

# DHP: Differentiable Meta Pruning via HyperNetworks

Yawei Li<sup>1</sup>, Shuhang Gu<sup>1,2</sup>, Kai Zhang<sup>1</sup>, Luc Van Gool<sup>1,3</sup>, Radu Timofte<sup>1</sup>

<sup>1</sup>Computer Vision Lab, ETH Zürich, <sup>2</sup>The University of Sydney, <sup>3</sup>KU Leuven  
{yawei.li, kai.zhang, vangool, radu.timofte}@vision.ee.ethz.ch  
shuhangu@gmail.com

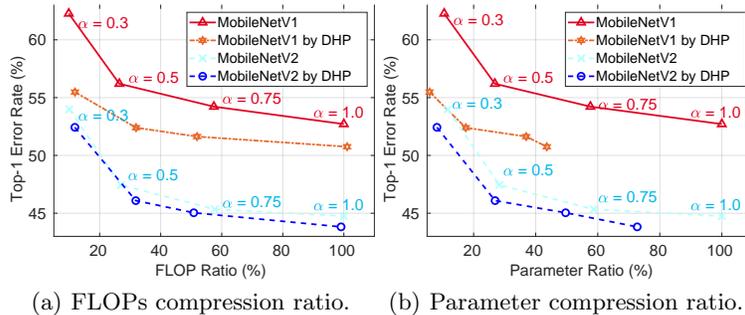
**Abstract.** Network pruning has been the driving force for the acceleration of neural networks and the alleviation of model storage/transmission burden. With the advent of AutoML and neural architecture search (NAS), pruning has become topical with automatic mechanism and searching based architecture optimization. Yet, current automatic designs rely on either reinforcement learning or evolutionary algorithm. Due to the non-differentiability of those algorithms, the pruning algorithm needs a long searching stage before reaching the convergence.

To circumvent this problem, this paper introduces a differentiable pruning method via hypernetworks for automatic network pruning. The specifically designed hypernetworks take latent vectors as input and generate the weight parameters of the backbone network. The latent vectors control the output channels of the convolutional layers in the backbone network and act as a handle for the pruning of the layers. By enforcing  $\ell_1$  sparsity regularization to the latent vectors and utilizing proximal gradient solver, sparse latent vectors can be obtained. Passing the sparsified latent vectors through the hypernetworks, the corresponding slices of the generated weight parameters can be removed, achieving the effect of network pruning. The latent vectors of all the layers are pruned together, resulting in an automatic layer configuration. Extensive experiments are conducted on various networks for image classification, single image super-resolution, and denoising. And the experimental results validate the proposed method. Code is available at <https://github.com/ofsoundof/dhp>.

**Keywords:** Network pruning, hypernetworks, meta learning, differentiable optimization, proximal gradient.

## 1 Introduction

These days, network pruning has become the workhorse for network compression, which aims at lightweight and efficient model for fast inference [12, 18, 38, 29]. This is of particular importance for the deployment of tiny artificial intelligence (Tiny AI) algorithms on smart phones and edge devices [1]. Since the emerging of network pruning a couple of methods have been proposed based on the analysis of gradients, Hessians or filter distribution [25, 13, 9, 40, 16]. With

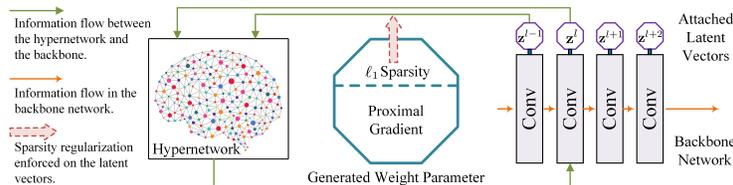


**Fig. 1.** Top-1 error *vs.* FLOPs and parameter compression ratio on MobileNets. The original model with different width multipliers  $\alpha$  is set as the baseline. The DHP operating points near 100% FLOPs ratio is obtained by pruning the two networks with  $\alpha = 2$ . The DHP models outperforms the original at all the operating points

the advent of AutoML and neural architecture search (NAS) [58, 6], a new trend of network pruning emerges, *i.e.* pruning with automatic algorithms and targeting distinguishing sub-architectures. Among them, reinforcement learning and evolutionary algorithm become the natural choice [17, 37]. The core idea is to search a certain fine-grained layer-wise distinguishing configuration among the all of the possible choices (population in the terminology of evolutionary algorithm). After the searching stage, the candidate that optimizes the network prediction accuracy under constrained budgets is chosen.

The advantage of these automatic pruning methods is the final layer-wise distinguishing configuration. Thus, hand-crafted design is no longer necessary. However, the main concern of these algorithms is the convergence property. For example, reinforcement learning is notorious for its difficulty of convergence under large or even middle level number of states [49]. Evolutionary algorithm needs to choose the best candidate from the already converged algorithm. But the dilemma lies in the impossibility of training the whole population till convergence and the difficulty of choosing the best candidate from unconverged population [37, 14]. A promising solution to this problem is endowing the searching mechanism with differentiability or resorting to an approximately differentiable algorithm. This is due to the fact that differentiability has the potential to make the searching stage efficient. Actually, differentiability has facilitated a couple of machine learning approaches and the typical one among them is NAS. Early works on NAS have insatiable demand for computing resources, consuming tens of thousands of GPU hours for a satisfactory convergence [58, 59]. Differentiable architecture search (DARTS) reduces the insatiable consumption to tens of GPU hours, which has boosted the development of NAS during the past year [36].

Another noteworthy direction for automatic pruning is brought by MetaPruning [37] which introduces hypernetworks [11] into network compression. The output of the so-called hypernetwork is used as the parameters of the backbone network. During training, the gradients are also back-propagated to the



**Fig. 2.** The workflow of the proposed differentiable pruning method. The latent vectors  $\mathbf{z}$  attached to the convolutional layers act as the handle for network pruning. The hypernetwork takes two latent vectors as input and emits output as the weight of the backbone layer.  $\ell_1$  sparsity regularization is enforced on the latent vectors. The differentiability comes with the hypernetwork tailored to pruning and the proximal gradient exploited to solve problem. After the pruning stage, sparse latent vectors are obtained which result in pruned weights after being passed through the hypernetwork

hypernetworks. This method falls in the paradigm of meta learning since the parameters in the hypernetwork act as the meta-data of the parameters in the backbone network. But the problem of this method is that the hypernetworks can only output fixed-size weights, which cannot serve as a layer-wise configuration searching mechanism. Thus, a searching algorithm such as evolutionary algorithm is necessary for the discovery of a good candidate. Although this is quite a natural choice, there is still one interesting question, namely, whether one can design a hypernetwork whose output size depends on the input (termed as latent vector in this paper) so that by only dealing with the latent vector, the backbone network can be automatically pruned.

