

Convolutional Neural Network Pruning with Structural Redundancy Reduction

Zi Wang¹, Chengcheng Li¹, Xiangyang Wang²

¹The University of Tennessee, Knoxville, TN, USA

²Sun Yat-sen University, Guangzhou, China

{zwang84, cli42}@vols.utk.edu, mcswxxy@mail.sysu.edu.cn

Abstract

Convolutional neural network (CNN) pruning has become one of the most successful network compression approaches in recent years. Existing works on network pruning usually focus on removing the least important filters in the network to achieve compact architectures. In this study, we claim that identifying structural redundancy plays a more essential role than finding unimportant filters, theoretically and empirically. We first statistically model the network pruning problem in a redundancy reduction perspective and find that pruning in the layer(s) with the most structural redundancy outperforms pruning the least important filters across all layers. Based on this finding, we then propose a network pruning approach that identifies structural redundancy of a CNN and prunes filters in the selected layer(s) with the most redundancy. Experiments on various benchmark network architectures and datasets show that our proposed approach significantly outperforms the previous state-of-the-art.

1. Introduction

Convolutional neural networks (CNNs) [22] have developed substantially in recent years and are widely used in various applications, such as object classification [2, 21], image synthesis, [8, 42], super-resolution [4], and game-playing [34, 40]. State-of-the-art performance are achieved by designing wider and deeper CNNs [41, 13, 18]. However, the over-parameterization problem of CNNs prevents them from being applied to resource-limited devices, such as mobile phones and robotics [39, 32]. Many approaches have been proposed to reduce the computation and storage cost of CNNs, such as quantization [10], matrix decomposition [48], network pruning [11, 24, 46, 44, 15], and knowledge distillation [16]. Network pruning is one of the most popular methods and attracts enormous attention.

Generally, network pruning can be categorized into weight (unstructured) pruning [11] and channel (structured) pruning [24, 35, 43, 46]. Weight pruning zeros out spe-

cific weights in filters and results in unstructured sparsities. To accelerate the pruned CNNs, specialized hardware and software have to be developed [9]. Channel pruning, which removes the whole convolutional filters, is a more flexible method without the need for special hardware. As the entire filters are deleted, a considerable pruning ratio can usually be achieved with little performance degradation. Many of the existing channel pruning approaches rely on finding and pruning the least important filters, or the filters that share the most similarities with others across all layers [24, 35, 15, 3]. For example, [35] uses the Taylor series to estimate the loss change after each filter’s removal and prune the filters that cause minimal training loss change. It has been a common belief that with a better filter ranking criterion, there is a better chance to drop the least important filters and get a compact network with less performance loss.

However, our studies on channel pruning contradict this common belief. Using statistical modeling to measure the redundancy in each convolutional layer, we theoretically show that (in certain cases, even randomly) pruning filters in the layer with the most redundancy outperforms pruning the least important filters across all layers. To our best knowledge, this is the first study that theoretically analyzes the rationale behind network pruning from a redundancy reduction perspective. With this finding, we propose a layer-adaptive channel pruning approach based on structural redundancy reduction (SRR), which is achieved by establishing a graph for each convolutional layer of a CNN and using two quantities associated with the graph, i.e., ℓ -covering number and quotient space size, as the measurement of the redundancy in each layer. After that, unimportant filters in the identified layer(s) with the most redundancy, rather than the least important filters across all layers, are pruned.

We summarize the contribution of this study as follows. (1) We theoretically analyze network pruning with statistical modeling from a perspective of redundancy reduction. We find that pruning in the layer(s) with the most redundancy outperforms pruning the least important filters across all layers. (2) We propose a layer-adaptive channel pruning approach based on structural redundancy reduction, which

builds a graph for each convolutional layer of a CNN to measure the redundancy existed in each layer. This approach prunes unimportant filters in the most redundant layer(s), rather than the filters with the least importance across all layers. (3) We validate the proposed approach on various network architectures and datasets. Experiment results demonstrate that our approach achieves state-of-the-art performance compared with recent channel pruning methods. More specifically, our pruned ResNet50 model on ImageNet can reduce 44.1% FLOPs while losing only 0.37% top-1 accuracy.

2. Related work

2.1. Early works and weight pruning

Network pruning is a long-standing topic that can be traced back to the 1990s [12, 23]. In the era of deep learning, [11] is one of the most famous early works that prunes weights below a threshold. After that, various weight pruning approaches have been proposed [1, 28, 50]. As mentioned before, weight pruning causes unstructured sparsities in a network, which is difficult to be used without specialized software and hardware [9].

2.2. Channel pruning

Channel pruning [24, 35, 15, 37] removes the entire filters in a network so that there is no need for specialized hardware. Among all channel pruning approaches, identifying and pruning the least important filters is one of the most popular branches, and can be further divided into three categories. (1) Ranking and pruning filters with a certain criterion. [24] and [38] prune the filters with small weight magnitudes or activation values in the corresponding feature maps. [17] uses the average percentage of zero (APoZ) activation neurons as the criterion and deletes the filters with small APoZ. First and second-order Taylor expansion are used to estimate the loss change after each filter’s removal and the filters that cause minimal loss change are removed [35, 49]. HRank [25] leverages the information in the feature maps to rank the filters. (2) Reconstruction error minimization. Thinet [31] and NISP [47] prune the filters whose removal leads to minimal reconstruction error of the next layer. (3) Similarity measurement. These approaches use various strategies, such as geometric median [15] and clustering [51, 6], to identify the most replaceable filters, or those functionally share the most similarity with others.

2.3. Pruning as network structure optimization

Recently, a number of empirical studies indicate that the network structure after pruning, rather than the removal of unimportant filters, plays a decisive role in maintaining the performance of a network. [29] trained several compact networks obtained by pruning approaches but with random

initialization. Surprisingly, comparable or even better performance can be achieved compared with fine-tuning the pruned models. [33] reports that a network’s performance can be recovered even after random pruning. Related to these works, we also find that pruning unimportant filters is not always essential. But beyond that, we theoretically show that pruning in the layers with large redundancy outperforms pruning the least important filters and propose to prune a network based on structural redundancy reduction.

