

Zen-NAS: A Zero-Shot NAS for High-Performance Image Recognition

Ming Lin

Alibaba Group

Bellevue, Washington, USA

ming.l@alibaba-inc.com

Pichao Wang

Alibaba Group

Bellevue, Washington, USA

pichao.wang@alibaba-inc.com

Zhenhong Sun

Alibaba Group

Hangzhou, Zhejiang, China

zhenhong.szh@alibaba-inc.com

Hesen Chen

Alibaba Group

Hangzhou, Zhejiang, China

hesen.chs@alibaba-inc.com

Xiuyu Sun

Alibaba Group

Hangzhou, Zhejiang, China

xiuyu.sxy@alibaba-inc.com

Qi Qian

Alibaba Group

Bellevue, Washington, USA

qi.qian@alibaba-inc.com

Hao Li

Alibaba Group

Hangzhou, Zhejiang, China

lihao.lh@alibaba-inc.com

Rong Jin

Alibaba Group

Hangzhou, Zhejiang, China

jinrong.jr@alibaba-inc.com

Abstract

Accuracy predictor is a key component in Neural Architecture Search (NAS) for ranking architectures. Building a high-quality accuracy predictor usually costs enormous computation. To address this issue, instead of using an accuracy predictor, we propose a novel zero-shot index dubbed Zen-Score to rank the architectures. The Zen-Score represents the network expressivity and positively correlates with the model accuracy. The calculation of Zen-Score only takes a few forward inferences through a randomly initialized network, without training network parameters. Built upon the Zen-Score, we further propose a new NAS algorithm, termed as Zen-NAS, by maximizing the Zen-Score of the target network under given inference budgets. Within less than half GPU day, Zen-NAS is able to directly search high performance architectures in a data-free style. Comparing with previous NAS methods, the proposed Zen-NAS is magnitude times faster on multiple server-side and mobile-side GPU platforms with state-of-the-art accuracy on ImageNet. Searching and training code as well as pre-trained models are available from <https://github.com/idstcv/ZenNAS>.

1. Introduction

The design of high-performance deep neural networks is a challenging task. Neural Architecture Search (NAS) methods facilitate this progress. There are mainly two key

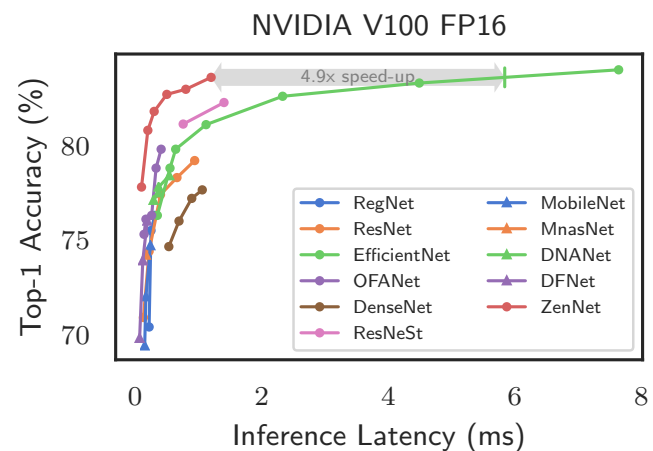


Figure 1. ZenNets top-1 accuracy v.s. inference latency (milliseconds per image) on ImageNet. Benchmarked on NVIDIA V100 GPU, half precision (FP16), batch size 64, searching cost 0.5 GPU day.

components, architecture generator and accuracy predictor, in existing NAS algorithms. The generator proposes potential high-performance networks and the predictor predicts their accuracies. Popular generators include uniform sampling [13], evolutionary algorithm [41] and reinforcement learning [30]. The accuracy predictors include brute-force methods [42, 57, 3, 41], predictor-based methods [30, 56, 29] and one-shot methods [26, 61, 69, 62, 57, 59, 6, 66, 54, 5].

A major challenge of building a high-quality accuracy predictor is the enormous computational cost. Both brute-

forced methods and predictor-based methods require to train considerable number of networks. The one-shot methods reduce the training cost via parameter sharing. Albeit being more efficient than brute-forced methods, the one-shot methods still need to train a huge supernet which is still computationally expensive. Recent studies also find that nearly all supernet-based methods suffer from model interfering [5, 63] which degrades the quality of accuracy predictor [46]. In addition, since the supernet must be much larger than the target network, it is difficult to search large target networks under limited resources. These issues make the one-shot methods struggling in designing high-performance networks.

To solve these problems, instead of using an expensive accuracy predictor, we propose an almost zero-cost proxy, dubbed Zen-Score, for efficient NAS. The Zen-Score measures the expressivity [39, 31] of a deep neural network and positively correlates with the model accuracy. The computation of Zen-Score only takes a few forward inferences on randomly initialized network using random Gaussian inputs, making it extremely fast, lightweight and data-free. Moreover, Zen-Score deals with the scale-sensitive problem caused by Batch Normalization (BN)[4, 35], making it widely applicable to real-world problems.

Based on Zen-Score, we design a novel Zen-NAS algorithm. It maximizes the Zen-Score of the target network within inference budgets. Zen-NAS is a **Zero-Shot** method since it does not optimize network parameters during search¹. We apply Zen-NAS to search optimal networks under various inference budgets, including inference latency, FLOPs (Floating Point Operations) and model size, and achieve the state-of-the-art (SOTA) performance on CIFAR-10/CIFAR-100/ImageNet, outperforming previous human-designed and NAS-designed models by a large margin. Zen-NAS is the first zero-shot method that achieves SOTA results on large-scale full-resolution ImageNet-1k dataset [12] by the time of writing this work [32, 1, 7].

Our approach is inspired by recent advances in deep learning studies [34, 11, 23, 39, 9, 28, 31, 44, 47, 14, 60] which show that deep models are superior than shallow ones since deep models are more expressive under the same number of neurons. According to the bias-variance trade-off in statistical learning theory [19], increasing the expressivity of a deep network implies smaller bias error. When the size n of training dataset is large enough, the variance error will diminish as $\mathcal{O}(1/\sqrt{n}) \rightarrow 0$. This means that the generalization error is dominated by the bias error which could be reduced by more expressive networks. These theoretical results are well-aligned with large-scale deep learning practices[36, 52, 37].

