

Training-free Transformer Architecture Search

Qinqin Zhou^{1*} Kekai Sheng^{2*} Xiawu Zheng³ Ke Li²
Xing Sun² Yonghong Tian⁴ Jie Chen^{4,3} Rongrong Ji^{1,3,5,6†}

¹Media Analytics and Computing Lab, School of Informatics, Xiamen University, ²Tencent Youtu Lab,
³Peng Cheng Laboratory, ⁴School of Electronic and Computer Engineering, Peking University,
⁵Institute of Artificial Intelligence, Xiamen University, ⁶Fujian Engineering Research
Center of Trusted Artificial Intelligence Analysis and Application, Xiamen University

Abstract

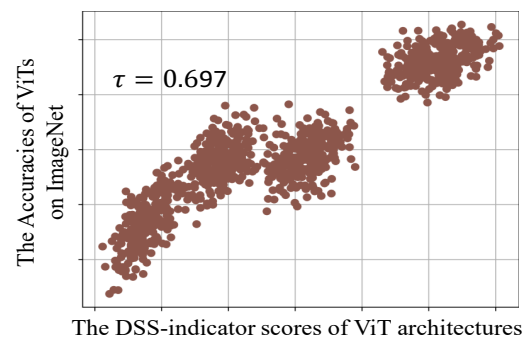
Recently, Vision Transformer (ViT) has achieved remarkable success in several computer vision tasks. The progresses are highly relevant to the architecture design, then it is worthwhile to propose Transformer Architecture Search (TAS) to search for better ViTs automatically. However, current TAS methods are time-consuming and existing zero-cost proxies in CNN do not generalize well to the ViT search space according to our experimental observations. In this paper, for the first time, we investigate how to conduct TAS in a training-free manner and devise an effective training-free TAS (TF-TAS) scheme. Firstly, we observe that the properties of multi-head self-attention (MSA) and multi-layer perceptron (MLP) in ViTs are quite different and that the synaptic diversity of MSA affects the performance notably. Secondly, based on the observation, we devise a modular strategy in TF-TAS that evaluates and ranks ViT architectures from two theoretical perspectives: synaptic diversity and synaptic saliency, termed as DSS-indicator. With DSS-indicator, evaluation results are strongly correlated with the test accuracies of ViT models. Experimental results demonstrate that our TF-TAS achieves a competitive performance against the state-of-the-art manually or automatically design ViT architectures, and it promotes the searching efficiency in ViT search space greatly: from about 24 GPU days to less than 0.5 GPU days. Moreover, the proposed DSS-indicator outperforms the existing cutting-edge zero-cost approaches (e.g., TE-score and NASWOT).

1. Introduction

Vision Transformer (ViT) [19,38,43] has shown its competitiveness in the computer vision community recently, and

*The first two authors contributed equally. This work was done when Qinqin Zhou was intern at Tencent Youtu Lab.

†Corresponding author: rrji@xmu.edu.cn.



Search Space	Method	#Param (M)	Top-1 Acc (%)	Cost Time (GPU Days)
AutoFormer	Supernet Pre-Train + Evolution Search	5.7	74.7	24
	TF-TAS (Ours)	5.9	75.3	0.5
PiT	Random Search	4.9	69.7	0.5
	TF-TAS (Ours)	4.6	73.7	0.5

Figure 1. **Top:** The evaluation results from DSS-indicator are positively correlated with the test accuracies of various ViT networks. **Bottom:** The design and computation cost of existing methods, i.e., AutoFormer [4] and PiT [17], and the proposed TF-TAS.

has been an important research hot-spot. With the emergence of manually-designed advanced ViT models [11, 37, 46], Transformer Architecture Search (TAS) [3, 4, 34, 35, 41] makes its grand debut and aims at searching multiple configurations of ViT architecture in an automated way. Although the one-shot NAS scheme [9, 13, 47, 48] is leveraged in TAS, it still demands a high computational cost (e.g., larger than 24 GPU days) to train a supernet for reliable performance estimations on various ViT architectures. Furthermore, since the magnitude of ViT search spaces (e.g., $\sim 10^{30}$ in GLiT [3]) far exceeds the magnitude of CNN search spaces (e.g., $\sim 10^{18}$ in DARTS [26]) and ViT models usually require more training epochs (e.g., 300), the search efficiency of one-shot based TAS is still unsatisfying.

Recall that, to promote the searching efficiency on

CNN search spaces, several proxies (*e.g.*, GraSP [40], TE-score [5], and NASWOT [29]) are proposed to evaluate the rank of different CNN architectures in a zero-cost manner. Technically, one typical CNN is mainly composed of convolution layers. On the other hand, the basic blocks of one ViT model, multi-head self-attention (MSA) and multi-layer perceptron (MLP), are mainly composed of linear layers. The difference makes it risky to apply existing zero-cost proxies that are verified on CNN directly to the ViT search space. Consequently, it is necessary and worthwhile to investigate the possibility of an effective zero-cost proxy that is more suitable for ranking ViT networks and facilitates the training efficiency of TAS. This problem motivates us to delve further into the ViT architecture to propose an effective approach to conduct TAS in a training-free manner.

To this end, we conduct a modular investigation on MSA and MLP in typical ViTs [4, 17] to devise effective performance indicator methods for MSA and MLP in ranking ViT networks. Based on numerical results, we observe that the MSA and MLP in ViT have different properties in indicating the performance of the model. When the MSA possesses higher *diversity* score or when the MLP has more *synaptic saliency* value, the corresponding ViT network always has better performance (see the *Top* of Fig. 1). Motivated by the significant insights, we propose an effective and efficient DSS-indicator and design a training-free TAS (TF-TAS). Specifically, we attempt to comprehensively rank various ViTs by exploiting the properties of the MSA and MLP mentioned above. DSS-indicator estimates the synaptic diversity of MSA and the synaptic saliency of MLP to generate effective evaluation scores for ViT architectures. The synaptic diversity measures the degree of rank collapse on one MSA, and the synaptic saliency estimates the amount of important parameter within one MLP. To the best of our knowledge, it is the first time to propose synaptic diversity of MSA and synaptic saliency of MLP as the proxy in evaluating ViT architectures. Besides, it should be noted that our TF-TAS is orthogonal to the search space design and weight sharing strategy. Thus, it is flexible to combine TF-TAS with other ViT search space or TAS methods to further promote the searching efficiency. Compared with the manually designed ViTs [6, 14, 27, 43] and the automatically searched ones [3, 4, 35], our TF-TAS achieves a competitive performance and accelerates the searching procedure from around 24 GPU days to less than 0.5 GPU days, about 48 times faster (see the *Bottom* of Fig. 1).

