

ResFormer: Scaling ViTs with Multi-Resolution Training

Rui Tian^{1,2} Zuxuan Wu^{1,2†} Qi Dai³ Han Hu³ Yu Qiao⁴ Yu-Gang Jiang^{1,2}

¹Shanghai Key Lab of Intell. Info. Processing, School of CS, Fudan University

²Shanghai Collaborative Innovation Center of Intelligent Visual Computing

³Microsoft Research Asia ⁴Shanghai AI Laboratory

Abstract

Vision Transformers (ViTs) have achieved overwhelming success, yet they suffer from vulnerable resolution scalability, i.e., the performance drops drastically when presented with input resolutions that are unseen during training. We introduce, ResFormer, a framework that is built upon the seminal idea of multi-resolution training for improved performance on a wide spectrum of, mostly unseen, testing resolutions. In particular, ResFormer operates on replicated images of different resolutions and enforces a scale consistency loss to engage interactive information across different scales. More importantly, to alternate among varying resolutions effectively, especially novel ones in testing, we propose a global-local positional embedding strategy that changes smoothly conditioned on input sizes. We conduct extensive experiments for image classification on ImageNet. The results provide strong quantitative evidence that ResFormer has promising scaling abilities towards a wide range of resolutions. For instance, ResFormer-B-MR achieves a Top-1 accuracy of 75.86% and 81.72% when evaluated on relatively low and high resolutions respectively (i.e., 96 and 640), which are 48% and 7.49% better than DeiT-B. We also demonstrate, moreover, ResFormer is flexible and can be easily extended to semantic segmentation, object detection and video action recognition.

1. Introduction

The strong track record of Transformers in a multitude of Natural Language Processing [53] tasks has motivated an extensive exploration of Transformers in the computer vision community. At its core, Vision Transformers (ViTs) build upon the multi-head self-attention mechanisms for feature learning through partitioning input images into patches of identical sizes and processing them as sequences for dependency modeling. Owing to their strong capabilities

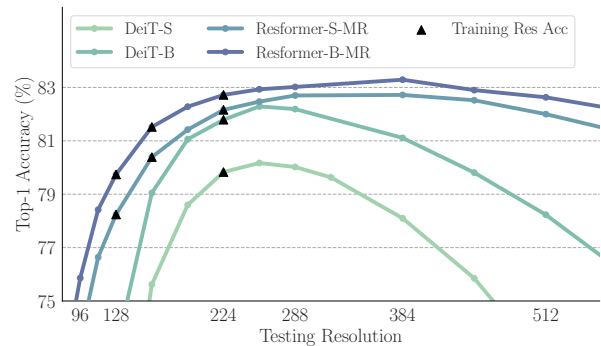


Figure 1. Comparisons between ResFormer and vanilla ViTs. ResFormer achieves promising results on a wide range of resolutions.

ities in capturing relationships among patches, ViTs and their variants demonstrate prominent results in versatile visual tasks, e.g., image classification [36, 50, 65, 70], object detection [4, 30, 55], vision-language modeling [25, 40, 54] and video recognition [3, 29, 37, 64].

While ViTs have been shown effective, it remains unclear how to scale ViTs to deal with inputs with varying sizes for different applications. For instance, in image classification, the *de facto* training resolution of 224 is commonly adopted [36, 50, 51, 65]. However, among works in pursuit of reducing the computational cost of ViTs [39, 43], shrinking the spatial dimension of inputs is a popular strategy [6, 32, 56]. On the other hand, fine-tuning with higher resolutions (e.g., 384) is widely used [15, 36, 48, 51, 59, 62] to produce better results. Similarly, dense prediction tasks such as semantic segmentation and object detection also require relatively high resolution inputs [1, 30, 35, 55].

Despite of the necessity for both low and high resolutions, limited effort has been made to equip ViTs with the ability to handle different input resolutions. Given a novel resolution that is different from that used during training, a common practice adopted for inference is to keep the patch size fixed and then perform bicubic interpolation on positional embeddings directly to the corresponding scale. As shown in Sec. 3, while such a strategy is able to scale ViTs to relatively larger input sizes, the results on low resolutions

[†]Corresponding author.

Note that we use resolution, scale and size interchangeably.

plunge sharply. In addition, significant changes between training and testing scales also lead to limited results (*e.g.*, DeiT-S trained on a resolution of 224 degrades by 1.73% and 7.2% when tested on 384 and 512 respectively).

Multi-resolution training, which randomly resizes images to different resolutions, is a promising way to accommodate varying resolutions at test time. While it has been widely used by CNNs for segmentation [22], detection [24] and action recognition [58], generalizing such an idea to ViTs is challenging and less explored. For CNNs, thanks to the stacked convolution design, all input images, regardless of their resolutions, share the same set of parameters in multi-resolution training. For ViTs, although it is feasible to share parameters for all samples, bicubic interpolations of positional embeddings, which are not scale-friendly, are still needed when iterating over images of different sizes.

In this paper, we posit that positional embeddings of ViTs should be adjusted smoothly across different scales for multi-resolution training. The resulting model then has the potential to scale to different resolutions during inference. Furthermore, as images in different scales contain objects of different sizes, we propose to explore useful information across different resolutions for improved performance in a similar spirit to feature pyramids, which are widely used in hierarchical backbone designs for both image classification [24, 36] and dense prediction tasks [22, 23, 33].

To this end, we introduce ResFormer, which takes in inputs as multi-resolution images during training and explores multi-scale clues for better results. Trained in a single run, ResFormer is expected to generalize to a large span of testing resolutions. In particular, given an image during training, ResFormer resize it to different scales, and then use all scales in the same feed-forward process. To encourage information interaction among different resolutions, we introduce a scale consistency loss, which bridges the gap between low-resolution and high-resolution features by self-knowledge distillation. More importantly, to facilitate multi-resolution training, we propose a global-local positional embedding strategy, which enforces parameter sharing and changes smoothly across different resolutions with the help of convolutions. Given a novel resolution at testing, ResFormer dynamically generates a new set of positional embeddings and performs inference.

