

A CLOSE LOOK AT TOKEN MIXER: FROM ATTENTION TO CONVOLUTION

Anonymous authors

Paper under double-blind review

ABSTRACT

There is an increasingly intensive debate about the effectiveness of ConvNets and Transformers in vision fields. Originating from the language processing community, Transformers show great promise for many vision tasks due to the insightful architecture design and attention mechanism. Nevertheless, we witnessed the revenge of ConvNets soon, surpassing Transformer variants in mainstream vision tasks. In this paper, we are not engaging in this debate; instead, we look into the details of attention and convolution. By looking into the self-attention responses in Transformers, we empirically find that **1.)** Vision Transformers present a query-irrelevant behavior in deep layers, where the attention maps exhibit nearly consistent contexts in global scope, regardless of the query patch position (also head-irrelevant). This phenomenon indicates that a global context may hide behind the self-attention mechanism. **2.)** The attention maps are intrinsically sparse; introducing the knowledge from ConvNets would largely smooth the attention and improve the performance. Motivated by these, we generalize self-attention formulation to abstract the query-irrelevant global context directly and further integrate the global context into convolutions. The resulting model, a Fully Convolutional Vision Transformer (*i.e.*, FCViT), purely consists of convolutional layers and firmly inherits the merits of both attention mechanism and convolutions, including dynamic property, weight sharing, and short- and long-range feature modeling, *etc.* Experimental results demonstrate the effectiveness of FCViT. With less than 14M parameters, our FCViT-S12 outperforms related work ResT-Lite by 3.7% top-1 accuracy on ImageNet-1K. When scaling FCViT to larger models, we still perform better than previous state-of-the-art ConvNeXt with even fewer parameters. FCViT-based models also demonstrate promising transferability to downstream tasks, like object detection, instance segmentation, and semantic segmentation. Codes and pre-trained models are available at: <https://anonymous.4open.science/r/FCViT-pytorch>.

1 INTRODUCTION

In the past few years, Vision Transformers Dosovitskiy et al. (2020); Touvron et al. (2021b) has dominated various visual tasks in the computer vision community. Although the costs (*i.e.*, parameters, and computations) are generally high, it is widely believed that Transformers are more likely to model better spatial relations and scale better with large model sizes and datasets compared with the conventional convolutions Li et al. (2021a). A common belief is that these gratifying virtues are credited to the self-attention mechanism. However, this plausible conjecture has been challenged recently. Studies Wightman et al. (2021); Ding et al. (2022); Liu et al. (2022) show that a ConvNet trained with powerful training recipes can also achieve competitive or even better performance, suggesting that a deep investigation of the feature modeling methods is worth further exploring.

To verify what has been learned by self-attention in vision transformers, we take a close look at the attention maps of the deep layers in several representative Vision Transformers, as shown in Fig. 1. Astonishingly, all these variants present a query-irrelevant behavior that reveals nearly consistent contexts in the global scope. This observation is a departure from the design philosophy of self-attention mechanism, indicating that a global context may be hidden behind the attention mechanism. Meanwhile, we notice that by introducing the knowledge from ConvNets to Vision Transformer, the attention map gets much smoother and pays more attention to objects. This phenomenon suggests that combining convolution and self-attention may lead to gratifying results.



Figure 1: Attention map visualizations of Vision Transformers Dosovitskiy et al. (2020); Touvron et al. (2021b). For each pair, we show the *query point* and its corresponding *attention map* (of last block and last head). All pre-trained models are downloaded from TIMM Wightman (2019). The right color bar identifies the value of normalized attention maps. Surprisingly, **the attention maps are almost the same, regardless of the query points**. Moreover, **the attention maps are intrinsically sparse if no knowledge from ConvNets is introduced**. Results from different models indicate that **these phenomena are common in Vision Transformers**. See supplementary for more examples and Sec. 1 for a more detailed analysis.

Motivated by aforementioned observations, we propose **Fully Convolutional Vision Transformer** (*i.e.*, FCViT) in this paper. Starting from the observation in Fig. 1, we progressively loosen the formulation of self-attention and abstract a global context that describes the global-range visual concepts. The global context is further dynamically introduced to local convolutional operations, making the efficient fusion of short- and long-range dependencies feasible. Intrinsically, FCViT is a pure ConvNet but a Transformer-like model that inherits the advantages of both Vision Transformer and ConvNet. Extensive experiments show that our FCViT consistently and significantly outperforms other methods with comparable costs. Remarkably, our FCViT achieves 80.9% top-1 accuracy on ImageNet-1K Deng et al. (2009) using only 14M parameters, outperforming ResT-Lite Zhang & Yang (2021) and PoolFormer-S12 by **3.7%**. Compared with state-of-the-art models like ConvNeXt Liu et al. (2022), we still present better results with even fewer parameters. Our FCViT also exhibits an excellent generalization ability. When transferred to downstream vision tasks like object detection, and instance segmentation, our FCViT consistently outperforms other methods.

2 RELATED WORK

Vision Transformers Dominate. Originating from natural language processing, Transformer Vaswani et al. (2017) targets to dynamically build mutual relationships for each token pair in a global scope. Motivated by the successes in language, tentative efforts have been made toward migrating Transformer to the vision community. The pioneering work ViT Dosovitskiy et al. (2020) has emerged as a promising approach for directly processing images using Transformer. Given an input image, ViT first tokenizes the input to non-overlapped patches and extracts token features via a stack of isotropic Transformer blocks. With an adequate training scheme, DeiT Touvron et al. (2021b) circumvents the problem of requiring large datasets. Since then, various Transformer variants have been springing up Liu et al. (2021); Yuan et al. (2021). Besides the aforementioned works, recent efforts Wang et al. (2021); Yang et al. (2022); Li et al. (2021b) mainly leverage the inductive biases from ConvNets to improve the performance. In this work, we push this trend further by connecting convolution and attention in FCViT. FCViT is a departure from standard Transformers, considering the fully convolutional operations; nevertheless, it unambiguously inherits the framework of the transformer architecture Yu et al. (2021) and enjoys the global receptive field as in Transformers.

ConvNets Strike Back. The necessity of multi-head self-attention in Transformers has been challenged, from language Jain & Wallace (2019); Dong et al. (2021) to vision community Yu et al.

(2021); Zhai et al. (2021); Tang et al. (2021). To put it another way, these critical investigations motivate lots of researchers to gush into the revitalization of ConvNets. A systematic study is ConvNeXt Liu et al. (2022), which reexamines the design spaces of a ConvNet by gradually modifying ResNet He et al. (2016) toward the standard vision Transformer Dosovitskiy et al. (2020) (*i.e.*, ViT). Astonishingly, ConvNeXt demonstrates that the resulting ConvNets compete favorably with Transformers in terms of both accuracy and scalability. Similar phenomena can also be observed in other ConvNets, like Bello et al. (2021); Ding et al. (2022). Redesigning the architectures *a la* the design philosophy of Transformers, ConvNets once again showed dominance in various visual tasks. This paper pushes the envelope further by inherently integrating virtues of ViTs with a fully convolutional network, which is where FCViT singularity lies.