For fair comparison and full investigation, we also present a reliable test-bed to evaluate the *state-of-the-art* zero-cost proxies (*e.g.*, TE-score [5] and NASWOT [29]) in the ViT search space. We construct a large proxy ViT benchmark, which is based on several pre-trained supernets of AutoFormer [4], to compare the relative performance of alternative zero-cost proxies on ViT architectures. With the

numerical observations of zero-cost proxies, we empirically verify the relative rankings of different zero-cost proxies in TAS, and our DSS-indicator outperforms the other counterparts. We also draw some practical insights in designing a better proxy for ranking ViT architectures.

Overall, our principal contributions are listed as follows:

- We propose a training-free TAS (TF-TAS), which includes a modular strategy to combine the synaptic diversity of the MSAs and synaptic saliency of the MLPs as a DSS-indicator in evaluating ViT architectures.
- Extensive experiments demonstrate that the proposed TF-TAS not only achieves a competitive search performance, but also improves the search efficiency in searching ViT architectures.
- We design a series of controlled experiments to compare the existing zero-cost proxies in TAS. The results provide some empirical insights in designing optimal proxy metric in evaluating ViT architectures.

2. Preliminary

Design & Search Transformer Architectures. Since ViT [38], the computer vision community has witnessed the emergence of many manually designed advanced ViT architectures [11, 17, 37, 44]. Technically, most of them consist of the same basic blocks that include MSA, Layer Normalization (LN), and MLP. Existing TAS approaches [3, 4, 35] search different dimensions in MSAs and MLPs, such as the number of heads in MSAs, the ratio of MSAs or MLPs. These methods are generally built on the one-shot NAS framework [9, 13, 47, 48]: train the supernet by training a subnet path in each epoch. The typical one-shot based TAS method AutoFormer [4] includes an entanglement strategy and divides the search space into three sub-supernets, each of which are trained in a one-shot manner for 500 epochs (about 24 GPU days) and requires 8 NVIDIA V100 GPUs. The size of AutoFormer search space is 1.7×10^{16} , which makes it time-consuming in training supernet.

In general, how to reduce the cost of searching the ViT architectures and ensuring the performance of the searched networks is a fundamental and challenging problem. In this paper, we try to find a way to maintain the performance of TAS and accelerate the searching efficiency.

Performance Evaluation via Zero-cost Proxy. There are two mainstreams of zero-cost proxy to reduce the cost of performance estimation and promote searching efficiency. The first one, inspired by the pruning community, sums up the saliency value of each model weight as the proxy of the corresponding CNN architecture with a single forward / backward propagation. The popular methods include Grad-norm [1], SNIP [23], and GraSP [40]. These proxies

follow a default assumption: the more salient the weight value is, the more important it is to the model; and the more salient weight one network has, the better performance the model does. The second one, such as TE-score [5], NAS-WOT [29], and Zen-Score [24], is designed specifically for CNNs. They analyze the important properties (e.g., expressivity) of the representations of CNNs. Mellor *et al.* [29] proposed jacobian covariance to sum up the saliency of each weight to rank CNNs. Chen *et al.* [5] applied two theory-inspired indicators as the proxy to find the best subnet.

Different from the existing literature, in this paper, for the first time, we identify the shortcoming of directly applying existing proxies in the ViT search space. Then, we propose a simple yet effective proxy to generate better performance and facilitate the searching efficiency of TAS.

3. Methodology

3.1. Motivation

The existing TAS methods [3, 4, 35] are relatively time-consuming, especially in performance estimation (e.g., 300 training epochs on 8 GPUs). Then, it is worthwhile to leverage the zero-cost proxies [5, 29, 40] to rank ViT architectures and reduce the computation cost in performance estimation. Nevertheless, the existing zero-cost proxies are specifically designed for the CNN search spaces (e.g., DARTS [26] and NAS-Bench 201 [8]). Obviously, the search space of ViT is quite different from that of CNN, then the existing proxies could not promise generalization on the ViT search space (see the results in Sec. 4.6). It motivates us to explore and exploit the useful properties of MSA and MLP within ViT, and design an effective ViT-oriented zero-cost proxy.

In this section, we first propose an effective method to calculate the synaptic diversity of MSA and generate evaluation results that are positively correlated with the classification accuracies of ViT networks. Then, we find that when the MLP module has more important weight parameters, *i.e.*, higher synaptic saliency value, the corresponding ViT network always yields better classification performance. Finally, we propose a DSS-indicator to comprehensively exploit the properties of MSAs and MLPs in ranking various ViT architectures effectively and efficiently, which promotes the searching efficiency of TAS.

3.2. Synaptic Diversity in MSA

Theoretical Analysis. MSA is a basic fundamental component of ViT architectures. Several works unveil one important property of MSA: its *diversity* [10, 49]. In particular, Dong *et al.* [10] pointed out that the MSA causes *rank collapse* in the learned representations. In specific, as the input propagates forward in the network and the depth continues to deepen, the outputs of MSAs in ViTs gradually converge to rank-1. And eventually, the output degenerates

into a matrix with a rank of 1, the value of each row becomes the same, *i.e.* the scarcity of diversity. Such rank collapse severely degenerates the performance of ViT.

However, estimating the rank collapse in the high dimension representation space requires a huge computation cost. Actually, Fazel *et al.* [12] demonstrates that the rank of a matrix contains representative cues of the diversity information within the features. Building on these understandings, the rank of the weight parameters in the MSA module could be adopted as an indicator to evaluate the ViT architecture.

