

The Heterogeneity Hypothesis: Finding Layer-Wise Differentiated Network Architectures

Yawei Li¹, Wen Li², Martin Danelljan¹, Kai Zhang¹, Shuhang Gu³, Luc Van Gool^{1,4}, Radu Timofte¹

¹Computer Vision Lab, ETH Zürich, ²UESTC, ³The University of Sydney, ⁴KU Leuven

{yawei.li, martin.danelljan, kai.zhang, vangool, radu.timofte}@vision.ee.ethz.ch

{liwenbnu, shuhangu}@gmail.com

Abstract

In this paper, we tackle the problem of convolutional neural network design. Instead of focusing on the design of the overall architecture, we investigate a design space that is usually overlooked, i.e. adjusting the channel configurations of predefined networks. We find that this adjustment can be achieved by shrinking widened baseline networks and leads to superior performance. Based on that, we articulate the “heterogeneity hypothesis”: with the same training protocol, there exists a layer-wise differentiated network architecture (LW-DNA) that can outperform the original network with regular channel configurations but with a lower level of model complexity.

The LW-DNA models are identified without extra computational cost or training time compared with the original network. This constraint leads to controlled experiments which direct the focus to the importance of layer-wise specific channel configurations. LW-DNA models come with advantages related to overfitting, i.e. the relative relationship between model complexity and dataset size. Experiments are conducted on various networks and datasets for image classification, visual tracking and image restoration. The resultant LW-DNA models consistently outperform the baseline models. Code is available at https://github.com/ofsoundof/Heterogeneity_Hypothesis.git.

1. Introduction

Since the advent of the deep learning era, convolutional neural network (CNN) [21] design has replaced the role of feature design in various computer vision tasks. Recently, neural network design has also evolved from manual design [44, 14, 18] to neural architecture search (NAS) [31, 45] and semi-automation [49, 16, 38]. State-of-the-art network designs focus on discovering the overall network architecture with regularly repeated convolutional

layers. This has been the golden standard of current CNN designs. For example, Ma *et al.* mentioned that a network should have equal channel width [33]. But their analysis is limited to minimizing the memory access cost given the FLOPs for a single pointwise convolution.

The motivation of this paper kind of contradicts the previous design heuristics. It investigates a design space that is usually overlooked and thus not fully explored, namely adjusting the layer-wise channel configurations. In this paper, the channel configuration of a network is defined as the vector that summarizes the output channels of the convolutional layers. We try to answer three questions: 1) whether there exists a layer-wise differentiated network architecture (LW-DNA) that can outperform the original one; 2) if so, how to identify it efficiently; and 3) why it can beat the regular configuration.

Question 1: The existence of LW-DNA. To answer the first question, we formally articulate the following hypothesis. **The Heterogeneity Hypothesis:** For a CNN, when trained with exactly the same training protocol (e.g. number of epochs, batch size, learning rate schedule), there exists a layer-wise differentiated network architecture (LW-DNA) that can outperform the original network with regular layer-wise channel configurations but with a lower model complexity in term of FLOPs and parameters.

To be specific, we aim at adjusting the numbers of channels of the convolutional layers in predefined CNNs. The other layer configurations such as kernel size and stride are not changed. Formally, consider an L -layer CNN $f(\mathbf{X}; \Theta, \mathbf{c})$, where $\mathbf{c} = (c_1, c_2, \dots, c_L)$ is the channel configuration of all of the convolutional layers, Θ denotes the parameters in the network, and \mathbf{X} is the input of the network. The heterogeneity hypothesis implies that there should exist a new channel configuration $\mathbf{c}' = (c'_1, c'_2, \dots, c'_L)$ such that the new architecture $f'(\mathbf{X}; \Theta', \mathbf{c}')$ performs no worse than the original one. After the adjustment, the channel configurations c'_l could be either larger or smaller than the original c_l . We try to answer this question by empirical experiments.

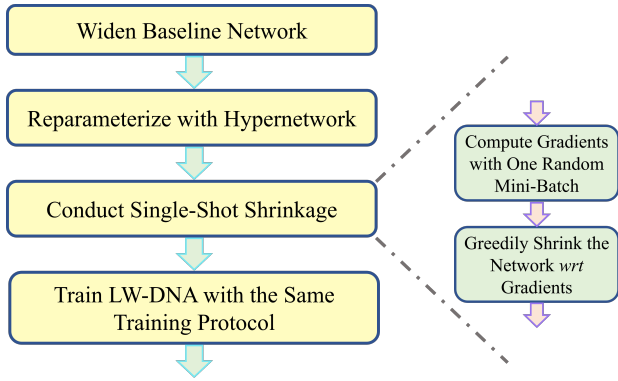


Figure 1: Pipeline of identifying LW-DNA models. Note that the single-shot shrinkage method only needs to run one random mini-batch. Then the network is shrunk after the single pass. Thus, almost no additional computational cost is introduced. This allows for fair comparison between the baseline model and the LW-DNA model.

Question 2: How to identify an LW-DNA efficiently?

Note that the focus of this paper is solely the network architectures. The influence of factors other than network architecture such as the training protocol are excluded. This choice allows for controlled experiments and a fair comparison between the possibly existing LW-DNA models and the baseline models. But we are in turn faced with the following problem. **Problem Statement:** *If the heterogeneity hypothesis is valid, how can we efficiently and reliably find an LW-DNA model for a CNN without additional computational cost and training time?*

To solve this problem, we are inspired by recent developments in network compression [25, 24, 28]. The pipeline of identifying LW-DNA models is shown in Fig. 1. In short, the LW-DNA models are identified by the single-shot shrinking of a widened and reparameterized version of the baseline network. The details are given in Sec. 4

Question 3: How to explain the benefits of LW-DNA?

