
DynaMixer: A Vision MLP Architecture with Dynamic Mixing

Ziyu Wang¹ Wenhao Jiang¹ Yiming Zhu² Li Yuan³ Yibing Song⁴ Wei Liu¹

Abstract

Recently, MLP-like vision models have achieved promising performances on mainstream visual recognition tasks. In contrast with vision transformers and CNNs, the success of MLP-like models shows that simple information fusion operations among tokens and channels can yield a good representation power for deep recognition models. However, existing MLP-like models fuse tokens through static fusion operations, lacking adaptability to the contents of the tokens to be mixed. Thus, customary information fusion procedures are not effective enough. To this end, this paper presents an efficient MLP-like network architecture, dubbed DynaMixer, resorting to dynamic information fusion. Critically, we propose a procedure, on which the DynaMixer model relies, to dynamically generate mixing matrices by leveraging the contents of all the tokens to be mixed. To reduce the time complexity and improve the robustness, a dimensionality reduction technique and a multi-segment fusion mechanism are adopted. Our proposed DynaMixer model (97M parameters) achieves 84.3% top-1 accuracy on the ImageNet-1K dataset without extra training data, performing favorably against the state-of-the-art vision MLP models. When the number of parameters is reduced to 26M, it still achieves 82.7% top-1 accuracy, surpassing the existing MLP-like models with a similar capacity. The code is available at <https://github.com/ziyuwang/DynaMixer>.

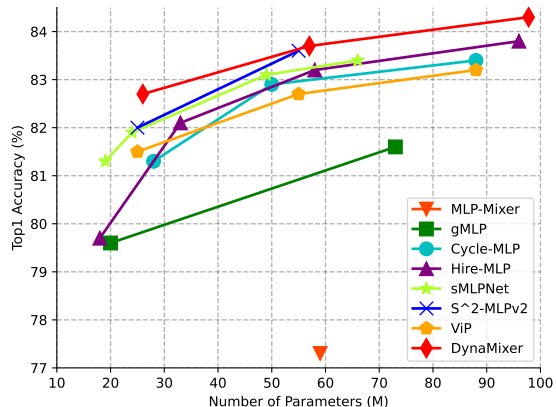


Figure 1. **ImageNet accuracy v.s. model capacity.** All models are trained on ImageNet-1K without extra data. DynaMixer outperforms all the existing MLP-like models, such as MLP-Mixer (Tolstikhin et al., 2021), Hire-MLP (Guo et al., 2021), S²-MLP-v2 (Yu et al., 2021a), ViP (Hou et al., 2021), and Cycle-MLP (Chen et al., 2021b).

1. Introduction

Convolutional neural networks (CNNs) have been the dominating solution for a wide range of computer vision tasks for a long time, *e.g.*, visual recognition and segmentation. CNNs aim at training the whole network end-to-end to remove hand-crafted visual features and inductive biases. And they heavily rely on convolution operators and pooling operators, which introduce locality and spatial invariances. Recently, Transformers (Vaswani et al., 2017), which have achieved overwhelming success in natural language processing, have also been introduced into the compute vision field. The Vision Transformers (ViTs) (Dosovitskiy et al., 2021) have achieved state-of-the-art performances on visual recognition tasks. In ViT, images are first divided into a sequence of non-overlapping patches, which can be seen as word tokens in the natural language processing field. Those patches are then fed into a stack of blocks based on self-attention to obtain the contextualized output tokens. ViT removes the convolution and pooling operators from CNNs, and employs a self-attention mechanism to model the relationships among image patches. Inspired by ViTs, many efforts have

¹Data Platform, Tencent ²Graduate school at ShenZhen, Tsinghua university ³School of Electrical and Computer Engineering, Peking University ⁴Tencent AI Lab. Correspondence to: Wenhao Jiang <cswhjiang@gmail.com>, Wei Liu <wl2223@columbia.edu>.

been made to designing simpler models in the setting with a large amount of data available.

More recently, a pure multilayer perceptron (MLP) architecture (Tolstikhin et al., 2021), called MLP-Mixer, was proposed to reduce the inductive biases further and showed competitive performance with CNN-based models and vision transformers. MLP-Mixer is much simpler than transformer-based models because it utilizes MLP as a building block, removing the need of invoking self-attention. In MLP-Mixer, each layer mainly relies on two steps to perform information interaction: token mixing step and channel mixing step, both based on MLP blocks. The token mixing step performs fusion among the tokens, while the channel mixing step performs fusion on the channel dimension. Inspired by this, a number of MLP-like models (Chen et al., 2021b; Hou et al., 2021; Yu et al., 2021a;b) have emerged for further improvements. However, these MLP-like methods rely on fixed static mixing matrices for patch communications, which may restrict the adaptability to the contents to be fused.

To solve the above limitation, we propose a vision MLP architecture with dynamic mixing, dubbed DynaMixer, which can generate mixing matrices dynamically for each set of tokens to be mixed by considering their contents. Note that mixing all the image tokens consumes a significant time cost. For a computational speedup, we mix tokens in a row-wise and column-wise way. During each mixing process, we reduce the feature dimensionality to generate the mixing matrices. We empirically find that we can reduce feature dimensionality significantly without undermining the performance much. *The feature dimensionality can even be reduced to 1 in our experiments.* Furthermore, we divide the feature channels into multiple segments and perform token mixing separately. This improves the mixing robustness as well as efficiency. As a result, our method achieves state-of-the-art performance among existing MLP-like vision models on the ImageNet-1K dataset (Deng et al., 2009), as summarized in Fig 1. Without any extra data, our DynaMixer accomplishes 84.3% top-1 accuracy with 97M parameters. When the model parameters are reduced to 26M, DynaMixer can still reach 82.7% top-1 accuracy, far exceeding other MLP-like models with a similar number of parameters.

2. Related Work

The deep neural networks for vision related tasks evolve from CNNs to ViTs and MLPs. In this section, we review these models and illustrate the relationships between our DynaMixer and these models.

2.1. Convolutional Neural Networks

The CNNs have drawn huge attention since AlexNet (Krizhevsky et al., 2012) improved image

classification by a large margin. Convolution operations, nonlinear activation, and pooling operations have become prevalent for CNN architecture designs. Based on the previous works on CNN, inceptions (Szegedy et al., 2015; 2016; 2017) were proposed by employing multiple parallel branches. A breakthrough was made by ResNet (He et al., 2016), in which the network depth was significantly increased. To ease training, skip connections with identity mapping were introduced into ResNets. They achieved state-of-the-art visual recognition performances at that time and inspired further research, including ResNeXt (Xie et al., 2017) and DenseNet (Huang et al., 2017). Afterward, attention mechanisms were gradually explored to benefit CNNs. Examples include SENet (Hu et al., 2018), non-local neural network (Wang et al., 2018), and local relationship network (Hu et al., 2019). CNNs have been widely applied to computer vision related tasks, such as visual recognition, image/video generation, object detection, semantic segmentation, and so on.

