

This CVPR 2020 paper is the Open Access version, provided by the Computer Vision Foundation. Except for this watermark, it is identical to the accepted version; the final published version of the proceedings is available on IEEE Xplore.

NeuralScale: Efficient Scaling of Neurons for Resource-Constrained Deep Neural Networks

Eugene Lee Chen-Yi Lee Institute of Electronics, National Chiao Tung University eugenelet.ee06q@nctu.edu.tw cylee@si2lab.org

Abstract

Deciding the amount of neurons during the design of a deep neural network to maximize performance is not intuitive. In this work, we attempt to search for the neuron (filter) configuration of a fixed network architecture that maximizes accuracy. Using iterative pruning methods as a proxy, we parameterize the change of the neuron (filter) number of each layer with respect to the change in parameters, allowing us to efficiently scale an architecture across arbitrary sizes. We also introduce architecture descent which iteratively refines the parameterized function used for model scaling. The combination of both proposed methods is coined as NeuralScale. To prove the efficiency of NeuralScale in terms of parameters, we show empirical simulations on VGG11, MobileNetV2 and ResNet18 using CIFAR10, CIFAR100 and TinyImageNet as benchmark datasets. Our results show an increase in accuracy of 3.04%, 8.56% and 3.41% for VGG11, MobileNetV2 and ResNet18 on CIFAR10, CIFAR100 and TinyImageNet respectively under a parameter-constrained setting (output neurons (filters) of default configuration with scaling factor of 0.25).

1. Introduction

The human brain contains around 100 billion of neurons [21] that are structured in such a way that they are utilized in an efficient manner. As the design of deep neural network (DNN) is inspired by the human brain, there's one key ingredient that is missing from the current design of DNNs: the efficient utilization of resources (parameters).

The success of DNN is a composition of many factors. On an architectural level, various architectures have been proposed to increase the accuracy of DNNs targeting efficiency in computational cost (FLOPs) and size (parameters). In a modern DNN architectures, hyperparameters like width (neurons/filters), depth, skip-connections and ac-



Figure 1: The input is represented as a green slab on the left, the output layer is the yellow bar on the right and intermediate layers are represented by 3D blocks with the width as its neuron (filter) number. The purple blocks are the final configuration of the neurons (filters). (a) shows our proposed method that non-linearly scales neurons (filters) across different layers to maximize performance. (b) is a uniform scaling method that is shown to be effective in [12].

tivation functions [38, 42] are the building blocks. Notable architectures that are constructed using those building blocks are: VGGNet [44], ResNet [18], DenseNet [24], GoogLeNet [45] and MobileNets [23, 43, 22]. Apart from the advances in architecture design, initialization of weights also helps in improving the accuracy of a DNN [11, 17].

We focus on optimizing the configuration of convolutional neural networks (CNNs) and shed light on the selection of the number of filters for each layer given a fixed architecture and depth. Our approach is complementary to the

Work supported by MOST of Taiwan: 107-2221-E-009-125-MY3.

modern variants of CNNs (VGGNet, ResNet, MobileNet, etc.) through the introduction of a guided approach in tuning its width instead of just blindly stacking additional layers to boost accuracy. Our approach also investigates the conventional wisdom on filter selection stating that as we go deeper into the network, more filters are required to capture high-level information embedded in the features and to compensate with the gradual reduction in the spatial dimension which has efficiency in FLOPs as a byproduct.

Intuitively speaking, the design of traditional CNNs is ad-hoc and introduces redundancy [29, 4]. This redundancy gives the opportunity for filter pruning techniques [15, 13, 29, 32] to strive, by conserving or improving accuracy using a lower parameter count. Current advances in pruning has led to a recent work by Liu et al. [33] that studies pruning in a new perspective. They show that pruning can be viewed as an architecture search method instead of just for removing redundancy. We incorporated this perspective along with the recent findings of EfficientNet [47] stating that through the search of an optimal ratio between the width, depth and resolution of a given architecture and dataset, the accuracy of a network can improve if scaled accordingly. Both these works led us to think that instead of finding the scale or ratio among the width, depth and resolution, we can scale the width of a CNN across several layers independently using global iterative pruning as a proxy. We hypothesize that if we are given a DNN with minimal redundant parameters, through the modeling of the change of neurons (filters) in each layer of the DNN with respect to the change in the total parameters of the DNN, we are able to scale the DNN across various sizes efficiently. Our approach can also be viewed as a variant of neural architecture search (NAS) where the search is on finding the optimal configuration of neurons (filters) across layers instead of searching for the optimal structure involving skipconnections or filter types [54, 30, 41, 50, 1, 31]. Our approach is comparatively light-weight as the only resource intensive task lies on the pruning of network. The gist of our proposed method is shown in Figure 1a.

The rest of the paper is structured as follows. We first show related work on available pruning techniques and the role of pruning for neural architecture search in Section 2. We then show the details of our approach in Section 3. Extensive experiments on our proposed method is shown in Section 4. We finally conclude our paper in Section 5.