We summarize our main contributions as follows:

- We propose a novel zero-shot proxy Zen-Score for NAS. The proposed Zen-Score is computationally efficient and is proved to be scale-insensitive in the present of BN. A novel NAS algorithm termed Zen-NAS is proposed to search for networks with maximal Zen-Score in the design space.
- Within half GPU day, the ZenNets designed by Zen-NAS achieve up to 83.6% top-1 accuracy on ImageNet that is as accurate as EfficientNet-B5 with inference speed magnitude times faster on multiple hardware platforms. To our best knowledge, **Zen-NAS is the first zero-shot method that outperforms training-based methods on ImageNet.**

2. Related Work

We briefly review the related works. For comprehensive review of NAS, the monograph [43] is referred to.

In the early days of NAS, brute-force methods are adopted to search architectures by directly training a network to obtain its accuracy. For example, the AmoebaNet [41] conducts structure search on CIFAR-10 using Evolutionary Algorithm (EA) [20] and then transfers the structure to ImageNet. It takes about 3150 GPU days of searching and achieves 74.5% top-1 accuracy on ImageNet. Inspired by the success of AmoebaNet, many EA-based NAS algorithms are proposed to improve the searching efficiency, such as EcoNAS [69], CARS [62], GeNet [57] and PNAS [25]. These methods search on down-sampled images or reduce the number of queries. Reinforced Learning is another popular generator (sampler) in NAS, including NAS-Net [70], Mnasnet [49] and MetaQNN [3].

Both EA and RL based methods require lots of network training. To address this problem, the predictor-based methods encode architectures into high dimensional vectors. A number of architectures are trained to obtain their accuracies [30, 29] and then are used as training data for learning accuracy predictor. The one-shot methods further reduce the training cost by training a big supernet. This framework is widely applied in many efficient NAS methods, including DARTS [26], SNAS [59], PC-DARTS [61], ProxylessNAS [6], GDAS [66], FBNetV2 [54], DNANet [21], Single-Path One-Shot NAS [13].

Although the above efforts have greatly reduced the searching cost, their top-1 accuracies on ImageNet are below 80.0%. The authors of OFANet [5] noted that weight-sharing suffers from model interfering. They propose a progressive-shrinking strategy to address this issue. The resultant OFANet achieves 80.1% accuracy after searching for 51.6 GPU days. EfficientNet [50] is another high precision network designed by NAS. It takes about 3800 GPU days to search EfficientNet-B7 whose accuracy is 84.4%. In comparison, Zen-NAS achieves 83.6% accuracy while using magnitude times fewer resources.

¹Obviously, the final searched architecture must be trained on the target dataset before deployment.

A few on-going works are actively exploring zero-shot proxies for efficient NAS. However, these efforts have not delivered the SOTA results. In a recent empirical study, [1] evaluates the performance of six zero-shot pruning proxies on NAS benchmark datasets. The synflow [51] achieves best results in their experiments. We compare synflow with Zen-Score under fair settings and show that Zen-Score achieves +1.1% better accuracy on CIFAR-10 and +8.2% better accuracy on CIFAR-100. The concurrent work TE-NAS [7] uses a combination of NTK-score and network expressivity as NAS proxy. Specifically, the TE-NAS estimates the expressivity by directly counting the number of active regions R_N on randomly sampled images. In comparison, Zen-Score not only considers the distribution of linear regions but also considers the Gaussian complexity of linear classifier in each linear region, giving a more accurate estimation of network expressivity. The computation of Zen-Score is 20 to 28 times faster than TE-NAS score. In terms of performance, TE-NAS achieves 74.1% top-1 accuracy on ImageNet, lagging behind SOTA baselines. Zen-NAS achieves +9.5% better accuracy within similar searching cost. Another concurrent work NASWOT [33] computes the architecture score according to the kernel matrix of binary activation patterns between mini-batch samples. It achieves similar top-1 accuracies on CIFAR-10/CIFAR-100 as TE-NAS.

It is important to distinguish Zen-NAS from unsupervised NAS (UnNAS) [24]. In UnNAS, the network is trained to predict the pre-text tasks therefore it still requires parameter training. In Zen-NAS, no parameter training is required during the search.

In this work, we mostly focus on the vanilla network space described in the next section. Several previous works design networks in a more general irregular design space, such as DARTS [26] and RandWire [58]. Zen-NAS cannot be applied to these irregular design spaces since Zen-Score is not mathematically well-defined in irregular design spaces. In practice, the vanilla network space is a large enough space which covers most SOTA networks, including but not limited to ResNet, MobileNet and EfficientNet. Particularly, Zen-NAS outperforms DARTS-based methods by a significant margin on ImageNet.

3. Expressivity of Vanilla Network

In this section, we discuss how to measure the expressivity of *vanilla convolutional neural network* (VCNN) family, an ideal prototype for theoretical studies. We show that the expressivity of a network can be efficiently measured by its expected Gaussian complexity, or Φ -score for short. In the next section, we further show that for very deep networks, directly computing Φ -score incurs numerical overflow. This overflow can be addressed by adding BN layers and then rescaling the Φ -score by a constant. This new score is named

as Zen-Score in Section 4.

3.1. Notations

An L -layer neural network is formulated as a function $f : \mathbb{R}^{m_0} \rightarrow \mathbb{R}^{m_L}$ where m_0 is the input dimension and m_L is the output dimension. $\mathbf{x}_0 \in \mathbb{R}^{m_0}$ denotes the input image. Correspondingly, the output feature map of the t -th layer is denoted by \mathbf{x}_t . The t -th layer has m_{t-1} input channels and m_t output channels. The convolutional kernel is $\theta_t \in \mathbb{R}^{m_t \times m_{t-1} \times k \times k}$. The image resolution is $H \times W$. The mini-batch size is B . The Gaussian distribution of mean μ and variance σ^2 is denoted by $\mathcal{N}(\mu, \sigma)$.

3.2. Vanilla Convolutional Neural Network

The **vanilla convolutional neural network** (VCNN) is a widely used prototype in theoretical studies [39, 47, 14]. The main body of a vanilla network is stacked by multiple convolutional layers. Each layer consists of one convolutional operator followed by RELU activation. All other components are removed from the backbone, including residual link and Batch Normalization. After the main body, global average pool layer (GAP) reduces the feature map resolution to 1x1, followed by a fully-connected layer. At the end a soft-max operation converts the network output to label prediction. Given the input \mathbf{x} and network parameters θ , $f(\mathbf{x}|\theta)$ refers to the output of the main body of the network, that is the feature map before the GAP layer (pre-GAP layer). We measure the network expressivity with respect to pre-GAP because it contains most of the information we need.