2.2. Vision Transformers

Transformer (Vaswani et al., 2017) was firstly proposed in the machine translation area. It was then adopted for computer vision applications in ViT (Dosovitskiy et al., 2021). Input images in ViT are divided into a sequence of non-overlapping patches, which are regarded as tokens and are fed into transformers. Unlike CNNs aggregating information within local windows, ViTs model global information interactions among visual tokens via self-attention. Inspired by ViTs, many variants are proposed to remove the drawbacks of original ViTs. In DeepViT (Zhou et al., 2021a) and CaiT (Touvron et al., 2021c), the problem that the performance of ViTs saturates fast when scaled to be deeper was addressed. In (Yuan et al., 2021b; Chen et al., 2021a; Han et al., 2021; Liu et al., 2021b; Wang et al., 2021b; Fan et al., 2021; Ge et al., 2021; Liang et al., 2022; Tong et al., 2022; Bai et al., 2022; Zhou et al., 2021b), multi-granularity was introduced into ViTs to improve the generalization ability. A complete survey on vision transformers can be found in (Han et al., 2020; Khan et al., 2021). All transformer-based models benefit from the self-attention mechanism for its flexibility and adaptiveness. Our model does not involve the self-attention mechanism while maintaining adaptability. Differences will be discussed in the following section.

2.3. Vision MLPs

Recently, MLP-based vision models (Tolstikhin et al., 2021; Touvron et al., 2021a) were proposed to reduce the inductive bias further and showed competitive performances with CNN-based models and ViTs. Existing MLP-like models share a similar macro framework but have different micro block designs. MLP-like models usually divide one input image into patches, like in vision transformers, and then per-

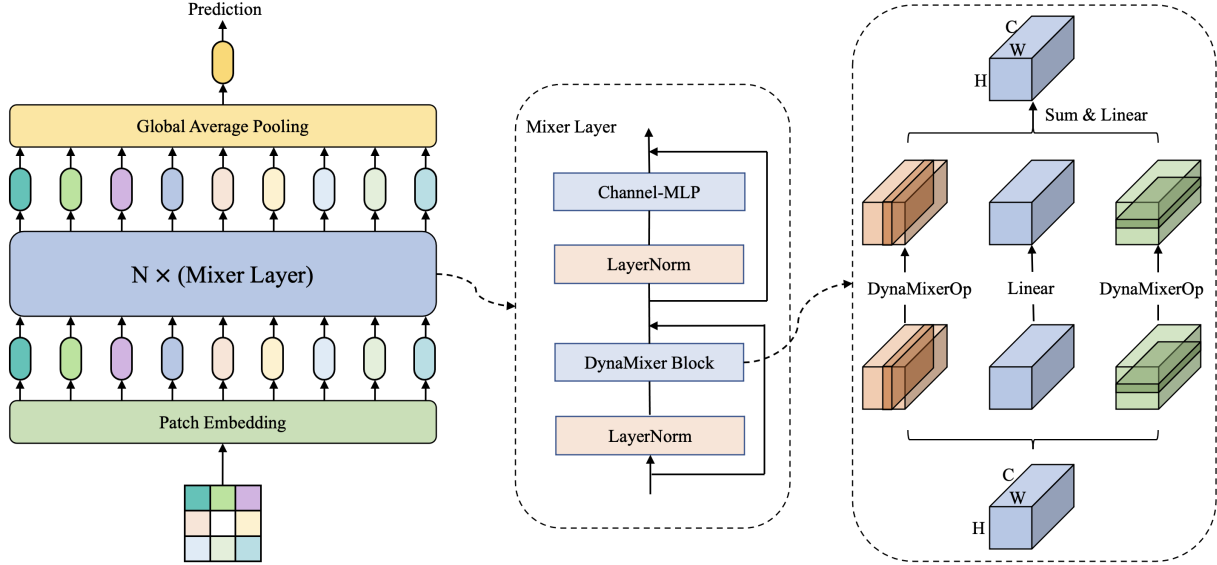


Figure 2. **The basic architecture of our proposed DynaMixer.** DynaMixer consists of a patch embedding layer, several mixer layers, a global average pooling layer, and a classifier head. The patch embedding layer transforms input non-overlapping patches into corresponding input tokens, which are fed into a sequence of mixer layers to generate the output tokens. All output tokens are averaged in the average pooling layer, and the final prediction is generated with a classifier head. The mixer layer (middle part) contains two layer-normalization layers, a DynaMixer block and a channel-MLP block. The DynaMixer block (right part) performs row mixing and column mixing via the DynaMixer operations (depicted in Fig. 3), and a simple channel-mixing via linear projection. The mixing results are summed and outputted with a linear transformation.

form two main steps: token mixing step and channel mixing step. The specific details of these two steps, especially token-mixing steps, are different among the existing methods. ViP (Hou et al., 2021) was proposed to perform token mixing along the height and width dimensions with simple linear projections to encode spatial information, which is different from other MLP-like models that fuse token information among all tokens in one step. S²-MLP (Yu et al., 2021a) replaced the token mixing step with a spatial-shift step to enable information interaction among tokens. CycleMLP (Chen et al., 2021b) samples local points along the channel dimension in a cyclical style, while AS-MLP (Lian et al., 2021) samples locations in a cross. Hire-MLP (Guo et al., 2021) mixes tokens within a local region and across local regions.

All these MLP-like methods rely on static mixing matrices for patch communications, which may limit the performance for lacking adaptiveness since it is natural that a different set of tokens should be mixed differently. Thus, we propose to generate the mixing matrices dynamically by considering the contents of tokens to be mixed. A recent work, Synthesizer (Tay et al., 2021), pointed out that either the attention matrix in ViTs or the mixing matrix in MLP can be seen as generated with a specially designed function. It is worth noting that although the Synthesizer (Dense) (Tay et al., 2021) adopts to generate the mixing matrix dynamically, the

mixing weights for a specific position are determined only by the contents of the token at the same position. However, the weights for it in our model are determined by all the tokens. We will show how to design such a generating function and how to reduce the computational complexity in the following section.