2. Related Work

Pruning of Deep Neural Network. Pruning of neural networks has been studied for decades with the goal of parameter reduction [27, 16]. Parameter reduction is important if a DNN needs to be deployed to targeted devices with limited storage capacity and computational resources. There are two ways to prune a network: structured

pruning [32, 20, 39, 28, 34, 52] and unstructured pruning [13, 14, 15, 9]. For structured pruning, entire filters or neurons are removed from a layer of a network. Such pruning method can be deployed directly to a general computing architecture, e.g. GPU, to see improvement in speed. For unstructured pruning, individual weights are pruned, resulting in a sparse architecture. A dedicated hardware is required to exploit the speed-up contributed by the sparsity of weights.

To prune a network, there are various criteria that have been studied. The most intuitive approach is to prune the weights based on its magnitude [15, 13, 29]. It was believed that the importance of weights is related to its magnitude. Although this approach is widely used in other works [9, 29, 19, 51], it is also shown on several works [35, 20, 27, 16] that magnitude pruning does not result in an architecture with the best performance. Pruning based on magnitude is adopted because of its simplicity when compared to more complicated pruning techniques, e.g. [16, 27] requires the computation of Hessian for pruning. A study proposed the use of geometric median [20] as a replacement of magnitude pruning for the criteria of network pruning. [35, 37] has also challenged the reliability of magnitude pruning and proposed the use of Taylor expansions to approximate the contribution of a filter to the final loss. Another intuitive way of pruning is through the addition of a regularizer to the loss function to force redundant weights to zero during training [15]. It has also been discovered that the scaling parameter used in Batch Normalization (BN) [25] can be used for structured pruning and yields performance better than pruning using magnitude [32]. A followup work that takes the shift term of BN into consideration for pruning is proposed in [52].

Neural Architecture Search via Pruning. There's a tremendous surge of efforts placed into the research of neural architecture search (NAS) techniques in the recent years on coming up with the most efficient architecture possible for a given task [54, 30, 41, 46, 3, 31]. NAS techniques are usually computationally expensive, limiting its applicability to research or corporate environment with limited computing resources. The search space of NAS is very broad and is defined distinctively across different works. Most of the search space involves the search of a suitable set of operations to be placed in a cell. The connections between different operations is also considered in the search space [2]. These cells are then stacked to increase network depth.

In our work, we focus on the decision of the number of neurons (filters) required for each layer in a DNN. We use existing pruning techniques as a proxy to tackle this problem. The idea of using pruning as an architecture search method is not novel and has been discussed in [33, 9, 10, 7] where its applicability can be seen in MorphNet [12]. Liu *et al.* [33] show that through pruning, we are removing redun-

dancy from a network and the resulting network is efficient in terms of parameters. They also show that training the pruned architecture from scratch has comparable, if not better, accuracy than fine-tuned networks, indicating that the accuracy gain is from the resulting efficient architecture obtained via pruning. For the case of unstructured pruning, it is studied in [9, 10, 53] that a Lottery Ticket (LT) can be found via iterative unstructured pruning. A LT is a sparse architecture that is the result of unstructured pruning and has accuracy better than the original network (usually found at a parameter count of an order less than the parameter count of the original network). This finding indicates that pruning does introduce inductive bias [5] and adds evidence on the suitability of using pruning as an architecture search technique. This idea is proven in MorphNet [12] where a pruned architecture is scaled uniformly to meet the targeted parameter count and is repeated for several iterations. A single iteration of [12] is illustrated in Figure 1b.

3. Method

3.1. Parameter Tracking via Iterative Pruning

To efficiently allocate neurons (filters) across different layers of a DNN that results in optimal accuracy given a parameter constraint, we model the change of neurons (filters) across layers with respect to the change of parameters. First, we need to begin with a network with minimal redundant parameters. To do so, a structured pruning method proposed by Molchanov et al. [35] that prunes iteratively is adopted. They proposed a pruning method that prunes neurons based on its importance. The importance of a parameter can be measured as the loss induced when it's removed from the network. They proposed the use of Taylor approximation as an efficient way to find parameters that are of less importance. A comprehensive comparison between their approach and an oracle (full combinatorial search over all parameters that results in minimum increase in loss) is done, proving its reliability. Here, we will give a brief introduction of their parameter pruning technique borrowed from their paper. The importance of a parameter is quantified by the error induced when it is removed from the network:

$$\mathcal{I}_m = (E(\mathcal{D}, \boldsymbol{W}) - E(\mathcal{D}, \boldsymbol{W} | w_m = 0))^2.$$
(1)

Here, $W = \{w_0, w_1, ..., w_M\}$ are the set of parameters of a neural network supported by a dataset $\mathcal{D} = \{(x_0, y_0), (x_1, y_1), ..., (x_K, y_K)\}$ of K independent samples composed of inputs x_k and outputs y_k . (1) can be approximated by the second-order Taylor expansion as:

$$\mathcal{I}_m^{(2)}(\boldsymbol{W}) = (g_m w_m - \frac{1}{2} w_m \boldsymbol{H}_m \boldsymbol{W})^2.$$
(2)

H is the Hessian matrix where H_m is the *m*-th row of it and $g_m = \frac{\partial E}{\partial w_m}$. (2) can be further approximated using the

first-order expansion:

$$\mathcal{I}_m^{(1)}(\boldsymbol{W}) = (g_m w_m)^2.$$
 (3)

To minimize computational cost, (3) will be used since it is shown in [35] that the performance is on par with the second-order expansion and the first-order Taylor expansion is often used to estimate the importance of DNN components (weights, kernels or filters) [48, 36, 8]. Consistent with their work, a gate z_m is placed after batch normalization layers [25] where the importance approximation is defined as:

$$\mathcal{I}_m^{(1)}(\boldsymbol{z}) = (\frac{\partial E}{\partial \boldsymbol{z}_m})^2 = (\sum_{s \in \mathcal{S}_m} g_m w_m)^2, \qquad (4)$$

where S_m corresponds to the set of weights $W_{s \in S_m}$ before the gate.