Synaptic Diversity. For the MSA module, it is still computationally complex to directly measure the rank of its weight matrix and hinders practical applications. To accelerate the calculation of synaptic diversity in MSA module, we leverage the Nuclear-norm of the MSA’s weight matrix to approximate its rank as the diversity indicator. Theoretically, the Nuclear-norm of a weight matrix can be treated as an equivalent substitution for its rank, when the Frobenius-norm of the weight matrix meets certain conditions. Specifically, we denote the weight parameter matrix of an MSA module as W_m . m indicates the m -th linear layer in an MSA module. The Frobenius-norm of W_m is defined as:

$$\|W_m\|_F = \sqrt{\sum_{i=1}^U \sum_{j=1}^V |w_{i,j}|^2}, \quad (1)$$

where U, V are the dimension of W_m , and $w_{i,j}$ denotes the the element in the i -th row and j -th column of W_m . According to inequality of arithmetic and geometric means, the upper-bound of W_m is calculated as:

$$\begin{aligned} \|W_m\|_F &\leq \sqrt{\sum_{i=1}^U (\sum_{j=1}^V w_{i,j}) \cdot (\sum_{j=1}^V w_{i,j})} \\ &= \sqrt{\sum_{i=1}^U 1 \cdot 1} = \sqrt{U}. \end{aligned} \quad (2)$$

This means the upper-bound of $\|W_m\|_F$ is the largest number of linear independent vectors of W_m , *i.e.*, the matrix rank. Given two randomly selected vectors w_m^i and w_m^j in W_m , $\|W_m\|_F$ could be larger when w_m^i and w_m^j are independent. This indicates that: the larger the Frobenius-norm of W_m goes, the closer the rank of W_m is to the diversity of W_m . And according to one *Theorem* proved by Fazel *et al.* [12] that when $\|W_m\|_F \leq 1$, the Nuclear-norm of W_m could be an approximation of the rank of W_m . Formally, the Nuclear-norm of W_m is defined as:

$$\|W_m\|_{nuc} = tr(\sqrt{W_m^T W_m}), \quad (3)$$

where $tr(*)$ denotes the trace of the corresponding matrix. In our case, we always have $\|W_m\|_F \leq \sqrt{U}$, thus the rank

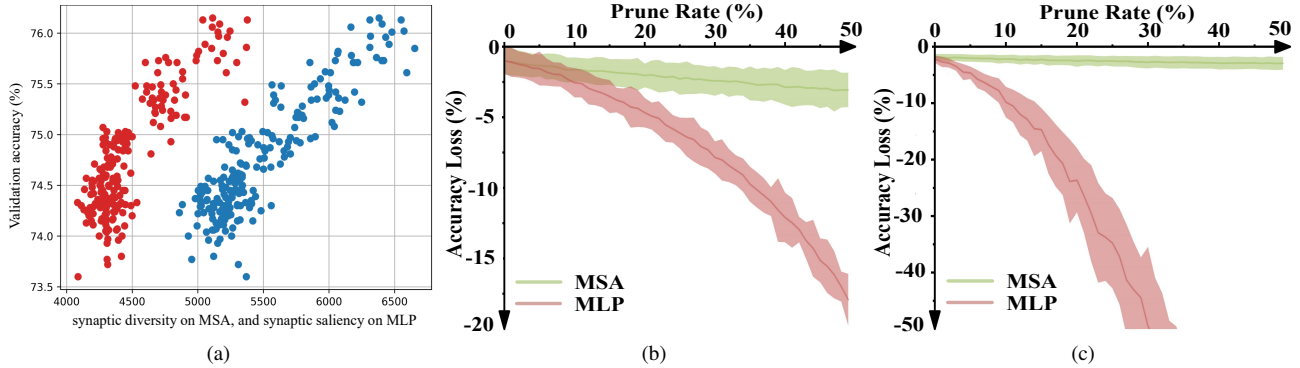


Figure 2. (a) Illustrations of the positive connection between the proxy score (D_{MSA}^l red points & S_{MLP}^l blue points) and the accuracy of the ViT model. (b) & (c) The sensitivity analysis of MSA and MLP to pruning on a flat ViT [4] and a deep-narrow ViT [17], respectively.

of W_m could be approximated as $\|W_m\|_{nuc}/\sqrt{U}$. Theoretically, $\|W_m\|_{nuc}$ is proportional to $\|W_m\|_{nuc}/\sqrt{U}$, which indicate that the Nuclear-norm of W_m could be used to measure the diversity of W_m . To better estimate the synaptic diversity of MSA modules from one ViT network that the weights are randomly initialized, we further consider the aforementioned procedure on the gradient matrix $\partial\mathcal{L}/\partial W_m$ (\mathcal{L} is the loss function) of each MSA module.

Overall, we define the synaptic diversity of the weight parameter in the l -th MSA module as follows:

$$D_{MSA}^l = \sum_m \left\| \frac{\partial\mathcal{L}}{\partial W_m} \right\|_{nuc} \odot \|W_m\|_{nuc}. \quad (4)$$

To verify the positive correlation between the synaptic diversity of MSA and the test accuracy of the given ViT architectures, we re-train 100 ViT networks sampled from AutoFormer [4] and obtain their corresponding classification performance and synaptic diversity of MSA modules. The Kendall's τ between them is 0.65 as shown in Fig. 2a. The results in Sec. 4 will also demonstrate the positive connection between the evaluation score of Eq.(4) and the performance for each input ViT architecture.

3.3. Synaptic Saliency in MLP

Theoretical Analysis. Network pruning [30, 33] has achieved many progresses for CNNs, and start showing its power on Transformer [31, 39, 42]. Several effective CNN pruning methods [22, 36, 45] have been proposed to measure the importance of the model weights at the early training stage. Tanaka *et al.* [36] measured the synaptic saliency in pruning CNNs without training. Wang *et al.* [42] find that different modules in Transformers show different degrees of redundancy even at the initialization stage, and tried to prune the different dimensions of Transformers. Similar to pruning, TAS focuses on searching several important dimensions, which include the number of attention heads, MSA and MLP ratio, *etc.* Inspired by these develop-

ments, we try to use the synaptic saliency to evaluate different ViTs. However, it has been validated [31, 39, 42] that the sensitivities of MSA and MLP to pruning are different. A large percentage of weight in the MSA is redundant [31, 39], which has little impact on the performance at test time. It implies that synaptic saliency might demonstrate different performances in MSA and MLP.