To solve the aforementioned problem, we propose the differentiable meta pruning approach via hypernetworks (DHP, D – Differentiable, H – Hyper, P – Pruning) shown in Fig. 2. A new design of hypernetwork is proposed to adapt to the requirements of differentiability. Each layer is endowed with a latent vector that controls the output channels of this layer. Since the layers in the network are connected, the latent vector also controls the input channel of the next layer. The hypernetwork takes as input the latent vectors of the current layer and previous layer that controls the output and input channels of the current layer respectively. Passing the latent vectors through the hypernetwork leads to outputs which are used as the parameters of the backbone network. To achieve the effect of automatic pruning,  $\ell_1$  sparsity regularizer is applied to the latent vectors. A pruned model is discovered by updating the latent vectors with proximal gradient. The searching stage stops when the compression ratio drops to the target level. After the searching stage, the latent vectors are sparsified. Accordingly, the outputs of the hypernetworks that are covariant with the latent vector are also compressed. The advantage of the proposed method is that it is only necessary to deal with the latent vectors, which automates network pruning without the other bells and whistles.

With the fast development of efficient network design and NAS, the usefulness of network pruning is frequently challenged. But by analyzing the performance on

MobileNetV1 [19] and MobileNetV2 [47] in Fig. 1, we conclude that automatic network pruning is of vital importance for further exploring the capacity of efficient networks. Efficient network design and NAS can only result in an overall architecture with building blocks endowed with the same sub-architecture. By automatic network pruning, the efficient networks obtained by either human experts or NAS can be further compressed, leading to layer-wise distinguishing configurations, which can be seen as a fine-grained architecture search.

Thus, the contribution of this paper is as follows.

- I A new architecture of hypernetwork is designed. Different from the classical hypernetwork composed of linear layers, the new design is tailored to automatic network pruning. By only operating on the input of the hypernetwork, the backbone network can be pruned.
- II A differentiable automatic networking pruning method is proposed. The differentiability comes with the designed hypernetwork and the utilized proximal gradient. It accelerates the convergence of the pruning algorithm.
- III By the experiments on various vision tasks and modern convolutional neural networks (CNNs) [15, 20, 19, 47, 56, 46, 26, 32], the potential of automatic network pruning as fine-grained architecture search is revealed.

## 2 Related Works

**Network pruning.** Aiming at removing the weak filter connections that have the least influence on the accuracy of the network, network pruning attracts increasing attention. Early attempts emphasize more on the storage consumption, various criteria have been explored to remove inconsequential connections in an unstructural manner [12, 34]. Despite their success in reducing network parameters, unstructural pruning leads to irregular weight parameters, limited in the actual acceleration of the pruned network. To further address the efficiency issue, structured pruning methods directly zero out structured groups of the convolutional filters. For example, Wen *et al.* [52] and Alvarez *et al.* [2] firstly proposed to resort to group sparsity regularization during training to reduce the number of feature maps in each layer. Since that, the field has witnessed a variety of regularization strategies [3, 54, 28, 50, 29]. These elaborately designed regularization methods considerably advance the pruning performance. But they often rely on carefully adjusted hyper-parameters selected for specific network architecture and dataset.

**AutoML.** Recently, there is an emerging trend of exploiting AutoML for automatic network compression [17, 37, 14, 7]. The rationality lies in the exploration among the total population of network configurations for a final best candidate. He *et al.* exploited reinforcement learning agents to prune the networks where hand-crafted design is not longer necessary [17]. Hayashi *et al.* utilized genetic algorithm to enumerate candidate in the designed hypergraph for tensor network decomposition [14]. Liu *et al.* trained a hypernetwork to generate weights of the backbone network and used evolutionary algorithm to search for the best

candidate [37]. The problem of these approaches is that the searching algorithms are not differentiable, which does not result in guaranteed convergence.

**NAS.** NAS automatizes the manual task of neural network architecture design. Optimally, searched networks achieve smaller test error, require fewer parameters and need less computations than their manually designed counterparts [58, 45]. But the main drawback of early strategies is their almost insatiable demand for computational resources. To alleviate the computational burden several methods [59, 35, 36] are proposed to search for a basic building block, *i.e.* cell, opposed to an entire network. Then, stacking multiple cells with equivalent structure but different weights defines a full network [44, 4]. Another recent trend in NAS is differentiable search methods such as DARTS [36]. The differentiability allows the fast convergence of the searching algorithm and thus boosts the fast development of NAS during the past year. In this paper we propose a differentiable counterpart for automatic network pruning.

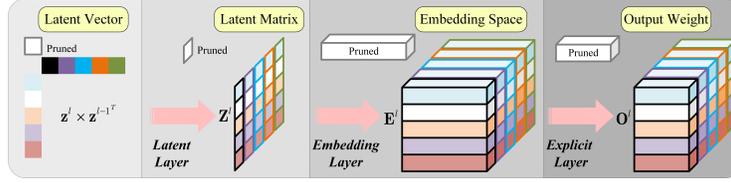
**Meta learning and hypernetworks.** Meta learning is a broad family of machine learning techniques that deal with the problem of learning to learn. An emerging trend of meta learning uses hypernetworks to predict the weight parameters in the backbone network [11]. Since the introduction of hypernetworks, it has found wide applications in NAS [4], multi-task learning [41], Bayesian neural networks [24], and also network pruning [37]. In this paper, we propose a new design of hypernetwork which is especially suitable for network pruning and makes differentiability possible for automatic network pruning.

### 3 Methodology

The pipeline of the proposed method is shown in Fig 2. The two cores of the whole pipeline are the designed hypernetwork and the optimization algorithm. In the forward pass, the designed hypernetwork takes as input the latent vectors and predicts the weight parameters for the backbone network. In the backward pass, the gradients are back-propagated to the hypernetwork. The  $\ell_1$  sparsity regularizer is enforced on the latent vectors and proximal gradient is used to solve the problem. The dimension of the output of the hypernetwork is covariant with that of the input. Due to this property, the output weights are pruned along with the sparsified latent vectors after the optimization step. The differentiability comes with the covariance property of the hypernetworks, the  $\ell_1$  sparsity regularization enforced on the latent vectors, and the proximal gradient used to solve the problem. The automation of pruning is due to the fact that all of the latent vectors are non-discriminatively regularized and that proximal gradient discovers the potential less important elements automatically.

