

Stitchable Neural Networks

Zizheng Pan Jianfei Cai Bohan Zhuang[†]

ZIP Lab, Monash University

<https://snet.github.io>




Figure 1. Compared with previous scalable deep learning frameworks. (a) Network compression shrinks a large network into a small one by techniques such as pruning, quantization and knowledge distillation, *etc.*, which is a one-to-one mapping. (b) One-shot neural architecture search first trains a supernet that supports diverse architectural settings and then specializes a subnet given the target resource constraint during deployment, which is a case of one-to-many. (c) Our proposed Stitchable Neural Network directly stitches the off-the-rack family of pretrained models and quickly obtains new networks for efficient model design and deployment in a novel many-to-many paradigm.

Abstract

The public model zoo containing enormous powerful pretrained model families (e.g., ResNet/DeiT) has reached an unprecedented scope than ever, which significantly contributes to the success of deep learning. As each model family consists of pretrained models with diverse scales (e.g., DeiT-Ti/S/B), it naturally arises a fundamental question of how to efficiently assemble these readily available models in a family for dynamic accuracy-efficiency trade-offs at runtime. To this end, we present *Stitchable Neural Networks (SN-Net)*, a novel scalable and efficient framework for model deployment. It cheaply produces numerous networks with different complexity and performance trade-offs given a family of pretrained neural networks, which we call *anchors*. Specifically, SN-Net splits the anchors across the blocks/layers and then stitches them together with simple stitching layers to map the activations from one anchor to another. With only a few epochs of training, SN-Net effectively interpolates between the performance of anchors with varying scales. At runtime, SN-Net can instantly adapt to dynamic resource constraints by switching the stitching positions. Extensive experiments on ImageNet classification demonstrate that SN-Net can obtain on-par or even better performance than many individually trained networks

while supporting diverse deployment scenarios. For example, by stitching Swin Transformers, we challenge hundreds of models in Timm model zoo with a single network. We believe this new elastic model framework can serve as a strong baseline for further research in wider communities.

1. Introduction

The vast computational resources available and large amount of data have driven researchers to build tens of thousands of powerful deep neural networks with strong performance, which have largely underpinned the most recent breakthroughs in machine learning and much broader artificial intelligence. Up to now, there are $\sim 81k$ models on HuggingFace [53] and ~ 800 models on Timm [52] that are ready to be downloaded and executed without the overhead of reproducing. Despite the large model zoo, a model family (e.g., DeiT-Ti/S/B [48]) that contains pretrained models with functionally similar architectures but different scales only covers a coarse-grained level of model complexity/performance, where each model only targets a specific resource budget (e.g., FLOPs). Moreover, the model family is not flexible to adapt to dynamic resource constraints since each individual model is not re-configurable due to the fixed computational graph. In reality, we usually need to deploy models to diverse platforms with different resource

[†]Corresponding author. E-mail: bohan.zhuang@gmail.com

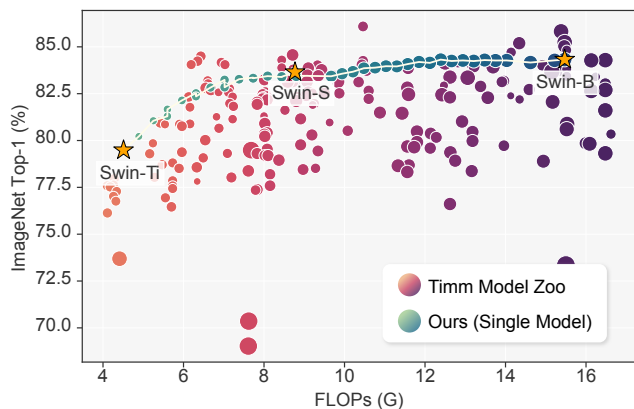


Figure 2. **One** Stitchable Neural Network vs. **200** models in Timm model zoo [52]. It shows an example of SN-Net by stitching ImageNet-22K pretrained Swin-Ti/S/B. Compared to each individual network, SN-Net is able to instantly switch network topology at runtime and covers a wide range of computing resource budgets. Larger and darker dots indicate a larger model with more parameters and higher complexity.

constraints (e.g., energy, latency, on-chip memory). For instance, a mobile app in Google Play has to support tens of thousands of unique Android devices, from a high-end Samsung Galaxy S22 to a low-end Nokia X5. Therefore, given a family of pretrained models in the model zoo, a fundamental research question naturally arises: *how to effectively utilise these off-the-shelf pretrained models to handle diverse deployment scenarios for Green AI [45]?*

To answer this question, a naive solution is to train individual models with different accuracy-efficiency trade-offs from scratch. However, such method has a linearly increased training and time cost with the number of possible cases. Therefore, one may consider the existing scalable deep learning frameworks, such as model compression and neural architecture search (NAS), to obtain models at different scales for diverse deployment requirements. Specifically, network compression approaches such as pruning [20, 22, 25], quantization [32, 42, 62] and knowledge distillation [7, 43, 47] aim to obtain a small model from a large and well-trained network, which however only target one specific resource budget (see Figure 1 (a)), thus not flexible to meet the requirements of real-world deployment scenarios. On the other hand, one-shot NAS [28, 37], a typical NAS framework that decouples training and specialization stages, seeks to train an over-parameterized supernet that supports many sub-networks for run-time dynamics (see Figure 1 (b)), but training the supernet is extremely time-consuming and computationally expensive (e.g., 1,200 GPU hours on 32 V100 GPUs in OFA [4]). To summarize, the existing scalable deep learning frameworks are still limited within a single model design space, which cannot inherit the rich knowledge from pretrained model families in a model zoo for better flexibility and accuracy. Besides, they also

require complicated training strategies to guarantee a good model performance.

