

Pushing the Efficiency Limit Using Structured Sparse Convolutions

Vinay Kumar Verma^{1*}, Nikhil Mehta^{1*}, Shijing Si³, Ricardo Henao¹, Lawrence Carin^{1,2}

¹Duke University ²KAUST Saudi Arabia ³SEF, Shanghai International Studies University

Abstract

Weight pruning is among the most popular approaches for compressing deep convolutional neural networks. Recent work suggests that in a randomly initialized deep neural network, there exist sparse subnetworks that achieve performance comparable to the original network. Unfortunately, finding these subnetworks involves iterative stages of training and pruning, which can be computationally expensive. We propose Structured Sparse Convolution (SSC), that leverages the inherent structure in images to reduce the parameters in the convolutional filter. This leads to improved efficiency of convolutional architectures compared to existing methods that perform pruning at initialization. We show that SSC is a generalization of commonly used layers (depthwise, groupwise and pointwise convolution) in “efficient architectures.” Extensive experiments on well-known CNN models and datasets show the effectiveness of the proposed method. Architectures based on SSC achieve state-of-the-art performance compared to baselines on CIFAR-10, CIFAR-100, Tiny-ImageNet, and ImageNet classification benchmarks. Our source code is publicly available at <https://github.com/vkvermaa/SSC>.

1. Introduction

Overparameterized deep neural networks (DNNs) are known to generalize well on the test data [4, 1]. However, overparameterization increases the network size, making DNNs resource-hungry and leading to extended training and inference time. This hinders the training and deployment of DNNs on low-power devices and limits the application of DNNs in systems with strict latency requirements. Several efforts have been made to reduce the storage and computational complexity of DNNs using model compression [19, 62, 51, 45, 33, 3, 34, 7, 44]. Network pruning is the most popular approach for model compression. In network pruning, we compress a large neural network by pruning redundant parameters while maintaining the model performance. The pruning approaches can be divided into two categories: unstructured and structured. Unstructured prun-

ing removes redundant connections in the kernel, leading to sparse tensors [27, 18, 64]. Unstructured sparsity produces sporadic connectivity in the neural architecture, causing irregular memory access [55] that adversely impacts the acceleration in hardware platforms. On the other hand, structured pruning involves pruning parameters that follow a high-level structure (*e.g.*, pruning parameters at the filter-level [29, 34, 12]). Typically, structure pruning leads to practical acceleration, as the parameters are reduced while the memory access remains contiguous. Existing pruning methods typically involve a three-stage pipeline: pretraining, pruning and finetuning, where the latter two stages are carried out in multiple stages until a desired pruning ratio is achieved. While the final pruned model leads to a low inference cost, the cost to achieve the pruned architecture remains high.

The lottery ticket hypothesis (LTH) [15, 16] showed that a randomly initialized overparametrized neural network contains a sub-network, referred to as the “winning ticket,” that when trained in isolation achieves the same test accuracy as the original network. Similar to LTH, there is compelling evidence [38, 39, 14, 13, 2, 1, 45] suggesting that overparameterization is not essential for high test accuracy, but is helpful to find a good initialization for the network [30, 65]. However, the procedure to find such sub-networks involves iterative pruning [15] making it computationally intensive. If we know the sub-network beforehand, we can train a much smaller and efficient model with only 1-10% of the parameters of the original network, reducing the computational cost involved during training.

An open research question concerns how to design a sub-network without undergoing the expensive multi-stage process of training, pruning and finetuning. There have been recent attempts [27, 52] to alleviate this issue, involving a one-time neural network pruning at initialization by solving an optimization problem for detecting and removing unimportant connections. Once the sub-network is identified, the model is trained without carrying out further pruning. This procedure of pruning only once is referred to as pruning at initialization or foresight pruning [52]. While these methods can find an approximation to the winning ticket, they have the following limitations hindering their practical applicability: (1) The initial optimization procedure still requires

*denotes equal contribution. Correspondence to
vinayugc@gmail.com, nick.mehta@duke.edu.

large memory, since the optimization process is carried out over the original overparameterized model. (2) The obtained winning ticket is specific to a particular dataset on which it is approximated, *i.e.*, a network pruned using a particular dataset may not perform optimally on a different dataset. (3) These pruning based methods lead to unstructured sparsity in the model. Due to common hardware limitations, it is very difficult to get a practical speedup from unstructured compression.

In this paper, we design a novel structured sparse convolution (SSC) filter for convolutional layers, requiring significantly fewer parameters compared to standard convolution. The proposed filter leverages the inherent spatial properties in the images. The commonly used deep convolutional architectures, when coupled with SSC, outperform other state-of-the-art methods that do pruning at initialization. Unlike typical pruning approaches, the proposed architecture is sparse by design and does not require multiple stages of pruning. The sparsity of the architecture is dataset agnostic and leads to better transfer ability of the model when compared to existing state-of-the-art methods that do pruning at initialization. We also show that the proposed filter has implicit orthogonality that ensures minimum filter redundancy at each layer. Additionally, we show that the proposed filter can be viewed as a generalization of existing efficient convolutional filters used in group-wise convolution (GWC) [59], point-wise convolution (PWC) [48], and depth-wise convolution (DWC) [50]. Extensive experiments and ablation studies on standard benchmarks depict the efficacy of the proposed filter. Moreover, we further compress existing efficient models such as MobileNetv2 [40] and ShuffleNetv2 [35] while achieving performance comparable to the original models.