3. A theoretic analysis of network pruning

We statistically formulate the channel pruning problem from a redundancy reduction perspective. In our context, layer redundancy refers to the number of filters in a convolutional layer. We will later show that the redundancy can be measured with other quantities in real applications. Suppose we have a two-layer CNN¹ with m and n filters, where $n \gg m$. Let $\{\xi_1, \xi_2, \dots, \xi_m\}$ and $\{\eta_1, \eta_2, \dots, \eta_n\}$ be one dimensional positive random variables (RVs) representing each filter’s contribution to the network performance. For example, a filter’s contribution can be represented as the absolute value of *training accuracy drop* or *training loss change* after pruning that filter. We call the two layers ξ layer and η layer for convenience. We first highlight our finding and then prove it from a statistical modeling perspective.

Claim: If a layer has much higher redundancy, pruning filters in that layer, either randomly or selectively, outperforms pruning the least important filters across all layers.

We choose positive constants $a, b > 0$, and use the random events $(\sum_{i=1}^m \xi_i \geq a)$ and $(\sum_{i=1}^n \eta_i \geq b)$ to describe the layers ξ and η “performing well”. Then the performance of a system (i.e., the whole neural network) p is measured by the sum of probabilities of the two events (see Equation (1)). We define one system (p_1) to perform better than another (p_2) if $p_1 > p_2$. A natural question is, if we prune a filter from the network, i.e., remove one variable from $\{\xi_1, \xi_2, \dots, \xi_m, \eta_1, \eta_2, \dots, \eta_n\}$, how does the system performance change? There are the following cases (the performances of the systems are listed in Equations (1)-(5)): (1) no pruning; (2) randomly pruning a filter in the η layer, without loss of generality, we assume the last one η_n is pruned; (3) pruning the least important filter $\eta = \min\{\eta_1, \dots, \eta_n\}$ in the η layer; (4) pruning the least important filter $\xi = \min\{\xi_1, \dots, \xi_m\}$ in the ξ layer; and (5) pruning the globally least important filter, i.e., $\min\{\xi, \eta\}$.

$$p_o = P\left(\sum_{i=1}^m \xi_i \geq a\right) + P\left(\sum_{i=1}^n \eta_i \geq b\right) \quad (1)$$

$$p_{\eta r} = P\left(\sum_{i=1}^m \xi_i \geq a\right) + P\left(\sum_{i=1}^{n-1} \eta_i \geq b\right) \quad (2)$$

¹This configuration can be extended to a multi-layer network (number of layers ≥ 3) with no difficulty.

$$p_{\underline{\eta}} = P\left(\sum_{i=1}^m \xi_i \geq a\right) + P\left(\sum_{i=1}^n \eta_i - \underline{\eta} \geq b\right) \quad (3)$$

$$p_{\underline{\xi}} = P\left(\sum_{i=1}^m \xi_i - \underline{\xi} \geq a\right) + P\left(\sum_{i=1}^n \eta_i \geq b\right) \quad (4)$$

$$p_g = \frac{m}{m+n} p_{\underline{\xi}} + \frac{n}{m+n} p_{\underline{\eta}} \quad (5)$$

It is worth mentioning that we consider the network performance from a perspective of redundancy (or capacity). That is why we do not divide the probabilities in Equations (1)-(5) by m or n . a or b can be considered as a threshold. As long as the total contribution of the filters in a layer is greater than the threshold, there is no performance loss. If a layer has too much redundancy (too many filters in our context), then it's very likely that the total contribution of the filters can still be greater than the threshold after pruning some of them.

Note that $0 \leq \eta_n - \underline{\eta} \leq \eta_n$, we have

$$P\left(\sum_{i=1}^{n-1} \eta_i \geq b\right) \leq P\left(\sum_{i=1}^n \eta_i - \underline{\eta} \geq b\right) \leq P\left(\sum_{i=1}^n \eta_i \geq b\right), \quad (6)$$

which indicates $p_{\eta_r} \leq p_{\underline{\eta}} \leq p_o$. For the filters in the η layer, we naturally assume that the contribution of a filter to the network's performance cannot be infinite, i.e., the variances of filters' contributions are uniformly bounded.

$$\exists C_1 > 0, \quad \text{s.t.} \quad \mathbb{D}\eta_i \leq C_1, i = 1, 2, \dots, n. \quad (7)$$

By Chebyshev's inequality, for any real number $\epsilon > 0$,

$$P\left(\frac{1}{n} \left| \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i) \right| \geq \epsilon\right) \leq \frac{\mathbb{D}(\sum_{i=1}^n \eta_i)}{\epsilon^2 n^2}. \quad (8)$$

With Equation (7), it is obvious that we have $\text{Cov}(\eta_i, \eta_j) \leq \sqrt{\mathbb{D}\eta_i \cdot \mathbb{D}\eta_j} \leq C_1$.

We further define that there are $C_2 n$ ($0 \leq C_2 \leq 1$) pairs of correlated filters in the η layer, i.e., $\#\{(i, j) : \text{Cov}(\eta_i, \eta_j) \neq 0, i \neq j, i, j = 1, \dots, n.\} \leq C_2 n$. Then we have,

$$\begin{aligned} \mathbb{D}\left(\sum_{i=1}^n \eta_i\right) &= \sum_{i=1}^n \mathbb{D}\eta_i + \sum_{i \neq j} \text{Cov}(\eta_i, \eta_j) \\ &\leq C_1 n + C_1 C_2 n = C_1(1 + C_2)n. \end{aligned}$$

By Equation (8),

$$P\left(\frac{1}{n} \left| \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i) \right| \geq \epsilon\right) \leq \frac{C_1(1 + C_2)}{\epsilon^2 n} \rightarrow 0.$$

This means $\frac{1}{n} \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i)$ converges in probability to zero, i.e., $\frac{1}{n} \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i) \xrightarrow{P} 0$. Suppose the number of filters in the η layer n is large enough, say $n > \frac{2b}{\epsilon_0}$.

