OVO: One-shot Vision Transformer Search with Online distillation

Zimian Wei[1], Hengyue Pan[1,*], Xin Niu[1], Dongsheng Li[1]
[1]College of computer, National University of Defense Technology

Abstract

Pure transformers have shown great potential for vision tasks recently. However, their accuracy in small or medium datasets is not satisfactory. Although some existing methods introduce a CNN as a teacher to guide the training process by distillation, the gap between teacher and student networks would lead to sub-optimal performance. In this work, we propose a new One-shot Vision transformer search framework with Online distillation, namely OVO. OVO samples sub-nets for both teacher and student networks for better distillation results. Benefiting from the online distillation, thousands of subnets in the supernet are well-trained without extra finetuning or retraining. In experiments, OVO-Ti achieves 73.32% top-1 accuracy on ImageNet and 75.2% on CIFAR-100, respectively.

1. Introduction

Vision transformer recently has drawn significant attention in computer vision due to its excellent model capability and superior potential in capturing long-range dependencies. ViT [9] is able to learn powerful visual representations from images and achieves very competitive performance compared to previous convolutional neural network models [12, 19].

However, their performance is not satisfactory when training from scratch on medium or small datasets. Extensive computing resources are involved in training transformers with a large private labeled image dataset (JFT-300M [28]). One of the most straightforward solutions is to apply knowledge distillation with a pre-trained strong convolutional network. DeiT [31] introduced a distillation token in the vision transformer ensuring that it learns from the convnet through attention. Unfortunately, the intrinsic gap between teacher and student networks would lead to sub-optimal results.

In this work, we present a new architecture search algorithm, named OVO, dedicated to finding optimal vision transformer models and distillation strategies simultaneously. Our approach mainly addresses two challenges in transformer search. 1) How to search for a good combination of the key factors in transformers, such as network depth, embedding dimension, and head number? 2) How to achieve better distillation results for different sub-networks?

To tackle the challenges, we first construct a large search space covering the main changeable dimensions of transformers, including embedding dimension, number of heads, query/key/value dimension, MLP ratio, and network depth. Moreover, we build a simple supernet of ResNet as the teacher network, which is trained mutually with the vision transformer to improve performance. The central idea is to enable the best-matched teacher and student pairs during the supernet training stage. This strategy is different from most one-shot NAS methods [11, 4, 34], in which only the weights of vision transformers are trainable.

We observe that when using the online distillation for transformer supernet training, the performance of these subnets is improved. This advantage allows our method to obtain better architectures while maintaining smaller computation costs. We perform an evolutionary search when the supernet finishes training. Experiments on ImageNet [6] demonstrate that our method achieves superior performance to the handcrafted state-of-the-art transformer models.

In summary, our major contributions are as follows. 1) We propose an online distillation method that automatically searches for optimal distillation strategy during the supernet training. 2) We propose a simple yet effective vision transformer framework, which produces thousands of high-quality transformers on tiny or medium datasets such as CIFAR-100 and ImageNet.

2. Background

Vision Transformer Transformer [9, 31] is originated from the natural language filed [33, 21, 8], which has shown its great potential for visual recognition tasks recently. To better understand our method, we revisit the basic architecture of the vision transformer as follows.

Given a 2D image, we first uniformly split it into a sequence of 2D patches, which is also referred as tokens in

* Corresponding author: hengyuepan@nudt.edu.cn.

natural language processing tasks. Then we flatten and transform the patches to D-dimension vectors, named patch embeddings, by either linear projection [9] or several CNN layers [37]. A learnable class embedding is injected into the head of the sequence to represent the whole image. Position embeddings are added to the patch embeddings to retain positional information. The combined embeddings are then fed to repeated transformer blocks. At last, a linear layer is used for the final classification.

A transformer block consists of multihead self-attention (MSA), and multi-layer perceptron (MLP) modules. LayerNorm (LN) [1] is applied before each module, and residual connections after every module. The details of MSA and MLP are given below.

