

“© 2020 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.”

Asymptotic Soft Filter Pruning for Deep Convolutional Neural Networks

Yang He, Xuanyi Dong, Guoliang Kang, Yanwei Fu, Chenggang Yan, and Yi Yang*

Abstract—Deeper and wider Convolutional Neural Networks (CNNs) achieve superior performance but bring expensive computation cost. Accelerating such over-parameterized neural network has received increased attention. A typical pruning algorithm is a three-stage pipeline, i.e., training, pruning, and retraining. Prevailing approaches fix the pruned filters to zero during retraining, and thus significantly reduce the optimization space. Besides, they directly prune a large number of filters at first, which would cause unrecoverable information loss. To solve these problems, we propose an Asymptotic Soft Filter Pruning (ASFP) method to accelerate the inference procedure of the deep neural networks. First, we update the pruned filters during the retraining stage. As a result, the optimization space of the pruned model would not be reduced but be the same as that of the original model. In this way, the model has enough capacity to learn from the training data. Second, we prune the network asymptotically. We prune few filters at first and asymptotically prune more filters during the training procedure. With asymptotic pruning, the information of the training set would be gradually concentrated in the remaining filters, so the subsequent training and pruning process would be stable. Experiments show the effectiveness of our ASFP on image classification benchmarks. Notably, on ILSVRC-2012, our ASFP reduces more than 40% FLOPs on ResNet-50 with only 0.14% top-5 accuracy degradation, which is higher than the soft filter pruning (SFP) by 8%.

Index Terms—Filter Pruning, Image Classification, Neural Networks

I. INTRODUCTION

CONVOLUTIONAL Neural Networks (CNNs) have demonstrated state-of-the-art performance in computer vision tasks [1]–[12]. The superior performance of deep CNNs usually comes from the deeper and wider architectures [1], [13]–[16], which cause the prohibitively expensive computation cost. The storage, memory, and computation of these cumbersome models significantly exceed the computing limitation of current mobile devices or drones. [17] shows that running a 1 billion connection neural network at 20Hz would require 12.8W just for DRAM access. Besides, the authors of [18], [19] claim that VGGNet [13] requires 321.1 MBytes memory and 236 mW power for a batch of three frames and processes 0.7 frame per second even when it is deployed on the optimized energy-efficient chip. Therefore, it is essential

Y. He, X. Dong, G. Kang and Y. Yang are with the Center for AI, University of Technology Sydney, Sydney, NSW 2007, Australia (e-mail: yang.he-1@student.uts.edu.au; xuanyi.dong@student.uts.edu.au; guoliang.kang@student.uts.edu.au; yi.yang@uts.edu.au.).

Y. Fu is with The School of Data Science, Fudan University, Shanghai 200433, China (e-mail: yanweif@fudan.edu.cn.)

C. Yan is with Hangdian University, Hangzhou 310000, China (cgyan@hdu.edu.cn)

*Corresponding Author: Yi Yang (yi.yang@uts.edu.au)

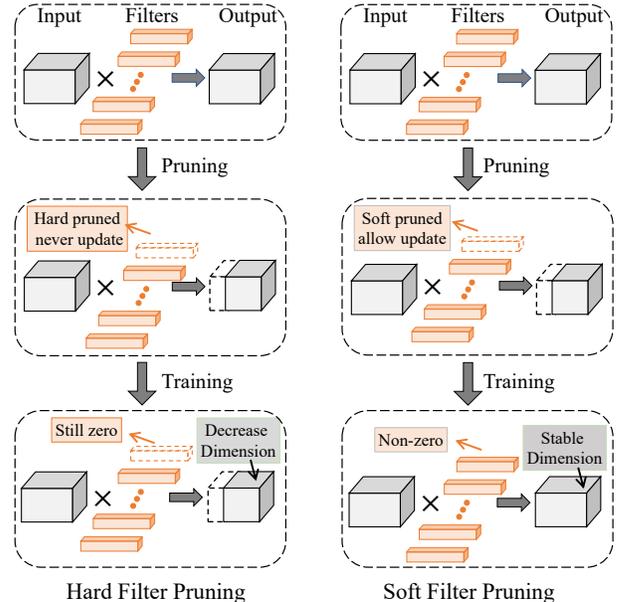


Fig. 1: Hard filter pruning v.s. soft filter pruning. We mark the pruned filter as the orange dashed box. For the hard filter pruning, the pruned filters are always *fixed* during the whole training procedure. Therefore, the model capacity is reduced and thus harms the performance because the dashed blue box is useless during training. On the contrary, our soft pruning method allows the pruned filters to be *updated* during the training procedure. In this way, the model capacity is recovered from the pruned model and thus leads a better accuracy.

to maintain the deep CNN models to have a relatively low computational cost but ensure high accuracy in real-world applications.

Pruning deep CNNs [20]–[22] is an important direction for accelerating the networks. Recent efforts have been made either on directly deleting some weight values of filters [17] (i.e., weight pruning) or totally discarding some filters (i.e., filter pruning) [23]–[25]. However, weight pruning results in the *unstructured sparsity* of filters. Since the unstructured model cannot leverage the existing high-efficiency BLAS (Basic Linear Algebra Subprograms) libraries, weight pruning is not efficient in saving the memory usage and computational cost. In contrast, filter pruning enables the model with *structured sparsity*, taking full advantage of BLAS libraries to achieve more efficient memory usage and more realistic acceleration. Therefore, filter pruning is more favored in accelerating the networks.

Nevertheless, most filter pruning algorithms suffer from two problems: (1) *the model capacity reduction* and (2) *the*

unrecoverable filter information loss. Specifically, as shown in Figure 1, most researchers conduct the “hard filter pruning (HFP)” [23]–[25], then the pruned filters are directly deleted and have no possibility to be recovered. The discarded filters will reduce the optimization space and model capacity, and thus it is unfavorable for the pruned network to learn enough knowledge. To alleviate this problem, they use pre-training to maintain a good performance, which in turn induces much more training time. Furthermore, existing methods directly prune a large number of filters, which contain information of training set, at first. This process leads to severe and unrecoverable information loss and thus inevitably degrades the performance.

To solve the above two problems, we propose ASFP, which prunes the convolutional filters dynamically. Particularly, before the first training epoch, the filters with small ℓ_2 -norm are selected and set to zero. Then we retrain the model, and the previously pruned filters could be updated. Before the next training epoch, we will prune a *new set* of filters with small ℓ_2 -norm. These training processes are continued until converged. Lastly, some filters with smallest ℓ_2 -norm will be selected and pruned without further updating. This soft manner enables the compressed network to have a larger optimization space and model capacity. Hence it is easier for the model to learn from the training data, and achieve higher accuracy even without the pre-training process.

In addition, we prune the network asymptotically — pruning few filters at first, and more filters at later training epochs. If few filters are pruned at first, little pre-trained information would be lost, so it is easy for the model to recover from pruning. With the following iterative training and soft pruning, the training set information would be gradually concentrated in some important filters. At the same time, the training and pruning process would be stable, as the information is lost gradually instead of suddenly.

We highlight the following three contributions of ASFP:

- (1) We propose a soft manner to allow the previous pruned filters to be reconstructed during training. This soft manner could significantly maintain the model capacity, which enables the network to be trained and pruned simultaneously from scratch.
- (2) To avoid severe information loss, we propose to asymptotically prune the filters, which makes the subsequent training and pruning process more stable.
- (3) Experiments on CIFAR-10 and ImageNet demonstrate the effectiveness and efficiency of the proposed ASFP.

II. RELATED WORK

CNN accelerating methods can be roughly divided into four categories, namely, *matrix decomposition*, *low-precision weights*, *weight pruning*, and *filter pruning*.

1) *Matrix Decomposition*: To reduce the computation costs of the convolutional layers, previous work propose to representing the weight matrix of the convolutional network as a low-rank product of two smaller matrices [26]–[31]. Then the calculation of production of one large matrix turns to the production of two smaller matrices. However, the computational cost of tensor decomposition operation is expensive,

which is not friendly to train deep CNNs. Besides, there exists an increasing usage of 1×1 convolution kernel in some recent neural networks, such as the bottleneck block structure of ResNet [16], cases where it is difficult to apply matrix decomposition.