We consider that a filter's contribution needs to be positive, but it could be infinitely small, i.e., the expectation of filters' contributions have a uniform positive lower bound.

$$\exists \epsilon_0 > 0, \quad \text{s.t.} \quad \mathbb{E}\eta_i \geq \epsilon_0, i = 1, 2, \dots, n. \quad (9)$$

With Equation (9), we have,

$$\begin{aligned} P\left(\frac{1}{n} \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i) > -\frac{\epsilon_0}{2}\right) &= P\left(\sum_{i=1}^n \eta_i > \sum_{i=1}^n \mathbb{E}\eta_i - \frac{\epsilon_0}{2}n\right) \\ &= P\left(\sum_{i=1}^n \eta_i > \frac{\epsilon_0}{2}n + \sum_{i=1}^n (\mathbb{E}\eta_i - \epsilon_0)\right) \\ &\leq P\left(\sum_{i=1}^n \eta_i > \frac{\epsilon_0}{2}n\right) \leq P\left(\sum_{i=1}^n \eta_i > b\right). \end{aligned}$$

Letting $n \rightarrow +\infty$, taking the limit and note that $\frac{1}{n} \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i) \xrightarrow{P} 0$, we have

$$\begin{aligned} \lim_{n \rightarrow \infty} P\left(\sum_{i=1}^n \eta_i > b\right) &\geq \lim_{n \rightarrow \infty} P\left(\frac{1}{n} \sum_{i=1}^n (\eta_i - \mathbb{E}\eta_i) > -\frac{\epsilon_0}{2}\right) = 1, \\ \lim_{n \rightarrow \infty} P\left(\sum_{i=1}^n \eta_i - \eta_r > b\right) &= \lim_{n \rightarrow \infty} P\left(\sum_{i=1}^n \eta_i - \underline{\eta} > b\right) = 1, \end{aligned}$$

and then we have $p_{\eta_r} \approx p_{\underline{\eta}} \approx p_o$ for n large enough. Note that $p_{\underline{\xi}} \leq p_o \approx p_{\underline{\eta}}$ and observe that p_g is the weighted average of $p_{\underline{\xi}}$ and $p_{\underline{\eta}}$. Hence we have $p_{\underline{\xi}} \leq p_g \leq p_{\underline{\eta}}$. It is worth mentioning that we cannot imply $p_g \approx p_{\underline{\eta}}$ from Equation (5) by letting $n \rightarrow \infty$ because we do not assume $m/n \rightarrow 0$.

In summary, we have $p_{\underline{\xi}} \leq p_g \leq p_{\eta_r} \leq p_{\underline{\eta}} \leq p_o$, which indicates that (even randomly) pruning a filter in the layer with much larger redundancy outperforms pruning the least important filter across all layers. Here we consider the CNN as a black-box and we do not assume any prior distribution for the RVs to achieve a good generalization. So the conclusion holds no matter how the RVs are distributed. Indeed, the conclusion relies on the assumption $n \rightarrow +\infty$. However, the assumption can be relaxed in real world applications such that $p_{\underline{\eta}} \geq p_g$ still holds on average (though not in every filter selection step). Appendix A presents some intuitive examples on a number of networks, which empirically provides evidence to validate the analysis above. As shown in Appendix A, the number of filters in a redundant layer does not need to be very large. Even when we use this naive strategy (randomly pruning filters in the layer with the most number of filters) with a standard AlexNet, it outperforms a number of popular pruning approaches. As we can see in the following, for more sophisticated architectures that contain less redundancy, such as ResNet, with a well designed metric to measure the layer redundancy, our proposed approach that prunes the least important filters in the layer(s) with larger redundancy outperforms recent pruning approaches that removes the least important filters across all layers.

4. Methodology

4.1. Notations and preliminaries

Network pruning. Suppose a CNN has L layers. For the i -th layer, the number of the input and output chan-




Figure 1. Overall workflow of the proposed approach. The number below each graph refers to the measurement of redundancy, which are just used for the illustration purpose and do not reflect the real measurements of the graphs.

nels are represented by N_i and N_{i+1} . Therefore, the CNN’s parameters \mathbf{W} can be represented as $\{\mathbf{W}^{(i)} \in \mathbb{R}^{N_{i+1} \times N_i \times h_i \times w_i}, i = 1, 2, \dots, L\}$, where h_i and w_i are the filter’s width and height. Channel pruning is formulated to find a set of parameters \mathbf{W}' optimized on certain objective functions such that $\|\mathbf{W}'\|_0 < K$, where $\|\cdot\|_0$ denotes the ℓ_0 norm and K limits the number of non-zero filters in \mathbf{W}' . Depending on different configurations, the objective functions can be minimizing a CNN’s cost function, drop of training accuracy, or the reconstruction error, etc.

Graph theory. Let X be a finite set. An undirected graph is a pair (X, E) , where E is a symmetric subset of $X \times X \setminus \{(x, x) : x \in X\}$. We call $x \in X$ a vertex (or a node) and $(x, y) \in E$ an edge. For $x, y \in X$, a path from x to y is a finite sequence $\{x_0, x_1, \dots, x_n\} \subset X$ such that $x_0 = x$, $x_n = y$ and $(x_i, x_{i+1}) \in E$. In general, the above path may not be unique if such a path exists. Denote $d(x, y)$ the minimal length of paths from x to y if x and y can be connected by a path; $d(x, y) = 0$ if $x = y$; and $d(x, y) = +\infty$ if x and y cannot be connected by a path. Then it is clear that $d(x, y)$ is an integer value metric. Recall that the degree of a vertex $x \in X$ is the total number of edges connected to x , i.e., $\deg(x) = \#\{(x, y) : (x, y) \in E\}$ ($\#A$ is the total number of elements in A).