For a network composed of L layers, we define the the set of neurons (filters) for the entire network as $\{\phi_l\}_{l=1}^L$. ϕ_l is the number of neurons (filters) of layer l. We then define the total number of parameters in a network as τ . As we are using an iterative pruning method, on every pruning iteration, we will obtain a set of τ 's and ϕ_l 's for the *l*-th layer which can be represented as $\xi_l = \{\tau, \phi_l\}$. After pruning for N iterations, we obtain $\boldsymbol{\xi}_l = \{\{\tau^{(n)}, \phi_l^{(n)}\}_{n=1}^N\}$. In our implementation, we start feeding ξ_l 's into ξ_l when all layers in a network is pruned by at least a single parameter. We conjecture that when all layers are pruned by at least a single parameter, most redundancy is removed and the residual parameters compose an efficient configuration. We stop pruning once the number of neurons (filters) is less than ϵ (we pick ϵ as 5% of the total neurons (filters) of the network in our implementation). Upon the completion of pruning, we have $\boldsymbol{\xi} \in \mathbb{R}^{2 \times N \times L}$. The entire pruning process begins once we have trained our network for P epochs (commonly known as network pre-training; we use the term pre-training epochs in our context instead) using a learning rate μ . The choice of P is studied and the conventional wisdom on the requirement of pre-training a network to convergence before pruning is investigated in the Supp. Section 6.3. The pruning algorithm is summarized in Algorithm 1.

3.2. Efficient Scaling of Parameters

The goal of this work is to scale the neurons (filters) of a network across different layers to satisfy the targeted total parameter size denoted as $\hat{\tau}$. For parameter scaling to match the targeted size, uniform scaling is used in MorphNet [12] and MobileNets [23] where all layers are scaled with a constant width multiplier. It is intuitive that the scale applied to neurons (filters) of different layers should be layer-dependent to maximize performance. In this work, we propose an efficient method to scale the number of neurons (filters) across different layers to maximize performance.

Algorithm 1 Iterative Prune				
1:	procedure IterativePrune(f, D)			
2:	for P epochs do			
3:	Update f using learning rate μ			
4:	while $\sum_{l=1}^{L} oldsymbol{z}_l _1 > \epsilon$ do			
5:	Train f for Q iterations			
6:	$\boldsymbol{\xi}' \leftarrow$ Prune f using criteria (4)			
7:	if all layers pruned at least once then			
8:	$\boldsymbol{\xi} \leftarrow \{\boldsymbol{\xi}, \boldsymbol{\xi}'\}$ \triangleright Record parameters			
9:	return ξ			



Figure 2: Growth in number of filters of different layers across various network sizes. Each color represents independent layers of the convolutional filters of VGG11. Solid line is the residual filters obtained using an iterative pruning method and dashed line represents our approach on curve fitting. It can be observed that there's a pattern in the change of filters w.r.t. parameters which is dataset dependent.

By using iterative pruning as a proxy, we parameterize the change of neurons (filters) across different layers with respect to the total parameters or simply put as $\frac{\partial \phi_l}{\partial \tau}$.

Modeling Parameter Growth. As discussed in the previous subsection, ξ_l collected for each layer resembles the efficient set of neurons (filters) for each layer at a given size constraint. We can use this as a proxy to model $\frac{\partial \phi_l}{\partial \tau}$. We first observe how the distribution of the residual neurons (filters) obtained using pruning scale across different τ 's, e.g. we use VGG11 as our network and CIFAR10/100 as our dataset to show the parameter growth across various sizes in Figure 2. We can observe that parameters across different layers don't scale linearly across different sizes, implying that uniform scaling is not an efficient scaling method. Figure 2 also shows us that the growth of the parameters resembles a power function that is monotonic. To fit the curves, we use:

$$\phi_l(\tau | \alpha_l, \beta_l) = \alpha_l \tau^{\beta_l},\tag{5}$$

where every layer is paramterized independently by α_l and β_l . To obtain these parameters, we can linearize the problem by taking ln on both sides of (5), giving us:

$$\ln \phi_l(\tau | \alpha_l, \beta_l) = \ln \alpha_l + \beta_l \ln \tau.$$
(6)

Since we pruned the network iteratively for N iterations, we obtain a set of linear equations which can be formulated into a matrix of the form:

$$\begin{bmatrix} 1 & \ln \tau^{(1)} \\ 1 & \ln \tau^{(2)} \\ \vdots & \vdots \\ 1 & \ln \tau^{(N)} \end{bmatrix} \underbrace{\begin{bmatrix} \ln \alpha_1 & \ln \alpha_2 & \dots & \ln \alpha_L \\ \beta_1 & \beta_2 & \dots & \beta_L \end{bmatrix}}_{\boldsymbol{\Theta}} = (7)$$

$$\boldsymbol{\tau}$$

$$\begin{bmatrix} \ln \phi_1^{(1)} & \ln \phi_2^{(1)} & \dots & \ln \phi_L^{(1)} \\ \ln \phi_1^{(2)} & \ddots & & \\ \vdots & & & \\ \ln \phi_1^{(N)} & & \ln \phi_L^{(N)} \end{bmatrix}.$$

