

Multi-Dimensional Vision Transformer Compression via Dependency Guided Gaussian Process Search

Zejiang Hou, Sun-Yuan Kung
Princeton University
{zejiangh, kung}@princeton.edu

Abstract

Vision transformers (ViT) have recently attracted considerable attentions, but the huge computational cost remains an issue for practical deployment. Previous ViT pruning methods tend to prune the model along one dimension solely, which may suffer from excessive reduction and lead to sub-optimal model quality. In contrast, we advocate a multi-dimensional ViT compression paradigm, and propose to harness the redundancy reduction from attention head, neuron and sequence dimensions jointly. Firstly, we propose a statistical dependence based pruning criterion that is generalizable to different dimensions for identifying the deleterious components. Moreover, we cast the multi-dimensional ViT compression as an optimization problem, objective of which is to learn an optimal pruning policy across the three dimensions while maximizing the compressed model’s accuracy under a computational budget. The problem is solved by an adapted Gaussian process search with expected improvement. Experimental results show that our method effectively reduces the computational cost of various ViT models. For example, our method reduces 40% FLOPs without top-1 accuracy loss for DeiT and T2T-ViT models on the ImageNet dataset, outperforming previous state-of-the-art ViT pruning methods.

1. Introduction

Vision transformers (ViT) [8, 26] have achieved substantial progress in prevalent computer vision tasks such as image classification, object detection and semantic segmentation. However, ViT models suffer from excessive computational and memory cost, impeding their deployment in resource-restricted or low-powered applications. Although various model compression algorithms have been proposed for convolutional neural networks (CNN), it is not immediately clear whether they are the same effective for vision transformers and there are only few works [5, 22, 23, 25, 33] on accelerating vision transformers.

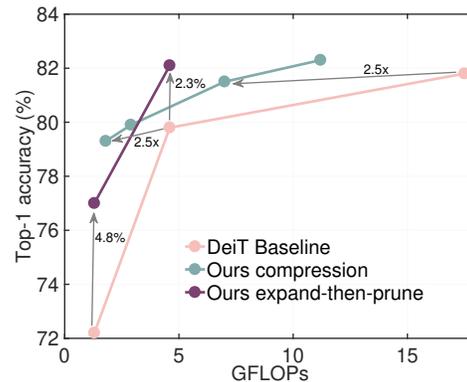


Figure 1. The accuracy-FLOPs curve of our compressed vision transformer models compared with the baseline models.

Prior works on compressing transformer models [6, 28, 32] in natural language processing (NLP) place their focus on addressing the quadratic complexity of the softmax-attention operation, presumably because the input has very long sequence in NLP tasks. However, for vision transformers, the softmax-attention operation constitutes a small fraction of the total FLOPs, as shown in Table 1. Instead, the projection layers are the major computation bottleneck. The computation complexity of these projection layers are affected by the number of attention heads, the number of neurons, and sequence length jointly.

Model	Softmax-attention	MHSA projections	FFN projections	Total
DeiT-S [26]	0.36G (8%)	1.39G (30%)	2.79G (61%)	4.6G
DeiT-B [26]	0.72G (4%)	5.58G (32%)	11.15G (64%)	17.5G

Table 1. FLOPs composition of DeiT-S/-B models. “MHSA”: multi-head self-attention. “FFN”: feed-forward network.

ViT models split the input image into a sequence of image tokens. Considering that not all tokens contribute to the final predictions [23] and the high similarity between tokens within a layer [25], recent ViT compression methods [23, 25, 33] adopt unstructured token pruning by removing the redundant and unimportant tokens. Since the self-

attention operator can process variable sequence length, these unstructured token pruning methods can achieve practical acceleration. However, to achieve more compelling computational cost reduction, excessive pruning of a single dimension (i.e., sequence length) leads to unacceptable accuracy loss, as will be discussed in Section 4.4. This motivates our study on how to search an optimal compression policy to reduce the computational cost from multiple dimensions jointly, in order to achieve better computation-accuracy trade-off. In this work, we consider reducing the number of attention heads in MHSA modules, the number of neurons in FFN modules, and the sequence length jointly.

Multi-dimensional ViT compression is challenging. Most of the current pruning algorithms are designed by the unique property of a single dimension, e.g., using the column mean of the attention matrix for sequence reduction [27]. These methods may hardly generalize to other pruning dimensions. Moreover, the large decision space stemming from three integrated dimensions makes it hard to decide how much of each dimension should be compressed.

Our contributions. To rectify the aforementioned problems, we firstly propose a general data-aware pruning criterion that is applicable to both structured neuron or head reduction and unstructured sequence reduction. The criterion measures the statistical dependency between the features of a dimension and the output predictions of the model based on the Hilbert-Schmidt norm of the cross-variance operator. Moreover, we formulate the multi-dimensional ViT compression as an optimization problem, seeking the optimal pruning policy (i.e, pruning ratios across the three dimensions) that maximize the compressed model’s accuracy under a target computational cost. Considering the non-differentiability of the problem and the optimization efficiency, we propose to use Gaussian process (GP) search with expected improvement to estimate the compressed model’s accuracy for different pruning policies, and solve the problem by non-linear programming solver. To fit a GP, we need to evaluate the actual accuracy of a small set of sampled pruning policies. We further design a weight sharing mechanism for fast accuracy evaluation without training each compressed model from scratch. Experimentally, our method achieves higher accuracy than previous state-of-the-art ViT compression methods under same FLOPs. When compressing DeiT [26] and T2T-ViT [37] on ImageNet [7], our method reduces 40% ~ 60% FLOPs and yields $1.3 \times \sim 2.2 \times$ practical speedup without significant accuracy drop.

2. Related Works

Model compression on vision transformer. To improve the efficiency of ViT models, [5, 38] applies structured neuron pruning or unstructured weight pruning. [22, 23, 25, 33] applies dynamic or static token sparsification. [26] proposes

a knowledge distillation method specific to transformer by introducing a distillation token. [20] uses post-training quantization to reduce the model size. However, multi-dimensional compression of ViT models has been rarely explored, and its effectiveness compared to uni-dimensional compression is unknown. In this work, we will show that excavating the redundancy from multiple dimensions is imperative to achieve more appealing FLOPs reductions, and our method achieves state-of-the-art pruning results compared to previous methods.

