

# Rethinking Spatial Dimensions of Vision Transformers

Byeongho Heo Sangdoon Yun Dongyoon Han Sanghyuk Chun Junsuk Choe Seong Joon Oh

NAVER AI Lab

## Abstract

*Vision Transformer (ViT) extends the application range of transformers from language processing to computer vision tasks as being an alternative architecture against the existing convolutional neural networks (CNN). Since the transformer-based architecture has been innovative for computer vision modeling, the design convention towards an effective architecture has been less studied yet. From the successful design principles of CNN, we investigate the role of the spatial dimension conversion and its effectiveness on the transformer-based architecture. We particularly attend the dimension reduction principle of CNNs; as the depth increases, a conventional CNN increases channel dimension and decreases spatial dimensions. We empirically show that such a spatial dimension reduction is beneficial to a transformer architecture as well, and propose a novel Pooling-based Vision Transformer (PiT) upon the original ViT model. We show that PiT achieves the improved model capability and generalization performance against ViT. Throughout the extensive experiments, we further show PiT outperforms the baseline on several tasks such as image classification, object detection and robustness evaluation. Source codes and ImageNet models are available at <https://github.com/naver-ai/pit>.*

## 1. Introduction

The architectures based on the self-attention mechanism have achieved great success in the field of Natural Language Processing (NLP) [35]. There have been attempts to utilize the self-attention mechanism in computer vision. Non-local networks [37] and DETR [4] are representative works, showing that the self-attention mechanism is also effective in video classification and object detection tasks, respectively. Recently, Vision Transformer (ViT) [9], a transformer architecture consisting of self-attention layers, has been proposed to compete with ResNet [13], and shows that it can achieve the best performance without convolution operation on ImageNet [8]. As a result, a new direction of network architectures based on self-attention mechanism, not

convolution operation, has emerged in computer vision.

ViT is quite different from convolutional neural networks (CNN). Input images are divided into  $16 \times 16$  patches and fed to the transformer network; except for the first embedding layer, there is no convolution operation in ViT, and the position interactions occur only through the self-attention layers. While CNNs have restricted spatial interactions, ViT allows all the positions in an image to interact through transformer layers. Although ViT is an innovative architecture and has proven its powerful image recognition ability, it follows the transformer architecture in NLP [35] without any changes. Some essential design principles of CNNs, which have proved to be effective in the computer vision domain over the past decade, are not sufficiently reflected. We thus revisit the design principles of CNN architectures and investigate their efficacy when applied to ViT architectures.

CNNs start with a feature of large spatial sizes and a small channel size and gradually increase the channel size while decreasing the spatial size. This dimension conversion is indispensable due to the layer called spatial pooling. Modern CNN architectures, including AlexNet [21], ResNet [13], and EfficientNet [33], follow this design principle. The pooling layer is deeply related to the receptive field size of each layer. Some studies [6, 27, 5] show that the pooling layer contributes to the expressiveness and generalization performance of the network. However, unlike the CNNs, ViT does not use a pooling layer and uses spatial tokens of the same size in all layers.

First, we verify the advantages of pooling layers on CNNs. Our experiments show that the pooling layer improves the model capability and generalization performance of ResNet. To extend the advantages of the pooling layer to ViT, we propose a Pooling-based Vision Transformer (PiT). PiT is a transformer architecture combined with a pooling layer. It enables the spatial size reduction in the ViT structure as in ResNet. We also investigate the benefits of PiT compared to ViT and confirm that the pooling layers also improve the performance of ViT. Finally, to analyze the effect of the pooling layers in ViT, we measure the spatial interaction ratio of ViT similar to the receptive field size of a convolutional architecture. We show that the pooling layer

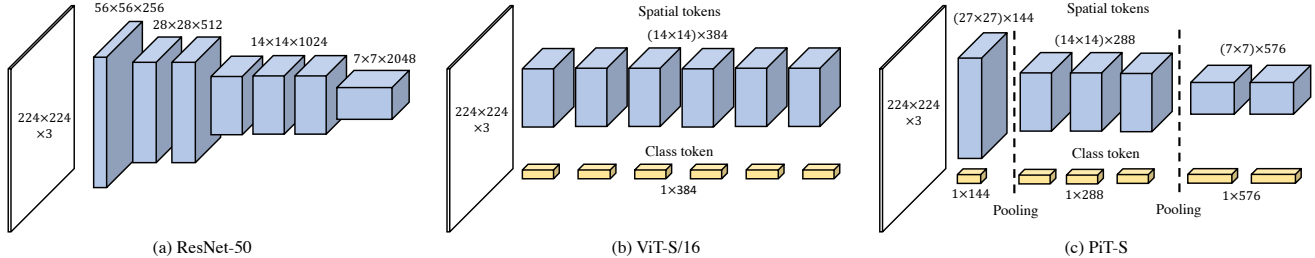


Figure 1. **Schematic illustration of dimension configurations of network architectures.** We visualize ResNet50 [13], Vision Transformer (ViT) [9], and our Pooling-based Vision Transformer (PiT); (a) ResNet50 gradually downsamples the features from the input to the output; (b) ViT does not use pooling layers, so the feature dimension is maintained for all the layers; (c) PiT involves pooling layers into ViT.

has the effect of controlling the size of spatial interaction occurring in the self-attention layer, which is similar to the receptive field control of the convolutional architecture.

