

Vision Transformer Slimming: Multi-Dimension Searching in Continuous Optimization Space

Arnav Chavan^{*1,3}, Zhiqiang Shen^{*2,3}, Zhuang Liu⁴, Zechun Liu⁵, Kwang-Ting Cheng⁶ and Eric Xing^{2,3}

¹IIT Dhanbad ²CMU ³MBZUAI ⁴UC Berkeley ⁵Reality Labs, Meta Inc. ⁶HKUST

arnav.18je0156@am.iitism.ac.in {zhiqians, zechunl}@andrew.cmu.edu zhuangl@berkeley.edu
timcheng@ust.hk epxing@cs.cmu.edu

Abstract

This paper explores the feasibility of finding an optimal sub-model from a vision transformer and introduces a pure vision transformer slimming (ViT-Slim) framework. It can search a sub-structure from the original model end-to-end across multiple dimensions, including the input tokens, MHSA and MLP modules with state-of-the-art performance. Our method is based on a learnable and unified ℓ_1 sparsity constraint with pre-defined factors to reflect the global importance in the continuous searching space of different dimensions. The searching process is highly efficient through a single-shot training scheme. For instance, on DeiT-S, ViT-Slim only takes ~ 43 GPU hours for the searching process, and the searched structure is flexible with diverse dimensionalities in different modules. Then, a budget threshold is employed according to the requirements of accuracy-FLOPs trade-off on running devices, and a re-training process is performed to obtain the final model. The extensive experiments show that our ViT-Slim can compress up to 40% of parameters and 40% FLOPs on various vision transformers while increasing the accuracy by $\sim 0.6\%$ on ImageNet. We also demonstrate the advantage of our searched models on several downstream datasets. Our code is available at <https://github.com/Arnav0400/ViT-Slim>.

1. Introduction

Transformer [49] has been a strong network model for various vision tasks, such as image classification [14, 28, 38, 46, 53], object detection [3, 43, 64], segmentation [50, 54, 61], etc. It is primarily composed of three underlying modules: Multi-Head Self-Attention (MHSA), Multi-Layer Perceptron (MLP) and Image Patching Mechanism. The major limitation of ViT in practice is the enormous model sizes

and excessive training and inference costs, which hinders its broader usage in real-world applications. Thus, a large body of recent studies has focused on compressing the vision transformer network through searching a stronger and more efficient architecture [6–8, 34, 42] for better deployment on various hardware devices.

However, many commonly-used searching strategies are recourse-consuming, such as the popular reinforcement learning and evolutionary searching approaches. Single-path one-shot (SPOS) [15] is an efficient searching strategy and promising solution for this task, while it still needs to train the *supernet* for hundreds of epochs and then evaluate thousands of *subnets* for finding out the optimal sub-architecture, which is still time-consuming, typically with more than tens of GPU days.

Recently, there are some works utilizing Batch Normalization (BN) scaling parameter as the indicator for the importance of operations to prune or search subnets, such as Network Slimming [27], SCP [23], BN-NAS [5], etc., since BN’s parameter is an existing factor and light-weight measure metric in the network for the importance of subnets. This searching method can offer $10\times$ speedup of training than SPOS in general. But in practice, not all networks contain BN layers, such as the transformers. Also, there are many unique properties in a transformer design like the dependency of input tokens from shallow layers to deep layers. Simply utilizing such a strategy is not necessarily practical or optimal for the newer transformer models.

Consequently, the main remaining problem is that there is no BN layer involved in conventional transformer architectures, so we cannot directly employ the scaling coefficient in BN as the indicator for searching. To address this, in this work we propose to incorporate the explicitly soft masks to indicate the global importance of dimensions across different modules of a transformer. We consider jointly searching on all three dimensions in a transformer end-to-end, including: layerwise tokens/patches, MHSA and MLP dimensions. In particular, we design additional differentiable soft masks on different modules for the in-

^{*}indicates equal contribution. This work was done when Arnav was a research assistant at MBZUAI, Zhiqiang Shen is the corresponding author.

Method	Target Arch	Searching Space	Searching Method	Searching Time	Inherit Pre-train	Reduce Params & FLOPs
GLiT [6]	SA + 1D-CNN	Discrete, Pre-defined	Two stage evolutionary	200-ep	No	Both
S ² ViTE [8]	ViT/DeiT family	Continuous, limited	Iterative Prune & Grow	510 g-hrs (600-ep)	Yes	Both
Dynamic-ViT [34]	ViT/DeiT family	Dynamic Patch Sel.	Layerwise Prediction module	26 g-hrs (30-ep)	Yes	FLOPs only
Patch-Slimming [45]	ViT/DeiT family	Layerwise Patch Sel.	Top-Down layerwise	31 g-hrs (36-ep)	Yes	FLOPs only
ViTAS [42]	Customized ViT	Discrete, Pre-defined	Evolutionary	300-ep	No	Both
AutoFormer [7]	Customized ViT	Discrete, Pre-defined	Evolutionary	500-ep	No	Both
ViT-Slim (Ours)	ViT/DeiT/Swin, etc.	Continuous, all modules	One-shot w/ ℓ_1-sparsity	43 g-hrs (50-ep)	Yes	Both

Table 1. Feature-by-feature comparison of compression and search approaches for vision transformers. “g-hrs” indicates GPU hours.

dividual dimension of a transformer, and the ℓ_1 -sparsity is also imposed to force the masks to be sparse during searching. We only need a few epochs to finetune these mask parameters (they are initialized to 1 so as to give equal importance to all dimensions at the beginning of search) together with the original model’s parameters for completing the whole searching process, which is extremely efficient. For the token search part, we apply a *tanh* over masks to avoid exploding mask values which is observed empirically.