In this work, we present Stitchable Neural Network (SN-Net), a novel scalable deep learning framework for efficient model design and deployment which quickly stitches an off-the-shelf pretrained model family with much less training effort to cover a fine-grained level of model complexity/performance for a wide range of deployment scenarios (see Figure 1 (c)). Specifically, SN-Net is motivated by the previous observations [2, 10, 21] that the typical minima reached by SGD can be stitched to each other with low loss penalty, which implies architectures of the same model family pretrained on the same task can be stitched. Based on this insight, SN-Net directly selects the well-performed pretrained models in a model family as “anchors”, and then inserts a few simple stitching layers at different positions to transform the activations from one anchor to its nearest anchor in terms of complexity. In this way, SN-Net naturally interpolates a path between neighbouring anchors of different accuracy-efficiency trade-offs, and thus can handle dynamic resource constraints *with a single neural network at runtime*. An example is shown in Figure 2, where a single Swin-based SN-Net is able to do what hundreds of models can do with only 50 epochs training on ImageNet-1K.

We systematically study the design principles for SN-Net, including the choice of anchors, the design of stitching layers, the stitching direction and strategy, along with a sufficiently simple but effective training strategy. With comprehensive experiments, we show that SN-Net demonstrates promising advantages: 1) Compared to the existing prevalent scalable deep learning frameworks (Figure 1), SN-Net is a new universal paradigm which breaks the limit of a single pretrained model or supernet design by extending the design space into a large number of model families in the model zoo, forming a “many-to-many” pipeline. 2) Different from NAS training that requires complex optimization techniques [4, 61], training SN-Net is as easy as training individual models while getting rid of the huge computational cost of training from scratch. 3) The final performance of stitches is almost predictable due to the interpolation-like performance curve between anchors, which implies that we can selectively train a number of stitches prior to training based on different deployment scenarios.

In a nutshell, we summarize our contributions as follows:

- We introduce Stitchable Neural Networks, a new universal framework for elastic deep learning by directly utilising the pretrained model families in model zoo via model stitching.
- We provide practical principles to design and train SN-Net, laying down the foundations for future research.
- Extensive experiments demonstrate that compared to training individual networks from scratch, e.g., a single

DeiT-based [48] SN-Net can achieve flexible accuracy-efficiency trade-offs at runtime while reducing 22× training cost and local disk storage.

2. Related Work

Model stitching. Model stitching was initially proposed by Lenc *et al.* [21] to study the equivalence of representations. Specifically, they showed that the early portion of a trained network can be connected with the last portion of another trained network by a 1×1 convolution stitching layer without significant performance drop. Most recently, Yamini *et al.* [2] revealed that neural networks, even with different architectures or trained with different strategies, can also be stitched together with small effect on performance. As a concurrent work to [2], Adrián *et al.* [10] studied using model stitching as an experimental tool to match neural network representations. They demonstrated that common similarity indices (*e.g.*, CKA [19], CCA [12], SVCCA [40]) are not correlated to the performance of the stitched model. Unlike these previous works which view model stitching as a tool to measure neural network representations, this paper unleashes the power of model stitching as a general approach for utilising the pretrained model families in the large-scale model zoo to obtain a single scalable neural network at a low cost that can instantly adapt to diverse deployment scenarios. More recently, Yang *et al.* proposed DeRy [58] to dissect and reassemble arbitrary pretrained models into a new network for a certain resource constraint (*e.g.*, FLOPs) one at a time. Unlike DeRy, the proposed SN-Net supports numerous sub-networks by stitching the off-the-shelf model families, being capable of handling diverse resource budgets at deployment time.

Neural architecture search. Neural architecture search (NAS) [64] aims to automatically search for the well-performed network architecture in a pre-defined search space under different resource constraints. In the early attempts [64, 65], NAS consumes prohibitive computational cost (*e.g.*, 500 GPUs across 4 days in [65]) due the requirement of training individual sub-networks until convergence for accurate performance estimation. To address this problem, one-shot NAS [5, 28, 30, 37, 56] has been proposed to improve NAS efficiency by weight sharing, where multiple subnets share the same weights with the supernet. However, training a supernet still requires intensive computing resources. Most recently, zero-shot NAS [1, 8, 18, 31] has been proposed to identify good architectures prior to training. However, obtaining the final model still requires training from scratch. Compared to NAS, our method builds upon the off-the-shelf family of pretrained models in model zoo, which exploits the large model design space and is able to assemble the existing rich knowledge from heterogeneous models for flexible and diverse model deployments.

Vision Transformers. Vision Transformers [11] are emerging deep neural networks which have challenged the de-facto standard of convolutional neural networks on vision tasks. The majority of the existing efforts focus on improving the performance of ViT as a single general vision backbone [29, 34, 49, 50, 57] or adopting ViT as a strong module for modeling global relationships to address downstream tasks [23, 46, 54]. Another line of works focus on improving ViT efficiency by token pruning [35, 41], quantization [24, 27] and dynamic inference [41, 51], *etc.* Most recently, large-scale self-supervised pretraining has helped ViTs achieve promising results on ImageNet, including contrastive learning [6, 9] and masked image modeling [3, 13, 17, 63]. However, these models are designed to be over-parameterized and have a fixed computational cost, which is inflexible at the inference stage and cannot adapt to diverse and dynamic deployment environment. Instead of proposing or pretraining a new ViT architecture, we utilize different pretrained ViTs or even CNNs [15, 55] to show that the proposed SN-Net is a general framework to assemble the existing model families.

3. Method

In this section, we first introduce the preliminary of model stitching at Section 3.1. Next, we describe the details of our proposed stitchable neural networks at Section 3.2.