We verify that PiT improves performances over ViT on various tasks. On ImageNet classification, PiT and outperforms ViT at various scales and training environments. Additionally, we have compared the performance of PiT with various convolutional architectures and have specified the scale at which the transformer architecture outperforms the CNN. We further measure the performance of PiT with a detection head on object detection. ViT- and PiT-based DETR [4] are trained on the COCO 2017 dataset [25] and the result shows that PiT is even better than ViT as a backbone architecture for a task other than image classification. Finally, we verify the performance of PiT in various environments through the robustness benchmark.

## 2. Related works

### 2.1. Dimension configuration of CNN

Convolutional neural networks use various types of pooling layers. We cover the variants in this section. In general, the term pooling is used for the max-pooling and average-pooling layers, but in this paper, the layer that increases the channel while reducing the spatial size is referred to as pooling. Pooling layers are found in AlexNet [21], which is early architecture of a convolutional neural network in computer vision. AlexNet uses three max-pooling layers. In max pooling layer, spatial size of the feature is reduced by half, and the channel size is increased by the convolution after the max-pooling. VGGnet [31] uses 5 spatial resolutions using 5 max-pooling. In the pooling layer, the spatial size is reduced by half and the channel size is doubled. So, even when the spatial resolution was changed, the floating point operations (FLOPs) of convolution layers are maintained. GoogLeNet [32] also used the pooling layer. In the paper [32], authors were concerned about the loss of spatial information due to the pooling layer, but argued that nevertheless the pooling layer is an essential part of a convolutional neural network. ResNet [13] performed spatial size

reduction using the convolution layer of stride 2 instead of max pooling. This is an improved pooling method and channel increase and spatial reduction are performed in a single layer. The convolution layer of stride 2 is also used as a pooling method in recent architectures (EfficientNet [33], MobileNet [30, 19]). The pooling layer used in this paper is also derived from the convolution layer of stride 2. Some papers focused on the channel increase due to the pooling layer. PyramidNet [11] pointed out that the channel increase occurs only in the pooling layer and proposed a method to gradually increase the channel size in layers other than the pooling layer. ReXNet [12] reported that the channel configuration of the network has a significant influence on the network performance. In summary, most of convolution networks use a dimension configuration based on a pooling layer. In addition, it is reported that the channel structure is important for performance of network. Our goal is to extend the effective channel configuration of CNN to ViT.

### 2.2. Self-attention mechanism

Transformer architecture [35] significantly increased the performance of the NLP task with self-attention mechanism. Funnel Transformer [7] improves the transformer architecture with reducing tokens by a pooling layer and skip-connection. However, because of the basic difference between the architecture of NLP and computer vision, the method of applying pooling is different from our method. Some studies are conducted to utilize the transformer architecture to the backbone network for computer vision task. Non-local network [37] adds a few self-attention layers to CNN backbone, and it shows that self-attention mechanism can be used in CNN. [29] replaced  $3 \times 3$  convolution of ResNet to local self-attention layer. [36] used an attention layer for each spatial axis. [2] enables self-attention of the entire spatial map by reducing the computation of the attention mechanism. Most of these methods replace  $3 \times 3$  convolution with self-attention or adds a few self-attention layers. Therefore, the basic structure of ResNet is inherited, that is, it has the convolution of stride 2 as ResNet, resulting in a network having a dimension configuration of ResNet.

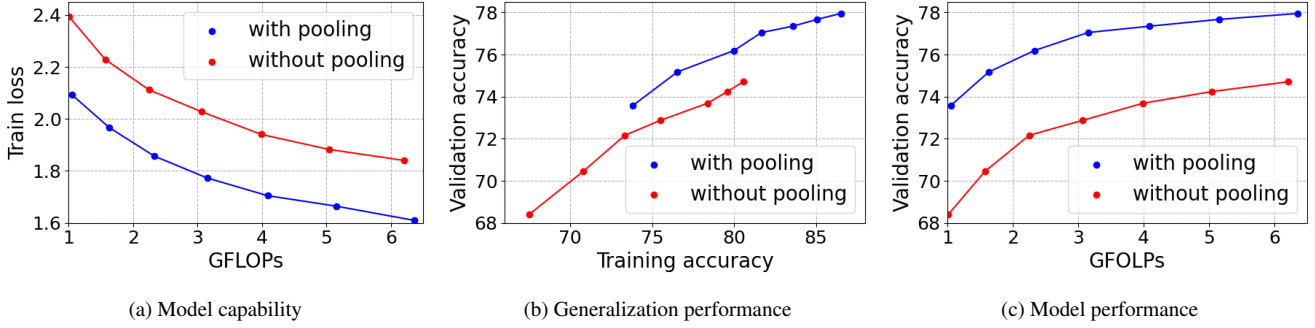


Figure 2. **Effects of the pooling layer in ResNet50 [13].** We perform an experiment to verify the effect of the pooling layer used in ResNet50. As shown in the figures, the pooling layer improves the model capability, generalization performance, and model performance of ResNet50.

Only the vision transformer uses a structure that uses the same spatial size in all layers. Although ViT did not follow the conventions of ResNet, it contains many valuable new components in the network architecture. In ViT, layer normalization is applied for each spatial token. Therefore, layer normalization of ViT is closer to positional normalization [22] than a layer norm of convolutional neural network [1, 39]. Although it overlaps with the lambda network [2], it is not common to use global attention through all blocks of the network. The use of class tokens instead of global average pooling is also new, and it has been reported that separating token increases the efficiency of distillation [34]. In addition, the layer configuration, the skip-connection position, and the normalization position of the Transformer are also different from ResNet. Therefore, our research can be seen as giving direction to a newly started architecture.