We call our method *ViT-Slim*, a joint sparse-mask based searching method with an implicit weight sharing mechanism for searching a better sub-transformer network. This is a more general and flexible design than previous BN-based approaches since we have no requirement of BN layers in the networks. This is more friendly to transformers and a feature-by-feature comparison with other ViT compression methods is shown in Table 1. One core advantage of our method is the efficiency of searching, we can inherit the pre-trained parameters and conduct a fast search upon it. Another advantage is the zero-cost subnet selection and high flexibility. In contrast to the SPOS searching that requires to evaluate thousands of subnets on validation data, once we complete the searching process, we can obtain countless subnetworks and the final structure can be determined by the requirement of accuracy-FLOPs trade-off of the real devices that we deploy our model on, without any extra evaluation. The last advantage is that we can search for a finer-grained architecture such as the different dimensionalities in different self-attention heads, as our search space is continuous in them. This characteristic allows us to find the architecture with unique individual dimensions and shapes in different layers and modules, which would always find out a better subnet than other counterparts.

Comprehensive experiments and ablation studies are conducted on ImageNet [13], which show that ViT-Slim can compress up to 40% of parameters and 40% FLOPs on various vision transformers like DeiT [47], Swin [28] without any compromising accuracy (in some circumstances our compressed model is even better than the original one). We also demonstrate the advantage of our searched models on several downstream datasets.

Our main contributions are:

- We introduce ViT-Slim, a framework that can jointly perform an efficient architecture search over all three modules - MHSA, MLP and Patching Mechanism in vision transformers. We stress that our method searches for

structured architectures which can bring practical efficiency on modern hardware (e.g., GPUs).

- We empirically explore various structured slimming strategies by sharing weights in candidate structures, and provide the best performing structure by employing a continuous search space in contrast to a pre-defined discrete search space in existing works.
- Our method can perform directly over pre-trained transformers by employing a single shot searching mechanism for all possible budgets, eliminating the need of search-specific pre-training of large models and performing repeated searching for different modules/budgets.
- We achieve state-of-the-art performance at different budgets on ImageNet across a variety of ViT compression and search variants. Our proposed ViT-Slim can compress up to 40% of parameters and 40% FLOPs while increasing the accuracy by $\sim 0.6\%$.

2. Related Work

Efficient Models and Architecture Search Neural network compression has been recognized as an important technology for applying deep neural network models to equipment with limited resources. The compression research extends from channel pruning [17, 27, 56], quantization and binarization [1, 9, 21, 29, 33, 63], knowledge distillation [18, 32, 36, 39–41], compact neural network design [10, 22, 31, 52, 60] to architecture search [35, 65, 66].

Specifically, MobileNets [19, 37] proposed to decompose the convolution filters to depth-wise convolution and point-wise convolution for reducing the parameters in convolutional neural networks (CNNs). EfficientNet [44] proposed to search the uniform scaling ratio among all dimensions of depth/width/resolution to achieve much better accuracy and efficiency. Network Slimming [27] used the BN parameters as the scaling factors to find the best sub-structure. JointPruning [30] jointly searched the layer-wise channel number choices together with the depth and resolution for finer-grained compression. NetAdapt [55] and AMC [16] adopted the feedback loop or the reinforcement learning method to search the channel numbers for the CNNs. Besides, many neural architecture search (NAS) methods are targeting at exploring the operation choices (e.g., 3×3 , 5×5 , 7×7 convolutions) for architectures. For instance, SPOS [15] built a supernet that contained all the possible choices and used the evolutionary search for a subnetwork

in the supernet. DARTS [26], FB-Net [51] and Proxyless-NAS [2] used gradient-based method to update the mask associating with each operation choices. However, these NAS methods defined on the discrete operation search spaces can hardly be generalized to tackle the problem of continuous channel number search.

Efficient Vision Transformers. There are several works exploring on this direction [6–8, 34, 42]. Patch-Slimming [45] explored the direction to improve the efficiency of transformers by sequentially pruning patches from top-to-bottom layers. Similarly, Dynamic-ViT [34] explored dynamic patch selection based on the input patches. They employed multiple hierarchical prediction modules to estimate the importance score of each patch. However, patch pruning did not improve parameter efficiency. ViTAS [42] used the evolutionary algorithm to search optimal architecture at a target budget. However, their search space was discrete, pre-defined and thus limited. GLiT [6] introduced a locality module to model local features along with the global features. But their method used CNNs along with the attention and performed an evolutionary search over global and local modules. BigNAS [57] introduced the single-stage method to generate highly efficient child models by slicing weight matrices. Based on this, AutoFormer [7] showed that weight entanglement was a better alternative than defining weight matrices for every possible sub module in architecture search and using an evolutionary algorithm to search for optimal sub-networks. But all of them had a limited discrete search space as [42] due to the adaptation of an evolutionary algorithm for searching. S²ViTE [8] provided an end-to-end sparsity exploration for vision transformers with an iterative pruning and growing strategy. Their structured pruning method eliminated complete attention heads and MLP neurons based on a scoring function calculated using the Taylor expansion for the loss function and ℓ_1 -norm, respectively. We argue that eliminating complete attention heads is a sub-optimal choice and limits the learning dynamics of a transformer. Allowing the model to determine optimal dimensions for every attention head (instead of eliminating complete heads) is a better alternative for pruning MHSA module.

3. Proposed Approach

3.1. Overview and Motivation

In this section, we start by working around some important questions raised by existing works, viz:

- Can we take advantage of the existence of only fully connected layers in transformers and make the search space continuous and hence much larger than existing works without larger memory or compute overhead?
- What is the optimal structure configuration that has to be searched in ViT family, and can joint search of architec-

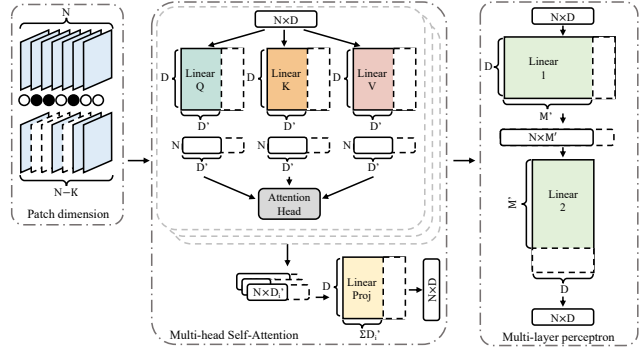


Figure 1. An overview of our proposed ViT-Slim framework. D' and M' are the dimensions after slimming operation.

ture (MHSA/MLP dimension search) and layer-wise data flow mechanism inside the architecture (layer-wise patch selection) be coupled together in a single shot setup?