3.1. Preliminaries of Model Stitching

Let θ be the model parameters of a pretrained neural network and f_i represent the function of the i -th layer. A typical feed-forward neural network with L layers can be defined as a composition of functions: $f_\theta = f_L \circ \dots \circ f_1$, where \circ indicates the composition, and $f_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ maps the inputs in an input space \mathcal{X} to the output space \mathcal{Y} . Let $\mathbf{X} \in \mathcal{X}$ be an input to the network. The basic idea of model stitching involves splitting a neural network into two portions of functions at a layer index l . The first portion of layers compose the front part that maps the input \mathbf{X} into the activation space of the l -th layer $\mathcal{A}_{\theta,l}$, which can be formulated as

$$H_{\theta,l}(\mathbf{X}) = f_l \circ \dots \circ f_1(\mathbf{X}) = \mathbf{X}_l, \quad (1)$$

where $\mathbf{X}_l \in \mathcal{A}_{\theta,l}$ denotes the output feature map at the l -th layer. Next, the last portion of layers maps \mathbf{X}_l into the final output

$$T_{\theta,l}(\mathbf{X}_l) = f_L \circ \dots \circ f_{l+1}(\mathbf{X}_l). \quad (2)$$

In this case, the original neural network function f_θ can be defined as a composition of the above functions $f_\theta = T_{\theta,l} \circ H_{\theta,l}$ for all layer indexes $l = 1, \dots, L - 1$.

Now suppose we have another pretrained neural network f_ϕ . Let $\mathcal{S} : \mathcal{A}_{\theta,l} \rightarrow \mathcal{A}_{\phi,m}$ be a stitching layer which implements a transformation between the activation space of the l -th layer of f_θ to the activation space of the m -th layer

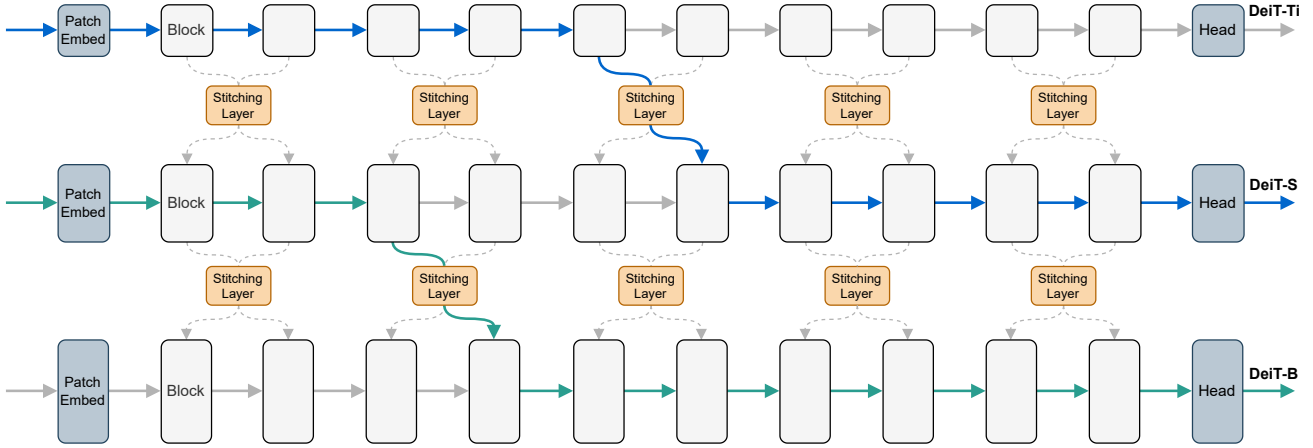


Figure 3. Illustration of the proposed **Stitchable Neural Network**, where three pretrained variants of DeiT’s are connected with simple stitching layers (1×1 convolutions). We share the same stitching layer among neighboring blocks (e.g., 2 blocks with a stride of 2 in this example) between two models. Apart from the basic anchor models, we obtain many sub-networks (stitches) by stitching the nearest pairs of anchors in complexity, e.g., DeiT-Ti and DeiT-S (the blue line), DeiT-S and DeiT-B (the green line). Best viewed in color.

of f_ϕ . The basic idea of model stitching to obtain a new network defined by \mathcal{S} can be expressed as

$$F_S(\mathbf{X}) = T_{\phi,m} \circ \mathcal{S} \circ H_{\theta,l}(\mathbf{X}). \quad (3)$$

By controlling the stitched layer indexes l and m , model stitching can produce a sequence of stitched networks. It has been observed by [21] that models of the same architecture but with different initializations (i.e., random seeds) can be stitched together with low loss penalty. Further experiments by [2, 10] have demonstrated that different architectures (e.g., ViTs and CNNs) may also be stitched without significant performance drop, regardless they are trained in different ways such as self-supervised learning or supervised learning.

3.2. Stitchable Neural Networks

Based on the insight of model stitching, we propose Stitchable Neural Networks (SN-Net), a new “many-to-many” elastic model paradigm. SN-Net is motivated by an increasing number of pretrained models in the publicly available model zoo [52], where most of the individually trained models are not directly adjustable to dynamic resource constraints. To this end, SN-Net inserts a few stitching layers to smoothly connect a family of pretrained models to form diverse stitched networks permitting run-time network selection. The framework of SN-Net is illustrated in Figure 3 by taking plain ViTs [11] as an example. For brevity, we will refer to the models that to be stitched as “anchors” and the derived models by stitching anchors as “stitches”. In the following, we describe the concrete approach in detail, including what, how and where to stitch, the stitching strategy and space, as well as an effective and efficient training strategy for SN-Net.