### 3. Revisiting spatial pooling

In order to introduce spatial pooling to ViT, we investigate spatial pooling from various perspectives. First, we verify the benefits of pooling-based dimension configuration in ResNet architecture. Although the pooling layer has been widely used for most convolutional architectures, its effectiveness is rarely verified. Before utilizing the pooling layer to ViT, we measure the benefits of the pooling layer on ResNet. Based on the findings, we propose a Pooling-based Vision Transformer (PiT) that applies the pooling layers to ViT. We propose a pooling layer for transformer architecture and design ViT with pooling (PiT). With PiT models, we verify whether the pooling layer brings advantages to ViT as in ResNet. In addition, we analyze the spatial location used in the attention layer of ViT to investigate the effect of spatial pooling in ViT. Finally, we introduce PiT architectures corresponding to various scales of ViT.

### 3.1. Effect of pooling on CNN

As shown in Figure 1 (a), most convolutional neural networks have pooling layers which reduces the spatial dimension while increases the channel dimension. In ResNet50, a stem layer reduces the spatial size of an image to  $56 \times 56$ . Convolution layers with stride 2 reduce the spatial dimension by half and double the channel dimension. The pooling operation using a convolution layer with stride 2 is a frequently used method in recent architectures [33, 30, 19, 12]. This is referred to as stride convolution or learnable pooling, but in this paper, we refer to it as a pooling layer for simplicity. There are various ways to perform pooling, but it is common for a convolutional neural network to include a pooling layer, which makes a design principle of the convolutional architecture.

We conduct an experiment to analyze the performance difference according to the presence or absence of a pooling layer in a convolutional architecture. ResNet50, one of the most widely used networks in ImageNet, is used for the architecture and is trained over 100 epochs without complex training techniques. In the case of ResNet without pooling, we change the stem layer to reduce the feature to  $14 \times 14$  and remove pooling layers to maintain feature dimensions like ViT. We measured the performance for several network sizes by changing the base channel size of ResNet.

First, we measured the relation between FLOPs and training loss with and without the pooling layers. As shown in Figure 2 (a), ResNet with pooling layers shows lower training loss over the same computation costs (FLOPs). It implies that the pooling layers increase the capability of architecture. Next, we analyzed the relation between training and validation accuracy, which represents the generalization performance of architecture. As shown in Figure 2 (b), ResNet with pooling layer achieves higher validation accuracy than Resnet without pooling layer. Therefore, the pooling layer is also helpful for the generalization performance of ResNet50. In summary, the pooling layers im-

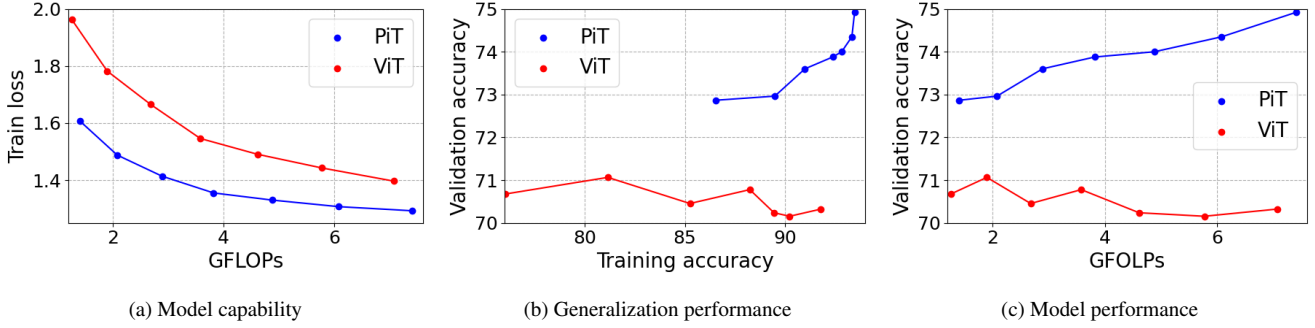


Figure 3. **Effects of the pooling layer in vision transformer (ViT) [9].** We compare our Pooling-based Vision Transformer (PiT) with original ViT at various aspects of network architecture. PiT outperforms ViT in capability, generalization performance, and model performance.

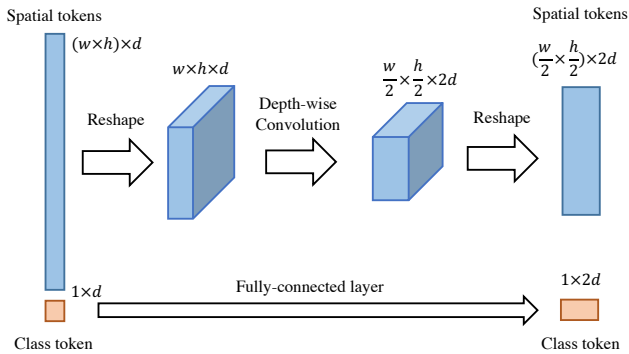


Figure 4. **Pooling layer of PiT architecture.** PiT uses the pooling layer based on depth-wise convolution to achieve channel multiplication and spatial reduction with small parameters.

prove the model capability and generalization performance of the architecture and consequently bring a significant improvement in validation accuracy as shown in Figure 2 (c).