2) *Low Precision*: Some other researchers focus on low-precision implementation to compress and accelerate CNN models [32]–[36]. Zhou *et al.* [33] propose trained ternary quantization to reduce the precision of weights in neural networks to ternary values. The authors of [34] present incremental network quantization, targeting to convert pre-trained full-precision CNN model into a low-precision version efficiently. In this situation, only low-precision weights are stored and used during the inference procedure, with the storage and computation cost being dramatically reduced.

3) *Weight Pruning*: Recent work [17], [32], [37] prunes weights of neural networks. For example, [17] proposed an iterative weight pruning method by discarding the small weights whose values are below the threshold. [38], [39] leveraged the sparsity property of feature maps or weight parameters to accelerate the CNN models. However, weight pruning always leads to unstructured models, so the model cannot leverage the existing efficient BLAS libraries in practice. Therefore, it is difficult for weight pruning to achieve realistic speedup. Meanwhile, Bayesian methods [40] are also applied to network pruning. However, these methods are evaluated on rather small datasets such as MNIST [41] and CIFAR-10 [42].

4) *Filter Pruning*: Pruning the filters [23]–[25], [43] leads to the removal of the corresponding feature maps, thus not only reducing the storage usage on devices but also decreasing the memory footprint consumption. Considering whether to utilize the training data to determine the pruned filters, the filter pruning methods are roughly divided into two categories, data dependent and data independent filter pruning. The latter method is more efficient than the former since training data may not be available during the pruning process.

Data Dependent Filter Pruning. Some approaches [24], [25], [43]–[49] utilize the training data to determine the pruned filters. The authors of [45] minimize the reconstruction error of activation maps to obtain a decomposition of convolutional layers. Luo *et al.* [25] adopt the statistics information from the next layer to guide the importance evaluation of filters.

Data Independent Filter Pruning. Concurrently with our work, some data independent filter pruning strategies [23], [50]–[52] have been explored. Li *et al.* [23] explore the sensitivity of layers for filter pruning and utilize a ℓ_1 -norm criterion to prune unimportant filters. Ye *et al.* [51] prune models by enforcing sparsity on the scaling parameters of batch normalization layers. However, for all these filter pruning methods, the representative capacity of the neural network after pruning is seriously affected by smaller optimization space. Besides, the information loss at the beginning is significant and unrecoverable.

III. METHODOLOGY

A. Preliminary

We formally introduce the symbol and notations in this section. The deep CNN network can be parameterized by

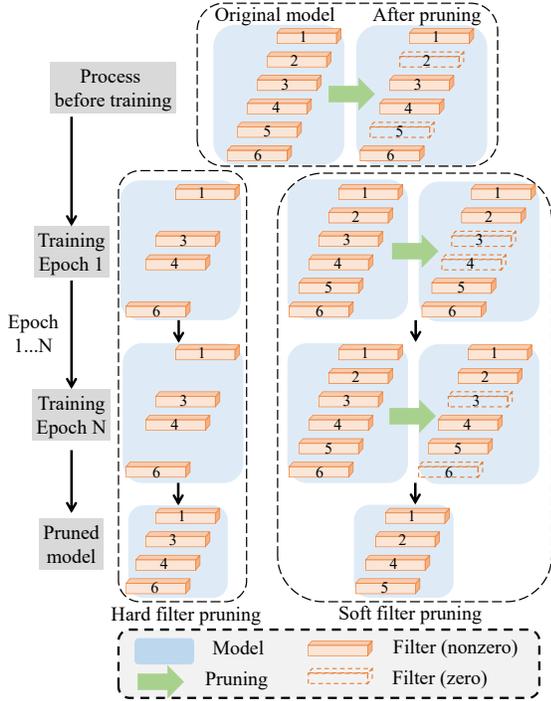


Fig. 2: Pruning and training schedule of HFP and SFP. Before training, we first select some filters with pre-defined importance evaluations. HFP directly deletes these filters before training, while for SFP, those are set to zero and kept. During training (epoch 1 to N), the model size is smaller than the original one for HFP. While for SFP, the zero value filters (filter 2 and 5) become non-zero after training epoch 1. Then we evaluate the importance of filters again and prune filter 3 and 4. The model size would not be the same as the original one. When training is finished, the final pruned model is the model at epoch N for HFP. While for SFP, we delete the zero value filters (filter 3 and 6) at epoch N to get the final pruned model.

$\{\mathbf{W}^{(i)} \in \mathbb{R}^{N_{i+1} \times N_i \times K \times K}, 1 \leq i \leq L\}$ ¹. $\mathbf{W}^{(i)}$ denotes a matrix of connection weights in the i_{th} layer. N_i denotes the number of input channels for the i_{th} convolution layer. L denotes the number of layers. The shapes of input tensor \mathbf{U} and output tensor \mathbf{V} are $N_i \times H_i \times W_i$ and $N_{i+1} \times H_{i+1} \times W_{i+1}$, respectively. The convolutional operation of the i_{th} layer can be written as:

$$\mathbf{V}_{i,j} = \mathcal{F}_{i,j} * \mathbf{U}, \text{ for } 1 \leq j \leq N_{i+1}, \quad (1)$$

where $\mathcal{F}_{i,j} \in \mathbb{R}^{N_i \times K \times K}$ represents the j_{th} filter of the i_{th} layer, and $\mathbf{V}_{i,j}$ represents the j_{th} output feature map of the i_{th} layer. $\mathbf{W}^{(i)}$ consists of $\{\mathcal{F}_{i,j}, 1 \leq j \leq N_{i+1}\}$.

Pruning filters can remove the output feature maps. In this way, the computational cost of the neural network will reduce remarkably. Let us assume the pruning rate is P_i for the i_{th} layer. The number of filters of this layer will be reduced from N_{i+1} to $N_{i+1}(1 - P_i)$, thereby the size of the output tensor $\mathbf{V}_{i,j}$ can be reduced to $N_{i+1}(1 - P_i) \times H_{i+1} \times W_{i+1}$.

B. Pruning with Hard Manner

Given a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^n$ and a desired sparsity level κ (i.e., the number of remaining non-zero filters), HFP

can be formulated as:

$$\begin{aligned} \min_{\mathcal{F}} \ell(\mathcal{F}; \mathcal{D}) &= \min_{\mathcal{F}} \frac{1}{n} \sum_{i=1}^n \ell(\mathcal{F}; (\mathbf{x}_i, \mathbf{y}_i)), \\ \text{s.t. } \mathcal{F} &\in \mathbb{R}^{N \times K \times K}, \quad N(\mathcal{F}) \leq \kappa. \end{aligned} \quad (2)$$

Here, $\ell(\cdot)$ is the standard loss function (e.g., cross-entropy loss), \mathcal{F} is the set of filters of the neural network, and N is the cardinality of the filter set. Typically, HFP firstly prunes filters of a single layer of a pre-trained model and fine-tune the pruned model to complement the performance degradation. Then they prune the next layer and fine-tune the model again until the last layer is pruned. Once the filters are pruned, HFP will not update these filters again.

The full training schedule of HFP is shown in the first column of Figure 2, and the detailed pruning process of HFP is shown in the first row of Figure 3. First, some of the filters with small ℓ_p -norm (filter 2 and 4 in HFP-a, marked in blue) are selected and pruned. After retraining, these pruned filters are not able to be updated again thus the ℓ_p -norm of these filters would be zero during all the training epochs (filter 2 and 4 in HFP-b). Meanwhile, the remaining filters (filter 1, 3 and N in HFP-c, marked in green) might be updated to another value after retraining to make up for the performance degradation due to pruning. After several epochs of retraining to converge the model, a compact model is obtained to accelerate the inference.

C. Pruning with Soft Manner

SFP can dynamically remove the filters in a soft manner. Specifically, the key is to keep updating the pruned filters in the training stage. The full training schedule of SFP is shown in the second column of Figure 2. Such an updating manner brings several benefits. It not only keeps the model capacity of the pruned models as the original models but also avoids the greedy layer by layer pruning procedure and enables pruning *all* convolutional layers at the same time. For SFP, the constrain in the Eq. 2 changes to:

$$\|\mathcal{F}\|_0 \leq \kappa, \quad N(\mathcal{F}) = N_{i+1}. \quad (3)$$

Here, $\|\cdot\|_0$ is the standard L_0 norm. After soft pruning, the number of filters $N(\mathcal{F})$ is still the same as that of the original model (N_{i+1}).