To validate the efficacy of ResFormer, we conduct comprehensive experiments on ImageNet-1K [13]. We observe that ResFormer makes remarkable gains compared with vanilla ViTs which are trained on single resolution. Given the testing resolution of 224, ResFormer-S-MR trained on resolutions of 128, 160 and 224 achieves a Top-1 accuracy of 82.16%, outperforming the 224-trained DeiT-S [50] by 2.24%. More importantly, as illustrated in Fig. 1, ResFormer surpasses DeiT by a large margin on unseen resolutions, *e.g.*, ResFormer-S-MR outperforms DeiT-S by

6.67% and 56.04% when tested on 448 and 80 respectively. Furthermore, we also validate the scalability of ResFormer on dense prediction tasks, *e.g.*, ResFormer-B-MR achieves 48.30 mIoU on ADE20K [72] and 47.6 AP^{box} on COCO [34]. We also show that ResFormer can be readily adapted for video action recognition with different sizes of inputs via building upon TimeSFormer [3].

2. Related Work

Scaling Vision Models. Many studies in recent literature [35, 44, 49, 67] discuss how to scale vision models, with most of them focusing on the capacity of deep neural networks. For instance, EfficientNet [49] studies how model width, depth and input resolution affect convolutional neural networks. RegNet [41] designs manual designing space for CNNs and finds simple linear correlation between the search space (*e.g.* width) and performance. ResNet-RS [2] presents how different scaling strategies on depth and input resolution can affect the model capacity.

Recent approaches have investigated scaling of transformers [35, 67]. For example, V-Moe [44] scales vision transformers to large model sizes with sparse mixture-of-experts. Several studies [17, 19, 27, 63] explore the aspect of data scaling under self-supervised framework, *i.e.*, how data sizes affect the model performance. In contrast, limited effort has been made towards the scaling abilities of models towards input resolutions. Attempts have been made by Liu *et al.* [35] to scale up to larger resolutions, while neglecting lower resolutions. Instead, our work takes the initiative to scale ViTs to various resolutions, both lower and higher, satisfying the practical needs from varied visual tasks.

Positional embedding. The self-attention architecture is clueless about spatial relationships among patches. Therefore, to overcome permutation-invariance, various positional embedding strategies have been proposed to enable Transformers to perceive the sequence order of input tokens. Absolute positional embeddings (APE) infuse global spatial information into Transformers, *e.g.*, sine-cosine APE [53] proposed for NLP tasks and learned APE adopted in the vanilla Vision Transformer [15]. Meanwhile, efficacy of relative position embeddings (RPB) is widely validated in both language [12, 45] and vision tasks [5, 35, 60]. For instance, Wu *et al.* replaces APE with the relative strategy of iRPE [60] for performance gains on classification and detection. ConViT [16] also suggests that adding gated relative positional embeddings to self-attention blocks brings about soft convolutional inductive biases. Moreover, dynamic positional embeddings are introduced to model local information from input tokens, *e.g.*, Twins [8] adopt conditional positional embeddings (CPE) [9] implemented by convolutions. In order to improve performance and scalability simultaneously, we propose to inject spatial embeddings




Figure 2. Top-1 accuracy of DeiT-S trained with 5 different resolutions and tested on resolutions varying from 80 to 576. During testing, we follow the common pre-processing steps (*i.e.*, `Resize` and `CenterCrop` in Pytorch implementation) and set the cropping rate to 0.875.

in ResFormer from both from global and local perspectives.

Multi-scale training. In early CNNs, multi-scale data augmentations [46] are employed for image classification by randomly sampling training images from a certain range of scales. Later in dense prediction tasks, multi-scale training and testing become a widely-adopted paradigm [4, 28, 47]. In addition, the idea has also been explored in action recognition. Wu *et al.* introduce a multigrid strategy [58], which enables efficient training by sampling data with different grids of temporal span, spatial span and temporal stride. Video ResKD [38] achieves excellent efficiency by employing high-resolution features from large models as teachers to improve low-resolution performance.

Most approaches in multi-scale training rely on CNNs, as convolutions can be readily applied to varying sizes of inputs. In contrast, vanilla ViTs are equipped with tokens of fixed-dimension, other related attempts lay emphasis on multi-scale spatial dimension of features instead of input [18, 20] or perform in an unsupervised way [42], yet limited effort has been made to explore multi-resolution supervised training for ViTs. In this paper, we make the first step to investigate such a strategy for ViTs which not only leads to good performance on training resolutions but can also generalize towards novel resolutions.

3. Resolution Generalization

In this section, we conduct a set of pilot experiments to show the scalability of ViTs towards different resolutions.

Generalizing to different resolutions. As revealed by previous work [52], CNNs suffer from distribution shifts between training and testing due to different pre-processing methods, *i.e.*, the *de facto* “random resizing and cropping” strategy for training and “center cropping” for testing result in different distributions of cropped regions in images. For ViTs, theoretically, the discrepancy persists since the

same pre-processing strategies of training and testing are employed. However, there lacks a comprehensive study on how ViTs behave towards input scales varied from the training process. To this end, we feed pre-trained ViTs with testing samples of varying sizes. In particular, we instantiate ViT models with DeiT-S [50] and initialize the model with weights pre-trained on ImageNet-1K. We then fine-tune the model on a resolution of 96, 128, 224, 288 and 384 respectively. These derived models are then tested on a broad spectrum of resolutions. Following [36, 50], we simply resize the position embeddings with bicubic interpolation on different testing resolutions. The results are shown in Fig. 2. We observe the following trends for scaling up or down:

- **Scaling down:** All models undergo severe performance drop when directly adapted to small-scale inputs, especially for ones pre-trained on larger resolutions. For example, the Top-1 accuracy of DeiT-S with a training resolution of 384 decreases by 6.18% when tested on 224 and even plummets below 30% when the testing resolution is further reduced to 160.
- **Scaling up:** Ideally, increasing testing resolutions results in improved accuracy, which is also suggested as byproduct of train-test distribution discrepancy in [52]. However, models yield unsatisfactory performance when gap enlarges. *e.g.*, DeiT-S with a low training resolution of 128 stops growing in accuracy when the testing resolution reaches 256. It achieves a Top-1 accuracy of 75.26% with testing resolution set to 224, which is 4.57% lower than model trained with a resolution of 224, directly.

Above all, ViTs are vulnerable to resolution discrepancies between training and testing, particularly when evaluated on low-resolution inputs. This motivates us to equip ViTs with scalability towards a wide range of test resolutions so as to meet the need of versatile applications.