What to stitch: the choice of anchors. In general, the large-scale model zoo determines the powerful representation capability of SN-Net as it is a universal framework for assembling the prevalent families of architectures. As shown in Section 4, SN-Net works for stitching representative ViTs and CNNs. However, intuitively, anchors that are pretrained on different tasks can learn very different representations (e.g., ImageNet [44] and COCO [26]) due to the large distribution gap of different domains [33], thus making it difficult for stitching layers to learn to transform activations among anchors. Therefore, the selected anchors should be consistent in terms of the pretrained domain.

How to stitch: the stitching layer and its initialization. Conceptually, the stitching layer should be as simple as possible since its aim is not to improve the model performance, but to transform the feature maps from one activation space to another [2]. To this end, the stitching layers in SN-Net are simply 1×1 convolutional layers. By default in PyTorch [36], these layers are initialized based on Kaiming initialization [14].

However, different from training a network from scratch as in most works [3, 29, 49, 50], SN-Net is built upon pretrained models. In this case, the anchors have already learned good representations, which allows to directly obtain an accurate transformation matrix by solving the following least squares problem

$$\|\mathbf{A}\mathbf{M}_o - \mathbf{B}\|_F = \min \|\mathbf{A}\mathbf{M} - \mathbf{B}\|_F, \quad (4)$$

where $\mathbf{A} \in \mathbb{R}^{N \times D_1}$ and $\mathbf{B} \in \mathbb{R}^{N \times D_2}$ are two feature maps of the same spatial size but with different number of hidden dimensions. N denotes the length of the input sequence and D_1, D_2 refer to the number of hidden dimensions. $\mathbf{M} \in$




Figure 4. Stitching direction: Fast-to-Slow vs. Slow-to-Fast.

$\mathbb{R}^{D_1 \times D_2}$ is the targeted transformation matrix.

One can observe that Eq. (4) indicates a closed form expression based on singular value decomposition, in which case the optimal solution can be achieved through an orthogonal projection in the space of matrices,

$$\mathbf{M}_o = \mathbf{A}^\dagger \mathbf{B}, \quad (5)$$

where \mathbf{A}^\dagger denotes the Moore-Penrose pseudoinverse of \mathbf{A} . To obtain \mathbf{M}_o requires only a few seconds on one CPU with hundreds of samples. However, we will show in Section 4.2 that directly using with least-squares solution achieves unstable performance for stitches, but it actually provides a good initialization for learning stitching layers with SGD. Therefore, the least-squares solution serves as the default initialization approach for the stitching layers in SN-Net.

Where to stitch: the stitching directions. Given anchors with different scales and complexities, there are two options to stitch them together: **Fast-to-Slow** and **Slow-to-Fast**. Taking two anchors as an example (Figure 4), Fast-to-Slow takes the first portion of layers (*i.e.*, Eq. (1)) from a smaller and faster model, and the last portion of layers (*i.e.*, Eq. (1)) from a larger and slower model, where Slow-to-Fast goes in a reverse direction. However, as Fast-to-Slow is more aligned with the existing model design principle (*i.e.*, increasing the network width as it goes deeper), we will show in Section 4.2 that it achieves more stable and better performance than Slow-to-Fast. In this case, we take Fast-to-Slow as the default stitching direction in SN-Net. Besides, as different anchors may reach very different minima, we propose a **nearest stitching** strategy by limiting the stitching between two anchors of the nearest model complexity/performance. Thus, each stitch in SN-Net assembles a pair of neighbouring anchors. We will show in Section 4.2 that stitching across anchors without the nearest stitching constraint achieves inferior performance.

Way to stitch: stitching as sliding windows. Our stitching strategy is inspired by the main observation: neighboring layers dealing with the same scale feature maps share similar representations [19]. To this end, we propose to stitch anchors as sliding windows, where the same window




Figure 5. Stitching as sliding windows, where paired stitching is proposed for stitching models with equal depth and unpaired stitching is utilised for models with unequal depth.

shares a common stitching layer, as shown in Figure 5. Let L_1 and L_2 be depth of two anchors. Then intuitively, there are two cases when stitching layers/blocks between the two anchors: **paired stitching** ($L_1 = L_2$) and **unpaired stitching** ($L_1 \neq L_2$). In the case of $L_1 = L_2$, the sliding windows can be controlled as sliding windows with a window size k and a stride s . Figure 5 left shows an example with $k = 2, s = 1$. However, in most cases we have unequal depth as different model architectures have different scales. Even though, matching L_1 layers to L_2 layers can be easily done by nearest interpolation where each layer from the shallower anchor can be stitched with more than one layers of the deeper anchor, as shown in Figure 5 right.

Stitching space. In SN-Net, we first split the anchors along the internal layers/blocks at each stage then apply our stitching strategy within each stage. As different anchors have different architectural configurations, the size of the stitching space can be variant based on the depth of the selected anchors and the stitching settings (*i.e.*, the kernel size k and stride s). For example, with $k = 2$ and $s = 1$, DeiT-based SN-Net can have 71 stitches under the constraint of our nearest stitching principle, or 732 stitches without this constraint, as shown in Figure 9 (b). We provide detailed illustrations for this figure in the supplementary material. More stitches can be obtained by choosing anchors with larger scales or configuring the sliding windows by using a larger window size or smaller stride. Overall, compared to one-shot NAS which can support more than 10^{20} sub-networks, SN-Net has a relatively smaller space (up to hundreds or thousands). However, we point out that even though NAS has a much larger architecture space, during deployment, it only focuses on the sub-networks on the Pareto frontier of performance and resource consumption [60]. Thus the vast majority of sub-networks are ignored. In contrast, we will show in Section 4.2 that the stitches in SN-Net smoothly distribute among the anchors, which indicates that the analogous performance curve can almost be estimated without too much searching cost, permitting fast deployment.

