

Isomorphic Pruning for Vision Models

Gongfan Fang¹, Xinyin Ma¹, Michael Bi Mi², and Xinchao Wang^{1*}

¹ National University of Singapore

² Huawei Technologies Ltd.

`gongfan@u.nus.edu`, `xinchao@nus.edu.sg`

Abstract. Structured pruning reduces the computational overhead of deep neural networks by removing redundant sub-structures. However, assessing the relative importance of different sub-structures remains a significant challenge, particularly in advanced vision models featuring novel mechanisms and architectures like self-attention, depth-wise convolutions, or residual connections. These heterogeneous substructures usually exhibit diverged parameter scales, weight distributions, and computational topology, introducing considerable difficulty to importance comparison. To overcome this, we present *Isomorphic Pruning*, a simple approach that demonstrates effectiveness across a range of network architectures such as Vision Transformers and CNNs, and delivers competitive performance across different model sizes. Isomorphic Pruning originates from an observation that, when evaluated under a pre-defined importance criterion, heterogeneous sub-structures demonstrate significant divergence in their importance distribution, as opposed to isomorphic structures that present similar importance patterns. This inspires us to perform isolated ranking and comparison on different types of sub-structures for more reliable pruning. Our empirical results on ImageNet-1K demonstrate that Isomorphic Pruning surpasses several pruning baselines dedicatedly designed for Transformers or CNNs. For instance, we improve the accuracy of DeiT-Tiny from 74.52% to 77.50% by pruning an off-the-shelf DeiT-Base model. And for ConvNext-Tiny, we enhanced performance from 82.06% to 82.18%, while reducing the number of parameters and memory usage. Code is available at <https://github.com/VainF/Isomorphic-Pruning>.

Keywords: Network Pruning · Vision Transformers · CNNs

1 Introduction

Deep Neural Networks are typically over-parameterized, containing numerous redundant sub-structures. Such redundancy can be effectively eliminated via structured pruning, without significantly sacrificing the network’s performance [16, 22, 25]. Pruning usually follows a “ranking-to-prune” framework: To facilitate compression while retaining the learned capabilities of networks, a carefully-designed criterion is utilized to assess the importance score of different structures [11, 22, 40], followed by a ranking process to identify and remove those

* Corresponding Author.




Fig. 1: ImageNet Top-1 Accuracy vs. Multiply-Accumulate Operations (MACs) of pruned DeiT [50] and pruned ConvNext [33]. The pruned models, marked as “★” have comparable or better latency yet superior performance compared to scratch training counterparts highlighted by “◇”. The MACs are in log scale for better visualization

relatively unimportant ones. After pruning, the compressed model is fine-tuned to recover its performance [11, 18].

However, ranking-based structured pruning presupposes an underlying assumption that the parameters from different structures are directly comparable. Classic deep neural networks, such as VGG [46], are built by stacking similar convolutional layers or blocks where the neuron connections share a similar computational topology. Under the circumstances, a straightforward global ranking is expected to reveal their relative importance and remove unimportant parameters [11, 22]. However, with recent innovation in network structure design, advanced vision models like Vision Transformers [10, 31, 50] and ConvNext [33] incorporate more sophisticated designs, such as complicated residual connections [19, 24], attention mechanisms [52] and depth-wise convolutions [6]. These novel structures increase the network’s internal complexity and heterogeneity, challenging the reliability of pruning importance criteria. For instance, a vanilla Vision Transformer [10] is composed of MLP, Multi-head Attention, and skip connections. Even in the same Multi-head attention layer, there are two pruning schemes to removing either entire attention heads [37, 53] or the within-head dimensions [58, 68]. These structures vary by their computational topologies and, as will be illustrated in the experiments, tend to exhibit diverged importance distributions. In this case, directly comparing the importance of these heterogeneous might be problematic.

To address the problem, we propose a simple pruning approach, termed Isomorphic Pruning, to alleviate the incomparability issue arising from heterogeneous sub-structures. The essence of this method lies in decomposing a network into multiple isomorphic groups, as illustrated in Figure 2 (d), based on the computational topology of the sub-structures. Each isomorphic group contains




Fig. 2: For a pre-trained network (a), we show three pruning strategies: (b) Local Pruning that compares parameter importance within current layers; (c) Global Pruning that performs global ranking for all parameters; (d) the proposed Isomorphic pruning that groups parameters by the computational topology and applies importance ranking within groups. In each group, the importance distributions are more comparable. Details about the distribution can be found in Fig. 4 of the experiments section.

elements with the same structure and tends to present a similar distribution of importance, making it more rational to compare their relative importance.

To validate the effectiveness of our method, we conducted experiments on a variety of visual models, including DeiT [50], ConvNext [33], ResNet [19], and MobileNet-V2 [44]. Pruning results, as shown in Figure 1, on ImageNet-1K [8] demonstrate that Isomorphic Pruning achieves comparable or even superior performance relative to pruning algorithms specifically designed for CNNs or transformers. For instance, by pruning a pre-trained large DeiT model, we enhanced the performance of the DeiT-T model from 74.52% to 77.50%, when compared to the scratch training one. Similarly, our algorithm also yields satisfactory results on CNNs. We constructed compact ConvNexts and ResNets via pruning, which has fewer MACs and parameters while surpassing a conventionally trained model. Additionally, we also studied the real-world latency and memory footprint of the compressed models on both GPU and CPU devices. Our findings indicate that models pruned with Isomorphic Pruning achieve significant actual speed-ups and are competitive to carefully designed efficient architectures.

Therefore, Our contribution lies in a practical and general pruning approach, which can handle various vision models with novel mechanisms and architectures. The pruned models show superior performance to several pruning competitors on ImageNet-1K benchmark.

2 Related Works

Network Pruning. Deep neural networks are typically over-parameterized and can be pruned without significant performance drop [16, 25]. The exploration of network pruning can be dated back to early works [1, 17, 25] about neural networks. Over the past few years, a large number of works have been conducted on convolutional networks to verify the effectiveness of pruning, which can be roughly categorized into unstructured [16, 29, 35] and structured pruning [5, 11, 18, 22, 56]. Unstructured pruning compresses a pre-trained model by forcing partial weights to be zero, leading sparse networks [29]. To accelerate

the sparse model, special patterns and devices will be used such as N:M sparsity [12, 42, 48]. Structured pruning provides a hardware-friendly approach by removing sub-structures such as channels [11] or blocks [14] in CNNs, which improves the model efficiency without dedicated accelerators. In recent years, a large number of works adapted the idea of network pruning to vision transformers. It demonstrated that even the vanilla vision transformers can be squeezed for a better efficiency-performance trade-off [2, 4, 45, 47, 58, 60–62, 64, 65, 68, 70]. In the literature, transformer pruning has been studied from several perspectives, such as width reduction [70], depth pruning [61], attention head removal [58], sparse training [4] and architecture transformation [18].