3. The Architecture of DynaMixer

In this section, we present our model designs. The overall framework of our network is illustrated in Fig. 2. Like in other MLP-based models, the input image is divided into non-overlapping patches. All patches are projected onto the embedding space with a shared matrix, which are then fed into a sequence of mixer layers to generate the output tokens. All output tokens are averaged in the average pooling layer, and then sent into the classifier head to yield the final prediction. Except for the layer-normalization layers and skip connections, the mixer layer contains a DynaMixer block and a channel-MLP block, which are responsible for fusing token information and channel information, respectively. The channel-MLP block is just a feed-forward layer as in Transformer (Vaswani et al., 2017). The DynaMixer block performs row mixing and column mixing via DynaMixer operations, and channel-mixing via simple linear projection, respectively. The mixing results are aggregated and

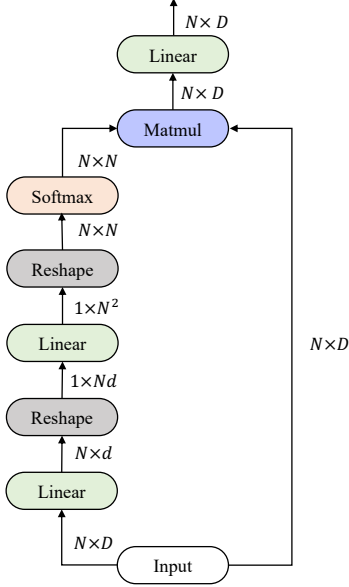


Figure 3. The procedure of our proposed DynaMixer operation for one segment.

outputted via a linear transformation. We first illustrate our DynaMixer operation, and then elaborate on the DynaMixer block. Finally, we analyze the differences from other MLP-like models.

3.1. The DynaMixer Operation

Dynamic mixing matrix generation. The principle of our design is to generate a dynamic mixing matrix P given a set of input tokens $X \in R^{N \times D}$ by considering their contents, where N is the number of tokens, and D is the feature dimensionality. Once we obtain P , we can mix the tokens by $Y = PX$ to obtain the output tokens Y . A simple way to get P is utilizing a linear function of all input token features to estimate it. Thus, we can simply flatten X into a vector and generate the mixing matrix as

$$P_i = \text{softmax}\left(\text{flat}(X)^T W^{(i)}\right), \quad (1)$$

where $\text{flat}(X) \in R^{ND \times 1}$ is a vector by flattening X , $\text{softmax}(\cdot)$ is the softmax operator performing on a row vector, $W^{(i)} \in R^{ND \times N}$, and P is the mixing matrix. P_i is the i -th row of P , and it contains the mixing weights for the i -th output token. However, the number of parameters of the above process is too large since ND is usually too big. Thus, we can perform dimensionality reduction first to reduce the number of parameters. Therefore, we have the following steps to generate the mixing matrix and perform

mixing on tokens:

$$\hat{X} = XW_d, \quad (2)$$

$$P_i = \text{softmax}\left(\text{flat}(\hat{X})W^{(i)}\right), \quad (3)$$

$$Y = PX, \quad (4)$$

where $\hat{X} \in R^{N \times d}$, and $d \ll D$ is a quite small number, say 1 or 2. The whole procedure for generation and mixing is described pictorially in Fig 3, which can be implemented easily with PyTorch (Paszke et al., 2019) or TensorFlow (Abadi et al., 2015).

Multi-segment fusion mechanism. To improve the robustness of our model, we divide features into S segments, perform the mixing operation separately, and combine the mixed results to obtain the final results:

$$\hat{X}^{(s)} = XW_d^{(s)}, \quad (5)$$

$$P_i^{(s)} = \text{softmax}\left(\text{flat}(\hat{X}^{(s)})W^{(s,i)}\right), \quad (6)$$

$$Y = [P^{(0)}X^{(0)}, \dots, P^{(S-1)}X^{(S-1)}]W_o, \quad (7)$$

where $\hat{X}^{(s)} \in R^{N \times d}$, $W^{(s,i)} \in R^{Nd \times N}$, $P_i^{(s)}$ is the i -th row of the token mixing matrix for the s -th segment, $X^{(s)}$ is the tokens in the s -th segment, $[\cdot, \cdot]$ is the concatenation operation, and $W_o \in R^{D \times D}$ is a feature fusion matrix for output. Note that, different dimensionality reduction matrices, *i.e.*, $W_d^{(s)}$, are used for different segments, which are beneficial for the expressive power of the model. Moreover, the mixing matrix for the s -th segment is also effected by the contents of other segments. Thus, the mixing behaviours for different segments are not independent.

Weights sharing among segments. To reduce the number of parameters, the matrices $W^{(s,i)}$ are shared among all segments. Thus, the number of parameters of the DynaMixer operation is $S \times D \times d + N^3 \times d + D^2$. In practice, we found that d can be quite small, *e.g.*, $d = 1$ or $d = 2$.

We use the following expression to denote the DynaMixer operation in the rest of this paper:

$$Y = \text{DynaMixerOp}(X). \quad (8)$$

3.2. The DynaMixer Block

The number of parameters of the above operation mainly depends on N , the number of tokens to be mixed, which is $H \times W$, where H and W are the numbers of patches in the height and width directions, respectively. Thus, to reduce the computational complexity, we adopt a strategy similar to ViP (Hou et al., 2021), which mixes tokens in a row or column at one step. As depicted in the right part of Fig 2, the DynaMixer block consists of three components, which

Algorithm 1 Pseudo-code for DynaMixer Block (PyTorch-like)

```
##### initializaiton #####
proj_c = nn.Linear(D, D)
proj_o = nn.Linear(D, D)

##### code in forward #####
def dyna_mixer_block(self, X):
    H, W, D = X.shape

    # row mixing
    for h = 1:H
        Y_h[h, :, :] = DynaMixerOp_h(X[h, :, :])

    # column mixing
    for w = 1:W
        Y_w[:, w, :] = DynaMixerOp_w(X[:, w, :])

    # channel mixing
    Y_c = proj_c(X)
    Y_out = Y_h + Y_w + Y_c

    return proj_o(Y_out)
```

are row mixing, column mixing, and channel mixing. In row mixing, the tokens that belong to the same row are mixed via DynaMixer operations. And the parameters of DynaMixer operations are shared among all rows. Similarly, the tokens that belong to the same column are mixed via DynaMixer operation in column mixing, in which the parameters of DynaMixer operations are also shared. The channel mixing is simply a linear transformation on features. We denote the outputs of the three components as Y_h , Y_w and Y_c , and the output of the DynaMixer block is

$$Y_{out} = (Y_h + Y_w + Y_c)W_o' \quad (9)$$

The procedure for the DynaMixer block is summarized in Algorithm 1, in which we denote the input as $X \in R^{H \times W \times D}$ for convenience.