To further validate the findings, we through a pruning sensitive experiment to show some quantitative results. As illustrated in Fig. 2b, we randomly sample 5 ViT architectures from AutoFormer search space to analyze the sensitivity of the MSA and MLP to pruning. We observe that the MLP is much more sensitive to pruning than the MSA. We also conduct the analysis on the deep-narrow ViT networks (*e.g.*, PiT [17]), and obtain similar observations (see Fig. 2c). Moreover, we adopt synaptic saliency on MSA and MLP modules as proxies to calculate the Kendall's τ on the proxy ViT benchmark, respectively. The Kendall's τ of synaptic saliency on the MLP is 0.47, which is better than on the MSA (0.24) and both of the MLP and MSA (0.41).

Since the synaptic saliency is generally calculated in the form of summation, the redundancy brings a cumulative effect. Specifically, the MSA module proves to be insensitive to pruning, which means the weight parameters of the MSAs have a higher redundancy. It has been proposed in the pruning community [1] that the values of the redundant weight parameters are much smaller than the non-redundant ones. Although the values of these redundant parameters are relatively small, redundancies of more than 50 percent tend to have a large cumulative effect, especially when distinguishing between similar architectures. For the cumulative effect, the redundant weight parameters of MSA are considered in the zero-cost proxies to measure the saliency, the cumulative forms in the zero-cost proxies lead to the cumulative effect in the MSA. The cumulative effect might make the zero-cost proxies give a higher rank to the network. Meanwhile, the synaptic saliency of the MLP modules is less affected by weight redundancy, which could be

adopted as an indicator for the MLP modules.

Synaptic Saliency. To evaluate the MLP in ViT, we resort to the *synaptic saliency*. It is extensively studied in network pruning to indicate the importance of model weights. There are several pruning-based zero-cost proxies [1,23,40] that could be directly used to measure the synaptic saliency of CNNs for the CNNs are mainly composed of convolution layers. On the other hand, ViT architectures are mainly composed of MLP and MSA modules, which have different pruning properties. Through the pruning sensitivity analysis of the MSA and MLP modules in Sec. 3.3, we validate that the MLP module is much more sensitive to pruning. Therefore, the differences in the importance of weights in MLP modules can be better reflected by the synaptic saliency. As a comparison, the MSA modules are relatively insensitive to the pruning, the synaptic saliency of which is often influenced by the redundant weights.

Building on the pruning sensitivity of MLP, we propose to measure the synaptic saliency in a modular manner. Specifically, the proposed modular strategy measures the synaptic saliency of MLPs as a part of the indicator for a ViT architecture. Formally, given a ViT architecture, the saliency score of the l -th MLP module is:

$$S_{MLP}^l = \sum_n \frac{\partial \mathcal{L}}{\partial W_n} \odot W_n. \quad (5)$$

where n denotes the number of linear layer in the l -th MLP in a specified ViT network, which usually set to 2.

Fig. 2a shows some qualitative results to verify the effectiveness of S_{MLP} in evaluating ViT architectures.

3.4. Training-free TAS

Based on the above analysis, we propose a training-free TAS (TF-TAS) with a modular strategy to further improve search efficiency. The modular strategy is proposed to consider the training-free evaluation of ViT architectures in two parts to form a DSS-indicator.

Combine the synaptic diversity of MSA and the saliency score of MLP, we formulate the DSS-indicator as follows:

$$S_{DSS}(\mathcal{A}) = \sum_l D_{MSA}^l + \sum_k S_{MLP}^k. \quad (6)$$

Overall, the DSS-indicator evaluates each ViT architecture from two different perspectives. TF-TAS calculates S_{DSS} after a forward and backward as the indicator of a specified ViT architecture. We keep each pixel of the input data being 1 to eliminate the affection of input data. Thus, S_{DSS} is invariant to random seed. Moreover, the loss at first iteration is expressed as: $\mathcal{L} = \mathbb{1}^T \left(\prod_l^L |\omega^{[l]}| \right) \mathbb{1}$, where $\mathbb{1}$ is the all ones vector. It makes Eq.(6) to take the inter-layer interactions of weight parameters into account to measure the diversity of MSAs and the saliency of MLPs.

Given a specified parameter constraint, we first randomly sample 8,000 subnets on one ViT search space. Then, the synaptic diversity score of MSAs and the saliency score of MLPs are calculated as the evaluation rank of each subnet. Based on the calculated DSS-indicator scores of each ViT architecture, we pick the networks with the highest proxy value as the optimal one. Finally, we retrain the searched optimal network to obtain its final test accuracy.

4. Experiments

4.1. Implementation Details

TF-TAS includes a search stage and a re-train stage. In the search stage, the number of sub-networks randomly sampled from the given ViT search space is set to 8,000, whose weights are randomly initialized. Their proxy scores are calculated for each sub-network. To compute the proposed DSS-indicator, the input is constructed with each pixel being 1. After the DSS-indicator of each sub-network is computed, we retrain the top-1 sub-network. In the re-train stage, we follow the training configuration in AutoFormer [4] to train the obtained optimal ViT networks: AdamW optimizer [28] with weight decay 0.05, initial learning rate 1×10^{-3} and minimal learning rate 1×10^{-5} with cosine scheduler, 5 epochs warmup, batch size of 256, and the models are trained with 300 epochs, *etc.* All experiments are implemented on NVIDIA Tesla V100 GPUs and the results are estimated on ImageNet [7], CIFAR-10/CIFAR-100 [21], and the COCO 2017 dataset [25]. The image resolution is 224×224 by default. We also use MindSpore to validate the generalizability of our method. Please refer to the supplementary for the pseudocode of DSS-indicator.