As a matter of examples, we identify LW-DNA versions of various state-of-the-art networks for three vision tasks, incl. image classification [14, 18, 17, 43, 16, 45, 46, 38], image restoration [23, 29, 52, 41], and visual tracking [3]. Interestingly, the identified LW-DNA models consistently outperform the baselines even with lower model complexities in terms of FLOPs and number of parameters. We try to explain this phenomenon from several perspectives.

1. CNNs are redundant. So it is possible to find a layer-wise specific channel configuration comparable with the baseline under lower model complexity.
2. As shown in Fig. 5, some layers of the LW-DNA models have more channels than the baseline. Indeed, the lower layers tend to be strengthened with more channels. It might be those layers that play the essential role in improving the network accuracy.

3. The accuracy gain of the LW-DNA models might be related to overfitting by the baseline models. We derive this conjecture from several observations. **I.** By comparing the training and testing curves of an LW-DNA model and its baseline in Fig. 4, we find that towards the end of the training, the identified LW-DNA model shows a higher training error but a lower testing error, *i.e.* improved generalization. This phenomenon is consistent across different datasets. This also matches the observations from the pioneering unstructured pruning, like a brain surgeon trying to boost network generalization after brain damage [22, 13]. **II.** The accuracy gain of an LW-DNA model is larger for smaller datasets (*i.e.* Tiny-ImageNet) that are easier to get overfitted to, compared with larger datasets (*i.e.* ImageNet). **III.** On the same dataset (*i.e.* ImageNet), it is easier to identify an LW-DNA model version for larger networks (*i.e.* ResNet50) than for smaller networks (*i.e.* MobileNetV3).

The contributions of this paper can be summarized as follows. **First**, it demonstrates the possibility of identifying a superior version of a network by only adjusting the channel configuration of the network. This could be used as a post-searching mechanism complementary to semi- or fully automated neural architecture search. **Secondly**, a method that can identify LW-DNA models almost without additional computational cost and training time is proposed. This method only needs the computation of one random batch. **Thirdly**, the possible reason for the improved performance of an LW-DNA is explained by observing the experimental results.

2. Related Work

The lottery ticket hypothesis (LTH). The heterogeneity hypothesis is reminiscent of the LTH [9], which addresses the existence of sparse subnetworks that can match the test accuracy of randomly-initialized dense networks. The winning ticket is identified by greedily pruning single elements of weight parameters with smallest magnitude. Following works try to extend [39], theoretically prove [34], understand [53], and improve the training process [40] of LTH. The unstructured pruning breaks the dynamical isometry in the network [24]. The core problem is the trainability of the sparse subnetworks and the gradient flow in the subnetworks [24]. In contrast, the heterogeneity hypothesis focuses on adjusting the channel configuration of the network. Since the weight elements of an entire channel are pruned together, there is no irregular kernel in the pruned network. Gradient flow is no longer a problem in this scenario.

NAS. NAS automatizes neural network design by searching in the design space [54, 30, 37]. Earlier works consume lots of computation [54, 30]. Recent develop-

ments accelerate the searching procedure by introducing differentiability into searching [31]. After the searching stage, the derived cells are repeated to construct the final network. Thus, the final network still has a regular architecture. In this paper, we try to adjust the channel configuration of the network, which can be regarded as a method complementary to NAS. Recent works try to push the frontier of NAS research by either redesigning the search space or proposing a more efficient search method [47, 42, 10, 5].

Hypernetworks. Hypernetworks are actually a kind of reparameterization of the backbone network [11]. Hypernetworks generate the weight parameters of the backbone network. The input of hypernetworks can be either static or dependent on the feature maps of the backbone network. In this sense, hypernetworks fall under the paradigm of meta learning. Recent developments bring hypernetworks to network compression [32, 28]. Earlier hypernetwork designs are just a stack of two linear layers. Thus, the outputs are fixed, which should be cropped before being used as weights of the backbone network. The recent hypernetworks [28] can adapt the outputs according to the length of the input latent vectors. This design naturally suits the task of network compression. This is one of the reasons why we select hypernetworks as our shrinkage agent.

Network shrinkage. Network shrinkage removes unimportant weight parameters in the network [22, 13, 12, 26, 27, 25, 28]. Since we want to purely investigate the importance of the architecture of the identified network, the other factors such as training protocol should be excluded. The network shrinkage procedure should also be simplified as much as possible. Inspired by [25, 24], the widened network is shrunk at initialization according to gradients. The network shrinkage procedure only needs one random batch.

Difference from network compression works. This paper is different from the previous network compression works in the following aspects. **Aim.** The aim of this paper is a proof of a concept that it is possible to benefit better from the computation and parameter budget by optimizing the architecture of the network. The identified LW-DNA model of a predefined network has improved accuracy and slightly reduced model complexity. Previous network compression works aim at improving the efficiency of networks. Accuracy drop is inevitable for the compact networks. **Method.** The single-shot method in [25] for unstructured pruning is transferred to network shrinkage by its collaboration with hypernetworks. There is no computational overhead for the network shrinkage method used in this paper. **Interpretation.** This paper tries to interpret where the benefit of the slightly reduced models comes from, which is not done by recent works.

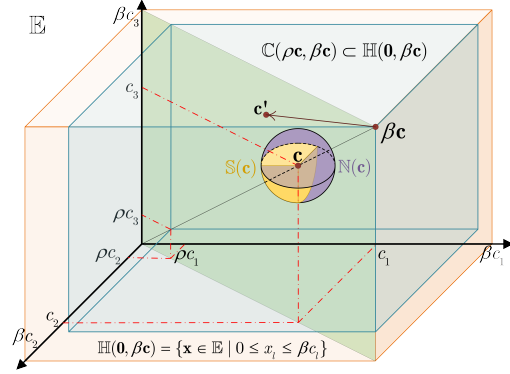


Figure 2: Illustration of the configuration space. The proposed method identifies layer-specific channel configurations within the enlarged and constrained subspace $\mathbb{C}(\rho c, \beta c)$. Compared with searching within the constrained neighborhood $\mathbb{S}(c)$ of c , the enlarged configuration space makes it possible to develop a straightforward shrinkage criterion.