4.2. Pruning with structural redundancy reduction

Overall architecture. We showed that pruning filters in the layer with larger redundancy outperforms pruning the least important filters across all layers. Our approach focuses on measuring how much redundancy exists in each layer and pruning filters from the most redundant layer(s) (Fig. 1). To measure the structural redundancy in a network, for each layer, we build an undirected graph in which each vertex represents a filter and the edges are defined with the distances between filter weights. We use two quantities associated with the graph, i.e., quotient space size and ℓ -covering number, as a measurement of how much redundancy exists in each graph, which is considered as the redundancy exists in each layer. At each time step, after the graph establishment and redundancy quantification, we randomly remove a vertex and its associated edges from the graph identified as with the most redundancy. Then we recalculate the redundancy after graph reconstruction for the next iteration. This process continues until a target is

reached (e.g., a certain number of filters are pruned). Finally, we prune the filters in each layer according to the remaining number of vertices in each graph with a certain filter selection criterion. Note that in the filter pruning phase, filters are ranked separately in each layer, rather than globally across all layers. Since the redundancy identification phase has selected a different number of filters in each layer, our approach is a layer-adaptive approach.

We present the details of our approach as follows: graph establishment, calculation of quotient space size and ℓ -covering number, intuition and quantification of graph redundancy, and filter selection.

Graph establishment. To illustrate how to build a graph for a convolutional layer, we use X to represent the filter weights of a certain layer $\mathbf{W}^{(i)}$ for simplicity. We first flatten and normalize the filter weights, which changes their lengths to 1. After that X becomes a finite subset of n dimensional unit sphere $S^n = \{x \in \mathbb{R}^n : |x| = 1\}$ in \mathbb{R}^n , where $n = N_i \times h_i \times w_i$ and $|x|$ is the length of x in \mathbb{R}^n . We define a graph on X as follows (assuming the elements in X are distinct). We choose a positive real number $\gamma > 0$ and define an edge set on X as

$$E = \{(x, y) \in X \times X \setminus \Delta : |x - y|/\sqrt{n} \leq \gamma\},$$

where $\Delta = \{(x, x) : x \in X\}$ is the diagonal of X , and $|x - y|$ is the Euclidean distance on \mathbb{R}^n . Then we get a graph (X, E) . By definition $(x, y) \in E$ implies x and y are approximately equal if γ is small.

ℓ -covering number. Recall that (X, d) is a metric space, where d is the graph metric defined previously. Let $\ell > 0$ be a fixed natural number, a subset $X_0 \subset X$ is called an ℓ -cover set of X , if $X \subset \bigcup\{B(x', \ell) : x' \in X_0\}$, where $B(x', \ell) = \{x \in X : d(x', x) \leq \ell\}$ is the ball centered at x' with radius ℓ . This means X is covered by the balls $\{B(x', \ell) : x' \in X_0\}$. We call the following quantity the ℓ -covering number of X :

$$N_\ell^c = (N_\ell^c(X) =) \min\{\#X_0 : X_0 \text{ is an } \ell\text{-cover set of } X\}.$$

Decomposition of a graph. We call a graph connected if for any $x \neq y$, there exists a path from x to y . In this case $d(x, y) < \infty$ for all $x, y \in X$. For an unconnected graph (X, E) , we define the notation “ \sim ” on X as follows: $x \sim y$ if and only if there exists a path from x to

y . Then it is clear that “ \sim ” is an equivalence relation. Let $X/\sim = \{X_1, X_2, \dots, X_k\}$ be the quotient space. This mathematical concept means that: using an equivalence relation, we can decompose the set X as a disjoint union $X = X_1 \cup X_2 \cup \dots \cup X_k$ such that the elements in the same X_i are equivalent. We call the number k (the total number of equivalence classes) the quotient space size. Intuitively, k is the number of unconnected sub-graphs of (X, E) .

Graph redundancy, intuition and quantification. Intuitively, larger values of the quotient space size and ℓ -covering number indicate a more complicated set of data (with less redundancy). In fact, $x \in B(x', \ell)$ if and only if $d(x, x') \leq \ell$, so x and x' are approximate equal. Hence the covering number can be approximately considered as the total number of vectors in X that are linearly independent. In our implementation we simply use $\ell = 1$, with the consideration of both performance and computation efficiency. Based on the above analysis, we define the graph (layer) redundancy as in Equation (10).

$$R(X) = \frac{N}{w_1 k + w_2 N_1^c}, \quad (10)$$

where $\{w_1, w_2\}$ is a probability weight that balances the importance of k and N_1^c , N is the number of filters. Besides the graph redundancy, we also investigate other criteria (i.e., the number of filters and principal component analysis (PCA)) to measure the structural redundancy in the ablation study.

Estimate of the 1-covering number. Since the calculation of ℓ -covering number is NP-hard and time-consuming in practice [7], we propose a lightweight method to estimate N_1^c . Let X_0 be the 1-cover set of a graph X , such that $\#X_0 = N_1^c$. We estimate $\#X_0$ as follows. Fix an integer $\ell (= 1 \text{ or } 2)$ and let $x_1^{(\ell)} \in X$, s.t. $\deg(x_1^{(\ell)}) = \max\{\deg(x) : x \in X\}$. We define a finite sequence $\{x_1^{(\ell)}, x_2^{(\ell)}, \dots, x_{n_\ell}^{(\ell)}\}$ by induction: If we have defined $x_k^{(\ell)}$, then there are two possible cases: (i) $X = \bigcup_{i=1}^k B(x_i^{(\ell)}, \ell)$, i.e., the family of balls $\{B(x_i^{(\ell)}, \ell) : 1 \leq i \leq k\}$ is an ℓ -cover of X . Then we stop the construction of the sequence and get $\{x_1^{(\ell)}, x_2^{(\ell)}, \dots, x_{n_k}^{(\ell)}\}$. (ii) Otherwise, choose (any) $x_{k+1}^{(\ell)} \in X \setminus \bigcup_{i=1}^k B(x_i^{(\ell)}, \ell)$, s.t.