4.2. Results on AutoFormer Search Space.

We first evaluate TF-TAS on the search space of AutoFormer, *i.e.*, AutoFormer search space \mathcal{S}_A . We compare the performance of searched optimal ViTs with that of *state-of-the-art* TAS methods [3,4,35] and manually designed CNNs and ViTs [11, 14, 18, 27, 37] on ImageNet.

As listed in Tab. 1, the searched optimal architectures TF-TAS (*i.e.*, TF-TAS-Ti, TF-TAS-S, and TF-TAS-B) outperforms manually designed CNNs (*e.g.*, ResNet [15], MobileNet [18], PVT [43], and T2T-ViT [46]) with a clear margin in all the three common model sizes (*i.e.*, tiny, small, and base). Compared with other manually designed ViT architectures [6, 11, 14, 27, 37], ours TF-TAS achieves competitive results. Specifically, the searched TF-TAS achieves a top-1 accuracy of 75.3%, which surpasses DeiT-tiny by 3.1 percent. Compared with other TAS approaches [3, 4, 35] that require larger than 24 GPU days to seek for optimal ViT architectures, the proposed DSS-indicator helps us achieves the comparable results with much fewer GPU

Table 1. The comparison results on the Autoformer search space. * denotes the results reported by [17].

Models	#Param (M)	FLOPS (B)	Top-1 (%)	Top-5 (%)	Model Type	Design Type	GPU Days
ResNet-18* [15]	11.7	1.8	72.5	-	CNN	Manual	-
MobileNet-V3 [18]	5.5	-	75.2	-	CNN	Manual	-
DeiT-Ti [37]	5.7	1.2	72.2	91.1	Transformer	Manual	-
TNT-Ti [14]	6.1	1.4	73.9	91.9	Transformer	Manual	-
ViT-Ti [11]	5.7	-	74.5	-	Transformer	Manual	-
CPVT-Ti [6]	6.0	-	74.9	92.6	Transformer	Manual	-
PVT-Tiny [43]	13.2	1.9	75.1	-	Transformer	Manual	-
ViTAS-C [35]	5.6	1.3	74.7	91.6	Transformer	Auto	32
AutoFormer-Ti [4]	5.7	1.3	74.7	92.6	Transformer	Auto	24
GLiT-Ti [3]	7.2	1.4	76.3	-	Hybrid	Auto	N/A
TF-TAS-Ti (Ours)	5.9	1.4	75.3	92.8	Transformer	Auto	0.5
ResNet-50* [15]	25.6	4.1	80.2	-	CNN	Manual	-
RegNetY-4GF [16]	20.6	-	79.4	-	CNN	Manual	-
DeiT-S [37]	22.1	4.7	79.9	95.0	Transformer	Manual	-
ViT-S/16 [11]	22.1	4.7	78.8	-	Transformer	Manual	-
PVT-Small [43]	24.5	3.8	79.8	-	Transformer	Manual	-
Swin-T [27]	29.0	4.5	81.3	-	Transformer	Manual	-
TNT-S [14]	23.8	5.2	81.5	95.7	Transformer	Manual	-
CPVT-S [6]	23.0	-	81.5	95.7	Transformer	Manual	-
T2T-ViT_t-14 [46]	21.5	-	81.7	-	Transformer	Manual	-
ViTAS-F [35]	27.6	6.0	80.5	95.1	Transformer	Auto	32
AutoFormer-S [4]	22.9	5.1	81.7	95.7	Transformer	Auto	24
GLiT-S [3]	24.6	4.4	80.5	-	Hybrid	Auto	N/A
TF-TAS-S (Ours)	22.8	5.0	81.9	95.8	Transformer	Auto	0.5
ResNet-152* [15]	60.2	11.5	81.9	-	CNN	Manual	-
RegNetY-16GF [32]	83.6	15.9	80.4	-	CNN	Manual	-
PVT-Large [43]	61.0	9.8	81.7	-	Transformer	Manual	-
DeiT-B [37]	86.0	18.0	81.8	95.6	Transformer	Manual	-
CPVT-B [6]	88.0	-	82.3	-	Transformer	Manual	-
TNT-B [14]	65.5	14.1	82.9	96.3	Transformer	Manual	-
Swin-B [27]	88.0	15.4	83.5	-	Transformer	Manual	-
T2T-ViT-24 [46]	64.1	-	82.6	-	Transformer	Manual	-
GLiT-B [3]	96.0	17.0	82.3	-	Hybrid	Auto	N/A
AutoFormer-B [4]	54.0	11.0	82.4	95.7	Transformer	Auto	24
TF-TAS-B (Ours)	54.0	12.0	82.2	95.6	Transformer	Auto	0.5

Days. Furthermore, our DSS-indicator comprehensively considers performance and efficiency in searching ViT architectures. Based on the estimation results of the proposed DSS-indicator for each input ViT architecture, we reduce lots of computation budgets in performance estimation and obtain optimal ViT networks with comparable performance within 0.5 GPU days. For the analysis of the searched optimal architectures, please refer to the supplementary.

4.3. Results on PiT Search Space.

To further investigate the versatility of our DSS-indicator, we build another search space: PiT search space \mathcal{S}_P . Without loss of generality, we propose \mathcal{S}_P on PiT [17] and include several important dimensions of ViT (e.g., depth, head number of MSA, MLP ratio), together with depth-wise convolution operations. For the detailed infor-

mation about \mathcal{S}_P , please refer to the supplementary.

As listed in Tab. 2, within the budget of 0.5 GPU days, the proposed DSS-indicator is still able to obtain the optimal ViT architectures that have the comparable or even better Top-1 classification accuracy to PiT-Ti and PiT-S. The searched networks outperform the randomly selected ones, PiT-Ti_{rand} and PiT-S_{rand}, by about 2.9 ~ 5%. We further conduct the transfer experiment on detection with COCO 2017 dataset [25] using a simplified setting [17]. As shown in Tab. 2, TF-TAS achieves competitive performance in detection task. These results help us ensure the generalization of the proposed DSS-indicator across different ViT search spaces. We also notice that the searched results of TF-TAS on PiT search space are lower than that on AutoFormer search space in Tab. 1. This observation implies that the search space is also an important part of TAS.