3. Preliminaries

3.1. Hints from network compression

Recent network compression methods shed light on the existence of advantageous layer-wise specific networks [32, 28, 7]. Those methods can result in shrunk networks with layer-wise specific channel configurations. Some works [32] report accuracy gains of the pruned network over the width-scaled versions of ResNet and MobileNets [14, 17, 43]. Yet, since the advantageous networks are identified in a network compression sense, thus with an accuracy drop compared with the uncompressed network, it still remains unknown whether there exists a layer-wise specific network that can compete with the original one. A recent work [28] reports an accuracy gain over uncompressed MobileNets on Tiny-ImageNet. Yet, further investigations on larger datasets are not conducted. Moreover, the compact networks are usually derived with training protocols different from those used for the baseline network, *e.g.* additional searching stage, larger batch size, or prolonged fine-tuning stage. It remains unknown how the layer-wise specific channel configurations benefit the network.

3.2. Notations and definitions

Notation. In this paper, bold lowercase letters such as c , x , and z are used to denote vectors while bold capital letters such as O , Z , W are used to denote matrices and higher dimensional tensors. The vectors, matrices, and higher dimensional tensors are indexed by subscripts. Greek letters such as α , β denote constant scalars. The configuration vector and configuration space are formally defined as follows.

Definition 1 (Channel configuration vector). Consider an L -layer CNN. The channel configuration vector of the

network is defined as an L -dimensional vector that summarizes the number of output channels of the network, *i.e.*

$$\mathbf{c} = (c_1, c_2, \dots, c_L), \quad (1)$$

where c_l denotes the number of output channels in the l -th layer.

Definition 2 (Configuration space). The configuration space \mathbb{E} is a subspace of Euclidean space that contains the allowable channel configuration vectors. (See Fig. 2 for one example of the configuration space.)

The dimension of the configuration vectors depends on the number of convolutional layers in the network. Take VGG11 for example. The configuration vector is an 8-dimensional vector, *i.e.*,

$$\mathbf{c}_{vgg} = (64, 128, 256, 256, 512, 512, 512, 512). \quad (2)$$

As in this example, the configuration vector is regular and its elements are dependent on each other in the sense that most of them are repeated. For image classification networks, the golden standard is to repeat building blocks with the same configuration up to the point where the spatial dimension of the feature map gets reduced. Some efficient designs for mobile devices introduce a width multiplier α to adapt to constrained resource requirements, which results in a scaled configuration vector, *i.e.*,

$$\mathbf{b} = (\alpha c_1, \alpha c_2, \dots, \alpha c_L), \alpha < 1. \quad (3)$$

3.3. Problem formulation and recast

Since the configuration vector is manually fixed, it is not guaranteed to be optimal. In this paper, we explore the corresponding configuration design space. The aim is to demonstrate that there is an irregular configuration vector \mathbf{c}' that can compete with the original, while offering reduced model complexity. To achieve that, we propose an algorithm which can adjust (increase or decrease) the elements of the configuration vector \mathbf{c} while controlling the model complexity. As shown in Fig. 2, such an adjustment procedure best searches in the neighborhood of the vector, *i.e.*

$$\mathbb{N}(\mathbf{c}) \subset \mathbb{E}. \quad (4)$$

After the adjustment, an element of the configuration vector \mathbf{c} can be either increased or decreased, which corresponds to growing or shrinking the l -th layer of the network. Shrinkage criteria can be defined on the existing network and network shrinkage algorithm could applied. The limitation of a shrinkage algorithm on the original network is that it can only explore a subspace of the neighborhood, *i.e.*

$$\mathbb{S}(\mathbf{c}) = \{\mathbf{x} \in \mathbb{N}(\mathbf{c}) | x_l \leq c_l\} \subset \mathbb{N}(\mathbf{c}). \quad (5)$$

But we do not want to be restricted to shrinkage only. Instead, it is desirable to do both network shrinkage and

growth at the same time for the configuration vector adjustment.

We circumvent this problem by recasting it as a shrinkage problem in a larger configuration space which is obtained by widening the width of the network with a width multiplier $\beta > 1$. The new searching space \mathbb{H} is a hyperrectangle delimited by the zero vector $\mathbf{0}$ and the up-scaled configuration vector $\beta\mathbf{c}$ in the high-dimensional space, *i.e.*

$$\mathbb{H}(\mathbf{0}, \beta\mathbf{c}) = \{\mathbf{x} \in \mathbb{E} | 0 \leq x_l \leq \beta c_l\} \subset \mathbb{E}. \quad (6)$$

The searching algorithm then starts from the up-scaled vector $\beta\mathbf{c}$ and reduces the value of its l -th element greedily according to the significance of the channels in the corresponding convolutional layer.

4. Methodology

After introducing the preliminaries and the designing considerations in the last section, the algorithm used to identify LW-DNA models is explained in this section. The pipeline is already shown in Fig. 1. The identifying procedure proceeds as follows. 1) Reparameterize the widened baseline network with hypernetworks. The outputs of the hypernetworks act as the weight parameters of the baseline network. The inputs of the hypernetwork serve as the handle to shrink the network. 2) Compute the gradients of the hypernetwork input, *i.e.* the latent vectors, with one random batch. 3) Sparsify the latent vectors greedily according to the magnitude of their gradients. 4) Compute the weight parameters with the sparsified latent vectors. 5) Train the resultant network from scratch with the same training protocol as the baseline network. And in the following, we explain some of the key steps in detail.