In Algorithm 1, the three mixing results are just summed to obtain the final result, as described in Eq. (9). To further improve the performance, importances of the three components can be estimated by a method similar to (Zhang et al., 2020; Hou et al., 2021). In the following experiments, the weighted DynaMixer block is used by default.

3.3. Discussion

DynaMixer v.s. Synthesizer (Dense). The weights for the i -th output token generated by Synthesizer (Dense) are only determined by the contents of the i -th input token. Thus, *it cannot model the effects from the other tokens, leading to inaccurate weight estimation.* However, our model considers all the tokens to generate more accurate mixing weights. Moreover, our model divides the channel into

non-overlapping segments and performs separate mixing operations on them. Such a design can improve the robustness and expressive power of our model.

DynaMixer v.s. ViP. Except that the mixing matrices in our model are generated dynamically, the mixing behaviors are also quite different. Suppose W tokens in the h -th row are to be mixed. Each token is divided into S segments along channel dimension, and the dimensionality of channel is denoted as D . In ViP, $W \times S$ vectors of dimensionality $\frac{D}{S}$ will be mixed by multiplying one fixed matrix. However, in DynaMixer, the mixing operation occurs S times. One time for one segment. And each mixing operation is performed on W vectors of dimensionality $\frac{D}{S}$ with different mixing matrices.

DynaMixer v.s. self-attention. The mixing matrix (attention map) in self-attention models the *pair-wise relationships between tokens*. Thus, it is a quadratic function of token features. However, in our method, the mixing matrix is estimated by computing the *relationships between tokens and a trainable matrix* ($W^{s,i}$ in Eq. (6)).

4. Experimental Results

In this section, we present the experimental results and analysis. First, we will give the configurations of our DynaMixer used in the experiments, and then the experimental settings and results on the ImageNet-1K dataset are provided. At last, the ablation studies are presented to provide a deep understanding of the designs in our model.

4.1. Configurations of DynaMixer

The configurations of our model used in the experiments are summarized in Table 1. There are three versions of DynaMixer, denoted as ‘‘DynaMixer-S’’, ‘‘DynaMixer-M’’, and ‘‘DynaMixer-L’’, according to the model sizes. In all our experiments, the input image size is 224×224 , and the input patch size is 7×7 . All our models have two stages, and each starts with a patch embedding layer. The patch size for the second stage is 2×2 . Thus, the corresponding patch embedding layer can be seen as a downsampling layer.

4.2. Experimental Configurations

We train our proposed DynaMixer on the public image classification benchmark ImageNet-1K dataset (Deng et al., 2009), which covers 1K categories of natural images and contains 1.2M training images and 50K validation images. Because the test set for this benchmark is unlabeled, we follow the common practice by evaluating the performance on the validation set. The code implementation is based on

DynaMixer: A Vision MLP Architecture with Dynamic Mixing

Specification	DynaMixer-S	DynaMixer-M	DynaMixer-L
Patch size	7×7	7×7	7×7
Hidden size	192	256	256
#Tokens	32×32	32×32	32×32
#Mixer Layers	4	7	8
S	8	8	8
Patch size	2×2	2×2	2×2
Hidden size	384	512	512
#Tokens	16×16	16×16	16×16
#Mixer Layers	14	17	28
S	16	16	16
#Layers	18	24	36
MLP Ratio	3	3	3
Stoch. Dep.	0.1	0.1	0.3
d	2	2	8
#Parameters	26M	57M	97M

Table 1. The configurations of DynaMixers used. We design three DynaMixers (DynaMixer-S, DynaMixer-M, and DynaMixer-L) of different sizes (Small, Medium, and Large) for the experiments, according to the number of parameters. The first two groups of rows show the configurations of the first and the second stage of the models, respectively. The third group of rows shows some other hyperparameters.

PyTorch (Paszke et al., 2019) and the TIMM¹ toolbox.

We train our model with one machine with 8 NVIDIA A100 GPUs with data parallelism. For model optimization, we adopt the AdamW optimizer (Loshchilov & Hutter, 2017). The learning rates for DynaMixer-S, DynaMixer-M, and DynaMixer-L are 0.002, and the corresponding batch sizes on one GPU are 256, 128, and 64, respectively. We set the weight decay rate to 0.05 and set the warmup learning rate to 10^{-6} to follow the settings in previous work (Touvron et al., 2021b; Jiang et al., 2021). We use automatic mixed-precision of the PyTorch version for training acceleration. Stochastic Depth (Huang et al., 2016) is used, and the drop rate varies with the size of the model. The model is trained for 300 epochs. For data augmentation methods, we use CutOut (Zhong et al., 2020), RandAug (Cubuk et al., 2020), MixUp (Zhang et al., 2017), and CutMix (Yun et al., 2019). Detailed settings about S , d , and the path drop rates for different versions of DynaMixer can be found in Table 1.

4.3. Results on ImageNet-1K

In this subsection, we compare our model with MLP-like models, CNN-based models, and transformer-based models to show the advantages of our dynamic mixing mechanism.

Comparisons with MLP-like Models. First, we compare our DynaMixer with MLP-like models. The top-1 accuracies of different MLP-like models on ImageNet-1K (Deng et al., 2009) without external data are shown in Table 2.