Modern networks use auxiliary structures such as residual link, Batch Normalization and self-attention block [16]. These structures will not significantly affect the representation power of networks. For example, BN layer can be merged into convolutional kernel via kernel fusion. Self-attention linearly combines existing feature maps hence spans the same subspace. Therefore, these structures are temporarily removed when measuring network expressivity and then added back in training and testing stages. For non-RELU activation functions, they are replaced by RELU in a similar way. These simple modifications make our method applicable to a majority of non-VCNN models widely used in practice. In fact, nearly all single-branch feed-forward networks can be converted to vanilla network by the aforementioned modifications.

3.3. Φ -Score as Proxy of Expressivity

Given a VCNN $f(\mathbf{x}|\theta)$, we propose a novel numerical index Φ -score as a proxy of its expressivity. The definition of Φ -score is inspired by recent theoretical studies on deep network expressivity [47, 60]. A key observation in these studies is that a vanilla network can be decomposed

into piece-wise linear functions conditioned on activation patterns [34]:

Lemma 1 ([34, 31]). *Denote the activation pattern of the t -th layer as $\mathcal{A}_t(\mathbf{x})$. Then for any vanilla network $f(\cdot)$,*

$$f(\mathbf{x}|\boldsymbol{\theta}) = \sum_{\mathcal{S}_i \in \mathcal{S}} \mathcal{I}_{\mathbf{x}}(\mathcal{S}_i) \mathbf{W}_{\mathcal{S}_i} \mathbf{x} \quad (1)$$

where \mathcal{S}_i is a convex polytope depending on $\{\mathcal{A}_1(\mathbf{x}), \mathcal{A}_2(\mathbf{x}), \dots, \mathcal{A}_L(\mathbf{x})\}$; \mathcal{S} is a finite set of convex polytopes in \mathbb{R}^{m_0} ; $\mathcal{I}_{\mathbf{x}}(\mathcal{S}_i) = 1$ if $\mathbf{x} \in \mathcal{S}_i$ otherwise zero; $\mathbf{W}_{\mathcal{S}_i}$ is a coefficient matrix of size $\mathbb{R}^{m_L \times m_0}$.

According to Lemma 1, any vanilla network is an ensemble of piece-wise linear functions segmented by convex polytopes $\mathcal{S} = \{\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_{|\mathcal{S}|}\}$ where $|\mathcal{S}|$ is the number of linear-regions (see Figure 2 in [14]. The number of linear regions $|\mathcal{S}|$ has been used as expressivity proxy in several theoretical studies [34, 47, 14, 67, 60]. However, directly using $|\mathcal{S}|$ incurs two limitations: a) Counting $|\mathcal{S}|$ for large network is computationally infeasible; b) The representation power of each $\mathbf{W}_{\mathcal{S}_i}$ is not considered in the proxy. The first limitation is due to fact that the number of linear regions grow exponentially for large networks [34, 60]. To understand the second limitation, we recall the Gaussian complexity [18] of linear classifiers:

Lemma 2 ([18]). *For linear function class $\{f : f(X) = WX \text{ s.t. } \|W\|_F \leq G\}$, its Gaussian complexity is upper bounded by $\mathcal{O}(G)$.*

In other words, Lemma 2 says that the expressivity of linear function class measured by Gaussian complexity is controlled by the Frobenius norm of its parameter matrix W . Inspired by Lemma 1 and Lemma 2, we define the following index for measuring network expressivity :

Definition 1 (Φ -score for VCNN). *The expected Gaussian complexity for a vanilla network $f(\cdot)$ is defined by*

$$\Phi(f) = \log \mathbb{E}_{\mathbf{x}, \boldsymbol{\theta}} \left\{ \sum_{\mathcal{S}_i \in \mathcal{S}} \mathcal{I}_{\mathbf{x}}(\mathcal{S}_i) \|\mathbf{W}_{\mathcal{S}_i}\|_F \right\} \quad (2)$$

$$= \log \mathbb{E}_{\mathbf{x}, \boldsymbol{\theta}} \|\nabla_{\mathbf{x}} f(\mathbf{x}|\boldsymbol{\theta})\|_F. \quad (3)$$

In Definition 1, we measure the network expressivity by its expected Gaussian complexity, or Φ -score for short. Since any VCNN is ensemble of linear functions, it is nature to measure its expressivity by averaging the Gaussian complexity of linear function in each linear region. To this end, we randomly sample \mathbf{x} and $\boldsymbol{\theta}$ from some prior distributions and then average $\|\mathbf{W}_{\mathcal{S}_i}\|_F$. This is equivalent to compute the expected gradient norm of f with respect to input \mathbf{x} . In our implementation, \mathbf{x} and $\boldsymbol{\theta}$ are sampled from standard Gaussian distribution which works well in practice. It

is important to note that in Φ -score, only the gradient of \mathbf{x} rather than $\boldsymbol{\theta}$ is involved. This is different to zero-cost proxies in [1] which compute gradient of $\boldsymbol{\theta}$ in their formulations. These proxies measure the trainability [55, 7] instead of the expressivity of networks.

4. Zen-Score and Zen-NAS

In this section, we show that directly computing Φ -score for very deep networks incurs numerical overflow due to the gradient explosion without BN layers. The gradient explosion could be resolved by adding BN layers back but the Φ -score will be adaptively re-scaled, making it difficult to compare Φ -score between different networks. The same phenomenon has been known as ‘scale-sensitive’ problem in deep learning complexity analysis [4, 35]. To address this open question, we propose to re-scale the Φ -score one more time by the product of BN layers’ variance statistics. This new score is denoted as Zen-Score in order to distinguish from the original Φ -score. The Zen-Score is proven to be scale-insensitive. Finally, we present Zen-NAS algorithm built on Zen-Score and demonstrate its effectiveness in the next section.

4.1. Overflow and BN-rescaling

When computing Φ -score for very deep vanilla networks, numerical overflow incurs almost surely. This is because BN layers are removed from the network and the magnitude of network output grows exponentially along depth. To see this, we construct a set of vanilla networks $P_{w/oBN}$ without BN layers. All networks have the same widths but different depths. Figure 2(a) plots the Φ -scores for $P_{w/oBN}$. After 30 layers, Φ -score overflows. To address the overflow, we add BN layers back and compute the Φ -scores in Figure 2(b). This time the overflow dismisses but the Φ -scores are scaled-down by a large factor. This phenomenon is termed as BN-rescaling.

To demonstrate that BN-rescaling disturbs architecture ranking, we construct another two set of networks, $Q_{w/oBN}$ and Q_{BN} , with and without BN respectively. All networks have two layers and have the same number of input and final output channels. The number of bottleneck channels, that is the width of the hidden layer, varies from 2 to 60. The corresponding Φ -score curves are plotted in Figure 2(d) and (e) respectively. When BN layer is presented, the Φ -score becomes nearly constant for all networks. This will confuse the architecture generator and drive the search to a wrong direction.