2. Methods

We propose Structured Sparse Convolution (SSC) filter, which is composed of layered spatially sparse $K \times K$ and 1×1 kernels. Unlike the typical CNN filters that have a kernel of fixed size, the SSC filter has three type of kernels, as shown in Figure 1. The heterogeneous nature of the kernels are designed to have varying receptive fields that can capture different features in the input. As shown in Section 2.1, heterogeneity in kernels allows the neural network layer to accumulate information from different spatial locations in the feature map while significantly reducing redundancy in the number of parameters. Consider layer l of a model with an input (h_{l-1}) of size $i_{l-1} \times i_{l-1} \times M$, where i_{l-1} corresponds to the spatial dimension (width and height) and M denotes the number of channels of the input. Assume that layer l has N filters, resulting in an output feature map h_l of size $i_l \times i_l \times N$. We represent the computational and memory cost at the l^{th} layer using the number of floating-point operations (F_l) and the number of parameters (P_l), respectively. The computational and memory cost associated with

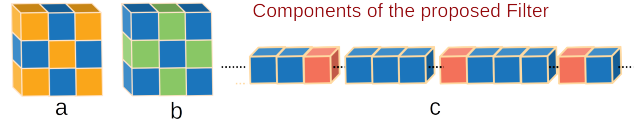


Figure 1. The three basic components that are used in the proposed SSC filter. Blue blocks indicates a zero-weight location. The red, orange and green blocks show the active weight location in three different type of kernels.

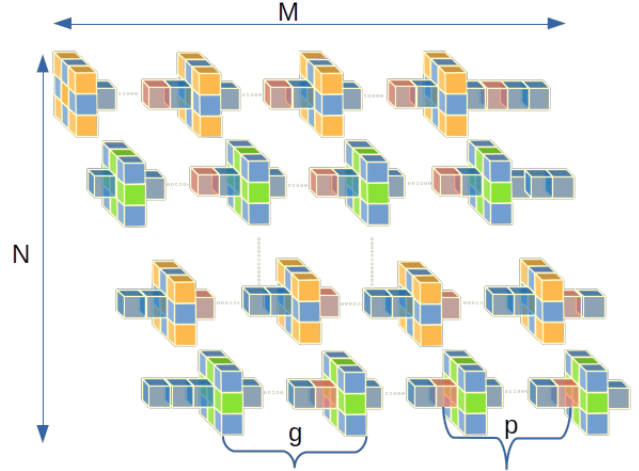


Figure 2. The proposed convolutional layer with N SSC filters. Blue blocks denote the zero-weight locations in 3×3 and 1×1 kernels, while other colors show active weights.

a standard convolutional layer with a $K \times K$ kernel is the following:

$$F_l = i_l^2 \times N \times (K^2 M), \quad (1)$$

$$P_l = N \times (K^2 M), \quad (2)$$

where $(K^2 M)$ represents the number of total parameters from all M channel-specific kernels. As is evident from (1) and (2), reducing $(K^2 M)$ directly reduces both the number of parameters and the computational cost of the model. This is indeed what our proposed method (SSC) achieves – we design two types of sparse kernels, which form the basic components of SSC.

Odd/Even $K \times K$ kernel: The two types of $K \times K$ kernels differ in terms of the location of the enforced sparsity. Considering $S \in R^{K^2}$ to be the flattened version of the $K \times K$ 2D kernel, we define the odd kernel as:

$$\begin{cases} S[i] = 0 & i \in \{2p \mid 0 < 2p < K^2, p \in \mathbb{N}\} \\ S[i] = w_i & i \in \{2p+1 \mid 0 < 2p+1 < K^2, p \in \mathbb{N}\} \end{cases}, \quad (3)$$

The even kernel is defined in a similar fashion, where the kernel is zero at odd coordinates and non-zero at even coordinates of the filter. Figure 1 (a-b) illustrates the odd and even kernel, respectively, when $K = 3$. These kernels replace the standard $K \times K$ kernels used in a convolutional layer.

SSC Filter: A convolutional layer having N SSC filters is shown in Figure 2. An SSC filter is referred to as odd (or even) filter, if it only contains odd (or even) kernels. For each convolutional layer, an equal number of odd and even filters are used. We make the following modifications to a standard convolution filter having M , $K \times K$ kernels:

1. Among the M different kernels, we replace each kernel at the $k * g$ location with an odd/even kernel, where g is a hyperparameter such that $0 < g < M$ and $k \in \{n \in \mathbb{N} \mid 0 < k * g \leq M\}$. Note that each filter has only one type (odd/even) of kernel. Each of the N filters have M/g such kernels. The computational cost (F_{sg}) and the memory cost (P_{sg}) for all the odd/even kernels in a filter:

$$F_{sg} = i_l^2 \times N \times \frac{(K^2 - c)M}{g}, \quad (4)$$

$$P_{sg} = N \times \frac{(K^2 - c)M}{g}, \quad (5)$$

$$c = \begin{cases} \lceil \frac{K^2}{2} \rceil & \text{Odd kernel} \\ K^2 - \lceil \frac{K^2}{2} \rceil & \text{Even kernel} \end{cases}, \quad (6)$$

where c represents the number of zeros in the kernel and $\lceil \cdot \rceil$ denotes the ceiling function.

2. Out of the remaining $M(1 - 1/g)$ kernel locations in the filter, we place a 1×1 kernel at a fixed interval of p as shown in Figure 1 (c). Each of the N filters have $M(1 - 1/g)/p$ 1×1 kernels. The computational and memory cost of these 1×1 kernels can be defined as:

$$F_{sp} = i_l^2 \times N \times \frac{M(1 - 1/g)}{p}, \quad (7)$$

$$P_{sp} = N \times \frac{M(1 - 1/g)}{p}. \quad (8)$$

3. The SSC filter is empty at the remaining $M(1 - 1/p)(1 - 1/g)$ locations, causing the filter to ignore the corresponding feature maps (input channels). Note that while a particular filter may not act on certain input features, other SSC filters of the convolutional layer will. This is enforced by the shifting procedure introduced below.