$$\deg(x_{k+1}^{(\ell)}) = \max\{\deg(x) : x \in X \setminus \bigcup_{i=1}^k B(x_i^{(\ell)}, \ell)\}.$$

We repeat the above process eventually, and get the sequence

$$\{x_1^{(\ell)}, x_2^{(\ell)}, \dots, x_{n_\ell}^{(\ell)}\}, \quad l = 1 \text{ or } 2.$$

It is obvious that we have $N_1^c = \#X_0 \leq n_1$ because the family $\{B(x_k^{(1)}, 1) : 1 \leq k \leq n_1\}$ is a 1-cover of X . Moreover, for any $i \neq j$ ($i, j \leq n_2$), we have $d(x_i^{(2)}, x_j^{(2)}) \geq 3$.

On the other hand, for each $x_0 \in X_0$, and any $x, y \in B(x_0, 1)$, we have $d(x, y) \leq d(x, x_0) + d(x_0, y) \leq 2$. Recall that X_0 is a 1-cover set of X , then for any $x_i^{(2)}$, there exists (may not unique) $x_0 \in X_0$ such that $x_i^{(2)} \in B(x_0, 1)$. Moreover, for $i \neq j$, $x_i^{(2)}$ and $x_j^{(2)}$ cannot be in the same ball $B(x, 1)$ (otherwise $d(x_i^{(2)}, x_j^{(2)}) \leq 2$, a contradiction). We see $n_2 \leq \#X_0 = N_1^c$. Hence $n_2 \leq N_1^c \leq n_1$.

We can use $\tilde{N}_1^c = \frac{1}{2}(n_1 + n_2)$ to estimate N_1^c , if $|n_1 - n_2|$ is acceptably small. Although we cannot theoretically find its upper bound, extensive experiments on various networks show that \tilde{N}_1^c is good enough as an estimation of N_1^c , and the computing time of \tilde{N}_1^c is negligibly small (see the Analysis and ablation study section).

Filter selection strategy. After identifying the layers with large redundancy, we prune unimportant filters from these layers. We can either train a pruned network architecture from scratch with random initialization or prune certain filters from the pre-trained network and do fine-tuning. There are various approaches for unimportant filter selection. In our study, we use a very common and simple strategy, i.e., pruning the filters with smaller absolute weights [24]. This method avoids feeding a large number of training samples into the CNN to get filter rankings, which is usually computationally intensive [38, 35]. But in general, our approach can be used together with any filter selection criterion.

5. Experiments

5.1. Experiment settings

We first evaluate our approach with the single-shot pruning scheme (pruning a large number of filters at one time), with two-widely used benchmark datasets (CIFAR-10 [20] and ImageNet ILSVRC-2012 [2]) on ResNet. We also present the results with the progressive pruning scheme (pruning a small number of filters and fine-tuning the remaining network for multiple times), which are presented in the Appendix due to the space limitation.

For single-shot pruning, we use ResNet{20,56} on the CIFAR-10 dataset and Resnet50 on ImageNet to evaluate the performance, in terms of accuracy drop and FLOPs reduction. We used the widely-used ResNet architecture as described in [13]. For the CIFAR-10 experiments, the models are trained following the setup in [15]. For the ImageNet experiments, pre-trained models from torchvision are used. We first evaluate the layer redundancy in the pre-trained models and identify the number of filters to be pruned in each layer, with our proposed approach. Then we prune the filter in each layer with the corresponding numbers identified and fine-tune the slimmed network. We follow the fine-tuning strategy in [14]. For CIFAR-10, we fine-tune each pruned network for 200 epochs, with a learning rate starting from 0.1, which is divided by 10 at the epochs 60, 120,

Model	Approach	Acc. before prune	Acc. after prune	Acc. drop	FLOPs drop
ResNet20	MW	92.35%	90.93%	1.42%	41.0%
	SFP	92.20%	90.83%	1.37%	42.4%
	GM	92.20%	91.09%	1.11%	42.2%
	TAS	-	92.88%	0.00%	45.0%
	SRR-GR	92.27%	92.48%	-0.21%	45.8%
ResNet56	MW	93.51%	92.90%	0.61%	51.5%
	NISP	-	93.01%	-	35.5%
	GAL	93.26%	93.38%	-0.12%	37.6%
	DCP	93.80%	93.49%	0.31%	49.8%
	HRank	93.26%	93.17%	0.09%	50.0%
	SCP	93.69%	93.23%	0.46%	51.5%
	SFP	93.59%	92.26%	1.33%	52.6%
	GM	93.59%	92.93%	0.66%	52.6%
	TAS	-	93.69%	0.77%	52.7%
	SRR-GR	93.38%	93.75%	-0.37%	53.8%

Table 1. Results of ResNet on CIFAR-10. MW results is with our own implementation. GR is graph redundancy.

and 160. For ImageNet, we fine-tune each pruned network for 150 epochs, with a learning rate starting from 0.1, which is divided by 10 every 30 epochs. For all the models, we use an SGD optimizer with a momentum of 0.9, a weight decay of $2e^{-5}$, and a batch size of 256. For the graph associated parameters, we use $w_1 = 0.35$, $w_2 = 0.65$ to emphasize the importance of ℓ -covering number. We choose $\gamma = 0.034$ to achieve the best performance. We implement the experiments with Pytorch 1.3 [36].

We compare the performance of our approach with several recent channel pruning methods, namely, minimum weight (MW) [24], Taylor expansion [35], average percentage of zero activation neurons (APoZ) [17], soft filter pruning (SFP) [14], discrimination-aware channel pruning (DCP) [52], neuron importance score propagation (NISP) [47], slimmable neural networks (SNN) [46], autopruner (AP) [30], generative adversarial learning (GAL) [27], geometric median (GM) [15], transformable architecture search (TAS) [5], cluster pruning (CUP) [6], ABC [26], trained rank pruning (TRP) [45], soft channel pruning (SCP) [19], and high-hank (HRank) [25].