#### 3.1 Hypernetwork design

**Notation:** Unless otherwise stated, we use the normal ( $x$ ), minuscule bold ( $\mathbf{z}$ ), and capital bold ( $\mathbf{Z}$ ) letters to denote scalars, vectors, and matrices/high-dimensional tensors. The elements of a matrix/tensor is indexed by the subscript



**Fig. 3.** Illustration of the hypernetwork designed for network pruning. It generates a weight tensor after passing the input latent vector through the latent layer, the embedding layer, and the explicit layer. If one element in  $\mathbf{z}^l$  is pruned, the corresponding slice of the output tensor is also pruned. See Subsec. 3.1 for details

as  $\mathbf{Z}_{i,j}$  which could be scalars or vectors depending on the the dimension of the indexed subject.

We first introduce the design of the hypernetwork shown in Fig. 3. In summary, the hypernetwork consists of three layers. The latent layer takes as input the latent vectors and computes a latent matrix from them. The embedding layer projects the elements of the latent matrix to an embedding space. The last explicit layer converts the embedded vectors to the final output. This design is inspired by fully connected layers in [11, 37] but differs from those designs in that the output dimension is covariant with the input latent vector. This design is applicable to all types of convolutions including the standard convolution, depth-wise convolution, point-wise convolution, and transposed convolution.

Suppose that the given is an  $L$ -layer CNN. The dimension of the weight parameter of the  $l$ -th convolutional layer is  $n \times c \times w \times h$ , where  $n$ ,  $c$ , and  $w \times h$  denote the output channel, input channel, and kernel size of the layer, respectively. Every layer is endowed with a latent vector  $\mathbf{z}^l$ . The latent vector has the same size as the output channel of the layer, *i.e.*,  $\mathbf{z}^l \in \mathbb{R}^n$ . Thus, the previous layer is given a latent vector  $\mathbf{z}^{l-1} \in \mathbb{R}^c$ . The hypernetwork receives the latent vectors  $\mathbf{z}^l$  and  $\mathbf{z}^{l-1}$  of the current and the previous layer as input. A latent matrix is first computed from the two latent vectors, namely,

$$\mathbf{Z}^l = \mathbf{z}^l \cdot \mathbf{z}^{l-1T} + \mathbf{B}_0^l, \quad (1)$$

where  $[T]$  and  $[\cdot]$  denote matrix transpose and multiplication,  $\mathbf{Z}^l, \mathbf{B}_0^l \in \mathbb{R}^{n \times c}$ . Then every element in the latent matrix is projected to an  $m$  dimensional embedding space, namely,

$$\mathbf{E}_{i,j}^l = \mathbf{Z}_{i,j}^l \mathbf{w}_1^l + \mathbf{b}_1^l, i = 1, \dots, n, j = 1, \dots, c, \quad (2)$$

where  $\mathbf{E}_{i,j}^l, \mathbf{w}_1^l, \mathbf{b}_1^l \in \mathbb{R}^m$ . The vectors  $\mathbf{w}_1^l$  and  $\mathbf{b}_1^l$  are element-wise unique and for the simplicity of notation, the subscript  $i,j$  is omitted.  $\mathbf{w}_1^l, \mathbf{b}_1^l$ , and  $\mathbf{E}_{i,j}^l$  can be aggregated as 3D tensors, namely  $\mathbf{W}_1^l, \mathbf{B}_1^l, \mathbf{E}^l \in \mathbb{R}^{n \times c \times m}$ . After the operation in Eqn. 2, the elements of  $\mathbf{Z}^l$  are converted to embedded vectors in the embedding space. The final step is to obtain the output that can be explicitly used as the weights of the convolutional layer. To achieve that, every embedded vector  $\mathbf{E}_{i,j}^l$

is multiplied by an explicit matrix, that is,

$$\mathbf{O}_{i,j}^l = \mathbf{w}_2^l \cdot \mathbf{E}_{i,j}^l + \mathbf{b}_2^l, i = 1, \dots, n, j = 1, \dots, c, \quad (3)$$

where  $\mathbf{O}_{i,j}^l, \mathbf{b}_2^l \in \mathbb{R}^{wh}$ ,  $\mathbf{w}_2^l \in \mathbb{R}^{wh \times m}$ . Again,  $\mathbf{w}_2^l$  and  $\mathbf{b}_2^l$  are unique for every embedded vector and the subscript  $i,j$  is omitted.  $\mathbf{w}_2^l, \mathbf{b}_2^l$ , and  $\mathbf{O}_{i,j}^l$  can also be aggregated as high-dimensional tensors, *i.e.*  $\mathbf{W}_2^l \in \mathbb{R}^{n \times c \times wh \times m}$  and  $\mathbf{B}_2^l, \mathbf{O}^l \in \mathbb{R}^{n \times c \times wh}$ . For the sake of simplicity, Eqn. 1, 2 and 3 can be abstracted as

$$\mathbf{O}^l = h(\mathbf{z}^l, \mathbf{z}^{l-1}; \mathbf{W}^l, \mathbf{B}^l), \quad (4)$$

where  $h(\cdot)$  denotes the functionality of the hypernetwork. The final output  $\mathbf{O}^l$  is used as the weight parameter of the  $l$ -th layer. The output  $\mathbf{O}^l$  is covariant with the input latent vector because pruning an element in the latent vector removes the corresponding slice of the output  $\mathbf{O}^l$  (See Fig. 3).

When designing the hypernetwork, we tried to add batch normalization and non-linear layers after the linear operation in Eqn. 2 and Eqn. 3. But it did not lead to clearly better results. Thus, we just kept to the simple design. This is also consistent with the previous designs [11, 37].

### 3.2 Sparsity regularization and proximal gradient

The core of differentiability comes with not only the specifically designed hypernetwork but also the mechanism used to search the the potential candidate. To achieve that, we enforce sparsity constraints to the latent vectors. The loss function of the aforementioned  $L$ -layer CNN is denoted as

$$\min_{\mathbf{W}, \mathbf{B}, \mathbf{z}} \mathcal{L}(\mathbf{y}, f(\mathbf{x}; h(\mathbf{z}; \mathbf{W}, \mathbf{B}))) + \gamma \mathcal{D}(\mathbf{W}) + \gamma \mathcal{D}(\mathbf{B}) + \lambda \mathcal{R}(\mathbf{z}), \quad (5)$$

where  $\mathcal{L}(\cdot, \cdot)$ ,  $\mathcal{D}(\cdot)$ , and  $\mathcal{R}(\cdot)$  are the loss function for a specific vision task, the weight decay term, and the sparsity regularization term,  $\gamma$  and  $\lambda$  are the regularization factors. For the simplicity of notation, the superscript  $l$  is omitted. The sparsity regularization takes the form of  $\ell_1$  norm, namely,