We can solve for Θ in (7) using the least-squares approach or by taking the pseudoinverse of \mathcal{T} and multiply it with Φ :

$$\boldsymbol{\Theta} = (\boldsymbol{\mathcal{T}}^T \boldsymbol{\mathcal{T}})^{-1} \boldsymbol{\mathcal{T}}^T \boldsymbol{\Phi}.$$
 (8)

By using this method, we are able to fit the curves or ξ_l 's obtained using iterative pruning of a network where the fitted results are shown in Figure 2. Our approach is a cost effective way of neural architecture parametrization and only takes two parameters (α and β) per layer to parameterize the non-linear growth of the neuron (filter) count across various parameter sizes or $\frac{\partial \phi_l}{\partial \tau}$. The simplicity of our approach also prevents the overfitting of noise embedded in the samples obtained via pruning. The search of parameters for efficient scaling is summarized in Algorithm 2.

Algorithm 2 Search Parameters					
1:	procedure SearchParam	$AS(\boldsymbol{\xi})$			
2:	$\mathcal{T}, oldsymbol{\Phi} \leftarrow oldsymbol{\xi}$	▷ Convert to matrix form			
3:	$oldsymbol{\Theta} \leftarrow (oldsymbol{\mathcal{T}}^Toldsymbol{\mathcal{T}})^{-1}oldsymbol{\mathcal{T}}^Toldsymbol{\Phi}$				
4:	return Θ				

Meeting Parameter Constraints. Since our approach fully parameterizes the independent scaling of network width across various sizes, we can meet tight parameter constraints during deployment of a DNN to devices with limited resource budget. For approaches like uniform scaling [12, 33, 23], only the number of output neurons (filters) can be scaled while the network size is a function of the input and output neurons (filters), hence meeting parameter constraints can only be done by performing an iterative grid search on the number of output neurons (filters) required.

For our approach, parameter scaling is intuitive as we can apply gradient descent on Φ w.r.t. τ . To do so, we define a DNN as $f(\boldsymbol{x}|\boldsymbol{W}, \Phi(\tau|\Theta))$ where f is a DNN architecture, \boldsymbol{x} is its input, \boldsymbol{W} are the weights of the DNN, and the

additional condition $\Phi(\tau|\Theta)$ is introduced to parameterize the number of neurons (filters) required for each layer of a DNN. We then define a function h that computes the number of parameters of a DNN. h is architecture dependent.

Given a parameter constraint $\hat{\tau}$ that needs to be met, we can generate an architecture having total parameters close to $\hat{\tau}$ by performing stochastic gradient descent (SGD) on $\Phi(\tau, \Theta)$ w.r.t. τ . Like other gradient descent problems, parameter initialization is important and the best way to do so is to fit $\hat{\tau}$ into (5) giving us:

$$\phi_l(\hat{\tau}|\alpha_l,\beta_l) = \alpha_l \hat{\tau}^{\beta_l}.$$
(9)

This gives us a good initial point, however there will still be a discrepancy between $h(f(\boldsymbol{x}|\boldsymbol{W}, \boldsymbol{\Phi}(\hat{\boldsymbol{\tau}}|\boldsymbol{\Theta})))$ and $\hat{\tau}$ which can be fixed by applying SGD on $\frac{1}{2}(\phi_l(\tau|\alpha_l, \beta_l) - \hat{\tau})^2$ w.r.t. to τ where the update of τ is given as:

$$\tau_i = \tau_{i-1} - \Delta \tau_{i-1} \tag{10}$$

$$= \tau_{i-1} - \eta(h(f(\boldsymbol{x}|\boldsymbol{W}, \boldsymbol{\Phi}(\tau_{i-1}|\boldsymbol{\Theta}))) - \hat{\tau}) \sum_{l=1}^{L} \beta_l \alpha_l \tau_{i-1}^{\beta_l - 1}$$
(11)

Here, the subscript of τ corresponds to the SGD iteration and the full proof of (11) is given in Supp. Section 6.1. Since the number of parameters of an architecture is a monotonic function of Φ , this problem is convex and converges easily. If the learning rate η is set carefully, we are able to obtain an architecture with total parameters close to $\hat{\tau}$. We summarize this procedure in Algorithm 3.

Algorithm 3 Generate Network Using Searched Parameters				
1:	1: procedure GENERATENET($\Theta, f, \hat{\tau}$)			
2:	$\tau \leftarrow \hat{\tau}$ > Initialize parameter			
3:	while not converged do			
4:	$\tau \leftarrow \tau - \Delta \tau \triangleright$ Update using SGD as in (11)			
5:	for $l \leftarrow 1$ to L do \triangleright Layer-wise architecture update			
6:	$f_l \leftarrow$ set output neurons (filters) as $\phi_l(\tau \boldsymbol{\Theta}_l)$ (5)			
7:	return f			

3.3. Architecture Descent for Model Refinement

Like any gradient descent algorithm, initialization plays an important role and affects the convergence of an algorithm. Our approach is similar in way where we attempt to search for the configuration of an architecture given an initial configuration, e.g. VGGNet [44] and ResNet [18] consist of a set of predefined filter numbers for different configurations. As our approach behaves similarly to gradient descent, we coin it *architecture descent* as there is no gradient involved and it is descending in the loss surface by making iterative changes to the architecture's configuration.