Multi-dimensional pruning methods have been proposed for compressing CNNs. [13, 17, 31] impose sparsity regularization, e.g., group LASSO, to prune channels and layers in CNNs. [10] uses L1 regularization to prune channels and feature-map spatial sizes in CNNs. In addition to the regularization-based methods, [21, 29] directly search the number of channels, layers and spatial sizes under a FLOPs budget by reinforcement learning or polynomial regression. In contrast, our method is specially designed for ViT compression by jointly pruning attention heads in the MHSA modules, neurons in the FFN modules, and sequence. We propose a dependency based pruning criterion and an efficient Gaussian process search to learn the optimal compression policy.

One-Shot NAS. Our multi-dimensional ViT compression method is also inspired by one-shot NAS [2, 11, 16]. In one-shot NAS, the architecture search space is encoded into a supernet, whose weight is shared among different architecture candidates. In the first stage, the weight of the supernet is trained by sampling different subnets during optimization. The second stage is to search the optimal architecture by ranking the performance of different subnets using the weight inherited from the supernet. AutoFormer [4] and GLiT [3] proposed one-shot NAS frameworks dedicated to vision transformer search. The differences of our method are summarized as follows: (1) our method is to compress an existing architecture; (2) we propose a Gaussian process search to solve the multi-dimensional ViT compression problem; (3) the weight sharing in our method is guided by our dependency based pruning criterion.

3. Methodology

3.1. Preliminary

ViT model contains interleaved multi-head self-attention (MHSA) and feed-forward network (FFN) modules. Denote the input features to MHSA and FFN in the l -th layer by $X^l, Z^l \in \mathbb{R}^{N \times d}$, where N is the sequence length and d is the embedding dimension. The MHSA module performs the following operations:

$$\text{MHSA}(X^l) = \sum_{h=1}^H \text{softmax}\left(\frac{Q_h K_h^T}{\sqrt{d_h}}\right) V_h W_h^o,$$

where the query, key and value features are computed by $Q_h = X^l W_h^Q, K_h = X^l W_h^K, V_h = X^l W_h^V$ ($W_h^Q, W_h^K, W_h^V \in \mathbb{R}^{d \times d_h}$), H is the number of attention heads, and $W_h^o \in \mathbb{R}^{d_h \times d}$ is the output projection. The FFN module performs the following operations:

$$\text{FFN}(Z^l) = \sigma(Z^l W_1 + b_1) W_2 + b_2,$$

where σ denotes the non-linear activation function and $W_1, W_2^T \in \mathbb{R}^{d \times d'}$, $b_1 \in \mathbb{R}^{d'}$, $b_2 \in \mathbb{R}^d$ are the projection matrices and biases.

3.2. Dependency based pruning criterion

Our goal is to accelerate ViT models by pruning multiple dimensions jointly, including the number of neurons in FFN modules, the number of heads in MHSA modules, and the sequence length. We need a general criterion to identify the deleterious features in different dimensions. Intuitively, the unimportant features contribute least to the output predictions. In other words, the output of the model has weak dependency on the unimportant features. Thus, we propose a dependency based pruning criterion, which evaluates the importance based on the statistical dependency between the features and the output predictions of the model.

Denote the random vector of the features by Z and the random vector of the model outputs by Y . Let $P_{Z,Y}$ be the joint distribution between the two random variables. To measure the dependence between Z and Y , we use the cross-covariance operator [1] defined as:

$$C_{zy} := \mathbb{E}_{zy}[(\Phi(z) - \mu_z) \otimes (\Psi(y) - \mu_y)], \quad (1)$$

where Φ (or Ψ) represents a kernel mapping from the feature space (or the model output space) to a reproducing kernel Hilbert space (RKHS), with mean vector μ_z (or μ_y). \otimes denotes the tensor product. To summarize the degree of dependence between Z and Y , we use the Hilbert-Schmidt norm of the cross-covariance operator C_{zy} , which is denoted by $\|C_{zy}\|_{HS}^2$ and is computed by the trace of $C_{zy} C_{zy}^T$. As shown in the following theorem by [9], $\|C_{zy}\|_{HS}^2$ can characterize the independence between random variables:

Theorem 1 (C_{zy} and Independence) *Given RKHSs with characteristic kernels. Then, $\|C_{zy}\|_{HS}^2 = 0$ if and only if Z and Y are independent.*

Characteristic kernels, such as Gaussian RBF kernel $k(x, x') = \exp(-\|x - x'\|_2^2 / (2\sigma^2))$, allows us to measure an arbitrary mode of dependence (including non-linear dependence) between Z and Y . Features with high $\|C_{zy}\|_{HS}^2$ value have high dependency with the outputs of the model, indicating that the features have considerable influence to the output predictions, thus they should be retained.

To use the dependency criterion in practice, we need an empirical estimate from a batch of training samples. Denote

the kernel functions by $k(z, z')$ and $l(y, y')$. Let \mathbf{K}, \mathbf{L} be the Gram matrices ($\mathbf{K}_{i,i'} = k(z_i, z_{i'})$, $\mathbf{L}_{i,i'} = l(y_i, y_{i'})$) computed over the features and model outputs of B training samples. An empirical estimator of $\|C_{zy}\|_{HS}^2$ is given by:

$$\widehat{\|C_{zy}\|_{HS}^2} := (B-1)^{-2} \text{tr}(\mathbf{K}\mathbf{L}\mathbf{C}\mathbf{L}\mathbf{C}), \quad (2)$$

where $\mathbf{C} = \mathbf{I}_B - (1/B)\mathbf{1}_B\mathbf{1}_B^T$ is the centering matrix ($\mathbf{I}_B, \mathbf{1}_B$ represent identity matrix and all-ones vector) and $\text{tr}(\cdot)$ is the matrix trace operation. As shown in [9], this empirical estimator converges sufficiently: with high probability, $|\widehat{\|C_{zy}\|_{HS}^2} - \|C_{zy}\|_{HS}^2|$ is bounded by a small constant.