Convolution Meets Attention. Another line of work on the topic of visual backbones learns to marriage the merits of both Transformers and ConvNets. That is, considering self-attention and convolution simultaneously in one token-mixer (mixing the spatial information). Dai *et al.* Dai et al. (2021) find that depthwise convolution and self-attention can be naturally unified in a token-mixer module. By doing so, CoAtNet Dai et al. (2021) effectively achieve better generalization and capacity. Similarly, CvT Wu et al. (2021) introduces convolutional token embedding and convolutional projection into vision transformers and empirically presents better performances. Works in Guo et al. (2021); Chen et al. (2022) bridge convolution module and attention module in a hybrid fashion. *Different from the aforementioned methods that connect convolution and self-attention in an explicit manner, without intrinsic dedicated designs involved, our FCViT subtly bridges global context and local structure in one unified operation.* Besides, tentative efforts have been made toward exploring the relationships between convolution and attention. In this vein, our analysis is similar to those taken by Cordonnier et al. (2020); Han et al. (2022) and allows us to further improve the performance.

3 FROM ATTENTION TO CONVOLUTION

By revisiting self-attention and convolution, we first bridge self-attention and convolution in one unified operation. A close look at the attention maps in ViTs also demonstrates the necessity of integrating convolution and attention. Motivated by this, a simple yet effective Transformer-like architecture is proposed to verify our findings. Fig. 2 intuitively shows one building block of FCViT.

3.1 A CLOSE LOOK AT SELF-ATTENTION AND CONVOLUTION

3.1.1 COMPARISON BETWEEN SELF-ATTENTION AND CONVOLUTION

As the main contribution in Transformers, self-attention effectively captures the long-distance dependencies and dynamically aggregates input features according to the query patch. Formally, given an input feature map $\mathbf{X} \in \mathbb{R}^{d \times n}$, where d is the embedding dimension, and n indicates the patch number, the self-attention mechanism adaptively aggregates global information for each patch by

$$y_i = \sum_{j=1}^n w(q_i, k_j) v_j, \quad \text{s.t., } w(q_i, k_j) = \text{softmax}(q_i^\top k_j) = \frac{\exp(q_i^\top k_j)}{\sum_{l=1}^n \exp(q_i^\top k_l)}, \quad (1)$$

where $q_i = \mathbf{W}_q x_i$, $k_i = \mathbf{W}_k x_i$, and $v_i = \mathbf{W}_v x_i$ are different embeddings from \mathbf{X} ; i, j , and l index a patch. For brevity, we ignore the positional encoding and dimensional scalar \sqrt{d} in Eq. 1.

Differently, the conventional convolution formulation can be written as:

$$y_i = \sum_{j \in \Omega} w_j x_j, \quad (2)$$

where Ω indicates a local receptive field (*e.g.*, a 3×3 kernel size).

Comparing the formulations, two distinct differences between attention and convolution would be:

- *The dynamic of weights aggregation:* self-attention dynamically aggregates the values v_j using $w(q_i, k_j)$, which is based on the input features x_i . On the contrary, convolution utilizes fixed weights w_j to fuse features, and the weights are shared across patches.
- *The size of receptive field:* self-attention can model long-distance dependencies among patches

directly, resulting in a global receptive field. Limited by the computations, convolution performs aggregation functions in a local region, like 3×3 and 7×7 . A sliding window pattern (with a small step size) guarantees the information flow among regions. By stacking multiple layers, convolution can progressively enlarge the receptive field.

3.1.2 ATTENTION MAP VISUALIZATION

Some conventional wisdom leans to credit the success of Vision Transformers to the attention mechanism, considering the two differences discussed above. We acknowledge the advantages of a dynamic scheme and large receptive field; however, a close look into the details is still worth further exploring. To this end, we visualize the attention maps of various Vision Transformers, including ViT-B, ViT-L, DeiT-B, and DeiT-B-Distill. Results in Fig. 1 present some interesting phenomena:

- *Observation 1*: the attention weights (see ViT-B, ViT-L, and DeiT-B) are relatively sparse, indicating that only several patches dominate the attention. By introducing the knowledge from convolution, the attention weights (see DeiT-B-Distill) are largely smoothed, and the performance is largely boosted as well (83.4% vs. 82.0% top-1 accuracy on ImageNet-1K).
- *Observation 2*: the attention maps consistently reveal a query-irrelevant behavior. Visually, the attention maps $w(q_i, k_j)$ are nearly the same for each testing model and image, regardless of the query patch q_i . This is a departure from the design philosophy of self-attention, where each patch should exhibit a particular attention map, different from each other.

These observations suggest that **1)** combining the knowledge from both self-attention and convolution can yield gratifying results; **2)** a query-irrelevant global context may work well for vision tasks.

3.1.3 FROM SELF-ATTENTION TO CONVOLUTION

Inspired by the aforementioned observations, we revisit the self-attention formulation and connect it to convolution. Given the Eq. 1, we first remove the querying patches as implied by observation 2. The resulting simplified attention can be written as

$$y = \sum_{j=1}^n w(k_j) v_j = \sum_{j=1}^n \frac{\exp(k_j)}{\sum_{l=1}^n \exp(k_l)} v_j = \mathcal{N}(\mathbf{K})\mathbf{V} = gc, \quad (3)$$

where $\mathcal{N}(\cdot)$ indicates a normalization function. Since Eq. 3 is purely based on the global input features \mathbf{X} and models the global information, we denote it as global context (*i.e.*, $gc \in \mathbb{R}^d$) for simplicity. Note that the objective of normalization function $\mathcal{N}(\cdot)$ is to generate a weight for v_i . Hence, it can be generalized to other forms besides softmax, like average or sigmoid function.

Furthermore, as indicated by observation 1, we are encouraged to marry the attention mechanism and convolution into a unified formulation. A simple implementation that bridges Eq. 3 and Eq. 2 is

$$y_i = \sum_{j \in \Omega} w_j(x_j + gc). \quad (4)$$

Conceptually, Eq. 4 introduces a global context into each patch and leverages a convolution to aggregate local information consequently. With Eq. 4, we efficiently introduce the global information into the local patches, avoiding the combination of two individual operations. Meanwhile, by employing the global context implemented in Eq 3, we significantly reduce the computational complexity of self-attention. Also, Eq. 4 can be easily implemented by a convolution layer, making our model concise yet effective. Next, we instantiate our FCViT based on the analysis above.