We define an iteration of architecture descent as a single run of Algorithm 1, 2 and 3 that corresponds to iterative pruning, parameter searching and network generation. Upon the completion of iterative pruning and parameter searching, we obtain a set of parameters that scales our network in a more efficient manner. We can then use this set of parameters to scale-up our network as shown in Algorithm 3 for further pruning. We then proceed with several iterations of architecture descent until the changes in the architecture configuration is minuscule, indicating convergence. By applying architecture descent, we are descending on the loss surface that is parameterized by Θ instead of a loss surface parameterized by its weights W performed in gradient descent. Architecture descent is summarized in Algorithm 4. NeuralScale is a composition of all algorithms we proposed as illustrated in Figure 1a.

Algorithm 4 Architecture Descent					
1:	1: procedure ARCHITECTUREDESCENT (f, \mathcal{D}, τ)				
2:	2: while not converged do				
3:	$\boldsymbol{\xi} \leftarrow \text{IterativePrune}(f, \mathcal{D}) \triangleright \text{Taylor FO} [35]$				
4:	: $\Theta \leftarrow \text{SearchParams}(\boldsymbol{\xi})$				
5:	$f \leftarrow \text{Gen}$	NERATENET $(\mathbf{\Theta}, f, au)$			
6:	return f	▷ Network with scaled parameters			

4. Experiments

In this section, we show experiments illustrating the importance of architecture descent. We then proceed with the benchmarking of our approach using public datasets, e.g. CIFAR10/100 [26] and TinyImageNet (subset of ImageNet [6] with images downsampled to 64×64 and consists of 200 classes having 500 training and 50 validation samples for each class) on commonly used architectures, e.g. VGG11 [44], MobileNetV2 [43] and ResNet18 [18]. All experiments are run on a single GTX1080Ti GPU using PyTorch [40] as our deep learning framework. We use SGD as our optimizer with an initial learning rate of 0.1, momentum set to 0.5 and a weight decay factor of 5^{-4} . Training of network that uses CIFAR10 and CIFAR100 are run for 300 epochs using a step decay of learning rate by at factor of 10 at epochs 100, 200 and 250 whereas network trained using TinyImageNet are run for 150 epochs with a decay in learning rate by a factor of 10 at epochs 50 and 100. For iterative pruning, we first train our network for P = 10 epochs using a learning rate of 0.1 and is decayed by a factor of 10 every 10 epochs. Source code is made available at https: //github.com/eugenelet/NeuralScale.

4.1. Importance of Architecture Descent

For all experiments, we run architecture descent for 15 iterations. We show configurations with total parameters

matching total parameters of network with its default set of filters uniformly scaled to a ratio, r. r = 0.25, 2 for CIFAR10 and CIFAR100. r = 0.25, 1 for TinyImageNet.

VGG11. Using a relatively shallow network, we demonstrate the application of architecture descent using CI-FAR10 and CIFAR100 as shown in Figure 3. By observing the resulting architecture configuration, we can make two conjectures. First, we show that conventional wisdom on network design that gradually increases the number of filters as we go deeper in a convolutional network does not guarantee optimal performance. It can be observed that the conventional wisdom on network design holds up to some level (layer 4) and bottlenecking of parameters can be observed up to the penultimate layer which is followed by a final layer which is comparatively larger. Second, the scaling of network should not be done linearly as was done in [23] and should follow a non-linear rule that we attempt to approximate using a power function. If we look closely, by applying architecture descent on datasets of higher complexity generates network with configuration that has more filters allocated toward the end. Our conjecture is that more resources are needed to capture the higher level features when the task is more difficult whereas for simple classification problem like CIFAR10, more resource is allocated to earlier layers to generate more useful features. These observations give us a better understanding on how resource should be allocated in DNNs and can be used as a guideline for deep learning practitioners in designing DNNs. A single iteration of architecture descent for VGG11 on CI-FAR10/100 is approximately 20 minutes.

MobileNetV2. We show the application of architecture descent on a more sophisticated architecture known as MobileNetV2 using CIFAR100 and TinyImageNet in Figure 3. Here, we only apply our search algorithm on deciding the size of the bottleneck layers while the size of the expansion layer follows the same expansion rule found in [43] where an expansion factor of $6 \times$ is used. The resulting configuration closely resembles the one found using a feedforward network like VGG. It can also be observed that resources are allocated toward the output for a more sophisticated dataset. A single iteration of architecture descent for MobileNetV2 on CIFAR100 and TinyImageNet is approximately 50 minutes and 1.2 hour respectively.

ResNet18. The application of architecture descent on ResNet18 using CIFAR100 and TinyImageNet is shown in Figure 3. We observe a different pattern of architecture configuration when compared to a simple feed forward network like VGG. This is an interesting observation as it agrees with the interpretation of residual networks as an ensemble of many paths of different lengths shown in [49]. An-



Figure 3: Shows the number of filters for each layer by running architecture descent for 15 iterations on various architecture-dataset pair. Vertical and horizontal axis of each plot corresponds to the filter number and architecture descent iteration respectively. r is the uniform scaling ratio applied to the default configuration. Best viewed in color.

other observation is that if we look only at a single layer of every residual block (each block consists of two layers), the searched configuration for ResNet follows the pattern found in VGG where there's a smooth gradient of filter progression across layers. A single iteration of architecture descent for ResNet18 on CIFAR100 and TinyImageNet is approximately 50 minutes and 45 minutes respectively.