The second row of Figure 3 explains the detailed process of SFP. First, the ℓ_2 -norms of all filters are computed for each weighted layer and used as our filter selection criterion. Second, some filters with a small ℓ_p -norm (filter 2 and 4 in SFP-a, marked in blue) are selected, and we prune those filters by setting the corresponding filter weights as zero (filter 2 and 4 in SFP-b). Then we retrain the model. As we allow the pruned filters to be updated during retraining, the pruned filters become nonzero again (filter 2 and 4 in SFP-c) due to back-propagation. After iterations of pruning and reconstruction to converge the model, we delete the unimportant filters and obtain a compact and efficient model.

¹Fully-connected layers can be viewed as convolutional layers with $k = 1$

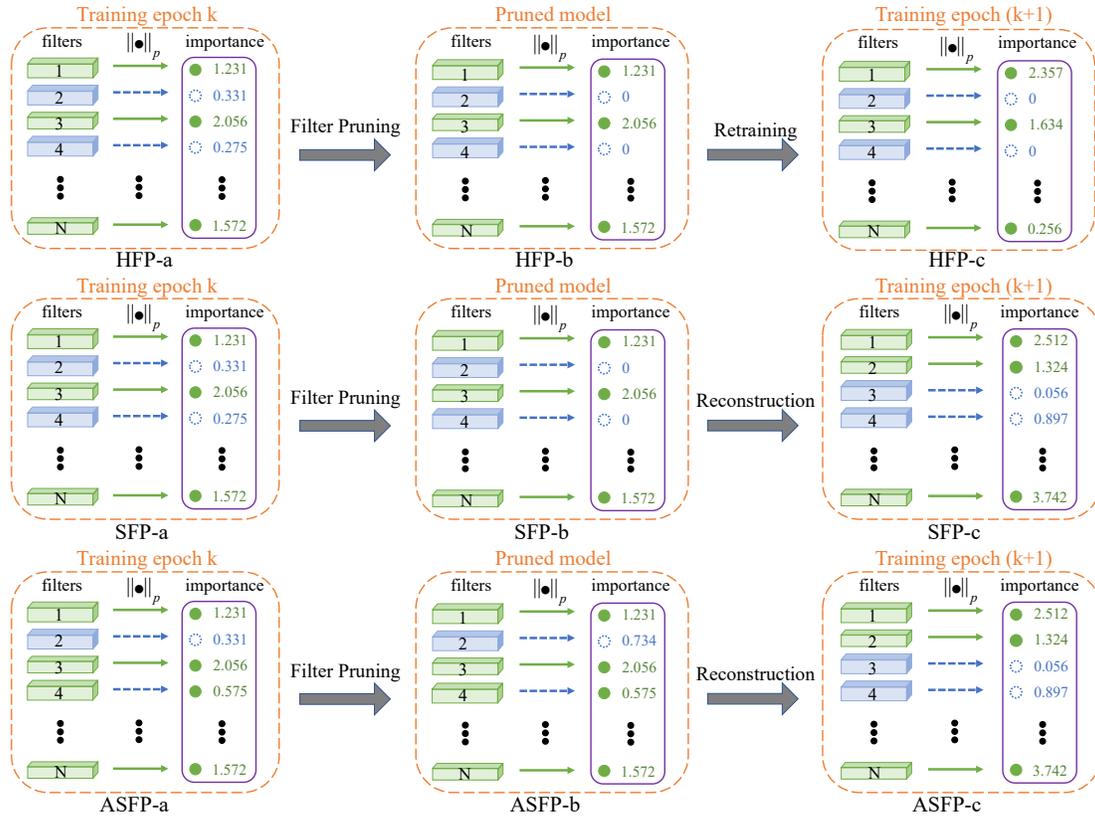


Fig. 3: Overview of HFP (first row), SFP (second row) and ASFP (third row). (a): filter instantiations before pruning. (b): filter instantiations after pruning. (c): filter instantiations after reconstruction. The filters are ranked by their ℓ_p -norms and the small ones (purple rectangles) are selected to be pruned.

D. Asymptotic Soft Filter Pruning (ASFP)

It is known that all the filters of the pre-trained model have the information of the training set. Therefore, pruning those informative filters would cause information loss. This situation is especially difficult for pruning the models that pre-trained on large datasets, or pruning a large number of filters of small models, as the lost information is rather massive. Therefore, we propose to pruning the neural network asymptotically.

Specifically, we use a small pruning rate at early epochs, and gradually increase the pruning rate later, until we reach the goal pruning rate P_i at the last retraining epoch. The optimization problem of ASFP:

$$\begin{aligned} \min_{\mathcal{F}} \ell(\mathcal{F}; \mathcal{D}) &= \min_{\mathcal{F}} \frac{1}{n} \sum_{i=1}^n \ell(\mathcal{F}; (\mathbf{x}_i, \mathbf{y}_i)), & (4) \\ \text{s.t. } \|\mathcal{F}\|_0 &\leq \kappa_{epoch}, \quad N(\mathcal{F}) = N_{i+1}. \end{aligned}$$

Here, κ_{epoch} means the sparsity level changes with the training epoch. Comparing Eq. 3 and Eq. 4, the difference between SFP and ASFP is whether the pruning rate is changing with epochs. For SFP, the κ in Eq. 3 is a pre-defined constant regarding the number of remaining non-zero filters, so it would not change during the training process. In contrast, κ_{epoch} in Eq. 4 is a function of $epoch$ and would change with the epoch number of training. Therefore, SFP is a special case of ASFP when the pruning rate P_i during all training epochs are the same. The details of ASFP is illustratively explained in Algorithm 1.

1) *Asymptotic Filter Selection*: We use the ℓ_p -norm to evaluate the importance of each filter as Eq. 5. In general, the convolutional results of the filters with the smaller ℓ_p -norm would lead to relatively lower activation values, and thus have a less numerical impact on the final prediction of deep CNN models. In term of this understanding, such filters of small ℓ_p -norm will be given higher priority of being pruned than those of higher ℓ_p -norm:

$$\|\mathcal{F}_{i,j}\|_p = \sqrt[p]{\sum_{n=1}^{N_i} \sum_{k_1=1}^K \sum_{k_2=1}^K |\mathcal{F}_{i,j}(n, k_1, k_2)|^p}. \quad (5)$$

In practice, we use the ℓ_2 -norm based on the empirical analysis.

Different from SFP [50] that the pruning rate equals the goal pruning rate P_i^{goal} during retraining, we use different pruning rate P'_i at every epoch. The definition of P'_i is list as follows:

$$P'_i = \mathcal{H}(P_i^{goal}, D, P_i^{min}, epoch), \quad (6)$$

where P_i^{goal} represents the goal pruning rate for the i_{th} layer, D and P_i^{min} are pre-defined parameter which will be explicitly explained later. As exponential parameter decay is widely used in optimization [53] to achieve a stable result, we change the pruning rate exponentially. The equation is as follows:

$$P'_i = a \times e^{-k \times epoch} + b, \quad (7)$$

In order to solve three parameters a, k, b of the above exponential equation, three points consists of (epoch, P'_i) pair are

Algorithm 1: Algorithm Description of ASFP

Input : training data \mathbf{X} and training epoch $epoch_{max}$;
 the model with parameters $\mathbf{W} = \{\mathbf{W}^{(i)}, 0 \leq i \leq L\}$;
 pruning rate P_i^{goal} , P_i^{min} and pruning rate decay D ;
Output: The compact model and its parameters \mathbf{W}^*

```

1 for  $epoch \leftarrow 1$  to  $epoch_{max}$  by 1 do
2   Update the model parameter  $\mathbf{W}$  based on  $\mathbf{X}$ ;
3   for  $i \leftarrow 1$  to  $L$  by 1 do
4     Calculate  $P_i'$  with Eq. 6;
5     Get the number of filters in  $i$  layer  $N_{i+1}$  with  $\mathbf{W}$ ;
6     Zeroize  $N_{i+1}P_i'$  filters with small  $\ell_2$ -norm;
7   end
8 end
9 Obtain the compact model with parameters  $\mathbf{W}^*$  from  $\mathbf{W}$ 
    
```

needed. Certainly, the first point is $(0, P_i^{min})$, which means the pruning rate is P_i^{min} for the first training epoch. In addition, to achieve the goal pruning rate P_i at the final retraining epoch, the point $(epoch_{max}, P_i^{goal})$ is essential. Now we have to define the third point $(epoch_{max} \times D, 3P_i^{goal}/4)$, to solve the equation. This means that when the epoch number is $epoch_{max} \times D$, the pruning rate increase to $3/4$ of the goal pruning rate P_i^{goal} .