4. Method

Our goal is to train a vision transformer that not only performs well on resolutions the network has seen during training, but more importantly it is able to adapt to a wide range of unseen resolutions without significant performance drop during testing. To this end, we first introduce a resolution scaling transformer in Sec. 4.1, ResFormer, which operates on input samples of multiple resolutions in the training stage. Since the size of objects varies in different scales, we also introduce a scale consistency loss to fully explore information from all resolutions for improved accuracy. Furthermore, as mentioned in Sec. 3, directly interpolating positional embeddings to unseen resolutions during inference produces unsatisfactory results. To mitigate this issue, ResFormer builds upon carefully designed global-local positional embeddings, which are generated conditioned on input resolutions, as will be described in Sec. 4.2.

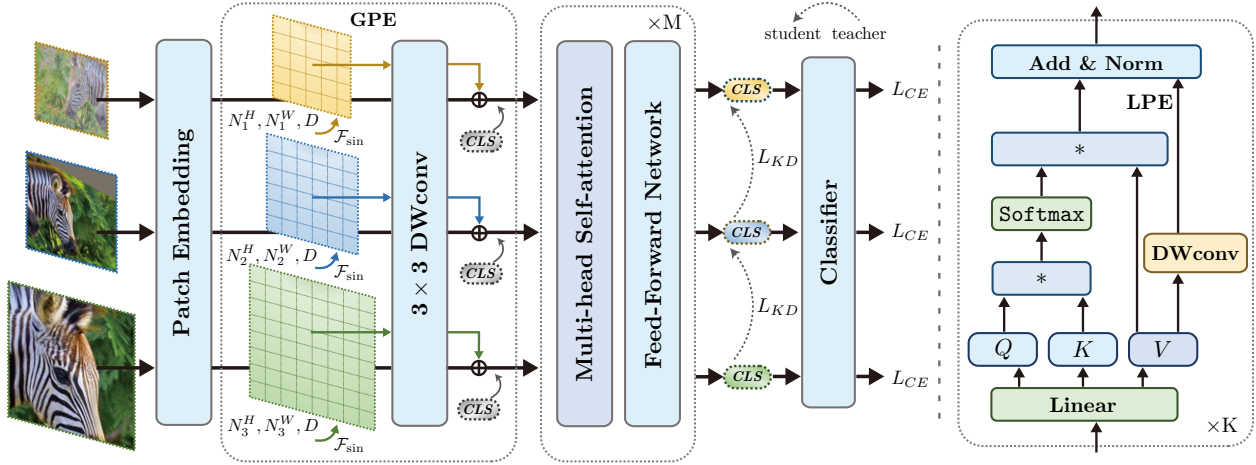


Figure 3. **Left:** The overview of ResFormer framework. **Right:** The pipeline of generating local positional embedding.

4.1. ResFormer

Following the vanilla ViT [15], given an input image X whose height and width are H and W , respectively, we first split it into $N^H \times N^W$ patches, where the patch size is set to t and $N_H = H/t, N_W = W/t$. Each image patch is projected into a D -dimension feature by patch embedding and is denoted as a “token”. Subsequently, a global class token cls is concatenated with image tokens before they are fed into Transformer blocks.

Unlike standard ViTs operating on single-scale images, ResFormer takes inputs of different resolutions during training so as to better model objects of varying sizes in different scales and generalize better during inference. More specifically, as shown in Fig. 3, we replicate a given training image for r times, where r denotes the number of resolutions used. For the i -th data replica, resizing and cropping operations¹ are applied to obtain a training sample X^i with a spatial size of $3 \times H_i \times W_i$. Afterwards, we apply random pre-processing strategies involved in ViTs training paradigm² on each scale of inputs separately. As a result, one mini-batch is composed of groups of multi-resolution inputs sharing identical labels, which is roughly equivalent to extending the base batch size by r times. In addition, for the input sample X^i , we feed the global class token output of the last transformer block as inputs into classification head to compute final predictions Y^i . Naturally, the classification losses can be written as:

$$L(\Theta) = \mathbb{E}_{(X,T) \sim \mathcal{D}} \sum_{i=1}^r L_{CE}(Y^i; X^i, T, \Theta), \quad (1)$$

where T denotes the ground-truth label and L_{CE} represents

¹In practice, we realize it with `RandomResizedCrop` in PyTorch.

²Random pre-processing includes Auto-Augment [10], RandAugment [11], random erasing [71], MixUp [68] and CutMix [66].

the cross-entropy loss. In addition, \mathcal{D} and Θ denote the training set and the parameters of the network, respectively.

Scale consistency loss. Given that larger inputs generally produce better recognition results compared to their smaller counterparts, we take advantage of knowledge distillation through enforcing consistencies among different resolutions. In particular, we use a smooth l_1 loss with feature whitening [57], denoted as L_{KD} , to transfer knowledge from the class token of a higher resolution to that of a lower resolution. This is achieved by serving cls_i as the teacher of cls_{i+1} with $H_i = W_i, H_i > H_{i+1}$. Combining with Eq. (1), the loss can be written as:

$$L(\Theta) = \mathbb{E}_{(X,T) \sim \mathcal{D}} \frac{1}{r} \left[\sum_{i=1}^r L_{CE}(Y^i; X^i, T, \Theta) + \sum_{i=1}^{r-1} L_{KD}(\text{cls}_{i+1}, \text{cls}_i) \right]. \quad (2)$$

Especially, teacher class tokens are detached from the gradient computational graph. At last, the loss is divided by r to ensure stability of training.

4.2. Global-Local Positional Embedding

The commonly-used positional embeddings highly depend on the size of input samples. As a result, when multiple resolutions are involved in the training process, positional embeddings need to be carefully adjusted when iterating images of different scales, as simple interpolations incur performance drops. Therefore, we propose to use conditional positional embeddings both globally and locally to bridge the resolution gap among a broad range of resolutions. Below, we first introduce the global positional embedding and then describe its local counterpart.

Global position embedding. To incorporate location information in patch embeddings, a typical way is to add absolute position embedding (APE). Given the input sample X , x^{img} refers to the output image tokens of the patch embedding, whose spatial dimension equals $N_H \times N_W$ and feature dimension is D . For simplicity, we denote x as concatenation of class token `cls` and image tokens x^{img} , the absolute positional embedding p can be expressed as,

$$x = x + p, \quad p \in \mathbb{R}^{1 \times (N_H \times N_W + 1) \times D}. \quad (3)$$

The most straightforward way is implementing p with learned parameters, as widely-adopted in [15, 50, 55]. Another common tactic is fulfilled with sinusoidal mapping $\mathcal{F}_{\text{sine}}$ [21, 53], through which p is generated on-the-fly by a fixed function dependent on N_H , N_W and D . Furthermore, compared with learned APE, the sine-cosine APE changes more weakly between different input scales, as displayed in Sec. 5.2. Therefore, we build our method upon sine-cosine APE with the assumption that smoother positional embedding would contribute to better resolution scalability.