¹<https://github.com/rwightman/pytorch-image-models>

Model	Date	Param	FLOPs	Top-1 (%)
Mixer-B/16 (Tolstikhin et al., 2021)	MAY, 2021	59M	12.7G	76.4
Mixer-B/16 [†] (Tolstikhin et al., 2021)		59M	12.7G	77.3
ResMLP-S12 (Touvron et al., 2021a)	MAY, 2021	15M	3.0G	76.6
ResMLP-S24 (Touvron et al., 2021a)		30M	6.0G	79.4
ResMLP-B24 (Touvron et al., 2021a)		116M	23.0G	81.0
gMLP-Ti (Liu et al., 2021a)	MAY, 2021	6M	1.4G	72.3
gMLP-S (Liu et al., 2021a)		20M	4.5G	79.6
gMLP-B (Liu et al., 2021a)		73M	15.8G	81.6
ViP-Small/14 (Hou et al., 2021)	JUN, 2021	30M	6.9G	80.7
ViP-Small/7 (Hou et al., 2021)		25M	6.9G	81.5
ViP-Medium/7 (Hou et al., 2021)		55M	16.3G	82.7
ViP-Large/7 (Hou et al., 2021)		88M	24.4G	83.2
AS-MLP-T (Lian et al., 2021)	JUL, 2021	28M	4.4G	81.3
AS-MLP-S (Lian et al., 2021)		50M	8.5G	83.1
AS-MLP-B (Lian et al., 2021)		88M	15.2G	83.3
S ² -MLPv2-Small/7 (Yu et al., 2021b)	AUG, 2021	25M	6.9G	82.0
S ² -MLPv2-Medium/7 (Yu et al., 2021b)		55M	16.3G	83.6
RaftMLP-36 (Tatsunami & Taki, 2021)	AUG, 2021	44M	9.0G	76.9
RaftMLP-12 (Tatsunami & Taki, 2021)		58M	12.0G	78.0
sMLPNet-T (Tang et al., 2021)	SEP, 2021	24M	5.0G	81.9
sMLPNet-S (Tang et al., 2021)		49M	10.3G	83.1
sMLPNet-B (Tang et al., 2021)		66M	14.0G	83.3
ConvMLP-S (Li et al., 2021)	SEP, 2021	9M	2.4G	76.8
ConvMLP-M (Li et al., 2021)		17M	3.9G	79.0
ConvMLP-L (Li et al., 2021)		43M	9.9G	80.2
CycleMLP-T (Chen et al., 2021b)	NOV, 2021	28M	4.4G	81.3
CycleMLP-S (Chen et al., 2021b)		50M	8.5G	82.9
CycleMLP-B (Chen et al., 2021b)		88M	15.2G	83.4
Hire-MLP-Ti (Guo et al., 2021)	NOV, 2021	18M	2.1G	79.7
Hire-MLP-S (Guo et al., 2021)		33M	4.2G	82.1
Hire-MLP-B (Guo et al., 2021)		58M	8.1G	83.2
Hire-MLP-L (Guo et al., 2021)		96M	13.4G	83.8
DynaMixer-S		26M	7.3G	82.7
DynaMixer-M		57M	17.0G	83.7
DynaMixer-L		97M	27.4G	84.3

Table 2. Image classification results of our DynaMixer and other MLP-like models on the ImageNet-1K benchmark without extra data. “Top-1” denotes Top-1 accuracy. “Param” and FLOPs denote the number of parameters and the number of floating point operations, respectively. Model with [†] was re-implemented by (Liu et al., 2021a). The date column means the corresponding initial release date on arXiv. The FLOPs are computed with the code of CycleMLP (Chen et al., 2021b).

The “Param” and “FLOPs” columns denote the number of parameters and the number of floating point operations, respectively. The results are extracted from the corresponding papers, and the FLOPs are computed with the code of CycleMLP². Firstly, it is obvious that our model consistently outperforms the other models of a similar number of parameters. In addition, we can see that our DynaMixer-S model with 26M parameters achieves top-1 accuracy of 82.7%, which has already surpassed most of the existing MLP-like models of all sizes and is even better than gMLP-B (Liu et al., 2021a) with 73M parameters. Increasing the number of parameters to 57M, our DynaMixer-M obtains accuracy 83.7%, which is superior to all MLP-like models. Further expanding the network to 97M parameters, DynaMixer-L can achieve top-1 accuracy of 84.3%, which is a new state-of-the-art among the MLP-like models.

²<https://github.com/ShoufaChen/CycleMLP>

By comparing with ViP (Hou et al., 2021), which is the most similar model to ours, we find that the number of parameters and FLOPs of our model are only slightly larger than those of ViP. However, the improvements of performance are not caused by the increment on model size. The model sizes of our model and ViP are similar. For example, the number of parameters and FLOPs of DynaMixer-M are 57M and 17G. And those of ViP-Medium/7 are 55M and 16.3G. Thus, the model size of ours is about 5% larger than ViP. But the performance is improved greatly from 82.7% to 83.7%. Therefore, the model structure should be attributed to such improvements. Specifically, the designs of dynamic mixing matrices and multi-segment fusion mechanism do help improve the performance. The mixing matrices generated dynamically could capture the contents of tokens to be mixed. Thus, different sets of tokens will be mixed with different mixing matrices. Therefore, token information could be fused more effectively. At last, the multi-segment fusion mechanism improves the robustness of the whole model. We also provide detailed ablation studies of these two designs in the following subsection for gaining a deeper understanding of our method.

Comparison with SOTA Classification Models. We also compare our model with well-known state-of-the-art models, which include CNN-based models, transformer-based models, and hybrid models. The results are shown in Table 3. We can see that our proposed models still achieve the best performance among models with a similar number of parameters. Specifically, our models achieve better performance than Swin Transformer (Liu et al., 2021b), which is the state-of-the-art transformer-based model. Our DynaMixer-S achieves accuracy 82.7% with slightly fewer parameters, while Swin-T achieved 81.3%. For medium-sized and large-sized models, our models are still better than Swin-S and Swin-B. Our model is also better than CrossFormer (Wang et al., 2021b), which employs a multi-scale strategy. We believe that by introducing multi-granularity, the performance of our method could be further improved. From the above comparisons, we can see that DynaMixer is a promising architecture for visual recognition tasks.

Model Latency Comparison. We test the throughput of our model, ViP, and ResMLP with a batch size set of 32 on a single NVIDIA V100 GPU. Table 4 shows these results. We can see that among models achieving similar top-1 accuracies, our model achieves the best throughput. For example, DynaMixer-S and ViP-M/7 have similar top-1 performance. The throughput of our model is 551 images/s, while the throughput of ViP-M/7 is only 390 images/s. Moreover, DynaMixer-S is also faster than ResMLP B24 with a higher top-1 accuracy. This indicates that compared to SOTA MLP-like models, our model achieves similar or better performance with fewer parameters.