- What is the impact of individual modules to the final model’s performance and can vision transformers which are originally designed to have an isotropic structure benefit from a highly non-uniform structure?

An overview of our framework is shown in Figure 1. In the following, we will discuss (i) How to achieve a continuous search space; (ii) Identify optimal search space; and (iii) Single-shot architecture search with the ℓ_1 -Sparsity.

3.2. Achieving A Continuous Search Space

One-shot NAS methods for CNNs [11, 15, 44] explicitly define multiple decoupled candidate blocks for every layer for training a central supernet. This strategy is suitable for CNNs as the candidate blocks at every layer come from a wide variety of sub-architectures so as to maintain the property of sampling diverse subnets from the supernet while searching. This is not true for transformers, as they are internally composed of multiple fully connected layers stacked in different configurations for different modules. The fact that the core blocks are all composed of fully connected layers opens the possibility to expand the search space by sharing weights between the candidate fully connected layers inside any block.

Consider a fully connected layer which has to be searched for optimal output dimension given an input dimension D_{in} , conventional method is to define multiple candidate layers with output dimensions from pre-defined search space and search the optimal layer from them with a suitable searching algorithm. However, there are a couple of disadvantages - 1) Longer search time as every candidate layer needs to be at least partially trained for searching 2) Larger memory footprint contributed by weight matrices of every candidate layer. We propose to solve these issues by weight sharing between all possible candidate layers. We fix the maximum permissible output dimension to D_{max}

and define a super weight matrix $W_{sup} \in D_{in} \times D_{max}$. The candidate layer weights can be easily sliced from W_{sup} .

To achieve a continuous search space, we adopt a single-stage method to rank the importance of every dimension of supernet weights, taking inspiration from [27] we use ℓ_1 -Sparsity to achieve it. We first pre-train the supernet until convergence (in practice, our method would work directly over pre-trained networks eliminating the need of supernet training). We then define masks corresponding to every dimension that has to be searched. The magnitude of mask values correspond to the importance score of respective dimension and thus we initialize all of them by 1. Once the pretrained weights are loaded we induce these masks into the model. Considering the previous example of a single fully-connected layer where the output dimension is to be searched of a layer with weight matrix W_{sup} , a mask $z \in D_{max}$ is defined. A dot product between W_{sup} and z gives the candidate weight matrix to be used in forwarding propagation. The search algorithm employs a loss function which is a combination of pre-training loss function (Cross-Entropy in classification tasks) and ℓ_1 -norm on masks. This combined loss drives the mask values towards 0 while minimizing the target loss function. In a way the optimization landscape implicitly drives the masks to rank themselves according to their impact on final performance.

3.3. Identifying Optimal Search Space

The fundamental question in NAS is defining a search space. This section presents our defined search space.

Flexible MHSA and MLP Modules. Recent works have followed two methodologies for defining the search space for MHSA module - 1) Searching for the number of heads in every distinct MHSA module [8, 25] and/or 2) Searching for a common feature dimension size from a pre-defined discrete sample space for all attention heads in any particular MHSA module [6, 7, 42]. These methods have shown some solid results but they are not completely flexible. A greater degree of flexibility can be achieved if every attention head can have a distinct feature dimension size. Assuming a super-transformer network with a maximum of L MHSA modules each with a maximum permissible number of heads set to H . This gives a total of $L \times H$ unique attention heads. If we fix the maximum permissible feature dimension size to d , the size of the equivalent search space is equal to $(d + 1)^{L \times H}$. Searching in such a massive search space is computationally very difficult. However, such a diverse search space has a couple of advantages: 1) The search algorithm can be more flexible in adapting the super-transformer to smaller architectures while maintaining the performance; and 2) The extracted architectures would be much more efficient as the search algorithm can push the feature dimensions of least important attention heads to even zero, decreasing the FLOPs substantially. We use the

complete $(d + 1)^{L \times H}$ search space in our method.

Similarly, for the MLP, there are a total of L modules throughout the network with existing works using discrete and limited search space to search the optimal feature dimension. If we fix the maximum permissible feature dimension size to M , the search space which can be explored by our method is equal to $(M + 1)^L$. Combining MHSA and MLP together in a single search mechanism is quite straightforward. Thus, it will generate a much more diverse set of child networks compared to all the existing works.

Patch Selection. Patch slimming [45] showed that MHSA aggregates patches and consequently cosine-similarity between the patches increases exponentially layer-wise becoming as large as 0.9 in the final layers. This opens the possibility to eliminate a large number of deeper patches and a few unimportant shallow patches too. Inducing patch-selection with MHSA and MLP search can extract more efficient architectures with reduced FLOPs at the same parameter count. While intuitively, a dynamic way is more promising on the patch dimension as the selected patches should be aligned with the input images to reflect the importance of different regions in the image.

3.4. Single Shot Arch. Search with ℓ_1 -Sparsity

The main objective of our search method is to rank the mask values according to their impact on final performance. Once ranked, we eliminate the dimensions having the least mask values. Let $f : \mathbb{R}^x \rightarrow \mathbb{R}^y$ denote a vision transformer network which learns to transform input \mathbf{x} into target outputs \mathbf{y} with weights \mathbf{W} and intermediate activations/tensors $\mathbf{T} \in \mathbb{R}^d$ constructed from weights \mathbf{W} and input \mathbf{x} . We define $\mathbf{z} \in \mathbb{R}^d$ as a set of sparsity masks, where $z_i \in \mathbf{z}$ refers to the mask associated with the intermediate tensor $\mathbf{t}_i \in \mathbf{T}$. To apply the mask, z_i is multiplied with corresponding entries of \mathbf{t}_i . The optimization objective can be stated as:

$$\min_{\mathbf{z}, \mathbf{z}} \mathcal{L}_{CE}(f(\mathbf{z} \odot \mathbf{T}(\mathbf{W}, \mathbf{x})), \mathbf{y}) + \|\mathbf{z}\|_1, \quad (1)$$