4.2. From Φ -Score to Zen-Score

In the above subsection, we showed that BN layer is necessary to prevent numerical overflow in computing Φ -score but comes with the side-effect of re-scaling. In this subsection, we design a new Zen-Score which is able to calibrate

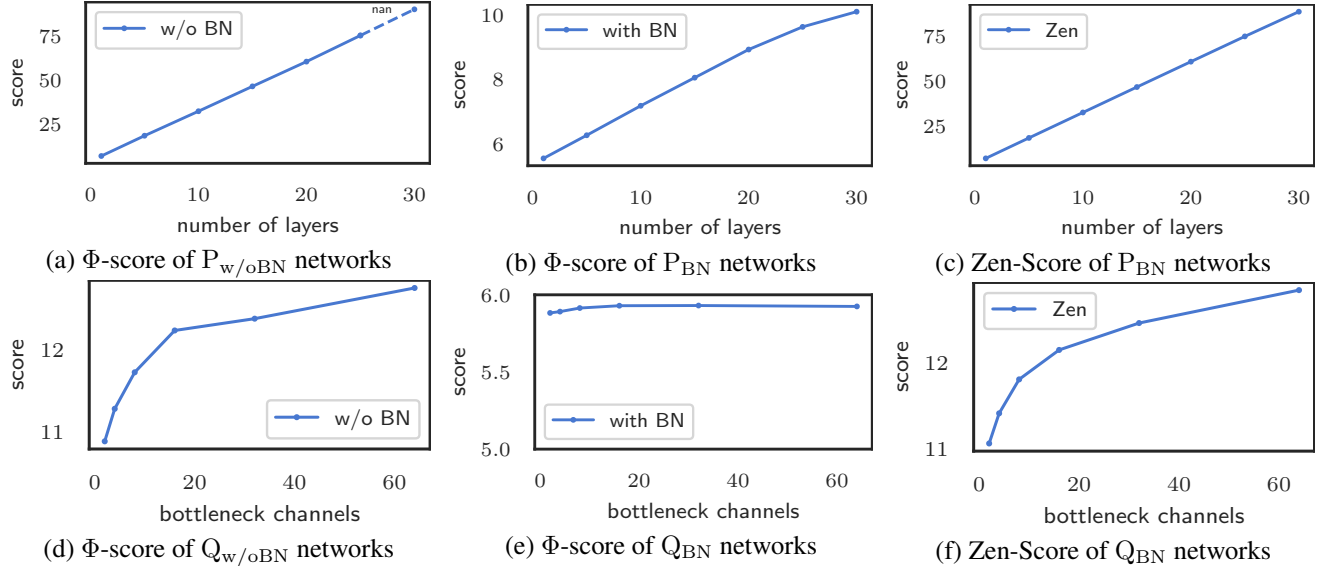


Figure 2. Φ -scores and Zen-Scores of networks, with different depths and bottleneck channels.

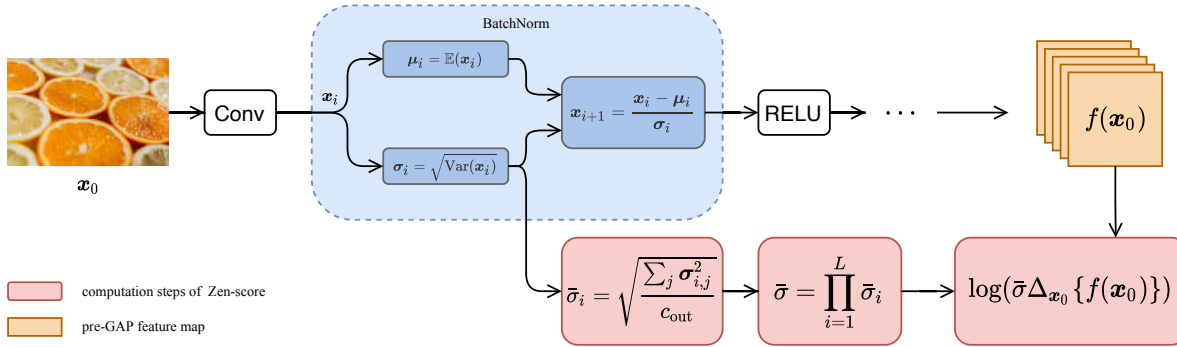


Figure 3. Zen-Score computational graph. \mathbf{x}_0 is one mini-batch of input images. For each BN layer, we extract its mini-batch deviation parameter σ_i . $\Delta_{\mathbf{x}_0} \{f(\mathbf{x}_0)\}$ is the differential of pre-GAP feature map $f(\mathbf{x}_0)$ with respect to \mathbf{x}_0 .

Algorithm 1 Zen-Score

Require: Network $\mathcal{F}(\cdot)$ with pre-GAP feature map $f(\cdot)$;
 $\alpha = 0.01$.

Ensure: Zen-Score $\text{Zen}(\mathcal{F})$.

- 1: Remove all residual links in \mathcal{F} .
- 2: Initialize all neurons in \mathcal{F} by $\mathcal{N}(0, 1)$.
- 3: Sample $\mathbf{x}, \epsilon \sim \mathcal{N}(0, 1)$.
- 4: Compute $\Delta \triangleq \mathbb{E}_{\mathbf{x}, \epsilon} \|f(\mathbf{x}) - f(\mathbf{x} + \alpha\epsilon)\|_F$.
- 5: For the i -th BN layer with m output channels, compute $\bar{\sigma}_i = \sqrt{\sum_j \sigma_{i,j}^2 / m}$ where $\sigma_{i,j}$ is the mini-batch standard deviation statistic of the j -th channel in BN.
- 6: $\text{Zen}(\mathcal{F}) \triangleq \log(\Delta) + \sum_i \log(\bar{\sigma}_i)$.

re-scaling when BN layer is present. The computation of Zen-Score is described in Algorithm 1. Figure 3 visualizes the computational graph of Algorithm 1.

In Algorithm 1, all residual links are removed from the

network as pre-processing. Then we randomly sample input vectors and perturb them with Gaussian noise. The perturbation of the pre-GAP feature map is denoted as Δ in Line 4. This step replaces the gradient of \mathbf{x} with finite differential Δ to avoid backward-propagation. To get Zen-Score, the scaling factor $\bar{\sigma}_i^2$ is averaged from the variance of each channel in BN layer. Finally, the Zen-Score is computed by the log-sum of Δ and $\bar{\sigma}_i$. The following theorem guarantees that the Zen-Score of network with BN layers approximates the Φ -score of the same network without BN layers. The proof is postponed to Supplementary I.