### 3.2. Pooling-based Vision Transformer (PiT)

Vision Transformer (ViT) performs network operations based on self-attention, not convolution operations. In the self-attention mechanism, the similarity between all locations is used for spatial interaction. Figure 1 (b) shows the dimension structure of this ViT. Similar to the stem layer of CNN, ViT divides the image by patch at the first embedding layer and embedding it to tokens. Basically, the structure does not include pooling and keeps the same number of spatial tokens overall layer of the network. Although the self-attention operation is not limited by spatial distance, the size of the spatial area participating in attention is affected by the spatial size of feature. Therefore, in order to adjust the size of the area covered by the layer like ResNet, a pooling layer is also required in ViT.

To utilize the advantages of the pooling layer to ViT, we propose a new architecture called Pooling-based Vision

Transformer (PiT). First, we designed the pooling layer for ViT. Our pooling layer is shown in Figure 4. Since ViT handles neuron responses in the form of 2D-matrix rather than 3D-tensor, the pooling layer should separate spatial tokens and reshape them into 3D-tensor with spatial structure. After reshaping, spatial size reduction and channel increase are performed by depth-wise convolution. And, the responses are reshaped into 2D-matrix for computation of transformer blocks. In ViT, there are parts that do not correspond to the spatial structure, such as a class token or distillation token [34]. For these parts, the pooling layer uses an additional fully-connected layer to adjust the channel size to match the spatial tokens. Our design aims to perform the pooling role with minimal operation. As a result, a depth-wise convolution layer conducts decreasing the spatial size and increasing the channel size with a small number of parameters. Our pooling layer enables pooling operation on ViT and is used for our PiT architecture as shown in Figure 1 (c). PiT includes two pooling layers and uses three scales of spatial tokens similar to ResNet.

Using PiT architecture, we performed an experiment to verify the effect of the pooling layer in ViT. The experiment setting is the same as the ResNet experiment. Figure 3 (a) represents the model capability of ViT with and without the pooling layer. At the same computation cost, ViT with pooling has a lower train loss than ViT without pooling. This result is similar to ResNet case. Using the pooling layers in ViT also improves the capability of architecture. The comparison between training accuracy and validation accuracy shows a significant difference. As shown in Figure 3 (b), ViT without pooling does not improve validation accuracy even if training accuracy increases. On the other hand, in ViT with pooling, validation accuracy increases as training accuracy increases. The big difference in generalization performance causes the performance difference between ViT with pooling and ViT without pooling as shown in Figure 3 (c). The phenomenon that ViT does not increase

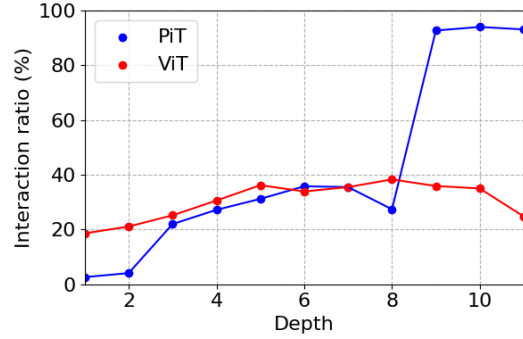
performance even when FLOPs increase in ImageNet is reported in ViT paper [9]. In the training data of ImageNet scale, ViT shows poor generalization performance, and our proposed pooling layer alleviates this. So, we can say that the pooling layer is necessary for the generalization of ViT. Using the training trick is also a way to improve the generalization performance of ViT in ImageNet. The combination of training trick and pooling layer is covered in the experiment section.

### 3.3. Spatial interaction

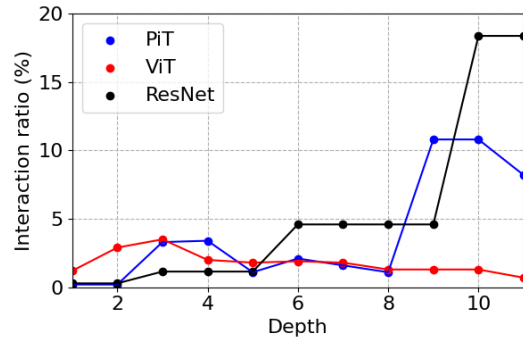
We investigate the effect of the pooling in ViT by analyzing the self-attention of the vision transformer. The role of pooling in a convolutional architecture is to adjust the receptive field. When the spatial size is large, a portion of a convolutional kernel is relatively small. On the other hand, as the spatial size decreases, the portion of the convolution kernel in the entire image increases. Therefore, an area of the spatial interaction in the layer can be adjusted through pooling layers. In the case of ViT, spatial interaction is performed regardless of spatial distance by the self-attention layer.

However, the self-attention layer also has limitations in a number of interacting tokens, and therefore, the interaction area is determined according to the spatial size. We measured the spatial interaction area of ViT and PiT with pre-trained models on ImageNet. The criterion for spatial interaction is based on the score after the soft-max of the attention matrix. We used 1% and 10% as thresholds, counted the number of spatial locations where interactions above the threshold occurred, and calculated a spatial interaction ratio by dividing the number of interaction locations by the total size of spatial tokens. Transformer uses multi-head attention, so there are as many attention maps as the number of heads. We measured the ratio of location that was used more than once among several heads. Figure 5 (a) shows the average ratio of locations that involved attention more than 1%.