Model	Family	Scale	Param	FLOPs	Top-1 (%)
ResNet18 (He et al., 2016)	CNN	224 ²	12M	1.8G	69.8
EffNet-B3 (Tan & Le, 2019)	CNN	300 ²	12M	1.8G	81.6
PVT-T (Wang et al., 2021a)	Trans	224 ²	13M	1.9G	75.1
GFNet-H-Ti (Rao et al., 2021)	FFT	224 ²	15M	2.0G	80.1
ResNet50 (He et al., 2016)	CNN	224 ²	26M	4.1G	78.5
RegNetY-4G (Radosavovic et al., 2020)	CNN	224 ²	21M	4.0G	80.0
DeiT-S (Touvron et al., 2021b)	Trans	224 ²	22M	4.6G	79.8
BoT-S1-50 (Srinivas et al., 2021)	Hybrid	224 ²	21M	4.3G	79.1
PVT-S (Wang et al., 2021a)	Trans	224 ²	25M	3.8G	79.8
T2T-14 (Yuan et al., 2021a)	Trans	224 ²	22M	4.8G	81.5
Swin-T (Liu et al., 2021b)	Trans	224 ²	29M	4.5G	81.3
GFNet-H-S (Rao et al., 2021)	FFT	224 ²	32M	4.5G	81.5
CrossFormer-T (Wang et al., 2021b)	Trans	224 ²	27M	2.9G	81.5
CrossFormer-S (Wang et al., 2021b)	Trans	224 ²	30M	4.9G	82.5
DynaMixer-S	MLP	224 ²	26M	7.3G	82.7
ResNet101 (He et al., 2016)	CNN	224 ²	45M	7.9G	79.8
RegNetY-8G (Radosavovic et al., 2020)	CNN	224 ²	39M	8.0G	81.7
BoT-S1-59 (Srinivas et al., 2021)	Hybrid	224 ²	34M	7.3G	81.7
PVT-M (Wang et al., 2021a)	Trans	224 ²	44M	6.7G	81.2
GFNet-H-B (Rao et al., 2021)	FFT	224 ²	54M	8.4G	82.9
Swin-S (Liu et al., 2021b)	Trans	224 ²	50M	8.7G	83.0
PVT-L (Wang et al., 2021a)	Trans	224 ²	61M	9.8G	81.7
CrossFormer-B (Wang et al., 2021b)	Trans	224 ²	52M	9.2G	83.4
T2T-24 (Yuan et al., 2021a)	Trans	224 ²	64M	13.8G	82.1
DynaMixer-M	MLP	224 ²	57M	17.0G	83.7
CrossFormer-L (Wang et al., 2021b)	Trans	224 ²	92M	16.1G	84.0
Swin-B (Liu et al., 2021b)	Trans	224 ²	88M	15.4G	83.3
DeiT-B (Touvron et al., 2021b)	Trans	224 ²	86M	17.5G	81.8
DeiT-B (Touvron et al., 2021b)	Trans	384 ²	86M	55.4G	83.1
DynaMixer-L	MLP	224 ²	97M	27.4G	84.3

Table 3. Top-1 accuracy comparisons with CNN-based models and transformer based models on the ImageNet-1K image classification benchmark. Note that all models are trained without external data. The ‘‘Scale’’ denotes the input size for training and testing stages. The ‘‘Param’’ and ‘‘FLOPs’’ columns denote the number of parameters and the number of floating point operations, respectively. The values in the column Family, *e.g.*, ‘‘CNN’’, ‘‘Trans’’, ‘‘Hybrid’’, and ‘‘FFT’’, mean CNN-based, transformer-based, CNN-and-transformer-based, and fast Fourier transform based models.

Limitations. The numbers of parameters and FLOPs of DynaMixer are slightly larger than those of other models, as the mixing matrix is generated dynamically. In practice, such an increment over model size will affect the training time. Compared with ViPs of a similar number of parameters, the FLOPs of our models are slightly bigger and the throughputs are about 16% ~ 18% lower. Moreover, the input image size for our model should be fixed, which restricts its applications on some downstream tasks, *e.g.*, object detection, and segmentation.

4.4. Ablation Study

In this subsection, we present ablation study results with DynaMixer-S to provide a deeper understanding of our methods.

Ablation on the number of segments S . To improve the robustness and generalization ability of our network, we introduce a multi-segment fusion mechanism in the DynaMixer operation. To show the effects of the number of segments, the performances with different values of segment dimensionalities, which is $\frac{D}{S}$, are reported in Table 5. Con-

DynaMixer: A Vision MLP Architecture with Dynamic Mixing

Model	# Param	FLOPs	Throughput image/s	Top-1 (%)
ResMLP S24	30M	6.0G	680	79.4
ResMLP B24	116M	23.0G	215	81.0
ViP-M/7	55M	16.3G	390	82.7
DynaMixer-S	26M	7.3G	551	82.7
ViP-L/7	88M	24.4G	276	83.2
DynaMixer-M	57M	17.0G	326	83.7

Table 4. The number of parameters, FLOPs and throughput of different models.

sidering the computational complexity, we only test 4 values of segment dimensionalities, which are D , 96, 48, and 24. We can see that even without multiple segments, our model can still obtain satisfying accuracy (82.2%). As S increases ($\frac{D}{S}$ decreases), the performance becomes better. Moreover, our network is not sensitive to segment dimensionalities in a quite large range.

Model	D/S	# parameters	Top-1 (%)
DynaMixer-S	D	26M	82.2
DynaMixer-S	96	26M	82.4
DynaMixer-S	48	26M	82.5
DynaMixer-S	24	26M	82.7

Table 5. The top-1 accuracies with different segment dimensionalities.

Ablation on weight generation methods. We study the effects of different methods to generate mixing matrices in this paragraph. The methods include ‘‘Synthesizer (Random)’’ and ‘‘Synthesizer (Dense)’’, which are described in (Tay et al., 2021). Once the mixing matrices are generated, they are applied to mixing tokens by following the procedure of DynaMixer, which means that tokens are mixed row by row and column by column with the multi-segment fusion mechanism. The mixing matrices generated by ‘‘Synthesizer (Random)’’ are actually trainable parameters. Thus it is similar to ViP (Hou et al., 2021), except that the information fusion among segments of channel is different. The mixing matrix generated by ‘‘Synthesizer (Dense)’’ does not consider the effects of other tokens. Thus, the mixing weight vector for a specific output token in ‘‘Synthesizer (Dense)’’ is generated merely with the information from the token at the same position. The performances of these methods are listed in Table 6. We can see that ‘‘Synthesize (Random)’’ achieved an accuracy of 81.5%, which is close to ViP’s. It is natural since their mixing matrices are generated in the same way. The accuracy achieved by ‘‘Synthesizer (Dense)’’, which is 20% larger than DynaMixer-S, is only 81.4%, which demonstrates that leveraging the contents of all tokens is necessary, and the technique of reducing the

computational complexity adopted in our model is effective.

Model	generating methods	# parameters	Top-1 (%)
DynaMixer-S	Synthesize(Random)	25M	81.5
DynaMixer-S	Synthesizer(Dense)	32M	81.4
DynaMixer-S	DynaMixer-S	26M	82.7

Table 6. The effects of different mixing matrix generation methods.