Theorem 1. Let $\bar{f}(\mathbf{x}_0) = \bar{\mathbf{x}}_L$ be an L -layer vanilla network without BN layers. $f(\mathbf{x}_0) = \mathbf{x}_L$ is its sister network with BN layers. For some constants $0 < \delta < 1$, $K_0 \leq \mathcal{O}[\sqrt{\log(1/\delta)}]$, when $BHW \geq \mathcal{O}[(LK_0)^2]$ is large

enough, with probability at least $1 - \delta$, we have

$$(1 - L\epsilon)^2 \leq \frac{(\prod_{t=1}^L \bar{\sigma}_t^2) \mathbb{E}_\theta \{\|\mathbf{x}_L\|^2\}}{\mathbb{E}_\theta \|\bar{\mathbf{x}}_L\|^2} \leq (1 + L\epsilon)^2 \quad (4)$$

where $\epsilon \triangleq \mathcal{O}(2K_0/\sqrt{BHW})$.

Informally speaking, Theorem 1 says that to compute $\|\tilde{f}(\cdot)\|$, we only need to compute $\|f(\cdot)\|$ then re-scale with $\prod_{t=1}^L \bar{\sigma}_t$. The approximation error is bounded by $L\epsilon$. By taking gradient of \mathbf{x} on both $\tilde{f}(\cdot)$ and $f(\cdot)$, we obtain the desired relationship between Zen-Score and Φ -score.

4.3. Zen-NAS For Maximizing Expressivity

Algorithm 2 Zen-NAS

Require: Search space \mathcal{S} , inference budget B , maximal depth L , total number of iterations T , evolutionary population size N , initial structure F_0 .

Ensure: NAS-designed ZenNet F^* .

- 1: Initialize population $\mathcal{P} = \{F_0\}$.
 - 2: **for** $t = 1, 2, \dots, T$ **do**
 - 3: Randomly select $F_t \in \mathcal{P}$.
 - 4: Mutate $\hat{F}_t = \text{MUTATE}(F_t, \mathcal{S})$
 - 5: **if** \hat{F}_t exceeds inference budget or has more than L layers **then**
 - 6: Do nothing.
 - 7: **else**
 - 8: Get Zen-Score $z = \text{Zen}(\hat{F}_t)$.
 - 9: Append \hat{F}_t to \mathcal{P} .
 - 10: **end if**
 - 11: Remove network of the smallest Zen-Score if the size of \mathcal{P} exceeds B .
 - 12: **end for**
 - 13: Return F^* , the network of the highest Zen-Score in \mathcal{P} .
-

Algorithm 3 MUTATE

Require: Structure F_t , search space \mathcal{S} .

Ensure: Randomly mutated structure \hat{F}_t .

- 1: Uniformly select a block h in F_t .
 - 2: Uniformly alternate the block type, kernel size, width and depth of h within some range.
 - 3: Return the mutated structure \hat{F}_t .
-

We design Zen-NAS algorithm to maximize the Zen-Score of the target network. The step-by-step description of Zen-NAS is given in Algorithm 2. The Zen-NAS uses Evolutionary Algorithm (EA) as architecture generator. It is possible to choose other generators such as Reinforced Learning or even greedy selection. The choice of EA is due to its simplicity.

In Algorithm 2, we randomly generate N structures. At each iteration step t , we randomly select a structure in the

population \mathcal{P} and mutate it. The mutation algorithm is presented in Algorithm 3. The width and depth of the selected layer is mutated in a given range. We choose $[0.5, 2.0]$ as the mutation range in this work, that is, within half or double of the current value. The new structure \hat{F}_t is appended to the population if its inference cost does not exceed the budget. The maximal depth of networks is controlled by L , which prevents the algorithm generate over-deep structures. Finally, we maintain the population size by removing networks with the smallest Zen-Scores. After T iterations, the network with the largest Zen-Score is returned as the output of Zen-NAS. We name the found architectures as ZenNets.

5. Experiments

In this section, experiments on CIFAR-10/CIFAR-100 [20] and ImageNet-1k [12] are conducted to validate the superiority of Zen-NAS. We first compare Zen-Score to several zero-shot proxies on CIFAR-10 and CIFAR-100, using the same search space, search policy and training settings. Then we compare Zen-NAS to the state-of-the-art methods on ImageNet. Zen-NAS on CIFAR-10/CIFAR-100 can be found in Supplementary D. Finally, we compare the searching cost of Zen-NAS with SOTA methods in subsection 5.3.

Due to space limitation, the inference speed on NVIDIA T4 and Google Pixel2 is reported in Supplementary C. The Zen-Scores of ResNets and accuracies under fair training settings are reported in Supplementary E. We enclose one big performance table of networks on ImageNet in Supplementary J.

To align with previous works, we consider the following two search spaces:

- **Search Space I** Following [15, 40], this search space consists of residual blocks and bottleneck blocks defined in ResNet.
- **Search Space II** Following [45, 38], this search space consists of MobileNet blocks. The depth-wise expansion ratio is searched in set $\{1, 2, 4, 6\}$.

Please see Supplementary A for datasets description and detail experiment settings.

In each trial, the initial structure is a randomly selected small network which is guaranteed to satisfy the inference budget. The kernel size is searched in set $\{3, 5, 7\}$. Following conventional designs, the number of stages is three for CIFAR-10/CIFAR-100 and five for ImageNet. The evolutionary population size is 256, number of evolutionary iterations $T = 96,000$. The resolution is 32x32 for CIFAR-10/CIFAR-100 and 224x224 for ImageNet.