2) *Filter Pruning*: We set the value of selected $N_{i+1}P_i'$ filters to zero (see the filter pruning step in Figure 3). This can temporarily eliminate their contribution to the network output. We prune *all* the weighted layers at the same time. In this way, we can prune each filter in parallel, which would cost negligible computation time. In contrast, previous methods [24], [25] always conduct a greedy layer by layer pruning and retraining, which would cost more computation time, especially when the model depth increases. Moreover, we use the *same* pruning rate for *all* the weighted layers, $P_i^{goal} = P$. Therefore, we only require one hyper-parameter P to balance acceleration and accuracy. This can avoid the inconvenient hyper-parameter search or the complicated sensitivity analysis shown in [23].

3) *Reconstruction*: After the pruning step, we train the network for one epoch to reconstruct the pruned filters. In order to keep the representative capacity and the high performance of the model, those pruned filters are updated to non-zero by back-propagation, as shown in Figure 3. In this way, the pruned model has the same capacity as the original model during the training. In contrast, hard filter pruning leads to a decrease of feature maps, so the model capacity is reduced and the performance is influenced. With large model capacity, we could integrate the pruning step into the normal training schema, that is, training and pruning the model synchronously.

4) *Obtaining Compact Model*: ASFP iterates over the filter selection, filter pruning, and reconstruction steps. After model convergence, we can obtain a sparse model containing many “zero filters”. The features maps corresponding to those “zero filters” will always be zero during the inference procedure. Therefore, there will be no influence to remove these filters as well as the corresponding feature maps.

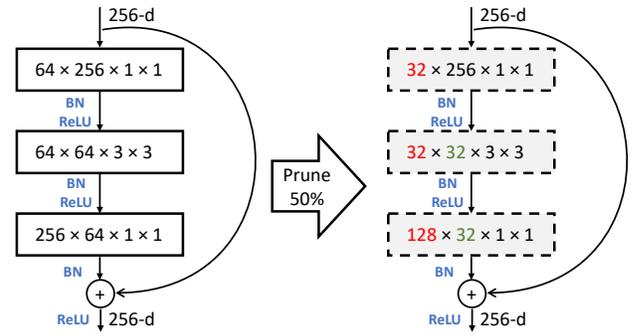


Fig. 4: Pruning residual block with pruning rate 50%. Red and green number means the remaining output channel number and input channel number after pruning, respectively. “BN” and “ReLU” represents the batch norm layer and non-linear layer, respectively.

E. Pruning Strategy for Convolutional Network

Pruning the traditional convolutional architectures [1], [13] is easy to understand, but it is elusive for pruning recent structural variants, such as ResNet [16]. In Figure 4, pruning the residual blocks is illustrated. If we set the pruning rate as 50%, the output dimension of three layers would decrease by 50% (the red numbers). The input dimension of the last two layers would change according to the output dimension of the previous layer (the green numbers). What we should especially care about is the element-wise additive of the residual connections and the output of convolutional layers. In Figure 4, the channel number of the residual connections and the output of convolutional layers is 256 and 128, respectively. The element-wise additive should change to:

$$\mathbf{O}_{index} = \begin{cases} \mathbf{R}_{index} + \mathbf{C}_{index} & \text{if } index \in \mathbf{I} \\ \mathbf{R}_{index} & \text{else} \end{cases} \quad (8)$$

where \mathbf{I} is the index of 128 remaining channels, \mathbf{C} is the convolutional output after the batch norm layer (128 channels), \mathbf{R} is the residual output (256 channels), and \mathbf{O} is the output of the whole residual block.

F. Computation Complexity Analysis

1) *Theoretical Speedup Analysis*: Suppose the filter pruning rate of the i_{th} layer is P_i , which means the $N_{i+1} \times P_i$ filters are set to zero and pruned from the layer, and the other $N_{i+1} \times (1 - P_i)$ filters remain unchanged. Besides, suppose the size of the input and output feature map of i_{th} layer is $H_i \times W_i$ and $H_{i+1} \times W_{i+1}$. Then after filter pruning, the dimension of useful output feature map of the i_{th} layer decreases from $N_{i+1} \times H_{i+1} \times W_{i+1}$ to $N_{i+1}(1 - P_i) \times H_{i+1} \times W_{i+1}$. Note that the output of i_{th} layer is the input of $(i + 1)_{th}$ layer. And we further prunes the $(i + 1)_{th}$ layer with a filter pruning rate P_{i+1} , then the calculation of $(i + 1)_{th}$ layer is decrease from $N_{i+2} \times N_{i+1} \times k^2 \times H_{i+2} \times W_{i+2}$ to $N_{i+2}(1 - P_{i+1}) \times N_{i+1}(1 - P_i) \times k^2 \times H_{i+2} \times W_{i+2}$. In other words, a proportion of $1 - (1 - P_{i+1}) \times (1 - P_i)$ of the original calculation is reduced, which will make the inference procedure of CNN models much faster.

2) *Realistic Speedup Analysis*: In theoretical speedup analysis, other operations such as batch normalization and pooling are negligible compared to the convolution operations. Therefore, we consider the FLOPs (Floating-point operations per second) of convolution operations for computation complexity comparison, which is commonly used in the previous work [23], [25]. In the real scenario, reduced FLOPs cannot bring the same level of realistic speedup because non-tensor layers (*e.g.*, batch normalization and pooling layers) also need the inference time on GPU [25]. In addition, the limitation of IO delay, buffer switch, and efficiency of BLAS libraries also lead to the wide gap between theoretical and realistic speedup ratio. We compare the theoretical and realistic speedup in Section IV-D.

IV. EXPERIMENT

A. Benchmark Datasets and Experimental Setting

Dataset. Our method is evaluated on two benchmarks: CIFAR-10 [42] and ILSVRC-2012 [55]. CIFAR-10 contains 50,000 training images and 10,000 test images, which are categorized into ten classes. ILSVRC-2012 is a large-scale dataset containing 1.28 million training images and 50k validation images of 1,000 classes.

Architecture. As discussed in [24], [25], [54], multiple-branch ResNet [16] is less redundant than VGGNet [13], so it is more difficult to accelerate ResNet. Therefore, we focus on pruning the challenging ResNet model. To validate our method on the single-branch network, we also prune the VGGNet following [23].

Training setting. In the CIFAR-10 experiments, we use the default parameter setting in [56] and follow the training schedule in [57]. For CIFAR-10 dataset, we test our ASFP on ResNet-56 and 110. On ILSVRC-2012, we follow the same parameter settings as [16], [56]. The data augmentation strategies are the same as PyTorch implementation [58]. We test ASFP on ResNet-18, 34, 50 and we use the pruning rate 30% for all the models. We also analyze the difference between pruning the pre-trained model and scratch model. For pruning the model from scratch, We use the normal training schedule without additional fine-tuning process. For pruning the pre-trained model, we reduce the learning rate to one-tenth of the original learning rate. To conduct a fair comparison of pruning from scratch and pre-trained models, we use the same training epochs to train/fine-tune the network. The previous work [23] uses fewer epochs to fine-tune the pruned model, but it converges too early and harms the accuracy, as shown in section IV-B.

Pruning setting. For VGGNet on CIFAR-10, we use the same pruning rate as [23]. For experiments on ResNet, we follow [50] and prune *all* the convolutional layers with the *same* pruning rate at the same time. We do not prune the projection shortcuts for simplification, which only need negligible time and do not affect the overall cost. Therefore, only one hyper-parameter, the pruning rate $P_i = P$ is used to balance acceleration and accuracy.

To asymptotically change the pruning rate, we set the parameters in Eq. 6 as $D = 1/8$ and $P_i^{min} = 0$. This setting is

denoted as ASFP-P0.² For example, if we use the goal pruning rate $P_i^{goal} = 30\%$, then the pruning rate curve according to epoch is shown in Figure 5. If we set $P_i^{min} = P_i^{goal}$, ASFP is same as SFP. The pruning operation is conducted at the end of every training epoch. We run some experiments three times and report the “mean \pm std”. The performance is compared with other state-of-the-art acceleration algorithms, *e.g.*, MIL [54], PFEC [23], CP [24], ThiNet [25], SFP [50], NISP [46]. We choose to directly cite the numbers from original papers for a fair comparison.