$$\mathcal{R}(\mathbf{z}) = \sum_{l=1}^L \|\mathbf{z}^l\|_1. \quad (6)$$

To solve the problem in Eqn. 5, the weights  $\mathbf{W}$  and and biases  $\mathbf{B}$  of the hypernetworks are updated with SGD. The gradients are back-propagated from the backbone network to the hypernetwork. Thus, neither the forward nor the backward pass challenges the information flow between them. The latent vectors are updated with proximal gradient algorithm, *i.e.* ,

$$\mathbf{z}[k+1] = \mathbf{prox}_{\lambda\mu\mathcal{R}}(\mathbf{z}[k] - \lambda\mu\nabla\mathcal{L}(\mathbf{z}[k])), \quad (7)$$

where  $\mu$  is the step size of proximal gradient and is set as the learning rate of SGD updates. As can be seen in the equation, the proximal gradient update contains

a gradient descent step and a proximal operation step. When the regularizer has the form of  $\ell_1$  norm, the proximal operator has closed-form solution, *i.e.*

$$\mathbf{z}[k+1] = \text{sgn}(\mathbf{z}[k+\Delta]) \left[ |\mathbf{z}[k+\Delta]| - \lambda\mu \right]_+, \quad (8)$$

where  $\mathbf{z}[k+\Delta] = \mathbf{z}[k] - \lambda\mu\nabla\mathcal{L}(\mathbf{z}[k])$  is the intermediate SGD update, the sign operator  $\text{sgn}(\cdot)$ , the thresholding operator  $[\cdot]_+$ , and the absolute value operator  $|\cdot|$  act element-wise on the vector. Eqn. 8 is the soft-thresholding function.

The latent vectors first get SGD updates along with the other parameters  $\mathbf{W}$  and  $\mathbf{B}$ . Then the proximal operator is applied. Due to the use of SGD updates and the fact that the proximal operator has closed-form solution, we recognize the whole solution as approximately differentiable (although the  $\ell_1$  norm is not differentiable at 0), which guarantees the fast convergence of the algorithm compared with reinforcement learning and evolutionary algorithm. The speed-up of proximal gradient lies in that instead of searching the best candidate among the total population it forces the solution towards the best sparse one.

The automation of pruning follows the way the sparsity applied in Eqn. 6 and the proximal gradient solution. First of all, all latent vectors are regularized together without distinguishment between them. During the optimization, information and gradients flows fluently between the backbone network and the hypernetwork. The proximal gradient algorithm forces the potential elements of the latent vectors to approach zero quicker than the others without any human effort and interference in this process. The optimization stops immediately when the target compression ratio is reached. In total, there are only two additional hyper-parameters in the algorithm, *i.e.* the sparsity regularization factor and the mask threshold  $\tau$  in Subsec. 3.3. Thus, running the algorithm is just like turning on the button, which enable the application of the algorithm to all of the CNNs without much interference of domain experts' knowledge.

### 3.3 Network pruning

Different from the fully connected layers, the proposed design of hypernetwork can adapt the dimension of the output according to that of the latent vectors. After the searching stage, sparse versions of the latent vectors are derived as  $\bar{\mathbf{z}}^{l-1}$  and  $\bar{\mathbf{z}}^l$ . For those vectors, some of their elements are zero or approaching zero. Thus, 1-0 masks can be derived by comparing the sparse latent vectors with a predefined small threshold  $\tau$ , *i.e.*  $\mathbf{m}^l = \mathcal{T}(\bar{\mathbf{z}}^l, \tau)$ , where the function  $\mathcal{T}(\cdot)$  element-wise compares the latent vector with the threshold and returns 1 if the element is not smaller than  $\tau$  and 0 otherwise. Then latent vector  $\bar{\mathbf{z}}^l$  is pruned according to the mask  $\mathbf{m}^l$ . See Subsec. 3.6 for more analysis.

### 3.4 Latent vector sharing

Due to the existence of skip connections in residual networks such as ResNet, MobileNetV2, SRResNet, and EDSR, the residual blocks are interconnected with

each other in the way that their input and output dimensions are related. Therefore, the skip connections are notoriously tricky to deal with. But back to the design of the proposed hypernetwork, a quite simple and straightforward solution to this problem is to let the hypernetworks of the correlated layers share the same latent vector. Note that the weight and bias parameters of the hypernetworks are not shared. Thus, sharing latent vectors does not force the correlated layers to be identical. By automatically pruning the single latent vector, all of the relevant layers are pruned together. Actually, we first tried to use different latent vectors for the correlated layers and applied group sparsity to them. But the experimental results showed that this is not a good choice because this strategy shot lower accuracy than the latent vector sharing strategy. (See details in the Supplementary).

### 3.5 Discussion on the convergence property

Compared with reinforcement learning and evolutionary algorithm, proximal gradient may not be the optimal solution for some problems. But as found by previous works [38, 37], automatic network pruning serves as an implicit searching method for the channel configuration of a network. In addition, the network is searched and trained from scratch in this paper. The important factor is the number of remaining channels of the convolutional layers in the network. Thus, it is relatively not important which filter is pruned as long as the number of pruned channels are the same. This reduces the number of possible candidates by orders of magnitude. In this case, proximal gradient works quite well.

### 3.6 Implementation consideration

**Compact representation of the hypernetwork.** Thanks to the default tensor operations in deep learning toolboxes [42], the operations in the hypernetwork could be represented in a compact form. The embedding operation in Eqn. 2 can be written as the following high-dimensional tensor operation

$$\mathbf{E}^l = \mathcal{U}^3(\mathbf{Z}^l) \circ \mathbf{W}_1^l + \mathbf{B}_1^l, \quad (9)$$

where  $\mathcal{U}^3(\mathbf{Z}^l) \in \mathbb{R}^{n \times c \times 1}$ ,  $[\circ]$  denotes the broadcastable element-wise tensor multiplication,  $\mathcal{U}^3(\cdot)$  inserts a third dimension for  $\mathbf{Z}^l$ . The operation in Eqn. 3 can be easily rewritten as batched matrix multiplication,

$$\mathbf{O}^l = \mathbf{W}_2^l * \mathbf{E}^l + \mathbf{B}_2^l, \quad (10)$$

where  $[*]$  denotes batched matrix multiplication.