Pruning by Global Ranking. In the context of network pruning, global pruning is recognized as a pivotal topic for exploration. This method entails a comprehensive ranking of all network layers, assigning adaptive pruning ratios to each layer according to the ranks [5, 11, 22, 58]. Different layers within a network often exhibit varying levels of redundancy. To this end, the advantage of global pruning lies in its ability to create a more parameter-efficient, non-uniformly compressed model [58]. However, a key challenge remains in the ranking: how to ensure the consistency and comparability across different layers? A widely-used solution for this issue is normalizing scores on a per-layer basis to even out the importance distribution across the network [30]. Alternatively, some approaches use sparse training to learn an effective target structure [11], which might be costly due to the additional optimization. In this work, we introduce a more straightforward approach termed isomorphic pruning to facilitate a reliable ranking for pruning.

3 Method

The core idea of Isomorphic Pruning is straightforward: we group sub-structures in a network by their network topology and perform importance ranking within groups. This imposes two questions 1) how to identify removable sub-structures in networks and 2) how to compare the topology graph for grouping and pruning.

3.1 Sub-structure in Network

Structured Pruning aims to remove sub-structures in networks, such as the dimensions of hidden states or features. To make the formulation clear, we primarily concentrate on stacked fully-connected networks, although the principles established herein can be readily extended to CNNs and transformers. Initially, we consider a single linear layer which is parameterized by a matrix $W \in \mathbb{R}^{m \times n}$. To articulate the pruning process, we introduce a pruning function $g(W, d, k)$ defined as follows:

$$\hat{w} = g(W, d, k) = \begin{cases} g(W, 0, k) = W_{[k,:]} \in \mathbb{R}^{1 \times n}, & \text{if pruning output dimension,} \\ g(W, 1, k) = W_{[:,k]} \in \mathbb{R}^{m \times 1}, & \text{if pruning input dimension,} \end{cases} \quad (1)$$

where \hat{w} refers to the weight column or row slated for removal, acquired via a slicing operation denoted by subscript $[:,k]$ or $[k,:]$. The pruning function $g(W, d, k)$ is defined by three parameters, the parameter matrix W , the axis $d \in \{0, 1\}$ indicating the dimension to be compressed, and the index k identifying the k -th rows or columns in W . In the context of a linear layer represented by a 2-D matrix $W \in \mathbb{R}^{m \times n}$, the axis d can assume values of either 0 or 1, with k having a permissible range of $[0, m]$ and $[0, n]$, respectively. This formulation encapsulates the pruning of a single layer; to extend it to the entire network, we also need to take the dependency [5, 11, 30] into account, where parameters from different layers are coupled and should be pruned simultaneously. As depicted in Figure 3, for two adjacent linear layers parameterized by W_1 and W_2 , the output dimension of the former is integrally coupled to the input dimension of the latter, thereby constituting a minimal sub-structure eligible for removal. If we treat each pruning operation $g(W, d, k)$ as vertices defined by a triplet (W, d, k) and the dependency as edges, then it is natural to model sub-structures as graphs.

3.2 Graph Modeling of Sub-structures

In this section, we delve deeper into the methodology for decomposing a network into several independent and minimal sub-structures in pruning, and model them as graphs. We can begin with the MLP network as illustrated in Figure 3, where the activation functions are omitted. The network comprises a sequence of 2-D parameter matrices $\{W_1, W_2, \dots, W_L\}$, with L denoting the network depth. Initially, all parameters are not assigned to any sub-structure. We randomly select one parameter W as the seed to find its associated substructures and build the corresponding graph. For instance, selecting parameter W_1 and setting the axis $d = 0$ allows us to establish an initial graph $G = (V = \{(W, 0, k)\}, E = \{\})$, focusing on the elimination of k -th the output dimension in W_1 . Then we iterate through all potential $(W', d', k') \notin V$ to ascertain any dependency to $(W, d, k) \in V$. Dependencies between layers can be ascertained through either manually-crafted patterns or automatic algorithms. In this work, we adhere to the approach of [11] to assess dependency based on two rudimentary rules:

Dependency Modeling. For a pair of triplets $(W', d', k'), (W, d, k)$, a dependency is established if either of the following criteria is met: 1) W' and W are adjacent layers with $d = d'$ and $k = k'$; 2) $W' = W$ and $g(W', d', k') = g(W, d, k)$. The first criterion addresses adjacent layers as exemplified in Figure 3, while the second pertains to specialized layers, such as normalization layers, which share the same pruning pattern for both outputs and inputs. Thereafter, we update the vertex set V and edge set E with the discovered dependency between (W', d', k') and (W, d, k) accordingly:

$$\begin{aligned} V &= V \cup \{(W', d', k')\}, \\ E &= E \cup \{((W, d, k), (W', d', k'))\}. \end{aligned} \quad (2)$$

We can further repeat the above process until no additional vertex is incorporated into the current graph. The resultant graph G then serves as a topological

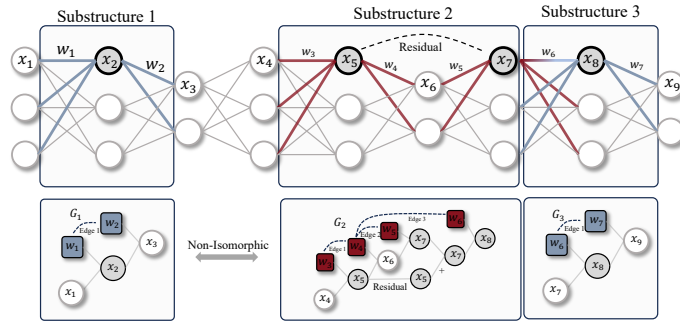


Fig. 3: Isomorphic Pruning models sub-structures as graphs, and apply isolated ranking and pruning with isomorphic structures. We highlight three removable substructures in an MLP and show their corresponding graphs below. Substructures 1 & 3 are isomorphic but 2 & 3 are non-isomorphic due to the additional residual connections.

representation of a single sub-structure associated with the initial seed parameter. To map all substructures onto different graphs, we subsequently embark on a recursive procedure, commencing from those nodes not yet allocated to any sub-structure. Figure 3 furnishes an exemplar of this process, showcasing three distinct substructures labeled as sub-structures 1, 2, and 3.