Explanation for Tables. The results of ResNet are listed in Table I and Table III. In “Pre-train?” column, “Y” and “N” indicate whether to use the pre-trained model as initialization or not. The “Accu. Drop” is the accuracy of the pruned model minus that of the baseline model, so negative number means the accelerated model has higher accuracy than the baseline model. A smaller number of “Accu. Drop” is better. For CIFAR-10, we run every experiment for three times to get the mean and standard deviation of the accuracy. For Imagenet, we just list the one-view accuracy.

Explanation for Baselines. The baseline network is the same for different pruning methods, and accuracy numbers are cited from the original paper. The different accuracies are due to different hyper-parameter settings (*e.g.* different data augmentations, different learning rate schedules, *etc.*) and different implementation frameworks (*e.g.* Caffe, TensorFlow and Pytorch). For example, in ThiNet [25], all the images are resized into 256×256 , then center-cropped to 224×224 . However, in CP [24], the images are resized such that the shorter side equals to 256. In this case, we choose to use the “Accu. Drop” (in Table I) rather than the “Accelerated Accu.” (in Table I) to fairly evaluate the effectiveness of our method.

Optimization Time of ASFP. In this paper, we care more about the acceleration during the inference time rather than the training time. However, we would like to show that the additional time cost of ASFP is negligible compared to a hard pruning method. Take the scratch model for an example. HFP and ASFP both need 200 epochs to training CIFAR-10 from scratch to converge, so the additional time cost of ASFP comes from the operation of pruning (Eq. 4). Two steps are needed in such a process. 1) Obtaining the pruning rate P'_i . 2) Ranking and pruning the filters. For step one, the exponent pattern of the pruning rate is pre-defined, and it could be directly accessed during training. For step two, after we get the importance scores of the filters, we zeroize the filters with smaller importance scores to conduct the pruning operation. All these steps bring minor computation cost. For HFP, it takes about 171.01s for training ResNet-110 for one epoch on GTX 1080. For ASFP, the total time for one epoch is 171.45s. Therefore, the time difference between soft pruning and hard pruning is negligible.

B. VGGNet on CIFAR-10

The result of pruning from scratch and pre-trained VGGNet is shown in Table II. Not surprisingly, ASFP achieves better

²For the following sections, the “ASFP” (without suffix) means “ASFP-P0” if not particularly indicated.

TABLE I: Overall performance of pruning ResNet on CIFAR-10.

Depth	Method	Pre-train?	Baseline Accu. (%)	Accelerated Accu. (%)	Accu. Drop (%)	FLOPs	Pruned FLOPs(%)
56	PFEC [23]	✗	93.04	91.31	1.75	9.09E7	27.6
	CP [24]	✗	92.80	90.90	1.90	-	50.0
	SFP [50]	✗	93.59 (±0.58)	92.26 (±0.31)	1.33	5.94E7	52.6
	ASFP (40%)	✗	93.59 (±0.58)	92.44 (±0.07)	1.15	5.94E7	52.6
	PFEC [23]	✓	93.04	93.06	-0.02	9.09E7	27.6
	CP [24]	✓	92.80	91.80	1.00	-	50.0
	SFP [50]	✓	93.59 (±0.58)	93.35 (±0.31)	0.24	5.94E7	52.6
	ASFP (40%)	✓	93.59 (±0.58)	93.12 (±0.20)	0.47	5.94E7	52.6
110	PFEC [23]	✗	93.53	92.94	0.61	1.55E8	38.6
	MIL [54]	✗	93.63	93.44	0.19	-	34.2
	SFP [50]	✗	93.68 (±0.32)	92.62 (±0.60)	1.04	1.21E8	52.3
	ASFP (20%)	✗	93.68 (±0.32)	93.94 (±0.56)	-0.24	1.82E8	28.2
	ASFP (40%)	✗	93.68 (±0.32)	93.20 (±0.10)	0.48	1.21E8	52.3
	PFEC [23]	✓	93.53	93.30	0.20	1.55E8	38.6
	SFP [50]	✓	93.68 (±0.32)	93.86 (±0.21)	-0.18	1.50E8	40.8
	SFP [50]	✓	93.68 (±0.32)	92.90 (±0.18)	0.78	1.21E8	52.3
	ASFP (30%)	✓	93.68 (±0.32)	93.37 (±0.12)	0.31	1.50E8	40.8
	ASFP (40%)	✓	93.68 (±0.32)	93.10 (±0.06)	0.58	1.21E8	52.3

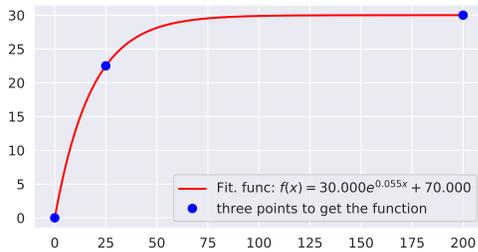


Fig. 5: Asymptotically changed pruning rate when the goal pruning rate is 30%. Three blue points are the three pairs to generate the exponential function of pruning rate (the solid curve).

performance than [23] in both settings. With our pruning criterion, we could achieve slightly better accuracy than [23] when pruning the random initialized VGGNet (93.37% vs. 93.31%). In addition, The pruned model without fine-tuning has better performance than [23] (81.66% vs. 77.45%). After fine-tuning 40 epochs, our model achieves similar accuracy with [23]. Notably, if more fine-tuning epochs (160) are utilized, the accuracy of [23] is almost unchanged (93.28% vs. 93.22%), which means their models have no much more capacity to learn. On the contrary, our method could achieve much better performance (94.02% vs. 93.28%) with more fine-tuning epochs, which shows the model capacity of our ASFP is much larger than [23].

TABLE II: Pruning from scratch and pre-trained VGGNet on CIFAR-10. “FT” means “fine-tuning” the pruned model.

Setting \ Acc (%)	PFEC [23]	Ours
Baseline	93.58 (±0.03)	93.58 (±0.03)
Prune from scratch	93.31 (±0.03)	93.37 (±0.08)
Prune from pre-train without FT	77.45 (±0.03)	81.66 (±0.03)
FT 40 epochs	93.22 (±0.03)	93.27 (±0.08)
FT 160 epochs	93.28 (±0.03)	94.02 (±0.15)

C. ResNet on CIFAR-10

Table I shows the results on CIFAR-10. Our ASFP could achieve a better performance than other state-of-the-art hard

filter pruning methods. For example, PFEC [23] accelerate ResNet-110 by 38.6% speedup ratio with 0.61% accuracy drop when pruning the scratch models. In contrast, our ASFP can accelerate ResNet-110 to 52.3% speed-up with only 0.48% accuracy drop. When pruning the pre-trained ResNet-110, the accuracy drop of our ASFP is smaller than PFEC [23] when pruning the same number of ratio. When pruning the scratch ResNet-56, we can achieve more acceleration ratio than CP [24] (52.6% vs. 50.0%) with less accuracy drop (1.15% vs. 1.90%) Notably, we can even improve 0.24% accuracy when pruning 28.2% FLOPs of scratch ResNet-56.

When the pruning rate is small, we find the performance of SFP [50] and ASFP is competitive. But ASFP outperforms SFP when a large portion of FLOPs are pruned. The comprehensive comparison is shown in Figure 6. This is because ASFP is suitable for the situation when a large quantity of the information is removed by pruning. These results validate the effectiveness of our ASFP algorithm, which can produce a more compressed model with comparable performance to the original model.

D. ResNet on ILSVRC-2012

Result Explanation. Table III shows that ASFP outperforms other state-of-the-art methods. For pruning from a random initialized ResNet-18, our ASFP achieves more inference speedup than MIL [54] (41.8% v.s. 34.6%), but the top-5 accuracy drop of our pruned model is less than that of their model (1.62% v.s. 2.30%). For pruning pre-trained ResNet-34, our ASFP achieves a much better acceleration than PFEC [23] (41.1% v.s. 24.2%) with comparable accuracy drop. For pre-trained ResNet-50, SFP leads to 8.27% top-5 accuracy drop for 41.8% speedup, but our ASFP could achieve negligible top-5 accuracy drop (0.14%) with the same speedup ratio. The maintained model capacity and asymptotic pruning of ASFP are the main reasons for the improved accuracy and efficiency.

