 (first pruned model), and PP-2 (second pruned model). We compare our results with the recent works




Figure 3: Our ResNet pruning strategy, where we pruned only first two convolutional layers in each block.

on filter pruning. Our approach consistently performs better as compared to Li-pruned [Li *et al.*, 2017], SBP [Neklyudov *et al.*, 2017], AFP [Ding *et al.*, 2018] as shown in Table 2.

4.2 Ablation Study for VGG-16 on CIFAR-10

This section shows a detailed analysis of the effect on the different component in the proposed approach.

Ablation Study on the Hyper-parameter α

We did an ablation study on the hyper-parameter α , i.e., how many filters are selected for partitioned \mathcal{U} . We experimented with $\alpha = 5, 10, 20, 30\%$. We found that if we take the lower value, it will not degrade the performance only it takes more epochs to prune. While if we take high α (say 30%) value, it starts degrading the performance of model early, hence we are unable to get high pruning rate. In our case, we find $\alpha = 10\%$ is a moderate value and at this rate, we can achieve a high pruning rate. This rate we set across all architecture like ResNet-50, VGG.

Pruning Iterations Vs Error

In Fig. 4 (left) we have shown that if we do the pruning in 1 shot, it has significant accuracy drop (8.03%) as compared to PP-2 on the same FLOPs pruning (84.5%). While if we prune the model iteratively, we have less error rate for the same FLOPs pruning.

Filter Importance Vs Layer Importance

Most of the previous approach [Ding *et al.*, 2018; Li *et al.*, 2017; Yu *et al.*, 2018] focus on the *How to prune* (filter importance/ranking), while it is also important for the better pruning rate to decide *Where to prune*. Our approach also decides the *where to prune* by considering layer importance. In the figure-4 (right) we are showing the ablation on our approach’s capability to decide the layers importance. In figure-4 (right) we are showing the error rate on the similar FLOPs pruning without/with considering layer importance for the four compressed model search (S1,S2,S3,S4) with our approach (PP-2) using same filter importance criteria.

Method	Error(%)	Params Pruned(%)	Pruned FLOPs(%)
Li-pruned [Li <i>et al.</i> , 2017]	6.60	64.0	34.20
SBP [Neklyudov <i>et al.</i> , 2017]	7.50	-	56.52
AFP-E [Ding <i>et al.</i> , 2018]	7.06	93.3	79.69
AFP-F [Ding <i>et al.</i> , 2018]	7.13	93.5	81.39
PP-1 (Ours)	6.54	92.5	82.8
PP-2 (Ours)	6.65	94.3	84.5

Table 2: Comparison of pruning VGG-16 on CIFAR-10 (the baseline accuracy is 93.49%).

Alpha value	Error(%)	Parameters	FLOPs
5	6.54	1.12×10^6	5.3×10^7
10	6.54	1.13×10^6	5.4×10^7
20	6.56	1.15×10^6	5.5×10^7
30	6.61	1.27×10^6	6.1×10^7

Table 3: Ablation study over the α values. Experimentally we found that $\alpha = 10$ is the most suitable.




Figure 4: (a) Left figure shows the effectiveness of iterative pruning (b) Right figure shows effect of layer importance on error for the same FLOPs pruning (84.5%)

4.3 ResNet-56 on CIFAR-10

We follow the same parameter settings and training schedule as [Li *et al.*, 2017; He *et al.*, 2018]. Our approach significantly outperforms various state-of-the-art approaches for ResNet-56 on CIFAR-10. The results are shown in Table 4. We achieve high pruning rate 68.4% with the 6.91% error rate, while AFP-G [Ding *et al.*, 2018] has the error rate of 7.06% with only 60.9% pruning.

4.4 VGG-16 On ILSVRC-2012

To show the effectiveness of our proposed approach, we also experimented with the large-scale dataset ILSVRC-2012 [Russakovsky *et al.*, 2015]. It contains 1000 classes with 1.5 million images. To make our validation set, randomly 10 images (from the training set) are selected from each class. This is used by PRC to calculate validation accuracy drop for adjusting the pruning rate. In this experiment, α is the same as the previous experiment. We follow the same setup and settings as [He *et al.*, 2017].