3.3 Ranking with Graph Isomorphism

With the graph-based representation of network sub-structures, we advance to the ranking phase. This stage involves estimating and ranking the significance of each sub-structure to identify redundancies. A pivotal component of this process is an importance criterion, which quantifies the relative significance of various sub-structures. Previous works in the literature presents numerous well-crafted criteria, including magnitude pruning and Taylor expansion. In this work, we focus on the general importance function $I(\cdot)$ for a single layer, represented as:

$$I(W, d, k) = \begin{cases} \|g(W, d, k)\|, & \text{if Magnitude Pruning} \\ \left\| \frac{\partial \mathcal{L}}{\partial g(W, d, k)} g(W, d, k) \right\|_2, & \text{if Taylor Pruning} \\ \dots & \text{Other importance Criteria.} \end{cases} \quad (3)$$

In structural pruning, our objective is to remove sub-structures encompassing parameters across multiple layers. Existing approaches [5, 11, 30] generally extend the single-layer importance to a holistic measure by aggregating scores across all constituents (W, d, k) for the vertexes set, denoted as $I^*(\cdot)$ as follows:

$$I^*(G(V, E)) = \sum_{(W, d, k) \in V} I(W, d, k), \quad (4)$$

thereby cumulating the importance of all elements within graph G . However, the inherent diversity in parameters and computational topology among different

substructures G introduces potential ambiguities in this metric. Experimental evidence, particularly in the context of a vision transformer as shown in Figure 4, elucidates this issue: the largest substructure encompasses all embedding dimensions, containing a substantial number of parameters, whereas the most minimal sub-structure comprises merely two linear layers. Consequently, such aggregation may fail to accurately reflect the true relative importance of these entities. To address the challenges posed by the heterogeneity of substructures, we delineate a simple approach designed to preserve the nuanced information encapsulated by the aggregated importance, as outlined in Equation 4, whilst alleviating the detrimental repercussions raised by the heterogeneity of substructures. Our principal strategy entails the clustering of sub-structures into various isomorphic groups, predicated on their topological logic and parameter configurations. In this case, applying aggregated importance within groups can be meaningful, since isomorphic sub-structures show one-to-one correspondence among their parameters, and how different parameters participate in computations is entirely identical. To facilitate this, the key step is to examine the isomorphism of two graphs G_1 and G_2 . Formally, we adopt the following definition to ascertain graph isomorphism.

Definition 1 (Graph Isomorphism). *Two labeled graphs G and G' are isomorphic if there exists a bijection between their vertices so that, all mapped vertices preserve the same labels, and are connected in the same way.*

To develop a general indicator $\text{Isomorphic}(g_1, g_2)$, we need to examine both the label and connectivity in two graphs.

Vertex Label and Connectivity. In our settings, each vertex is labeled by the triplet (W, d, k) , encompassing the parameter matrix \mathbf{W} , axis d , and index k . We follow a simple principle for labling: Two vertices (\mathbf{W}, d, k) and (\mathbf{W}', d', k') are labeled identically if and only if the parameter \mathbf{W}' originates from the same layer type such as Linear or Conv, with the same pruning dimension $d = d'$. We permit the index k to differ, since within the same parameter matrix, the weight vectors sliced along the d -th dimension are consistently homogeneous and thus can be considered isomorphic. After determining the label for each vertex, we further explore the connectivity in two graphs G_1 and G_2 . As illustrated in Figure 3, we show three substructures alongside their respective graphs G_1 , G_2 , and G_3 , where G_1 and G_3 incorporate a singular edge, whereas G_2 encompasses three edges. If we check the edges following the execution order of the computational graphs, it’s natural to observe the isomorphism of G_1 and G_3 , since they share the same topology. To make this process more formal, we assume the edges in graph G are already sorted according to their execution order in forwarding, and check the isomorphism by

$$\text{Isomorphic}(G, G') = \mathbb{1}\{|G| = |G'| \wedge \text{label}(V_i) = \text{label}(V'_i); \forall i \in \{1, 2, \dots, |G'|\}\}, \quad (5)$$

where $\mathbb{1}$ is an indicator function. $|G'|$ denotes the number of edges in the graph and \wedge signifies the logical AND operation. The first condition ensures the equiva-

Algorithm 1 Isomorphic Pruning

```

1: Input: Network with parameters  $\{W_1, W_2, \dots, W_n\}$ , Pruning ratio  $p\%$ .
2: Output: Pruned Network  $\{W'_1, W'_2, \dots, W'_n\}$ 
3: procedure IDENTIFYSUBSTRUCTURES( $\{W_1, W_2, \dots, W_n\}$ )
4:   Initialize empty list of sub-structures  $\mathbf{S} = \{\}$ 
5:   for each layer  $W_i$  in  $\{W_1, W_2, \dots, W_n\}$  do
6:     if  $W_i$  is not assigned to any substructure then
7:       Initialize a graph  $G(V = \{W_i\}, E = \{\})$ .
8:       Find and add coupled  $W_j$  to  $G$  until there is no new dependency.
9:        $\mathbf{S} = \mathbf{S} \cup \{G\}$ 
10:  return  $\mathbf{S}$ 
11: procedure PRUNEBYGROUPEDRANKING( $\mathbf{S}, p$ )
12:  for each substructure  $G_i$  in  $\{G_1, G_2, \dots, G_n\}$  do
13:    if  $G_i$  is not assigned to any group then
14:      Initialize a Group  $\mathbf{R} = \{G_i\}$ .
15:      Find and add  $G_j$  that satisfies  $\text{Isomorphism}(G_i, G_j) = 1$  to  $\mathbf{R}$ .
16:      Calculate the importance scores  $\{I^*(G) | G \in \mathbf{R}\}$ .
17:      Remove  $p\%$  substructures in  $\mathbf{R}$  by ranking the importance scores.
18:  return  $\mathbf{S}$ 
19:  $\mathbf{S} \leftarrow \text{IDENTIFYSUBSTRUCTURES}(\{W_1, W_2, \dots, W_n\})$ 
20:  $\mathbf{S} \leftarrow \text{PRUNEBYGROUPEDRANKING}(\mathbf{S}, p)$ 

```

lence of the number of vertices in the graphs, while the second condition confirms the matching labels of vertices on either side of the edge pairs V_i and V'_i .

Using Equation 5, it is easy to cluster substructures into different groups and perform rankings within each group exclusively. Specifically, we examine a collection of substructures $\{G_1, G_2, \dots, G_N\}$ that share identical graph topologies. We estimate the aggregated significance of each graph independently using Equation 4, resulting in an N -dimensional importance vector $[I^*(G_1), \dots, I^*(G_N)]$. This facilitates a straightforward ranking of the importance vector, followed by the elimination of the $p\%$ unimportant substructures based on their scores. This ranking method is deemed reliable since all substructures within a group possess the same architectural design, parameter scale, and computational graph. A similar principle can be uniformly applied to other isomorphic groups for pruning, as detailed in Algorithm 1.