Realistic Acceleration. In order to test the realistic speedup ratio, we measure the forward time of the pruned models on one GTX 1080 Ti GPU with a batch size of 64. The results are shown in Table IV. The gap between theoretical and realistic speed may come from non-tensor layers and the

TABLE III: Overall performance of pruning ResNet on ImageNet.

Depth	Method	Pre-train?	Top-1 Accu.	Top-1 Accu.	Top-5 Accu.	Top-5 Accu.	Top-1 Accu.	Top-5 Accu.	Pruned FLOPs(%)
			Baseline(%)	Accelerated(%)	Baseline(%)	Accelerated(%)	Drop(%)	Drop(%)	
18	MIL [54]	✗	69.98	66.33	89.24	86.94	3.65	2.30	34.6
	SFP [50]	✗	70.23 (± 0.06)	67.25 (± 0.13)	89.51 (± 0.10)	87.76 (± 0.06)	2.98	1.75	41.8
	ASFP (30%)	✗	70.23 (± 0.06)	67.41	89.51 (± 0.10)	87.89	2.82	1.62	41.8
	SFP [50]	✓	70.23 (± 0.06)	60.79	89.51 (± 0.10)	83.11	9.44	6.40	41.8
	ASFP (30%)	✓	70.23 (± 0.06)	68.02	89.51 (± 0.10)	88.19	2.21	1.32	41.8
34	SFP [50]	✗	73.92	71.83	91.62	90.33	2.09	1.29	41.1
	ASFP (30%)	✗	73.92	71.72	91.62	90.65	2.20	0.97	41.1
	PFEC [23]	✓	73.23	72.17	-	-	1.06	-	24.2
	SFP [50]	✓	73.92	72.29	91.62	90.90	1.63	0.72	41.1
	ASFP (30%)	✓	73.92	72.53	91.62	91.04	1.39	0.58	41.1
50	SFP [50]	✗	76.15	74.61	92.87	92.06	1.54	0.81	41.8
	ASFP (30%)	✗	76.15	74.88	92.87	92.39	1.27	0.48	41.8
	CP [24]	✓	-	-	92.20	90.80	-	1.40	50.0
	ThiNet [25]	✓	72.88	72.04	91.14	90.67	0.84	0.47	36.7
	NISP [46]	✓	-	-	-	-	-	0.89	44.0
	SFP [50]	✓	76.15	62.14	92.87	84.60	14.01	8.27	41.8
ASFP (30%)	✓	76.15	75.53	92.87	92.73	0.62	0.14	41.8	

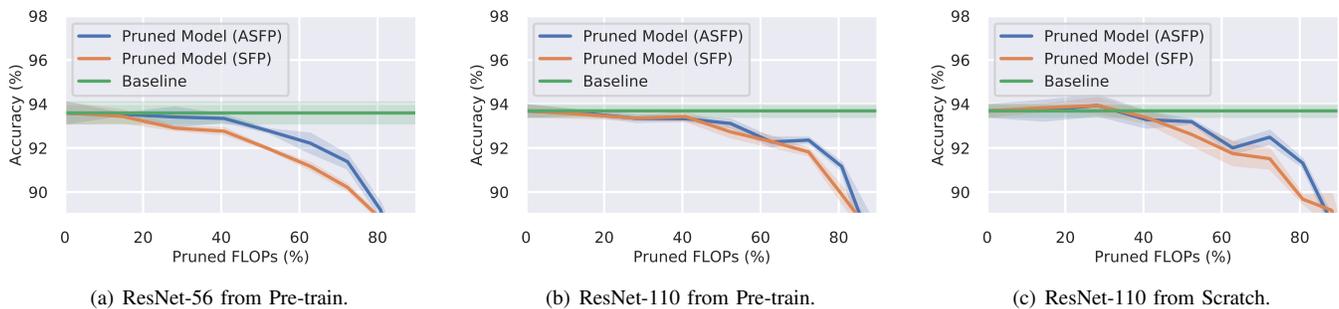


Fig. 6: Model performance regarding different ratio of pruned FLOPs. The green line indicates the model without pruning. The blue and the orange lines represent the model under ASFP and SFP, respectively.

TABLE IV: Comparison of the theoretical and realistic speedup. We only count the time consumption of the forward procedure.

Model	Baseline time (ms)	Pruned time (ms)	Realistic Speed-up(%)	Theoretical Speed-up(%)
ResNet-18	37.10	26.97	27.4	41.8
ResNet-34	63.97	45.14	29.4	41.1
ResNet-50	135.01	94.66	29.8	41.8

limitation of IO delay, buffer switch and the efficiency of BLAS libraries [54].

E. Comparing SFP and ASFP

Performance Regrading Ratio of Pruned FLOPs. In Figure 6, we test the accuracy of ResNet-56 and ResNet-110 under different ratios of pruned FLOPs. For pruning pre-trained initialization, as shown in Figure 6(a) and Figure 6(b), ASFP could obtain better performance than SFP on almost all ratio of pruned FLOPs. Even for pruning models with the random initialization, as shown in Figure 6(c), our method could still outperform SFP. All the results verify that ASFP provides a more effective way to reduce the information loss and thus improves the network performance.

Stable Training Process of ASFP. The model accuracies during training for SFP and ASFP are shown in Figure 7. We run this comparison experiment on ResNet-18 and ResNet-50, and the pruning rate is 30%. We find that the performance gap is not stable for SFP during almost all the 100 retraining

epochs. On the contrary, the performance gap of ASFP is much more stable than that of SFP. For SFP, directly pruning a large number of filters leads to severe information loss. It is difficult for the network to recover from this, and it leads to an unstable training process. In contrast, ASFP would result in a small amount of information loss, consequently increasing the stability of the pruning process.

F. Ablation Study

Extensive ablation study is also conducted to further analyze each component of our model.

1) *Filter Selection Criteria:* The magnitude based criteria such as ℓ_p -norm are widely used to filter selection because computational resources cost is small [23]. We compare the ℓ_2 -norm and ℓ_1 -norm, and the results are shown in Table V. We find that the performance of ℓ_2 -norm criteria are slightly better than that of ℓ_1 -norm criteria. The result of ℓ_2 -norm is dominated by the largest element, while the result of ℓ_1 -norm is also largely affected by other small elements. Therefore, filters with some large weights would be preserved by the ℓ_2 -norm criteria. Consequently, the corresponding discriminative features are kept so the accuracy of the pruned model is better.

2) *Varying Pruned FLOPs:* We evaluate the accuracy of different pruned FLOPs for ResNet-110, and show the results in Figure 6. When the ratio of pruned FLOPs is less than 40%, ASFP and SFP achieve similar accuracy. However, when the

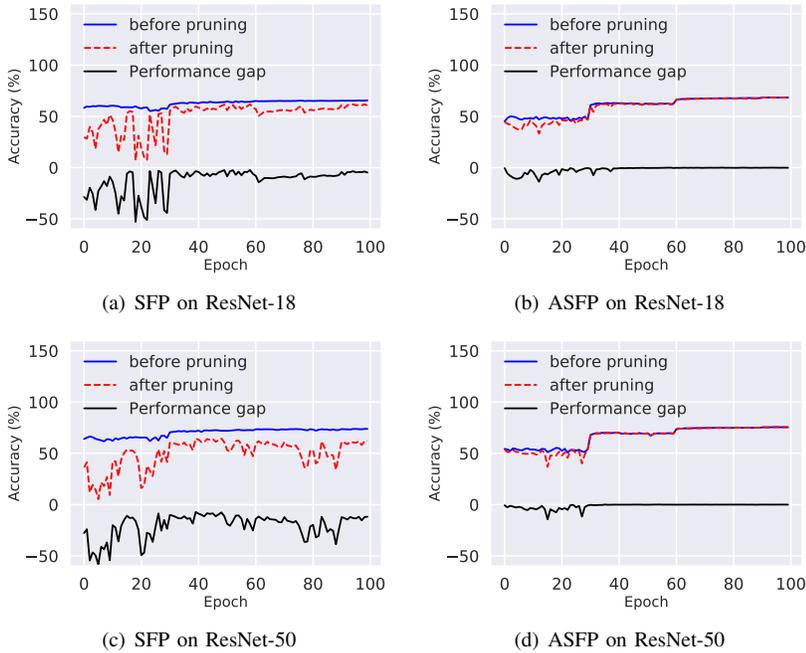


Fig. 7: The training process of ResNet-18 and ResNet-50 and on ImageNet regarding SFP and ASFP. The solid blue line and red dashed line indicate the accuracy of the model before and after pruning, respectively. The black line is the performance gap due to pruning, which is calculated by the accuracy after pruning subtracting that before pruning.