**Pruning analysis.** Analyzing the three layers of the hypernetworks together with the masked latent vectors leads to a direct impression on how the backbone

layers are automatically pruned. That is,

$$\bar{\mathbf{O}}^l = \mathbf{W}_2^l * [\mathcal{U}^3 ((\mathbf{m}^l \circ \bar{\mathbf{z}}^l) \cdot (\mathbf{m}^{l-1} \circ \bar{\mathbf{z}}^{l-1})^T) \circ \mathbf{W}_1^l] \quad (11)$$

$$= \mathbf{W}_2^l * [\mathcal{U}^3(\mathbf{m}^l \cdot \mathbf{m}^{l-1T}) \circ \mathcal{U}^3(\bar{\mathbf{z}}^l \cdot \bar{\mathbf{z}}^{l-1T}) \circ \mathbf{W}_1^l] \quad (12)$$

$$= \mathcal{U}^3(\mathbf{m}^l \cdot \mathbf{m}^{l-1T}) \circ [\mathbf{W}_2^l * (\mathcal{U}^3(\bar{\mathbf{z}}^l \cdot \bar{\mathbf{z}}^{l-1T}) \circ \mathbf{W}_1^l)] \quad (13)$$

The equality follows the broadcastability of the the operations  $[\circ]$  and  $[*]$ . As shown in the above equations, applying the masks on the latent vectors has the same effect of applying them on the final output. Note that in the above analysis the bias terms  $\mathbf{B}_0^l$ ,  $\mathbf{B}_1^l$ , and  $\mathbf{B}_2^l$  are omitted since they have a really small influence on the output of the hypernetwork. In conclusion, the final output can be pruned according to the same criterion for the latent vectors.

**Initialization of the hypernetwork.** All biases are initialized as zero, the latent vector with standard normal distribution, and  $\mathbf{W}_1^l$  with Xavier uniform [10]. The weight of the explicit layer  $\mathbf{W}_2^l$  is initialized with Hyperfan-in which guarantees stable backbone network weights and fast convergence [5].

## 4 Experimental Results

To validate the proposed method, extensive experiments were conducted on various CNN architectures including ResNet [15], DenseNet [20] for CIFAR10 [23] image classification, MobileNetV1 [19], MobileNetV2 [47] for Tiny-ImageNet [8] image classification, SRResNet [26], EDSR [32] for single image super-resolution, and DnCNN [56], UNet [46] for gray image denoising. The proposed DHP algorithm starts from a randomly initialized network with the initialization method detailed in Subsec. 3.6. After pruning, the training of the pruned network continues with the same training protocol used for the original network. All of the experiments are conducted on NVIDIA TITAN Xp GPUs.

**Hyperparameters.** A target FLOPs compression ratio is set for the pruning algorithm. When the difference between the target compression ratio and the actual compression ratio falls below 2%, the automatic pruning procedure stops. The parameter space of hypernetwork increases in proportion to the dimension  $m$  of the embedding space. Thus,  $m$  should not be too large. In this paper,  $m$  is set to 8. The step size  $\mu$  of the proximal operator is set as the learning rate of SGD updates. The sparsity regularization factor  $\lambda$  is set by empirical studies. The value is chosen such that the searching epochs constitute arounds 5% – 10% of the whole training epochs. This guarantees acceptable convergence during searching while not introducing too much additional computation. Please refer to the supplementary for the detailed training and testing protocol.

### 4.1 Image classification

The compression results on image classification networks are shown in Table 1. ‘DHP-\*\*\*’ denotes the proposed method with the target FLOPs ratio during

**Table 1.** Results on image classification networks. The FLOPs ratio and parameter ratio of the pruned networks are reported. DHP outperforms the compared methods under comparable model complexity. On Tiny-ImageNet, DHP-24-2 shoots lower error rates than the original model

Network Top-1 Error (%)	Compression Method	Top-1 Error (%)	FLOPs Ratio (%)	Parameter Ratio (%)
CIFAR10				
ResNet-20 7.46	[53]	9.10	52.60	62.78
	<b>DHP-50 (Ours)</b>	8.46	51.80	56.13
	FPGM [16]	9.38	46.00	–
ResNet-56 7.05	Variational [57]	7.74	79.70	79.51
	Pruned-B [27]	6.94	72.40	86.30
	GAL-0.6 [33]	6.62	63.40	88.20
	NISP [55]	6.99	56.39	57.40
	<b>DHP-50 (Ours)</b>	6.42	50.96	58.42
	CaP [39]	6.78	50.20	–
	ENC [22]	7.00	50.00	–
	AMC [17]	8.10	50.00	–
	KSE [31]	6.77	48.00	45.27
	FPGM [16]	6.74	47.70	–
	GAL-0.8 [33]	8.42	39.80	34.10
<b>DHP-38 (Ours)</b>	7.06	39.07	41.10	
ResNet-110 5.31	<b>DHP-62 (Ours)</b>	5.37	63.66	63.2
	Variational [57]	7.04	63.56	58.73
	Pruned-B [27]	6.70	61.40	67.60
	GAL-0.5 [33]	7.26	51.50	55.20
<b>DHP-20 (Ours)</b>	6.61	21.63	22.40	
ResNet-164 4.97	Hinge [29]	5.40	53.61	70.34
	SSS [21]	5.78	53.53	84.75
	<b>DHP-50 (Ours)</b>	5.22	51.67	50.97
	Variational [57]	6.84	50.92	43.30
	<b>DHP-20 (Ours)</b>	6.30	21.78	20.46
DenseNet-12-40 5.26	Variational [57]	6.84	55.22	40.33
	<b>DHP-38 (Ours)</b>	6.06	39.80	63.76
	<b>DHP-28 (Ours)</b>	6.51	29.52	26.01
	GAL-0.1 [33]	6.77	28.60	25.00
Tiny-ImageNet				
MobileNetV1 52.71	<b>DHP-24-2 (Ours)</b>	50.75	101.08	43.58
	MobileNetV1-0.75	54.22	57.42	57.64
	MetaPruning [37]	54.48	56.77	88.14
	<b>DHP-50 (Ours)</b>	51.63	51.91	36.95
MobileNetV2 44.75	<b>DHP-24-2 (Ours)</b>	43.82	99.09	72.72
	<b>DHP-10 (Ours)</b>	52.43	11.92	6.50
	MetaPruning [37]	56.72	11.00	90.27
	MobileNetV2-0.3	53.99	10.09	11.64

pruning stage. As in Fig. 1, the operating point DHP-24-2 is derived by compressing the widened mobile networks with  $\alpha = 2$  and the target FLOPs ratio 24%. For ResNet-56, the proposed method is compared with 9 different network compression methods and achieves the best performance, *i.e.* 6.42% Top-1 error rate on the most intensively investigated 50% compression level. The compression of DenseNet-12-40 is reasonable compared with the other methods. The accuracy of the operating points DHP-62 of ResNet-110 and DHP-50 of ResNet-164