5.1. Zen-Score v.s. Other Zero-Shot Proxies

Following [1, 7], we compare Zen-Score to five zero-shot proxies: FLOPs, gradient-norm (grad) of network parameters, synflow [51], TE-NAS score (TE-Score) [7] and NASWOT [33]. For each proxy, we replace Zen-Score by

proxy	CIFAR-10	CIFAR-100
Zen-Score	96.2%	80.1%
FLOPs	93.1%	64.7%
grad	92.8%	65.4%
synflow	95.1%	75.9%
TE-Score	96.1%	77.2%
NASWOT	96.0%	77.5%
Random	93.5±0.7%	71.1±3.1%

Table 1. Top-1 accuracies on CIFAR-10/CIFAR-100 for five zero-shot proxies. Budget: model size $N \leq 1$ M. ‘Random’: average accuracy \pm std for random search.

proxy	model	N	time	speed-up
TE-Score	ResNet-18	16	0.34	1/28x
	ResNet-50	16	0.77	1/20x
NASWOT†	ResNet-18	16	0.040	1/3.3x
	ResNet-50	16	0.059	1/1.6x
Zen-Score	ResNet-18	16	0.012	1.0
	ResNet-50	16	0.037	1.0

Table 2. Time cost (in seconds) of computing Zen/TE-Score for ResNet-18/50 at resolution 224x224. The statistical error is within 5%. ‘time’: time for computing Zen/TE-score for N images, measured in seconds, averaged over 100 trials. ‘speed-up’: speed-up rate of TE-Score v.s. Zen-Score.

†: The official implementation outputs Inf score for ResNet-18/50.

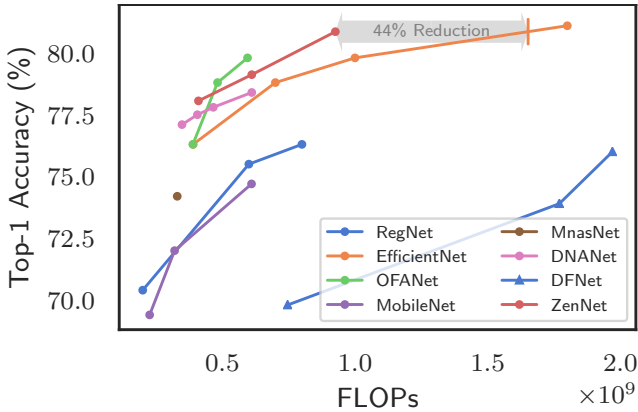


Figure 4. ZenNets optimized for FLOPs.

that proxy in Algorithm 2 and then run Algorithm 2 for $T = 96,000$ iterations to ensure convergence. Since synflow is the smaller the better, we use its negative value in Algo-

NAS	Method	Top-1 (%)	GPU Day
AmoebaNet-A [41]	EA	74.5	3150†
EcoNAS [69]	EA	74.8	8
CARS-I [62]	EA	75.2	0.4
GeNet [57]	EA	72.1	17
DARTS [26]	GD	73.1	4
SNAS [59]	GD	72.7	1.5
PC-DARTS [61]	GD	75.8	3.8
ProxylessNAS [6]	GD	75.1	8.3
GDAS [66]	GD	74	0.8
FBNetV2-L1 [54]	GD	77.2	25
NASNet-A [70]	RL	74	1800
Mnasnet-A [49]	RL	75.2	-
MetaQNN [3]	RL	77.4	96
PNAS [25]	SMBO	74.2	224
SemiNAS [29]	SSL	76.5	4
TE-NAS [7]	ZS	74.1	0.2
OFANet [5]	PS	80.1	51.6
EfficientNet-B7 [50]	Scaling	84.4	3800‡
Zen-NAS	ZS	83.6	0.5

Table 3. NAS searching cost comparison. ‘Top-1’: top-1 accuracy on ImageNet-1k. ‘Method’: ‘EA’ is short for Evolutionary Algorithm; ‘GD’ is short for Gradient Descent; ‘RL’ is short for reinforcement Learning; ‘ZS’ is short for Zero-shot; ‘SMBO’, ‘SSL’, ‘PS’ and ‘Scaling’ are special searching methods/frameworks. †: Running on TPU; ‡: The cost is estimated by [54];

rithm 2. Following convention, we search for best network on CIFAR-10/CIFAR-100 within model size $N \leq 1$ M. The convergence curves are plotted in Supplementary C. In these figures, all six scores improves monotonically along iterations.

After the above NAS step, we train the network of the best score for each proxy under the same training setting. To provide a random baseline, we randomly generate networks. The width of the layer varies in range [4, 512], and the depth of each stage varies in range [1, 10]. If the network size is larger than 1 M, we shrink its width by factor 0.75 each time until it satisfies the budget. 32 random networks are generated and trained in total.

The top-1 accuracy is reported in Table 1. Zen-Score significantly outperforms the other five proxies on both CIFAR-10 and CIFAR-100. TE-Score and NASWOT are the runner-up proxies with similar performance, followed by synflow. It is not surprise to see that naive proxies, such as FLOPs and gradient-norm, perform poorly, even worse than random search.

To compare the computational efficiency of Zen-Score and TE-score, we compute two scores for ResNet-18 and ResNet-50 at 224x224 resolution. The expected time cost is averaged over 100 trials. We find that averaging Zen/TE-

Score over $N = 16$ random images is sufficient to reduce the statistical error below 5%. The results are reported in Table 2. The computation of Zen-Score is 20 \sim 28 times faster than TE-Score.

We tried our best to benchmark NASWOT for ResNet-18/50 using the official code. However, the official code always outputs Inf for ResNet-18/50 at resolution 224. Despite of the Inf issue, Zen-Score is 3.3x times faster than NASWOT on ResNet-18 and 1.6x times faster on ResNet-50.

5.2. Zen-NAS on ImageNet

We use Zen-NAS to search efficient network (ZenNet) on ImageNet. We consider the following popular networks as baselines: (a) manually-designed networks, including ResNet [15], DenseNet [17], ResNeSt [65], MobileNet-V2 [45] (b) NAS-designed networks for fast inference on GPU, including OFANet-9ms/11ms [5], DFNet [22], RegNet [40]; (c) NAS-designed networks optimized for FLOPs, including OFANet-389M/482M/595M [5], DNANet [21], EfficientNet [50], Mnasnet [49].

Among these networks, EfficientNet is a popular baseline in NAS-related works. EfficientNet-B0/B1 are suitable for mobile device for their small FLOPs and model size. EfficientNet-B3~B7 are large models that are best to be deployed on a high-end GPU. Although EfficientNet is optimized for FLOPs, its inference speed on GPU is within top-tier ones. Many previous works compare to EfficientNet by inference speed on GPU [65, 5, 40].

Searching Low Latency Networks Following previous works [5, 22, 40], we use Zen-NAS to optimize network inference speed on NVIDIA V100 GPU. We use Search Space I in this experiment. The inference speed is tested at batch size 64, half precision (float16). We search for networks of inference latency within 0.1/0.2/0.3/0.5/0.8/1.2 milliseconds (ms) per image. For testing inference latency, we set batch-size=64 and do mini-batch inference 30 times. The averaged inference latency is recorded. The top-1 accuracy on ImageNet v.s. inference latency is plotted in Figure 1. Clearly, ZenNets outperform baseline models in both accuracy and inference speed by a large margin. The largest model ZenNet-1.2ms achieves 83.6% top-1 accuracy which is between EfficientNet-B5 and B6. It is about 4.9x faster than EfficientNet at the same accuracy level.