Our proposed dependency based pruning criterion can be applied to prune different dimensions, including attention heads, neurons, and sequence length, introduced next:

FFN neuron reduction. We prune neurons in the intermediate layer of FFN modules. Denote the features from the intermediate layer in the l -th FFN by $Z^l \in \mathbb{R}^{B \times N \times d'}$, where each neuron $j \in [d']$ has features $Z_{:,j}^l \in \mathbb{R}^{B \times N}$. Given a neuron pruning ratio κ^l , we retain $\lceil (1 - \kappa^l)d' \rceil$ important neurons. We compute the dependency score for each neuron in the layer by $\psi_j^l = \text{tr}(\mathbf{K}_j^l \mathbf{C}\mathbf{L}\mathbf{C})$, where $\mathbf{K}_j^l \in \mathbb{R}^{B \times B}$ is the Gram matrix defined over $Z_{:,j}^l$, i.e., $[\mathbf{K}_j^l]_{i,i'} = k(Z_{i,:}^l, Z_{i',:}^l)$. Then, we rank the dependency scores in descending order, and identify the important neurons by $\text{ArgTopK}(\{\psi_1^l \dots \psi_{d'}^l\}; \lceil (1 - \kappa^l)d' \rceil)$, which gives the neuron indices with the top $\lceil (1 - \kappa^l)d' \rceil$ dependency scores. The remaining bottom-ranking neurons are pruned.

Attention head reduction. Denote the output features of the self-attention operator in the l -th MHSA by $Z^l \in \mathbb{R}^{B \times N \times H \times d_h}$. Given a head pruning ratio ζ^l , the head pruning procedure is similar to neuron pruning, except that each head $h \in [H]$ has output features $Z_{:,h,:}^l \in \mathbb{R}^{B \times N \times d_h}$. To construct the feature Gram matrix, we perform mean pooling along the embedding dimension to obtain $\tilde{Z}_{:,h}^l = \text{mean}(Z_{:,h,:}^l; \text{dim}=-1)$ (-1 means the last tensor dimension), which has a size of $\mathbb{R}^{B \times N}$. Then, we compute the Gram matrix \mathbf{K}_h^l over $\tilde{Z}_{:,h}^l$ and the head dependency score ψ_h^l . We keep those heads with the top $\lceil (1 - \zeta^l)H \rceil$ dependency scores and prune the other bottom-ranking heads.

Sequence reduction. To achieve unstructured sequence reduction, we insert token selection layer (TSL) (no extra parameters) after the MHSA module and before the FFN module at each transformer layer. Let $Z^l \in \mathbb{R}^{B \times N \times d}$ be the input features to l -th TSL. Given a sequence reduction ratio ν^l , TSL outputs the selected $\lceil (1 - \nu^l)N \rceil$ important tokens from $Z^l \in \mathbb{R}^{B \times N \times d}$ according to indices $\text{ArgTopK}(\{\psi_1^l \dots \psi_N^l\}; \lceil (1 - \nu^l)N \rceil)$. That is, TSL extracts a sub-tensor from the input features, the output has size $\mathbb{R}^{B \times \lceil (1 - \nu^l)N \rceil \times d}$. The token importance scores $\psi_n^l, n \in$

$[N]$ are computed based on the Gram matrix \mathbf{K}_n^l defined over the token features $Z_{:,n,:}^l \in \mathbb{R}^{B \times d}$. After TSL, the sequence length in subsequent layers becomes $\lceil (1 - \nu^l)N \rceil$. And the subsequent TSLs select tokens from what are preserved by the previous TSLs.

3.3. Multi-dimensional ViT compression via Gaussian process search

Formulation. We aim to strike a balance among the three pruning dimensions so that the compressed model has the best accuracy under a computation constraint. To this end, the multi-dimensional compression is formulated as:

$$\begin{aligned} \max_{\{\kappa^l, \zeta^l, \nu^l\}_{l=1}^L} & \text{Accuracy}(\{\kappa^l, \zeta^l, \nu^l\}_{l=1}^L), \\ \text{s.t. } & \mathcal{C}(\{\kappa^l, \zeta^l, \nu^l\}_{l=1}^L) \leq \mathcal{T}, \end{aligned} \quad (3)$$

where $\text{Accuracy}(\cdot)$ gives the accuracy of the compressed model with pruning ratios $\{\kappa^l, \zeta^l, \nu^l\}_{l=1}^L$, $\mathcal{C}(\cdot)$ represents the FLOPs of the compressed model, and \mathcal{T} is the constraint.

Gaussian process search. To solve problem (3), we resort to Bayesian optimization, as it provides an efficient framework for objectives that may not be differentiable or expressed in a closed-form. We use Gaussian process (GP) [24] with expected improvement (EI) to estimate the accuracy function in closed form, so that problem (3) can be transformed to a simpler and solvable constrained non-linear optimization. As a guideline, the pseudo-code of our method is provided in Algorithm 1.