4 Experiments

This section reports our empirical results on the ImageNet-1K dataset [8], which comprises 1,281,167 images for training and 50,000 for validation. We deploy the proposed Isomorphic Pruning to compress several vision models, including ConvNext [33], ResNet [36], MobileNetv2 [44], and Vision Transformers [10, 50]. We report the parameter amount (#Params), MACs, and Top-1 validation accuracy in Table 1 and 2, and report the actual latency and peak memory consumption in Table 3.

4.1 Settings for ImageNet-1K

Pruning. For simplicity, we adopt the Taylor-based pruning framework [39], which applies a data-driven Taylor expansion to the loss function. For this purpose, we randomly sampled 100 mini-batches from the ImageNet dataset, each containing 64 images. These images were resized to 256x256 pixels and then center-cropped to 224x224 pixels. We normalize all images using the mean and standard deviation specific to ImageNet. Notably, while some methods, such as ConvNext such as ConvNext [33] and DeiT [50] employ additional training techniques like label smoothing and knowledge distillation, our approach only leverages a straightforward cross-entropy loss for the Taylor expansion. We accumulate gradients from all the batches for a more reliable importance estimation and apply Isomorphic pruning in a one-shot manner without iterative pruning and fine-tuning [40].

Finetuning. During the fine-tuning process of pruned models, we adhere to the original training protocol for a fair comparison. While maintaining most of the original settings unchanged, we adjusted the learning rates and batch sizes to make the training suitable for A5000 GPUs, following the linear LR scaling rule proposed in [13]. For example, to compressing a pre-trained DeiT models [50], we adopt the same augmentation strategies for training, including RandAugmentation [7], Mixup [67], CutMix [66], Random Erasing [69], Repeated Augmentation [23], label smoothing [49]. The pruned models are distilled with a pre-trained RegNetY [43, 50] for 300 epochs using the AdamW [34] optimizer, with a learning rate of 0.0005 and a total batch size of 2048. A weight decay of 0.05 and a cosine annealing scheduler is deployed for training. For other models like ConvNext [33], ResNet [19] and MobileNet-v2 [44], the same principle will be applied. More details of hyper-parameters, pruning, and finetuning can be found in the appendix.

Evaluation. We report the optimal classification accuracy achieved on the ImageNet-1K validation set, employing the standard resize-and-crop protocol [19, 33] for evaluation. To further validate the efficiency of models, we profiled the latency and the peak memory consumption of the pruned models on both GPU (a single RTX A5000) and CPU. To guarantee a fair and standardized comparison, we also tested the latency of several baseline models sourced from the Pytorch-Image Models [57] and Torchvision [38], following the same testing procedure. The latency was estimated with a batch size of 256 for GPU and 8 for CPU, averaged on 100 repeats. Appropriate GPU warmup was deployed in our test.

4.2 Pruning Convolutional Neural Networks

Convolutional Neural Network is a critical focus in the area of network pruning. In this section, we concentrate on pruning three popular ConvNets: ConvNext [33], ResNet [19] and MobileNet-v2 [44].

Table 1: Pruning CNNs on ImageNet-1K.

Pruning Method	#Params (M)	MACs (G)	Base Acc.	Final Acc.	Δ Acc.
ConvNext-B [†] [33]	88.59	15.39	-	83.83	-
ConvNext-S [†] [33]	50.22	8.71	-	83.14	-
ConvNext-S (Ours)	47.36	8.48	83.83	83.17	-0.66
ConvNext-T [†] [33]	28.59	4.47	-	82.06	-
ConvNext-T (Ours)	25.32	4.19	83.83	82.19	-1.64
ResNet-152 [†]	60.19	11.58	-	78.31	-
AOFPE2 [9]	-	4.13	77.37	77.00	-0.37
ABCPruner [28]	24.07	4.31	78.31	77.12	-1.19
EPruner-0.63 [27]	21.56	4.08	78.31	76.83	-1.48
ResNet-152-4.0G (Ours)	23.11	4.05	78.31	77.84	-0.47
ResNet-101 [†]	44.55	7.85	-	77.38	-
IE [39]	31.2	4.70	77.37	77.35	-0.02
FPGM [21]	-	4.51	77.37	77.32	-0.05
SFP [20]	-	4.51	77.37	77.51	+0.14
ISTA [59]	-	4.47	76.40	75.27	-1.13
Res101-4.5G (Ours)	29.14	4.48	77.38	77.56	+0.16
GFP [30]	28.02	3.90	78.29	78.33	+0.04
AOFPE [9]	-	3.89	76.63	76.40	-0.23
Res101-3.8G (Ours)	24.87	3.85	77.38	77.43	+0.05
ResNet-50 [†] [19]	25.56	4.13	-	76.13	-
Taylor [39]	14.20	2.25	76.18	74.50	-1.68
CCP [41]	-	2.11	76.15	75.50	-0.65
GFP [30]	19.42	2.04	76.79	76.42	-0.37
AutoSlim [63]	20.60	2.00	76.10	75.60	-0.50
DepGraph [11]	-	1.99	76.15	75.83	-0.32
ResNet50-2G (Ours)	15.05	2.06	76.13	75.91	-0.22
Mobv2 [†]	3.50	0.32	-	71.88	-
Meta [32]	-	0.14	74.70	68.20	-6.50
GFP [30]	-	0.15	75.74	69.16	-6.58
DepGraph [11]	-	0.15	71.90	68.49	-3.41
Mobv2-0.15G (Ours)	2.19	0.15	71.88	68.91	-2.97

ConvNext. Our experiments in Table 1 begin with the evaluation on an advanced CNN, ConvNext [33], which achieves impressive performance on several benchmarks. We use the official ConvNext-Base [38] as the pre-trained model for pruning, which achieves an accuracy of 83.83% on the ImageNet-1K validation set. The pruning of ConvNext models is a relatively underexplored topic. Hence, we establish our baselines using the official variants with fewer parameters, ConvNext-S and ConvNext-T. For ConvNext, we also generalize our method to depth pruning, which removes unimportant layers with the minimal average scores on substructures, so that the pruned models have the same depth as baselines. The pruned models are fine-tuned for 300 epochs following the same training protocol as [33]. During fine-tuning, we observed that the pruned ConvNext with half of the parameters removed, achieves an accuracy of 80.09 in only 20 epochs, which effectively validates the efficiency of pruning. Moreover, as illustrated in Table 1, the final model after 300 full fine-tuning, achieves competitive performance (82.19% vs. 82.06%) to pre-trained models with fewer parameters (25.32M vs. 28.59M) and MACs (4.19G vs. 4.47G).

ResNet. In line with prior studies [5, 11, 22, 30], we demonstrate the efficacy of our method by structurally pruning pre-trained ResNet MACs. ResNet-

Table 2: Pruning DeiT_s on ImageNet-1K.