Shift operation: If we naively use SSC filters in a convolutional layer, there will be a loss of information, as all N filters will ignore the same input feature maps. To ensure that each SSC filter attends to a different set of feature maps, we shift the location of all kernels ($K \times K$ and 1×1) by $(n \bmod q)$ at initialization¹, where $n \in \{1, \dots, N\}$ denotes the index of the filter and $q := \max(g, p)$. The shift operation across N filters can be visualized in Figure 2. We can divide the N filters into sets of disjoint filters such that all

¹The shift operation is applied only once before training begins.

the filters in a particular set attend to distinct input feature maps. Formally, let the collection of sets be defined as:

$$\mathcal{Q} := \{(0, q), [q, 2q), \dots, [N - (N \bmod q), N)\}, \quad (9)$$

where $[a, a + q)$ denotes the set of filters a through $a + q - 1$. Then $\forall f, f' \in [a, a + q)$, f and f' tend to disjoint input feature maps if $f \neq f'$. Moreover, f and f' are “near-orthogonal” ($f^T f' \approx 0$), since they attend to non-overlapping regions of the input feature maps. As discussed in Section 2.1, the orthogonal property of a layer is of independent interest and allows the network to learn uncorrelated filters. Note that the design of the SSC filter induces structural sparsity as the sparse region is predetermined and fixed, which is in contrast to the unstructured pruning method [15, 27, 52].

We can quantify the total reduction in the number of floating-point operations (R_F) and the number of parameters (R_p) with respect to the standard convolutional layer:

$$R_F = \left(1 - \frac{F_{sg} + F_{sp}}{F_l}\right) \times 100\% \quad (10)$$

$$= \left(1 - \frac{(1 - c/K^2)}{g} - \frac{(1 - 1/g)}{K^2 p}\right) \times 100\%, \quad (11)$$

$$R_p = \left(1 - \frac{P_{sg} + P_{sp}}{P_l}\right) \times 100\% \quad (12)$$

$$= \left(1 - \frac{(1 - c/K^2)}{g} - \frac{(1 - 1/g)}{K^2 p}\right) \times 100\%, \quad (13)$$

where $0 < g$ and $p < M$. The hyperparameters p and g are set to achieve the desired sparsity in the architectures; we use R_p as guiding principle behind choosing p and g . One can also achieve a desired reduction in floating-point operations (R_F) to determine the corresponding hyperparameters. However, in our experiments we consider sparsity constraints only.

2.1. Implicit Orthogonality

Recent work [42, 58] shows that deep convolutional networks learn correlated filters in overparameterized regimes. This implies filter redundancy and correlated feature maps when working with deep architectures. The issue of correlation across multiple filters in the convolutional layer has been addressed by incorporating an explicit orthogonality constraint to the filter of each layer [5, 53]. Consider a 2D matrix $W \in \mathbb{R}^{J \times N}$ containing all the filters: $W = [f_1, f_2, \dots, f_N]$, where $f_n \in \mathbb{R}^J$ is the vector containing all the parameters in the n^{th} filter, and $J = K^2 M$ for a standard convolutional layer. The soft-orthogonality (SO) constraint on a layer l with the corresponding 2D matrix W_l is defined as:

$$L_{SO} = \lambda \|W_l^T W_l - I\|_F^2, \quad (14)$$

where $I \in \mathbb{R}^{N \times N}$ is the identity matrix, λ controls the degree of orthogonality and $\|\cdot\|_F^2$ is the Frobenius norm.

However, note that the columns of W can only be mutually orthogonal if W_l is an undercomplete or a square matrix ($J \geq N$), which may not be the case in practice. For overcomplete settings ($J < N$), $W_l^T W_l$ can be far from the identity, since the rank of W_l will be upper-bounded by J . This makes the optimization in (14) a biased objective. To alleviate this issue, double soft orthogonality (DSO) regularization has been adopted [5], which covers both overcomplete and undercomplete cases:

$$L_{DSO} = \lambda (||W_l^T W_l - I||_F^2 + ||W_l W_l^T - I||_F^2). \quad (15)$$

Both the constraints L_{SO} and L_{DSO} are commonly used regularization techniques to encourage filter diversity. While the above regularization methods provide a reasonable result by decreasing the correlation across filters, they have several limitations: (1) The objectives in (14-15) are computationally expensive and have to be computed over all the layers of the convolutional network. (2) The above objectives do not enforce uncorrelated filters but only encourage filter diversity, making such regularization dependent on the dataset complexity. For example, a dataset with fewer training examples may lead to a large number of redundant filters compared to a complex dataset having many training examples.

In contrast, the SSC filters induce group-wise “near-orthogonality”. In particular, for each set $[a, a+q)$ as defined in (9), the q filters are implicitly pairwise orthogonal as they operate at non-overlapping regions of the input feature maps. In particular, the odd/even kernels combined with the shift operation lead to implicit orthogonality. This can be easily visualized from Figure 2. The only source of potential redundancy in SSC filters could arise from the 1×1 kernels, whose receptive field may overlap with odd/even kernels. However, as we show in Section 4.3, the average pairwise correlation is significantly lower compared to the correlation in a standard convolutional layer. Moreover, a key advantage of using SSC filters is that we do not require expensive explicit layer-wise regularization objectives (14) and (15). The reduction in the correlation among filters is a byproduct of the sparsity enforced in the SSC filters, which is complemented by low computational and memory costs as shown in (11) and (13). We provide a comparison of the correlation obtained from the proposed filter with the correlation obtained in a standard model with SO and DSO constraints in the experiments.

2.2. Connection to other efficient filters

The existing efficient filters used for groupwise convolution (GWC), depthwise convolution (DWC) and pointwise convolution (PWC) can be seen as special cases of the proposed SSC filters. The GWC, DWC and PWC have proven to be essential components for the state-of-the-art architectures designed for low-end computational devices, such as mobile devices [40, 35].

For convenience, we denote the hyperparameter setting $g = 0$ to the case where no odd/even kernel is included in the convolutional layer, *i.e.*, there is no sparsity within the kernel. Similarly, $p = 0$ denotes the setting where no 1×1 kernel is included. If we set $p = 0$, the SSC filters can be used for GWC, where each filter will be of size $K \times K \times M/g$ and the kernels operate on a group of feature maps that are separated by an interval of g . The DWC operation can be achieved by using a SSC filter with $p = 0$ and $g = M$, where each filter will act on a single input feature map. The PWC operation can be achieved with the SSC filter by setting $g = 0$ and $p = 1$. While the SSC operation provides a generic framework that can be reduced to GWC, DWC and PWC, the SSC filters are inherently sparse. One can get the standard (non-sparse) version of these efficient filters by using the standard $K \times K$ kernels instead of odd/even kernels defined in Section 2.