3.2 FULLY CONVOLUTIONAL VISION TRANSFORMER

3.2.1 GENERAL FRAMEWORK

Our FCViT follows the framework of MetaFormer Yu et al. (2021), which adopts hierarchical architecture with 4 stages that progressively reduce the spatial size. Given an input image, we first utilize overlapping patch embedding to tokenize images by linear mapping. In each stage, a series of isotropic FCViT blocks are utilized to extract features. FCViT block involves two independent modules, the token-mixer and channel-mixer (as well as residual connections and Layer

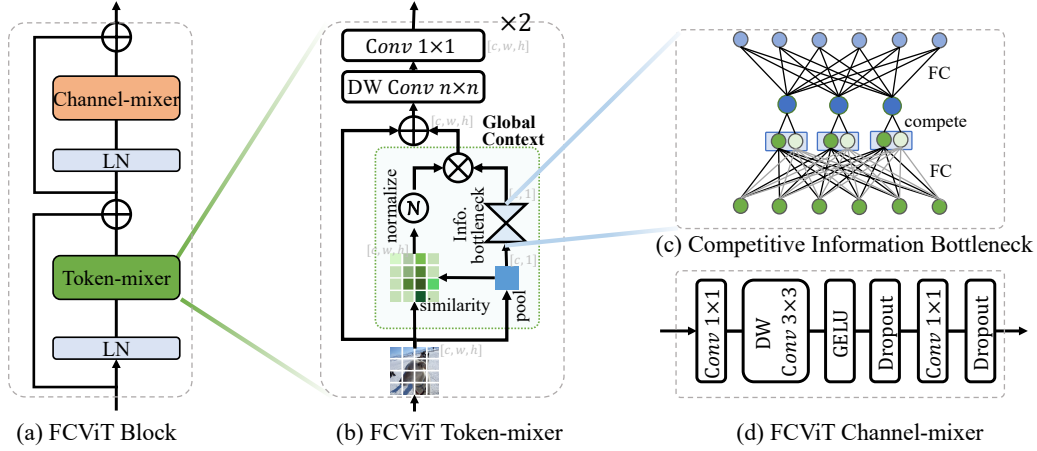


Figure 2: Illustration of an FCViT block. FCViT considers the block as a combination of token-mixer and channel-mixer. In the token-mixer, we dynamically integrate the global context with input tokens by the token-global similarity. A depth-wise convolution is employed to fuse local information. To improve the generalization of the global context, we introduce a competition-driven information bottleneck structure. Overall FCViT configurations are presented in the supplementary.

Normalization), as shown in Fig. 2. In the end, a classifier is employed to generate the classification logits. Varying the number of blocks in each stage and the channel number, we instantiate FCViT by FCViT-Tiny, FCViT-B12, and FCViT-B24, *etc.* Detailed configurations can be found in the supplementary. Next, we describe the detailed designs in our FCViT.

3.2.2 ENHANCED GLOBAL CONTEXT IN TOKEN-MIXER

Following Eq. 4, our token-mixer can be implemented by

$$\mathbf{Y} = \text{conv}_1 (\text{conv}_k (\mathbf{X} + gc)) = \text{conv} (\mathbf{X} + gc), \quad (5)$$

where conv_1 is a point-wise convolution and conv_k is a depth-wise convolution with kernel size of k . For convenience, we use conv to denote the combination of the two convolutions. We repeat the operation in Eq. 5 twice to achieve the best model size and accuracy trade-off. Specifically, we normalize the global context by average-pooling instead of softmax for simplicity. That is:

$$gc = \sum_{i=1}^n \frac{\mathbf{W}_v x_i}{n} = \mathbf{W}_v \sum_{i=1}^n \frac{x_i}{n}. \quad (6)$$

Dynamic Global Context. Directly fusing the global context $gc \in \mathbb{R}^d$ with input feature $\mathbf{X} \in \mathbb{R}^{d \times w \times h}$ may lead to limited improvements since gc is broadcasted equally for each patch in \mathbf{X} . In other words, Eq. 5 can be rewritten as

$$\mathbf{Y} = \text{conv} \left(\mathbf{X} + \mathbf{W}_v \sum_{i=1}^n \frac{x_i}{n} \right) = \text{conv} (\mathbf{X}) + \mathbf{W}_{conv}^\top \mathbf{W}_v \sum_{i=1}^n \frac{x_i}{n}, \quad (7)$$

where \mathbf{W}_{conv}^\top is the weights in the convolution layer. Eq. 7 explicitly demonstrates that directly fusing gc with \mathbf{X} is equal to fusing gc outside the convolution. Hence, we expect a dynamic fusion scheme that fuses the global context with bias based on \mathbf{X} .

We circumvent this problem by promoting token-global similarity. For clarity, we denote $\sum_{i=1}^n \frac{x_i}{n}$ as $\bar{\mathbf{X}}$. The similarity score $\mathbf{S} \in \mathbb{R}^{w \times h}$ is given by $\mathbf{S} = \mathbf{X} \bar{\mathbf{X}}$, which calculates the relations between each patch x_i and the global average pooling $\bar{\mathbf{X}}$. We further re-scale the similarity and update gc by

$$gc' = \left(\alpha \frac{\mathbf{S} - \mu_{\mathbf{S}}}{\sigma_{\mathbf{S}} + \epsilon} + \beta \right) gc = \mathbf{S}' gc, \quad (8)$$

where α and β are learnable scalars; $\mu_{\mathbf{S}}$ and $\sigma_{\mathbf{S}}$ are mean and standard deviation of \mathbf{S} ; $\epsilon = 1e^{-5}$ is for numerical stability. We use \mathbf{S}' and gc' to present normalized \mathbf{S} and updated gc , respectively. By doing so, each patch dynamically integrates the global context according to the similarity.

Multi-Group Similarity. Inspired by multi-head attention, we further extend our token-global similarity to a multi-group fashion in pursuit of better diversity. We first divide gc , \mathbf{X} , and $\bar{\mathbf{X}}$ into g groups along the channel dimension, respectively. We then perform the operation of Eq. 8 for each group and merge the outputs via concatenation. That is, $gc' = \text{concat}([\mathbf{S}'_1 gc_1, \mathbf{S}'_2 gc_2, \dots, \mathbf{S}'_g gc_g])$. Different from the multi-head attention that maps all channels into multiple heads, we leverage the group operation to reduce computational overheads. Ablation study in Fig. 3 visually shows the differences between our multi-group similarity and multi-head attention.

Competitive Information Bottleneck. Motivated by the information bottleneck method Tishby et al. (2000); Srivastava et al. (2013), we next introduce a competitive information bottleneck to further improve the generalization ability of gc and better describe the global context. We first squeeze $\bar{\mathbf{X}}$ by $\mathbf{W}_{s1} \in \mathbb{R}^{\frac{d}{r} \times d}$ and $\mathbf{W}_{s2} \in \mathbb{R}^{\frac{d}{r} \times d}$, respectively. Then, the two squeezed vectors compete to generate a bottleneck representation by maxout Goodfellow et al. (2013). Lastly, we recover the representation to original dimension by $\mathbf{W}_r \in \mathbb{R}^{d \times \frac{d}{r}}$. Our competitive information bottleneck can be written as $gc = \mathbf{W}_r \text{maxout}(\mathbf{W}_{s1} \bar{\mathbf{X}}, \mathbf{W}_{s2} \bar{\mathbf{X}})$. By default, we set r to 8. Then, the multi-group dynamic gc can be applied consequently. While the design is similar to SENet Hu et al. (2018b), we are in pursuit of competitive information and parameters reduction; no attention design is involved.

3.2.3 CONDUCTIVE IMPLEMENTATION DETAILS

While the token-mixer is the core of FCViT, some detailed designs are also instrumental. Different from the non-overlapping tokenization in ViT Dosovitskiy et al. (2020) and ConvNeXt Liu et al. (2022), we consider the overlapped patch embedding as presented in Wang et al. (2022), where each patch has small overlapping to ease the communication among patches. Also, a depth-wise convolution is introduced between the two point-wise convolution layers, as shown in Fig. 2.