Ablation on the reduced dimensionality d . In our model, dimensionality is reduced to d by a liner projection to decrease the number of parameters. Ideally, d should be large to provide enough information. However, we found that d can be extremely low in practice. The top-1 accuracies with different values of d are shown in Table 7. We can see that our model with $d = 1$ can still achieve an accuracy of 82.4%, which is an absolute improvement of 0.9% compared with ViP. Thus, even dimensionality is reduced to 1, the reduced token feature could provide information of all tokens to some extent. The performances with values of d set to 2, 4, and 8 are similar. Thus, d is set to 2 for DynaMixer-S in our experiments.

Model	d	# parameters	Top-1 (%)
DynaMixer-S	1	26M	82.4
DynaMixer-S	2	26M	82.7
DynaMixer-S	4	27M	82.6
DynaMixer-S	8	29M	82.7

Table 7. The effects of reduced dimensionality.

Ablation on the designs of DynaMixer Block. In the DynaMixer block, the final results are from three components: row mixing, column mixing, and channel mixing. The final results are a weighted sum of the above three. In this paragraph, we study the effects of row mixing, column mixing, channel mixing, and reweighting by removing one of them in turn from our model while keeping the other components unchanged. The results are shown in Table 8. We can see that without row mixing or column mixing, the performance will decrease from 82.7% to about 79%. Moreover, if the channel mixing component is removed, the performance will drop 0.6% absolutely (from 82.7% to 82.1%). Without the reweighting step, the performance will also become worse. Thus, all the components in our DynaMixer block are necessary for our model. Finally, we evaluate the performance of our model when the parameters of DynaMixer operations for row mixing and column mixing are shared, which is a straightforward way to reduce the number of parameters. We can see that the accuracy only drops to 82.2%, which is still better than the other vision MLPs of similar size.

Model	# parameters	Top-1 (%)
DynaMixer-S	26M	82.7
- column mixing	23M	79.2
- row mixing	23M	79.4
- channel fusion	23M	82.1
- reweighting	24M	81.9
+ sharing DynaMixer op	23M	82.2

Table 8. The effects of different components in the DynaMixer block.

4.5. Transfer Learning

In this subsection, we show the advantages of our models on the transfer learning tasks. We transfer our pre-trained DynaMixer-S to downstream datasets such as CIFAR10 and CIFAR100. We finetune our DynaMixer-S with 60 epochs by using the SGD optimizer and cosine learning rate decay. The results are given in Table 9. We find that our DynaMixer can achieve higher accuracies than ViP and the original ViT with similar model sizes on the downstream datasets.

Model	dataset	# parameters	Top-1 (%)
ViT-S/16 (Dosovitskiy et al., 2021)	CIFAR-10	49M	97.1
T2T-14 (Yuan et al., 2021a)	CIFAR-10	22M	97.5
ViP-S/7 (Hou et al., 2021)	CIFAR-10	25M	98.0
DynaMixer-S	CIFAR-10	26M	98.2
ViT-S/16 (Dosovitskiy et al., 2021)	CIFAR-100	49M	87.1
T2T-14 (Yuan et al., 2021a)	CIFAR-100	22M	88.4
ViP-S/7 (Hou et al., 2021)	CIFAR-100	25M	88.4
DynaMixer-S	CIFAR-100	26M	88.6

Table 9. The results of fine-tuning the pre-trained DynaMixer on ImageNet to downstream datasets: CIFAR10 and CIFAR100.

5. Conclusion

MLP-based networks have recently been designed for vision recognition, targeting at introducing less inductive biases. The performances of MLP-based networks have been shown better than or comparable with transformer-based models. Existing MLP-based vision models utilize static trainable mixing matrices to perform token information fusion. In this paper, we proposed a novel MLP-based vision architecture, called DynaMixer, which further improves the visual recognition performance. In DynaMixer, the mixing matrices are generated dynamically by leveraging the contents of all the tokens to be mixed. To reduce the time complexity and meanwhile improve the robustness, we exerted a dimensionality reduction technique and a multi-segment fusion mechanism. Our DynaMixer model has been demonstrated to achieve state-of-the-art performance on the ImageNet recognition task. We also provided an insightful analysis for gaining a deeper understanding of our model.

References

- Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL <https://www.tensorflow.org/>.
- Bai, J., Yuan, L., Xia, S.-T., Yan, S., Li, Z., and Liu, W. Improving vision transformers by revisiting high-frequency components. *arXiv preprint arXiv:2204.00993*, 2022.
- Chen, C.-F., Fan, Q., and Panda, R. Crossvit: Cross-attention multi-scale vision transformer for image classification. *arXiv preprint arXiv:2103.14899*, 2021a.
- Chen, S., Xie, E., Ge, C., Liang, D., and Luo, P. Cyclemlp: A mlp-like architecture for dense prediction. *arXiv preprint arXiv:2107.10224*, 2021b.
- Cubuk, E. D., Zoph, B., Shlens, J., and 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.
- Deng, J., Dong, W., Socher, R., Li, L.-J., Li, K., and 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.
- 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. In *International Conference on Learning Representations*, 2021.
- Fan, H., Xiong, B., Mangalam, K., Li, Y., Yan, Z., Malik, J., and Feichtenhofer, C. Multiscale vision transformers. *arXiv preprint arXiv:2104.11227*, 2021.
- Ge, C., Liang, Y., Song, Y., Jiao, J., Wang, J., and Luo, P. Revitalizing cnn attention via transformers in self-supervised visual representation learning. *Advances in Neural Information Processing Systems*, 2021.
- Guo, J., Tang, Y., Han, K., Chen, X., Wu, H., Xu, C., Xu, C., and Wang, Y. Hire-mlp: Vision mlp via hierarchical rearrangement. *arXiv preprint arXiv:2108.13341*, 2021.