We further improve the sine-cosine positional embedding with conditional computation such that the embeddings are tailored to the model during training. As illustrated in Fig. 3, a simple yet effective depth-wise convolution is applied so as to generate the final positional embedding conditioned on sinusoidal encoding. Since convolutions should be performed in 2-D dimension, we leave out the class token by concatenating a zero padding shaped of $\mathbb{R}^{1 \times 1 \times D}$ with output embeddings of DWconv. In general, the strategy introduced above aims at injecting smooth spatial information of global context into ViT, thereby we denote it as global positional embedding (GPE).

Local positional embedding. Positional embeddings introduced in [14, 29, 69] share the same design philosophy since they are both dynamically generated by input tokens and carry spatial information of local neighbourhood. It has been unveiled that such strategies can effectively introduce translation invariance into ViTs and hence facilitate generalizing to various resolutions. We refer to the positional embedding conditioned on local input feature as local positional embedding (LPE) and hypothesize that LPE is orthogonal with GPE in modelling spatial information of image tokens. Consequently, the combination of LPE and GPE may result in best resolution scalability.

To this end, we incorporate local positional embeddings into attention blocks in a similar fashion to [14]. Given a multi-head self-attention block, a query Q , a key K and a value V are obtained through a linear projection, and the output z can be derived as:

$$z = \text{Softmax}(QK^T/\sqrt{D})V. \quad (4)$$

In particular, local spatial information of V is utilized. We first set the class token aside and reshape the value ma-

trix to get $V' \in \mathbb{R}^{M \times D' \times N_H \times N_W}$, where M denotes the number of attention heads and D' satisfies $D = D' \cdot M$. Inspired by [35, 36], we generate dynamic positional embeddings conditioned on V' separately for each head. Therefore, a 3×3 depth-wise convolution is implemented to obtain the LPE for each head. The above operations can be denoted as mapping \mathcal{H} conditioned on V . Therefore, Eq. (4) can be re-written as:

$$z = \text{Softmax}(QK^T/\sqrt{D})V + \mathcal{H}(V). \quad (5)$$

By virtue of convolutions, LPE can be dynamically generated regardless of input scales. Eventually, in ResFormer, global and local positional embeddings are combined to ensure better generalization to novel resolutions.

5. Experiments

Implementation details. We instantiate ResFormer with DeiT [50] due to its simplicity. Given an input image, we resize it to 128, 160 and 224, respectively, for multi-resolution training throughout the experiments, unless specified otherwise. The resulting images are then used as inputs to ResFormer. For image classification, we use AdamW [26] as our optimizer and apply a cosine decay learning rate scheduler. Small and tiny models are trained with a batch size of 1024 and a learning rate of $5e^{-4}$, yet a learning rate of $8e^{-4}$ is used for the base model. We keep all augmentation and regularization settings in [50] for fair comparisons. For all experiments, we follow the official training and testing split as well as the evaluation metrics. For testing, we report results on a wide range of resolutions. Note that ResFormer only uses a single scale during testing.

5.1. Main Results

Effectiveness of ResFormer in image classification. Tab. 1 presents the results of ResFormer and comparisons with DeiT [50] using various settings. In particular, we use ResFormer-M-R to denote a variant of ResFormer, where M represents the model size (*i.e.*, T, S, B for tiny, small and base models respectively) and R indicates the resolution used for training (*i.e.*, MR denotes multiple resolution; if R is a number, it represents the resolution itself).

When evaluated with a testing resolution of 224, ResFormers achieve highly competitive results—ResFormer-S-MR and ResFormer-B-MR offers an accuracy of 82.16% and 82.72%, respectively, outperforming their DeiT counterparts by 2.33% and 0.93%. We also see from Tab. 1 that ResFormer trained with multi-resolution images outperforms models trained with single scale inputs with clear margins on all “seen” resolutions. For instance, ResFormer-S-MR outperforms ResFormer-S-128, ResFormer-S-160, ResFormer-S-224 by 2.77%, 2.33% and 1.33% respectively. Similar trends can also be found for ResFormer-

Table 1. Top-1 Accuracy of DeiT and ResFormer on ImageNet-1K. Columns highlighted with grey background refer to the training resolutions of given models. Specifically, ResFormer adopts training resolutions of 128, 160 and 224 for multi-resolution training.

Model	Testing resolution										
	96	112	128	160	192	224	288	384	448	512	640
DeiT-T [50]	8.06	34.22	52.16	65.68	70.18	72.14	73.1	71.29	67.43	66.07	59.31
ResFormer-T-MR	61.40	64.93	67.78	71.09	72.97	73.85	74.85	75.04	74.39	73.77	71.65
DeiT-S [50]	17.55	54.34	67.02	75.62	78.60	79.83	80.02	78.10	75.85	72.63	63.86
ResFormer-S-128	70.25	73.91	75.47	77.06	77.48	76.89	74.78	69.55	64.54	58.34	45.25
ResFormer-S-160	67.34	72.26	75.05	78.06	78.94	79.19	78.25	74.86	71.38	66.65	54.77
ResFormer-S-224	57.80	66.36	71.35	76.99	79.63	80.83	81.42	80.65	79.28	77.73	73.26
ResFormer-S-MR	73.59	76.64	78.24	80.39	81.42	82.16	82.70	82.72	82.52	82.00	80.72
DeiT-B [50]	27.86	64.46	73.18	79.05	81.06	81.79	82.19	81.11	79.81	78.23	74.23
ResFormer-B-MR	75.86	78.42	79.74	81.52	82.28	82.72	83.02	83.29	82.9	82.63	81.72

Table 2. Results and comparisons of different backbones on ADE20K. All backbones are pre-trained on ImageNet-1k, among which MAE [21] uses unsupervised pre-training.

Backbone	#Param	Lr schd	mIoU	ms + flip
DeiT-S [50]	52.1M	80k	42.96	43.79
XCiT-S12/16 [1]	52.4M	160k	45.90	46.72
ResFormer-S-224	51.7M	80k	45.47	46.61
ResFormer-S-MR	51.7M	80k	46.31	47.45
DeiT-B [50]	120.6M	160k	45.36	47.16
XCiT-S24/16 [1]	109.0M	160k	47.69	48.57
ViT-B + MAE [21]	176.5M	160k	48.13	48.70
ResFormer-B-MR	119.8M	160k	48.30	49.28

B and ResFormer-T. This highlights the effectiveness of multi-resolution training.