4 EXPERIMENTS

We validate FCViT on ImageNet-1K Deng et al. (2009), MS COCO Lin et al. (2014), and ADE20K Zhou et al. (2017) datasets. We first demonstrate the effectiveness of FCViT on ImageNet-1K classification task. Extensive ablation studies provide a close look at the internal operations. Then, we transfer pre-trained models to object detection, instance segmentation, and semantic segmentation tasks to examine the generalization ability of FCViT.

4.1 IMAGE CLASSIFICATION ON IMAGENET-1K

Experimental Settings. We train FCViT models on the ImageNet-1K training set (with around 1.3M images) and evaluate upon the validation set. We follow the common training recipe in Dai et al. (2021); Wightman (2019); Touvron et al. (2021b); Yu et al. (2021). All our models are trained for 310 epochs using AdamW Loshchilov & Hutter (2019) with a momentum of 0.9 and a weight decay of 0.05. The learning rate is initialized to 0.001 and adjusted by cosine scheduler Loshchilov & Hutter (2017). By default, the models are trained on 8 A100 GPUs with a mini-batch size of 128 (1024 in total). Similar to previous works Guo et al. (2022); Touvron et al. (2021b), we employ Exponential Moving Average (EMA) to improve the training. Results are reported in Table 1.

Performance Analysis. Empirically, our FCViT outperforms related work by a clear margin. FCViT-Tiny outperforms DeiT-Tiny by **2.7%** (74.9% vs. 72.2%) top-1 accuracy, using much fewer parameters and FLOPs. When increasing the model size, FCViT consistently achieves a leading performance. FCViT-B12 achieves 80.9% accuracy with 14M parameters, outperforming ResT-Lite and PoolFormer-S12 by **3.7%**. Comparing with the state-of-the-art model ConvNeXt, we still present an improvement of 0.4% (82.5% vs. 82.1%) accuracy, with fewer parameters but similar FLOPs.

4.2 ABLATION STUDY

We next conduct extensive ablation studies to better understand and evaluate the utility of FCViT. Detailed analyses demonstrate the contribution of introducing the dynamic global context to local

Table 1: Comparison with SOTA backbones on ImageNet-1k benchmark. Throughput (images / s) is measured on a single A100 GPU with a batch size of 128. All models are trained and evaluated on 224x224 resolution. We use dots with different colors to present different types of token-mixer, **attention-based**, **convolution-based**, **MLP-based**, **Att&Conv.-based**, and **other** token-mixers. The best results are marked in **bold**. For more results of larger models, please see the supplementary.

Method	Param.	GFLOPs	Top-1	Top-5	Throughputs (images/s)
● PVTv2-B0 Wang et al. (2022)	3.4	0.6	70.5	-	-
● T2T-ViT-7 Yuan et al. (2021)	4.3	1.1	71.7	-	-
● DeiT-Tiny/16 Touvron et al. (2021b)	5.7	1.3	72.2	91.1	767.07
● FCViT-Tiny (Ours)	4.6	0.8	74.9	92.6	759.79
● PVT-Tiny Wang et al. (2021)	13.2	1.9	75.1	92.4	-
● ResT-Lite Zhang & Yang (2021)	10.49	1.4	77.2	93.7	-
● PoolFormer-S12 Yu et al. (2021)	11.9	2.0	77.2	-	764.03
● ResNet18 He et al. (2016)	12	1.8	69.8	89.1	926.73
● FCViT-B12 (Ours)	14	2.5	80.9	95.5	771.56
● DeiT-Small/16 Touvron et al. (2021b)	22.1	4.6	79.8	95.0	762.19
● T2T-ViT _t -14 Yuan et al. (2021)	21.5	6.1	81.7	-	-
● PVT-Small Wang et al. (2021)	24.5	3.8	79.8	94.9	-
● Swin-T Liu et al. (2021)	28.3	4.5	81.3	95.5	758.84
● ResT-Base Zhang & Yang (2021)	30.28	4.3	81.6	95.7	-
● ResMLP-24 Touvron et al. (2021a)	30.0	6.0	79.4	94.5	756.88
● AS-MLP-T Lian et al. (2022)	28	4.4	81.3	-	-
● PoolFormer-S24 Yu et al. (2021)	21.4	3.6	80.3	-	763.77
● CoAtNet-0 Dai et al. (2021)	25	4.2	81.6	-	-
● CvT-13 Wu et al. (2021)	20	4.5	81.6	-	-
● Conformer-Ti Peng et al. (2021)	23.5	5.2	81.3	-	-
● ResNet50 He et al. (2016)	26	4.1	80.4	94.6	770.34
● I-D-DW-Conv.-T Han et al. (2022)	26	4.4	81.8	-	-
● ConvNeXt-T Liu et al. (2022)	28.6	4.5	82.1	95.9	747.27
● FCViT-B24 (Ours)	25.7	4.7	82.5	96.2	706.97

regions. Component ablations further disentangle the effectiveness of each module. In this subsection, we conduct all experiments based on FCViT-T, and we train all models with a mini-batch size of 256.

Kernel Size. We vary the kernel size of convolutions in the token-mixer from 3 to 13. Different from previous work Liu et al. (2022),

we did not observe a linear correlation between kernel size and performance, and the performance gap is relatively small, ranging from 73.2% to 73.5%. One possible reason would be that we explicitly introduced the global context to the local tokens, which already enlarges the receptive field to a global range. Hence, the influence of kernel size is largely diluted. By default, we set the value to 11.

Global Context. The kernel contribution of FCViT is to introduce the global context to local tokens, converting self-attention to local convolution. Here, we evaluate the necessity of our global context. We first build a plain FCViT version as our baseline, where the global context is removed. Then, we integrate the dynamic version and competitive information bottleneck progressively. Table 2 shows the effectiveness of each component. Without Global Context, a plain FCViT architecture reaches 72.8% top-1 accuracy, which is still a leading result in the first block of Table 1. By introducing the global context, we slightly improve the plain network by 0.2%. When

Size	3	5	7	9	11	13
Acc.	73.2	73.2	73.4	73.2	73.5	73.2

Table 2: Ablation of Global Context. We test the necessity of global context, dynamic scheme and competitive information bottleneck, individually.