Current state-of-the-art efficient architectures MobileNetv2 [40] and ShuffleNetv2 [35] reduce the number of parameters by using a two-layer sequential step: DWC followed by a PWC operation. While these operations are less computationally expensive compared to the standard convolutional layer, the latency of the model increases due to the two-layer sequential step. In contrast, the proposed SSC filter can achieve the two operations (GWC/DWC followed by PWC) using a single layer without increasing the number of parameters. We can perform a composition of GWC and PWC by using $p = 1, g > 0$. Similarly, a composition of DWC and PWC can be done using $p = 1$ and $g = M$. In Table 5 we show the results of above composition.

3. Related Work

The area of model compression has seen immense progress in recent years. Below, we highlight some recently proposed techniques for learning sparse neural networks. We divide the model compression and pruning techniques into three categories based on the computational cost involved.

Pruning After Training (PAT) [47, 18, 19, 15, 17, 55, 60, 46, 44] is the most widely used pruning method. [19] proposed pruning the weight parameters using iterative thresholding. Knowledge distillation based methods [41, 24, 8] attempt to train a compressed student model that mimics the behaviour of the full-sized original model. [23] and [10] proposed low-rank approximation of the weight tensors to reduce the memory and time complexity at training and testing time. These approaches tend to accumulate errors in the prediction when multiple layers are compressed. PAT-based pruning methods are typically costly and time-consuming, since they require pretraining of the original overparameterized model.

Dynamic Pruning (DP) [31, 54, 37, 36] involves pruning and training the model simultaneously; as the training progresses the model size decreases. Soft Filter Pruning

(SFP) [22] prunes the filter after each epoch, but updates the pruned filters when training the model. Deep Rewiring (DeepR) [6] and Dynamic Sparse [37] prune and re-grow the architecture periodically, but are more computationally expensive when pruning large networks. Sparse momentum (SM) [11] also follows the prune and re-grow approach, however, it uses smoothed gradients to accelerate the training. These approaches can be trained more efficiently compared to PAT-based methods, since a model trained with DP shrinks with training.

Pruning Before Training (PBT) is the most challenging and yet the most practically useful setting among the three categories. Despite its importance, there have been only a few attempts that explore this setting. These approaches find a small sub-network before the training begins and thus require less computational resources and training time. While there have been some attempts to prune the deep neural network before the training [27, 28, 52, 56, 45], they still require multiple forward/backward passes through the full model to detect unimportant connections in the network. Single-Shot Network Pruning (SNIP) [27] tries to identify a sparse sub-network by solving a minimization problem that preserves the loss of the pruned network at initialization. [28] studied the pruning problem from a signal propagation perspective and proposed using an orthogonal initialization to ensure faithful signal propagation. Recently, [52] proposed Gradient Signal Preservation (GraSP) to prune the network at initialization by preserving gradient flow. Interspace Pruning (IP) [56] was recently proposed to overcome the bias introduced in PBT methods, which as a result improves the generalization of existing unstructured pruning methods (including SNIP and GraSP). Although these methods are developed to find a sub-network at initialization, they still require optimizing the original overparameterized model using the training dataset, which can be expensive for low-end devices. Moreover, the sub-networks found are specific to a particular dataset, hindering knowledge transfer across multiple tasks. Synaptic Flow (SynFlow) [49] prunes weights using a information throughput criterion to find a sparse network. Our proposed method also falls into the PBT category, where the sparse network is identified at the initialization. In contrast to previous methods, the proposed method is sparse by design and does not require solving an optimization problem to find a task-specific sub-network.

4. Experiments

We demonstrate the efficacy and efficiency of our proposed filter with an extensive set of experiments. First, we evaluate the performance of commonly used deep convolutional networks (ResNet-32/50 and VGG-19 [21, 43]) with SSC filters on four classification benchmarks: CIFAR-10/100 [25], Tiny-ImageNet [57] and ImageNet [9]. We show that the proposed SSC filters achieve state-of-the-art accuracy in

most settings. We also apply SSC filters to existing state-of-the-art “efficient” architectures: MobileNetV2 [40] and ShuffleNetV2 [35]. In Section 4.2, we show that these architectures can be further compressed by 47-48% with SSC, while achieving high accuracy on the CIFAR-10 benchmark.

Next, we conduct experiments to analyze the properties of SSC filters. Specifically, we find that SSC leads to significantly lower layer-wise filter correlation, implying fewer redundant filters compared to alternative methods that use explicit regularizers. We also test the robustness of SSC to overfitting when trained with limited training data. Finally, we evaluate the ability of SSC filters in transfer learning.

For all the experiments, we follow the training setup (optimizer, learning rate, training epochs) used in [52]. We include all hyperparameters in the supplementary material. In addition to the aforementioned experiments, we carry out extensive ablations on choosing the hyperparameters g and p . In general, the hyperparameters p and g are set to achieve the desired sparsity in the architectures; we use R_p (as defined in 13) as our guiding principle behind choosing p and g . For brevity, we defer the ablation studies and further discussion to the supplementary.

Baselines: We include a number of baselines to compare the performance of the proposed approach. The baselines include methods that prune the architecture after complete training, such as OBD [20], MLPrune [63] and LT [15]. We also consider methods that do dynamic pruning: DSR [37], SET [36] and Deep-R [6]. From the class of methods that do pruning before training (PBT), we consider SNIP [27], GraSP [52], and SynFlow [49]. All the above baselines, except PBT methods, have an added advantage of training a large overparameterized network compared to our proposed method that trains on a highly-sparse architecture from initialization. While we report performance of all the above baselines, a fair comparison of our method can only be done with SynFlow, SNIP and GraSP. We also consider the improvements introduced with IP [56] for PBT methods.