We denote a pruning policy by $\omega = \{\kappa^l, \zeta^l, \nu^l\}_{l=1}^L$. A Gaussian process is described by:

$$f \sim \mathcal{GP}(\mu(\cdot), k(\cdot, \cdot)),$$

where the GP has a mean function $\mu(\omega) = \mathbb{E}[f(\omega)]$ and a covariance kernel $k(\omega, \omega') = \mathbb{E}[(f(\omega) - \mu(\omega))(f(\omega') - \mu(\omega'))]$. We sample m different pruning policies $\Omega = \{\omega_i\}_{i=1}^m$ satisfying the constraint \mathcal{T} , obtain m compressed models using our dependency based pruning, evaluate their actual accuracy $\mathcal{A}(\omega_i)$ on a hold-out set, and fit the GP model by the set $\{\omega_i, \mathcal{A}(\omega_i)\}_{i=1}^m$. At a new pruning policy $\hat{\omega}$, the posterior of f at this point is given by:

$$\begin{aligned} \hat{f} & \sim \mathcal{N}(\hat{\mu}, \hat{\Sigma}), \\ \hat{\mu} & = \mu(\hat{\omega}) + k(\hat{\omega}, \Omega)k(\Omega, \Omega)^{-1}(\mathcal{A}(\Omega) - \mu(\Omega)), \\ \hat{\Sigma} & = k(\hat{\omega}, \hat{\omega}) - k(\hat{\omega}, \Omega)k(\Omega, \Omega)^{-1}k(\Omega, \hat{\omega}). \end{aligned}$$

The expected improvement in accuracy at this new policy $\hat{\omega}$ is computed in closed form by:

$$\text{EI}(\hat{\omega}) = (\mathcal{A}^* - \hat{\mu})\Phi((\mathcal{A}^* - \hat{\mu})/\hat{\Sigma}) + \hat{\Sigma}\phi((\mathcal{A}^* - \hat{\mu})/\hat{\Sigma}),$$

Algorithm 1: Multi-dimensional ViT compression.

- 1 **Input:** An L -layer ViT model with weights $\mathcal{W} = \{\mathbf{w}^1, \dots, \mathbf{w}^L\}$; pre-training iterations T_{pre} ; target computational cost \mathcal{T} ; population size m ; GP search iterations T_{gp} ; finetuning iterations T_{ft} ; training set \mathcal{D}_{tr} ; hold-out validation set \mathcal{D}_{val} ;
 - 2 **Output:** Compressed ViT model satisfying the constraint \mathcal{T} and its optimal weights \mathcal{W}^* ;
 - /* Pre-training with Eq. (5) */
 - 3 Randomly initialize the model weights \mathcal{W} ;
 - 4 **for** each training iteration $t \in [T_{\text{pre}}]$ **do**
 - 5 Sample a mini-batch (x, y) from \mathcal{D}_{tr} , sample a pruning policy ω from $U_{\mathcal{T}}$, and select weights $\mathcal{W}(\omega)$ by weight sharing as described in Sec.3.3;
 - 6 Compute training loss $\mathcal{L}(y|x; \mathcal{W}(\omega))$, backprop and update \mathcal{W} ;
 - /* GP search as described in Sec.3.3 */
 - 7 Randomly sample m different pruning policies $\{\omega_i\}_{i=1}^m$ satisfying \mathcal{T} , get m compressed models (with weights $\mathcal{W}(\omega_i)$) by dependency based pruning as described in Sec.3.2, evaluate their actual accuracy $\mathcal{A}(\omega_i)$ on \mathcal{D}_{val} , and fit a GP model with $\Omega = \{\omega_i, \mathcal{A}(\omega_i)\}_{i=1}^m$;
 - 8 **for** each search iteration $t \in [T_{\text{gp}}]$ **do**
 - 9 Solve the non-linear programming Eq.(4) by SQP to get pruning policy ω_t^* ;
 - 10 Evaluate the actual accuracy of ω_t^* on \mathcal{D}_{val} ;
 - 11 Augment $\{\omega_t^*, \mathcal{A}(\omega_t^*)\}$ to Ω and refine the GP model;
 - /* Final pruning and finetuning */
 - 12 Compress the ViT model with the optimal pruning policy $\omega_{T_{\text{gp}}}^*$ using dependency based pruning;
 - 13 Finetune the compressed model by T_{ft} iterations;
-

where Φ, ϕ represent the CDF and PDF of the standard normal distribution, and \mathcal{A}^* is the accuracy of the best policy in Ω . Therefore, the most promising policy ω^* to evaluate is the solution to the following non-linear programming:

$$\max_{\hat{\omega}} \text{EI}(\hat{\omega}), \quad \text{s.t. } \mathcal{C}(\hat{\omega}) \leq \mathcal{T}. \quad (4)$$

Since both the objective and constraints¹ in (4) have closed-form formulas, the problem can be solved by standard constrained optimization solver, and we use sequential quadratic programming (SQP) [15]. We iterate the search process by using the obtained ω^* and its actual accuracy $\mathcal{A}(\omega^*)$ to refine the GP model, and find the next (more optimal) policy until no accuracy improvement is observed. With the final optimal pruning policy, we apply our dependency based pruning, and finetune the compressed model.

¹FLOPs of a ViT model can be computed by closed-form formula with the pruning ratios in three dimensions.

Accuracy evaluation in GP search. Our GP search involves evaluating the actual accuracy of different compressed models. Instead of training many models with different pruning policies from scratch, we apply weight sharing [11] for efficient accuracy evaluation. We pre-train the full model weights \mathcal{W} in a way that the accuracy of sub-models with inherited weights are predictive for the accuracy obtained by training them independently. Our pre-training objective is given by:

$$\min_{\mathcal{W}} \mathbb{E}_{(x,y) \sim \mathcal{D}, \omega \sim U_{\mathcal{T}}} [\mathcal{L}(y|x; \mathcal{W}(\omega))] \quad (5)$$

where \mathcal{D} is the training set, $U_{\mathcal{T}}$ is a constrained uniform sampling distribution² of pruning policies, and \mathcal{L} is the loss function. $\mathcal{W}(\omega)$ means selecting weights from \mathcal{W} to form the compressed model with pruning policy ω . For neuron and head dimensions, we keep the weights corresponding to neurons or heads with the top dependency scores as described in Section 3.2. Reducing sequence dimension does not require modification to the model weights, thus all parameters are shared. Thanks to the weight sharing, we only need to train one set of weights and directly evaluate the accuracy of different pruning policies by inheriting weights from the full model. This takes much less time than training each compressed model from scratch and makes our GP search efficient.