Our large-scale experiment for VGG-16 [Simonyan and Zisserman, 2015] on the ImageNet [Russakovsky *et al.*, 2015] shows the state-of-art result over the other approaches for model compression. Channel-pruning (CP) [He *et al.*, 2017] has the 80.0% model FLOPs compression with the top-5 ac-

Method	Error(%)	Pruned FLOPs(%)
Li-B [Li <i>et al.</i> , 2017]	6.94	27.6
NISP [Yu <i>et al.</i> , 2018]	6.99	43.6
CP [He <i>et al.</i> , 2017]	8.20	50.0
SFP [He <i>et al.</i> , 2018]	6.65	52.6
AFP-G [Ding <i>et al.</i> , 2018]	7.06	60.9
PP-1 (Ours)	6.91	68.4

Table 4: Comparison of pruning ResNet-56 on CIFAR-10 (the baseline accuracy is 93.1%).

Method	Top-5 Accu.(%)	Pruned FLOPs(%)
RNP (3X)[Lin <i>et al.</i> , 2017a]	87.57	66.67
ThiNet-70 [Luo <i>et al.</i> , 2017]	89.53	69.04
CP [He <i>et al.</i> , 2017]	88.20	80.00
PP-1 (Ours)	89.81	80.20

Table 5: Comparison of pruning VGG-16 on ImageNet (the baseline accuracy is 90.0%).

Method	Top-5 Accu.(%)	Parameters	Pruned FLOPs(%)
ThiNet [Luo <i>et al.</i> , 2017]	90.7	16.94M	36.9
SFP [He <i>et al.</i> , 2018]	92.0	-	41.8
PP-1 (Ours)	92.0	15.7M	44.1
CP [He <i>et al.</i> , 2017]	90.8	~18M	50.0
PP-2 (Ours)	91.4	13.7M	52.2

Table 6: Comparison of pruning ResNet-50 on ImageNet (the baseline accuracy is 92.2%).

curacy 88.2%, while we have same FLOPs pruning (80.2%) with the top-5 accuracy 89.81%. Refer to table-5 for the detail comparison results. Our compressed model (PP-1) is obtained after 38 epochs.

4.5 ResNet-50 ON ILSVRC-2012

In ResNet, there exist restrictions on the few layers due to its identity mapping (skip connection). Since for $output = f(x) + x$ we need to sum two vector, therefore we need x and $f(x)$ should have same dimension. Hence we cannot change the output dimension freely. Hence only two convolutional layers can be pruned for each block (see Fig-3). Unlike the previous work [Luo *et al.*, 2017] where they explicitly set $p=q$, we have not imposed any such restriction which results in more compression with better accuracy. We prune ResNet-50 from block 2a to 5c continuously as described in the proposed approach. If the filter is pruned, then the corresponding channels in the batch-normalization layer and all dependencies to that filter are also removed. We follow the same settings as [He *et al.*, 2017].

Our results on ResNet are shown in Table 6. We are iteratively pruning convolutional filters in each epoch as described earlier. PP-1 is obtained after 34 epochs. Similarly, PP-2 is obtained after 62 epochs. We have experimentally shown that our approach reduces FLOPs and Parameters without any significant drop in accuracy.

4.6 Practical Speedup

The practical speedup is sometimes very different by the result reported in terms of FLOPs prune percentage. The practical speedup depends on the many other factors, for example, intermediate layers bottleneck, availability of data (batch size) and the number of CPU/GPU cores available.

Model	data	Avg. Precision, IoU:		
		0.5:0.95	0.5	0.75
F-RCNN original	trainval35K	30.3	51.3	31.8
F-RCNN pruned	trainval35K	30.3	51.1	31.7

Table 7: Generalization results on MS-COCO. Pruned ResNet-50 (PP-2) used as a base model for Faster-RCNN.




Figure 5: Speedup corresponding to CPU (i7-4770 CPU@3.40GHz) and GPU (GTX-1080) over the different batch size for VGG-16 on CIFAR-10.

For VGG-16 architecture with the 512 batch size, we have 4.02X practical GPU speedup, while the theoretical speedup is 6.43X (figure-5). This gap is very close on the CPU, and our approach gives the 6.24X practical CPU speedup compare to 6.43X theoretical (Fig. 5).

4.7 Generalization Ability

To show the generalization ability of our proposed approach, we also experiment on the object detection architecture. In this experiment we have taken the most popular architecture Faster-RCNN [Ren *et al.*, 2015] on MS-COCO [Lin *et al.*, 2014] dataset.