4.1. Performance on Classification benchmarks

We evaluate the proposed method on the CIFAR10 and CIFAR100 classification benchmark by training commonly used architectures (VGG-19 and ResNet-32) with the standard convolution filters replaced by the proposed SSC filters. We report the test accuracy using VGG-19 and ResNet-32 under three pruning ratios, namely 90%, 95% and 98% in Table 1. We observed that SSC performs better than SNIP and GraSP in 8 of the 12 different settings considered, especially when we use ResNet32 (outperforming in 5 of 6 settings). Note that even in the settings where SSC is inferior to SNIP or GraSP, the performance is highly competitive to the best method. Moreover, in the extreme setting of 98% sparsity on the CIFAR-100 dataset, SSC outperforms the next best

Table 1. Test accuracy of pruned VGG19 and ResNet32 on CIFAR-10 and CIFAR-100 datasets. The bold number is the higher one between the accuracy of GraSP, SNIP and SSC.

Dataset	CIFAR-10			CIFAR-100		
Pruning ratio	90%	95%	98%	90%	95%	98%
VGG19 (Unpruned)	94.23	-	-	74.16	-	-
OBD [26]	93.74	93.58	93.49	73.83	71.98	67.79
MLPrune [63]	93.83	93.69	93.49	73.79	73.07	71.69
LT (original initialization)	93.51	92.92	92.34	72.78	71.44	68.95
LT (reset to epoch 5)	93.82	93.61	93.09	74.06	72.87	70.55
DSR [37]	93.75	93.86	93.13	72.31	71.98	70.70
SET [36]	92.46	91.73	89.18	72.36	69.81	65.94
Deep-R [6]	90.81	89.59	86.77	66.83	63.46	59.58
SNIP [27]	93.63±0.06	93.43±0.20	92.05±0.28	72.84±0.22	71.83±0.23	58.46±1.10
GraSP [52]	93.30±0.14	93.04±0.18	92.19±0.12	71.95±0.18	71.23±0.12	68.90±0.47
SynFlow [49]	92.99±0.18	92.23±0.15	91.01±0.17	-.-	-.-	-.-
IP-SynFlow [56]	93.17±0.20	92.46±0.14	92.11±0.24	-.-	-.-	-.-
SSC (Ours)	93.68±0.11	93.45±0.17	91.16±0.20	71.65±0.18	70.71±0.21	68.95±0.44
ResNet32 (Unpruned)	94.80	-	-	74.64	-	-
OBD [26]	94.17	93.29	90.31	71.96	68.73	60.65
MLPrune [63]	94.21	93.02	89.65	72.34	67.58	59.02
LT (original initialization)	92.31	91.06	88.78	68.99	65.02	57.37
LT (reset to epoch 5)	93.97	92.46	89.18	71.43	67.28	58.95
DSR [37]	92.97	91.61	88.46	69.63	68.20	61.24
SET [36]	92.30	90.76	88.29	69.66	67.41	62.25
Deep-R [6]	91.62	89.84	86.45	66.78	63.90	58.47
SNIP [27]	92.59±0.10	91.01±0.21	87.51±0.31	68.89±0.45	65.22±0.69	54.81±1.43
GraSP [52]	92.38±0.21	91.39±0.25	88.81±0.14	69.24±0.24	66.50±0.11	58.43±0.43
SSC (Ours)	93.29±0.20	91.66±0.13	88.39±0.24	71.13±0.12	67.31±0.15	63.01±0.31

Table 2. Test accuracy of pruned VGG19 and ResNet32 on Tiny-ImageNet dataset. The bold number is the higher one between the accuracy of GraSP and that of SNIP.

Network	VGG19			ResNet32		
Pruning ratio	90%	95%	98%	85%	90%	95%
VGG19/ResNet32 (Unpruned)	61.38	-	-	62.89	-	-
OBD [26]	61.21	60.49	54.98	58.55	56.80	51.00
MLPrune [63]	60.23	59.23	55.55	58.86	57.62	51.70
LT (original initialization)	60.32	59.48	55.12	56.52	54.27	49.47
LT (reset to epoch 5)	61.19	60.57	56.18	60.31	57.77	51.21
DSR [37]	62.43	59.81	58.36	57.08	57.19	56.08
SET [36]	62.49	59.42	56.22	57.02	56.92	56.18
Deep-R [6]	55.64	52.93	49.32	53.29	52.62	52.00
SNIP [27]	61.02±0.41	59.27±0.39	48.95±1.73	56.33±0.24	55.43±0.14	49.57±0.44
GraSP [52]	60.76±0.23	59.50±0.33	57.28±0.34	57.25±0.11	55.53±0.11	51.34±0.29
SSC (Ours)	61.53±0.25	59.72±0.23	56.81±0.39	57.92±0.16	57.36±0.19	53.70±0.21

method using ResNet32 with a 7.8% relative improvement.

Tiny-ImageNet is a medium-scale dataset containing images of 200 classes from the ImageNet. Again, we chose VGG-19 and ResNet-32 as the base architectures with varying pruning ratio. The results are reported in the Table 2. Once again, we observe that our approach outperforms the baselines methods SNIP and GraSP in 5 of the 6 settings. In the extreme setting of 95% sparsity in ResNet-32, we

observe that the proposed method shows a relative improvement of 4.6% (absolute improvement of 2.36%) over GraSP.

We also conduct a large scale experiment on the ImageNet dataset using the ResNet-50 architecture. Results are shown in Table 3. The model is trained under two pruning ratios: 60% and 80%. We report the top-1 and top-5 accuracy using SNIP, GraSP and our proposed method. Despite the added advantage in SNIP and GraSP that use memory intensive

Table 3. The test accuracy using ResNet-50 architecture on the ImageNet benchmark with pruning ratios 60% and 80%. The baselines SNIP and GraSP have additional overhead compared to SSC. Nevertheless, SSC performs comparably/better than SNIP and GraSP.

Pruning ratios	60%		80%	
Accuracy	top-1	top-5	top-1	top5
ResNet-50 (Unpruned)	75.70	92.81	-	-
SNIP [27]	73.95	91.97	69.67	89.24
GraSP [52]	74.02	91.86	72.06	90.82
SSC (Ours)	73.76	91.83	71.27	90.32

Table 4. The test accuracy using ResNet-18 architecture on the ImageNet benchmark with pruning ratios 60% and 75%. All the baselines have additional overhead compared to SSC.