Training strategy. Given the anchors with different accuracy-efficiency trade-offs from the model zoo, our aim is to train an elastic joint network that covers a large number of stitches in a highly efficient way so that it can fit diverse resource constraints with low energy cost. The detailed training algorithm is provided in Algorithm 1 with PyTorch style, where we firstly define a configuration set that contains all possible stitches and initialize all stitching layers with least-squares matching by solving Eq. (4). Next, at each training iteration, we randomly sample a stitch and follow the standard training process as in common practices [29, 50]. To further improve the performance of stitches, we also adopt knowledge distillation with RegNetY-160 [39] as the teacher model. The overall training process requires only a few epochs (e.g., 50) on ImageNet, which is far less than the supernet training in NAS [4, 5, 56] and other techniques [59, 61] that train networks from scratch. Moreover, as anchors are already well-trained, we do not observe significant interference [4] among stitches, as shown in the experiments.

Algorithm 1 Training Stitchable Neural Networks

Require: M pretrained anchors to be stitched. Configuration set $E = \{e_1, \dots, e_Q\}$ with Q stitching positions.

- 1: Initialize all stitching layers by least-squares matching
 - 2: **for** $i = 1, \dots, n_{iters}$ **do**
 - 3: Get next mini-batch of data \mathbf{X} and label \mathbf{Y} .
 - 4: Clear gradients, $optimizer.zero_grad()$.
 - 5: Randomly sample a stitching e_q from set E .
 - 6: Execute the current stitch, $\hat{\mathbf{Y}} = F_{e_q}(\mathbf{X})$.
 - 7: Compute loss, $loss = criterion(\hat{\mathbf{Y}}, \mathbf{Y})$.
 - 8: Compute gradients, $loss.backward()$.
 - 9: Update weights, $optimizer.step()$.
 - 10: **end for**
-

4. Experiment

Implementation details. We conduct all experiments on ImageNet-1K [44], a large-scale image dataset which contains ~ 1.2 M training images and 50K validation images from 1K categories. Model performance is measured by Top-1 accuracy. Furthermore, we report the FLOPs and throughput as indicators of theoretical complexity and real speed on hardware, respectively. We study stitching plain ViTs, hierarchical ViTs, CNNs, and CNN with ViT. We choose the representative model families as anchors: DeiT [48], Swin Transformer [29] and ResNet [15]. By default, we randomly sample 100 images from the training set to initialize the stitching layers. For paired stitching, we set the default sliding kernel size as 2 and the stride as 1. For unpaired stitching, we match layers by nearest interpolation. Unless otherwise specified, all experiments adopt a total batch size of 1,024 on 8 V100 GPUs. We train

DeiT/Swin with 50 epochs with an initial learning rate of 1×10^{-4} . For the experiments with ResNet, we train with 30 epochs based on the training scripts from timm [52] with an initial learning rate of 0.05. All other hyperparameters adopt the default setting as in [29, 48, 52]. For hierarchical models, we scale the learning rate of the anchor parameters by 1/10 compared to that of stitching layers.

4.1. Main Results

Stitching plain ViTs. Based on Algorithm 1, we first generate a stitching configuration set by assembling ImageNet-1K pretrained DeiT-Ti/S/B, which contains 71 stitches including 3 anchors. Then we jointly train the stitches in DeiT-based SN-Net on ImageNet with 50 epochs. The whole training and evaluation process takes only around 110 and 3 GPU hours on V100 GPUs, respectively. In Figure 6 left, we visualize the performance of all 71 stitches, including the anchors DeiT-Ti/S/B (highlighted as yellow stars). In general, SN-Net achieves a wide range of successful stitches, where they achieve smoothly increased performance when stitching more blocks from a larger anchor. We also observe a phenomenon of **model-level interpolation** between two anchors: *with the architectures of the stitches become more similar to the nearest larger anchor, the performance also gradually gets closer to it.*

Moreover, we compare individually trained models from scratch and select stitches from our jointly optimized SN-Net. For brevity, we denote “Ti-S” as the stitches with DeiT-Ti/S as anchors and “S-B” as the stitches with DeiT-S/B as anchors. The results are shown in Table 1. As it indicates, compared to individually trained “S-B” stitches, SN-Net achieves even better performance. It is worth noting that some stitches can fail to converge when training from scratch. However, due to all anchors in SN-Net have been well-trained, the stitches can be easily interpolated among them. Also note that “Ti-S” stitches achieve inferior performance than individual ones. We speculate that due to a slightly larger performance gap between DeiT-Ti/S compared to DeiT-S/B, training Ti-S stitches from scratch may help to find a better local optimum. We also notice a performance drop for anchor DeiT-Ti, for which we assume a more intelligent stitch sampling strategy can help in future works. Overall, a single SN-Net can cover a wide range of accuracy-efficiency trade-offs while achieving competitive performance with models that trained from scratch. To be emphasized, SN-Net reduces around $22\times$ training cost (71×300 epochs vs. $3 \times 300 + 50$ epochs) and local disk storage (2,630M vs. 118M) compared to training and saving all individual networks.

Stitching hierarchical ViTs. Furthermore, we conduct experiment by stitching hierarchical ViTs. In particular, we assemble Swin-Ti/S/B trained on ImageNet-22K by stitching the blocks at the first three stages. Note that we do

Table 1. Performance comparisons on ImageNet-1K between individually trained models from scratch with **300 epochs** and stitches selected from our proposed SN-Net trained with **50 epochs**. A single SN-Net with 118.4M parameters can include all possible stitches. We denote “# Ti/S/B Blocks” as the number of stitched blocks chosen from DeiT-Ti/S/B, respectively. “failed” means training such stitched model from scratch fails to converge and incurs “loss is nan”. Throughput is measured on one RTX 3090 and averaged over 30 runs, with a batch size of 64 and input resolution of 224×224 .