4.1. Reparameterizing with hypernetworks

The network shrinkage method is explained in this section. Instead of directly shrinking the baseline network, we first widen it and reparameterize it with hypernetworks [32, 28]. The reparameterization is adopted based on the following considerations. The hypernetworks bring the shrinkage problem into a latent space. Removing a channel is equivalent to deleting a single element of the latent vector, which converts the problem of dealing with elements in the whole channel to an easier one of dealing with a single element in the latent vector. In addition, it provides a straightforward extension of single-shot shrinkage [25] to channel pruning (See Subsec 4.2). And single-shot shrinkage is the core of avoiding additional computational cost when identifying LW-DNA models. The latent vector sharing mechanism in the hypernetworks also makes it possible to deal with various state-of-the-art networks.

Consider the L -layer CNN that is brought to the larger configuration space $\mathbb{H}(\mathbf{0}, \beta\mathbf{c})$ as in Eqn. (6). The weight

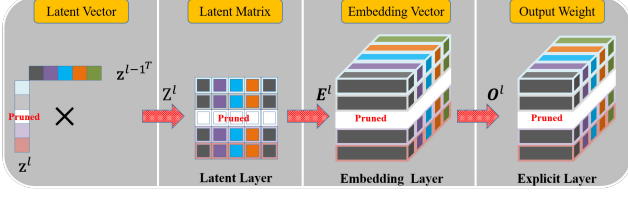


Figure 3: Illustration of the operations in the hypernetworks.

parameter of the l -th convolutional layer of the CNN has the dimension of $\beta_{c_l} \times \beta_{c_{l-1}} \times w \times h$, where $\beta_{c_l}, \beta_{c_{l-1}}$, and $w \times h$ denotes the output channel, input channel, and kernel size of the layer. For the simplicity of notation, let $n = \beta_{c_l}$ and $c = \beta_{c_{l-1}}$ in the following.

To reparameterize the l -th layer of the network, a latent vector $\mathbf{z}^l \in \mathbb{R}^n$ is first attached to it. The latent vector controls the output channel of the layer and removing an element of the latent vector is equivalent to deleting an output channel of the layer. Since the output channel of the current layer acts as the input channel of the next layer, the latent vectors are shared between consecutive layers. Then as shown in Fig. 3, the hypernetwork takes as input the latent vector of the previous layer and the current layer. It computes a latent matrix, *i.e.* $\mathbf{Z}^l = \mathbf{z}^l \cdot \mathbf{z}^{l-1T}$, which forms a grid used for network shrinkage. Every element of the latent matrix is transformed to a vector by two consecutive linear operations, *i.e.*

$$\mathbf{O}_{i,j}^l = \mathbf{W}_2^l \cdot (\mathbf{Z}_{i,j}^l \mathbf{W}_1^l), \quad (7)$$

where $\mathbf{W}_1^l \in \mathbb{R}^{m \times 1}$ transforms the elements into a higher-dimensional embedding space and $\mathbf{W}_2^l \in \mathbb{R}^{wh \times m}$ converts the tensors vectors in the embedding space to the output weight. Note that \mathbf{W}_1^l and \mathbf{W}_2^l are unique for each element $\mathbf{Z}_{i,j}^l$ and for the simplicity of notation the subscript i, j is omitted. The output could be assembled into a 3D tensor $\mathbf{O}^l \in \mathbb{R}^{n \times c \times wh}$ which can be used as the weight parameter of the convolutional layer. The latent matrix acts as a handle to shrink the layer. As shown in Fig. 3, if a single element of the latent vector \mathbf{z}^l is nullified, the entire row of the latent matrix is masked out. As a result, the corresponding output channel is removed. Similarly, removing an element of \mathbf{z}^{l-1} corresponds to removing the output channel of the $l-1$ -th layer and the input channel of the l -th layer. Details of handling different layers such as depth-wise convolution and batch normalization are given in the supplementary.

4.2. Single-shot shrinkage

After reparameterizing the network with hypernetworks, the parameters in the network are first randomly initialized [4]. Then the single-shot shrinkage method is used to adjust the width of the network.

Consider a single mini-batch $\{\mathbf{X}_i, \mathbf{Y}_i\}$ from the dataset. The output of the network is computed as

$$\hat{\mathbf{Y}}_i = f(\mathbf{X}_i; \Theta, \mathbf{z}), \quad (8)$$

where \mathbf{z} denotes the latent vectors and Θ is the parameter set that contains \mathbf{W}_1 and \mathbf{W}_2 . The loss is computed as

$$\mathcal{L} = \mathcal{L}(\mathbf{Y}_i, f(\mathbf{X}_i; \Theta, \mathbf{z})). \quad (9)$$

Then the gradients of the loss function with respect to the latent vectors are computed as

$$\nabla \mathcal{L} = \frac{\partial \mathcal{L}(\mathbf{Y}_i, f(\mathbf{X}_i; \Theta, \mathbf{z}))}{\partial \mathbf{z}}. \quad (10)$$

The magnitude of the gradients is used as the criterion to sparsify the latent vectors. The elements whose gradient magnitude is smaller than a threshold are removed. The threshold is determined by a binary search algorithm, which allows the resultant network to reach a predefined FLOP target. The resultant network is the final LW-DNA model and is trained from scratch with the same training protocol as the baseline model.

The single-shot shrinkage method is inspired by single-shot pruning of weight elements [25]. But the original method is single element oriented. It removes single weight parameters in the network and results in unstructured kernels. It remains to be explored how to transform the single-shot method to network shrinkage. The hypernetworks provide such a connection. By resorting to hypernetworks, the shrinkage is conducted on the latent space whose elements correspond to channels in the network and serve as the agent for shrinkage. Deleting an element of the latent vector is equivalent to remove a channel in the network. Thus, sparsifying the latent vectors according to their gradients is a natural transferring of the single-shot method in [25].