In the case of the ViT, the interactions are between 20%-40% on average, and since there is no pooling layer, the numerical value does not change significantly depending on the layer. PiT reduces the number of tokens while increasing the head through pooling. So, as shown in Figure 5 (a), the early layers have a small interaction ratio, but the latter layer shows a nearly 100% interaction ratio. For comparison with ResNet, we changed the threshold to 10%, and the result is as shown in Figure 5 (b). In the case of ResNet, 3x3 convolution means 3x3 spatial interactions. Therefore, we divide 3x3 by the spatial size and use it as an approximation to compare with the interaction ratio of attention. While the interaction ratio of ViT is similar across the layers, the interaction ratio of ResNet and PiT increases as it passes through the pooling layer. In summary, as the receptive field of ResNet increases through the pooling layer, the



(a) Spatial interaction (over 1%)



(b) Spatial interaction (over 10%)

Figure 5. **Spatial interaction ratio.** We investigate the attention matrix of the self-attention layer and record the ratio of spatial locations that influenced more than 1% or 10%, which is called the interaction ratio. The figures show the average spatial interaction ratio of self-attention or convolution layer. PiT makes a transformer have interaction ratios similar to that of ResNet.

spatial interaction area of the transformer becomes wider through the pooling layer. We believe that the control of the interaction area is closely related to the performance of the architecture.

### 3.4. Architecture design

The architectures proposed in ViT paper [9] aimed at datasets larger than ImageNet. These architectures (ViT-Large, ViT-Huge) have an extremely large scale than general ImageNet networks, so it is not easy to compare them with other networks. So, following the previous study [34] of Vision Transformer on ImageNet, we design the PiT at a scale similar to the small-scale ViT architectures (ViT-Base, ViT-Small, ViT-Tiny). In the DeiT paper [34], ViT-Small and ViT-Tiny are named as DeiT-S and DeiT-Ti, but in order to avoid confusion due to the model name change, we use ViT for all models. Corresponding to the three scales of ViT (tiny, small, and base), we design four scales of PiT (tiny, extra small, small, and base). Detail architectures are described in Table 1. For convenience, we abbreviate



Network	Spatial size	# of blocks	# of heads	Channel size	FLOPs
ViT-Ti [34]	14 x 14	12	3	192	1.3B
PiT-Ti	27 x 27	2	2	64	0.7B
	14 x 14	6	4	128	
	7 x 7	4	8	256	
PiT-XS	27 x 27	2	2	96	1.4B
	14 x 14	6	4	192	
	7 x 7	4	8	384	
ViT-S [34]	14 x 14	12	6	384	4.6B
PiT-S	27 x 27	2	3	144	2.9B
	14 x 14	6	6	288	
	7 x 7	4	12	576	
ViT-B [9]	14 x 14	12	12	768	17.6B
PiT-B	31 x 31	3	4	256	12.5B
	16 x 16	6	8	512	
	8 x 8	4	16	1024	

Table 1. **Architecture configuration.** The table shows spatial sizes, number of blocks, number of heads, channel size, and FLOPs of ViT and PiT. The structure of PiT is designed to be as similar as possible to ViT and to have less GPU latency.

the model names: Tiny - Ti, eXtra Small - XS, Small - S, Base - B FLOPs and spatial size were measured based on  $224 \times 224$  image. Since PiT uses a larger spatial size than ViT, we reduce the stride size of the embedding layer to 8, while patch-size is 16 as ViT. A total of two pooling layers are used for PiT, and the channel increase of the pooling layer was implemented by increasing the number of heads of multi-head attention. We design PiT to have the similar depth to ViT, and adjust the channels and the heads to have smaller FLOPs, parameter size, and GPU latency than those of ViT. We want to clarify that PiT is not designed by large-scale parameter search such as NAS [26, 3], so PiT can be further improved through a network architecture search algorithm.

## 4. Experiments

We verified the performance of PiT through various experiments. First, we compared PiT at various scales with ViT in various training environments of ImageNet training. And, we extended the ImageNet comparison to architectures other than Transformer. In particular, we focus on the comparison of the performance of ResNet and PiT, and investigate whether PiT can beat ResNet. We also applied PiT to an object detector based on DETR [4], and compared the performance as a backbone architecture for object detection. To analyze PiT in various views, we evaluated the performance of PiT on robustness benchmarks.

### 4.1. ImageNet classification

We compared the performance of PiT models of Table 1 with corresponding ViT models. To clarify the computation time and size of the network, we measured FLOPs, the number of parameters, and GPU throughput (images / sec) of each network. The GPU throughput was measured on NVIDIA V100 single GPU with 128 batch-size. We trained the network using four representative training environments. The first is a vanilla setting that trains the network without complicated training techniques. The vanilla setting is the same setting as the experiments in Figure 2 and Figure 3, and it is the setting with the lowest performance due to the lack of techniques to help generalization performance. The second is a training setting using CutMix [41] data augmentation. CutMix can be used without any other technique, and it significantly improves the performance of the network. The third is the DeiT [34] setting, which is a compilation of training techniques to train ViT on ImageNet. DeiT setting includes various training techniques and parameter tuning, and we used the same training setting through the open-source code. However, in the case of Repeated Augment [18], we confirmed that it had a negative effect in a small model, and it was used only for base models. The last is a DeiT setting with knowledge distillation. The distillation setting is reported as the best performance setting in DeiT [34] paper. The network uses an additional distillation token and trained with distillation loss [17] using RegNetY-16GF [28] as a teacher network. We used AdamP [16] optimizer for all settings, and the learning rate, weight decay, and warmup were set equal to DeiT [34] setting. The cosine learning rate decay was used as the learning rate schedule. We train models over 100 epochs for Vanilla and CutMix settings, and 300 epochs for DeiT settings.