We introduce uniform masking to search optimal dimension size for each distinct head in MHSA modules, dimension size for each MLP module and layerwise most important patches. Consider a transformer network with L layers of MHSA+MLP blocks and each block consisting of H self-attention heads. The input tensor to each MHSA layer $\mathbf{t}_{a_l} \in \mathbb{R}^{N \times D}$ where N is number of patches and D is global feature dimension size. Inside every head i of MHSA module, \mathbf{t}_{a_l} is transformed with fully connected layers to $\mathbf{q}_i \in \mathbb{R}^{N \times d}$, $\mathbf{k}_i \in \mathbb{R}^{N \times d}$ and $\mathbf{v}_i \in \mathbb{R}^{N \times d}$, d denoting feature dimension of each self-attention head. We define mask $\mathbf{z}_a \in \mathbb{R}^{L \times H}$, and corresponding vectors in \mathbf{z}_a , $z_{a_l, h} \in \mathbb{R}^d$ corresponding to l^{th} layer and h^{th} head. The total possibilities for MHSA modules across the network that can be explored by our method is thus $(d + 1)^{L \times H}$. The computa-

tions inside MHSA module with sparsity masks is:

$$A_i = \text{softmax}((q_i \odot z_{a_i,h}) \times (k_i \odot z_{a_i,h})^T / \sqrt{d}) \quad (2)$$

$$O_i = A_i \times (v_i \odot z_{a_i,h}) \quad (3)$$

$$\mathbf{t}_{m_1} = \text{projection}([O_1, O_1, \dots, O_H]) \quad (4)$$

where $\mathbf{t}_{m_1} \in \mathbb{R}^{N \times D}$ is the output from MHSA block which in turn becomes the input to MLP block. Inside MLP block, \mathbf{t}_{m_1} is projected to a higher-dimensional space through a fully connected layer f_1 to form an intermediate tensor $\mathbf{t}_{e_1} \in \mathbb{R}^{N \times M}$ which is again projected back to $\mathbb{R}^{N \times D}$ through another fully connected layer f_2 . We define mask $\mathbf{z}_m \in \mathbb{R}^L$, and corresponding vectors in \mathbf{z}_m , $z_{m_l} \in \mathbb{R}^M$ corresponding to l^{th} layer. The total possibilities for MLP modules across the network is thus $(M + 1)^L$ by our method. The following computation shows the interaction of masks with MLP module:

$$\mathbf{t}_{e_1} = f_1(\mathbf{t}_{m_1}) \odot z_{m_l}, \mathbf{t}_{a_{l+1}} = f_2(\mathbf{t}_{e_1}) \quad (5)$$

Solving Patch Dependency Across Layers in Patch Dim.

For patch selection, we define a distinct mask value corresponding to each patch in every layer and eliminate the patches corresponding to a lower mask value. A problem arises due to the global single shot search that there may be anomalous instances where the same patch is eliminated in a shallower layer but exists in the deeper layer. In practice, such anomalous instances are limited indicating that ℓ_1 -Sparsity based search strategy aligns the patches as per their importance and a shallow eliminated patch implicitly forces its deeper counterpart to have less importance and consequently eliminates it. To counter these limited anomalous patches, once a patch is eliminated from an earlier layer, we eliminate it from further layers too, while imposing budget. Additionally, we apply a \tanh ¹ activation function over patch-specific masks before taking a dot product with the patches to stop mask values from exploding.

3.4.1 Searching Time Analysis

Our method directly works over pre-trained models, eliminating the need of training a search-specific model. We induce sparsity masks on a pre-trained model and jointly optimize the masks and model weights with a combination of CE loss and ℓ_1 -norm of masks. In our setup, we fix the search epochs to 50 for all searches. This translates to ~ 43 GPU hrs for DeiT-S and ~ 71 GPU hrs for DeiT-B. At the end of search, masks are ranked according to their values.

3.4.2 Retraining with Implicit Budget Imposition

Once ranked, depending upon the target budget, low-rank dimensions/patches are eliminated from the network. For MHSA+MLP joint search, budget approximately translates

¹initialized to 3.0 which is equivalent to ~ 1.0 after \tanh .

to FLOPs and Params budget too due to the linear relation between the number of dimensions and FLOPs/Params, and inducing patch selection further decreases FLOPs. Once the budget-specific structure is extracted, it is retrained with exactly the same setup as the pre-trained model. This allows the weights to adjust from continuous masks in searching space to binary/non-existing masks in the final structure.

4. Experiments

In this section, we first explore the contribution of each individual component in the final performance of the ViT [14]/DeiT [46] model and search for the optimal unidimensional searched model. We then move to joint search combining all three components and show that our method outperforms all existing architecture search and pruning methods. We also show the applicability of our method to other transformer architectures, such as Swin [28]. Finally, we further show the performance of searched models on the downstream datasets in a transfer learning setup.

4.1. Training Procedures and Settings

There are three steps in our workflow for the ViT-Slim framework, including:

(i) One-shot Searching. We use pre-trained weights to initialize existing vision transformer models and use them as our supernet. We then induce sparsity masks into the model depending upon the dimension to search and jointly train the weights and masks with the loss function given in Equation 1 for 50 epochs with a constant learning rate of $5e-04$ and AdamW optimizer with a weight decay of $1e-03$. The batch size is set to 1024 for DeiT-S and 512 for DeiT-B. We also employ stochastic depth [20], cutmix [58], mixup [59], randaugment [12], random erasing [62], etc., as augmentations following DeiT while searching.

(ii) Budget Selection. Once the search is complete, the masks are ranked according to their values after the searching step. Depending upon the target budget of compression, low-rank dimensions are eliminated from the supernet to extract the final structure of searched model.

(iii) Re-training Finally, we retrain the extracted compressed structure for 300 epochs with the same setting on which it was originally pre-trained in [28, 46].