4.3. Constraining model complexity

Model complexity is measured in terms of FLOP and parameter count. The target is to find a model that has both fewer FLOPs and parameters while achieving improved accuracy. Yet, the two metrics are not always consistent with each other. For example, when the FLOPs target is set, a parameter over-pruned model might be observed in some of the experiments, which could lead to inferior performance. Thus, a new hyper-parameter ρ is introduced which controls the minimum percentage of remaining channels in convolutional layers. In this way, the search space $\mathbb{C}(\rho\mathbf{c}, \beta\mathbf{c})$ is a confined subspace of the original search space $\mathbb{H}(\mathbf{0}, \beta\mathbf{c})$, *i.e.*

$$\mathbb{C}(\rho\mathbf{c}, \beta\mathbf{c}) = \{\mathbf{x} \in \mathbb{E} | \rho c_l \leq x_l \leq \beta c_l\} \subset \mathbb{H}(\mathbf{0}, \beta\mathbf{c}). \quad (11)$$

A similar hyper-parameter τ is introduced for the final linear layers of image classification networks. The hyper-parameters ρ and τ are termed convolutional percentage

Dataset	Network	Method	Top-1 Error (%)	FLOPs [G] / Ratio (%)	Params [M] / Ratio (%)
ImageNet [6]	ResNet50 [14]	Baseline	23.28	4.1177 / 100.0	25.557 / 100.0
		MutualNet [48]	21.40	4.1177 / 100.0	25.557 / 100.0
		LW-DNA	23.00	3.7307 / 90.60	23.741 / 92.90
		MetaPruning [32]	23.80	3.0000 / 72.86	–
		AutoSlim [50]	24.00	3.0000 / 72.86	23.100 / 90.39
	RegNet [38] X-4.0GF	Baseline	23.05	4.0005 / 100.0	22.118 / 100.0
		LW-DNA	22.74	3.8199 / 95.49	15.285 / 69.10
	MobileNetV3 small [16]	Baseline	34.91	0.0612 / 100.0	3.108 / 100.0
		LW-DNA	34.84	0.0605 / 98.86	3.049 / 98.11
Tiny-ImageNet	MobileNetV1 [17]	Baseline	51.87	0.0478 / 100.0	3.412 / 100.0
		Baseline KD	48.00	0.0478 / 100.0	3.412 / 100.0
		DHP KD	46.70	0.0474 / 99.16	2.267 / 66.43
		LW-DNA	46.44	0.0460 / 96.23	1.265 / 37.08
	MobileNetV2 [43]	Baseline	44.38	0.0930 / 100.0	2.480 / 100.0
		Baseline KD	41.25	0.0930 / 100.0	2.480 / 100.0
		DHP KD	41.06	0.0896 / 96.34	2.662 / 107.34
		LW-DNA	40.74	0.0872 / 93.76	2.230 / 89.90
	MobileNetV3 small [16]	Baseline	47.55	0.0207 / 100.0	2.083 / 100.0
		Baseline KD	41.52	0.0207 / 100.0	2.083 / 100.0
		DHP KD	41.46	0.0192 / 92.75	1.078 / 51.76
		LW-DNA	41.35	0.0178 / 85.99	1.799 / 86.36
	MnasNet [45]	Baseline	51.79	0.0271 / 100.0	3.359 / 100.0
		Baseline KD	48.17	0.0271 / 100.0	3.359 / 100.0
		DHP KD	48.10	0.0264 / 97.42	2.512 / 74.79
		LW-DNA	46.85	0.0250 / 92.25	1.258 / 37.45
CIFAR100	RegNet [38] Y-400MF	Baseline	21.65	0.4585 / 100.0	3.947 / 100.0
		LW-DNA	18.65	0.4468 / 97.45	2.466 / 62.48
	RegNet [38] X-400MF	Baseline	21.75	0.4698 / 100.0	4.810 / 100.0
		LW-DNA	18.81	0.4610 / 98.13	4.404 / 91.56
	EfficientNet [46]	Baseline	20.74	0.4161 / 100.0	4.136 / 100.0
		LW-DNA	19.54	0.3850 / 92.53	2.121 / 51.28
	DenseNet40 [18]	Baseline	26.00	0.2901 / 100.0	1.100 / 100.0
		LW-DNA	22.46	0.2638 / 90.93	1.016 / 92.35
CIFAR10 [20]	DenseNet40 [18]	Baseline	5.50	0.2901 / 100.0	1.059 / 100.0
		LW-DNA	4.87	0.2632 / 90.73	0.963 / 90.87
	ResNet56 [14]	Baseline	5.74	0.1274 / 100.0	0.856 / 100.0
		LW-DNA	5.49	0.1262 / 99.06	0.536 / 62.62

Table 1: Image classification results. Baseline and Baseline KD denote the original network trained without and with knowledge distillation respectively.

and linear percentage in this paper, respectively. During the pruning, the FLOP budget is fixed. By tuning the hyperparameters ρ and τ , the algorithm is able to find networks with the same FLOPs but varying parameter budgets.

5. Experimental Results

The experimental results are shown in this section. We try to identify LW-DNA for various state-of-the-art networks including ResNet [14], RegNet [38], MobileNets [17, 43, 16], EfficientNet [46], MnasNet [45], DenseNet [18], SRResNet [23], EDSR [29], DnCNN [52], and U-Net [41]. The identified LW-DNA model and the baseline network

are trained with exactly the same training protocol. The details of the training protocol for different tasks are given in the supplementary. Knowledge distillation [15] is used for image classification on CIFAR [20] and Tiny-ImageNet[6] (Baseline KD, DHP KD [28], and LW-DNA model). The balancing hyperparameter and temperature are set to 0.4 and 4, respectively. The teacher is the pretrained widened version of the baseline network. Knowledge distillation is not used for experiments on ImageNet because the execution of the teacher network in this case also consumes considerable time and GPU resources.