4.2. Benchmarking of NeuralScale

Here, we compare the accuracy of NeuralScale with the first (Iteration=1, to compare with MorphNet [12]) and last (Iteration=15) iteration of architecture descent with a uniform scaling (baseline) method and a method where a network is first pruned using Taylor-FO [35] until it has 50%

of filters left and then scaled uniformly (resembling the first iteration of MorphNet [12] and the use of [35] is to match our pruning method for a fair comparison), named as MorphNet (Taylor-FO) in our comparison tables and plots. The accuracy is obtained by averaging across the maximum test accuracy of 5 independent simulations. For the accuracy plots in Figure 4 and 5, the output filters of the original network are scaled to the ratios from 0.25 to 2 with an increment of 0.25 for CIFAR10/100 along with 0.25, 0.5, 0.75 and 1 for TinyImageNet. For the comparison tables, only the ratios 0.25, 0.75, 2 for CIFAR10/100 and 0.25, 0.75 for TinyImageNet are reported. Comparisons are also made with a structured magnitude pruning method [29] where we first pre-train our network using the same prescription for other methods and proceed with 40 and 20 epochs of fine-tuning for CIFAR10/100 and Tiny-ImageNet respectively, using a learning rate of 0.001. We only show results for VGG11-CIFAR10, MobileNetV2-CIFAR100 and ResNet18-TinyImageNet in Table 1 and the rest are deferred to Supp. Section 6.2. Note that all methods are trained from scratch and only [29] is trained using the pretrain-prune-finetuning pipeline. Results show that the hypothesis in [33] holds (training from scratch performs better). As our approach is designed for platforms with structured parallel computing capability like GPUs, we report the latency of different methods instead of FLOPs. Note that our approach isn't optimized for latency. Here, latency is defined as the time required for an input to propagate to the output. All latencies reported in Table 1 are based on a batch size of 100 where 20 batches are first fed for warm-up of cache and is proceeded with 80 batches which are averaged to give the final latencies. As a comparison of latency based solely on the scale of parameter in Table 1 is unintuitive, we show a plot comparing accuracy of different methods against latency in Figure 5.

VGG11. By observing the comparison plot shown in Figure 4a and 4b, our approach compares favourably in terms of parameter efficiency for CIFAR10 and CIFAR100. As shown in Table 1, at the lowest parameter count, an accuracy gain of 3.04% is obtained for CIFAR10. Efficiency in latency of our approach is also comparable with the baseline approach as shown in Figure 5a and 5b. Diminishing returns are observed when the network increases in size. We conjecture that as the network grows larger, more subspaces are covered, hence the network can still adapt to the suboptimal configuration by suppressing redundant parameters. Another observation is that the performance gain is more substantial on a more complicated dataset which is intuitive as inductive bias is introduced in an architectural level.

MobileNetV2. The application to MobileNetV2 is to show the extensibility of our approach to a delicately hand-



Figure 4: All plots are organized such that each row corresponds to a single architecture, e.g. (a),(b) corresponds to VGG11, (c),(d) corresponds to MobileNetV2 and (e),(f) corresponds to ResNet18. Each plot consists of the accuracy comparison of different scaling methods (applied on the width), plotted against different parameter scales. The shaded region of each line corresponds to the maximum and minimum accuracy across 5 independent simulations.

crafted architecture. Our approach is superior in parameter efficiency (most cases) when compared to other methods as shown in Figure 4c and 4d. As shown in Table 1, an accuracy gain of 8.56% for CIFAR100 relative to baseline is observed at a scaling ratio of 0.25. Our approach is also efficient in latency as shown in Figure 5c. An unintuitive observation can be seen on the experiment for TinyImageNet where the accuracy at iteration 1 outperforms iteration 15 for NeuralScale at ratio 0.25. The accuracy is below the baseline when more parameters are used. As the default configuration of MobileNetV2 has consistent filters, the default shortcut connections are identity mappings. We hypothesize that the switch from identity mapping to a convolutional mapping for the shortcut connection is the culprit. Empirical study is done in Supp. Section 6.4 with results supporting our hypothesis and explaining the observations.



(e) ResNet18 on CIFAR100.

(f) ResNet18 on TinyImageNet.

Figure 5: The layout of these plots are structured as shown in Figure 4. The only difference is that the accuracies of different methods are plotted against latencies.

ResNet18. As shown from the accuracy comparison plot in Figure 4e and 4f, substantial accuracy gain under different parameter scales is observed. From Table 1, an accuracy gain of 3.41% is observed for TinyImageNet at a scale of 0.25. Accuracy gain using architecture descent is also more pronounced here. The accuracy gain here is in contradictory with the results in [33] (no gain in accuracy observed) probably due to the better pruning technique we use [35]. In Figure 5e and 5f our approach is comparable and in most cases better than the baseline configuration in latency.