TABLE V: Accuracy of CIFAR-10 on ResNet-110 under different pruning rate with different filter selection criteria.

Pruning rate(%)	10	20	30
ℓ_1 -norm	93.68 ± 0.60	93.68 ± 0.76	93.34 ± 0.12
ℓ_2 -norm	93.89 ± 0.19	93.93 ± 0.41	93.38 ± 0.30

ratio of pruned FLOPs is more than 40%, ASFP could obtain much better performance than SFP. This is because pruning a large number of filters leads to severe information lose and ASFP is especially effective for this case. In contrast, when only a small portion of the information is lost, the maintained model capacity of SFP is enough for good results. For the pruning rate between 0% and about 23%, the accuracy of the accelerated model is higher than the baseline model. This shows that our ASFP and SFP both have a regularization effect on the neural network.

3) *Selection of the Pruned Layers*: Previous work always prunes a portion of the layers of the network. Besides, different layers always have different pruning rates. For example, [23] only prunes insensitive layers, [25] skips the last layer of every block of the ResNet, and [25] prunes more aggressively for shallower layers and prune less for deep layers.

Similarly, we compare the performance of pruning the first and second layer of all basic blocks of ResNet-110. We set the pruning rate as 30%. The model with all the first layers of blocks pruned has an accuracy of $93.96 \pm 0.13\%$, while the model with all the second layers of blocks pruned has an accuracy of $93.38 \pm 0.44\%$. If we carefully select the pruned layers based on the sensitivity, performance improvement may be potentially obtained. However, tuning these hyper-parameters is not the focus of this manuscript.

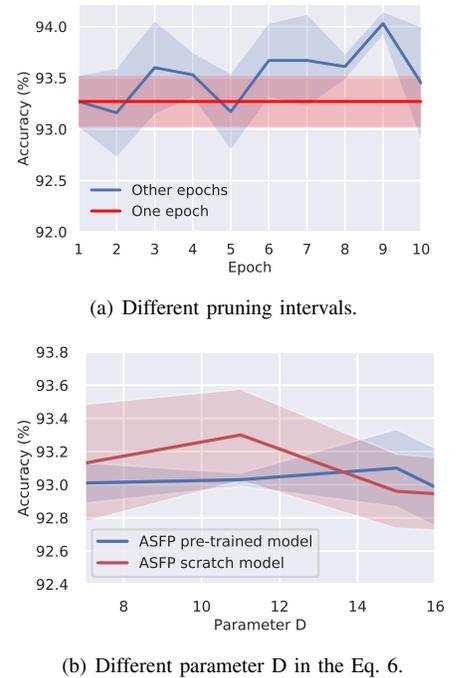


Fig. 8: Ablation study of ASFP. (Solid line and shadow denote the mean and standard deviation of three experiments, respectively.)

4) *Sensitivity of the ASFP Interval*: By default, we conduct our ASFP operation at the end of every training epoch, we call the ASFP interval equals one under this setting. However, different ASFP intervals may lead to a different performance, so we explore the sensitivity of ASFP interval. We use ResNet-110 under a pruning rate of 30% as a baseline, and change the ASFP interval from one epoch to ten epochs. The result is shown in Figure 8(a). We find the model accuracy of most (80%) intervals surpasses the accuracy of one epoch interval. Therefore, we can even achieve better performance if we fine-tune this parameter.

5) *Sensitivity of Parameter D of ASFP*: We change the parameter D in the Eq. 6 to comprehensively understand ASFP, and the results are shown in the Figure 8(b). We prune scratch and pre-trained ResNet-56 on CIFAR-10 and set the pruning rate as 40%. When changing the parameter D from 7 to 16, we find the model accuracy has no large fluctuation ($< 0.3\%$). This shows that the final result of pruning is not sensitive to the parameter D.

V. CONCLUSION & FUTURE WORK

In this paper, we propose an asymptotic soft filter pruning approach (ASFP) to accelerate the deep CNNs. As the training procedures go, we allow the pruned filters to be updated and asymptotically adjust the pruning rate. The soft manner could maintain the model capacity, and the asymptotic pruning could make the pruning process more stable. Therefore, our ASFP could achieve superior performance. Remarkably, without using the pre-trained model, our ASFP can achieve competitive performance compared to the state-of-the-art approaches.

Moreover, by leveraging the pre-trained model, our ASFP achieves better results.

Although changing the pruning rate asymptotically would be beneficial, it has several limitations. First, several additional hyper-parameters are necessary to define the pruning rate during training. Second, the schedule of the pruning rate is hand-crafted and may not be the best schedule. Third, the theoretical demonstration of the information loss of the network after pruning need to be investigated. Furthermore, ASFP can be combined with other acceleration algorithms, e.g., matrix decomposition and low-precision weights, to further improve the performance. We will explore these directions in our future work.