Method	#Params (M)	MACs (G)	Acc (%)	Method	#Params (M)	MACs (G)	Acc (%)
DeiT-B [†] [50]	87.34	17.69	83.32	DeiT-T [†] [50]	5.91	1.27	74.52
DeiT-S [†] [50]	22.44	4.64	81.17	DeiT-T-600EP [58]	5.91	1.27	75.00
DeiT-S-600EP [58]	22.44	4.64	81.80	WDPruning [61]	13.30	2.60	78.38
VTP - 40% [70]	48.00	10.00	80.70	X-Pruner [64]	-	2.40	78.93
WDPruning [61]	55.30	9.90	80.76	UVC [65]	-	2.32	78.82
UVC [65]	-	8.00	80.57	ViT-Slim [70]	11.40	2.30	77.94
X-Pruner [64]	-	8.50	81.02	ViT-Slim [70]	11.40	2.30	77.94
NViT-B [58]	34.00	6.80	83.29	NViT-T [58]	6.90	1.30	76.21
SAViT [68]	25.40	5.30	81.66	SAViT+KD [68]	6.60	1.30	76.95
UP-DeiT-S [62]	22.10	-	81.56	UP-DeiT-T [62]	5.70	-	75.79
NViT-S [58]	21.00	4.20	82.19	DeiT-T (Ours)	5.74	1.21	77.50
U _{pop} [45]	19.90	4.10	81.10	VTC-LFC [55]	5.10	-	71.60
DeiT-S (Ours)	20.69	4.16	82.41	SViTE-T [68]	4.21	0.99	70.12
GOHSP [60]	14.40	2.99	79.98	GOHSP [60]	4.00	0.91	70.24
GOHSP [60]	11.10	2.81	79.86	SAViT [68]	4.20	0.90	70.72
ViT-Slim [2]	13.50	2.80	79.20	CP-ViT [47]	-	0.74	71.24
U _{pop} [45]	13.50	2.80	79.60	WDPruning [61]	3.50	0.70	70.34
ViT-Slim [2]	13.50	2.80	79.20	X-Pruner [64]	-	0.60	71.10
CP-ViT [47]	-	2.66	79.08	UVC [65]	-	0.51	70.60
DeiT-2.6G (Ours)	13.07	2.62	81.13	DeiT-0.6G (Ours)	3.08	0.62	72.60

50/101/152 consists of four residual blocks, each formed by several Bottle-neck structures and residual connections [19]. The proposed Isomorphic Pruning identifies five kinds of distinct isomorphic groups within ResNet-50, four targeting the block structures and one associated with bottleneck structures [19]. It is worthwhile to note that different sub-structures vary in their parameter and MAC counts. For example, residual blocks naturally contain more parameters than the bottleneck structures which are only composed of a few convolutional layers. The conventional global pruning, employing a naive global ranking threshold, faces the potential risks of over-pruning certain layers due to the biased importance distribution. Isomorphic Pruning, however, employs a more stable scheme for pruning. As illustrated in Table 1, we report both the accuracy and the accuracy drop compared to their base models. Isomorphic pruning achieves lossless compression on ResNet-101, under both 4.5G and 3.8G settings.

MobileNet-v2. We also conduct experiments on lightweight models, MobileNet-v2, by compressing its MACs from 0.32 G to 0.15 G. In MobileNet-v2, we can also discover 5 isomorphic groups in MobileNet. The pruned model achieves better accuracy compared to DepGraph [11] which requires sparse training.

4.3 Pruning Vision Transformers.

Vision Transformers (ViT) has profoundly reshaped the paradigm of visual modeling over the past few years. Nevertheless, their inherent over-parameterization often imposes huge overheads to both training and inference. Compared to CNNs that are composed of similar convolution layers or blocks, Vision Transformers are usually built with MLP, Multi-head Attention, and skip connections, which are more challenging for pruning due to the heterogeneous internal structures.

Table 2 presents our pruning results on DeiT, a plain vision transformer enhanced with knowledge distillation [50]. We benchmark our approach with

Table 3: Profiling the pruned models and other off-the-shelf vision models from PyTorch-Image-Models [57] on a single RTX A5000 GPU and CPU. For NViT [58], we report their best speed up in the original paper, tested on a single V100 or RTX 3080.

Method	#Params ↓	MACs ↓	Peak Mem. ↓	GPU Latency (256) ↓	CPU Latency (8) ↓	Acc. (%)
DeiT-B [50]	86.40 M	17.60 G	2909.78 MB	802.97 ms (1.00×)	197.73 ms (1.00×)	83.32
DeiT-S [50]	22.44 M	4.64 G	1363.47 MB	241.81 ms (3.32×)	80.64 ms (2.45×)	81.17
Swin-T [31]	28.29 M	4.51 G	4257.98 MB	375.45 ms (2.14×)	203.86 ms (0.97×)	81.19
NViT-S [58]	21.00 M	4.20 G	-	- (2.52×)	-	82.19
TNT-S [15]	23.76 M	4.85 G	1457.56 MB	631.02 ms (1.27×)	175.58 ms (1.13×)	81.50
CaiT-XS24 [51]	26.56 M	5.40 G	1496.42 MB	557.41 ms (1.44×)	342.28 ms (0.58×)	82.00
CrossViT-S [3]	26.86 M	5.64 G	1546.94 MB	315.02 ms (2.55×)	115.51 ms (1.71×)	81.00
EfficientFormer-L3 [26]	31.41 M	3.94 G	2681.75 MB	249.05 ms (3.22×)	167.16 ms (1.18×)	82.40
PVTv2-B2 [54]	25.36 M	4.00 G	4478.18 MB	364.00 ms (2.21×)	212.44 ms (0.93×)	82.00
DeiT-S (Ours)	20.69 M	4.16 G	1547.97 MB	230.84 ms (3.48×)	44.10 ms (4.48×)	82.41
DeiT-2.6G (Ours)	13.07 M	2.62 G	1203.07 MB	189.68 ms (4.23×)	56.59 ms (3.49×)	81.13
DeiT-T [50]	5.91 M	1.27 G	720.22 MB	93.78 ms (8.56×)	46.83 ms (4.22×)	74.52
NViT-T [58]	6.90 M	1.30 G	-	- (4.97×)	-	76.21
DeiT-T (Ours)	5.71 M	1.20 G	837.16 MB	93.17 ms (8.62×)	47.89 ms (4.13×)	77.50
DeiT-0.6G (Ours)	3.08 M	0.62 G	629.28 MB	63.23 ms (12.70×)	40.98 ms (4.83×)	72.60
ConvNext-B [33]	88.59 M	15.36 G	5394.10 MB	789.24 ms (1.00×)	262.58 ms (1.00×)	83.83
ConvNext-S [33]	50.22 M	8.69 G	3958.45 MB	508.20 ms (1.57×)	188.07 ms (1.40×)	83.14
ConvNext-S (Ours)	47.36 M	8.48 G	3788.65 MB	535.95 ms (1.59×)	179.72 ms (1.46×)	83.17
ConvNext-T [33]	28.59 M	4.46 G	3785.37 MB	291.02 ms (2.82×)	126.17 ms (2.08×)	82.06
ConvNext-T (Ours)	25.32 M	4.19 G	3607.06 MB	292.77 ms (2.83×)	97.82 ms (2.68×)	82.19
ResNet152 [19]	60.19 M	11.58 G	3320.24 MB	386.65 ms (1.00×)	272.21 ms (1.00×)	78.31
ResNet152-4G (Ours)	23.11 M	4.06 G	1938.91 MB	238.12 ms (1.62×)	180.65 ms (1.51×)	77.84
ResNet101 [19]	44.55 M	7.85 G	3196.77 MB	271.34 ms (1.00×)	202.66 ms (1.00×)	77.38
ResNet101-4.5G (Ours)	29.14 M	4.48 G	2639.40 MB	220.72 ms (1.23×)	162.33 ms (1.25×)	77.56
ResNet101-3.8G (Ours)	24.87 M	3.85 G	2286.36 MB	194.39 ms (1.40×)	158.98 ms (1.28×)	77.43
ResNet50 [19]	25.56 M	4.13 G	3035.84 MB	171.54 ms (1.00×)	114.86 ms (1.00×)	76.13
ResNet50-2.0G (Ours)	15.05 M	2.06 G	2238.15 MB	78.75 ms (1.33×)	80.88 ms (1.42×)	75.91
Mobv2 [44]	3.50 M	0.32 G	2855.44 MB	71.42 ms (1.00×)	74.46 ms (1.00×)	71.88
Mobv2-0.15G (Ours)	2.19 M	0.15 G	1354.82 MB	43.48 ms (1.64×)	38.63 ms (1.93×)	68.91