The results are shown in Table 2. Comparing the PiT and ViT of the same name, the PiT has fewer FLOPs and faster speed than ViT. Nevertheless, PiT shows higher performance than ViT. In the case of vanilla and CutMix settings, where a few training techniques are applied, the performance of PiT is superior to the performance of ViT. Even in the case of a DeiT and distill settings, PiT shows comparable or better performance to ViT. Therefore, PiT can be seen as a better architecture than ViT in terms of performance and computation. The generalization performance issue of ViT in Figure 3 can also be observed in this experiment. Like ViT-S in the Vanilla setting and ViT-B in the CutMix setting, ViT often shows no increase in performance even when the model size increases. On the other hand, the performance of PiT increases according to the model size in all training settings. It seems that the generalization performance problem of ViT is alleviated by the pooling layers.

We also compared the performance of PiT with the convolutional architectures. In the previous experiment, we per-

Architecture	FLOPs	# of params	Throughput (imgs/sec)	Vanilla	+CutMix [41]	+DeiT [34]	+Distill [34]
ViT-Ti [34]	1.3 B	5.7 M	2564	68.7%	68.5%	72.2%	74.5%
PiT-Ti	0.7 B	4.9 M	3030	71.3%	72.6%	73.0%	74.6%
PiT-XS	1.4 B	10.6 M	2128	72.4%	76.8%	78.1%	79.1%
ViT-S [34]	4.6 B	22.1 M	980	68.7%	76.5%	79.8%	81.2%
PiT-S	2.9 B	23.5 M	1266	73.3%	79.0%	80.9%	81.9%
ViT-B [9]	17.6 B	86.6 M	303	69.3%	75.3%	81.8%	83.4%
PiT-B	12.5 B	73.8 M	348	76.1%	79.9%	82.0%	84.0%

Table 2. **ImageNet performance comparison with ViT.** We compare the performances of ViT and PiT with some training techniques on ImageNet dataset. PiT shows better performance with low computation compared to ViT.

Network	# of params	Throughput (imgs/sec)	Accuracy
ResNet18 [13, 42]	11.7M	4545	72.5%
MobileNetV2 [30]	3.5M	3846	72.0%
MobileNetV3 [19]	5.5M	3846	75.2%
EfficientNet-B0 [33]	5.3M	2857	77.1%
ViT-Ti [34]	5.7M	2564	74.5%
<b>PiT-Ti</b>	4.9M	3030	74.6%
ResNet34 [13, 38]	21.8M	2631	75.1%
ResNet34D [14, 38]	21.8M	2325	77.1%
EfficientNet-B1 [33]	7.8M	1754	79.1%
<b>PiT-XS</b>	10.6M	2128	79.1%
ResNet50 [13, 42]	25.6M	1266	80.2%
ResNet101 [13, 42]	44.6M	757	81.6%
ResNet50D [14, 38]	25.6M	1176	80.5%
EfficientNet-B2 [33]	9.2M	1333	80.1%
EfficientNet-B3 [33]	12.2M	806	81.6%
RegNetY-4GF [28]	20.6M	1136	79.4%
ResNeSt50 [43]	27.5M	877	81.1%
ViT-S [34]	22.1M	980	81.2%
<b>PiT-S</b>	23.5M	1266	81.9%
ResNet152 [13, 42]	60.2M	420	81.9%
ResNet101D [14, 38]	44.6M	354	83.0%
ResNet152D [14, 38]	60.2M	251	83.7%
EfficientNet-B4 [33]	19.3M	368	82.9%
EfficientNet-B5 [33]	30.4M	179	83.6%
RegNetY-16GF [28]	83.6M	352	80.4%
ResNeSt101 [43]	48.3M	398	83.0%
ResNeSt200 [43]	70.2M	144	83.9%
ViT-B [9, 34]	86.6M	303	83.4%
<b>PiT-B</b>	73.8M	348	84.0%

Table 3. **ImageNet performance of various models.** We compare our PiT-(Ti, XS, S, and B) models with the counterparts which have similar number of parameters.

formed performance comparison in the same training setting using the similarity of architecture. However, when comparing various architectures, it is infeasible to unify with a setting that works well for all architectures. There-

Backbone	Avg. Precision at IOU			Params.	Throughput (imgs/sec)
	AP	AP <sub>50</sub>	AP <sub>75</sub>		
ResNet50	36.7	56.8	38.1	41.3 M	22.8
ViT-S	32.1	52.4	32.3	39.3 M	26.1
PiT-S	34.4	54.7	35.3	40.6 M	26.3

Table 4. **COCO detection performance based on DETR [4].** We evaluate the performance of PiT as a pretrained backbone for object detection.

fore, we performed the comparison based on the best performance reported for each architecture. But, it was limited to the model trained using only ImageNet images. When the paper that proposed the architecture and the paper that reported the best performance was different, we cite both papers. When the architecture is different, the comparison of FLOPs often fails to reflect the actual throughput. Therefore, we re-measured the GPU throughput and number of params on a single V100 GPU, and compared the top-1 accuracy for the performance index. Table 3 shows the comparison result. In the case of the PiT-B scale, the transformer-based architecture (ViT-B, PiT-B) outperforms the convolutional architecture. Even in the PiT-S scale, PiT-S shows superior performance than convolutional architecture (ResNet50) or outperforms in throughput (EfficientNet-b3). However, in the case of PiT-Ti, the performance of convolutional architectures such as ResNet34 [13], MobileNetV3 [19], and EfficientNet-b0 [33] outperforms ViT-Ti and PiT-Ti. Overall, the transformer architecture shows better performance than the convolutional architecture at the scale of ResNet50 or higher, but it is weak at a small scale. Creating a light-weight transformer architecture such as MobileNet is one of the future works of ViT research.