4.2. Unidimensional Search

To show the impact of MHSA and MLP modules independently in the final model’s performance, we search for optimal dimensions in both of them separately. We induce sparsity masks in respective modules of DeiT-S and search two supernets each with a sparsity weight of $1e-04$ for MLP and MHSA dimensions. The total number of masks induced in MLP is $4 \times$ that of MHSA due to the existence of $4 \times$ more dimensions in MLP module. The post-search accuracy and final accuracy at different budgets are shown in Ta-

Budget (%)	Module	#Params (M)	FLOPs (B)	Top-1 (%)	Top-5 (%)
100	-	22.0	4.6	79.90	95.01
Post-Search		22.0	4.6	76.85	93.70
70	MHSA	19.9	4.1	80.90	95.44
60		19.2	3.9	80.63	95.31
50		18.5	3.7	80.10	95.07
40		17.8	3.5	79.61	94.73
Post-Search		22.0	4.6	76.85	93.70
70	MLP	17.8	3.8	80.80	95.37
60		16.4	3.5	80.39	95.28
50		15.0	3.2	79.89	95.05
40		13.5	2.9	79.20	94.75

Table 2. Performance of **DeiT-S** [47] for MHSA and MLP dimension search. Budget indicates the % of active dimensions of respective search modules across the network.

W_1	W_2	Post Search Accuracy		Final Accuracy	
		Top-1 (%)	Top-5 (%)	Top-1 (%)	Top-5 (%)
1e-04	1e-04	76.54	93.37	79.10	94.65
3e-04	1e-04	76.34	93.36	79.00	94.70
4e-04	1e-04	76.40	93.39	78.71	94.57
2e-04	4e-05	76.69	93.61	79.17	94.72
2e-04	5e-05	76.68	93.64	79.20	94.75

Table 3. W_1 (MHSA sparsity weight) and W_2 (MLP sparsity weight) grid search for **DeiT-S** [47] (**ViT-Slim_{PS}**). *Post search* indicates post search soft accuracy of supernet. *Final accuracy* indicates final accuracy of the compressed models at 60% budget.

ble 2. At higher budgets, compressed models even perform better than pre-trained model giving as much as 1% boost in accuracy. MHSA search performs better than MLP search at the same budgets, but MLP achieves a better degree of parameter and FLOPs compression. MHSA at 40% and MLP at 60% have the same FLOPs but MLP outperforms MHSA. Similarly, MHSA at 40% and MLP at 70% having the same number of parameters, MLP outperforms MHSA by a fair margin of 1.2%. This clearly shows that it is much easier to compress MLP dimensions as compared to MHSA dimensions to achieve the same target FLOPs/Parameters, indicating the importance of MHSA to be greater than MLP.

4.3. Partial Combination for Parameter Search

Next, we combine MHSA and MLP in a single supernet search. The most important hyperparameter to control is the sparsity weight for each of different modules. Based on the fact that MLP is easy to compress and has $4\times$ more dimensions than MHSA, we expect the optimal sparsity weight to be in a similar ratio. We start with an equal sparsity weight of $1e-04$ and do a thorough grid search to achieve the optimal performance as shown in Table 3. We search and retrain at 60% budget for a fair comparison. From the results, it is clear that post-search accuracy directly reflects the final model accuracy. As expected, the optimal sparsity weights $2e-04$ and $5e-05$ are in the ratio of 4:1.

We retrain the optimal searched model at multiple budgets as shown in Table 4. The performance of the model is intact and even better than the pre-trained model up to a budget of 70% after which the accuracy starts to drop. This

B-MLP	B-MHSA	#Params (M)	FLOPs (B)	Top-1 (%)	Top-5 (%)
100	100	22.0	4.6	79.90	95.01
80	80	17.7	3.7	80.60	95.29
70	70	15.6	3.3	80.03	95.05
60	60	13.5	2.8	79.20	94.75
50	50	11.4	2.3	77.94	94.14

Table 4. Performance of **DeiT-S** [47] (**ViT-Slim_{PS}**) for MLP and MHSA joint dimension search. Budget indicates the % of active MHSA and MLP dimensions across the network respectively.

W_1	W_2	W_3	Post Search Accuracy	
			Top-1 (%)	Top-5 (%)
2e-04	5e-05	1e-04	76.92	93.54
2e-04	5e-05	2e-04	76.76	93.54
2e-04	5e-05	5e-05	76.79	93.55

Table 5. W_1 (MHSA sparsity weight), W_2 (MLP sparsity weight) and W_3 (Patch sparsity weight) grid search for **DeiT-S** (**ViT-Slim_{JS}**).

Budget	#Params (M)	FLOPs (B)	Top-1 (%)	Top-5 (%)
100	15.6	3.3	80.03	95.01
80	15.6	3.1	79.91	94.97
70	15.6	2.9	79.72	94.91
60	15.6	2.8	79.51	94.75

Table 6. Performance of **DeiT-S** [47] (**ViT-Slim_{JS}**) for MLP, MHSA and patch selection joint search. Budget indicates the % of **active patches** across the network. Budget for both MHSA and MLP is fixed at 70% across models.

translates to better performance with 30% FLOPs and parameter reduction. However, going from 60% to 50% budget shows a drastic drop in performance. We name these budgeted model family as **ViT-Slim_{PS}** indicating direct parameter search or partial search without patch selection.

Discussion: Budgeted FLOPs and Parameter Reduction

The budget indicates the number of active dimensions that exist in the final searched model for the respective modules. However, for MHSA and MLP joint search, budget translates to final FLOPs and parameter budget due to a linear relation between them². This can be helpful to deploy large transformer networks as per the target FLOPs/Parameter budget on specific hardware.

4.4. Multidimensional Joint Search

Finally, a joint search over all three dimensions - MHSA, MLP and Patch Selection is performed. Once we identify the optimal MHSA and MLP sparsity weights as $2e-04$ and $5e-05$, respectively, we do a grid search over patch sparsity as shown in Table 5. Optimal patch sparsity weight used in all our experiments is $1e-04$. To show the effect of patch selection at different budgets, we fix the MHSA and MLP budget at 70% (as 70% keeps the performance intact as shown in Table 4) and retrain at multiple patch-

²FLOPs and Parameters are rounded to 1st decimal place creating a small difference between dimension budget and FLOPs/Parameters budget in Table 4.