Searching Lightweight Networks Following previous works [5, 50], we use Zen-NAS to search lightweight networks with small FLOPs. We use Search Space II in this experiment. We search for networks of computational cost within 400/600/900 M FLOPs. Similar to OFANet and EfficientNet, we add SE-blocks after convolutional layers. The top-1 accuracy v.s. FLOPs is plotted in Figure 4.

Again, ZenNets outperform most models by a large margin. ZenNet-900M-SE achieves 80.8% top-1 accuracy which is comparable to EfficientNet-B3 with 43% fewer FLOPs. The runner-up is OFANet whose efficiency is similar to ZenNet.

5.3. Searching Cost of Zen-NAS v.s. SOTA

The major time cost of Zen-NAS is the computation of Zen-Score. The network latency is predicted by an in-house latency predictor whose time cost is nearly zero. According to Table 2, the computation of Zen-Score for ResNet-50 only takes 0.15 second. This means that scoring 96,000 networks similar to ResNet-50 only takes 4 GPU hours, or 0.17 GPU day.

We compare Zen-NAS searching cost to SOTA NAS methods in Table 3. Since every NAS method uses different settings, it is difficult to make a fair comparison that everyone agrees with. Nevertheless, we only concern about the best model reported in each paper and the corresponding searching cost. This gives us a rough impression of the efficiency of these methods and their practical ability of designing high-performance models.

From Table 3, for conventional NAS methods, it takes hundreds to thousands GPU days to find a good structure of accuracy better than 78.0%. Many one-shot methods are very fast. For most one-shot methods, the best accuracy is below 80%. In comparison, Zen-NAS achieves 83.6% top-1 accuracy within 0.5 GPU day. Among methods achieving above 80.0% top-1 accuracy in Table 3, the searching speed of Zen-NAS is nearly 100 times faster than OFANet and 7800 times faster than EfficientNet. TE-NAS uses less GPU day than Zen-NAS in Table 3. This does not conflict with Table 2 because the total number of networks evaluated by the two methods are different.

6. Conclusion

We proposed Zen-NAS, a zero-shot neural architecture search framework for designing high performance deep image recognition networks. Without optimizing network parameters, Zen-NAS ranks networks via network expressivity which can be numerically measured by Zen-Score. The searching speed of Zen-NAS is dramatically faster than previous SOTA methods. The ZenNets automatically designed by Zen-NAS are significantly more efficient in terms of inference latency, FLOPs and model size, in multiple recognition tasks. We wish the elegance of Zen-NAS will inspire more theoretical researches towards a deeper understanding of efficient network design.

References

- [1] Mohamed S. Abdelfattah, Abhinav Mehrotra, Łukasz Dudziak, and Nicholas D. Lane. Zero-Cost Proxies for Lightweight NAS. In *ICLR*, 2021.

- [2] Gustavo Aguilar, Yuan Ling, Yu Zhang, Benjamin Yao, Xing Fan, and Chenlei Guo. Knowledge Distillation from Internal Representations. In *AAAI*, 2020.
- [3] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing Neural Network Architectures using Reinforcement Learning. In *ICLR*, 2017.
- [4] Peter L. Bartlett, Dylan J. Foster, and Matus J. Telgarsky. Spectrally-normalized margin bounds for neural networks. In *NIPS*, 2017.
- [5] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-All: Train One Network and Specialize it for Efficient Deployment on Diverse Hardware Platforms. In *ICLR*, 2020.
- [6] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct Neural Architecture Search on Target Task and Hardware. In *ICLR*, 2019.
- [7] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *ICLR*, 2021.
- [8] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive DARTS: Bridging the Optimization Gap for NAS in the Wild. In *ICCV*, 2019.
- [9] Nadav Cohen and Amnon Shashua. Inductive Bias of Deep Convolutional Networks through Pooling Geometry. In *ICLR*, 2017.
- [10] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. AutoAugment: Learning Augmentation Policies from Data. In *CVPR*, 2019.
- [11] Amit Daniely, Roy Frostig, and Yoram Singer. Toward Deeper Understanding of Neural Networks: The Power of Initialization and a Dual View on Expressivity. In *NIPS*, 2016.
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. ImageNet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [13] 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 *ECCV*, 2020.
- [14] Boris Hanin and David Rolnick. Complexity of Linear Regions in Deep Networks. In *ICML*, 2019.
- [15] K. He, X. Zhang, S. Ren, and J. Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [16] Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Enhua Wu. Squeeze-and-Excitation Networks. In *CVPR*, 2018.
- [17] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely Connected Convolutional Networks. In *CVPR*, 2017.
- [18] Sham M. Kakade, Karthik Sridharan, and Ambuj Tewari. On the Complexity of Linear Prediction: Risk Bounds, Margin Bounds, and Regularization. In *NIPS*, 2008.
- [19] Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems*, volume 2033. Springer Science & Business Media, 2011.
- [20] Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009.
- [21] Changlin Li, Jiefeng Peng, Liuchun Yuan, Guangrun Wang, Xiaodan Liang, Liang Lin, and Xiaojun Chang. Blockwisely Supervised Neural Architecture Search with Knowledge Distillation. In *CVPR*, 2020.
- [22] Xin Li, Yiming Zhou, Zheng Pan, and Jiashi Feng. Partial Order Pruning: For Best Speed/Accuracy Trade-off in Neural Architecture Search. In *CVPR*, 2019.
- [23] Shiyu Liang and R. Srikant. Why Deep Neural Networks for Function Approximation? In *ICLR*, 2016.
- [24] Chenxi Liu, Piotr Dollár, Kaiming He, Ross Girshick, Alan Yuille, and Saining Xie. Are Labels Necessary for Neural Architecture Search? In *ECCV*, 2020.
- [25] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Li Fei-Fei, Alan L. Yuille, Jonathan Huang, and Kevin Murphy. Progressive Neural Architecture Search. In *ECCV*, 2018.
- [26] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable Architecture Search. In *ICLR*, 2019.
- [27] Ilya Loshchilov and Frank Hutter. SGDR: Stochastic Gradient Descent with Warm Restarts. In *ICLR*, 2017.
- [28] Zhou Lu, Hongming Pu, Feicheng Wang, Zhiqiang Hu, and Liwei Wang. The Expressive Power of Neural Networks: A View from the Width. In *NIPS*, 2017.
- [29] Renqian Luo, Xu Tan, Rui Wang, Tao Qin, Enhong Chen, and Tie-Yan Liu. Semi-Supervised Neural Architecture Search. In *NIPS*, 2020.
- [30] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural Architecture Optimization. In *NIPS*, 2018.
- [31] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the Expressive Power of Deep Neural Networks. In *ICML*, 2017.
- [32] Joseph Mellor, Jack Turner, Amos Storkey, and Elliot J. Crowley. Neural Architecture Search without Training. *arXiv:2006.04647 [cs, stat]*, 2021.
- [33] Joseph Mellor, Jack Turner, Amos Storkey, and Elliot J. Crowley. Neural Architecture Search without Training. In *ICML*, 2021.
- [34] Guido Montúfar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. On the Number of Linear Regions of Deep Neural Networks. In *NIPS*, 2014.
- [35] Behnam Neyshabur, Zhiyuan Li, Srinadh Bhojanapalli, Yann LeCun, and Nathan Srebro. The role of over-parametrization in generalization of neural networks. In *ICML*, 2018.
- [36] Thao Nguyen, Maithra Raghu, and Simon Kornblith. Do Wide and Deep Networks Learn the Same Things? Uncovering How Neural Network Representations Vary with Width and Depth. In *ICLR*, 2021.
- [37] Hieu Pham, Zihang Dai, Qizhe Xie, Minh-Thang Luong, and Quoc V. Le. Meta Pseudo Labels. *arXiv:2003.10580 [cs, stat]*, 2021.
- [38] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient Neural Architecture Search via Parameters Sharing. In *ICML*, 2018.
- [39] Ben Poole, Subhaneil Lahiri, Maithra Raghu, Jascha Sohl-Dickstein, and Surya Ganguli. Exponential expressivity in deep neural networks through transient chaos. In *NIPS*, 2016.