5.2. Performance evaluation

CIFAR-10. Results of pruning ResNet20 and ResNet56 on CIFAR-10 are presented in Table 1. Our approach prunes a large percent of FLOPs from both architectures without performance degradation, which outperforms the previous state-of-the-art with an obvious margin. For ResNet20, we prune 45.8% FLOPs and the test accuracy is increased by 0.21%. Our pruned ResNet56 model reduces 53.8% FLOPs and achieves a test accuracy of 93.75%, which outperforms the baseline by 0.37%.

ImageNet. Results of ResNet50 on ImageNet are shown in Table 2. Our pruned model with 44.1% FLOPs reduction only loses 0.37% top-1 accuracy and 0.19% top-5 accuracy. When pruning comparable FLOPs, the top-1 accuracy

of the previous state-of-the-art approaches usually drop by more than 1%. As the pruning ratio increases to 55.1%, the proposed approach can still achieve a promising test accuracy (1.02% and 0.51% drop for top-1 and top-5 accuracy), which is the best performance compared with recent works. These results verify the effectiveness of our approach on the single-shot pruning scheme. It is worth mentioning that the MW approach can be considered as a baseline of our approach because we add a redundancy identification stage before using MW for filter pruning. It is observed that pruning filters uniformly with MW results in unsatisfactory results (71.24% top-1 accuracy). After identifying the redundancy and pruning the corresponding number of filters in each layer, the performance is significantly improved by around 4%.

6. Analysis and ablation study

6.1. 1-covering number estimate and computation time

We first show the effectiveness of our approach for 1-covering number estimate. We build a series of graphs for each layer of a pre-trained AlexNet and VGG16 by changing γ from 0.001 to 0.3 and visualize n_1 and n_2 (defined in the Methodology section) for the illustration purpose (Fig. 2(a-b)). Obviously, $n_1 \approx n_2$ in nearly all cases. The same trend is also observed in ResNet. Actually, in all of our experiments, we do not observe any large deviations between n_1 and n_2 . It is with negligible influence to use $\tilde{N}_1^c = \frac{1}{2}(n_1 + n_2)$ as an estimate of N_1^c , whose values is between n_1 and n_2 .

We also evaluate the computing time for estimating the 1-covering number N_1^c . To obtain N_1^c , a complete search of all combinations of vertices has to be done to see if all the vertices are covered by the selected balls, We name this approach as the oracle approach. We measure the computation

Approach	Top-1 acc. baseline	Top-5 acc. baseline	Top-1 acc. after prune	Top-5 acc. after prune	Top-1 acc. ↓	Top-5 acc. ↓	FLOPs ↓
MW	76.13%	92.86%	71.24%	90.38%	4.89%	2.48%	41.8%
SFP	76.15%	92.87%	74.61%	92.06%	1.54%	0.81%	41.8%
GM	76.15%	92.87%	75.03%	92.40%	1.12%	0.47%	42.2%
GAL	76.15%	92.87%	71.95%	90.94%	4.20%	1.93%	43.0%
TAS	-	-	76.20%	93.07%	1.26%	0.48%	43.5%
HRank	76.15%	92.87%	74.98%	92.33%	1.17%	0.54%	43.8%
SNN	-	-	74.90%	-	1.10%	-	43.9%
SRR-GR	76.13%	92.86%	75.76%	92.67%	0.37%	0.19%	44.1%
TRP	-	-	74.06%	92.07%	-	-	44.4%
AP	76.15%	92.87%	74.76%	92.15%	1.39%	0.72%	51.2%
GM	76.15%	92.87%	74.13%	91.94%	2.02%	0.93%	53.5%
ABC	76.01%	92.96%	73.86%	91.69%	2.15%	1.27%	54.0%
SCP	75.89%	92.98%	74.20%	92.00%	1.69%	0.98%	54.3%
CUP	-	-	-	-	1.47%	0.88%	54.5%
GAL	76.15%	92.87%	71.80%	90.82%	4.35%	2.05%	55.0%
SRR-GR	76.13%	92.86%	75.11%	92.35%	1.02%	0.51%	55.1%

Table 2. ResNet50 results on ImageNet. Results of MW, APoZ, and Taylor are based on our own implementation.




Figure 2. (a-b) The value of n_1 and n_2 by changing γ from 0.001 to 0.3. Black solid line refers to $n_1 = n_2$. (c-d) The network structure comparison when 40% filters are pruned from AlexNet using different γ s.

Num of filters	1-covering number	Time with oracle method	Time with our method
64	1	0.001	0.0015
	2	0.045	0.0013
	3	1.384	0.0020
	4	28.21	0.0026
192	1	0.023	0.0014
	2	0.974	0.0027
	3	90.29	0.0028

Table 3. Average computation time of the oracle and our proposed method for the 1-covering number calculation/estimate (in secs).

time of the oracle approach and our proposed lightweight approach for calculating 1-covering number. It is clear that with the oracle method, the computation time grows drastically as the number of filters and the actual value of N_1^c increase. It is even not temporally feasible to use the oracle approach when $N_1^c > 4$. In contrast, the time used with our method for N_1^c is negligibly short and is merely influenced by the number of filters in the layer and the value of N_1^c . These results indicate that our proposed approach for the estimate of N_1^c is valid and efficient. Therefore, the real

Approach	MW	MA	Taylor	Random
Accuracy	75.99%	75.84%	75.95%	75.82%

Table 4. Performance with different filter selection criteria after pruning 30% FLOPs of AlexNet.

running time of the proposed approach is almost the same as existing methods with the same filter selection criteria.