Budget	Model	#Params (M)	FLOPs (B)	Top-1 (%)	Top-5 (%)
100	DeiT-B	86.6	17.5	81.8	95.6
-	PS-ViT-B [45]	86.6	10.5	81.5	-
-	S ² ViTE-B [8]	56.8	11.7	82.2	-
-	GLiT-B [6]	96.1	17.0	82.3	-
-	AutoFormer-B [7]	54.0	11.0	82.4	95.7
60	ViT-Slim-B	52.6	10.6	82.4	96.1
100		28.3	4.5	81.3	95.5
80	Swin-T	22.3	3.8	81.3	95.5
70		19.4	3.4	80.7	95.4

Table 7. Performance of **DeiT-B** [47] (**ViT-Slim-B_{PS}**) and **Swin-T** [28] for MLP and MHSA joint dimension search. Budget indicates the % of active MHSA and MLP dimensions across the network respectively.

selection budgets as shown in Table 6. Eliminating as much as 40% patches doesn’t cause massive degradation in performance. At 80% budget, performance matches that of pre-trained DeiT-S (79.9% Top-1). We name this model family with multidimensional search as ViT-Slim_{JS} indicating joint search.

4.5. Other Architectures

We further show the efficacy of our method on DeiT-B (ViT-Slim-B_{PS}) with the same set of hyperparameters. We show a thorough comparison with the existing method in Table 7. Our model’s budget is set to 60% which is equivalent to 40% drop in FLOPs and parameters. ViT-Slim outperforms all existing methods with substantially fewer FLOPs and parameters and increases accuracy by $\sim 0.6\%$ of the pre-trained DeiT-B. Note that although our accuracy is comparable to AutoFormer-B, our searching resource is only **1/10** of it with smaller model size and FLOPs.

We also perform a search with the same hyperparameters as ViT-Slim_{PS} on Swin-T [28] and present the results in Table 7. The models are retrained with the same policy as in [28]. Final accuracy is intact at 80% budget but drops at 70% budget. This is partly because of the fact that Swin is a carefully designed hierarchical architecture which already maximises dimensions at every layer and partly because we don’t do a thorough hyperparameter search for sparsity weights specifically for Swin-T. But, still we achieve a fair amount of compression while keeping the accuracy intact.

4.6. Comparison to State-of-the-art Approaches

We compare both our model families - ViT-Slim_{PS} and ViT-Slim_{JS} with existing efficient transformer architecture search and compression methods as shown in Table 8. Our method outperforms all of them at different target parameters and FLOPs. GLiT-S [6] improves accuracy (80.5%) over baseline DeiT-S but with an additional parameter increase and minimal FLOPs reduction. Our ViT-Slim_{PS} achieves better accuracy (80.6%) and substantially decreases both FLOPs and Parameters at the same time. Our ViT-Slim_{JS} model allows for up to 30% parameter reduction and greater than 30% FLOPs reduction while matching the performance of it’s supernet (DeiT-S). Even fur-

Model	#Params (M)	FLOPs (B)	Top-1 (%)	Top-5 (%)
DeiT - S [47]	22.0	4.6	79.9	95.0
GLiT - S [6]	24.6	4.4	80.5	-
DynamicViT - S [34]	22.0	4.0	79.8	-
ViT-Slim_{PS}	17.7	3.7	80.6	95.3
S ² ViTE - S [8]	14.6	3.1	79.2	-
DynamicViT - S [34]	22.0	2.9	79.3	-
PS-ViT-S [45]	22.0	2.7	79.4	-
ViTAS - E [42]	12.6	2.7	77.4	93.8
S ² ViTE+ - S [8]	14.6	2.7	78.2	-
ViT-Slim_{JS}	15.7	3.1	79.9	95.0
ViT-Slim_{JS}	15.7	2.8	79.5	94.6

Table 8. Comparison with SOTA ViT search and compression methods on ImageNet-1k. “**PS**” is the partial search or parameter search, i.e., MHSA+MLP dimensions, since *Patch Selection* will not affect the number of parameters but only help reduce FLOPs in our models. “**JS**” is our final joint-dimension search.

Model	#Params (M)	FLOPs (B)	C100	C10	iNat-19	iNat-18
DEiT-S	22.0	4.6	87.80	98.56	75.35	69.02
ViT-Slim_{PS}	15.6	3.3	88.16	98.70	76.67	69.83

Table 9. Transfer learning accuracy of **DeiT-S** [47] and **ViT-Slim_{PS}** on CIFAR-100 (C100), CIFAR-10 (C10) [24], iNaturalist-2018 (iNat-18) and iNaturalist-2019 [48] (iNat-19) datasets. All models are searched and pre-trained on ImageNet.

ther reducing FLOPs doesn’t cause a drastic degradation in performance of ViT-Slim_{JS} in contrast to others. PS-ViT-S [45] and Dynamic-ViT-S [34] performs comparably to our method with respect to FLOPs, but they don’t bring in any improvement in parameter efficiency.

4.7. Transfer Learning on Downstream Datasets

We analyse the performance of our searched and re-trained model on various downstream classification tasks. We provide results on CIFAR-10, CIFAR-100 [24], inaturalist-2018 [48] and inaturalist-2019 datasets as shown in Table 9. ViT-Slim_{PS} is retrained at 70% budget and they consistently outperform DeiT-S baseline across datasets. An important point to note is that ViT-Slim architectures were searched on ImageNet [13] and not on respective downstream datasets directly, which shows the ability of the same ViT-Slim architectures to transfer well on other downstream tasks too.