**Table 2.** Results on image super-resolution networks. The upscaling factor is  $\times 4$ . Runtime is averaged for Urban100. Maximum GPU memory consumption is reported for Urban100. FLOPs is reported for a  $128 \times 128$  image patch. DHP achieves significant reduction of runtime

Network	Method	PSNR [dB]					FLOPs	Params	Run-time	GPU Mem
		Set5	Set14	B100	Urban100	DIV2K	[G]	[M]	[ms]	[GB]
SRResNet [26]	Baseline	32.03	28.50	27.52	25.88	28.85	32.83	1.54	34.73	0.6773
	Clustering [48]	31.93	28.44	27.47	25.71	28.75	32.83	0.34	31.07	0.8123
	Factor-SIC3 [51]	31.86	28.38	27.40	25.58	28.65	20.83	0.81	102.51	1.4957
	<b>DHP-60 (Ours)</b>	31.97	28.47	27.48	25.76	28.79	20.29	0.95	27.91	0.5923
	Basis-32-32 [30]	31.90	28.42	27.44	25.65	28.69	19.77	0.74	45.73	0.9331
	Factor-SIC2 [51]	31.68	28.32	27.37	25.47	28.58	18.38	0.66	74.66	1.1201
	Basis-64-14 [30]	31.84	28.38	27.39	25.54	28.63	17.49	0.60	36.75	0.6741
	<b>DHP-40 (Ours)</b>	31.90	28.45	27.47	25.72	28.75	13.71	0.64	22.71	0.4907
	<b>DHP-20 (Ours)</b>	31.77	28.34	27.40	25.55	28.60	7.77	0.36	14.74	0.3795
EDSR [32]	Baseline	32.10	28.55	27.55	26.02	28.93	90.37	3.70	49.73	1.3276
	Clustering [48]	31.93	28.47	27.48	25.77	28.80	90.37	0.82	50.51	1.2838
	Factor-SIC3 [51]	31.96	28.47	27.49	25.81	28.81	65.49	2.19	125.10	1.5007
	Basis-128-40 [30]	32.03	28.45	27.50	25.81	28.82	62.65	2.00	48.19	1.3219
	Factor-SIC2 [51]	31.82	28.40	27.43	25.63	28.70	60.90	1.90	94.94	1.3209
	Basis-128-27 [30]	31.95	28.42	27.46	25.76	28.76	58.28	1.74	45.84	1.3209
	<b>DHP-60 (Ours)</b>	31.99	28.52	27.53	25.92	28.88	55.67	2.28	45.11	0.6950
	<b>DHP-40 (Ours)</b>	32.01	28.49	27.52	25.86	28.85	37.77	1.53	33.50	0.9650
	<b>DHP-20 (Ours)</b>	31.94	28.42	27.47	25.69	28.77	19.40	0.79	22.63	1.1588

is quite close to that of the baseline. More results on ResNet-110 and ResNet-164 are shown in the Supplementary. A comparison between  $\ell_1$  and  $\ell_2$  regularization in Eqn. 6 is done. The results in the Supplementary shows that  $\ell_1$  regularization is better than  $\ell_2$  regularization.

On Tiny-ImageNet, DHP achieves lower Top-1 error rates than MetaPruning [37] under the same FLOPs constraint. DHP results in models that are more accurate than the uniformly scaled networks. The error rate of DHP-10 is 1.56% lower than that of MobileNetV2-0.3 with slightly fewer FLOPs and 5.14% fewer parameters. On MobileNetV1, the accuracy gain of DHP-50 over MobileNetV1-0.75 goes to 2.59% with over 5% fewer FLOPs and 10% fewer parameters. Based on this, we hypothesized that it is possible to derive a model which is more accurate than the original version by pruning the widened mobile networks. And this is confirmed by comparing the accuracy of the operating points DHP-24-2 with the baseline accuracy.

## 4.2 Super-resolution

The results on image super-resolution networks are shown in Table 2. DHP is compared with factorized convolution (Factor) [51], learning filter basis method (Basis) [30], and K-means clustering method (Clustering) [48]. ‘SIC\*’ denotes the number of SIC layers in Factor [51]. The practical FLOPs instead of the theoretical FLOPs is reported for Clustering [48]. To fairly compare the methods and

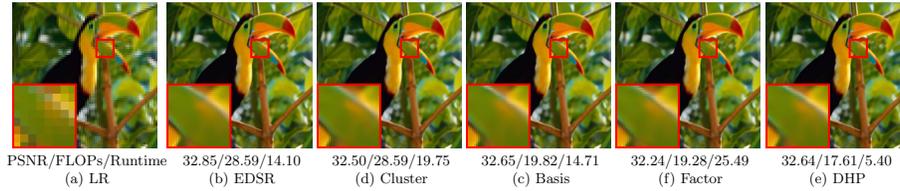
**Table 3.** Results on image denoising networks. The noise level is 70. Runtime and maximum GPU memory are reported for BSD68. FLOPs is reported for a  $128 \times 128$  image. DHP achieves significant reduction of runtime

Network	Method	PSNR [ <i>dB</i> ]		FLOPs	Params	Runtime	GPU Mem
		BSD68	DIV2K	[ <i>G</i> ]	[ <i>M</i> ]	[ <i>ms</i> ]	[ <i>GB</i> ]
DnCNN [56]	Baseline	24.93	26.73	9.13	0.56	23.38	0.1534
	Clustering [48]	24.90	26.67	9.13	0.12	21.97	0.2973
	<b>DHP-60 (Ours)</b>	24.91	26.69	5.65	0.34	18.90	0.1443
	<b>DHP-40 (Ours)</b>	24.89	26.65	3.83	0.23	14.62	0.1194
	Factor-SIC3 [51]	24.97	26.83	3.54	0.22	125.46	0.5910
	Group [43]	24.88	26.64	3.34	0.20	25.69	0.1807
	Factor-SIC2 [51]	24.93	26.76	2.38	0.15	84.17	0.4149
	<b>DHP-20 (Ours)</b>	24.84	26.58	2.01	0.12	10.72	0.0869
UNet [46]	Baseline	25.17	27.17	3.41	7.76	8.73	0.1684
	Clustering [48]	25.01	26.90	3.41	1.72	10.01	0.6704
	<b>DHP-60 (Ours)</b>	25.14	27.11	2.11	4.76	6.86	0.4992
	Factor-SIC3 [51]	25.04	26.94	1.56	3.42	39.84	0.1889
	Group [43]	25.13	27.08	1.49	2.06	11.20	0.1481
	<b>DHP-40 (Ours)</b>	25.12	27.08	1.43	3.24	4.50	0.4992
	Factor-SIC2 [51]	25.01	26.90	1.22	2.51	30.16	0.1855
	<b>DHP-20 (Ours)</b>	25.04	26.97	0.75	1.61	3.93	0.4992