6.2. Filter selection criteria

In previous experiments, we use the minimum weight criterion to prune filters. We further investigate whether different filter selection criteria have any influence on the performance. We use AlexNet on CIFAR-10 for illustration, by pruning 30% FLOPs and fine-tuning the remaining networks for 100 epochs with a learning rate of $1e^{-4}$. Results (Table 4) indicate that choosing filters with the minimum weight strategy achieves the best accuracy (75.99%). However, using other filter selection criteria results in a similar performance, and the accuracy only drops 0.17% even we randomly prune filters in the layers identified as with large redundancy by our approach. Therefore, our proposed approach is not sensitive to filter selection criteria, which fur-

Approach	Top-1 acc. baseline	Top-5 acc. baseline	Top-1 acc. after prune	Top-5 acc. after prune	Top-1 acc. ↓	Top-5 acc. ↓	FLOPs ↓
SRR-NOF			74.88%	92.27%	1.25%	0.59%	44.0%
SRR-PCA	76.13%	92.86%	75.19%	92.48%	0.94%	0.38%	44.1%
SRR-GR			75.76%	92.67%	0.37%	0.19%	44.1%

Table 5. Performance of the pruned ResNet50 networks with different redundancy reduction measurements.

ther validates the fact that pruning filters in the layers with large redundancy is more essential than identifying unimportant filters.

6.3. The value of gamma

We change the distance threshold γ for graph establishment to analyze its influence on the performance. We keep using AlexNet on CIFAR-10 as an example by pruning 40% of the filters with $\gamma = 0.003, 0.034$, and 0.3 . Results (Fig. 2(c-d)) shows that with a large γ , the last four layers remain the same number of filters, which indicates that the layer with the largest number of filters are identified as the redundant layer at each time (Fig. 2(c)). When γ is small, nearly the same percent of filters are removed from all layers (Fig. 2(d)). With a suitable value of γ (0.034), our approach identified Layer 3 as the most redundant layer. Layer 4 and 5 are also redundant to some extent but Layer 4 is with a little more redundancy. Different from other existing works, our approach suggests that Layer 1 should not be pruned if we only aim to remove 40% of the filters from AlexNet. These results are consistent with the definition of layer redundancy (in the Methodology section). For a layer with n filters, when $\gamma \rightarrow 0$, it is clear that $k = n$, $N_1^c = n$, and $R(X) = 1$. Therefore, all layers have the same level of redundancy and the approach becomes a uniform pruning. When $\gamma \rightarrow +\infty$, $k = 1$, $N_1^c = 1$, and $R(X) = n$, which indicates that the layer with the most number of filters is with the largest redundancy. Our approach can be considered as a dynamic architecture search approach controlled by γ .

6.4. Other criteria for structural redundancy identification

Since there are few studies that consider network pruning from the perspective of structural redundancy reduction, we further investigate the effectiveness of structural redundancy reduction for channel pruning with the following layer redundancy measurement metrics. (1) SRR-NOF: This strategy simply uses the number of filters in the convolutional layers as the measurement of layer redundancy. The layer with more filters is considered as with more redundancy. In the redundancy identification phase, for each iteration, a filter from the layer with the most number of filters are removed. If there exist more than one layer containing the same number of the most filters, a filter is removed from a randomly chosen layer. When the requirement is reached, the network is pruned with the minimum

weight criterion according to the remaining number of filters in each layer. (2) SRR-PCA: This strategy uses principal component analysis (PCA) on the intermediate feature maps of a network to measure the correlation between filters. We first feed training samples to the CNN and record the flattened feature maps of each convolutional layer. Then we fit PCA on these flattened feature maps and get a list of percentage of variance explained by each of them for all convolutional layers. We select the N smallest percentage of variances across all layers and count how many items are selected in each convolutional layer. Finally, we prune the filters in each layer accordingly, with the minimum weight ranking criterion. (3) SRR-GR: This strategy uses graph redundancy as described in the Methodology section. The training and pruning configuration are the same as in the Experiment section.

Experiment results (Table 5) show that by pruning 44% FLOPs from ResNet50, even with a naive layer redundancy measurement (i.e., the number of filters in the layer), the performance is comparable to recent studies. PCA identifies the layer redundancy better than NOF and with SRR-PCA, the drops of top-1 and top-5 accuracy further decrease to 0.94% and 0.38%. With the graph redundancy-based approach, the pruning performance is significantly improved. These results validate that (1) structural redundancy reduction is an efficient approach for channel pruning, and (2) the proposed graph redundancy-based approach is a promising way for layer redundancy measurement.

7. Conclusion

We theoretically studied the rationale behind network pruning from a perspective of redundancy reduction via a statistical modeling and discovered that pruning filters in the layer(s) with the most structural redundancy plays a more essential role than pruning the least important filters across all layers. We proposed to identify the level of redundancy existed in each convolutional layer of a CNN via a graph establishment for each layer and two graph-related quantities as the measurement of the redundancy. After that, filters are pruned from the selected layer(s) by a simple filter selection criterion. Experimental results validated that our approach improved the state-of-the-art on image classification tasks. We believe that the proposed approach can be effective on more complicated tasks such as object detection and image synthesis, which is left for future research.