Compression for Object Detection

The experiments are performed on COCO detection datasets with 80 object categories [Lin *et al.*, 2014]. Here all 80k train images and a 35k val images are used for training (trainval35K) [Lin *et al.*, 2017b]. We are reporting the detection accuracies over the 5k unused val images (minival). In this first, we trained Faster-RCNN with the ImageNet pre-trained ResNet-50 base model. The results are shown in table-7. In this experiment we used our pruned ResNet-50 model (PP-2) as given in Table-6 as a base network in Faster-RCNN. We found that the pruned model shows similar performances in all cases. In the Faster-RCNN implementation, we use ROI Align and use stride 1 for the last block of the convolutional layer (layer4) in the base network.

5 Conclusion

We proposed a Play and Prun Filter Pruning (PP) framework to prune CNNs. Our approach follows a min-max game between two modules (AFP and PRC). Since our approach can prune the entire convolution filter, there is a significant reduction in FLOPs and the number of model parameters. Our approach does not require any special hardware/software support, is generic, and practically usable. We evaluated it on various architectures like VGG and Resnet. Our approach can also be used in conjunction with pruning methods such as binary/quantized weights, weight pruning, etc.

References

- [Alvarez and Salzman, 2016] Jose M Alvarez and Mathieu Salzman. Learning the number of neurons in deep networks. In *NIPS*, pages 2270–2278, 2016.
- [Chen *et al.*, 2015] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, and Yixin Chen. Compressing neural networks with the hashing trick. In *ICML*, pages 2285–2294, 2015.
- [Chollet, 2017] Francois Chollet. Xception: Deep learning with depthwise separable convolutions. In *CVPR*, pages 1251–1258, 2017.
- [Ding *et al.*, 2018] Xiaohan Ding, Guiguang Ding, Jungong Han, and Sheng Tang. Auto-balanced filter pruning for efficient convolutional neural networks. *AAAI*, 2018.
- [Han *et al.*, 2015a] 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*, 2015.
- [Han *et al.*, 2015b] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *NIPS*, pages 1135–1143, 2015.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016.
- [He *et al.*, 2017] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *ICCV*, volume 2, page 6, 2017.
- [He *et al.*, 2018] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *IJCAI*, 2018.
- [Hu *et al.*, 2016] 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.
- [Krizhevsky and Hinton, 2009] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. 2009.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, pages 1097–1105, 2012.
- [LeCun *et al.*, 1998] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, pages 2278–2324, 1998.
- [Li *et al.*, 2017] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *ICLR*, 2017.
- [Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, pages 740–755. Springer, 2014.
- [Lin *et al.*, 2017a] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. Runtime neural pruning. In *Advances in Neural Information Processing Systems*, pages 2181–2191, 2017.
- [Lin *et al.*, 2017b] Tsung-Yi Lin, Piotr Dollár, Ross B Girshick, Kaiming He, Bharath Hariharan, and Serge J Belongie. Feature pyramid networks for object detection. In *CVPR*, page 4, 2017.
- [Luo and Wu, 2017] Jian-Hao Luo and Jianxin Wu. An entropy-based pruning method for cnn compression. *arXiv preprint arXiv:1706.05791*, 2017.
- [Luo *et al.*, 2017] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. *ICCV*, 2017.
- [Molchanov *et al.*, 2017] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *ICLR*, 2017.
- [Neklyudov *et al.*, 2017] Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, and Dmitry P Vetrov. Structured bayesian pruning via log-normal multiplicative noise. In *NIPS*, pages 6778–6787, 2017.
- [Rastegari *et al.*, 2016] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, pages 525–542. Springer, 2016.
- [Ren *et al.*, 2015] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. In *NIPS*, pages 91–99, 2015.
- [Russakovsky *et al.*, 2015] 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. *IJCV*, 115(3):211–252, 2015.
- [Simonyan and Zisserman, 2015] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.
- [Wen *et al.*, 2016] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *NIPS*, pages 2074–2082, 2016.
- [Yu *et al.*, 2018] Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S Davis. Nisp: Pruning networks using neuron importance score propagation. *CVPR*, 2018.
- [Zhou *et al.*, 2016] Hao Zhou, Jose M Alvarez, and Fatih Porikli. Less is more: Towards compact cnns. In *ECCV*, pages 662–677. Springer, 2016.