Furthermore, for “unseen” resolutions, ResFormer demonstrates clear scaling capabilities. In particular, given a test resolution of 384, ResFormer-S-224 achieves a Top-1 accuracy of 80.65%, which is 2.55% higher than its DeiT-S counterpart (78.10%). This suggests that global-local positional embeddings can indeed improve generalization of different resolutions. ResFormer-S-MR further boosts the accuracy to 82.72%, demonstrating the benefit of multi-resolution training. Besides, ResFormer consistently generalize well to lower resolutions. Compared with DeiT, ResFormer-S-MR and ResFormer-B-MR increase by 56.24% and 48.00% when evaluated on a resolution of 80, highlighting the effectiveness of ResFormer when dealing with significant resolution shifts during inference.

Semantic segmentation. To show flexibility of ResFormer, We evaluate for semantic segmentation on ADE20K [72] with UperNet [61]. As shown in Tab. 2, ResFormer-S-224 improves DeiT-S by 2.51 measured by mIoU. Both ResFormer-S-MR and ResFormer-B-MR, which are pre-trained with the multi-resolution strategy, achieve better re-

Table 3. Results and comparisons of different backbones on the mini-val set of COCO2017 using Mask R-CNN [22] and 3× training schedule. All backbones are pre-trained on ImageNet-1k in the supervised setting. Part of results are credited to [1, 7].

Backbone	#Param	AP ^b	AP ^b ₅₀	AP ^b ₇₅	AP ^m	AP ^m ₅₀	AP ^m ₇₅
PVT-Small [55]	44.1M	43.0	65.3	46.9	39.9	62.5	42.8
XCiT-S12/16 [1]	44.3M	45.3	67.0	49.5	40.8	64.0	43.8
ViT-S [31]	43.8M	44.0	66.9	47.8	39.9	63.4	42.2
ViTDet-S [30]	45.7M	44.5	66.9	48.4	40.1	63.6	42.5
ResFormer-S-MR	45.6M	46.4	68.5	50.4	40.7	64.7	43.4
PVT-Large [55]	81.0M	44.5	66.0	48.3	40.7	63.4	43.7
XCiT-M24/16 [1]	101.1M	46.7	68.2	51.1	42.0	65.6	44.9
ViT-B [31]	113.6M	45.8	68.2	50.1	41.3	65.1	44.4
ViTDet-B [30]	121.3M	46.3	68.6	50.5	41.6	65.3	44.5
ResFormer-B-MR	115.3M	47.6	69.0	52.0	41.9	65.9	44.4

sults. In particular, ResFormer-S-MR reaches up to 47.45 and ResFormer-B-MR hits the peak of 49.28 mIoU. This suggests that ResFormer effectively models multi-scale and high-resolution features for pixel-level dense predictions. Note that ResFormer-B-MR and ResFormer-S-MR are directly used as pre-trained backbones and we do not perform multi-resolution fine-tuning on ADE20K, since segmentation tasks already require images with a size of 512×512 as inputs, and multi-resolution training would be computationally expensive. Nonetheless, results in Tab. 2 demonstrate the great potential of transferring models that are pre-trained with multiple resolutions for dense prediction tasks.

Object detection. We further explore performance of ResFormer on COCO2017 [34] for object detection and instance segmentation, following the designs of ViTDet [30] by appending simple feature pyramids on the feature maps of last-layer outputs and using both non-shifted window attention and global self-attention blocks. In addition, To adapt from ResFormer pre-trained on ImageNet-1K, we also adopt global positional embedding and inject local po-

Table 4. Top-1 Accuracy of TimeSformer on Kinetics-400. MR stands for multi-resolution training.

Model	Testing resolution				
	96	128	160	224	288
TimeSformer [3]	26.28	61.94	70.60	75.54	75.45
ResFormer-B-224	58.61	68.50	73.09	76.32	76.78
ResFormer-B-160	67.28	71.64	74.56	75.98	75.18
ResFormer-B-128	64.66	72.32	74.13	74.19	72.51
ResFormer-B-MR	70.56	74.33	76.38	77.32	77.56

sitional embeddings into all attention blocks. According to results reported in Tab. 3, ResFormer achieves promising results, *e.g.* ResFormer-S-MR outperforms ViTDet-S by 2.0 box AP and 0.6 mask AP and ResFormer-B-MR surpasses ViTDet-B by 1.3 box AP and 0.3 mask AP. We believe that the improved resolution scalability of ResFormer contributes to better performance on object detection.

Video action recognition. We also evaluate ResFormer for video action recognition on Kinetics400. For an easy adaptation from our pre-trained image models to the video domain, we choose the TimeSformer [3] framework with a divided spatial and temporal attention design. In particular, we initialize the backbone with weights of model pre-trained on ImageNet-1K and conduct multi-resolution training on Kinetics400 with clip sizes set to $8 \times 224 \times 224$, $8 \times 160 \times 160$ and $8 \times 128 \times 128$ respectively. As Tab. 4 demonstrates, ResFormer-B fine-tuned with single resolution outweighs vanilla TimeSformer by 0.78% on a testing resolution of 224 and generalizes better to clips of both higher and lower resolutions. On top of that, by implementing multi-resolution training on video samples, ResFormer improves performance on each training resolution by a large margin, *e.g.*, the Top-1 accuracy on testing resolution of 160 grows from 73.09% to 76.08%.

5.2. Discussion

Training resolutions. We experiment with 3 different settings using a small model, *i.e.*, (128, 160, 224), (160, 224, 288) and (128, 224, 384), which we denoted as (a), (b) and (c) respectively. The results are summarized in Fig. 4. We see that ResFormer achieves outstanding performance on a wide range of resolutions. More specifically, compared with (a), (b) adopts higher training resolutions, consequently reflecting on performance rise in high-resolution inputs and incurring a drop on low resolution. Regarding differences between setting of (b) and (c), the spectrum of training resolutions expands in both directions. Despite the wide range between 128 and 384, we witness a all-around improvement of (c) over (b), highlighting that ResFormer is able to deal with significant resolution variations.