4. Experiments

4.1. Setup

We conduct experiments on the ImageNet dataset [7] with representative ViT models, DeiT [26] and T2T-ViT [37], which were also used by previous ViT compression methods [5, 22, 23, 25, 33]. All experiments run on PyTorch framework with Nvidia A100 GPUs. We firstly pre-train the models from scratch with Eq.(5), and follow the same training hyper-parameters as the paper of DeiT and T2T-ViT. Then, we conduct Gaussian process (GP) search to obtain the optimal pruning policy. The target computational costs are listed as the FLOPs reductions in Table 2. The initial population size to fit a GP model is 100, and the GP search runs for 100 iterations. We randomly sample 50k images (50 images per class) from the training set of ImageNet for accuracy evaluation during GP search. Our GP search process is computationally efficient, taking less than 1 hour on a single A100 GPU for all cases. Based on the optimal pruning policy, we compress the pre-trained models along head, neuron and sequence dimensions using our dependency based pruning. The compressed model is finetuned following the same training strategy as [25, 26]. To compute the dependency score in Eq.(2), we randomly

²We sample the pruning policy repeatedly until the compressed model FLOPs satisfies the constraint \mathcal{T} .

Method	Top-1	Top-1 drop	FLOPs	FLOPs reduction
<i>DeiT-Small model</i>				
Baseline [26]	79.8%	-	4.6G	-
SPViT [12]	78.3%	1.5%	3.3G	29%
IA-RED ² [22]	79.1%	0.7%	3.1G	32%
S ² ViTE [5]	79.2%	0.6%	3.1G	32%
Evo-ViT [33]	79.4%	0.4%	2.9G	37%
DynamicViT [23]	79.3%	0.5%	2.9G	37%
Ours	79.9%	-0.1%	2.9G	37%
UVC [36]	78.4%	1.4%	2.4G	48%
Ours	79.3%	0.5%	1.8G	60%
<i>DeiT-Base model</i>				
Baseline [26]	81.8%	-	17.5G	-
VTP [38]	81.3%	0.5%	13.8G	22%
IA-RED ² [22]	80.3%	1.5%	11.8G	33%
S ² ViTE [5]	82.2%	-0.4%	11.8G	33%
SPViT [12]	81.6%	0.2%	11.7G	33%
Evo-ViT [33]	81.3%	0.5%	11.7G	33%
DynamicViT [23]	81.3%	0.5%	11.2G	36%
Ours	82.3%	-0.5%	11.2G	36%
UVC [36]	80.6%	1.2%	8.0G	55%
Ours	81.5%	0.3%	7.0G	60%
<i>T2T-ViT-14 model</i>				
Baseline [37]	81.5%	-	4.8G	-
PatchSlim [25]	81.1%	0.4%	2.9G	40%
Ours	81.7%	-0.2%	2.9G	40%

Table 2. Comparison of our compressed ViT models versus baselines and previous methods on ImageNet. Results of ‘‘Ours’’ on DeiT-S/B are obtained by applying the proposed method to reduce 40% and 60% FLOPs respectively, in order to compare with different methods. Negative ‘‘Top-1 drop’’ means that the accuracy improves over the baseline.

sample a mini-batch of 256 images from the training set and use Gaussian kernel with bandwidth $\sigma = 1$. In contrast to [23, 36], our method does not use knowledge distillation.

4.2. Comparison with the state-of-the-art ViT compression methods.

We compare with the latest ViT model compression methods, including sequence reduction methods (DynamicViT [23], IA-RED² [22], PatchSlim [25], Evo-ViT [33]), weight pruning methods (VTP [38], S²ViTE [5]), unified ViT compression method UVC [36], and NAS-based ViT pruning method [12]. The results are shown in Table 2.

Our method achieves noticeably higher accuracy than previous methods under same FLOPs. For example, our pruned DeiT-S model with 37% FLOPs reduction outperforms DynamicViT and Evo-ViT by 0.6% and 0.5% accuracy, respectively. On the other hand, at the same target accuracy, our method achieves higher FLOPs reduction. For example, our pruned DeiT-B model yields 60% FLOPs re-

Method	DeiT-Ti		DeiT-S	
	Top-1	FLOPs	Top-1	FLOPs
Baseline	72.2%	1.3G	79.8%	4.6G
Baseline _{2×}	73.9%	1.3G	81.0%	4.6G
Manifold [14]	75.1%	1.3G	81.5%	4.6G
UP-DeiT [35]	75.8%	1.3G	81.6%	4.6G
NViT-DeiT [34]	76.2%	1.3G	82.1%	4.6G
Ours	77.0%	1.3G	82.1%	4.6G

Table 3. Results of expand-then-compress on ImageNet. “2×” means doubling the training epochs when training the baseline.

Method	Top-1	FLOPs	Top-1	FLOPs
GLiT [3]	76.3%	1.4G	80.5%	4.4G
AutoFormer [4]	74.7%	1.3G	81.7%	5.1G
Ours	77.0%	1.3G	82.1%	4.6G

Table 4. Comparison with NAS results on ImageNet.

duction with 81.5% top-1 accuracy, compared to DynamicViT and Evo-ViT yielding less than 36% FLOPs reduction. These results clearly evidence the advantage of pruning multiple dimensions in the ViT models, when we aim to achieve compelling FLOPs reduction without compromising too much accuracy.

4.3. Results of expand-then-compress.

Apart from compressing the model for faster inference, our method can improve existing models for higher accuracy under the same FLOPs, as shown in Table 3. More exactly, we apply our method to compress a scaled-up DeiT-Ti model (width scaled by 2×), with the goal of reducing its FLOPs to the same level as the original DeiT-Ti. Notably, the obtained model achieves 77% top-1 accuracy at 1.3 GFLOPs, outperforming the original DeiT-Ti (trained with longer training epochs) by 3.1%. Same phenomenon also applies to DeiT-S where our method achieves 82.1% top-1 accuracy, improving the baseline by 1.1%. These results suggest that heavily compressed larger ViT models may achieve higher accuracy than small models. In summary, our method achieves better Pareto frontier compared to existing models as shown in Figure 1.