a range of innovative pruning baselines specifically tailored for vision transformers [45, 47, 50, 55, 58, 60, 61, 64, 65, 68, 70]. A DeiT model consists of various components, such as MLP and attention layers. Within the attention layers, we can even find two feasible pruning schemes: pruning 1) the number of heads or 2) within-head dimensions. We implement Isomorphic Pruning to compress a pre-trained DeiT-Base model by pruning both heads, head dimensions, and the embedding sizes following [58]. This is achieved through a straightforward first-order Taylor expansion for importance estimation, followed by a grouped ranking in Isomorphic Pruning. The details of pruning vision transformers can be found in the appendix.

We developed four lightweight vision transformers: DeiT-S, DeiT-2.6G, DeiT-T and DeiT-0.6G. The pruned DeiT-S and DeiT-T models achieve comparable latencies comparable to the standard DeiT-S and DeiT-T with manually designed architecture. The empirical results are reported in Table 2, which shows that our method yields greater performance gains in complex models with extensive heterogeneous internal structures, especially in comparison to CNNs. Specifically, the optimized DeiT-T achieves an accuracy of 77.50%, with only 1.2G MACs, which is better than the results (74.52%) of the pre-trained DeiT-T with uniform width across layers. Detailed configurations, such as base models and pruning ratios, are available in the appendix.

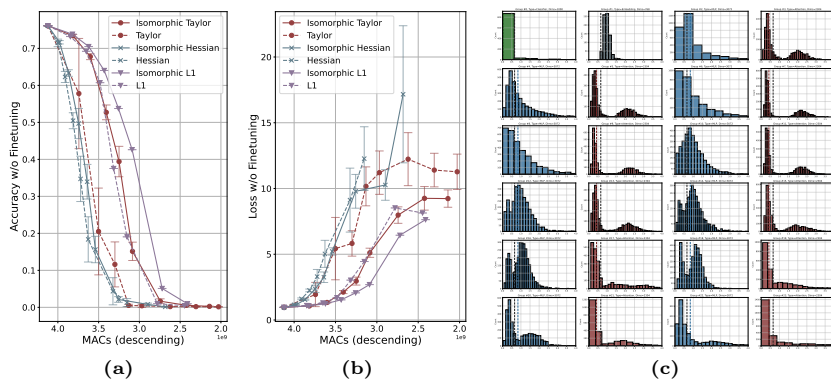


Fig. 4: (a-b) Top-1 Accuracy and loss of pruning ResNet-50 on ImageNet-1K validation set without finetuning. We report the mean and standard deviation of 10 experiments. IsomorphicP Pruning consistently improves the performance of pruned models. (c) Histogram of Taylor importance score of DeiT sub-structures. Isomorphic structures are highlighted with the same color. Zoom in for more details about each structure. We visualize the threshold of 50% pruning ratio for naive global pruning (the black dash) and isomorphic pruning (the colored dash).

4.4 Analysis on the Pruned Models

On-device Latency Table 3 measures the actual latency and peak memory consumptions on a single RTX A5000 GPU and CPU. We evaluate all models with a batch size of 256 for GPU and 8 for CPU, respectively. The average latency over 100-round experiments are reported, with a 20-step warmup process before evaluation. Firstly, results show that the pruned DeiT models can achieve comparable performance (82.41% vs. 82.40%) to some advanced transformers like EfficientFormer [26], while preserving a competitive latency. Additionally, we observe that the peak memory consumption is primarily determined by the largest layer width. Therefore, under the same budget, a uniform DeiT may require less peak memory (1363 MB for DeiT-S) for inference. In contrast, a non-uniform architecture may require slightly more memory (1547 MB for the DeiT-S Pruned) under the given budget. But non-uniform pruning can achieve better accuracy, which reveals an important trade-off between accuracy and memory consumption for transformers. In addition, for CNNs like ResNets, the pruned model requires much less memory due to the compressed feature dimensions. Furthermore, we note that the actual speed-up may not be linearly correlated with the MAC reductions. This is particularly evident in ResNet-50, where only a $1.33\times$ actual speed-up is achieved despite the removal of half of the MACs. In contrast, compression on Mobv2 is more effective, achieving a $1.64\times$ acceleration with a similar compression ratio on MACs.

Distribution of Importance Scores. To further explore how Isomorphic Pruning affects the pruning behavior, we study the distribution of importance

score in a DeiT-Base model [50]. As shown in Figure 4 (c), the importance histogram of different sub-structures is highlighted with different colors, each corresponding to one kind of isomorphic structure. It is natural to find that, the MLP layer and the attention layer have diverged importance distributions due to their different computation process. Besides, when we apply a simple 50% global pruning on the DeiT model, most parameters in Attention #5, #9, and #13 will be removed, since the importance scores in attentions are centered around 0. In contrast, the embedding number as #1 will not be even been pruned in this case, due to the relatively large importance scores. This inevitably leads to an imbalanced pruning of different structures. Notably, the embedding structure, as illustrated in the appendix, contains massive parameters in the DeiT, where biased pruning may lead to very limited acceleration.