Pruning ratios	60%		75%	
Accuracy	top-1	top-5	top-1	top5
ResNet-18 (Unpruned)	69.77	89.07	-	-
SNIP [27]	66.56	87.32	64.33	85.74
GraSP [52]	65.87	86.71	63.57	85.34
SynFlow [49]	66.93	87.56	63.85	85.61
IP-SNIP [56]	67.29	87.65	64.93	86.18
IP-GraSP [56]	66.92	87.48	64.94	86.20
IP-SynFlow [56]	67.31	87.85	64.63	86.11
SSC (Ours)	67.55	87.89	64.86	85.83

Table 5. Compression of MobileNet and ShuffleNet architectures on the CIFAR-10 benchmark. Note that baselines require pruning after training, which is more computationally expensive than SSC.

	Method	Param(%)	Acc (%)
MobileNetV2	Unpruned	100.0	94.3
	L_1P [29]	51.3	84.4
	Slimming [32]	43.8	91.5
	AutoSlim [61]	46.0	91.9
	SSC (Ours)	52.4	92.9
ShuffleNetV2	Unpruned	100.0	91.0
	Slimming [32]	55.8	89.5
	SSC (Ours)	52.2	89.7

“foresight pruning” on the large ResNet-50 architecture, our proposed SSC performs comparably. For the 60% pruning ratio, the top-1 accuracy of SSC only lags behind by 0.26% and the top-5 by 0.14% in absolute difference. This makes SSC based deep convolutional models appealing for devices that cannot train large models as SSC does not require iterative pruning on a large model. We also conducted the ImageNet experiment over the ResNet18 architecture and the results are shown in the Table 4. In spite of the additional overhead used in baseline methods, we find SSC to be better or comparable to the baselines.

4.2. Compression of the efficient Architecture

We apply the proposed SSC filter to the MobileNetV2 [40] and ShuffleNetV2 [35] models, which are the state-of-the-art architectures for low-end devices. As discussed in Section 2.2, we replace the DWC and PWC operations with the proposed SSC operation, allowing us to further compress the efficient architectures [40, 35]. We set the hyperparameters as follows: (1) For all PWC filters, we fix $p = 2$ and $g = 0$. This reduces the total number of PWC parameters in the standard architecture by 50%. (2) For the DWC filters, we set $p = 0$ and $g = M$. We do not use the odd/even kernel in the SSC filter as there is only a single kernel in DWC filters. We report the results in Table 5. We consider L_1P [29], Slimming [32] and AutoSlim [61] as the baseline models. Unlike the proposed method, all three baselines require a pretrained model to prune the model. Despite this advantage over SSC, the SSC filter shows significant improvement over the baselines for the MobileNetV2 and ShuffleNetV2 architectures.

4.3. Parameter Redundancy

In Section 2.1, we discussed the implicit orthogonal nature of the SSC filters. We compare the average absolute correlation at each layer of the ResNet-32 architecture trained with and without using the SSC filters. The pairwise correlation (ρ) between the i^{th} and j^{th} filter can be calculated as: $\rho(f_i, f_j) = \mathbb{E}[(f_i - \mu_i)(f_j - \mu_j)] / (\sigma_i \sigma_j)^{1/2}$, where $f_i \in \mathbb{R}^{K^2 M}$ represents the 1D vector of all the parameters in the i^{th} filter and μ_i is the mean of all the parameters in f_i and $\sigma_i = \sum_{i=1}^{K^2 M} (f_i - \mu_i)$. We define the average absolute correlation at layer l : $C^l = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N |\rho(f_i^l, f_j^l)|$. In Figure 3, we report the correlation measure C^l at each layer of the ResNet-32 architecture trained on the CIFAR-100 benchmark. We also compare the correlation measures obtained when a standard ResNet-32 model is trained with explicit orthogonal constraints described in (14) and (15).

As seen in Figure 3, the correlation measure of the standard convolutional filters is higher compared to other methods, implying filter redundancy in the standard ResNet-32 architecture. We observe that for 31 layers out of the total 33, the SSC filters have lower correlation measure compared to standard filters. In comparison to methods with explicit regularization constraints, the SSC filters have lower correlation measure for 21 layers. This shows that SSC-based ResNet architecture can utilize parameters more effectively by learning diverse filters compared to the baselines. Moreover, the SSC filters circumvent the under/over-complete issues (Section 2.1), while avoiding expensive computations (14-15).

4.4. Robustness to overfitting

An overparameterized deep learning model is known to overfit in a low-data regime. As shown empirically in Section 4.3,

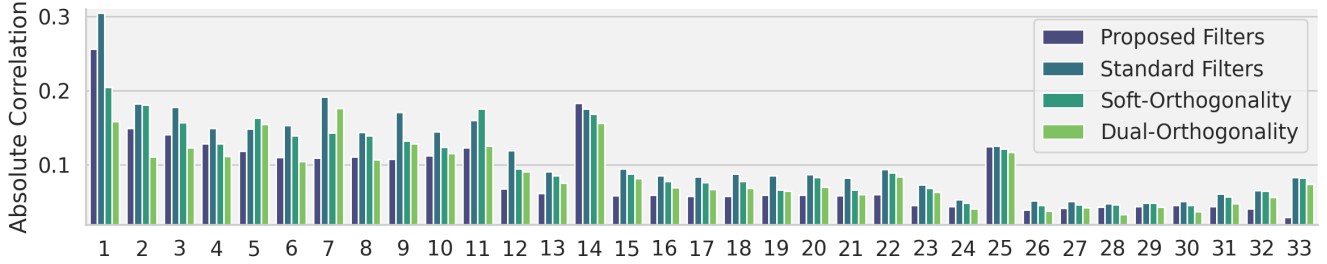


Figure 3. The filter correlation (Y axis) in ResNet-32 (X axis: Layer) for the proposed SSC (10% parameters) and baselines.