5. Conclusion

In this work, we propose a method to efficiently scale the neuron (filter) number of an architecture. We hypothesize that the scaling of network should follow a non-linear rule and is shown empirically that through our approach, networks efficient in parameters across different scales can be generated using this rule. Our empirical results on architecture descent also shed light on the efficient allocation of parameters in a deep neural network. As our approach is computationally-efficient and is complementary to most

TT 1 1 1	a ·	C	•		•
Table 1.	(omnarison	OT.	Various	network_dataset	naire
rable r.	Comparison	O1	various	network under	pans.

Method	Params	Latency	Accuracy (%)		
	VGG11 C	IFAR10			
Uniform Coolo	0.58M	1.29ms	88.18 ± 0.16		
(Pasalina)	5.20M	4.31ms	91.64 ± 0.10		
(Baseline)	36.89M	18.86ms	92.96 ± 0.09		
Li <i>et al</i> . [29] [†]	5.20M	4.75ms	91.12 ± 0.02		
MorphNet [12]	0.58M	2.07ms	90.60 ± 0.11		
(Taylor-FO [35])	5.22M	8.43ms	92.60 ± 0.09		
(Taylor-1 O [33])	36.72M	48.52ms	93.18 ± 0.11		
NouralSaala	0.58M	2.56ms	91.13 ± 0.07		
(Iteration - 1)	5.20M	8.89ms	92.61 ± 0.15		
(Iteration = 1)	36.88M	37.39ms	$\textbf{93.31} \pm \textbf{0.05}$		
Name 10 and a	0.58M	2.94ms	$\textbf{91.22} \pm \textbf{0.15}$		
NeuralScale	5.20M	12.52ms	$\textbf{92.63} \pm \textbf{0.12}$		
(Iteration = 15)	36.90M	53.35ms	93.29 ± 0.09		
Mo	bileNetV2	CIFAR100)		
Uniform Carls	0.20M	5.37ms	57.80 ± 0.31		
(Pasalina)	1.42M	7.46ms	67.85 ± 0.38		
(Baseline)	9.30M	19.69ms	72.40 ± 0.22		
Li <i>et al</i> . [29] [†]	1.42M	7.71ms	67.12 ± 0.08		
MorphNat [12]	0.20M	6.14ms	57.51 ± 0.36		
(Taylor FO [35])	1.42M	10.95ms	67.51 ± 0.48		
(Taylor-10 [33])	9.30M	26.53ms	72.29 ± 0.28		
NauralCaala	0.19M	5.69ms	66.00 ± 0.12		
NeuralScale $(Iteration = 1)$	1.40M	11.73ms	70.23 ± 0.25		
(Iteration = 1)	9.21M	21.32ms	72.37 ± 0.12		
N	0.19M	7.84ms	$\textbf{66.36} \pm \textbf{0.28}$		
NeuralScale	1.41M	17.89ms	$\textbf{71.94} \pm \textbf{0.45}$		
(Iteration = 15)	9.27M	40.48ms	$\textbf{74.73} \pm \textbf{0.26}$		
ResNet18 TinyImageNet					
Uniform Scale	0.73M	3.02ms	50.54 ± 0.37		
(Baseline)	6.36M	11.56ms	56.68 ± 0.28		
Li et al. [29] [†]	6.36M	11.93ms	54.72 ± 0.24		
MorphNet [12]	0.72M	3.80ms	50.79 ± 0.38		
(Taylor-FO [35])	6.39M	14.83ms	56.78 ± 0.85		
NeuralScale	0.72M	5.96ms	$\overline{51.66\pm0.80}$		
(Iteration = 1)	6.42M	14.58ms	57.89 ± 0.28		
NeuralScale	0.72M	6.42ms	$\textbf{53.95} \pm \textbf{0.53}$		
(Iteration = 15)	6.40M	17.52ms	$\textbf{58.40} \pm \textbf{0.54}$		

[†] Fine-tuned using pre-trained network (not trained from scratch).

techniques and architectures, the inclusion to existing deep learning framework is cost-effective and results in a model of higher accuracy under the same parameter constraint.

References

- Irwan Bello, Barret Zoph, Vijay Vasudevan, and Quoc V Le. Neural optimizer search with reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 459–468. JMLR. org, 2017. 2
- [2] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *International Conference on Machine Learning*, pages 549–558, 2018. 2
- [3] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. arXiv preprint arXiv:1812.00332, 2018. 2
- [4] Michael Cogswell, Faruk Ahmed, Ross Girshick, Larry Zitnick, and Dhruv Batra. Reducing overfitting in deep networks by decorrelating representations. arXiv preprint arXiv:1511.06068, 2015. 2
- [5] Nadav Cohen and Amnon Shashua. Inductive bias of deep convolutional networks through pooling geometry. arXiv preprint arXiv:1605.06743, 2016. 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 2009 IEEE conference on computer vision and pattern recognition, pages 248–255. Ieee, 2009. 5
- [7] Xiaohan Ding, Guiguang Ding, Yuchen Guo, Jungong Han, and Chenggang Yan. Approximated oracle filter pruning for destructive cnn width optimization. arXiv preprint arXiv:1905.04748, 2019. 2
- [8] Xiaohan Ding, Xiangxin Zhou, Yuchen Guo, Jungong Han, Ji Liu, et al. Global sparse momentum sgd for pruning very deep neural networks. In *Advances in Neural Information Processing Systems*, pages 6379–6391, 2019. 3
- [9] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. arXiv preprint arXiv:1803.03635, 2018. 2, 3
- [10] Jonathan Frankle, Gintare Karolina Dziugaite, Daniel M Roy, and Michael Carbin. The lottery ticket hypothesis at scale. arXiv preprint arXiv:1903.01611, 2019. 2, 3
- [11] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256, 2010. 1
- [12] Ariel Gordon, Elad Eban, Ofir Nachum, Bo Chen, Hao Wu, Tien-Ju Yang, and Edward Choi. Morphnet: Fast & simple resource-constrained structure learning of deep networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1586–1595, 2018. 1, 2, 3, 4, 6, 7, 8, 11, 17
- [13] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. arXiv preprint arXiv:1510.00149, 2015. 2
- [14] Song Han, Jeff Pool, Sharan Narang, Huizi Mao, Shijian Tang, Erich Elsen, Bryan Catanzaro, John Tran, and William J Dally. Dsd: regularizing deep neural networks with dense-sparse-dense training flow. arXiv preprint arXiv:1607.04381, 3(6), 2016. 2