Table 2. The comparison results on the PiT search space. † indicates the results we reproduce.

Models	#Param (M)	FLOPs (B)	Top-1 (%)	Top-5 (%)
PiT-Ti† [17]	4.9	0.7	73.8	91.7
PiT-Ti _{rand}	4.9	0.7	69.7	89.1
TF-TAS-Ti (Ours)	4.6	0.6	73.7	91.7
PiT-XS† [17]	10.6	1.4	78.2	94.0
PiT-XS _{rand}	10.5	1.8	74.8	92.2
TF-TAS-XS (Ours)	10.0	1.8	77.7	93.8
PiT-S† [17]	23.5	2.9	79.9	94.4
PiT-S _{rand}	24.2	3.3	75.1	92.4
TF-TAS-S (Ours)	23.8	3.2	80.5	94.9
Backbone	#Param (M)	Avg. Precision at IOU		
		AP	AP ₅₀	AP ₇₅
ResNet-50 [15]	41.0	41.5	60.5	44.3
ViT-S [11]	34.9	36.9	57.0	38.0
Ours	36.0	39.7	60.9	40.4

Table 3. Evaluation results of the proposed zero-cost proxy on the state-of-the-art ViT architectures.

Models	#Param (M)	Top-1 (%)	Proxy
PiT [17]			
PiT-Ti	4.9	73.8	2.9×10^4
PiT-XS	10.6	78.2	3.4×10^4
PiT-S	23.5	80.5	4.3×10^4
T2T-ViT [46]			
T2T-ViT-7	4.3	71.7	1.1×10^5
T2T-ViT-10	5.9	75.2	1.3×10^5
T2T-ViT-12	6.9	76.5	1.6×10^5
T2T-ViT-14	21.5	81.5	5.8×10^5
T2T-ViT-19	39.2	81.9	1.1×10^6
T2T-ViT-24	64.1	82.3	2.0×10^6
XCiT [2]			
XCiT-tiny-24-p16	12.0	79.4	3.4×10^4
XCiT-small-24-p16	48.0	82.6	9.2×10^4
XCiT-medium-24-p16	84.0	82.7	1.4×10^5
XCiT-large-24-p16	189.0	82.9	2.4×10^5

4.4. On Evaluating Popular Architectures.

To further investigate the effectiveness and versatility of the proposed proxy, we also conduct evaluation experiments on other popular *state-of-the-art* ViT architectures [2, 17, 46]. As shown in Tab. 3, our DSS-indicator can evaluate the right rank of ViT architectures in their corresponding search space. Interestingly, the values of the proposed proxy obtained across different search spaces are not comparable. This might be caused by several factors. For example, the different ways of model initialization, and the search space itself contains several different modules which make it difficult to achieve a fair comparison.

4.5. Transfer Learning Results.

To test the transferability of the searched optimal ViT networks, we conduct some transfer learning experiments.

Table 4. The results (%) on downstream classification datasets. † 384 means that the model is fine-tuned with 384×384 resolution.

Models	#Param	ImageNet	C-10	C-100
ViT-B/16 [11]	86M	77.9	98.1	87.1
DeiT-B [37]† 384	86M	83.1	99.1	90.8
AutoFormer-S [4]† 384	23M	83.4	99.1	91.1
TF-TAS-S† 384 (Ours)	23M	83.5	99.1	91.2

Table 5. The Kendall τ values between various evaluation metrics and the final classification accuracy on the inherit networks randomly sampled from three pre-trained AutoFormer supernets.

Proxy	#Param (M)		
	5 - 7	15 - 19	23 - 25
SNIP [23]	0.481	0.028	-0.282
GraSP [40]	0.053	-0.022	-0.029
TE-score [5]	-0.039	-0.248	-0.075
NASWOT [29]	0.378	0.171	0.208
DSS-indicator (Ours)	0.697	0.615	0.306

We follow the same settings as DeiT [37] and finetune the TF-TAS-S (see Tab. 1) in CIFAR-10 (C-10) and CIFAR-100 (C-100) [21]. The results are listed in Tab. 4. As we observe that the optimal ViT architectures found by DSS-indicator in a training-free manner have a similar fine-tuning performance as that of the networks searched by AutoFormer [4].

4.6. Comparison of Zero-Cost Proxies.

For full investigation, we compare our DSS-indicator with alternative *state-of-the-art* zero-cost proxies [5, 23, 29, 40] on CNN search spaces. To build a reliable test-bed to evaluate these zero-cost proxies, we need a ViT benchmark and we resort to AutoFormer. We call the search space of AutoFormer [4] as \mathcal{S}_A for simplicity. Empirically, Chen *et al.* [4] find that, the subnet from \mathcal{S}_A with its weights inherited from the pre-trained supernet can achieve the performance comparable to that of the retrained one. Building on this observation, we sample 3,000 subsets from \mathcal{S}_A and obtain their accuracies when they inherit their weights from the pre-trained supernet. Without loss of generality, we sample the subnet with the amount of parameter in three common ranges: 5M \sim 7M, 15M \sim 19M, and 23M \sim 25M. With this ViT benchmark, we compare our DSS-indicator with four cutting-edge zero-cost proxy methods: SNIP [23], GraSP [40], NASWOT [29] and TE-score [5].