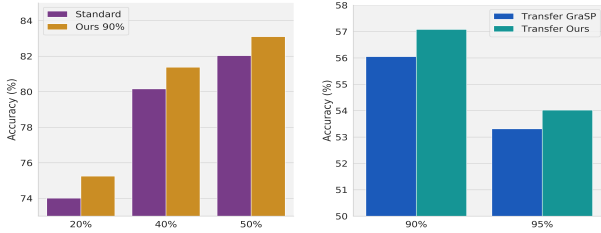


Figure 4. **Left:** Accuracy with ResNet-32 on the CIFAR-10 dataset when only 20%, 40% and 50% samples are used. **Right:** Accuracy of sparse networks when transferred to a different task.

the proposed SSC filter can learn diverse filters. In this section, we analyze the robustness of the SSC filters when the training data is scarce. To test if the SSC filters are robust to overfitting, we train the ResNet-32 architecture on partial training data. We consider three scenarios that involve learning the model using only 20%, 40% and 50% images of the total training set. The performance of the three scenarios is reported in Figure 4 (Left). The hyperparameters g and p are chosen such that the neural network has only 10% of the total parameters in a standard ResNet-32 architecture. As expected, we observe that the model with the SSC filters are more robust to overfitting when compared to the model with the standard convolutional filters. In our experiments, we also observed that the difference in performance becomes more significant when the training data decreases.

4.5. Transfer Learning

In this section, we study the generalization of the sparse architectures on two different tasks. Recall from Section 2 that the pruning in SSC is governed by choosing the hyperparameters g and p ; thus, the obtained sparsity in SSC is independent of the training data. Whereas for the baseline methods, the sparse network is found by iterative procedures performed on the training data. We compare the transfer ability of the SSC architecture with the best performing baseline (*i.e.*, GraSP). To study the transfer ability across two tasks, we first train GraSP on CIFAR-100 benchmark to prune the original architecture and then transfer (fine-tune) the obtained sparse network on the Tiny-ImageNet dataset. Similarly, for the proposed SSC-based architecture, we first

find the sparse network by tuning the hyperparameters g and p on the CIFAR-100 benchmark, and then use the obtained sparse network for training on the Tiny-ImageNet dataset. To keep the complexity of the Tiny-ImageNet task same as CIFAR-100, we only use 100 classes from the Tiny-ImageNet dataset. The performance of the two methods (GraSP and Ours) under pruning ratios of 90% and 95% is shown in Figure 4 (right). We observe that the SSC-based architectures performs better than GraSP, depicting better transfer ability.

5. Conclusions

We have proposed structured sparse convolutional (SSC) for deep convolutional neural networks. SSC is based on efficient filters composed of novel odd/even kernels and 1×1 kernels. The proposed kernels leverage the spatial dependencies in the input features to reduce the floating-point operations and the number of parameters in the deep neural network from initialization. Through a series of experiments, we demonstrate the efficacy of the proposed SSC when applied to commonly used deep convolutional networks. A key attribute of SSC filters is that unlike existing approaches, SSC requires no additional pruning during/after training. We also show that the SSC filters generalize other efficient filters (GWC, DWC, and PWC) and demonstrate the applicability of SSC-filters on existing efficient architectures like MobileNet and ShuffleNet.

While this work has demonstrated potential in designing sparse filters without additional steps, our proposed filter has limited practical speedup as existing deep learning libraries do not support efficient operations on structured tensors. However, an efficient CUDA implementation to incorporate such operations for structured sparsity can easily solve this issue. We strongly believe that the proposed method could be highly beneficial to the broader community in training powerful deep systems with only a fraction of computational resources. Regardless, there is still little that we understand about identifying good sparse models that can generalize. We hope that this work will motivate the research community to identify efficient filters that can naturally lead to highly-sparse and effective deep learning models.

References

- [1] Zeyuan Allen-Zhu, Yuanzhi Li, and Yingyu Liang. Learning and generalization in overparameterized neural networks, going beyond two layers. In *Advances in neural information processing systems*, pages 6158–6169, 2019.
- [2] Zeyuan Allen-Zhu, Yuanzhi Li, and Zhao Song. A convergence theory for deep learning via over-parameterization. In *International Conference on Machine Learning*, pages 242–252. PMLR, 2019.
- [3] Jose M Alvarez and Mathieu Salzmann. Compression-aware training of deep networks. In *Advances in Neural Information Processing Systems*, pages 856–867, 2017.
- [4] Sanjeev Arora, Simon S Du, Wei Hu, Zhiyuan Li, and Ruosong Wang. Fine-grained analysis of optimization and generalization for overparameterized two-layer neural networks. In *Proceedings of the 36th International Conference on Machine Learning*, pages 322–332, 2019.
- [5] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep networks? In *Advances in Neural Information Processing Systems*, pages 4261–4271, 2018.
- [6] Guillaume Bellec, David Kappel, Wolfgang Maass, and Robert Legenstein. Deep rewiring: Training very sparse deep networks. In *International Conference on Learning Representations*, 2018.
- [7] Changan Chen, Frederick Tung, Naveen Vedula, and Greg Mori. Constraint-aware deep neural network compression. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 400–415, 2018.
- [8] Yen-Chun Chen, Zhe Gan, Yu Cheng, Jingzhou Liu, and Jingjing Liu. Distilling knowledge learned in bert for text generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7893–7905, 2020.
- [9] 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, 2009.
- [10] Emily L Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. In *Advances in neural information processing systems*, pages 1269–1277, 2014.
- [11] Tim Dettmers and Luke Zettlemoyer. Sparse networks from scratch: Faster training without losing performance. 2019.
- [12] Xiaohan Ding, Guiguang Ding, Jungong Han, and Sheng Tang. Auto-balanced filter pruning for efficient convolutional neural networks. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- [13] Simon S Du and Jason D Lee. On the power of over-parametrization in neural networks with quadratic activation. *arXiv preprint arXiv:1803.01206*, 2018.
- [14] Simon S Du, Xiyu Zhai, Barnabas Poczos, and Aarti Singh. Gradient descent provably optimizes over-parameterized neural networks. *arXiv preprint arXiv:1810.02054*, 2018.
- [15] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *International Conference on Learning Representations*, 2019.
- [16] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carbin. Stabilizing the lottery ticket hypothesis. *arXiv preprint arXiv:1903.01611*, 2019.
- [17] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carbin. The lottery ticket hypothesis at scale. 2020.
- [18] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations*, 2015.
- [19] 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.
- [20] Babak Hassibi, David G Stork, and Gregory J Wolff. Optimal brain surgeon and general network pruning. In *IEEE international conference on neural networks*, pages 293–299. IEEE, 1993.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. *Computer Vision and Pattern Recognition*, 2016.
- [22] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence, IJCAI’18*, pages 2234–2240. AAAI Press, 2018.
- [23] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. Speeding up convolutional neural networks with low rank expansions. In *BMVC*, 2014.
- [24] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.
- [25] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009.
- [26] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pages 598–605, 1990.
- [27] Namhoon Lee, Ajanthan, Thalaiyasingam, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. *International Conference on Learning Representations*, 2019.
- [28] Namhoon Lee, Thalaiyasingam Ajanthan, Stephen Gould, and Philip H. S. Torr. A signal propagation perspective for pruning neural networks at initialization. In *International Conference on Learning Representations*, 2020.
- [29] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *ICLR*, 2016.
- [30] Yuanzhi Li and Yingyu Liang. Learning overparameterized neural networks via stochastic gradient descent on structured data. In *Advances in Neural Information Processing Systems*, pages 8157–8166, 2018.