Positional embedding. We evaluate the performance of ResFormer with different positional embedding strategies

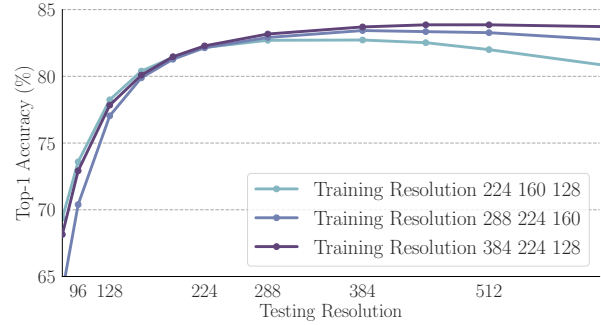


Figure 4. Top-1 Accuracy of ResFormer-S-MR with different training resolutions on ImageNet-1K.

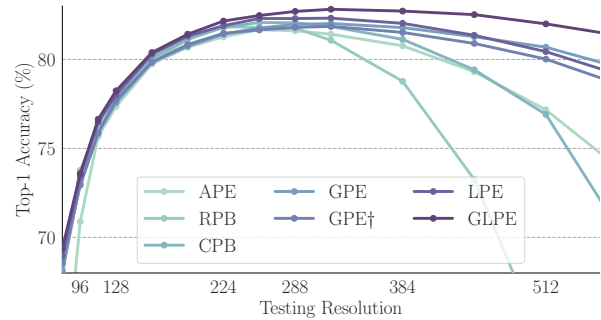


Figure 5. Results of different positional embedding strategies on a broader range of resolutions.

using a small model. In particular, we compare with (1) APE, which stands for vanilla absolute positional embedding in DeiT [50]. In practice, we set the spatial dimension of APE according to the highest training resolution and downsample it for lower resolutions. For inference, the position embedding can be re-scaled to any test resolution with bicubic interpolation; (2) APE* which uses an individual APE for each training resolution; (3) RPB, which is introduced in [36] and we use the RPB of highest resolution for interpolation during inference; (4) CPB, which is a resolution-agnostic strategy [35] and images of arbitrary scales can be input into ViTs with CPB directly; (5) GPE, which is our global positional embedding; (6) GPE†, which represents the plain sine-cosine absolute positional embedding without convolutional enhancement; (7) LPE, which is our local positional embedding; (8) GLPE, which is the combination our GPE and LPE.

Tab. 5 shows the results of different positional embeddings on training resolutions. We see that interpolating APE makes no differences compared with maintaining multiple APES (*i.e.*, APE*), which suggests that ViTs can be trained to deal with different scales of inputs with shared spatial information. Furthermore, all positional embeddings coupled with multi-resolution training demonstrate better results compared to their counterparts trained with single resolutions, *i.e.*, steady gains are made by all positional embeddings when testing on 128, 160 and 224, (gains are shown in

Table 5. Results of ResFormer-S-MR with different positional embedding (PE) strategies on ImageNet-1K. The performance gain compared to single-resolution training is indicated in the bracket.

PE	Testing resolution		
	128	160	224
APE	77.36 ($\uparrow 3.99$)	79.74 ($\uparrow 2.46$)	81.27 ($\uparrow 1.44$)
APE*	77.31 ($\uparrow 3.93$)	79.58 ($\uparrow 2.21$)	81.42 ($\uparrow 1.59$)
RPB	77.90 ($\uparrow 2.74$)	79.92 ($\uparrow 2.04$)	81.84 ($\uparrow 1.27$)
CPB	77.64 ($\uparrow 1.33$)	79.77 ($\uparrow 1.77$)	81.74 ($\uparrow 1.13$)
GPE \dagger	77.73 ($\uparrow 2.73$)	79.83 ($\uparrow 2.27$)	81.47 ($\uparrow 1.29$)
GPE	77.57 ($\uparrow 2.76$)	79.63 ($\uparrow 2.05$)	81.42 ($\uparrow 1.40$)
LPE	78.02 ($\uparrow 2.62$)	80.29 ($\uparrow 2.29$)	81.90 ($\uparrow 1.28$)
GLPE	78.24 ($\uparrow 2.77$)	80.39 ($\uparrow 2.33$)	82.16 ($\uparrow 1.33$)

Table 6. Results of ResFormer-S-MR with different distillation strategies on ImageNet-1K for 100ep. Performance gains over result of training **without distillation** are shown in the bracket.

Distillation		Testing resolution		
Target	Loss	128	160	224
logit	KL	73.50 ($\leftrightarrow 0.0$)	76.45 ($\uparrow 0.07$)	78.82 ($\uparrow 0.26$)
cls	L_2	74.71 ($\uparrow 1.21$)	77.27 ($\uparrow 0.89$)	79.33 ($\uparrow 0.77$)
cls	smooth L_1	74.71 ($\uparrow 1.21$)	77.39 ($\uparrow 1.01$)	79.68 ($\uparrow 1.12$)

the bracket in Tab. 5). Fig. 5 further presents the results of generalizing to more resolutions. Clear performance drops can also be observed in Fig. 5 when APE, RPB and CPB are scaled up to unseen large resolutions, especially RPB. In contrast, LPE and GPE decreases slowly towards extremely large resolutions. GLPE, the combination of LPE and GPE, offers the best results.

Knowledge distillation. To strengthen the interaction between different resolutions, we use a smooth- L_1 loss to distill information from class tokens. We also experiment with a L_2 loss (*i.e.*, Mean squared error). Further, as inputs of different resolutions output features with different scales, we additionally follow the practice in DeiT [50] by distilling logits with a Kullback-Leibler divergence loss. The experiments are conducted on ResFormer-S-MR for 100 epochs for efficiency purposes. Tab. 6 shows the ablation results. We observe the efficacy of distilling with class tokens compared to logits. In addition, the smooth L_1 loss have similar performance with L_2 loss with slightly better results on high resolutions (*i.e.*, 224).

Training strategies. We also explore a widely-used multi-resolution training strategy [22,24] without cross-scale consistency loss, where one iteration consists of randomly sampled images of one certain resolution. In particular, we feed samples of different scales (*i.e.*, 128, 160, 224) iteratively based on two settings: (1) iteration-based, where each mini-batch uses one resolution and resolutions vary for different training iterations; (2) epoch-based, where a fixed resolution is used for each epoch and the change of resolu-

Table 7. Results of ResFormer-S-MR with different training strategies on ImageNet-1K. We append performance gain/drop compared with single-resolution training in the bracket.