Importance Criteria. Note that Isomorphic Pruning can be combined with different pruning criteria like Taylor [39], Hessian [25] and even L1-norm [22]. We quantify the performance gain of Isomorphic Pruning by visualizing the accuracy and loss curve with different pruning ratios in Figure 4 (a) and (b). Specifically, we gradually increase the pruning ratio for each criterion and report the average accuracy and loss without finetuning over 10 repeated runs with random inputs for Taylor expansion [25, 39]. Note that L1 pruning is deterministic and thus presents 0 std. We find the Isomorphic Pruning, marked as solid lines, consistently improves the counterpart baselines highlighted in dashed lines.

Transfer Learning. We also evaluate the performance of transfer learning based on our pruned models. Table 4 reports the accuracy of transferring different models to CIFAR-10 and CIFAR-100. We resize all images to 224×224 and apply the same training recipe for transfer. For the DeiT models, we did not include the supervision from the teacher models. It is observed that our method achieves better transfer accuracy compared to the scratch training model and baseline methods.

Table 4: Transfer learning on pruned models.

Models	CIFAR-10	CIFAR-100
DeiT-S [39]	98.52	87.07
NViT-S [58]	98.78	87.90
DeiT-S (Ours)	99.02	90.09
DeiT-T	93.93	85.66
NViT-T [58]	98.31	85.88
DeiT-S (Ours)	98.28	87.08
ConvNext-S	98.77	89.15
ConvNext-S (Ours)	98.78	89.36

5 Conclusion

In this paper, we introduce Isomorphic Pruning, a practical approach designed to compress vision networks with novel mechanisms and architectures. This method focuses on mitigating the challenges posed by heterogeneous sub-structures in networks, thereby enhancing the reliability of ranking and pruning processes. Our empirical results on the ImageNet-1k dataset, showcase the effectiveness of Isomorphic Pruning across several vision models like ConvNext, ResNet, MobileNet-v2, and Vision Transformers.

Acknowledgement

This project is supported by the National Research Foundation Singapore under its Medium Sized Center for Advanced Robotics Technology Innovation, and the Advanced Research and Technology Innovation Centre (ARTIC), the National University of Singapore under Grant (project number: A-0005947-21-00, project reference: ECT-RP2).

References

1. Baum, E., Haussler, D.: What size net gives valid generalization? *Advances in neural information processing systems* **1** (1988)
2. Chavan, A., Shen, Z., Liu, Z., Liu, Z., Cheng, K.T., Xing, E.P.: Vision transformer slimming: Multi-dimension searching in continuous optimization space. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 4931–4941 (2022)
3. Chen, C.F.R., Fan, Q., Panda, R.: Crossvit: Cross-attention multi-scale vision transformer for image classification. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 357–366 (2021)
4. Chen, T., Cheng, Y., Gan, Z., Yuan, L., Zhang, L., Wang, Z.: Chasing sparsity in vision transformers: An end-to-end exploration. *Advances in Neural Information Processing Systems* **34**, 19974–19988 (2021)
5. Chen, T., Liang, L., Ding, T., Zhu, Z., Zharkov, I.: Otov2: Automatic, generic, user-friendly. *arXiv preprint arXiv:2303.06862* (2023)
6. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 1251–1258 (2017)
7. Cubuk, E.D., Zoph, B., Shlens, J., Le, Q.V.: Randaugment: Practical automated data augmentation with a reduced search space. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*. pp. 702–703 (2020)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. *Ieee* (2009)
9. Ding, X., Ding, G., Guo, Y., Han, J., Yan, C.: Approximated oracle filter pruning for destructive cnn width optimization. In: *International Conference on Machine Learning*. pp. 1607–1616. *PMLR* (2019)
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
11. Fang, G., Ma, X., Song, M., Mi, M.B., Wang, X.: Depgraph: Towards any structural pruning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 16091–16101 (2023)
12. Frantar, E., Alistarh, D.: Sparsegpt: Massive language models can be accurately pruned in one-shot (2023)
13. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677* (2017)

14. Guo, J., Ouyang, W., Xu, D.: Multi-dimensional pruning: A unified framework for model compression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1508–1517 (2020)
15. Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., Wang, Y.: Transformer in transformer. *Advances in Neural Information Processing Systems* **34**, 15908–15919 (2021)
16. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *International Conference on Learning Representations (ICLR)* (2016)
17. Hassibi, B., Stork, D.G., Wolff, G.J.: Optimal brain surgeon and general network pruning. In: *IEEE international conference on neural networks*. pp. 293–299. IEEE (1993)
18. He, H., Liu, J., Pan, Z., Cai, J., Zhang, J., Tao, D., Zhuang, B.: Pruning self-attentions into convolutional layers in single path. *arXiv preprint arXiv:2111.11802* (2021)
19. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)
20. He, Y., Kang, G., Dong, X., Fu, Y., Yang, Y.: Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866* (2018)
21. He, Y., Liu, P., Wang, Z., Hu, Z., Yang, Y.: Filter pruning via geometric median for deep convolutional neural networks acceleration. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 4340–4349 (2019)
22. He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. In: *Proceedings of the IEEE international conference on computer vision*. pp. 1389–1397 (2017)
23. Hoffer, E., Ben-Nun, T., Hubara, I., Giladi, N., Hoefler, T., Soudry, D.: Augment your batch: Improving generalization through instance repetition. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8129–8138 (2020)
24. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 4700–4708 (2017)
25. LeCun, Y., Denker, J., Solla, S.: Optimal brain damage. *Advances in neural information processing systems* **2** (1989)
26. Li, Y., Yuan, G., Wen, Y., Hu, J., Evangelidis, G., Tulyakov, S., Wang, Y., Ren, J.: Efficientformer: Vision transformers at mobilenet speed. *Advances in Neural Information Processing Systems* **35**, 12934–12949 (2022)
27. Lin, M., Ji, R., Li, S., Wang, Y., Wu, Y., Huang, F., Ye, Q.: Network pruning using adaptive exemplar filters. *IEEE Transactions on Neural Networks and Learning Systems* **33**(12), 7357–7366 (2021)
28. Lin, M., Ji, R., Zhang, Y., Zhang, B., Wu, Y., Tian, Y.: Channel pruning via automatic structure search. *arXiv preprint arXiv:2001.08565* (2020)
29. Liu, B., Wang, M., Foroosh, H., Tappen, M., Pensky, M.: Sparse convolutional neural networks. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 806–814 (2015)
30. Liu, L., Zhang, S., Kuang, Z., Zhou, A., Xue, J.H., Wang, X., Chen, Y., Yang, W., Liao, Q., Zhang, W.: Group fisher pruning for practical network compression. In: *International Conference on Machine Learning*. pp. 7021–7032. PMLR (2021)
31. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings*