REFERENCES

- [1] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "ImageNet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2012, pp. 1097–1105.
- [2] Y. Yang, Z. Ma, Y. Yang, F. Nie, and H. T. Shen, "Multitask spectral clustering by exploring intertask correlation," *IEEE Trans. Cybern.*, vol. 45, no. 5, pp. 1083–1094, 2015.
- [3] S. Ren, K. He, R. Girshick, and J. Sun, "Faster r-cnn: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 91–99.
- [4] X. Chang, Z. Ma, Y. Yang, Z. Zeng, and A. G. Hauptmann, "Bi-level semantic representation analysis for multimedia event detection," *IEEE Trans. Cybern.*, vol. 47, no. 5, pp. 1180–1197, 2017.
- [5] X. Dong and Y. Yang, "Searching for a robust neural architecture in four gpu hours," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 1761–1770.
- [6] M. Luo, X. Chang, L. Nie, Y. Yang, A. G. Hauptmann, and Q. Zheng, "An adaptive semisupervised feature analysis for video semantic recognition," *IEEE Trans. Cybern.*, vol. 48, no. 2, pp. 648–660, 2018.
- [7] B. Du, W. Xiong, J. Wu, L. Zhang, L. Zhang, and D. Tao, "Stacked convolutional denoising auto-encoders for feature representation," *IEEE Trans. Cybern.*, vol. 47, no. 4, pp. 1017–1027, 2017.
- [8] Y. Wei, Y. Zhao, C. Lu, S. Wei, L. Liu, Z. Zhu, and S. Yan, "Cross-modal retrieval with cnn visual features: A new baseline," *IEEE Trans. Cybern.*, vol. 47, no. 2, pp. 449–460, 2017.
- [9] L. Zhang and P. N. Suganthan, "Visual tracking with convolutional random vector functional link network," *IEEE Trans. Cybern.*, vol. 47, no. 10, pp. 3243–3253, 2017.
- [10] Y. Wu, Y. Lin, X. Dong, Y. Yan, W. Bian, and Y. Yang, "Progressive learning for person re-identification with one example," *IEEE Trans. Image Process.*, vol. 28, no. 6, pp. 2872–2881, June 2019.
- [11] X. Dong, Y. Yan, M. Tan, Y. Yang, and I. W. Tsang, "Late fusion via subspace search with consistency preservation," *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 518–528, Jan 2019.
- [12] G. Kang, J. Li, and D. Tao, "Shakeout: A new approach to regularized deep neural network training," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 5, pp. 1245–1258, 2017.
- [13] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," in *Proc. Int. Conf. Learn. Represent.*, 2015.
- [14] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1–9.
- [15] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna, "Rethinking the inception architecture for computer vision," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2818–2826.
- [16] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 770–778.
- [17] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Proc. Adv. Neural Inf. Process. Syst.*, 2015, pp. 1135–1143.
- [18] Y.-H. Chen, T. Krishna, J. S. Emer, and V. Sze, "Eyeriss: An energy-efficient reconfigurable accelerator for deep convolutional neural networks," *IEEE J. Solid-State Circuits*, vol. 52, no. 1, pp. 127–138, 2017.
- [19] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proc. IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [20] B. Hassibi and D. G. Stork, "Second order derivatives for network pruning: Optimal brain surgeon," in *Proc. Adv. Neural Inf. Process. Syst.*, 1993, pp. 164–171.
- [21] R. Reed, "Pruning algorithms—a survey," *IEEE Trans. Neural Netw.*, vol. 4, no. 5, pp. 740–747, 1993.
- [22] L. Zeng and X. Tian, "Accelerating convolutional neural networks by removing interspatial and interkernel redundancies," *IEEE Trans. Cybern.*, 2018.
- [23] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, "Pruning filters for efficient convnets," in *Proc. Int. Conf. Learn. Represent.*, 2017.
- [24] Y. He, X. Zhang, and J. Sun, "Channel pruning for accelerating very deep neural networks," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 1398–1406.
- [25] J.-H. Luo, J. Wu, and W. Lin, "ThiNet: A filter level pruning method for deep neural network compression," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 5068–5076.
- [26] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," in *Proc. Brit. Mach. Vis. Conf.*, 2014.
- [27] X. Zhang, J. Zou, K. He, and J. Sun, "Accelerating very deep convolutional networks for classification and detection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 10, pp. 1943–1955, 2016.
- [28] X. Zhang, J. Zou, X. Ming, K. He, and J. Sun, "Efficient and accurate approximations of nonlinear convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2015, pp. 1984–1992.
- [29] C. Tai, T. Xiao, Y. Zhang, X. Wang et al., "Convolutional neural networks with low-rank regularization," in *Proc. Int. Conf. Learn. Represent.*, 2016.
- [30] J. Park, S. Li, W. Wen, P. T. P. Tang, H. Li, Y. Chen, and P. Dubey, "Faster cnns with direct sparse convolutions and guided pruning," in *Proc. Int. Conf. Learn. Represent.*, 2017.
- [31] J.-T. Chien and Y.-T. Bao, "Tensor-factorized neural networks," *IEEE Trans. Neural Netw. Learn. Syst.*, vol. 29, no. 5, pp. 1998–2011, 2018.
- [32] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding," in *Proc. Int. Conf. Learn. Represent.*, 2015.
- [33] C. Zhu, S. Han, H. Mao, and W. J. Dally, "Trained ternary quantization," *Proc. Int. Conf. Learn. Represent.*, 2017.
- [34] A. Zhou, A. Yao, Y. Guo, L. Xu, and Y. Chen, "Incremental network quantization: Towards lossless cnns with low-precision weights," *Proc. Int. Conf. Learn. Represent.*, 2017.
- [35] I. Hubara, M. Courbariaux, D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 4107–4115.
- [36] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "XNOR-Net: Imagenet classification using binary convolutional neural networks," in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 525–542.
- [37] Y. Guo, A. Yao, and Y. Chen, "Dynamic network surgery for efficient DNNs," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 1379–1387.
- [38] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2016, pp. 2074–2082.
- [39] V. Lebedev and V. Lempitsky, "Fast convnets using group-wise brain damage," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2016, pp. 2554–2564.
- [40] K. Ullrich, E. Meeds, and M. Welling, "Soft weight-sharing for neural network compression," in *Proc. Int. Conf. Learn. Represent.*, 2017.
- [41] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proc. of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [42] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.
- [43] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang, "Learning efficient convolutional networks through network slimming," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2017, pp. 2755–2763.
- [44] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz, "Pruning convolutional neural networks for resource efficient transfer learning," in *Proc. Int. Conf. Learn. Represent.*, 2017.
- [45] A. Dubey, M. Chatterjee, and N. Ahuja, "Coreset-based neural network compression," in *Proc. Eur. Conf. Comput. Vis.*, 2018.
- [46] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2018.
- [47] Z. Zhuang, M. Tan, B. Zhuang, J. Liu, Y. Guo, Q. Wu, J. Huang, and J. Zhu, "Discrimination-aware channel pruning for deep neural networks," in *Proc. Adv. Neural Inf. Process. Syst.*, 2018, pp. 883–894.

- [48] Y. He and S. Han, “ADC: Automated deep compression and acceleration with reinforcement learning,” in *Eur. Conf. Comput. Vis.*, 2018.
- [49] Q. Huang, K. Zhou, S. You, and U. Neumann, “Learning to prune filters in convolutional neural networks,” in *Proc. IEEE Winter Conf. Appl. Comput. Vis.*, 2018, pp. 709–718.
- [50] Y. He, G. Kang, X. Dong, Y. Fu, and Y. Yang, “Soft filter pruning for accelerating deep convolutional neural networks,” in *Proc. Int. Joint Conf. Artif. Intell.*, 2018, pp. 2234–2240.
- [51] J. Ye, X. Lu, Z. Lin, and J. Z. Wang, “Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers,” in *Proc. Int. Conf. Learn. Represent.*, 2018.
- [52] Y. He, P. Liu, Z. Wang, Z. Hu, and Y. Yang, “Filter pruning via geometric median for deep convolutional neural networks acceleration,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2019, pp. 4340–4349.
- [53] D. P. Kingma and J. Ba, “Adam: A method for stochastic optimization,” in *Proc. Int. Conf. Learn. Represent.*, 2015.
- [54] X. Dong, J. Huang, Y. Yang, and S. Yan, “More is less: A more complicated network with less inference complexity,” in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, 2017, pp. 1895–1903.
- [55] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein *et al.*, “ImageNet large scale visual recognition challenge,” *Int. J. Comput. Vis.*, vol. 115, no. 3, pp. 211–252, 2015.
- [56] K. He, X. Zhang, S. Ren, and J. Sun, “Identity mappings in deep residual networks,” in *Proc. Eur. Conf. Comput. Vis.*, 2016, pp. 630–645.
- [57] S. Zagoruyko and N. Komodakis, “Wide residual networks,” in *Proc. Brit. Mach. Vis. Conf.*, 2016, pp. 87.1–87.12.
- [58] 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 *Proc. Adv. Neural Inf. Process. Syst. Workshops*, 2017.



Yanwei Fu received the PhD degree from Queen Mary University of London in 2014, and the MEng degree in the Department of Computer Science & Technology at Nanjing University in 2011, China. He worked as a Post-doc in Disney Research at Pittsburgh from 2015-2016. He is currently an Assistant Professor at Fudan University. His research interest is image and video understanding, and life-long learning.



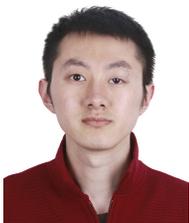
Chenggang Yan received the B.S. degree from Shandong University, Jinan, China, in 2008, and the Ph.D. degree from the Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China, in 2013, both in computer science. He is currently a Professor with Hangzhou Dianzi University, Hangzhou, China. Before that, he was an Assistant Research Fellow with Tsinghua University. His research interests include machine learning, image processing, computational biology, and computational photography. He has authored or co-authored more than 30 refereed journal and conference papers. As a co-author, Prof. Yan was the recipient of the Best Paper Awards in the International Conference on Game Theory for Networks 2014, and SPIE/COS Photonics Asia Conference 2014, and the Best Paper Candidate in the International Conference on Multimedia and Expo 2011.



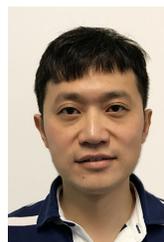
Yang He received the BS degree and MSc from University of Science and Technology of China, Hefei, China, in 2014 and 2017, respectively. He is currently working toward the PhD degree with Center of AI, Faculty of Engineering and Information Technology, University of Technology Sydney. His research interests include deep learning, computer vision, and machine learning.



Xuanyi Dong received the B.E. degree in Computer Science and Technology from Beihang University, Beijing, China, in 2016. He is currently a Ph.D. student in the Center for Artificial Intelligence, University of Technology Sydney, Australia, under the supervision of Prof. Yi Yang.



Guoliang Kang received the BS degree in automation from Chongqing University, Chongqing, China, in 2011 and the MSc degree in pattern recognition and intelligent system from Beihang University, Beijing, China, in 2014. He is currently working toward the PhD degree with Center of AI, Faculty of Engineering and Information Technology, University of Technology Sydney. His research interests include deep learning, computer vision, and statistical machine learning.



Yi Yang received the Ph.D. degree in computer science from Zhejiang University, Hangzhou, China, in 2010. He is currently a professor with University of Technology Sydney, Australia. He was a Post-Doctoral Research with the School of Computer Science, Carnegie Mellon University, Pittsburgh, PA, USA. His current research interest include machine learning and its applications to multimedia content analysis and computer vision, such as multimedia indexing and retrieval, surveillance video analysis and video semantics understanding.