Training Strategy	Testing resolution		
	128	160	224
MR (iter)	75.70 ($\uparrow 0.23$)	78.31 ($\uparrow 0.25$)	80.32 ($\downarrow 0.51$)
MR (epoch)	75.18 ($\downarrow 0.29$)	78.05 ($\downarrow 0.01$)	80.26 ($\uparrow 0.16$)
MR w/o KD	77.72 ($\uparrow 2.25$)	79.66 ($\uparrow 1.60$)	81.77 ($\uparrow 0.94$)
MR	78.24 ($\uparrow 2.77$)	80.39 ($\uparrow 2.33$)	82.16 ($\uparrow 1.33$)

tions only occur at the epoch-level. As Tab. 7 shows, both iteration-based and epoch-based multi-resolution training generate worse results compared to single resolution training. In contrast, our strategy demonstrates strong advantages on all training resolutions by clear margins, even without the scale-consistency loss, highlighting the importance of enforcing consistencies of all resolutions in a mini-batch.

Qualitative visualizations. We visualize two positional embeddings on resolutions of 128, 160, 224 and 384, respectively. As shown in Fig. 6, compared with APEs shifted by interpolation, our GPEs that are generated with convolutions demonstrate a smoother variations among input scales. In addition, Fig. 5 suggests that GPE generalizes better to higher resolutions unseen in training.




Figure 6. Heatmaps of different PE averaged on each token. (a): Absolute Positional Embeddings (APE), (b): Global Positional Embeddings (GPE).

6. Conclusion

We introduced ResFormer, a ViT framework to encourage excellent all-round performance on a wide range of resolutions. In particular, ResFormer was motivated by training on sample of different scales and aided by a scale-consistency loss. A global-local positional embedding strategy was also introduced to facilitate better generalization on unseen resolutions. Extensive experiments demonstrated promising scalabilities of ResFormer in a broad range of resolutions. We also observe that ResFormer can be readily adapted to downstream tasks, *e.g.*, semantic segmentation, object detection and video action recognition.

Acknowledgement This project was supported by NSFC under Grant No. 62102092 and No. 62032006.

References

- [1] Alaaeldin Ali, Hugo Touvron, Mathilde Caron, Piotr Bojanowski, Matthijs Douze, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Jakob Verbeek, et al. Xcit: Cross-covariance image transformers. In *NeurIPS*, 2021. 1, 6
- [2] Irwan Bello, William Fedus, Xianzhi Du, Ekin Dogus Cubuk, Aravind Srinivas, Tsung-Yi Lin, Jonathon Shlens, and Barret Zoph. Revisiting resnets: Improved training and scaling strategies. In *NeurIPS*, 2021. 2
- [3] Gedas Bertasius, Heng Wang, and Lorenzo Torresani. Is space-time attention all you need for video understanding? In *ICML*, 2021. 1, 2, 7
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, 2020. 1, 3
- [5] Chun-Fu Chen, Rameswar Panda, and Quanfu Fan. Regionvit: Regional-to-local attention for vision transformers. In *ICLR*, 2022. 2
- [6] Mengzhao Chen, Mingbao Lin, Ke Li, Yunhang Shen, Yongjian Wu, Fei Chao, and Rongrong Ji. Cf-vit: A general coarse-to-fine method for vision transformer. *arXiv preprint arXiv:2203.03821*, 2022. 1
- [7] Zhe Chen, Yuchen Duan, Wenhai Wang, Junjun He, Tong Lu, Jifeng Dai, and Yu Qiao. Vision transformer adapter for dense predictions. In *ICLR*, 2023. 6
- [8] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. In *NeurIPS*, 2021. 2
- [9] Xiangxiang Chu, Zhi Tian, Bo Zhang, Xinlong Wang, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Conditional positional encodings for vision transformers. *arXiv preprint arXiv:2102.10882*, 2021. 2
- [10] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation policies from data. In *CVPR*, 2019. 4
- [11] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *CVPR Workshops*, 2020. 4
- [12] Zihang Dai, Zhilin Yang, Yiming Yang, Jaime G Carbonell, Quoc Viet Le, and Ruslan Salakhutdinov. Transformer-xl: Attentive language models beyond a fixed-length context. In *ACL*, 2019. 2
- [13] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2
- [14] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *CVPR*, 2022. 5
- [15] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 1, 2, 4, 5
- [16] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. In *ICML*, 2021. 2
- [17] Alaaeldin El-Nouby, Gautier Izacard, Hugo Touvron, Ivan Laptev, Hervé Jegou, and Edouard Grave. Are large-scale datasets necessary for self-supervised pre-training? *arXiv preprint arXiv:2112.10740*, 2021. 2
- [18] Haoqi Fan, Bo Xiong, Kartikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale vision transformers. In *ICCV*, 2021. 3
- [19] Priya Goyal, Mathilde Caron, Benjamin Lefaudeux, Min Xu, Pengchao Wang, Vivek Pai, Mannat Singh, Vitaliy Liptchinsky, Ishan Misra, Armand Joulin, et al. Self-supervised pretraining of visual features in the wild. *arXiv preprint arXiv:2103.01988*, 2021. 2
- [20] Jiaqi Gu, Hyoukjun Kwon, Dilin Wang, Wei Ye, Meng Li, Yu-Hsin Chen, Liangzhen Lai, Vikas Chandra, and David Z Pan. Multi-scale high-resolution vision transformer for semantic segmentation. In *CVPR*, 2022. 3
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, 2022. 5, 6
- [22] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017. 2, 6, 8
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Spatial pyramid pooling in deep convolutional networks for visual recognition. *TPAMI*, 2015. 2
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 8
- [25] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *ICML*, 2021. 1
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 5
- [27] Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In *ECCV*, 2020. 2
- [28] Hei Law and Jia Deng. Cornernet: Detecting objects as paired keypoints. In *ECCV*, 2018. 3
- [29] Kunchang Li, Yali Wang, Peng Gao, Guanglu Song, Yu Liu, Hongsheng Li, and Yu Qiao. Uniformer: Unified transformer for efficient spatiotemporal representation learning. In *ICLR*, 2022. 1, 5
- [30] Yanghao Li, Hanzi Mao, Ross Girshick, and Kaiming He. Exploring plain vision transformer backbones for object detection. *arXiv preprint arXiv:2203.16527*, 2022. 1, 6
- [31] Yanghao Li, Saining Xie, Xinlei Chen, Piotr Dollár, Kaiming He, and Ross Girshick. Benchmarking detection transfer learning with vision transformers. *arXiv preprint arXiv:2111.11429*, 2021. 6