# Ti Blocks	# S Blocks	# B Blocks	FLOPs (G)	Throughput (images/s)	Individually Trained		SN-Net	
					Params (M)	Top-1 (%)	Params (M)	Top-1 (%)
12	0	0	1.3	2,839	5.7	72.1		70.6
9	3	0	2.1	2,352	10.0	75.9		72.6
6	6	0	2.9	1,963	14.0	78.2		76.5
3	9	0	3.8	1,673	18.0	79.4		78.2
0	12	0	4.6	1,458	22.1	79.8	118.4	79.5
0	9	3	7.9	1,060	38.7	79.4		80.0
0	6	6	11.2	828	54.6	failed		81.5
0	3	9	14.3	679	70.6	80.3		82.0
0	0	12	17.6	577	86.6	81.8		81.9




Figure 6. Performance of SN-Net by stitching DeiT-Ti/S/B and Swin-Ti/S/B.

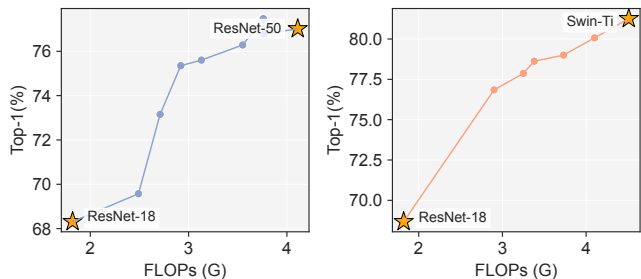


Figure 7. Effect of stitching CNNs and CNN-ViT.

not choose ImageNet-1K pretrained Swin models due to the minor performance gap (83.1% vs. 83.5%) but the large difference in FLOPs (8.7G vs. 15.4G) between Swin-S/B. We visualize the results at Figure 6 right. It shows that the Swin-based SN-Net also achieves flexible accuracy-efficiency trade-offs among the three anchors. This strongly demonstrates that the proposed SN-Net is a general solution for both plain and hierarchical models.

Stitching CNNs and CNN-ViT. We show that SN-Net also works for stitching CNN models and even connecting CNNs with ViTs. As Figure 7 shows, with only 30 epochs of training, the stitches by assembling from ResNet-18 [15] to ResNet-50/Swin-Ti perform favourably, which again emphasizes that SN-Net can be general for both CNNs and

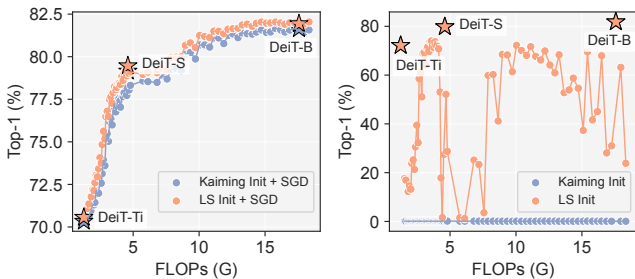


Figure 8. Different learning strategies for stitching layers.

ViTs. Also note that ResNet-18/50 and Swin-Ti are shallow models, so we obtain a small number of stitches.

4.2. Ablation Study

In this section, we ablate the design principles for SN-Net. Unless otherwise specified, our experiments are based on stitching DeiT-Ti/S/B with $k = 2, s = 1$ and knowledge distillation with RegNetY-160. By default, the training strategy is the same as in Section 4.1, e.g., 50 epochs on ImageNet. We provide more ablation studies in the supplementary material, such as the effect of kernel size and stride for controlling the sliding windows during stitching, etc.

Effect of different stitching layer learning strategies. To study the effect of different learning strategies for stitching




Figure 9. From left to right, Figure (a) shows the effect of different stitching directions. Figure (b) presents the effect of nearest stitching based on DeiT, where “Ti”, “S”, “B” denote the stitched anchors. For example, “Ti-S-B” refers to a stitch that defined by connecting the tiny, small and base variants of DeiT, sequentially. Figure (c) shows the comparison of full model tuning vs. tuning stitching layers only.

layers, we consider 4 cases: 1) **Kaiming Init**, the default initialization method in PyTorch. 2) **Least-squares (LS) Init**, the LS solution by solving Eq. (4). 3) **Kaiming Init + SGD**, learning with gradients update on ImageNet after Kaiming Init. 4) **LS Init + SGD**, learning with gradients update on ImageNet after LS Init. We report the experiment results in Figure 8. Overall, we find LS Init serves as a good starting point for learning stitching layers compared to the default Kaiming Init. Interestingly, we observe some stitches by directly matching with LS solution perform quite well compared to Kaiming Init, as shown in Figure 8 right. However, in general, directly matching with LS solution results in an unstable performance curve. This indicates that LS Init is not fully aware of the final performance of the stitches and updating the stitching layers is essential.

Effect of stitching directions. In Figure 9 (a), we compare the stitching directions of Fast-to-Slow and Slow-to-Fast based on DeiT. In general, Fast-to-Slow helps to ensure a better performance for most stitches. On the other hand, Slow-to-Fast obtains a more unstable performance curve, especially for stitching DeiT-S/B. Compared to Fast-to-Slow which increases feature representation capacity by expanding the hidden dimension of activations from a narrower model to a wider one, Slow-to-Fast shrinks the hidden dimension, which contradicts to the existing model design principle [15, 16, 39] that gradually expands the hidden dimension to encode rich semantics as the network goes deeper. Therefore, the resulting information loss of Slow-to-Fast may increase the optimization difficulty.