Comparison with NAS results. We also compare the expand-then-compress results obtained by our method against NAS results [3, 4]. As shown in Table 4, compressing a larger ViT model to smaller FLOPs can achieve higher accuracy than searching ViT architecture from scratch.

4.4. Ablation study.

Effect of multi-dimensional compression. We investigate the effectiveness of multi-dimensional compression in Table 5, where we compare with reducing the neuron, head or sequence individually for DeiT-B on ImageNet.

Model	Neuron	Head	Sequence	Top-1	FLOPs
DeiT-B	-	-	-	81.8%	17.5G
	✓	-	-	79.4%	7.4G
	-	-	✓	79.5%	7.0G
	✓	✓	-	80.4%	7.1G
	✓	-	✓	80.7%	7.0G
	-	✓	✓	80.2%	7.0G
	✓	✓	✓	81.5%	7.0G

Table 5. Ablation study on the effectiveness of multi-dimensional ViT compression. “✓” means that pruning is conducted along the corresponding dimension, while “-” means no pruning along the dimension. In the last row, jointly pruning along the neuron, head and sequence dimensions achieves the best accuracy. Results are obtained with DeiT-B model on ImageNet.

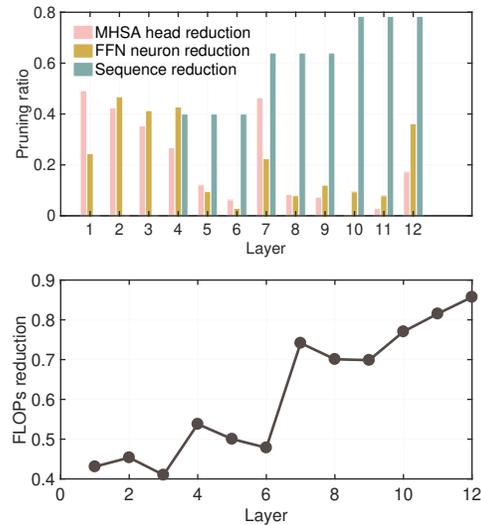


Figure 2. **Top:** Layer-wise pruning ratio in our learned multi-dimensional pruning policy by GP search. **Bottom:** Layer-wise FLOPs reduction. Results are obtained by applying our method to DeiT-B model on ImageNet.

Firstly, certain uni-dimensional compression method (e.g., head pruning alone) cannot yield significant FLOPs reduction, since the FLOPs of all the MHA modules only account for 40% of the total FLOPs. Secondly, although each dimension is prunable to some extent, excessive pruning of whichever dimension causes unacceptable accuracy loss, even for the fine-grained sequence reduction. In contrast, our multi-dimensional compression (last row of Table 5) achieves more FLOPs reduction with better accuracy. Searching the optimal policy to balance the FLOPs reduction from different dimensions is of vital importance if we aim to achieve significant acceleration on vision transformers. In Figure 2, we also visualize the learned pruning policy by plotting the layer-wise pruning ratio and FLOPs reduction. Most of the head and neuron reductions come

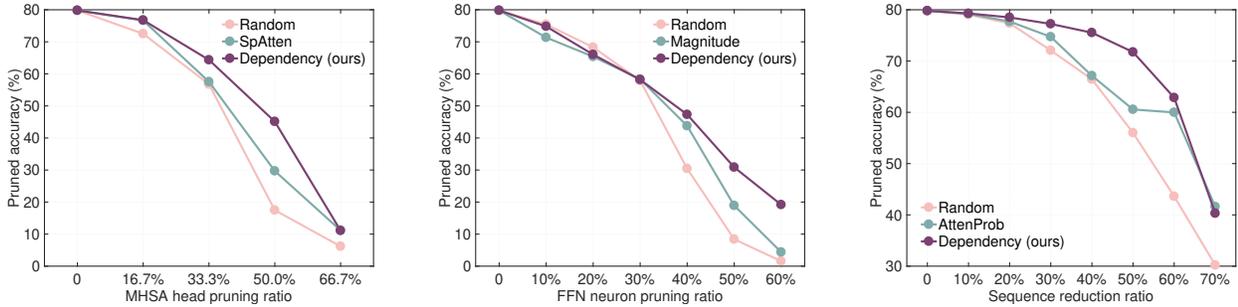


Figure 3. Compare our dependency criterion versus other metrics for pruning each dimension individually. The plots show the pruned accuracy (before finetuning) versus pruning ratios along each dimension. Results are obtained with DeiT-S on ImageNet.

Model	DeiT-S		DeiT-B		T2T-ViT-14	
	Base	Pruned	Base	Pruned	Base	Pruned
(FLOPs reduction)	(0%)	(37% / 60%)	(0%)	(36% / 60%)	(0%)	(40%)
Top-1(%)	79.8	79.9 / 79.3	81.8	82.3 / 81.5	81.5	81.7
Throughput (img/s)	2773	4050 / 5523	1239	1792 / 2649	1940	2527

Table 6. Compare the throughput of our compressed models with baseline models, measured on one Nvidia A100 GPU.

Model	Method	Top-1	FLOPs
DeiT-S	Random search	76.4%	1.8G
	GP search (ours)	79.3%	1.8G

Table 7. Compare GP search with random search.

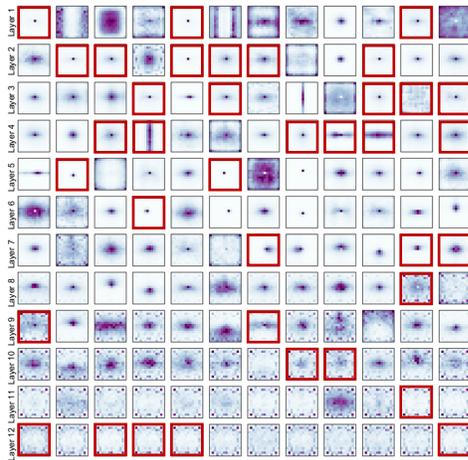


Figure 4. Visualization of the attention-maps (averaged over 256 images) produced by all heads in the DeiT-B model. Red box means the head is pruned based on our dependency criterion. Number of heads removed follow our pruning policy in Table 5.

from shallower layers while most of the sequence reductions occur at deeper layers. Moreover, deeper layers tend to have more redundancy, reflected by the increased layer-wise FLOPs reduction.