Multihead Self-Attention (MSA). In a standard self-attention module, the input sequence $z \in \mathbb{R}^{N\times D}$ will be first linearly transformed to queries $Q \in \mathbb{R}^{N\times D_h}$, keys $K \in \mathbb{R}^{N\times D_h}$ and values $V \in \mathbb{R}^{N\times D_h}$, where N is the number of tokens, D is the embedding dimension, D_h is the Q-K-V dimension. Then we compute the weighted sum over all values for each element in the sequence. The weights or attention are based on the pairwise similarity between two elements of the sequence:

$$\text{Attention}(Q,K,V) = \text{softmax}\Big(\frac{QK^T}{\sqrt{d_h}}\Big)V, \tag{1}$$

where $\frac{1}{\sqrt{d_h}}$ is the scaling factor. Lastly, a fully connected layer is applied. Multihead self-attention splits the queries, keys and values into different heads and performs self-attention in parallel and projects their concatenated outputs.

Multi-Layer Perceptron (MLP). The MLP block consists two fully connected layers with an activation function, usually GELU [15]. In this work, we focus on finding optimal choices of the MLP ratios in each layer.

One-Shot NAS One-shot NAS typically adopts a weight sharing strategy to avoid training each subnet from scratch [11, 25]. The architecture search space $\mathcal{A}$ is encoded in a supernet, denoted as $\mathcal{N}(\mathcal{A}, W)$, where W is the weight of the supernet. W is shared across all the architecture candidates, *i.e*., subnets $\alpha \in \mathcal{A}$ in $\mathcal{N}$. The search of the optimal architecture α^* in one-shot NAS is usually formulated as a two-stage optimization problem. The first-stage is to optimize the weight W by

$$W_{\mathcal{A}} = \underset{W}{\arg\min}\, \mathcal{L}_{train}(\mathcal{N}(\mathcal{A}, W)), \tag{2}$$

where $\mathcal{L}_{train}$ represents the loss function on the training dataset. To reduce memory usage, one-shot methods usually sample subnets from $\mathcal{N}$ for optimization. The second-stage is to search architectures via ranking the performance of subnets $\alpha \in \mathcal{A}$ based on the learned weights in $W_{\mathcal{A}}$:

$$\alpha^* = \underset{\alpha\in\mathcal{A}}{\arg\max}\ \text{Acc}_{val}\left(\mathcal{N}(\alpha, w)\right), \tag{3}$$

where the sampled subnet α inherits weight w from $W_{\mathcal{A}}$, and Acc_{val} indicates the top-1 accuracy of the architecture α on the validation dataset. Since it is impossible to enumerate all the architectures $\alpha \in \mathcal{A}$ for evaluation, prior works resort to random search [23, 2], evolution algorithms [27, 11] or reinforcement learning [26, 29] to find the most promising one.

Knowledge Distillation Knowledge distillation is proposed by Hinton et al. [17] and has become a commonly used technology for model compression and acceleration. As mentioned in [10], the knowledge to be distilled can be divided into three kinds, *i.e.*, response-based knowledge [17], feature-based knowledge [14] and relation-based knowledge [24, 32]. Our method is highly related to feature-based knowledge distillation. Most applications of knowledge distillation are based on the setting of a strong teacher model and a weak student model, to achieve model compression and acceleration. Different from them, our goal of using CNN teacher is making the learning process of vision transformer easier so that the VT can be trained better. In our setting, the performance of the teacher will not be the performance bottleneck of the VT, since the VT is still learning by itself and can play to the advantage of the transformer. Therefore, a lightweight CNN teacher would suffice. A recent proposed method DeiT [31] also uses knowledge distillation, which makes the VT learn the classification results of the CNN teacher. However, a CNN of comparable size to the VT is required in DeiT. By comparison, our method can achieve much higher performance with just a lightweight CNN.

3. OVO

We give a detailed description of our method. In Section 3.1, we briefly introduce the dilemma of supernet training and the motivation of our method. In Section 3.2, we propose an online distillation method during the supernet training, which automatically samples the teacher network and student network for distillation. Finally, in Section 3.3, we present details of the search pipeline to obtain our models.

3.1. The dilemma of supernet training for ViT

During supernet training in AutoFormer, a subnet is uniformly sampled in each training iteration $\alpha = (\alpha^{(1)}, ...\alpha^{(i)}, ...\alpha^{(l)})$. The sampled weights $w = (w^{(1)}, ...w^{(i)}, ...w^{(l)})$ in the supernet $W_{\mathcal{A}}$ are updated while the rests are frozen. However, the supernet training is not stable for vision transformers, and it takes a long training period (500 epochs) for each supernet to obtain satisfactory performance of its assembled sub-networks.