- Han, K., Wang, Y., Chen, H., Chen, X., Guo, J., Liu, Z., Tang, Y., Xiao, A., Xu, C., Xu, Y., et al. A survey on visual transformer. *arXiv preprint arXiv:2012.12556*, 2020.
- Han, K., Xiao, A., Wu, E., Guo, J., Xu, C., and Wang, Y. Transformer in transformer. *arXiv preprint arXiv:2103.00112*, 2021.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Hou, Q., Jiang, Z., Yuan, L., Cheng, M.-M., Yan, S., and Feng, J. Vision permutator: A permutable mlp-like architecture for visual recognition. *arXiv preprint arXiv:2106.12368*, 2021.
- Hu, H., Zhang, Z., Xie, Z., and Lin, S. Local relation networks for image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3464–3473, 2019.
- Hu, J., Shen, L., and Sun, G. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.
- Huang, G., Sun, Y., Liu, Z., Sedra, D., and Weinberger, K. Q. Deep networks with stochastic depth. In *European conference on computer vision*, pp. 646–661. Springer, 2016.
- Huang, G., Liu, Z., Van Der Maaten, L., and Weinberger, K. Q. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.
- Jiang, Z.-H., Hou, Q., Yuan, L., Zhou, D., Shi, Y., Jin, X., Wang, A., and Feng, J. All tokens matter: Token labeling for training better vision transformers. *Advances in Neural Information Processing Systems*, 34, 2021.
- Khan, S., Naseer, M., Hayat, M., Zamir, S. W., Khan, F. S., and Shah, M. Transformers in vision: A survey. *arXiv preprint arXiv:2101.01169*, 2021.
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105, 2012.
- Li, J., Hassani, A., Walton, S., and Shi, H. Convmlp: Hierarchical convolutional mlps for vision. *arXiv preprint arXiv:2109.04454*, 2021.
- Lian, D., Yu, Z., Sun, X., and Gao, S. As-mlp: An axial shifted mlp architecture for vision. *arXiv preprint arXiv:2107.08391*, 2021.
- Liang, Y., Ge, C., Tong, Z., Song, Y., Wang, J., and Xie, P. Not all patches are what you need: Expediting vision transformers via token reorganizations. *arXiv preprint arXiv:2202.07800*, 2022.
- Liu, H., Dai, Z., So, D. R., and Le, Q. V. Pay attention to mlps. *arXiv preprint arXiv:2105.08050*, 2021a.
- Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., and Guo, B. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021b.
- Loshchilov, I. and Hutter, F. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. 2019.
- Radosavovic, I., Kosaraju, R. P., Girshick, R., He, K., and Dollár, P. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10428–10436, 2020.
- Rao, Y., Zhao, W., Zhu, Z., Lu, J., and Zhou, J. Global filter networks for image classification. *arXiv preprint arXiv:2107.00645*, 2021.
- Srinivas, A., Lin, T.-Y., Parmar, N., Shlens, J., Abbeel, P., and Vaswani, A. Bottleneck transformers for visual recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 16519–16529, 2021.
- Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., and Rabinovich, A. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1–9, 2015.
- Szegedy, C., Vanhoucke, V., Ioffe, S., Shlens, J., and 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.
- Szegedy, C., Ioffe, S., Vanhoucke, V., and Alemi, A. A. Inception-v4, inception-resnet and the impact of residual connections on learning. In *Thirty-first AAAI conference on artificial intelligence*, 2017.

- Tan, M. and Le, Q. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114. PMLR, 2019.
- Tang, C., Zhao, Y., Wang, G., Luo, C., Xie, W., and Zeng, W. Sparse mlp for image recognition: Is self-attention really necessary? *arXiv preprint arXiv:2109.05422*, 2021.
- Tatsunami, Y. and Taki, M. Raftmlp: Do mlp-based models dream of winning over computer vision? *CoRR*, abs/2108.04384, 2021.
- Tay, Y., Bahri, D., Metzler, D., Juan, D.-C., Zhao, Z., and Zheng, C. Synthesizer: Rethinking self-attention for transformer models. In *International Conference on Machine Learning*, pp. 10183–10192. PMLR, 2021.
- Tolstikhin, I., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.
- Tong, Z., Song, Y., Wang, J., and Wang, L. Videomae: Masked autoencoders are data-efficient learners for self-supervised video pre-training. *arXiv preprint arXiv:2203.12602*, 2022.
- Touvron, H., Bojanowski, P., Caron, M., Cord, M., El-Nouby, A., Grave, E., Izacard, G., Joulin, A., Synnaeve, G., Verbeek, J., et al. Resmlp: Feedforward networks for image classification with data-efficient training. *arXiv preprint arXiv:2105.03404*, 2021a.
- Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., and Jégou, H. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pp. 10347–10357. PMLR, 2021b.
- Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., and Jégou, H. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239*, 2021c.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Wang, W., Xie, E., Li, X., Fan, D.-P., Song, K., Liang, D., Lu, T., Luo, P., and Shao, L. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. *arXiv preprint arXiv:2102.12122*, 2021a.
- Wang, W., Yao, L., Chen, L., Lin, B., Cai, D., He, X., and Liu, W. Crossformer: A versatile vision transformer hinging on cross-scale attention. *arXiv preprint arXiv:2108.00154*, 2021b.
- Wang, X., Girshick, R., Gupta, A., and He, K. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7794–7803, 2018.
- Xie, S., Girshick, R., Dollár, P., Tu, Z., and He, K. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.
- Yu, T., Li, X., Cai, Y., Sun, M., and Li, P. S²-mlp: Spatial-shift mlp architecture for vision. *arXiv preprint arXiv:2106.07477*, 2021a.
- Yu, T., Li, X., Cai, Y., Sun, M., and Li, P. S²-mlpv2: Improved spatial-shift mlp architecture for vision. *arXiv preprint arXiv:2108.01072*, 2021b.
- Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.-H., Tay, F. E., Feng, J., and Yan, S. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 558–567, October 2021a.
- Yuan, L., Hou, Q., Jiang, Z., Feng, J., and Yan, S. Volo: Vision outlooker for visual recognition. *arXiv preprint arXiv:2106.13112*, 2021b.
- Yun, S., Han, D., Oh, S. J., Chun, S., Choe, J., and 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.
- Zhang, H., Cisse, M., Dauphin, Y. N., and Lopez-Paz, D. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Zhang, H., Wu, C., Zhang, Z., Zhu, Y., Lin, H., Zhang, Z., Sun, Y., He, T., Mueller, J., Manmatha, R., et al. Resnest: Split-attention networks. *arXiv preprint arXiv:2004.08955*, 2020.
- Zhong, Z., Zheng, L., Kang, G., Li, S., and Yang, Y. Random erasing data augmentation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pp. 13001–13008, 2020.
- Zhou, D., Kang, B., Jin, X., Yang, L., Lian, X., Jiang, Z., Hou, Q., and Feng, J. Deepvit: Towards deeper vision transformer. *arXiv preprint arXiv:2103.11886*, 2021a.
- Zhou, D., Shi, Y., Kang, B., Yu, W., Jiang, Z., Li, Y., Jin, X., Hou, Q., and Feng, J. Refiner: Refining self-attention for vision transformers. *arXiv preprint arXiv:2106.03714*, 2021b.