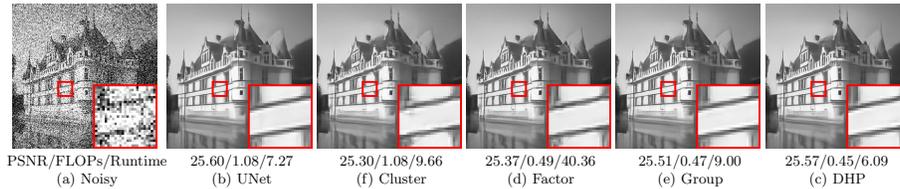
measure the practical compression effectiveness, five metrics are involved including Peak Signal-to-Noise Ratio (PSNR), FLOPs, number of parameters, runtime and GPU memory consumption. Several conclusion can be drawn. **I.** Previous methods mainly focus on the reduction of FLOPs and number of parameter without paying special attention to the actual acceleration. Although Clustering can reduce substantial parameters while maintaining quite good PSNR accuracy, the actual computing resource requirement (GPU memory and runtime) is remained. **II.** Convolution factorization and decomposition methods result in additional CUDA kernel calls, which is not efficient for the actual acceleration. **III.** For the proposed method, the two model complexity metrics, *i.e.* FLOPs and parameters change consistently across different operating points, which leads to consistent reduction of computation resources. **IV.** DHP results in both inference-efficient (DHP-20) and accuracy-preserving (DHP-60) models. The visual results are shown in Fig. 4. As can be seen, the visual quality of the images of DHP is almost indistinguishable from that of the baseline.

### 4.3 Denoising

The results for image denoising networks are shown in Table 3. The same metrics as super-resolution are reported for denoising. An additional method, *i.e.* filter group approximation (Group) [43] is included. In addition to the same conclusion in Subsec. 4.2, another two conclusions are drawn here. **I.** Group [43] fails to reduce the actual computation resources although with quite good accuracy and satisfactory reduction of FLOPs and number of parameters. This might due to the additional  $1 \times 1$  convolution and possibly the inefficient implementation of group convolution in current deep learning toolboxes. **II.** For DnCNN, one



**Fig. 4.** Single image super-resolution visual results. PSNR and FLOPs measured on the image. Runtime averaged on Set5



**Fig. 5.** Image denoising visual results. PSNR and FLOPs measured on the image. Runtime averaged on B100

interesting phenomenon is that Factor [51] is more accurate than the baseline but has larger appetite for other resources. This is due to two facts. Firstly, Factor [51] has skip connections within the SIC layer. The higher accuracy of Factor [51] just validates the effectiveness of skip connections. Secondly, the SIC layer of Factor [51] introduces more convolutional layers. So Factor-SIC3 has five times more convolutional layers than the baseline, which definitely slows down the execution. The visual results are shown in Fig. 5.

## 5 Conclusion and Future Work

In this paper, we proposed a differentiable automatic meta pruning method via hypernetwork for network compression. The differentiability comes with the specially designed hypernetwork and the proximal gradient used to search the potential candidate network configurations. The automation of pruning lies in the uniformly applied  $\ell_1$  sparsity on the latent vectors and the proximal gradient that solves the problem. By pruning mobile network with width multiplier  $\alpha = 2$ , we obtained models with higher accuracy but lower computation complexity than that with  $\alpha = 1$ . We hypothesize this is due to the per-layer distinguishing configuration resulting from the automatic pruning. Future work might be investigating whether this phenomenon reoccurs for the other networks.

## Acknowledgements

This work was partly supported by the ETH Zürich Fund (OK), a Huawei Technologies Oy (Finland) project, an Amazon AWS grant, and an Nvidia grant.

## References

1. MIT technology review: 10 breakthrough technologies 2020. <https://www.technologyreview.com/lists/technologies/2020/#tiny-ai>, accessed: 2020-03-01
2. Alvarez, J.M., Salzmann, M.: Learning the number of neurons in deep networks. In: *Proce. NeurIPS*. pp. 2270–2278 (2016)
3. Alvarez, J.M., Salzmann, M.: Compression-aware training of deep networks. In: *Proc. NeurIPS*. pp. 856–867 (2017)
4. Brock, A., Lim, T., Ritchie, J., Weston, N.: SMASH: One-shot model architecture search through hypernetworks. In: *Proc. ICLR* (2018)
5. Chang, O., Flokas, L., Lipson, H.: Principled weight initialization for hypernetworks. In: *Proc. ICLR* (2020)
6. Chen, C., Tung, F., Vedula, N., Mori, G.: Constraint-aware deep neural network compression. In: *Proc. ECCV*. pp. 400–415 (2018)
7. Chin, T.W., Ding, R., Zhang, C., Marculescu, D.: Towards efficient model compression via learned global ranking. In: *Proc. CVPR*. pp. 1518–1528 (2020)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: ImageNet: A large-scale hierarchical image database. In: *Proc. CVPR*. pp. 248–255. IEEE (2009)
9. Dong, X., Chen, S., Pan, S.: Learning to prune deep neural networks via layer-wise optimal brain surgeon. In: *Proc. NIPS*. pp. 4857–4867 (2017)
10. Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: *Proc. AISTATS*. pp. 249–256 (2010)
11. Ha, D., Dai, A., Le, Q.V.: HyperNetworks. In: *Proc. ICLR* (2017)
12. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In: *Proc. ICLR* (2015)
13. Hassibi, B., Stork, D.G.: Second order derivatives for network pruning: Optimal brain surgeon. In: *Proc. NeurIPS*. pp. 164–171 (1993)
14. Hayashi, K., Yamaguchi, T., Sugawara, Y., Maeda, S.i.: Einconv: Exploring unexplored tensor decompositions for convolutional neural networks. In: *Proc. NeurIPS*. pp. 5553–5563 (2019)
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proc. CVPR*. pp. 770–778 (2016)
16. He, Y., Liu, P., Wang, Z., Hu, Z., Yang, Y.: Filter pruning via geometric median for deep convolutional neural networks acceleration. In: *Proc. CVPR*. pp. 4340–4349 (2019)
17. He, Y., Lin, J., Liu, Z., Wang, H., Li, L.J., Han, S.: AMC: AutoML for model compression and acceleration on mobile devices. In: *Proc. ECCV*. pp. 784–800 (2018)
18. He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. In: *Proc. ICCV*. pp. 1389–1397 (2017)
19. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
20. Huang, G., Liu, Z., van der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proc. CVPR*. pp. 2261–2269 (2017)
21. Huang, Z., Wang, N.: Data-driven sparse structure selection for deep neural networks. In: *Proc. ECCV*. pp. 304–320 (2018)