Other methods [36, 35] use a sandwich sampling strategy, which samples multiple sub-networks (including the largest, smallest, and two random choices), then aggregates the gradients together for each mini-batch. The heavy computation cost of supernet training is more severe when a sandwich sampling strategy is applied.

3.2. Supernet Training with Online Distillation

We update one randomly sampled sub-network at each iteration as proposed in [11] when training both teacher and student networks. We train sub-networks from the teacher supernet with ground truth labels, while training sampled student networks with KD following [22].

Equipped with online distillation, one-shot NAS is capable of searching transformer architectures in an efficient and effective manner. Compared with classical one-shot NAS methods, our method has two advantages. 1) *Faster convergence*. knowledge from the CNN provides inductive bias , which helps each transformer block converges faster than the previous independent training. 2) *Better subnets performance*. The subnets trained with online distillation could achieve better performance on tiny or medium datasets.

3.3. Search Pipeline

After the supernet finishes training, we conduct evolutionary search to pick sub-networks with maximize classification accuracy. At the beginning of the evolution search, we randomly sample N architectures as seeds. All sub-networks are evaluated on the validation dataset based on the inherited weights from the supernet. The top k architectures are selected as parents to generate the next generation by crossover and mutation. For a crossover, two randomly selected parent networks are crossed to produce a new one during each generation. When a mutation is conducted, a parent network first mutates its depth with probability P_d, then each block is mutated with a probability of P_m to produce a new architecture.

4. Experiments

In this section, we first present the implementation details of OVO. Then we present the performance of OVO with comparisons with other state-of-the-art models designed manually or automatically.

4.1. Implementation Details

We conduct OVO on AutoFormer search space [3]. The datasets that we applied to evaluate OVO are ImageNet [7] and CIFAR-100 [20]. OVO includes a supernet training stage and a search stage. In the supernet training stage, We train the supernet using a similar recipe as DeiT [31]. Specifically, RandAugment [5], Cutmix [38], Mixup [39] and random erasing are adopted as data augmentation techniques. Notably, in each iteration we sample one random path from both teacher and student networks, then train the paths using one batch data. Images are split into patches of size 16x16. Before conducting the evolutionary search, we formulate validation set by sub-sampling 100 images per class from training examples, while reserving the ImageNet validation set for testing. We set the population size to 50 and number of generations to 20. In each generation, the top 10 architectures are selected as the parents to generate child networks by mutation and crossover. The mutation probability P_d and P_m are set to 0.2 and 0.4.

4.2. Results on ImageNet

We compare the performance of the searched optimal model with that of state-of-the-art CNNs and ViTs on ImageNet. We train the supernet of OVO on ImageNet-1K and search the target transformer model with a specified parameter size. After the supernet finishes training, subnets inherit weights directly, without extra retraining and other post-processing. The performance is reported in Table 1. It is clear that OVO achieves higher accuracy than the other state-of-the-art models.

4.3. Results on CIFAR-100

We validate the effectiveness of OVO on smaller datasets such as CIFAR-100 by conducting the same experiment as in ImageNet-1K. Specifically, the supernet is trained from scratch on the CIFAR-100 dataset, in which candidates of different subnets share the weights in a slimmable manner [36]. The results are shown in Table 2. Our OVO outperforms recent excellent methods [31, 16], proving the effectiveness of our method on relatively small datasets.

5. Conclusion

In this work, we propose a new one-shot transformer search method, OVO, which is equipped with online distillation by automatically sampled teacher and student networks. Under this strategy, the performance of sub-nets within the supernet are improved. Extensive experiments demonstrate the proposed strategy can improve the training of supernet and find promising architectures. Our searched OVOs achieve state-of-the-art results on ImageNet and tiny dataset CIFAR-100, which show the effectiveness of online distillation.

References

[1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 2

[2] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *ICML*, 2018. 2

[3] Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. In *Proceedings of the IEEE/CVF International Con-*

Table 1. Results on ImageNet with comparisons to state-of-the-arts. † refers to our implementation.