## 4.2. Object detection

We validate our backbones’ performances through object detection on COCO dataset [25] in DETR [4] which

	Standard	Occ	IN-A [15]	BGC [40]	FGSM [10]
PiT-S	80.8	74.6	21.7	21.0	29.5
ViT-S [34]	79.8	73.0	19.1	17.6	27.2
ResNet50 [13]	76.0	52.2	0.0	22.3	7.1
ResNet50 <sup>†</sup> [38]	79.0	67.1	5.4	32.7	24.7

Table 5. **ImageNet robustness benchmarks.** We compare three comparable architectures, PiT-B, ViT-S, and ResNet50 on various ImageNet robustness benchmarks, including center occlusion (Occ), ImageNet-A (IN-A), background challenge (BGC), and fast sign gradient method (FGSM) attack. We evaluate two ResNet50 models from the official PyTorch repository, and the well-optimized implementation [38], denoted as <sup>†</sup>.

is a transformer-based detector. We train each DETR with the different backbones including ResNet50, ViT-S, and our PiT-S. We follow the training setup of the original paper [4], but changed the image resolution to  $600 \times 400$  and training epoch 150. To match spatial sizes of backbones’ feature, we use dilated convolution [24] in the pooling layer of PiT and the last block of ResNet.

Table 4 shows the measured AP score on val2017. The detector based on PiT-S outperforms the detector with ViT-S. It shows that the pooling layer of PiT is effective not only for ImageNet classification but also for pretrained backbone for object detection. ViT detector and PiT detector have almost the same throughput. Since PiT-S originally has a higher throughput than ViT-S, it is possible to maintain the same throughput even when the computation of PiT is increased by the dilated convolution at the pooling layer. Although PiT detector cannot beat the performance of the ResNet50 detector, PiT detector has higher throughput, and the training setting was also not tuned for PiT detector. Additional investigation on the training settings for PiT detector would improve the performance of the PiT detector.

### 4.3. Robustness benchmarks

In this subsection, we investigate the effectiveness of the proposed architecture in terms of robustness against input changes. We presume that the existing ViT design concept, which keeps the spatial dimension from the input layer to the last layer, has two conceptual limitations: *Lack of background robustness* and *sensitivity to the local discriminative visual features*. We, therefore, presume that PiT, our new design choice with the pooling mechanism, performs better than ViT for the background robustness benchmarks and the local discriminative sensitivity benchmarks.

We employ four different robustness benchmarks. **Occlusion benchmark** measures the ImageNet validation accuracy where the center  $112 \times 112$  patch of the images is zero-ed out. This benchmark measures whether a model only focuses on a small discriminative visual feature or not. **ImageNet-A (IN-A)** is a dataset constructed by collecting

the failure cases of ResNet50 from the web [15] where the collected images contain unusual backgrounds or objects with very small size [23]. From this benchmark, we can infer how a model is less sensitive to unusual backgrounds or object size changes. However, since IN-A is constructed by collecting images (queried by 200 ImageNet subclasses) where ResNet50 predicts a wrong label, this dataset can be biased towards ResNet50 features. We, therefore, employ **background challenge (BGC)** benchmark [40] to explore the explicit background robustness. The BGC dataset consists of two parts, foregrounds, and backgrounds. This benchmark measures the model validation accuracy while keeping the foreground but adversarially changing the background from the other image. Since BGC dataset is built upon nine subclasses of ImageNet, the baseline random chance is 11.1%. Lastly, we tested adversarial attack robustness using the fast gradient sign method (FGSM) attack [10].

Table 5 shows the results. First, we observe that PiT shows better performances than ViT in all robustness benchmarks, despite they show comparable performances in the standard ImageNet benchmark (80.8 vs. 79.8). It supports that our pooling design choice makes the model less sensitive to the backgrounds and the local discriminative features. Also, we found that the performance drops for occluded samples by ResNet50 are much dramatic than PiT;  $80.8 \rightarrow 74.6$ , 5% drops for PiT,  $79.0 \rightarrow 67.1$ , 15% drops for ResNet50. This implies that ResNet50 focuses more on the local discriminative areas, by the nature of convolutional operations. Interestingly, in Table 5, ResNet50 outperforms vision transformer variants in the background challenge dataset (32.7 vs. 21.0). This implies that the self-attention mechanism unintentionally attends more backgrounds comparing to ResNet design choice. Overcoming this potential drawback of vision transformers will be an interesting research direction.

## 5. Conclusion

In this paper, we have shown that the design principle widely used in CNNs - the spatial dimensional transformation performed by pooling or convolution with strides, is not considered in transformer-based architectures such as ViT; ultimately affects the model performance. We have first studied with ResNet and found that the transformation in respect of the spatial dimension increases the computational efficiency and the generalization ability. To leverage the benefits in ViT, we propose a PiT that incorporates a pooling layer into ViT, and PiT shows that these advantages can be well harmonized to ViT through extensive experiments. Consequently, while significantly improving the performance of the ViT architecture, we have shown that the pooling layer by considering spatial interaction ratio is essential to a self-attention-based architecture.