5. Visualization and Analysis

5.1. Searched Architecture

Layerwise attention head dimensions at various budgets. Figure 2 shows the MHSA modules across DeiT-S searched model. There are a total of 12 MHSA modules each with 6 attention heads. Numbers inside the grids indicate the dimension size of that particular head. It can be seen that at low budgets, deeper layers have the least dimension sizes. Most of the dimensions are intact in the middle of the network and a moderate reduction in dimensions is done in the beginning of the network. Self-attention mech-

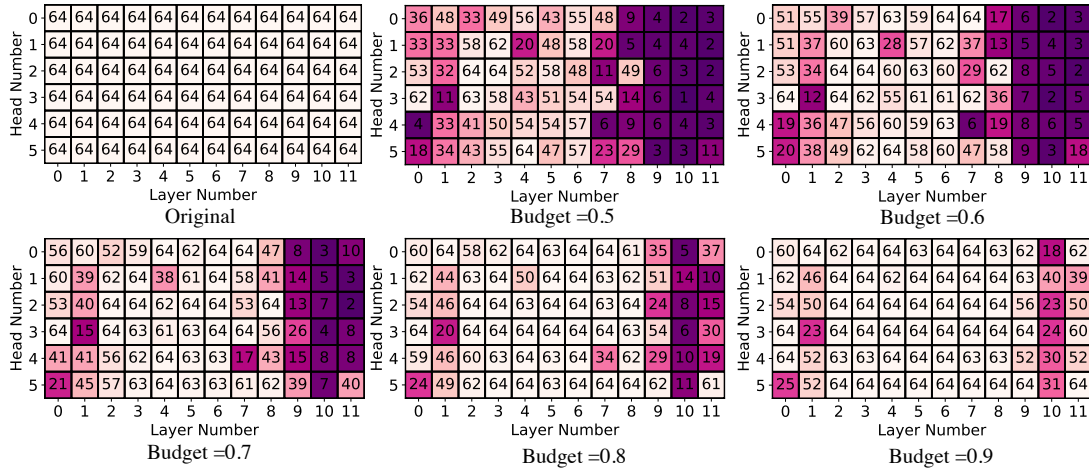


Figure 2. Layerwise Attention Head Dimensions at various budgets (50%, 60%, 70%, 80% and 90%) for ViT-Slim model.




Figure 3. Layerwise MLP dimensions at various budgets (50%, 60%, 70% and 80%) for ViT-Slim model.

anism is required in the middle and to some extent at the beginning of the network when the patches are distinct and information exchange between them is required.

Layerwise MLP dimensions at various budgets. Similarly, Figure 3 shows the MLP module dimensions across DeiT-S searched model at various budgets. The pattern across layers is similar to that of MHSA, where deeper layers have a greater degree of reduced dimensions as compared to the earlier layers. The deeper layers have maximum dimensions removed, while the middle layers have maximum dimensions intact. This is again in agreement with the fact that the majority of features are already learnt in the earlier layers, enabling deeper layers to consist of smaller dimension sizes.

5.2. Attention Maps Visualization

We adopt the method presented in [4] which employs Deep Taylor Decomposition to calculate local relevance and then propagates these relevancy scores through the layers to generate a final relevancy map. Class-wise visualisation of randomly chosen ImageNet images is shown in Figure 4. ViT-Slim_{JS} focuses better on class-specific important areas as compared to DeiT-S and thus achieves better performance. This also shows that ViT-Slim_{JS} has better inter-




Figure 4. Class wise visualisation of ImageNet images with method presented in [4]. First row are the original images, second row represent visualisations from DeiT-S and third row represents ViT-Slim_{JS} at 70% budget.

pretability and thus helps in domains where interpretability and explainability of deep networks are important.

6. Conclusion

We have presented ViT-Slim, a flexible and efficient searching strategy for subnet discovery on vision transformers leveraging the model sparsity. The proposed method can jointly search on all three dimensions in a ViT, including: layerwise tokens/patches, MHSA and MLP modules end-to-end. We identified that the global importance factors are crucial, and designed additional differentiable soft masks on different modules to reflect the individual importance of dimension. Moreover, the ℓ_1 -sparsity is imposed to force the masks to be sparse during searching. Extensive experiments are conducted on ImageNet and downstream datasets using a variety of ViT architectures to demonstrate the efficiency and effectiveness of our proposed method.

References

- [1] Adrian Bulat and Georgios Tzimiropoulos. Xnor-net++: Improved binary neural networks. *British Machine Vision Conference*, 2019. [2](#)
- [2] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. In *International Conference on Learning Representations*, 2018. [3](#)
- [3] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. [1](#)
- [4] Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 782–791, June 2021. [8](#)
- [5] Boyu Chen, Peixia Li, Baopu Li, Chen Lin, Chuming Li, Ming Sun, Junjie Yan, and Wanli Ouyang. Bn-nas: Neural architecture search with batch normalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 307–316, 2021. [1](#)
- [6] Boyu Chen, Peixia Li, Chuming Li, Baopu Li, Lei Bai, Chen Lin, Ming Sun, Junjie Yan, and Wanli Ouyang. Glit: Neural architecture search for global and local image transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12–21, 2021. [1](#), [2](#), [3](#), [4](#), [7](#)
- [7] Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. 2021. [1](#), [2](#), [3](#), [4](#), [7](#)
- [8] Tianlong Chen, Yu Cheng, Zhe Gan, Lu Yuan, Lei Zhang, and Zhangyang Wang. Chasing sparsity in vision transformers: An end-to-end exploration. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. [1](#), [2](#), [3](#), [4](#), [7](#)
- [9] Jungwook Choi, Zhuo Wang, Swagath Venkataramani, Pierce I-Jen Chuang, Vijayalakshmi Srinivasan, and Kailash Gopalakrishnan. Pact: Parameterized clipping activation for quantized neural networks. *arXiv preprint arXiv:1805.06085*, 2018. [2](#)
- [10] François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1251–1258, 2017. [2](#)
- [11] Xiangxiang Chu, Bo Zhang, and Ruijun Xu. Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12239–12248, 2021. [3](#)
- [12] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 702–703, 2020. [5](#)
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#), [7](#)
- [14] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [1](#), [5](#)
- [15] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *European Conference on Computer Vision*, pages 544–560. Springer, 2020. [1](#), [2](#), [3](#)
- [16] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. Amc: Automl for model compression and acceleration on mobile devices. In *ECCV*, 2018. [2](#)
- [17] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1389–1397, 2017. [2](#)
- [18] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. [2](#)
- [19] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. [2](#)
- [20] Gao Huang, Yu Sun, Zhuang Liu, Daniel Sedra, and Kilian Q Weinberger. Deep networks with stochastic depth. In *European conference on computer vision*, pages 646–661. Springer, 2016. [5](#)
- [21] Itay Hubara, Matthieu Courbariaux, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. Quantized neural networks: Training neural networks with low precision weights and activations. *The Journal of Machine Learning Research*, 18(1):6869–6898, 2017. [2](#)
- [22] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016. [2](#)
- [23] Minsoo Kang and Bohyung Han. Operation-aware soft channel pruning using differentiable masks. In *International Conference on Machine Learning*, pages 5122–5131. PMLR, 2020. [1](#)
- [24] Alex Krizhevsky, Vinod Nair, and Geoffrey Hinton. Cifar-10 (canadian institute for advanced research). [7](#)
- [25] Zi Lin, Jeremiah Z. Liu, Zi Yang, Nan Hua, and Dan Roth. Pruning redundant mappings in transformer models via spectral-normalized identity prior. In *EMNLP (Findings)*, pages 719–730, 2020. [4](#)
- [26] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2018. [3](#)
- [27] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient

- convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, pages 2736–2744, 2017. 1, 2, 4
- [28] 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. *arXiv preprint arXiv:2103.14030*, 2021. 1, 2, 5, 7
- [29] Zechun Liu, Zhiqiang Shen, Marios Savvides, and Kwang-Ting Cheng. Reactnet: Towards precise binary neural network with generalized activation functions. In *European Conference on Computer Vision*, pages 143–159. Springer, 2020. 2
- [30] Zechun Liu, Xiangyu Zhang, Zhiqiang Shen, Yichen Wei, Kwang-Ting Cheng, and Jian Sun. Joint multi-dimension pruning via numerical gradient update. *IEEE Transactions on Image Processing*, 30:8034–8045, 2021. 2
- [31] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 116–131, 2018. 2
- [32] Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. When does label smoothing help? *Advances in Neural Information Processing Systems*, 32:4694–4703, 2019. 2
- [33] Hai Phan, Zechun Liu, Dang Huynh, Marios Savvides, Kwang-Ting Cheng, and Zhiqiang Shen. Binarizing mobilenet via evolution-based searching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 13420–13429, 2020. 2
- [34] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 1, 2, 3, 7
- [35] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 4780–4789, 2019. 2
- [36] Adriana Romero, Nicolas Ballas, Samira Ebrahimi Kahou, Antoine Chassang, Carlo Gatta, and Yoshua Bengio. Fitnets: Hints for thin deep nets. *arXiv preprint arXiv:1412.6550*, 2014. 2
- [37] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4510–4520, 2018. 2
- [38] Zhiqiang Shen, Zechun Liu, and Eric Xing. Sliced recursive transformer. *arXiv preprint arXiv:2111.05297*, 2021. 1
- [39] Zhiqiang Shen, Zechun Liu, DeJia Xu, Zitian Chen, Kwang-Ting Cheng, and Marios Savvides. Is label smoothing truly incompatible with knowledge distillation: An empirical study. In *International Conference on Learning Representations*, 2021. 2
- [40] Zhiqiang Shen and Marios Savvides. Meal v2: Boosting vanilla resnet-50 to 80%+ top-1 accuracy on imagenet without tricks. *arXiv preprint arXiv:2009.08453*, 2020. 2
- [41] Zhiqiang Shen and Eric Xing. A fast knowledge distillation framework for visual recognition. *arXiv preprint arXiv:2112.01528*, 2021. 2
- [42] Xiu Su, Shan You, Jiyang Xie, Mingkai Zheng, Fei Wang, Chen Qian, Changshui Zhang, Xiaogang Wang, and Chang Xu. Vision transformer architecture search. *arXiv preprint arXiv:2106.13700*, 2021. 1, 2, 3, 4, 7
- [43] Zhiqing Sun, Shengcao Cao, Yiming Yang, and Kris M Kitani. Rethinking transformer-based set prediction for object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3611–3620, 2021. 1
- [44] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019. 2, 3
- [45] Yehui Tang, Kai Han, Yunhe Wang, Chang Xu, Jianyuan Guo, Chao Xu, and Dacheng Tao. Patch slimming for efficient vision transformers. *arXiv preprint arXiv:2106.02852*, 2021. 2, 3, 4, 7
- [46] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pages 10347–10357. PMLR, 2021. 1, 5
- [47] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, volume 139, pages 10347–10357, July 2021. 2, 6, 7
- [48] Grant Van Horn, Oisín Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8769–8778, 2018. 7
- [49] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 1
- [50] Yuqing Wang, Zhaoliang Xu, Xinlong Wang, Chunhua Shen, Baoshan Cheng, Hao Shen, and Huaxia Xia. End-to-end video instance segmentation with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8741–8750, 2021. 1
- [51] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10734–10742, 2019. 3
- [52] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, and Kurt Keutzer. Shift: A zero flop, zero parameter alternative to spatial convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 9127–9135, 2018. 2
- [53] Bichen Wu, Chenfeng Xu, Xiaoliang Dai, Alvin Wan, Peizhao Zhang, Zhicheng Yan, Masayoshi Tomizuka, Joseph

- Gonzalez, Kurt Keutzer, and Peter Vajda. Visual transformers: Token-based image representation and processing for computer vision. *arXiv preprint arXiv:2006.03677*, 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. *arXiv preprint arXiv:2105.15203*, 2021. 1
- [55] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. Nektadap: Platform-aware neural network adaptation for mobile applications. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 285–300, 2018. 2
- [56] Jianbo Ye, Xin Lu, Zhe Lin, and James Z Wang. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *arXiv preprint arXiv:1802.00124*, 2018. 2
- [57] Jiahui Yu, Pengchong Jin, Hanxiao Liu, Gabriel Bender, Pieter-Jan Kindermans, Mingxing Tan, Thomas Huang, Xiaodan Song, Ruoming Pang, and Quoc Le. Bignas: Scaling up neural architecture search with big single-stage models. In *European Conference on Computer Vision*, pages 702–717. Springer, 2020. 3
- [58] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019. 5
- [59] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 5
- [60] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018. 2
- [61] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6881–6890, 2021. 1
- [62] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13001–13008, 2020. 5
- [63] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016. 2
- [64] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable detr: Deformable transformers for end-to-end object detection. In *International Conference on Learning Representations*, 2020. 1
- [65] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578*, 2016. 2
- [66] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018. 2