The results of Kendall τ [20] are illustrated in Tab. 5 and Fig. 3. Overall, the relative ranking of the proxies is: Ours > NASWOT > SNIP > GraSP > TE-score. Our DSS-indicator outperforms the others in ranking various ViT architectures. The results also provide practical insights to design an effective zero-cost proxy for TAS: 1) Both MSA and MLP should be taken into consideration to rank ViT

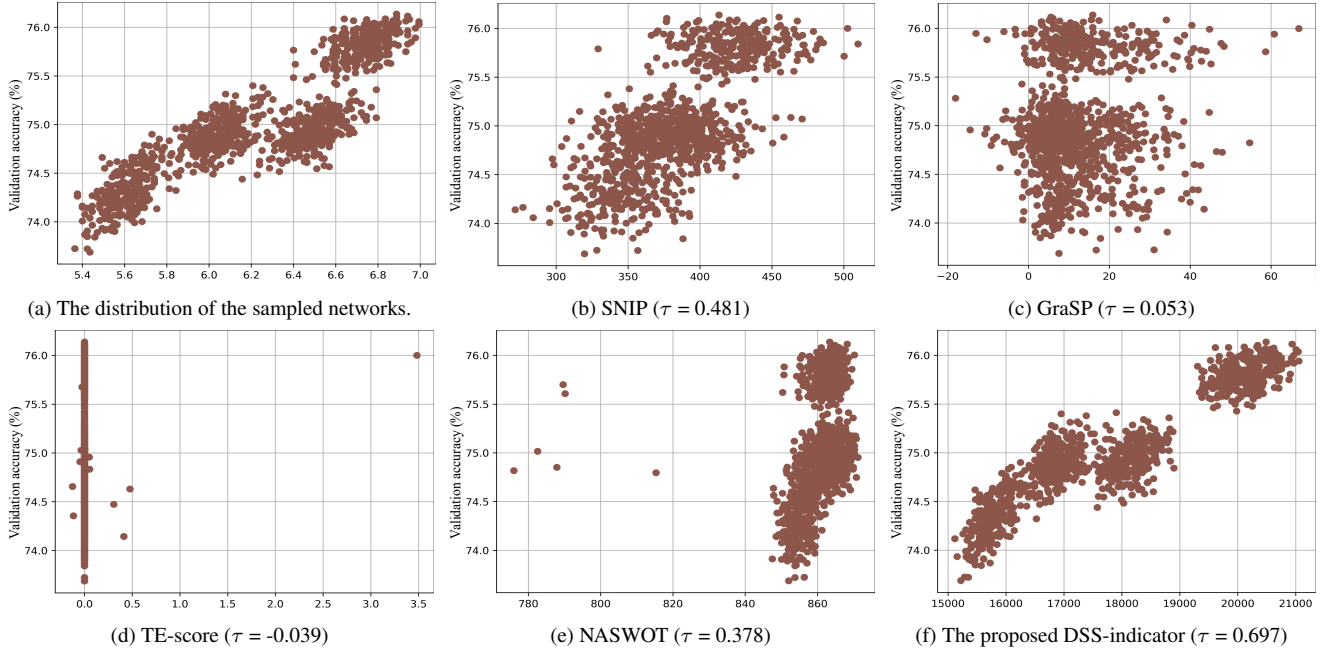


Figure 3. (a): The distribution plot of the 1,000 subnets ($\#param \in [5M, 7M]$) sampled from \mathcal{S}_A . Based on this proxy ViT benchmark, we compare the proposed DSS-indicator with the other counterparts. (b) - (f): The Kendall’s τ rank correlation between the evaluation scores of zero-cost metrics and the classification accuracy on the inherit networks from one pre-trained AutoFormer-Tiny supernet.

Table 6. The affection of different initialization seeds on the evaluation results of various zero-cost proxies.

Proxy	random seed				AVG	STD
	0	1	2	3		
SNIP [23]	0.481	0.530	0.486	0.507	0.501	0.019
GraSP [40]	0.053	0.126	0.138	0.152	0.117	0.038
TE-score [5]	-0.039	-0.003	-0.04	0.013	-0.017	0.023
NASWOT [29]	0.378	0.332	0.394	0.421	0.381	0.032
DSS-indicator (Ours)	0.697	0.697	0.697	0.697	0.697	0

effectively. And that is the main reason why our DSS-indicator is better than the other alternatives in ranking ViT networks. 2) Based on the results of SNIP [23] and our DSS-indicator, it is clear that the gradient matrix from the initialized ViT network contains rich information to evaluate the corresponding model. 3) Based on the performance of GraSP [40] and TE-score [5], we find that: despite its practical value in CNNs [40], the Hessian matrix of ViT is not easy to use and requires further efforts.

4.7. Consistency across different random seeds.

To check the stability of various proxies, we generate the results with four random seeds and calculate the statics. For simplicity, the experiments are conducted when the amount of parameter of ViT network is among $5 \sim 7M$. As listed in Tab. 6, there are certain fluctuations on several proxies under different seeds. Our DSS-indicator is invariant to dif-

ferent seeds (as mentioned in Sec. 3.4). For different seed influences the sampling of input data, the proxies that use the sampled data as the input could be unstable.

5. Conclusion

To accelerate the searching efficiency of TAS, for the first time, we propose an effective zero-cost proxy in evaluating ViT architectures. In specific, a ViT-oriented performance indicator, *i.e.*, DSS-indicator, is proposed. It is built on two theoretical perspectives: synaptic diversity and synaptic saliency. Based on these two dimensions, the proposed indicator measures the synaptic diversity on MSAs and the synaptic saliency on MLPs, respectively. Compared to other cutting-edge TAS methods, the random search guided by our DSS-indicator achieves a competitive performance among different popular ViT search spaces. And significantly, we greatly promote the searching efficiency of TAS: it merely takes 0.5 GPU days to seek relatively optimal ViT architecture, compared with 24 GPU days by the existing counterparts.

Acknowledgments: This work was supported by the National Science Fund for Distinguished Young Scholars (No.62025603), the National Natural Science Foundation of China (No.U21B2037, No.62176222, No.62176223, No.62176226, No.62072386, No.62072387, No.62072389, and No.62002305), Guangdong Basic and Applied Basic Research Foundation (No.2019B1515120049), the Natural Science Foundation of Fujian Province of China (No.2021J01002), and CAAI-Huawei MindSpore Open Fund.