## Acknowledgement

We thank NAVER AI Lab members for valuable discussion and advice. NAVER Smart Machine Learning (NSML) [20] has been used for experiments.

## References

- [1] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016. 3
- [2] Irwan Bello. Lambdanetworks: Modeling long-range interactions without attention. In *International Conference on Learning Representations*, 2021. 2, 3
- [3] Han Cai, Ligeng Zhu, and Song Han. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332*, 2018. 6
- [4] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *European Conference on Computer Vision*, pages 213–229. Springer, 2020. 1, 2, 6, 7, 8
- [5] Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on learning theory*, pages 698–728. PMLR, 2016. 1
- [6] Nadav Cohen and Amnon Shashua. Inductive bias of deep convolutional networks through pooling geometry. In *International Conference on Learning Representations*, 2016. 1
- [7] Zihang Dai, Guokun Lai, Yiming Yang, and Quoc V Le. Funnel-transformer: Filtering out sequential redundancy for efficient language processing. *arXiv preprint arXiv:2006.03236*, 2020. 2
- [8] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 1
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 1, 2, 4, 5, 6, 7
- [10] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. 8
- [11] Dongyoon Han, Jiwhan Kim, and Junmo Kim. Deep pyramidal residual networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5927–5935, 2017. 2
- [12] Dongyoon Han, Sangdoon Yun, Byeongho Heo, and Youngjoon Yoo. Rexnet: Diminishing representational bottleneck on convolutional neural network. *arXiv preprint arXiv:2007.00992*, 2020. 2, 3
- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 2, 3, 7, 8
- [14] Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 7
- [15] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. *arXiv preprint arXiv:1907.07174*, 2019. 8
- [16] Byeongho Heo, Sanghyuk Chun, Seong Joon Oh, Dongyoon Han, Sangdoon Yun, Gyuwan Kim, Youngjung Uh, and Jungwoo Ha. Adamp: Slowing down the slowdown for momentum optimizers on scale-invariant weights. In *International Conference on Learning Representations*, 2021. 6
- [17] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015. 6
- [18] Elad Hoffer, Tal Ben-Nun, Itay Hubara, Niv Giladi, Torsten Hoefer, and Daniel Soudry. Augment your batch: Improving generalization through instance repetition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8129–8138, 2020. 6
- [19] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019. 2, 3, 7
- [20] Hanjoo Kim, Minkyu Kim, Dongjoo Seo, Jinwoong Kim, Heungseok Park, Soeun Park, Hyunwoo Jo, KyungHyun Kim, Youngil Yang, Youngkwan Kim, Nako Sung, and Jungwoo Ha. NSML: meet the mlaas platform with a real-world case study. *CoRR*, abs/1810.09957, 2018. 9
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 1, 2
- [22] Boyi Li, Felix Wu, Kilian Q Weinberger, and Serge Belongie. Positional normalization. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d’Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 3
- [23] Xiao Li, Jianmin Li, Ting Dai, Jie Shi, Jun Zhu, and Xiaolin Hu. Rethinking natural adversarial examples for classification models. *arXiv preprint arXiv:2102.11731*, 2021. 8
- [24] Yi Li, Haozhi Qi, Jifeng Dai, Xiangyang Ji, and Yichen Wei. Fully convolutional instance-aware semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2359–2367, 2017. 8
- [25] 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 *European conference on computer vision*, pages 740–755. Springer, 2014. 2, 7
- [26] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018. 6
- [27] Filip Radenović, Giorgos Tolias, and Ondřej Chum. Fine-tuning cnn image retrieval with no human annotation. *IEEE*

- transactions on pattern analysis and machine intelligence*, 41(7):1655–1668, 2018. 1
- [28] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10428–10436, 2020. 6, 7
- [29] Prajit Ramachandran, Niki Parmar, Ashish Vaswani, Irwan Bello, Anselm Levskaya, and Jon Shlens. Stand-alone self-attention in vision models. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. 2
- [30] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018. 2, 3, 7
- [31] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [32] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. 2
- [33] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019. 1, 2, 3, 7
- [34] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. *arXiv preprint arXiv:2012.12877*, 2020. 3, 4, 5, 6, 7, 8
- [35] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. Attention is all you need. *arXiv preprint arXiv:1706.03762*, 2017. 1, 2
- [36] Huiyu Wang, Yukun Zhu, Bradley Green, Hartwig Adam, Alan Yuille, and Liang-Chieh Chen. Axial-deeplab: Stand-alone axial-attention for panoptic segmentation. In *European Conference on Computer Vision*, pages 108–126. Springer, 2020. 2
- [37] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7794–7803, 2018. 1, 2
- [38] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. 7, 8
- [39] Yuxin Wu and Kaiming He. Group normalization. In *Proceedings of the European conference on computer vision (ECCV)*, pages 3–19, 2018. 3
- [40] Kai Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or signal: The role of image backgrounds in object recognition. *arXiv preprint arXiv:2006.09994*, 2020. 8
- [41] 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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019. 6, 7
- [42] Sangdoon Yun, Seong Joon Oh, Byeongho Heo, Dongyoon Han, Junsuk Choe, and Sanghyuk Chun. Re-labeling imagenet: from single to multi-labels, from global to localized labels. *arXiv preprint arXiv:2101.05022*, 2021. 7
- [43] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Mueller, R Manmatha, et al. Resnest: Split-attention networks. *arXiv preprint arXiv:2004.08955*, 2020. 7