Actual inference speedup. We compare the throughput of our compressed models over baselines on a single Nvidia A100 GPU with a fixed batch size of 256. As shown in Table 6, our compressed models achieve $1.3\times \sim 2.2\times$ throughput improvement without significant accuracy loss.

Different pruning criteria. In Figure 3, we compare our dependency based pruning criterion with previous metrics for a specific pruning dimensions: SpAtten [27] for head pruning, attention probability [27] (AttenProb) for sequence reduction, and row-wise norm of the weight matrix (Magnitude) for neuron pruning. All criteria (including random selection) perform well when the pruning rate is small ($< 20\%$), suggesting that the redundancy indeed exists in each dimension. However, our dependency based pruning achieves relatively higher accuracy at larger pruning rates.

In Figure 4, by visualizing the attention-maps produced by all the heads in DeiT-B model, we observe that dependency based pruning indeed removes the redundant heads. In Figure 5, by visualizing the attention-maps produced by the top-ranking and bottom-ranking heads in the last block of the DeiT-B model, we see that our proposed dependency criterion can identify the heads that are more important to the model prediction.

GP search versus random search. Compared to random search [16] which determines the pruning policy by selecting the candidate with best validation accuracy from a random population, our GP search obtains better compressed model with 2.9% higher accuracy at the same FLOPs, as shown in Table 7.

GP search process is visualized in Figure 6. Our GP search firstly randomly samples 100 populations (pruning policies that satisfy the computation target) to fit GP model, which are shown on the left of the figure. On the right, the plot shows the validation accuracy of the pruning policy obtained by solving Eq.(4) at each search iteration. As the search process iterates, the obtained pruning policy gradu-

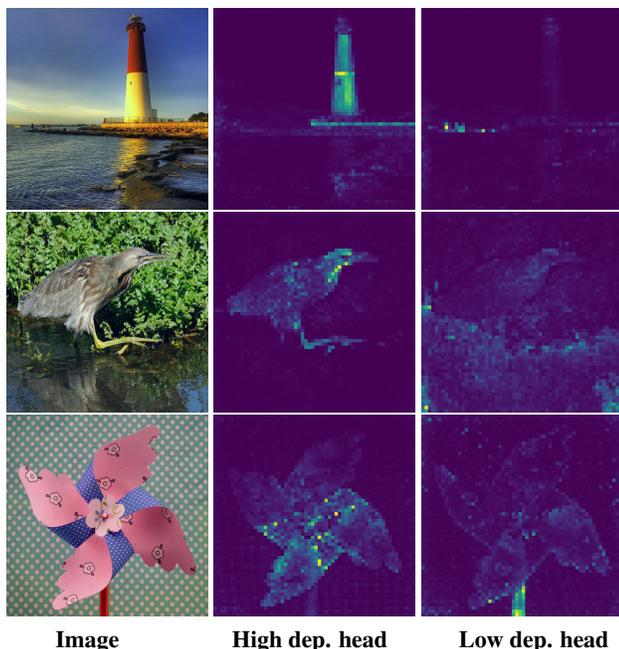


Figure 5. Examples of the attention-maps produced by top-ranking and bottom-ranking (in terms of our dependency based pruning criterion) attention head in the last block of DeiT-B model for difference input images.

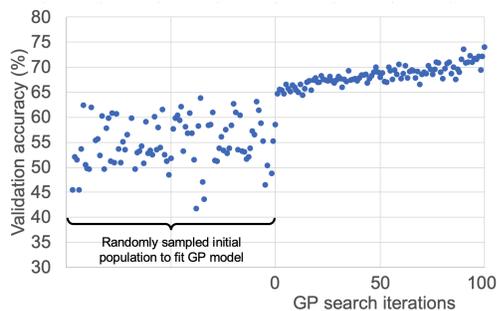


Figure 6. GP search process with DeiT-B model on ImageNet by plotting the validation accuracy over search iterations.

ally improves in terms of the top-1 accuracy until convergence. And the final policy achieves notably higher accuracy than the best result in the initial random sampling.

4.5. Results of object detection

For object detection, we apply our method to compress RetinaNet [18] with PVT-Small [30] backbone on COCO2017 dataset [19]. Following [30], models are trained on COCO train2017 (118k images) and evaluated on val2017 (5k images). We use the same training hyper-parameters as [30] to finetune the compressed model: AdamW optimizer with a batch size of 16, an initial learning rate of 1×10^{-4} , standard $1 \times$ training schedule (12

Method	FLOPs reduction	AP	AP ₅₀	AP ₅₀	AP _S	AP _M	AP _L
PVT-Small [30]	0%	38.7	59.3	40.8	21.2	41.6	54.4
Ours	30%	38.6	58.8	40.6	21.0	41.5	54.4

Table 8. Results of RetinaNet with PVT-Small backbone on COCO2017 data for object detection task.

epochs), learning rate is decayed by 10 at epoch 8 and 11. The training image is resized to have a shorter side of 640 pixels, and during testing the shorter side of the input image is fixed to 640 pixels.

PVT-Small is an efficient ViT models. It already has progressive sequence shrinking strategy and als performs Q-K reduction in the self-attention operators. Accordingly, we modify our sequence reduction strategy for PVT-Small. The token selection layer no longer discard the unimportant tokens (those with bottom-ranking dependency scores). Instead, they are directly skipped without participating any computations. The important tokens selected based on our dependency based pruning criterion participate the computations in the MHSA and FFN modules, and the output will be concatenated with the skipped unimportant tokens.

The results are shown in Table 8. Our compressed model achieves almost the same AP values compared to the base-line model, while yielding 30% FLOPs reduction.