Effect of nearest stitching. In SN-Net, we adopt the nearest stitching strategy which limits a stitch to connect with a pair of anchors that have the nearest model complexity/performance. However, it is possible to simultaneously stitching more than two anchors (e.g., stitching all DeiT-Ti/S/B sequentially) or anchors with a large gap in complexity/performance (e.g., stitching DeiT-Ti with DeiT-B). With the same 50 epochs training, this approach helps to produce $10\times$ more stitches than our default settings (732 vs. 71). However, as shown in Figure 9 (b), even though Ti-B and Ti-S-B achieve good interpolated performance among

the anchors (i.e., they are stitchable), most of them cannot outperform Ti-S and S-B stitches. In the case of Ti-B, we speculate that without a better minima as a guide in the middle (e.g., DeiT-S), the local minima that found by stitching layers can be sub-optimal due to the large complexity/performance gap between two anchors. Besides, stitching more than two anchors simultaneously does not bring obvious gain at this stage, which we leave for future work.

Effect of tuning full model vs. stitching layers only. In SN-Net, the role of stitching layers is to map the feature maps from one activation space to another. However, since the anchors have been well-trained, one question is how the performance changes if we only update the stitching layers during training. In Figure 9 (c), we show that tuning stitching layers is only promising for some stitches. In contrast, we observe that the performance of stitches can be improved by tuning the full model. Therefore, we by default make SN-Net to be fully updated during training.

5. Conclusion

We have introduced Stitchable Neural Networks, a novel general framework for developing elastic neural networks that directly inherit the rich knowledge from pretrained model families in the large-scale model zoo. Extensive experiments have shown that SN-Net can deliver fast and flexible accuracy-efficiency trade-offs at runtime with low cost, fostering the massive deployment of deep models for real-world applications. With the rapid growth of the number of large-scale pretrained models [13, 38], we believe our work paves a new way for efficient model development and deployment, yielding a significant step towards Green AI. In future works, SN-Net can be extended into more tasks, such as natural language processing, dense prediction and transfer learning.

Limitations and societal impact. Our current training strategy randomly samples a stitch at each training iteration, which implies that with a much larger stitching space, the stitches may not be sufficiently trained unless using more training epochs. We leave this for future work.

References

- [1] Mohamed S. Abdelfattah, Abhinav Mehrotra, Lukasz Dudziak, and Nicholas Donald Lane. Zero-cost proxies for lightweight NAS. In *ICLR*, 2021. [2](#), [3](#), [4](#)
- [2] Yamini Bansal, Preetum Nakkiran, and Boaz Barak. Revisiting model stitching to compare neural representations. In *NeurIPS*, pages 225–236, 2021. [2](#), [3](#), [4](#)
- [3] Hangbo Bao, Li Dong, Songhao Piao, and Furu Wei. Beit: BERT pre-training of image transformers. In *ICLR*, 2022. [3](#), [4](#)
- [4] Han Cai, Chuang Gan, Tianzhe Wang, Zhekai Zhang, and Song Han. Once-for-all: Train one network and specialize it for efficient deployment. In *ICLR*, 2020. [2](#), [6](#)
- [5] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *ICLR*, 2019. [3](#), [6](#)
- [6] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. In *ICCV*, 2021. [3](#)
- [7] Pengguang Chen, Shu Liu, Hengshuang Zhao, and Jiaya Jia. Distilling knowledge via knowledge review. In *CVPR*, pages 5008–5017, 2021. [2](#)
- [8] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four GPU hours: A theoretically inspired perspective. In *ICLR*, 2021. [3](#)
- [9] Xinlei Chen, Saining Xie, and Kaiming He. An empirical study of training self-supervised vision transformers. In *ICCV*, pages 9620–9629, 2021. [3](#)
- [10] Adrián Csizsárik, Péter Korösi-Szabó, Ákos K. Matszangosz, Gergely Papp, and Dániel Varga. Similarity and matching of neural network representations. In *NeurIPS*, pages 5656–5668, 2021. [2](#), [3](#), [4](#)
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. *ICLR*, 2021. [3](#), [4](#)
- [12] David R. Hardoon, Sándor Szedmák, and John Shawe-Taylor. Canonical correlation analysis: An overview with application to learning methods. *Neural Computing*, 16(12):2639–2664, 2004. [3](#)
- [13] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross B. Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 15979–15988, 2022. [3](#), [8](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, pages 1026–1034, 2015. [4](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [3](#), [6](#), [7](#), [8](#)
- [16] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *CVPR*, pages 2261–2269, 2017. [8](#)
- [17] Lang Huang, Shan You, Mingkai Zheng, Fei Wang, Chen Qian, and Toshihiko Yamasaki. Green hierarchical vision transformer for masked image modeling. In *NeurIPS*, 2022. [3](#)
- [18] Roxana Istrate, Florian Scheidegger, Giovanni Mariani, Dimitrios S. Nikolopoulos, Constantine Bekas, and Adelmo Cristiano Innocenza Malossi. TAPAS: train-less accuracy predictor for architecture search. In *AAAI*, pages 3927–3934, 2019. [3](#)
- [19] Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey E. Hinton. Similarity of neural network representations revisited. In *ICML*, volume 97, pages 3519–3529, 2019. [3](#), [5](#)
- [20] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip H. S. Torr. Snip: single-shot network pruning based on connection sensitivity. In *ICLR*, 2019. [2](#)
- [21] Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *CVPR*, pages 991–999, 2015. [2](#), [3](#), [4](#)
- [22] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. In *ICLR*, 2017. [2](#)
- [23] Yanghao Li, Hanzi Mao, Ross B. Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. In *ECCV*, 2022. [3](#)
- [24] Yanjing Li, Sheng Xu, Baochang Zhang, Xianbin Cao, Peng Gao, and Guodong Guo. Q-vit: Accurate and fully quantized low-bit vision transformer. In *NeurIPS*, 2022. [3](#)
- [25] Mingbao Lin, Rongrong Ji, Yan Wang, Yichen Zhang, Baochang Zhang, Yonghong Tian, and Ling Shao. Hrank: Filter pruning using high-rank feature map. In *CVPR*, pages 1526–1535, 2020. [2](#)
- [26] Tsung-Yi Lin, Michael Maire, Serge J. Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C. Lawrence Zitnick. Microsoft COCO: common objects in context. In David J. Fleet, Tomás Pajdla, Bernt Schiele, and Tinne Tuytelaars, editors, *ECCV*, pages 740–755, 2014. [4](#)
- [27] Yang Lin, Tianyu Zhang, Peiqin Sun, Zheng Li, and Shuchang Zhou. Fq-vit: Post-training quantization for fully quantized vision transformer. In *IJCAI*, pages 1173–1179, 2022. [3](#)
- [28] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: differentiable architecture search. In *ICLR*, 2019. [2](#), [3](#)
- [29] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, pages 9992–10002, 2021. [3](#), [4](#), [6](#)
- [30] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. In *NeurIPS*, pages 7827–7838, 2018. [3](#)
- [31] Joe Mellor, Jack Turner, Amos J. Storkey, and Elliot J. Crowley. Neural architecture search without training. In *ICML*, volume 139, pages 7588–7598, 2021. [3](#)
- [32] Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Data-free quantization through weight equalization and bias correction. In *ICCV*, pages 1325–1334, 2019. [2](#)