- [32] Mingbao Lin, Mengzhao Chen, Yuxin Zhang, Ke Li, Yunhang Shen, Chunhua Shen, and Rongrong Ji. Super vision transformer. *arXiv preprint arXiv:2205.11397*, 2022. [1](#)
- [33] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *CVPR*, 2017. [2](#)
- [34] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014. [2](#), [6](#)
- [35] Ze Liu, Han Hu, Yutong Lin, Zhuliang Yao, Zhenda Xie, Yixuan Wei, Jia Ning, Yue Cao, Zheng Zhang, Li Dong, et al. Swin transformer v2: Scaling up capacity and resolution. In *CVPR*, 2022. [1](#), [2](#), [5](#), [7](#)
- [36] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. [1](#), [2](#), [3](#), [5](#), [7](#)
- [37] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. In *CVPR*, 2022. [1](#)
- [38] Chuofan Ma, Qiushan Guo, Yi Jiang, Zehuan Yuan, Ping Luo, and Xiaojuan Qi. Rethinking resolution in the context of efficient video recognition. In *NeurIPS*, 2022. [3](#)
- [39] Lingchen Meng, Hengduo Li, Bor-Chun Chen, Shiyi Lan, Zuxuan Wu, Yu-Gang Jiang, and Ser-Nam Lim. Advait: Adaptive vision transformers for efficient image recognition. In *CVPR*, 2022. [1](#)
- [40] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021. [1](#)
- [41] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, 2020. [2](#)
- [42] Kanchana Ranasinghe, Muzammal Naseer, Salman Khan, Fahad Shahbaz Khan, and Michael S Ryoo. Self-supervised video transformer. In *CVPR*, 2022. [3](#)
- [43] Yongming Rao, Wenliang Zhao, Benlin Liu, Jiwen Lu, Jie Zhou, and Cho-Jui Hsieh. Dynamicvit: Efficient vision transformers with dynamic token sparsification. In *NeurIPS*, 2021. [1](#)
- [44] Carlos Riquelme, Joan Puigcerver, Basil Mustafa, Maxim Neumann, Rodolphe Jenatton, André Susano Pinto, Daniel Keysers, and Neil Houlsby. Scaling vision with sparse mixture of experts. In *NeurIPS*, 2021. [2](#)
- [45] Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. Self-attention with relative position representations. In *NAACL*, 2018. [2](#)
- [46] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In *ICLR*, 2015. [3](#)
- [47] Bharat Singh and Larry S Davis. An analysis of scale invariance in object detection snip. In *CVPR*, 2018. [3](#)
- [48] Andreas Peter Steiner, Alexander Kolesnikov, Xiaohua Zhai, Ross Wightman, Jakob Uszkoreit, and Lucas Beyer. How to train your vit? data, augmentation, and regularization in vision transformers. *TMLR*, 2022. [1](#)
- [49] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *ICML*, 2019. [2](#)
- [50] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, 2021. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [51] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022. [1](#)
- [52] Hugo Touvron, Andrea Vedaldi, Matthijs Douze, and Hervé Jégou. Fixing the train-test resolution discrepancy. In *NeurIPS*, 2019. [3](#)
- [53] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *NeurIPS*, 2017. [1](#), [2](#), [5](#)
- [54] Junke Wang, Dongdong Chen, Zuxuan Wu, Chong Luo, Lu-wei Zhou, Yucheng Zhao, Yujia Xie, Ce Liu, Yu-Gang Jiang, and Lu Yuan. Omnivl: One foundation model for image-language and video-language tasks. In *NeurIPS*, 2022. [1](#)
- [55] Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In *ICCV*, 2021. [1](#), [5](#), [6](#)
- [56] Yulin Wang, Rui Huang, Shiji Song, Zeyi Huang, and Gao Huang. Not all images are worth 16x16 words: Dynamic transformers for efficient image recognition. In *NeurIPS*, 2021. [1](#)
- [57] Yixuan Wei, Han Hu, Zhenda Xie, Zheng Zhang, Yue Cao, Jianmin Bao, Dong Chen, and Baining Guo. Contrastive learning rivals masked image modeling in fine-tuning via feature distillation. *arXiv preprint arXiv:2205.14141*, 2022. [4](#)
- [58] Chao-Yuan Wu, Ross Girshick, Kaiming He, Christoph Feichtenhofer, and Philipp Krahenbuhl. A multigrid method for efficiently training video models. In *NeurIPS*, 2020. [2](#), [3](#)
- [59] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *ICCV*, 2021. [1](#)
- [60] Kan Wu, Houwen Peng, Minghao Chen, Jianlong Fu, and Hongyang Chao. Rethinking and improving relative position encoding for vision transformer. In *ICCV*, 2021. [2](#)
- [61] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified perceptual parsing for scene understanding. In *ECCV*, 2018. [6](#)
- [62] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Jianmin Bao, Zhuliang Yao, Qi Dai, and Han Hu. Simmim: A simple framework for masked image modeling. In *CVPR*, 2022. [1](#)
- [63] Zhenda Xie, Zheng Zhang, Yue Cao, Yutong Lin, Yixuan Wei, Qi Dai, and Han Hu. On data scaling in masked image modeling. In *CVPR*, 2023. [2](#)
- [64] Zhen Xing, Qi Dai, Han Hu, Jingjing Chen, Zuxuan Wu, and Yu-Gang Jiang. Svformer: Semi-supervised video transformer for action recognition. In *CVPR*, 2023. [1](#)

- [65] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis EH Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-token vit: Training vision transformers from scratch on imagenet. In *ICCV*, 2021. 1
- [66] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, 2019. 4
- [67] Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transformers. In *CVPR*, 2022. 2
- [68] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 4
- [69] Linfeng Zhang, Chenglong Bao, and Kaisheng Ma. Self-distillation: Towards efficient and compact neural networks. In *NeurIPS*, 2021. 5
- [70] Xiaosong Zhang, Yunjie Tian, Lingxi Xie, Wei Huang, Qi Dai, Qixiang Ye, and Qi Tian. Hivit: A simpler and more efficient design of hierarchical vision transformer. In *ICLR*, 2023. 1
- [71] Zhun Zhong, Liang Zheng, Guoliang Kang, Shaozi Li, and Yi Yang. Random erasing data augmentation. In *AAAI*, 2020. 4
- [72] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 2019. 2, 6