Models	Top-1 Acc.	Top-5 Acc.	#Param	FLOPs	Resolution	Model Type	Design Type
MobileNetV3-Large1.0 [18]	75.2%	-	5.4M	0.22G	224^2	CNN	Auto
EfficietNet-B0 [30]	77.1%	93.30%	5.4M	0.39G	224^2	CNN	Auto
DeiT-tiny [31]	72.2%	91.10%	5.7M	1.2G	224^2	Transformer	Manual
AutoFormer-tiny† [3]	73.11%	91.10%	5.7M	1.3G	224^2	Transformer	Auto
OVO-Ti	73.32%	91.13%	5.7M	1.3G	224^2	Transformer	Auto

Table 2. Results on CIFAR100.

Models	Top-1 Acc.	Resolution	Model Type
ResNet-56 [13]	70.43%	32^2	CNN
ResNet-18 [13]	79.00%	224^2	CNN
DeiT-Ti [31]	65.08%	224^2	Transformer
PiT-Ti [16]	73.58%	224^2	Transformer
OVO-Ti	75.20%	224^2	Transformer

ference on Computer Vision*, pages 12270–12280, 2021. 3, 4

[4] Xiangxiang Chu, Bo Zhang, Ruijun Xu, and Jixiang Li. Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search. *arXiv preprint arXiv:1907.01845*, 2019. 1

[5] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPR Workshops*, 2020. 3

[6] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 1

[7] 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. 3

[8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *NAACL*, 2019. 1

[9] 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. *ICLR*, 2021. 1, 2

[10] Jianping Gou, Baosheng Yu, Stephen J Maybank, and Dacheng Tao. Knowledge distillation: A survey. *International Journal of Computer Vision*, 129(6):1789–1819, 2021. 2

[11] 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. *ECCV*, 2020. 1, 2, 3

[12] Kai Han, Yunhe Wang, Hanting Chen, Xinghao Chen, Jianyuan Guo, Zhenhua Liu, Yehui Tang, An Xiao, Chunjing Xu, Yixing Xu, et al. A survey on visual transformer. *arXiv preprint arXiv:2012.12556*, 2020. 1

[13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 4

[14] Tong He, Chunhua Shen, Zhi Tian, Dong Gong, Changming Sun, and Youliang Yan. Knowledge adaptation for efficient semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 578–587, 2019. 2

[15] Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*, 2016. 2

[16] Byeongho Heo, Sangdoo Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 11936–11945, 2021. 3, 4

[17] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 2

[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. 4

[19] Salman Khan, Muzammal Naseer, Munawar Hayat, Syed Waqas Zamir, Fahad Shahbaz Khan, and Mubarak Shah. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021. 1

[20] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 3

[21] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. Albert: A lite bert for self-supervised learning of language representations. *ICLR*, 2020. 1

[22] Kehan Li, Runyi Yu, Zhennan Wang, Li Yuan, Guoli Song, and Jie Chen. Locality guidance for improving vision transformers on tiny datasets. In *European Conference on Computer Vision*, pages 110–127. Springer, 2022. 3

[23] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *UAI*, 2019. 2

[24] Nikolaos Passalis, Maria Tzelepi, and Anastasios Tefas. Heterogeneous knowledge distillation using information flow modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2339–2348, 2020. 2

[25] Houwen Peng, Hao Du, Hongyuan Yu, Qi Li, Jing Liao, and Jianlong Fu. Cream of the crop: Distilling prioritized paths for one-shot neural architecture search. *NeurIPS*, 2020. 2

[26] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *ICML*, 2018. 2

[27] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI*, 2019. 2

[28] Chen Sun, Abhinav Shrivastava, Saurabh Singh, and Abhinav Gupta. Revisiting unreasonable effectiveness of data in deep learning era. In *Proceedings of the IEEE international conference on computer vision*, pages 843–852, 2017. 1

[29] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *CVPR*, 2019. 2

[30] Mingxing Tan and Quoc V. Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. 4

[31] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herv'e J'egou. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020. 1, 2, 3, 4

[32] Frederick Tung and Greg Mori. Similarity-preserving knowledge distillation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1365–1374, 2019. 2

[33] 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

[34] 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 *CVPR*, 2019. 1

[35] 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. *NeurIPS*, 2020. 2

[36] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas Huang. Slimmable neural networks. *ICLR*, 2019. 2, 3

[37] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. *arXiv preprint arXiv:2101.11986*, 2021. 2

[38] Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 3

[39] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 3