GC	Dy. GC.	Comp. Info.	params	FLOPs	top-1 (%)
✗	✗	✗	4.2M	0.8G	72.8 _{±0.0}
✓	✗	✗	4.4M	0.8G	73.0 _{±0.2}
✓	✓	✗	4.4M	0.8G	73.5 _{±0.7}
✓	✓	✓	4.5M	0.8G	73.7 _{±0.9}

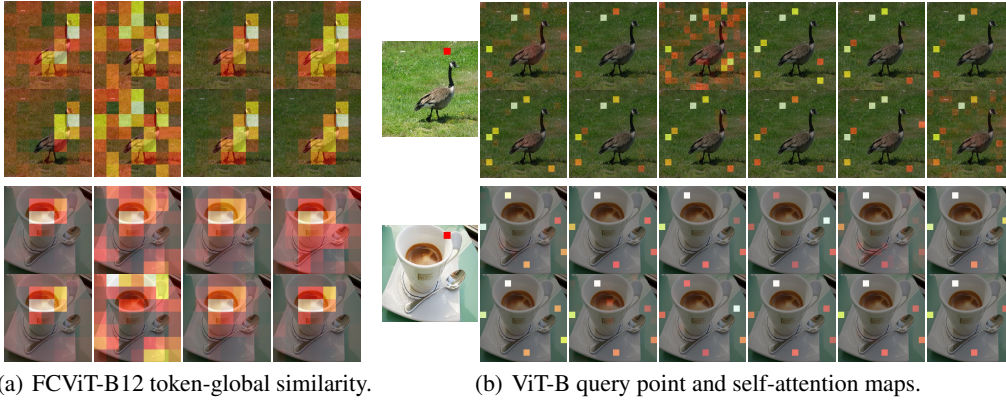


Figure 3: Visual comparisons of FCViT-B12 similarity and ViT-B attention map. We plot all the outputs of the last block for the two models (8 groups for FCViT and 12 heads for ViT). Compared to ViT, the results indicate that: **1)** our FCViT focuses more on the objects; **2)** FCViT presents more diversities than multi-head attention, whose attention maps from different heads are nearly the same.

integrating the global context as presented in Eq. 8, the performance is increased to 73.5%, indicating the effectiveness of the dynamic scheme. Remarkably, the computation and parameter overheads are negligible, leading to almost no additional inference time. When introducing the competitive information bottleneck, we further improve the performance to 73.7%, exhibiting 0.9% improvements over the plain network without global context. In the following experiments, we will consider the global context with the dynamic scheme and competitive information bottleneck by default.

Group Number for Token-Global Similarity.

As introduced, we intentionally copy the design philosophy of multi-head attention to emphasize the similarity in different groups. The table on the right reveals the benefits of multi-group similarity. When the group number is increased to 8, we achieve the best result (0.3% improvements). Note that only memory operation is introduced for multi-group similarity (computational complexity and operational redundancy are almost the same), and the running speed is not increased. Empirically, we set the group number to 8 in all other experiments.

Group	1	4	8	16	32
Acc.	73.2	73.4(↑)	73.5	73.3	73.3

4.3 MULTI-GROUP TOKEN-GLOBAL SIMILARITY vs. MULTI-HEAD SELF-ATTENTION.

Similar to the multi-head self-attention mechanism, we consider the token-global similarity in a grouped-channel fashion. In Fig. 3(a), we showcase the similarity scores of FCViT-B12 for all 8 groups in the last block. Similarly, we present the attention maps of ViT-B for all 12 heads in Fig. 3(b). We randomly fix the query point for all validating images since Fig. 1 demonstrated that the attention map of the last block is nearly query-irrelevant.

The visualization results exhibit some interesting insights. **1)** Besides the query-irrelevance, we notice that the attention map of ViT is head-irrelevant as well, as shown in Fig. 3(b). This is a departure from the design philosophy of multi-head attention mechanism, where different heads suppose to exhibit different attention. Differently, our FCViT presents diverse similarities among the groups. This phenomenon can be explained by the formulations. In multi-head self-attention, all channels are employed to map to multiple heads. Without constraints (*e.g.*, loss functions), the heads are hard to present meaningful diversities, especially when the model converges in the last layers. As a comparison, our channel-group similarity calculates the similarity for a group of channels individually, making each group self-contained and preventing the similarities collapsed. **2)** Our FCViT focuses more on the objects, while ViT cannot exhibit such a desirable property. Among all testing examples and groups, FCViT explicitly emphasizes the tokens within the objects. Recall that the similarity is calculated between local tokens and the global context. This result highlights the validity of our global context. Meanwhile, by dynamically fusing the global context, our FCViT further enhances the representational ability of local tokens in turn.

Table 3: COCO object detection and instance segmentation results using Mask-RCNN (1×).

backbone	Params	AP ^{box}	AP ^{box} ₅₀	AP ^{box} ₇₅	AP ^{mask}	AP ^{mask} ₅₀	AP ^{mask} ₇₅
ResNet-18	31.2M	34.0	54.0	36.7	31.2	51.0	32.7
PoolFormer-S12	31.6M	37.3	59.0	40.1	34.6	55.8	36.9
PVT-Tiny	32.9M	36.7	59.2	39.3	35.1	56.7	37.3
FCViT-B12	34.3M	42.3	64.2	46.2	38.6	61.1	41.3
ResNet-50	44.2M	38.0	58.6	41.4	34.4	55.1	36.7
PoolFormer-S24	41.0M	40.1	62.2	43.4	37.0	59.1	39.6
PVT-Small	44.1M	40.4	62.9	43.8	37.8	60.1	40.3
FCViT-B24	43.1M	44.1	65.4	48.4	39.9	62.4	42.7

4.4 OBJECT DETECTION AND INSTANCE SEGMENTATION ON MS COCO

We next probe the transferability of FCViT on downstream tasks, including object detection and instance segmentation. We conduct experiments on MS COCO 2017 benchmark Lin et al. (2014). We train and evaluate Mask R-CNN He et al. (2017) with FCViT backbone initialized with classification pre-trained weights. For a fair comparison, we follow the settings in PVT Wang et al. (2021) and PoolFormer Yu et al. (2021) that adopts the 1× training schedule (*i.e.*, 12 epochs) based on the MMDetection Chen et al. (2019) framework. Results in Table 3 demonstrate that our FCViT significantly outperforms related work by a clear margin. Particularly, our FCViT brings **5.0 points of mAP^{box}** and **4.0 points of mAP^{mask}** against PoolFormer-S12 at comparable settings. Even compared with larger models (parameters in backbones are doubled) like PoolFormer-S24 and ResNet-50, our FCViT-Tiny still exhibits significantly better results, showing gratifying transferability.

4.5 SEMANTIC SEGMENTATION ON ADE20K

We also evaluate our FCViT equipped with Semantic FPN Kirillov et al. (2019) on ADE20K Zhou et al. (2017) dataset for the semantic segmentation task. ADE20K contains 150 semantic categories and consists of 20k, 2k, and 3k images for training, validation, and testing, respectively. We use the ImageNet-pretrained backbone model taken from Table 1. Following Yu et al. (2021), we train all our models for 40k iterations with a batch size of 32. All our models are trained using AdamW optimizer with an initial learning rate of 2×10^{-4} . We decay the learning rate by a polynomial decay schedule with a power of 0.9.

The results compared with previous work are presented in Table 4. Empirically, our FCViT achieves promising performance on the semantic segmentation task. Compared with PoolFormer-S12, FCViT-B12 improves the performance by 6.1% mIoU using a similar number of parameters. Compared with Twins-PCPVT-S, FCViT-B24 improves the performance by 1.2% mIoU using even fewer parameters. In line with the results in object detection and instance segmentation tasks, FCViT also demonstrated promising transferability for the semantic segmentation task.

Table 4: Semantic segmentation performance of different backbones with Semantic FPN on the ADE20K validation set.