Image classification. The results of image classification networks are compared in Table 1. A complete version of

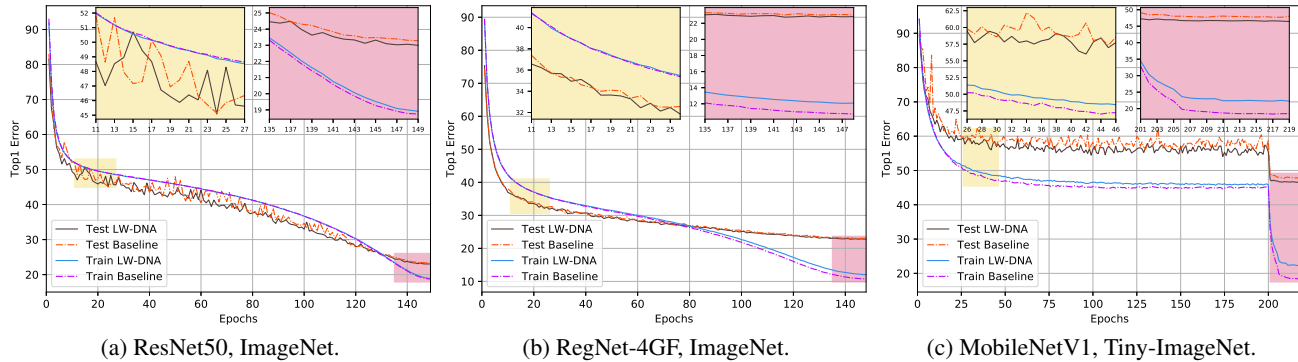


Figure 4: Training and testing log of the LW-DNA models and the baseline models.

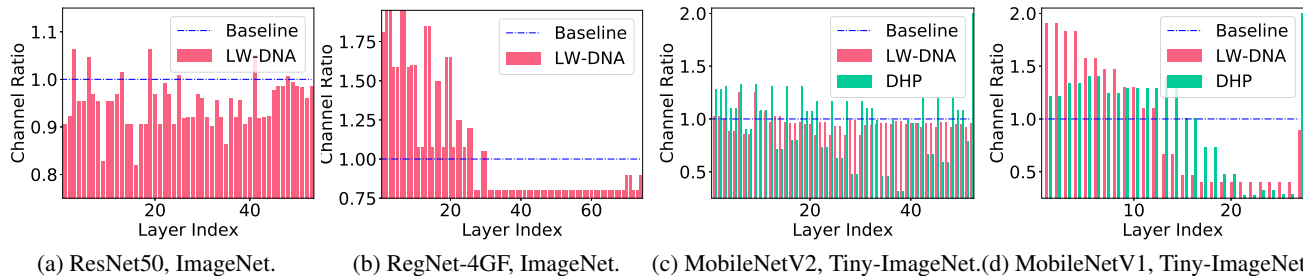


Figure 5: Percentage of remaining output channels of LW-DNA models over the baseline network.

the results is given in the supplementary. We have several key observations. **I.** The identified LW-DNA models outperform the original network (denoted as Baseline or Baseline KD when knowledge distillation is used) with lower model complexity in terms of both FLOPs and number of parameters. This is a direct support for the Heterogeneity Hypothesis. **II.** The accuracy of the baseline network can be improved by knowledge distillation. Yet, the improved baseline still performs worse than LW-DNA. This shows the robustness of LW-DNA, *i.e.* not affected by a specific training technique. **III.** The improvement of LW-DNA scales up to large-scale datasets, *i.e.* ImageNet. For the ImageNet experiment, we set $\rho = 0.4$ and $\tau = 0.45$ by the ablation study on Tiny-ImageNet shown in the supplementary. This hyperparameter combination works well across the three investigated networks. The success on ImageNet and the robustness of the hyper-parameters imply the wide existence of LW-DNA models and the ease of finding them. **IV.** MutualNet is a training scheme when applied to a specific network, which could be combined with our work.

Proximal gradient descent vs. single-shot shrinkage. Besides single-shot shrinkage, there are also other candidate methods to prune networks, *e.g.* proximal gradient descent (PGD). The choice of single-shot shrinkage is based on the following considerations. First, it is extremely computation-efficient. Only one random batch is used to identify the LW-DNA models. This meets the design re-

quirements of introducing no computational cost. This consistency makes it possible to identify the importance of the architecture of LW-DNA models while controlling the other factors. Secondly, by analyzing the closed-form solution to the proximal operator with ℓ_1 regularization, *i.e.* the soft-thresholding operator, we find that PGD tends to diminish the elements of the latent vectors with the approximately consistent speed. As a results, the final magnitude of the elements has some kind of relationship with the initial magnitude. Therefore, if the initialization of an element is large, it is likely that the final magnitude is still relatively large. The distribution of the latent vectors during the PGD optimization is shown in the supplementary. The final distribution is related to the initialization. Thus, it becomes reasonable to shrink the latent vectors at initialization.