References

- [1] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, and Yixin Chen. Compressing neural networks with the hashing trick. In *International Conference on Machine Learning*, pages 2285–2294, 2015. 2
- [2] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1, 5
- [3] Xiaohan Ding, Guiguang Ding, Yuchen Guo, Jungong Han, and Chenggang Yan. Approximated oracle filter pruning for destructive cnn width optimization. *arXiv preprint arXiv:1905.04748*, 2019. 1
- [4] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 1
- [5] Xuanyi Dong and Yi Yang. Network pruning via transformable architecture search. In *Advances in Neural Information Processing Systems*, pages 759–770, 2019. 6
- [6] Rahul Duggal, Cao Xiao, Richard Vuduc, and Jimeng Sun. Cup: Cluster pruning for compressing deep neural networks. *arXiv preprint arXiv:1911.08630*, 2019. 2, 6
- [7] Matthew Richard Gibson. Clusters and covers: geometric set cover algorithms. 2010. 5
- [8] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014. 1
- [9] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. Eie: efficient inference engine on compressed deep neural network. In *Computer Architecture (ISCA), 2016 ACM/IEEE 43rd Annual International Symposium on*, pages 243–254. IEEE, 2016. 1, 2
- [10] 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. 1
- [11] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pages 1135–1143, 2015. 1, 2
- [12] Stephen José Hanson and Lorien Y Pratt. Comparing biases for minimal network construction with back-propagation. In *Advances in neural information processing systems*, pages 177–185, 1989. 2
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 5
- [14] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018. 5, 6
- [15] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019. 1, 2, 5, 6
- [16] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 1
- [17] 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. 2, 6
- [18] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017. 1
- [19] Minsoo Kang and Bohyung Han. Operation-aware soft channel pruning using differentiable masks. *arXiv preprint arXiv:2007.03938*, 2020. 6
- [20] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009. 5
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pages 1097–1105, 2012. 1
- [22] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *nature*, 521(7553):436, 2015. 1
- [23] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pages 598–605, 1990. 2
- [24] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016. 1, 2, 5, 6
- [25] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: Filter pruning using high-rank feature map. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1529–1538, 2020. 2, 6
- [26] Mingbao Lin, Rongrong Ji, Yuxin Zhang, Baochang Zhang, Yongjian Wu, and Yonghong Tian. Channel pruning via automatic structure search. *arXiv preprint arXiv:2001.08565*, 2020. 6
- [27] Shaohui Lin, Rongrong Ji, Chenqian Yan, Baochang Zhang, Liujuan Cao, Qixiang Ye, Feiyue Huang, and David Doermann. Towards optimal structured cnn pruning via generative adversarial learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2790–2799, 2019. 6
- [28] Xingyu Liu, Jeff Pool, Song Han, and William J Dally. Efficient sparse-winograd convolutional neural networks. *arXiv preprint arXiv:1802.06367*, 2018. 2
- [29] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018. 2
- [30] Jian-Hao Luo and Jianxin Wu. Autopruner: An end-to-end trainable filter pruning method for efficient deep model inference. *arXiv preprint arXiv:1805.08941*, 2018. 6

- [31] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. *arXiv preprint arXiv:1707.06342*, 2017. **2**
- [32] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 116–131, 2018. **1**
- [33] Deepak Mittal, Shweta Bhardwaj, Mitesh M Khapra, and Balaraman Ravindran. Recovering from random pruning: On the plasticity of deep convolutional neural networks. *arXiv preprint arXiv:1801.10447*, 2018. **2**
- [34] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Andrei A Rusu, Joel Veness, Marc G Bellemare, Alex Graves, Martin Riedmiller, Andreas K Fiedjeland, Georg Ostrovski, et al. Human-level control through deep reinforcement learning. *Nature*, 518(7540):529–533, 2015. **1**
- [35] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. *arXiv preprint arXiv:1611.06440*, 2016. **1, 2, 5, 6**
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems*, pages 8024–8035, 2019. **6**
- [37] Hanyu Peng, Jiaxiang Wu, Shifeng Chen, and Junzhou Huang. Collaborative channel pruning for deep networks. In *International Conference on Machine Learning*, pages 5113–5122, 2019. **2**
- [38] Adam Polyak and Lior Wolf. Channel-level acceleration of deep face representations. *IEEE Access*, 3:2163–2175, 2015. **2, 5**
- [39] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. *arXiv preprint arXiv:1801.04381*, 2018. **1**
- [40] David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *nature*, 529(7587):484, 2016. **1**
- [41] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. **1**
- [42] Sergey Tulyakov, Ming-Yu Liu, Xiaodong Yang, and Jan Kautz. Mocogan: Decomposing motion and content for video generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1526–1535, 2018. **1**
- [43] Chaoqi Wang, Roger Grosse, Sanja Fidler, and Guodong Zhang. Eigendamage: Structured pruning in the kronecker-factored eigenbasis. *arXiv preprint arXiv:1905.05934*, 2019. **1**
- [44] Zi Wang, Chengcheng Li, Xiangyang Wang, and Dali Wang. Towards efficient convolutional neural networks through low-error filter saliency estimation. In *Pacific Rim International Conference on Artificial Intelligence*, pages 255–267. Springer, 2019. **1**
- [45] Yuhui Xu, Yuxi Li, Shuai Zhang, Wei Wen, Botao Wang, Yingyong Qi, Yiran Chen, Weiyao Lin, and Hongkai Xiong. Trp: Trained rank pruning for efficient deep neural networks. *arXiv preprint arXiv:2004.14566*, 2020. **6**
- [46] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *arXiv preprint arXiv:1812.08928*, 2018. **1, 6**
- [47] 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. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9194–9203, 2018. **2, 6**
- [48] Xiyu Yu, Tongliang Liu, Xinchao Wang, and Dacheng Tao. On compressing deep models by low rank and sparse decomposition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 7370–7379, 2017. **1**
- [49] Wenyuan Zeng and Raquel Urtasun. Mlprune: Multi-layer pruning for automated neural network compression. 2018. **2**
- [50] Tianyun Zhang, Shaokai Ye, Kaiqi Zhang, Jian Tang, Wujie Wen, Makan Fardad, and Yanzhi Wang. A systematic dnn weight pruning framework using alternating direction method of multipliers. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 184–199, 2018. **2**
- [51] Zhengguang Zhou, Wengang Zhou, Houqiang Li, and Richang Hong. Online filter clustering and pruning for efficient convnets. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 11–15. IEEE, 2018. **2**
- [52] Zhuangwei Zhuang, Mingkui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu. Discrimination-aware channel pruning for deep neural networks. In *Advances in Neural Information Processing Systems*, pages 875–886, 2018. **6**