- of the IEEE/CVF international conference on computer vision. pp. 10012–10022 (2021)
32. Liu, Z., Mu, H., Zhang, X., Guo, Z., Yang, X., Cheng, K.T., Sun, J.: Metapruning: Meta learning for automatic neural network channel pruning. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3296–3305 (2019)
 33. Liu, Z., Mao, H., Wu, C.Y., Feichtenhofer, C., Darrell, T., Xie, S.: A convnet for the 2020s. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11976–11986 (2022)
 34. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101 (2017)
 35. Louizos, C., Welling, M., Kingma, D.P.: Learning sparse neural networks through l_0 regularization. arXiv preprint arXiv:1712.01312 (2017)
 36. Ma, X., Yuan, G., Lin, S., Li, Z., Sun, H., Wang, Y.: Resnet can be pruned 60×: Introducing network purification and unused path removal (p-rm) after weight pruning. In: 2019 IEEE/ACM International Symposium on Nanoscale Architectures (NANOARCH). pp. 1–2. IEEE (2019)
 37. Ma, X., Fang, G., Wang, X.: Llm-pruner: On the structural pruning of large language models. In: Advances in Neural Information Processing Systems (2023)
 38. maintainers, T., contributors: Torchvision: Pytorch’s computer vision library. GitHub repository (2016)
 39. Molchanov, P., Mallya, A., Tyree, S., Frosio, I., Kautz, J.: Importance estimation for neural network pruning. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11264–11272 (2019)
 40. Molchanov, P., Tyree, S., Karras, T., Aila, T., Kautz, J.: Pruning convolutional neural networks for resource efficient inference. arXiv preprint arXiv:1611.06440 (2016)
 41. Peng, H., Wu, J., Chen, S., Huang, J.: Collaborative channel pruning for deep networks. In: International Conference on Machine Learning. pp. 5113–5122. PMLR (2019)
 42. Pool, J., Sawarkar, A., Rodge, J.: Accelerating inference with sparsity using the nvidia ampere architecture and nvidia tensorrt. NVIDIA Developer Technical Blog, <https://developer.nvidia.com/blog/accelerating-inference-with-sparsity-using-ampere-and-tensorrt> (2021)
 43. Radosavovic, I., Kosaraju, R.P., Girshick, R., He, K., Dollár, P.: Designing network design spaces. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 10428–10436 (2020)
 44. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Mobilenetv2: Inverted residuals and linear bottlenecks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4510–4520 (2018)
 45. Shi, D., Tao, C., Jin, Y., Yang, Z., Yuan, C., Wang, J.: UPop: Unified and progressive pruning for compressing vision-language transformers. In: Proceedings of the 40th International Conference on Machine Learning. vol. 202, pp. 31292–31311. PMLR (2023)
 46. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. arXiv preprint arXiv:1409.1556 (2014)
 47. Song, Z., Xu, Y., He, Z., Jiang, L., Jing, N., Liang, X.: Cp-vit: Cascade vision transformer pruning via progressive sparsity prediction. arXiv preprint arXiv:2203.04570 (2022)
 48. Sun, M., Liu, Z., Bair, A., Kolter, J.Z.: A simple and effective pruning approach for large language models. arXiv preprint arXiv:2306.11695 (2023)

49. Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., Wojna, Z.: Rethinking the inception architecture for computer vision. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 2818–2826 (2016)
50. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jegou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning. vol. 139, pp. 10347–10357 (July 2021)
51. Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., Jégou, H.: Going deeper with image transformers. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 32–42 (2021)
52. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
53. Voita, E., Talbot, D., Moiseev, F., Sennrich, R., Titov, I.: Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. arXiv preprint arXiv:1905.09418 (2019)
54. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pvt v2: Improved baselines with pyramid vision transformer. *Computational Visual Media* **8**(3), 415–424 (2022)
55. Wang, Z., Luo, H., Wang, P., Ding, F., Wang, F., Li, H.: Vtc-lfc: Vision transformer compression with low-frequency components. *Advances in Neural Information Processing Systems* **35**, 13974–13988 (2022)
56. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. *Advances in neural information processing systems* **29** (2016)
57. Wightman, R.: Pytorch-image-models. <https://github.com/huggingface/pytorch-image-models>
58. Yang, H., Yin, H., Shen, M., Molchanov, P., Li, H., Kautz, J.: Global vision transformer pruning with hessian-aware saliency. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 18547–18557 (2023)
59. Ye, J., Lu, X., Lin, Z., Wang, J.Z.: Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. arXiv preprint arXiv:1802.00124 (2018)
60. Yin, M., Uzkent, B., Shen, Y., Jin, H., Yuan, B.: Gohsp: A unified framework of graph and optimization-based heterogeneous structured pruning for vision transformer. arXiv preprint arXiv:2301.05345 (2023)
61. Yu, F., Huang, K., Wang, M., Cheng, Y., Chu, W., Cui, L.: Width & depth pruning for vision transformers. In: AAAI Conference on Artificial Intelligence (AAAI). vol. 2022 (2022)
62. Yu, H., Wu, J.: A unified pruning framework for vision transformers. *Science China Information Sciences* **66**(7), 1–2 (2023)
63. Yu, J., Huang, T.: Autoslim: Towards one-shot architecture search for channel numbers. arXiv preprint arXiv:1903.11728 (2019)
64. Yu, L., Xiang, W.: X-pruner: explainable pruning for vision transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 24355–24363 (2023)
65. Yu, S., Chen, T., Shen, J., Yuan, H., Tan, J., Yang, S., Liu, J., Wang, Z.: Unified visual transformer compression. arXiv preprint arXiv:2203.08243 (2022)
66. Yun, S., Han, D., Oh, S.J., Chun, S., Choe, J., Yoo, Y.: Cutmix: Regularization strategy to train strong classifiers with localizable features. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 6023–6032 (2019)
67. Zhang, H., Cisse, M., Dauphin, Y.N., Lopez-Paz, D.: mixup: Beyond empirical risk minimization. arXiv preprint arXiv:1710.09412 (2017)

68. Zheng, C., Zhang, K., Yang, Z., Tan, W., Xiao, J., Ren, Y., Pu, S., et al.: Savit: Structure-aware vision transformer pruning via collaborative optimization. *Advances in Neural Information Processing Systems* **35**, 9010–9023 (2022)
69. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 34, pp. 13001–13008 (2020)
70. Zhu, M., Tang, Y., Han, K.: Vision transformer pruning. *arXiv preprint arXiv:2104.08500* (2021)