5. Conclusion

In this paper, we present a novel ViT model compression framework which prunes a pre-trained ViT model from attention head, neuron, and sequence dimensions jointly. We propose a statistical dependency based pruning criterion based on the Hilbert-Schmidt norm of the cross-covariance operator in order to identify the deleterious features in different dimensions. Moreover, our framework learns the optimal pruning policy by casting multi-dimensional compression as a constrained non-linear optimization and using Gaussian process search with expected improvement to solve it. Our results on ImageNet with various ViT models outperform previous state-of-the-art ViT pruning methods under same computational budget.

Limitation. One limitation is that our method does not explicitly incorporate model depth into the compression space. Instead, depth reduction is implicitly covered since pruning all the neurons or heads would remove the layer except for the skip connection. However, we did not observe such case in our optimal pruning policy (cf. Figure 2), suggesting that layer removal may be too aggressive for current ViT models. Nevertheless, this study demonstrates the potential of jointly pruning multiple dimensions for accelerating ViT models, and we hope our results to be useful for future research.

References

- [1] Charles R Baker. Joint measures and cross-covariance operators. *Transactions of the American Mathematical Society*, 1973. 3
- [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] 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 *ICCV*, 2021. 2, 6
- [4] Minghao Chen, Houwen Peng, Jianlong Fu, and Haibin Ling. Autoformer: Searching transformers for visual recognition. In *ICCV*, 2021. 2, 6
- [5] Tianlong Chen, Yu Cheng, Zhe Gan, Lu Yuan, Lei Zhang, and Zhangyang Wang. Chasing sparsity in vision transformers: An end-to-end exploration. In *NeurIPS*, 2021. 1, 2, 5
- [6] Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019. 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 *CVPR*, 2009. 2, 5
- [8] 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*, 2021. 1
- [9] Arthur Gretton, Olivier Bousquet, Alex Smola, and Bernhard Schölkopf. Measuring statistical dependence with hilbert-schmidt norms. *ALT*, 2005. 3
- [10] Jinyang Guo, Wanli Ouyang, and Dong Xu. Multi-dimensional pruning: A unified framework for model compression. In *CVPR*, 2020. 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. In *ECCV*, 2020. 2, 5
- [12] Haoyu He, Jing Liu, Zizheng Pan, Jianfei Cai, Jing Zhang, Dacheng Tao, and Bohan Zhuang. Pruning self-attentions into convolutional layers in single path. *arXiv preprint arXiv:2111.11802*, 2021. 5
- [13] Zehao Huang and Naiyan Wang. Data-driven sparse structure selection for deep neural networks. In *ECCV*, 2018. 2
- [14] Ding Jia, Kai Han, Yunhe Wang, Yehui Tang, Jianyuan Guo, Chao Zhang, and Dacheng Tao. Efficient vision transformers via fine-grained manifold distillation. *arXiv preprint arXiv:2107.01378*, 2021. 6
- [15] Dieter Kraft et al. A software package for sequential quadratic programming. 1988. 4
- [16] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *UAI*, 2020. 2, 7
- [17] Shaohui Lin, Rongrong Ji, Chenqian Yan, Baochang Zhang, Liujuan Cao, Qixiang Ye, Feiyue Huang, and David Doermann. Towards optimal structured cnn pruning via generative adversarial learning. In *CVPR*, 2019. 2
- [18] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *CVPR*, 2017. 8
- [19] 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. 8
- [20] Zhenhua Liu, Yunhe Wang, Kai Han, Wei Zhang, Siwei Ma, and Wen Gao. Post-training quantization for vision transformer. In *NeurIPS*, 2021. 2
- [21] Zechun Liu, Xiangyu Zhang, Zhiqiang Shen, Zhe Li, Yichen Wei, Kwang-Ting Cheng, and Jian Sun. Joint multi-dimension pruning. *arXiv preprint arXiv:2005.08931*, 2020. 2
- [22] Bowen Pan, Rameswar Panda, Yifan Jiang, Zhangyang Wang, Rogerio Feris, and Aude Oliva. Ia-red²: Interpretability-aware redundancy reduction for vision transformers. In *NeurIPS*, 2021. 1, 2, 5
- [23] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, 2021. 1, 2, 5
- [24] Carl Edward Rasmussen. Gaussian processes in machine learning. 2006. 4
- [25] 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. 1, 2, 5
- [26] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. 1, 2, 5
- [27] Hanrui Wang, Zhekai Zhang, and Song Han. Spatten: Efficient sparse attention architecture with cascade token and head pruning. In *HPCA*, 2021. 2, 7
- [28] Sinong Wang, Belinda Z Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020. 1
- [29] Wenxiao Wang, Minghao Chen, Shuai Zhao, Long Chen, Jinming Hu, Haifeng Liu, Deng Cai, Xiaofei He, and Wei Liu. Accelerate cnns from three dimensions: A comprehensive pruning framework. In *ICML*, 2021. 2
- [30] 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. 8
- [31] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *NeurIPS*, 2016. 2
- [32] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nystromformer: A nystrom-based algorithm for approximating self-attention. In *AAAI*, 2021. 1
- [33] 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. *arXiv preprint arXiv:2108.01390*, 2021. 1, 2, 5

- [34] Huanrui Yang, Hongxu Yin, Pavlo Molchanov, Hai Li, and Jan Kautz. Nvit: Vision transformer compression and parameter redistribution. *arXiv preprint arXiv:2110.04869*, 2021. [6](#)
- [35] Hao Yu and Jianxin Wu. A unified pruning framework for vision transformers. *arXiv preprint arXiv:2111.15127*, 2021. [6](#)
- [36] Shixing Yu, Tianlong Chen, Jiayi Shen, Huan Yuan, Jianchao Tan, Sen Yang, Ji Liu, and Zhangyang Wang. Unified visual transformer compression. In *ICLR*, 2022. [5](#)
- [37] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis E.H. Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *ICCV*, 2021. [2](#), [5](#)
- [38] Mingjian Zhu, Kai Han, Yehui Tang, and Yunhe Wang. Visual transformer pruning. In *KDD*, 2021. [2](#), [5](#)