References

- [1] Mohamed S Abdelfattah, Abhinav Mehrotra, Łukasz Dudziak, and Nicholas D Lane. Zero-cost proxies for lightweight nas. In *ICLR*, 2021. 2, 4, 5
- [2] Alaaeldin Ali, Hugo Touvron, Mathilde Caron, Piotr Bojanowski, Matthijs Douze, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Jakob Verbeek, et al. Xcit: Cross-covariance image transformers. *NeurIPS*, 2021. 7
- [3] Boyu Chen, Peixia Li, Chuming Li, Baopu Li, Lei Bai, Chen Lin, Ming Sun, Wanli Ouyang, et al. Glit: Neural architecture search for global and local image transformer. In *ICCV*, 2021. 1, 2, 3, 5, 6
- [4] Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. In *ICCV*, 2021. 1, 2, 3, 4, 5, 6, 7
- [5] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *ICLR*, 2021. 2, 3, 7, 8
- [6] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *Arxiv preprint 2102.10882*, 2021. 2, 5, 6
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 5
- [8] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *ICLR*, 2019. 3
- [9] Xuanyi Dong and Yi Yang. One-shot neural architecture search via self-evaluated template network. In *ICCV*, 2019. 1, 2
- [10] Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *ICML*, 2021. 3
- [11] 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. In *ICLR*, 2020. 1, 2, 5, 6, 7
- [12] Maryam Fazel. *Matrix rank minimization with applications*. PhD thesis, PhD thesis, Stanford University, 2002. 3
- [13] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *ECCV*, 2020. 1, 2
- [14] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. In *NeurIPS*, 2021. 2, 5, 6
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 5, 6, 7
- [16] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *CVPR*, 2019. 6
- [17] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. In *ICCV*, 2021. 1, 2, 4, 6, 7
- [18] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *ICCV*, 2019. 5, 6
- [19] Jie Hu, Liujuan Cao, Yao Lu, ShengChuan Zhang, Yan Wang, Ke Li, Feiyue Huang, Ling Shao, and Rongrong Ji. Istr: End-to-end instance segmentation with transformers. *arXiv preprint arXiv:2105.00637*, 2021. 1
- [20] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 1938. 7
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 5, 7
- [22] Namhoon Lee, Thalaiyasingam Ajanthan, Stephen Gould, and Philip HS Torr. A signal propagation perspective for pruning neural networks at initialization. In *ICLR*, 2019. 4
- [23] Namhoon Lee, Thalaiyasingam Ajanthan, and Philip HS Torr. Snip: Single-shot network pruning based on connection sensitivity. In *ICLR*, 2019. 2, 5, 7, 8
- [24] Ming Lin, Pichao Wang, Zhenhong Sun, Heseng Chen, Xiuyu Sun, Qi Qian, Hao Li, and Rong Jin. Zen-nas: A zero-shot nas for high-performance deep image recognition. In *ICCV*, 2021. 3
- [25] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. 5, 6
- [26] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *ICLR*, 2018. 1, 3
- [27] 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*, 2021. 2, 5, 6
- [28] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2018. 5
- [29] Joe Mellor, Jack Turner, Amos Storkey, and Elliot J Crowley. Neural architecture search without training. In *ICML*, 2021. 2, 3, 7, 8
- [30] Fanxu Meng, Hao Cheng, Ke Li, Huixiang Luo, Xiaowei Guo, Guangming Lu, and Xing Sun. Pruning filter in filter. In *NeurIPS*, 2020. 4
- [31] Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? In *NeurIPS*, 2019. 4
- [32] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, 2020. 6
- [33] Russell Reed. Pruning algorithms—a survey. *IEEE transactions on Neural Networks*, 1993. 4
- [34] David So, Quoc Le, and Chen Liang. The evolved transformer. In *ICML*, 2019. 1
- [35] 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, 5, 6

- [36] Hidenori Tanaka, Daniel Kunin, Daniel LK Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In *NeurIPS*, 2020. 4
- [37] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*. PMLR, 2021. 1, 2, 5, 6, 7
- [38] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. 1, 2
- [39] Elena Voita, David Talbot, Fedor Moiseev, Rico Sennrich, and Ivan Titov. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *ACL*, 2019. 4
- [40] Chaoqi Wang, Guodong Zhang, and Roger Grosse. Picking winning tickets before training by preserving gradient flow. In *ICLR*, 2020. 2, 3, 5, 7, 8
- [41] Hanrui Wang, Zhanghao Wu, Zhijian Liu, Han Cai, Ligeng Zhu, Chuang Gan, and Song Han. Hat: Hardware-aware transformers for efficient natural language processing. In *ACL*, 2020. 1
- [42] Wenxuan Wang and Zhaopeng Tu. Rethinking the value of transformer components. In *Proceedings of the 28th International Conference on Computational Linguistics*, 2020. 4
- [43] 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*, 2021. 1, 2, 5, 6
- [44] Yifan Xu, Zhijie Zhang, Mengdan Zhang, Kekai Sheng, Ke Li, Weiming Dong, Liqing Zhang, Changsheng Xu, and Xing Sun. Evo-vit: Slow-fast token evolution for dynamic vision transformer. *AAAI*, 2022. 2
- [45] Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Richard G Baraniuk, Zhangyang Wang, and Yingyan Lin. Drawing early-bird tickets: Towards more efficient training of deep networks. In *ICLR*, 2019. 4
- [46] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zihang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *ICCV*, 2021. 1, 5, 6, 7
- [47] Xiawu Zheng, Rongrong Ji, Yuhang Chen, Qiang Wang, Baochang Zhang, Jie Chen, Qixiang Ye, Feiyue Huang, and Yonghong Tian. Migo-nas: Towards fast and generalizable neural architecture search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):2936–2952, 2021. 1, 2
- [48] Xiawu Zheng, Yang Zhang, Sirui Hong, Huixia Li, Lang Tang, Youcheng Xiong, Jin Zhou, Yan Wang, Xiaoshuai Sun, Pengfei Zhu, et al. Evolving fully automated machine learning via life-long knowledge anchors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):3091–3107, 2021. 1, 2
- [49] Daquan Zhou, Bingyi Kang, Xiaojie Jin, Linjie Yang, Xiaochen Lian, Zihang Jiang, Qibin Hou, and Jiashi Feng. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021. 3