22. Kim, H., Umar Karim Khan, M., Kyung, C.M.: Efficient neural network compression. In: Proc. CVPR (June 2019)
23. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Tech. rep., Citeseer (2009)
24. Krueger, D., Huang, C.W., Islam, R., Turner, R., Lacoste, A., Courville, A.: Bayesian hypernetworks. arXiv preprint arXiv:1710.04759 (2017)
25. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: Proc. NeurIPS. pp. 598–605 (1990)
26. Ledig, C., Theis, L., Huszár, F., Caballero, J., Cunningham, A., Acosta, A., Aitken, A., Tejani, A., Totz, J., Wang, Z., et al.: Photo-realistic single image super-resolution using a generative adversarial network. In: Proc. CVPR. pp. 105–114 (2017)
27. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. In: Proc. ICLR (2017)
28. Li, J., Qi, Q., Wang, J., Ge, C., Li, Y., Yue, Z., Sun, H.: OICSR: Out-in-channel sparsity regularization for compact deep neural networks. In: Proc. CVPR. pp. 7046–7055 (2019)
29. Li, Y., Gu, S., Mayer, C., Van Gool, L., Timofte, R.: Group sparsity: The hinge between filter pruning and decomposition for network compression. In: Proc. CVPR (2020)
30. Li, Y., Gu, S., Van Gool, L., Timofte, R.: Learning filter basis for convolutional neural network compression. In: Proc. ICCV. pp. 5623–5632 (2019)
31. Li, Y., Lin, S., Zhang, B., Liu, J., Doermann, D., Wu, Y., Huang, F., Ji, R.: Exploiting kernel sparsity and entropy for interpretable CNN compression. In: Proc. CVPR (2019)
32. Lim, B., Son, S., Kim, H., Nah, S., Lee, K.M.: Enhanced deep residual networks for single image super-resolution. In: Proc. CVPRW. pp. 1132–1140 (2017)
33. Lin, S., Ji, R., Yan, C., Zhang, B., Cao, L., Ye, Q., Huang, F., Doermann, D.: Towards optimal structured cnn pruning via generative adversarial learning. In: Proc. CVPR. pp. 2790–2799 (2019)
34. Liu, B., Wang, M., Foroosh, H., Tappen, M., Pensky, M.: Sparse convolutional neural networks. In: Proc. CVPR. pp. 806–814 (2015)
35. Liu, C., Zoph, B., Neumann, M., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K.: Progressive neural architecture search. In: Proc. ECCV (September 2018)
36. Liu, H., Simonyan, K., Yang, Y.: DARTS: Differentiable architecture search. In: Proc. ICLR (2019)
37. Liu, Z., Mu, H., Zhang, X., Guo, Z., Yang, X., Cheng, T.K.T., Sun, J.: MetaPruning: Meta learning for automatic neural network channel pruning. In: Proc. ICCV (2019)
38. Liu, Z., Sun, M., Zhou, T., Huang, G., Darrell, T.: Rethinking the value of network pruning. In: Proc. ICLR (2019)
39. Minnehan, B., Savakis, A.: Cascaded projection: End-to-end network compression and acceleration. In: Proc. CVPR (June 2019)
40. Molchanov, P., Mallya, A., Tyree, S., Frosio, I., Kautz, J.: Importance estimation for neural network pruning. In: Proc. CVPR. pp. 11264–11272 (2019)
41. Pan, Z., Liang, Y., Zhang, J., Yi, X., Yu, Y., Zheng, Y.: Hyperst-net: Hypernetworks for spatio-temporal forecasting. arXiv preprint arXiv:1809.10889 (2018)
42. Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., Lerer, A.: Automatic differentiation in PyTorch (2017)

43. Peng, B., Tan, W., Li, Z., Zhang, S., Xie, D., Pu, S.: Extreme network compression via filter group approximation. In: Proc. ECCV. pp. 300–316 (2018)
44. Pham, H., Guan, M., Zoph, B., Le, Q., Dean, J.: Efficient neural architecture search via parameters sharing. In: Proc. ICML. pp. 4095–4104 (2018)
45. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. In: Proc. AAAI. vol. 33, pp. 4780–4789 (2019)
46. Ronneberger, O., Fischer, P., Brox, T.: U-Net: Convolutional networks for biomedical image segmentation. In: Proc. MICCAI. pp. 234–241. Springer (2015)
47. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: MobileNetV2: Inverted residuals and linear bottlenecks. In: Proc. CVPR. pp. 4510–4520 (2018)
48. Son, S., Nah, S., Mu Lee, K.: Clustering convolutional kernels to compress deep neural networks. In: Proc. ECCV. pp. 216–232 (2018)
49. Sutton, R.S., Barto, A.G.: Reinforcement learning: An introduction. MIT press (2018)
50. Torfi, A., Shirvani, R.A., Soleymani, S., Nasrabadi, N.M.: GASL: Guided attention for sparsity learning in deep neural networks. arXiv preprint arXiv:1901.01939 (2019)
51. Wang, M., Liu, B., Foroosh, H.: Factorized convolutional neural networks. In: Proc. ICCVW. pp. 545–553 (2017)
52. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: Proc. NeurIPS. pp. 2074–2082 (2016)
53. Ye, J., Lu, X., Lin, Z., Wang, J.Z.: Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. In: Proc. ICLR (2018)
54. Yoon, J., Hwang, S.J.: Combined group and exclusive sparsity for deep neural networks. In: Proc. ICML. pp. 3958–3966. JMLR. org (2017)
55. Yu, R., Li, A., Chen, C.F., Lai, J.H., Morariu, V.I., Han, X., Gao, M., Lin, C.Y., Davis, L.S.: NISP: Pruning networks using neuron importance score propagation. In: Proc. CVPR. pp. 9194–9203 (2018)
56. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. IEEE TIP **26**(7), 3142–3155 (2017)
57. Zhao, C., Ni, B., Zhang, J., Zhao, Q., Zhang, W., Tian, Q.: Variational convolutional neural network pruning. In: Proc. CVPR. pp. 2780–2789 (2019)
58. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. In: Proc. ICLR (2017)
59. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: Proc. CVPR (June 2018)