The benefits of LW-DNA models are analyzed by several observations of the experimental results. **I.** The percentage of remaining channels is shown in Fig. 5. Some layers of the LW-DNA networks are strengthened. This might contribute to the improved performance of LW-DNA. **II.** As shown in Fig. 4, towards the end of the training, the LW-DNA models shoot a lower test error with increased training error. The improved generalization on the test set comes with reduced model complexity and lower training accuracy. This phenomenon is consistent with the pioneering unstructured pruning methods [22, 13] that try to balance model complexity and overfitting. The same phenomenon

Network	Method	PSNR [dB]					FLOPs [G] / Ratio (%)	Params [M] / Ratio (%)
		Set5 [2]	Set14 [51]	B100 [35]	Urban100 [19]	DIV2K [1]		
SRResNet [23]	Baseline	32.02	28.50	27.52	25.88	28.84	32.81 / 100.0	1.53 / 100.0
	LW-DNA	32.07	28.51	27.52	25.88	28.85	28.79 / 87.75	1.36 / 88.43
EDSR [29]	Baseline	32.10	28.55	27.55	26.02	28.93	90.37 / 100.0	3.70 / 100.0
	LW-DNA	32.13	28.61	27.59	26.09	28.99	55.44 / 61.34	2.84 / 76.94

Table 2: Results on single image super-resolution networks. The upscaling factor is $\times 4$.

on both unstructured pruning and structured pruning points to a common underlying factor. **III.** The accuracy gain of LW-DNA on Tiny-ImageNet is larger than ImageNet. As known, smaller datasets are easier to be overfitted to. **IV.** On ImageNet, it is easier to identify LW-DNA models for ResNet50 and RegNet than MobileNetV3. Since the larger models ResNet50 and RegNet contain more redundancy, it is easier for them to overfit the dataset. Based on the above observations, we conjecture that the improvement of LW-DNA model might be related to model overfitting.

Visual tracking. To validate the generalization ability of the identified LW-DNA, we apply the LW-DNA and baseline version of ResNet50 to visual tracking. State-of-the-art tracking workflow DiMP [3] is used as the test bed. For a fair comparison, the LW-DNA and the baseline are trained with the same protocol. They are first pretrained on ImageNet then finetuned following the DiMP workflow. In Table 3, the networks are compared on two datasets, *i.e.* TrackingNet [36] and LaSOT [8]. On the smaller dataset TrackingNet, LW-DNA version slightly beats the baseline while on the larger dataset LaSOT, LW-DNA outperforms the baseline elegantly. The success plot on LaSOT is shown in Fig. 6. As shown there, DiMP-LW-DNA is consistently better than DiMP-Baseline and other state-of-the-art tracking methods across the range of overlap threshold. In conclusion, the results show that **the benefits of LW-DNA can be transferred to other vision tasks.**

Image Restoration. Table 2 shows the results on super-resolution networks. For EDSR, the LW-DNA models perform better than the baseline but with significant reduction of model complexity. On the large test dataset Urban100 and DIV2K, the LW-DNA model of EDSR leads to nearly 0.1dB PSNR gain over the baseline. For SRResNet, LW-DNA achieves slightly reduction of model complexity without drop of PSNR. More results on image denoising are shown in the supplementary. In conclusion, the results **validate the existence of LW-DNA models for low-level vision networks.**

6. Conclusion

In this paper, we state the heterogeneity hypothesis which in essence is the existence of advantageous LW-DNA models for a predefined network architecture. We try to val-

Metric	DiMP-Baseline	DiMP-LW-DNA
TrackingNet [36]		
Precision	68.06	68.27
Norm. Prec. (%)	79.70	79.64
Success (AUC) (%)	73.77	73.83
LaSOT [8]		
Precision	54.97	57.30
Norm. Prec. (%)	63.70	65.82
Success (AUC) (%)	55.87	57.43

Table 3: Tracking test results. DiMP-LW-DNA and DiMP-Baseline use the identified LW-DNA and baseline version of ResNet50, respectively.

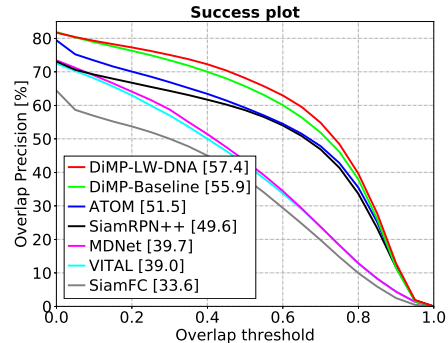


Figure 6: Success plot on the LaSOT dataset for visual tracking.

idate the hypothesis by empirical studies. In order to single out the importance of the network architecture, the training protocol is kept the same for the baseline and the LW-DNA models. This is achieved by converting the problem of identifying LW-DNA to a network shrinkage problem and designing an efficient shrinkage algorithm. The experiments on various network architectures and vision tasks demonstrate the benefits of the identified LW-DNA models. By examining the results, we conjecture that the advantage of the LW-DNA model might be related to model overfitting.

Acknowledgement: This work is partially supported by the Major Project for New Generation of AI under Grant No. 2018AAA0100400, the ETH Zürich Fund (OK), a Huawei Technologies Oy (Finland) project, an Amazon AWS grant, and an Nvidia grant.