Backbone	Params	mIoU(%)
ResNet18	15.5M	32.9
PVT-Tiny	17.0M	35.7
PoolFormer-S12	15.7M	37.2
FCViT-B12	17.8M	43.3
ResNet50	28.5M	36.7
PVT-Small	28.2M	39.8
Poolformer-S24	23.2M	40.3
Twins-PCPVT-S	28.4M	44.3
Twins-SVT-S	28.3M	42.6
FCViT-B24	25.3M	45.5

5 CONCLUSION

In this paper, we first take a close look at two foundational token-mixers, attention and convolution. Detailed comparisons and visual analysis motivate us to connect attention with convolution by abstracting a global context and dynamically fusing it with local tokens. The resulting model, Fully Convolutional Vision Transformer (FCViT), explicitly embraces the advantages of both convolution and attention, and exhibits promising efficiency and performance. The multi-group token-global similarity and competitive information bottleneck further boost the representational ability of our method. Experiments on multiple tasks validate the utility of FCViT and confirm our analysis. We hope that our FCViT can bring new insights to the vision community.

ETHICS STATEMENT

We adhere closely to the ICLR ethical research rules and guidelines in our paper. All of the datasets we used are openly accessible, and all relevant articles and source codes are properly cited.

REPRODUCIBILITY STATEMENT

We follow ICLR reproducibility guidelines and take several steps to guarantee the reproducibility of our work, including:

- We upload our codes, pre-trained models, and the training log files to an anonymous GitHub repository (<https://anonymous.4open.science/r/FCViT-pytorch>), as stated in the abstract. Clear descriptions of each individual tasks and experiments are presented to ease the reviewing work.
- We clearly present the design of our method in Section 3.2.
- Detailed framework and more experiments are presented in appendix.

REFERENCES

- 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. *NeurIPS*, 2021.
- Kai Chen, Jiaqi Wang, Jiangmiao Pang, Yuhang Cao, Yu Xiong, Xiaoxiao Li, Shuyang Sun, Wansen Feng, Ziwei Liu, Jiarui Xu, Zheng Zhang, Dazhi Cheng, Chenchen Zhu, Tianheng Cheng, Qijie Zhao, Buyu Li, Xin Lu, Rui Zhu, Yue Wu, Jifeng Dai, Jingdong Wang, Jianping Shi, Wanli Ouyang, Chen Change Loy, and Dahua Lin. MMDetection: Open mmlab detection toolbox and benchmark. *arXiv preprint arXiv:1906.07155*, 2019.
- Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobile-former: Bridging mobilenet and transformer. *CVPR*, 2022.
- Jean-Baptiste Cordonnier, Andreas Loukas, and Martin Jaggi. On the relationship between self-attention and convolutional layers. In *ICLR*, 2020.
- Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. *NeurIPS*, 2021.
- 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.
- Xiaohan Ding, Xiangyu Zhang, Yizhuang Zhou, Jungong Han, Guiguang Ding, and Jian Sun. Scaling up your kernels to 31x31: Revisiting large kernel design in cnns. *CVPR*, 2022.
- Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *ICML*. PMLR, 2021.
- 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.
- Ian Goodfellow, David Warde-Farley, Mehdi Mirza, Aaron Courville, and Yoshua Bengio. Maxout networks. In *ICML*. PMLR, 2013.
- Jianyuan Guo, Kai Han, Han Wu, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. Cmt: Convolutional neural networks meet vision transformers. *arXiv preprint arXiv:2107.06263*, 2021.
- Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual attention network. *arXiv preprint arXiv:2202.09741*, 2022.

- Qi Han, Zejia Fan, Qi Dai, Lei Sun, Ming-Ming Cheng, Jiaying Liu, and Jingdong Wang. On the connection between local attention and dynamic depth-wise convolution. In *ICLR*, 2022.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. In *ICCV*, 2017.
- Jie Hu, Li Shen, Samuel Albanie, Gang Sun, and Andrea Vedaldi. Gather-excite: Exploiting feature context in convolutional neural networks. *NeurIPS*, 2018a.
- Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *CVPR*, 2018b.
- Sarthak Jain and Byron C. Wallace. Attention is not explanation. In *NAACL-HLT (1)*, pp. 3543–3556, 2019. URL <https://aclweb.org/anthology/papers/N/N19/N19-1357/>.
- Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *CVPR*, 2019.
- Yanhao Li, Saining Xie, Xinlei Chen, Piotr Dollar, Kaiming He, and Ross Girshick. Benchmarking detection transfer learning with vision transformers. *arXiv preprint arXiv:2111.11429*, 2021a.
- Yawei Li, Kai Zhang, Jiezhong Cao, Radu Timofte, and Luc Van Gool. Localvit: Bringing locality to vision transformers. *arXiv preprint arXiv:2104.05707*, 2021b.
- Dongze Lian, Zehao Yu, Xing Sun, and Shenghua Gao. As-mlp: An axial shifted mlp architecture for vision. *ICLR*, 2022.
- 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*. Springer, 2014.
- 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.
- Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CVPR*, 2022.
- Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *ICLR*, 2017.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *ICLR*, 2019.
- Zhiliang Peng, Wei Huang, Shanzhi Gu, Lingxi Xie, Yaowei Wang, Jianbin Jiao, and Qixiang Ye. Conformer: Local features coupling global representations for visual recognition. In *ICCV*, 2021.
- Rupesh K Srivastava, Jonathan Masci, Sohrab Kazerounian, Faustino Gomez, and Jürgen Schmidhuber. Compete to compute. *NeurIPS*, 2013.
- Chuanxin Tang, Yucheng Zhao, Guangting Wang, Chong Luo, Wenxuan Xie, and Wenjun Zeng. Sparse mlp for image recognition: Is self-attention really necessary? *arXiv preprint arXiv:2109.05422*, 2021.
- Naftali Tishby, Fernando C Pereira, and William Bialek. The information bottleneck method. *arXiv preprint physics/0004057*, 2000.
- Ilya O Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, et al. Mlp-mixer: An all-mlp architecture for vision. *NeurIPS*, 2021.
- Hugo Touvron, Piotr Bojanowski, Mathilde Caron, Matthieu Cord, Alaaeldin El-Nouby, Edouard Grave, Gautier Izacard, Armand Joulin, Gabriel Synnaeve, Jakob Verbeek, et al. Resmlp: Feedforward networks for image classification with data-efficient training. *arXiv preprint arXiv:2105.03404*, 2021a.

- Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Herve Jegou. Training data-efficient image transformers and distillation through attention. In *ICML*, 2021b.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *NeurIPS*, 2017.
- 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.
- Wenhai Wang, Enze Xie, Xiang Li, Deng-Ping Fan, Kaitao Song, Ding Liang, Tong Lu, Ping Luo, and Ling Shao. Ptv2: Improved baselines with pyramid vision transformer. *Computational Visual Media*, 8(3):1–10, 2022.
- Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019.
- Ross Wightman, Hugo Touvron, and Hervé Jégou. Resnet strikes back: An improved training procedure in timm. *arXiv preprint arXiv:2110.00476*, 2021.
- Sanghyun Woo, Jongchan Park, Joon-Young Lee, and In So Kweon. Cbam: Convolutional block attention module. In *ECCV*, 2018.
- Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. Cvt: Introducing convolutions to vision transformers. In *ICCV*, pp. 22–31, 2021.
- Jianwei Yang, Chunyuan Li, Pengchuan Zhang, Xiyang Dai, Bin Xiao, Lu Yuan, and Jianfeng Gao. Focal self-attention for local-global interactions in vision transformers. *NeurIPS*, 2021.
- Jianwei Yang, Chunyuan Li, and Jianfeng Gao. Focal modulation networks. *arXiv preprint arXiv:2203.11926*, 2022.
- Tan Yu, Xu Li, Yunfeng Cai, Mingming Sun, and Ping Li. S2-mlp: Spatial-shift mlp architecture for vision. In *WACV*, 2022.
- Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and Shuicheng Yan. Metaformer is actually what you need for vision. *CVPR*, 2021.
- 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.
- Shuangfei Zhai, Walter Talbott, Nitish Srivastava, Chen Huang, Hanlin Goh, Ruixiang Zhang, and Josh Susskind. An attention free transformer. *arXiv preprint arXiv:2105.14103*, 2021.
- Qinglong Zhang and Yu-Bin Yang. Rest: An efficient transformer for visual recognition. *NeurIPS*, 2021.
- Bolei Zhou, Hang Zhao, Xavier Puig, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Scene parsing through ade20k dataset. In *CVPR*, 2017.

A FCViT ARCHITECTURE & CONFIGURATIONS

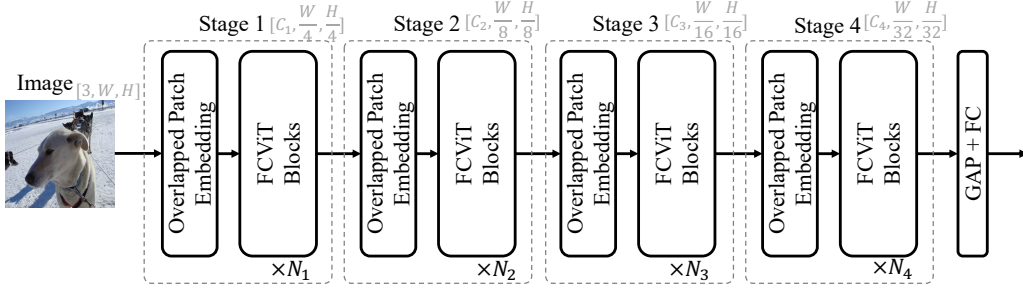


Figure 4: The overall architecture of FCViT. Given an input image, FCViT considers 4 stages to process and extract features. In each stage, one overlapped patch embedding is employed to reduce the spatial size, and a stack of FCViT blocks is employed to model feature relations.

Table 5: Model configurations for our FCViT. We introduce four configurations FCViT-Tiny, FCViT-B12, FCViT-B24, and FCViT-B48, with different model capacities.

Stage	Size	Layer	FCViT-Tiny	FCViT-B12	FCViT-B24	FCViT-B48
S1	56x56	Patch Embed.	patch_size = 7 stride = 4 dim = 32	patch_size = 7 stride = 4 dim = 64	patch_size = 7 stride = 4 dim = 64	patch_size = 7 stride = 4 dim = 64
	56x56	FCViT Blocks	[mlp_r. = 8 dim = 32] × 3	[mlp_r. = 8 dim = 64] × 2	[mlp_r. = 8 dim = 64] × 4	[mlp_r. = 8 dim = 64] × 8
S2	28x28	Patch Embed.	patch_size = 3 stride = 2 dim = 64	patch_size = 3 stride = 2 dim = 128	patch_size = 3 stride = 2 dim = 128	patch_size = 3 stride = 2 dim = 128
	28x28	FCViT Blocks	[mlp_r. = 8 dim = 64] × 3	[mlp_r. = 8 dim = 128] × 2	[mlp_r. = 8 dim = 128] × 4	[mlp_r. = 8 dim = 128] × 8
S3	14x14	Patch Embed.	patch_size = 3 stride = 2 dim = 160	patch_size = 3 stride = 2 dim = 320	patch_size = 3 stride = 2 dim = 320	patch_size = 3 stride = 2 dim = 320
	14x14	FCViT Blocks	[mlp_r. = 4 dim = 160] × 5	[mlp_r. = 4 dim = 320] × 6	[mlp_r. = 4 dim = 320] × 12	[mlp_r. = 4 dim = 320] × 24
S4	7x7	Patch Embed.	patch_size = 3 stride = 2 dim = 320	patch_size = 3 stride = 2 dim = 512	patch_size = 3 stride = 2 dim = 512	patch_size = 3 stride = 2 dim = 512
	7x7	FCViT Blocks	[mlp_r. = 4 dim = 320] × 2	[mlp_r. = 4 dim = 512] × 2	[mlp_r. = 4 dim = 512] × 4	[mlp_r. = 4 dim = 512] × 8

A.1 FCViT ARCHITECTURE

We first describe the high-level architecture overview of our FCViT. Following hierarchical Transformers Liu et al. (2021) and classical ConvNets He et al. (2016), our FCViT includes four stages to reduce the spatial size gradually. Specifically, we reduce the spatial size by a factor of 4 in the first stage and a factor of 2 for the rest. In the end, a classifier (which is combined by a global averaging pooling and a fully connected layer) is employed to give the classification logits.

Next, we exhaustively present the details of overlapped patch embedding. Different from the non-overlapping patch embedding used in the original ViT and ConvNeXt, our overlapped patch embedding introduces appropriate overlapping between neighbor patches. In our implementation, we achieve this by a convolution operation, where the stride size is smaller than the kernel size. For example, the kernel size is set to 7, and the step size is set to 4 in the first stage. As a result, half of the pixels in a patch are overlapped with neighbor patches. To facilitate the training, we add Layer Normalization after the convolution operation, which is consistent with previous work.

Table 6: Comparison with larger SOTA backbones on ImageNet-1k benchmark. All models are trained and evaluated on 224×224 resolution. We use dots with different colors to present different types of token-mixer, **attention-based**, **convolution-based**, **MLP-based**, **Att&Conv.-based**, and **other** token-mixers. The best results are marked in **bold**.

Method	Params. (M)	GFLOPs	Top-1 (%)	Top-5 (%)
● T2T-ViT _t -19 Yuan et al. (2021)	39.2	9.8	82.4	-
● PVT-Medium Wang et al. (2021)	44.2	6.7	81.2	95.6
● Swin-S Liu et al. (2021)	49.6	8.7	83.0	96.2
● PVTv2-B3 Wang et al. (2022)	45.2	6.9	83.2	-
● Focal-S Yang et al. (2021)	51.1	9.1	83.5	-
● T2T-ViT _t -24 Yuan et al. (2021)	64.0	15.0	82.3	-
● AS-MLP-S Lian et al. (2022)	50	8.5	83.1	-
● Mixer-B/16 Tolstikhin et al. (2021)	59.0	11.6	76.4	-
● PoolFormer-M36 Yu et al. (2021)	56	9.1	82.1	-
● PoolFormer-M48 Yu et al. (2021)	73	11.9	82.5	-
● S2-MLP-deep Yu et al. (2022)	51	10.5	80.7	95.4
● CoAtNet-1 Dai et al. (2021)	42	8.4	83.3	-
● CvT-21 Wu et al. (2021)	32	7.1	82.5	-
● Conformer-S Peng et al. (2021)	37.7	10.6	83.4	-
● ResNet101 He et al. (2016)	45	7.9	81.3	-
● ConvNeXt-S Liu et al. (2021)	50.1	8.7	83.1	-
● FCViT-B48 _(ours)	49.1	9.2	83.6	96.5

A.2 FCViT CONFIGURATIONS

The detailed configurations of our FCViT variants are shown in Table 5. We instantiate four variants of FCViT for different model capacities. Empirically, we set the MLP-ratio in the channel-mixer (FFN) to 8 for the first and second stages. For the blocks number, we follows the design in Liu et al. (2022); Yu et al. (2021): the number is set to n , n , $3n$, and n for each stage.