- [15] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In Advances in neural information processing systems, pages 1135–1143, 2015. 2
- [16] Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeon. In Advances in neural information processing systems, pages 164–171, 1993. 2
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. 1
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 770–778, 2016. 1, 5
- [19] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 784– 800, 2018. 2
- [20] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019. 2
- [21] Suzana Herculano-Houzel. The human brain in numbers: a linearly scaled-up primate brain. *Frontiers in human neuro-science*, 3:31, 2009.
- [22] 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. arXiv preprint arXiv:1905.02244, 2019. 1
- [23] 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. 1, 3, 4, 6
- [24] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In Proceedings of the IEEE conference on computer vision and pattern recognition, pages 4700–4708, 2017. 1
- [25] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. arXiv preprint arXiv:1502.03167, 2015. 2, 3
- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. Technical report, Citeseer, 2009. 5
- [27] Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In Advances in neural information processing systems, pages 598–605, 1990. 2
- [28] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. arXiv preprint arXiv:1810.02340, 2018. 2
- [29] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710, 2016. 2, 7, 8, 11

- [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 Proceedings of the European Conference on Computer Vision (ECCV), pages 19–34, 2018. 2
- [31] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 2
- [32] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In Proceedings of the IEEE International Conference on Computer Vision, pages 2736–2744, 2017. 2
- [33] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. *arXiv preprint arXiv:1810.05270*, 2018. 2, 4, 7, 8
- [34] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pages 5058–5066, 2017. 2
- [35] Pavlo Molchanov, Arun Mallya, Stephen Tyree, Iuri Frosio, and Jan Kautz. Importance estimation for neural network pruning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11264–11272, 2019. 2, 3, 5, 6, 7, 8, 11, 17
- [36] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. arXiv preprint arXiv:1611.06440, 2016. 3
- [37] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient transfer learning. arXiv preprint arXiv:1611.06440, 3, 2016. 2
- [38] Vinod Nair and Geoffrey E Hinton. Rectified linear units improve restricted boltzmann machines. In *Proceedings of the* 27th international conference on machine learning (ICML-10), pages 807–814, 2010. 1
- [39] Kirill Neklyudov, Dmitry Molchanov, Arsenii Ashukha, and Dmitry P Vetrov. Structured bayesian pruning via log-normal multiplicative noise. In Advances in Neural Information Processing Systems, pages 6775–6784, 2017. 2
- [40] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in PyTorch. In *NIPS Autodiff Workshop*, 2017.
 5
- [41] Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. arXiv preprint arXiv:1802.03268, 2018. 2
- [42] Prajit Ramachandran, Barret Zoph, and Quoc V Le. Searching for activation functions. arXiv preprint arXiv:1710.05941, 2017. 1
- [43] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 1, 5, 6

- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556, 2014. 1, 5
- [45] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [46] 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 Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pages 2820–2828, 2019. 2
- [47] Mingxing Tan and Quoc V Le. Efficientnet: Rethinking model scaling for convolutional neural networks. arXiv preprint arXiv:1905.11946, 2019. 2
- [48] Lucas Theis, Iryna Korshunova, Alykhan Tejani, and Ferenc Huszár. Faster gaze prediction with dense networks and fisher pruning. arXiv preprint arXiv:1801.05787, 2018. 3
- [49] Andreas Veit, Michael J Wilber, and Serge Belongie. Residual networks behave like ensembles of relatively shallow networks. In Advances in neural information processing systems, pages 550–558, 2016. 6
- [50] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. arXiv preprint arXiv:1812.09926, 2018. 2
- [51] Tien-Ju Yang, Yu-Hsin Chen, and Vivienne Sze. Designing energy-efficient convolutional neural networks using energyaware pruning. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 5687– 5695, 2017. 2
- [52] Chenglong Zhao, Bingbing Ni, Jian Zhang, Qiwei Zhao, Wenjun Zhang, and Qi Tian. Variational convolutional neural network pruning. In *Proceedings of the IEEE Conference* on Computer Vision and Pattern Recognition, pages 2780– 2789, 2019. 2
- [53] Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. arXiv preprint arXiv:1905.01067, 2019. 3
- [54] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. arXiv preprint arXiv:1611.01578, 2016. 2