- [33] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *TKDE*, pages 1345–1359, 2010. 4
- [34] Zizheng Pan, Bohan Zhuang, Haoyu He, Jing Liu, and Jianfei Cai. Less is more: Pay less attention in vision transformers. In *AAAI*, pages 2035–2043, 2022. 3
- [35] Zizheng Pan, Bohan Zhuang, Jing Liu, Haoyu He, and Jianfei Cai. Scalable visual transformers with hierarchical pooling. In *ICCV*, pages 377–386, 2021. 3
- [36] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *NeurIPS*, pages 8024–8035, 2019. 4
- [37] Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *ICML*, volume 80, pages 4092–4101, 2018. 2, 3
- [38] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, pages 8748–8763, 2021. 8
- [39] Ilija Radosavovic, Raj Prateek Kosaraju, Ross B. Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, pages 10425–10433, 2020. 6, 8
- [40] Maithra Raghu, Justin Gilmer, Jason Yosinski, and Jascha Sohl-Dickstein. SVCCA: singular vector canonical correlation analysis for deep learning dynamics and interpretability. In *NeurIPS*, pages 6076–6085, 2017. 3
- [41] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, pages 13937–13949, 2021. 3
- [42] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*, pages 525–542, 2016. 2
- [43] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. In *ICLR*, 2015. 2
- [44] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, pages 211–252, 2015. 4, 6
- [45] Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. Green ai. *Commun. ACM*, 63(12):54–63, nov 2020. 2
- [46] Robin Strudel, Ricardo Garcia Pinel, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *ICCV*, pages 7242–7252, 2021. 3
- [47] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for BERT model compression. In *EMNLP-IJCNLP*, pages 4322–4331, 2019. 2
- [48] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357, 2021. 1, 3, 6
- [49] Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. In *ICCV*, pages 32–42, 2021. 3, 4
- [50] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, pages 548–558, 2021. 3, 4, 6
- [51] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. In *NeurIPS*, pages 11960–11973, 2021. 3
- [52] Ross Wightman. Pytorch image models, 2019. 1, 2, 4, 6
- [53] Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. Transformers: State-of-the-art natural language processing. In *EMNLP*, pages 38–45, 2020. 1
- [54] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M. Alvarez, and Ping Luo. Segformer: Simple and efficient design for semantic segmentation with transformers. In *NeurIPS*, pages 12077–12090, 2021. 3
- [55] Saining Xie, Ross B. Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, pages 5987–5995, 2017. 3
- [56] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. SNAS: stochastic neural architecture search. In *ICLR*, 2019. 3, 6
- [57] Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. In *NeurIPS*, pages 30008–30022, 2021. 3
- [58] Xingyi Yang, Daquan Zhou, Songhua Liu, Jingwen Ye, and Xinchao Wang. Deep model reassembly. *NeurIPS*, 2022. 3
- [59] Jiahui Yu and Thomas S. Huang. Universally slimmable networks and improved training techniques. In *ICCV*, pages 1803–1811, 2019. 6
- [60] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas S. Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In *ECCV*, pages 702–717, 2020. 5
- [61] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas S. Huang. Slimmable neural networks. In *ICLR*, 2019. 2, 6
- [62] Dongqing Zhang, Jiaolong Yang, Dongqiangzi Ye, and Gang Hua. Lq-nets: Learned quantization for highly accurate and compact deep neural networks. In *ECCV*, 2018. 2
- [63] Jinghao Zhou, Chen Wei, Huiyu Wang, Wei Shen, Cihang Xie, Alan L. Yuille, and Tao Kong. Image BERT pre-training with online tokenizer. In *ICLR*, 2022. 3

- [64] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017. 3
- [65] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V. Le. Learning transferable architectures for scalable image recognition. In *CVPR*, pages 8697–8710, 2018. 3