B MORE EXPERIMENTS

We present more experiments, including large model comparisons and more ablation studies.

Larger Model. First, we introduce the performance of larger models as a supplementary. Clearly, our FCViT-B48 still dominates the classification performance with relatively few parameters for the larger models. This promising result indicates that our FCViT not only surpasses related work significantly with a small model capability, but also enjoys a good scalability.

Integration of Attention Module. Some works introduce attention modules laterally to enhance the model performance with minimal additional overheads Hu et al. (2018b). We examine the effectiveness of combining attention module with our FCViT to verify if attention modules can boost our FCViT as well.

To achieve this, we first build two fully convolutional attention modules, the spatial-attention module and the channel-attention module. Then, we plug the two modules into token-mixer and channel-mixer, respectively. Fig. 5 show the attention module augmented FCViT block. In spatial attention, we introduce two depth-wise convolutions

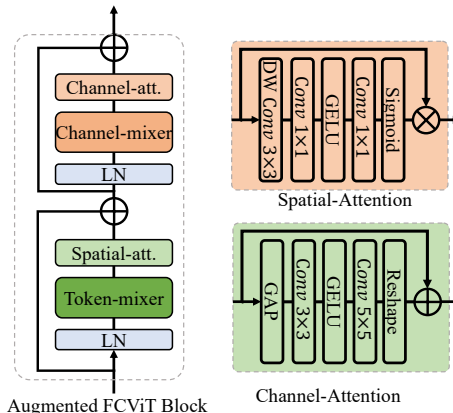


Figure 5: Attention module augmented FCViT block, and details in spatial- and channel-attention.

Table 8: Ablation of each component in enhanced global context.

GC	Dy. GC	Compet. Info	Group Sim.	top1
✓	✓	✓	✓	73.7
✗	✗	✗	✗	72.8 (↓ 0.9)
✓	✗	✓	✓	73.2 (↓ 0.5)
✓	✓	✗	✓	73.5 (↓ 0.2)
✓	✓	✓	✗	73.3 (↓ 0.4)

Table 9: Replacing Enhanced Global Context to Convolution or Attention. See the description for detail modifications.

Token-mixer	Parameters	FLOPs	Blocks	Top-1
Self-attention	4.8M	0.8G	[3,3,5,2]	74.7
Conv-3x3	4.8M	0.9G	[3,3,9,2]	74.6
Conv-11x11	4.8M	1.0G	[3,3,7,2]	74.5
FCViT-Tiny	4.6M	0.8G	[3,3,5,2]	74.9

and two point-wise convolutions to model the local spatial attention. In channel-attention, we first abstract the global feature via global average pooling, then two 1D convolutions are applied along the channel dimension to module the channel dependencies. Note that we employ the summation for the channel-attention in our implementation. Note that both our spatial- and channel-attention are based on convolutional operation, aligning with our motivation and design philosophy.

Results in Tale 7 exhibit some interesting phenomena. Different from previous works Hu et al. (2018b); Woo et al. (2018); Hu et al. (2018a), additional attention modules improve our FCViT marginally, only 0.1% to 0.2% top-1 accuracy and no improvements on top-5 accuracy. The following two aspects can explain this: 1) our FCViT explicitly inherits both global- and local-range feature modeling abilities, making additional attention modules show limited improvements; 2) the training recipe further limits the contributions of attention modules. Unlike conventional ConvNets trained with 100 epochs with simple data augmentation methods, recent works train all models by 300 epochs and with more and better training strategies. Besides our tailored convolutional attention modules, SE module can also be employed as a channel-attention module in our method, but no improvements were observed. Considering the computation and parameter overheads, we use our channel-attention module, which introduces 10 parameters. Considering the limited contribution of additional attention modules and strong baselines achieved by FCViT, we did not include convolutional attention modules in our architecture.

Table 7: Ablation on attention modules.

Spatial-att.	Channel-att.	top-1	top-5
✗	✗	74.9	92.6
✓	✗	74.9	92.6
✗	✓	75.0	92.6
✓	✓	75.1	92.5

Repeated token-mixer. As shown in the FCViT Token-mixer, we repeat the global context and depth-wise convolution by two times. By reducing this operation repetition time to one, the performance decreases 0.9%. More repetitions would largely increase the computational overhead. Hence, we set the repetition to 2 in our implementation.

Depth-wise Convolution in Channel-Mixer(FFN). Results indicate that this simple modification can improve the performance by 1.2% (73.7% vs. 74.9%). Empirically, we set the kernel size to 3. Other values may lead to better performance, but not the key contributions to our work.

Removing each component in enhanced global context. Our enhanced global context has three components: dynamic global context, multi-group similarity, and competitive information bottleneck. We remove each component individually and present the results in Table 8. Clearly, each component can contribute to our FCViT. Combining them together, we achieve the best performance.

Replacing enhanced global context with conv or attention. Next, we replace the global context module with multi-head attention or convolution. It is relatively hard to compare FCViT and the

related replacements fairly since the parameters and flops may change a lot. To this end, we tailor each variant to achieve a roughly fair comparison. For the self-attention variant, we use adaptive-average-pooling (to a resolution of 7) to reduce the complexity. We employ 4 heads, and each head has 32 dimensions. For the Convolution variants, we repeat convolution twice in each block and increase the block number in stage 3 to achieve similar computations.

The results are presented in Table 9. Using even fewer parameters and FLOPs, our FCViT-Tiny still achieves better results than the self-attention and convolution variants. Notice that convolution variants are deeper, and the self-attention variant introduces avg-pooling (which can be considered as convolution). All these variants train slower than our FCViT (around 1.2x to 1.3x time cost). Results showcased that our FCViT still outperforms all related variants, demonstrating the effectiveness of our method.

C MORE VISUALIZATIONS

Similarity map in FCViT. We also present more token-global similarity results on different examples, as shown in Fig 6. We use the pre-trained FCViT-B24 for illustration.




Figure 6: More token-global similarity visualization. Obviously, each group exhibits distinct similarities, and the similarities in all groups consistently emphasize the critical regions.

Attention Map in ViTs. To further support our motivation, we plot more attention maps in ViTs on more examples. The results are shown in Fig. 7.




Figure 7: More attention maps of ViT variants. ViTs show a query-irrelevant behavior in general, while some bias may happen (very few, *e.g.*, second example in ViT-L).