- [31] Tao Lin, Sebastian U Stich, Luis Barba, Daniil Dmitriev, and Martin Jaggi. Dynamic model pruning with feedback. *ICLR*, 2020.
- [32] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2736–2744, 2017.
- [33] Christos Louizos, Karen Ullrich, and Max Welling. Bayesian compression for deep learning. In *Advances in neural information processing systems*, pages 3288–3298, 2017.
- [34] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pages 5058–5066, 2017.
- [35] 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.
- [36] Decebal Constantin Mocanu, Elena Mocanu, Peter Stone, Phuong H Nguyen, Madeleine Gibescu, and Antonio Liotta. Scalable training of artificial neural networks with adaptive sparse connectivity inspired by network science. *Nature communications*, 9(1):2383, 2018.
- [37] Hesham Mostafa and Xin Wang. Parameter efficient training of deep convolutional neural networks by dynamic sparse reparameterization. *International Conference on Machine Learning*, 2019.
- [38] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. The role of over-parametrization in generalization of neural networks. In *International Conference on Learning Representations*, 2018.
- [39] Behnam Neyshabur, Ryota Tomioka, and Nathan Srebro. In search of the real inductive bias: On the role of implicit regularization in deep learning. *arXiv preprint arXiv:1412.6614*, 2014.
- [40] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [41] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [42] Wenling Shang, Kihyuk Sohn, Diogo Almeida, and Honglak Lee. Understanding and improving convolutional neural networks via concatenated rectified linear units, 2016.
- [43] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *International Conference on Learning Representations*, 2014.
- [44] Pravendra Singh, Vinay Kumar Verma, Piyush Rai, and Vinay Namboodiri. Leveraging filter correlations for deep model compression. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 835–844, 2020.
- [45] Pravendra Singh, Vinay Kumar Verma, Piyush Rai, and Vinay P Namboodiri. Hetconv: Heterogeneous kernel-based convolutions for deep cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4835–4844, 2019.
- [46] Pravendra Singh, Vinay Kumar Verma, Piyush Rai, and Vinay P Namboodiri. Play and prune: adaptive filter pruning for deep model compression. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence*, pages 3460–3466, 2019.
- [47] Suraj Srinivas and R Venkatesh Babu. Data-free parameter pruning for deep neural networks. *arXiv preprint arXiv:1507.06149*, 2015.
- [48] C. Szegedy, Wei Liu, Yangqing Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, 2015.
- [49] Hidenori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 33:6377–6389, 2020.
- [50] Vincent Vanhoucke. Data-free parameter pruning for deep neural networks. *ICLR invited talk*, 2014.
- [51] Vinay Kumar Verma, Pravendra Singh, Vinay Namboodiri, and Piyush Rai. A” network pruning network” approach to deep model compression. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 3009–3018, 2020.
- [52] Chaoqi Wang, Zhang, Guodong, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. *ICLR*, 2020.
- [53] Jiayun Wang, Yubei Chen, Rudras Chakraborty, and Stella X Yu. Orthogonal convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11505–11515, 2020.
- [54] Yulong Wang, Xiaolu Zhang, Xiaolin Hu, Bo Zhang, and Hang Su. Dynamic network pruning with interpretable layer-wise channel selection. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6299–6306, 2020.
- [55] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *NIPS*, pages 2074–2082, 2016.
- [56] Paul Wimmer, Jens Mehnert, and Alexandru Condurache. Interspace pruning: Using adaptive filter representations to improve training of sparse cnns. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12527–12537, 2022.
- [57] Jiayu Wu, Qixiang Zhang, and Guoxi Xu. Tiny imagenet challenge. Technical report.
- [58] Di Xie, Jiang Xiong, and Shiliang Pu. All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6176–6185, 2017.
- [59] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on*

- computer vision and pattern recognition*, pages 1492–1500, 2017.
- [60] Jaehong Yoon and Sung Ju Hwang. Combined group and exclusive sparsity for deep neural networks. volume 70 of *Proceedings of Machine Learning Research*, pages 3958–3966, International Convention Centre, Sydney, Australia, 06–11 Aug 2017. PMLR.
 - [61] Jiahui Yu and Thomas Huang. Autoslim: Towards one-shot architecture search for channel numbers. *ICLR*, 2019.
 - [62] 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.
 - [63] Wenyuan Zeng and Raquel Urtasun. MLPrune: Multi-layer pruning for automated neural network compression, 2019.
 - [64] 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.
 - [65] Difan Zou, Yuan Cao, Dongruo Zhou, and Quanquan Gu. Gradient descent optimizes over-parameterized deep relu networks. *Machine Learning*, 109(3):467–492, 2020.