References

- [1] Eirikur Agustsson and Radu Timofte. NTIRE 2017 challenge on single image super-resolution: Dataset and study. In *Proc. CVPRW*, July 2017. 8
- [2] Marco Bevilacqua, Aline Roumy, Christine Guillemot, and Marie Line Alberi-Morel. Low-complexity single-image super-resolution based on nonnegative neighbor embedding. In *Proc. BMVC*, 2012. 8
- [3] Goutam Bhat, Martin Danelljan, Luc Van Gool, and Radu Timofte. Learning discriminative model prediction for tracking. In *Proc. ICCV*, pages 6182–6191, 2019. 2, 8
- [4] Oscar Chang, Lampros Flokas, and Hod Lipson. Principled weight initialization for hypernetworks. In *Proc. ICLR*, 2020. 5
- [5] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1294–1303, 2019. 3
- [6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *Proc. CVPR*, pages 248–255. IEEE, 2009. 6
- [7] Xiaohan Ding, Tianxiang Hao, Ji Liu, Jungong Han, Yuchen Guo, and Guiguang Ding. Lossless cnn channel pruning via gradient resetting and convolutional re-parameterization. *arXiv preprint arXiv:2007.03260*, 2020. 3
- [8] Heng Fan, Liting Lin, Fan Yang, Peng Chu, Ge Deng, Sijia Yu, Hexin Bai, Yong Xu, Chunyuan Liao, and Haibin Ling. LaSOT: A high-quality benchmark for large-scale single object tracking. In *Proc. CVPR*, pages 5374–5383, 2019. 8
- [9] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018. 2
- [10] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *European Conference on Computer Vision*, pages 544–560. Springer, 2020. 3
- [11] David Ha, Andrew Dai, and Quoc V Le. HyperNetworks. In *Proc. ICLR*, 2017. 3
- [12] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *Proc. ICLR*, 2015. 3
- [13] Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Proc. NeurIPS*, pages 164–171, 1993. 2, 3, 7
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proc. CVPR*, pages 770–778, 2016. 1, 2, 3, 6
- [15] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 6
- [16] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proc. ICCV*, pages 1314–1324, 2019. 1, 2, 6
- [17] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. MobileNets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 2, 3, 6
- [18] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proc. CVPR*, pages 2261–2269, 2017. 1, 2, 6
- [19] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proc. CVPR*, pages 5197–5206, 2015. 8
- [20] Alex Krizhevsky and Geoffrey Hinton. Learning multiple layers of features from tiny images. Technical report, Cite-seer, 2009. 6
- [21] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 1
- [22] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Proc. NeurIPS*, pages 598–605, 1990. 2, 3, 7
- [23] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proc. CVPR*, pages 105–114, 2017. 2, 6, 8
- [24] Namhoon Lee, Thalaiyasingam Ajanthan, Stephen Gould, and Philip HS Torr. A signal propagation perspective for pruning neural networks at initialization. *arXiv preprint arXiv:1906.06307*, 2019. 2, 3
- [25] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. SNIP: Single-shot network pruning based on connection sensitivity. *arXiv preprint arXiv:1810.02340*, 2018. 2, 3, 4, 5
- [26] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *Proc. ICLR*, 2017. 3
- [27] Yawei Li, Shuhang Gu, Christoph Mayer, Luc Van Gool, and Radu Timofte. Group sparsity: The hinge between filter pruning and decomposition for network compression. In *Proc. CVPR*, 2020. 3
- [28] Yawei Li, Shuhang Gu, Kai Zhang, Luc Van Gool, and Radu Timofte. DHP: Differentiable meta pruning via hypernetworks. *arXiv preprint arXiv:2003.13683*, 2020. 2, 3, 4, 6
- [29] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proc. CVPRW*, pages 1132–1140, 2017. 2, 6, 8
- [30] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proc. ECCV*, September 2018. 2
- [31] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *Proc. ICLR*, 2019. 1, 3

- [32] Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Tim Kwang-Ting Cheng, and Jian Sun. MetaPruning: Meta learning for automatic neural network channel pruning. In *Proc. ICCV*, 2019. 3, 4, 6
- [33] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. ShuffleNet V2: Practical guidelines for efficient cnn architecture design. In *Proc ECCV*, pages 116–131, 2018. 1
- [34] Eran Malach, Gilad Yehudai, Shai Shalev-Shwartz, and Ohad Shamir. Proving the lottery ticket hypothesis: Pruning is all you need. 2020. 2
- [35] D. Martin, C. Fowlkes, D. Tal, and J. Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proc. ICCV*, volume 2, pages 416–423, July 2001. 8
- [36] Matthias Muller, Adel Bibi, Silvio Giancola, Salman Alsubaihi, and Bernard Ghanem. Trackingnet: A large-scale dataset and benchmark for object tracking in the wild. In *Proc. ECCV*, pages 300–317, 2018. 8
- [37] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *Proc. ICML*, pages 4095–4104, 2018. 2
- [38] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. *arXiv preprint arXiv:2003.13678*, 2020. 1, 2, 6
- [39] Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What’s hidden in a randomly weighted neural network? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11893–11902, 2020. 2
- [40] Alex Renda, Jonathan Frankle, and Michael Carbin. Comparing rewinding and fine-tuning in neural network pruning. *arXiv preprint arXiv:2003.02389*, 2020. 2
- [41] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *Proc. MICCAI*, pages 234–241. Springer, 2015. 2, 6
- [42] Binxin Ru, Pedro Esperanca, and Fabio Carlucci. Neural architecture generator optimization. *arXiv preprint arXiv:2004.01395*, 2020. 3
- [43] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted residuals and linear bottlenecks. In *Proc. CVPR*, pages 4510–4520, 2018. 2, 3, 6
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1
- [45] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *Proc. CVPR*, pages 2820–2828, 2019. 1, 2, 6
- [46] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946*, 2019. 2, 6
- [47] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guojun Qi, Qi Tian, and Hongkai Xiong. PC-DARTS: Partial channel connections for memory-efficient architecture search. *arXiv preprint arXiv:1907.05737*, 2019. 3
- [48] Taojiannan Yang, Sijie Zhu, Chen Chen, Shen Yan, Mi Zhang, and Andrew Willis. MutualNet: Adaptive convnet via mutual learning from network width and resolution. 2020. 6
- [49] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. NetAdapt: Platform-aware neural network adaptation for mobile applications. In *Proc. ECCV*, pages 285–300, 2018. 1
- [50] Jiahui Yu and Thomas Huang. Autoslim: Towards one-shot architecture search for channel numbers. *arXiv preprint arXiv:1903.11728*, 2019. 6
- [51] Roman Zeyde, Michael Elad, and Matan Protter. On single image scale-up using sparse-representations. In *International Conference on Curves and Surfaces*, pages 711–730. Springer, 2010. 8
- [52] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a Gaussian denoiser: residual learning of deep CNN for image denoising. *IEEE TIP*, 26(7):3142–3155, 2017. 2, 6
- [53] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. In *NeurIPS*, pages 3592–3602, 2019. 2
- [54] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *Proc. ICLR*, 2017. 2