- [40] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing Network Design Spaces. In *CVPR*, 2020.
- [41] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V. Le. Regularized Evolution for Image Classifier Architecture Search. In *AAAI*, 2019.
- [42] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Quoc V. Le, and Alex Kurakin. Large-Scale Evolution of Image Classifiers. In *ICML*, 2017.
- [43] Pengzhen Ren, Yun Xiao, Xiaojun Chang, Po-Yao Huang, Zhihui Li, Xiaojiang Chen, and Xin Wang. A Comprehensive Survey of Neural Architecture Search: Challenges and Solutions, 2020.
- [44] David Rolnick and Max Tegmark. The power of deeper networks for expressing natural functions. In *ICLR*, 2018.
- [45] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. MobileNetV2: Inverted Residuals and Linear Bottlenecks. In *CVPR*, 2018.
- [46] Christian Sciuto, Kaicheng Yu, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the Search Phase of Neural Architecture Search. *arXiv:1902.08142 [cs, stat]*, 2019.
- [47] Thiago Serra, Christian Tjandraatmadja, and Srikumar Ramalingam. Bounding and Counting Linear Regions of Deep Neural Networks. In *ICML*, 2018.
- [48] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jonathon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *CVPR*, 2016.
- [49] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V. Le. MnasNet: Platform-Aware Neural Architecture Search for Mobile. In *CVPR*, 2019.
- [50] Mingxing Tan and Quoc Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *ICML*, 2019.
- [51] Hidenori Tanaka, Daniel Kunin, Daniel L. K. Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In *NIPS*, 2020.
- [52] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv:2012.12877 [cs]*, 2021.
- [53] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 2018.
- [54] Alvin Wan, Xiaoliang Dai, Peizhao Zhang, Zijian He, Yuan-dong Tian, Saining Xie, Bichen Wu, Matthew Yu, Tao Xu, Kan Chen, Peter Vajda, and Joseph E. Gonzalez. FBNetV2: Differentiable Neural Architecture Search for Spatial and Channel Dimensions. In *CVPR*, 2020.
- [55] Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking Winning Tickets Before Training by Preserving Gradient Flow. In *ICLR*, 2019.
- [56] Wei Wen, Hanxiao Liu, Hai Li, Yiran Chen, Gabriel Bender, and Pieter-Jan Kindermans. Neural Predictor for Neural Architecture Search. In *ECCV*, 2020.
- [57] Lingxi Xie and Alan Yuille. Genetic CNN. In *ICCV*, 2017.
- [58] Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring Randomly Wired Neural Networks for Image Recognition. *arXiv:1904.01569 [cs]*, 2019.
- [59] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. SNAS: Stochastic neural architecture search. In *ICLR*, 2018.
- [60] H. Xiong, L. Huang, M. Yu, L. Liu, F. Zhu, and L. Shao. On the Number of Linear Regions of Convolutional Neural Networks. In *ICML*, 2020.
- [61] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. PC-DARTS: Partial Channel Connections for Memory-Efficient Architecture Search. In *ICLR*, 2019.
- [62] Zhaohui Yang, Yunhe Wang, Xinghao Chen, Boxin Shi, Chao Xu, Chunjing Xu, Qi Tian, and Chang Xu. CARS: Continuous Evolution for Efficient Neural Architecture Search. In *CVPR*, 2020.
- [63] Chris Ying, Aaron Klein, Esteban Real, Eric Christiansen, Kevin Murphy, and Frank Hutter. NAS-Bench-101: Towards Reproducible Neural Architecture Search. In *ICML*, 2019.
- [64] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. Mixup: Beyond Empirical Risk Minimization. In *ICLR*, 2018.
- [65] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Mueller, R. Manmatha, Mu Li, and Alexander Smola. ResNeSt: Split-Attention Networks, 2020.
- [66] Miao Zhang, Huiqi Li, Shirui Pan, Xiaojun Chang, and Steven Su. Overcoming Multi-Model Forgetting in One-Shot NAS With Diversity Maximization. In *CVPR*, 2020.
- [67] Xiao Zhang and Dongrui Wu. Empirical Studies on the Properties of Linear Regions in Deep Neural Networks. In *ICLR*, 2019.
- [68] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random Erasing Data Augmentation. In *AAAI*, 2020.
- [69] Dongzhan Zhou, Xinchu Zhou, Wenwei Zhang, Chen Change Loy, Shuai Yi, Xuesen Zhang, and Wanli Ouyang. EcoNAS: Finding Proxies for Economical Neural Architecture Search. In *CVPR*, 2020.
- [70] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning Transferable Architectures for Scalable Image Recognition. In *CVPR*, 2018.