

# SWAT: Spatial Structure Within and Among Tokens

Kumara Kahatapitiya and Michael S. Ryoo

Stony Brook University

{kkahatapitiy, mryoo}@cs.stonybrook.edu

## Abstract

Modeling visual data as tokens (i.e., image patches) using attention mechanisms, feed-forward networks or convolutions has been highly effective in recent years. Such methods usually have a common pipeline: a tokenization method, followed by a set of layers/blocks for information mixing, both within and among tokens. When image patches are converted into tokens, they are often flattened, discarding the spatial structure within each patch. As a result, any processing that follows (eg: multi-head self-attention) may fail to recover and/or benefit from such information. In this paper, we argue that models can have significant gains when spatial structure is preserved during tokenization, and is explicitly used during the mixing stage. We propose two key contributions: (1) Structure-aware Tokenization and, (2) Structure-aware Mixing, both of which can be combined with existing models with minimal effort. We introduce a family of models (**SWAT**), showing improvements over the likes of DeiT, MLP-Mixer and Swin Transformer, across multiple benchmarks including ImageNet classification and ADE20K segmentation. Our code is available at [github.com/kkahatapitiya/SWAT](https://github.com/kkahatapitiya/SWAT).

## 1 Introduction

Convolutional architectures (CNNs) [He *et al.*, 2016] have been dominant in computer vision for a while now. When they were first introduced for large-scale training in image domain, their benefits were quickly realized over Multi-layer Perceptrons (MLPs). In addition to efficient weight sharing, the inductive bias generated by exploring the local structure in images was one of the key factors for its success [LeCun *et al.*, 2015]. In language domain however, CNNs were less effective due to lack of such strong local structure. Consequently, attention mechanisms emerged dominant, exploring long-range relationships and modeling language as a sequence [Dauphin *et al.*, 2017]. More recently, attention models— specifically Transformers [Vaswani *et al.*, 2017], have been extended to represent visual data [Dosovitskiy *et al.*, 2021], with the key concept of tokenizing an input image to create a sequence

(or a set), often discarding their structure. Within a short period of time, such *token-based models* (i.e., class of models such as ViTs [Dosovitskiy *et al.*, 2021] and MLP-Mixers [Tolstikhin *et al.*, 2021]) have outperformed CNNs on most visual tasks. However, we ask, could the spatial structure— when preserved, benefit token-based models and further improve their performance?

Token-based models in computer vision are rapidly evolving. From Vision Transformers [Dosovitskiy *et al.*, 2021] to MLP-Mixers [Tolstikhin *et al.*, 2021] and hybrid-architectures [Peng *et al.*, 2021; Wu *et al.*, 2021], intriguing concepts are being introduced and tested on tasks including classification [Dosovitskiy *et al.*, 2021; Touvron *et al.*, 2021b; Liu *et al.*, 2021], detection [Zhu *et al.*, 2020; Dai *et al.*, 2021] and segmentation [Xie *et al.*, 2021; Duke *et al.*, 2021], to name a few. All such models can be framed with two main components: (1) *Tokenization*, which converts image patches into tokens, and (2) *Mixing* (attention-based as in Multi-head Self Attention (MHSA), MLP-based or convolution-based), which shares information within and among tokens. In general, during tokenization, an image patch is directly mapped into a token, not preserving the spatial structure within a token. After this mapping, models usually focus on global patterns among tokens, without capturing local spatial structure *within* tokens.

Structure is an important cue in visual data. In images, 2D spatial structure preserves geometry and object-part relationships. Simply put, structure gives meaning to visual data in human perspective. However, in machine perspective, if a jumbled set of image patches are tokenized and processed through a token-based model, it can give the same classification performance (as it is a set operator), even though the input is really meaningless to a human [Naseer *et al.*, 2021]. This is in fact a drawback of token-based models (eg: can be prone to such an adversarial attack), which could be addressed by structure-aware modeling. Not only the structure among tokens, but also the structure within tokens is equally-important which is often discarded during tokenization. It is particularly beneficial to maintain the structure within tokens for fine-grained prediction tasks such as segmentation.

In this paper, we propose to preserve and make use of the spatial structure both within and among tokens. To do this we focus on two components: (1) *Structure-aware Tokeniza-*

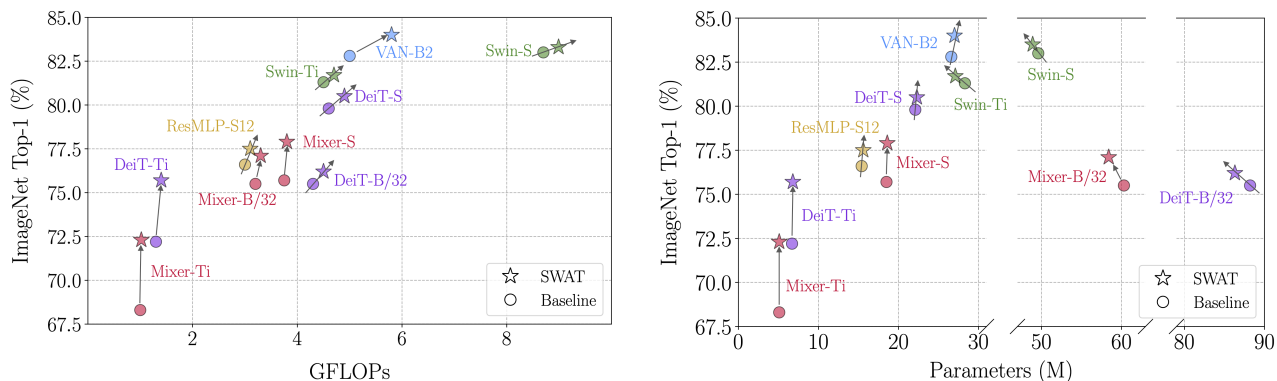


Figure 1: **Performance vs. Complexity** on ImageNet-1K [Deng *et al.*, 2009]. We implement our proposed (1) *Structure-aware Tokenization*, and (2) *Structure-aware Mixing* in common token-based architectures including DeiT [Touvron *et al.*, 2021b], Swin [Liu *et al.*, 2021], MLP-Mixer [Tolstikhin *et al.*, 2021], ResMLP [Touvron *et al.*, 2021a] and VAN [Guo *et al.*, 2022]. The resulting family of **SWAT** models consistently outperform their counterparts, with minimal increase in complexity. We consider the system-agnostic metrics such as FLOPs and Parameters as the complexity measures here.

tion and (2) *Structure-aware Mixing*<sup>1</sup>, both of which can be adopted in existing token-based architectures with minimal effort. Our Structure-aware Tokenization converts image patches to tokens, but *preserves the spatial structure within a patch as channel segments of the corresponding token*. Our Structure-aware Mixing benefits from the preserved structure by *considering local neighborhoods both within and among tokens*, based on 2D convolutions. We also refer to this as token mixing with channel structure and channel mixing with token structure. With these two contributions, we introduce a family of models: **SWAT**, and compare against common baselines such as DeiT [Touvron *et al.*, 2021b], Swin Transformer [Liu *et al.*, 2021], MLP-Mixer [Tolstikhin *et al.*, 2021], ResMLP [Touvron *et al.*, 2021a] and VAN [Guo *et al.*, 2022]. Our models show consistent improvements over baseline models on multiple benchmarks including ImageNet-1K [Deng *et al.*, 2009] classification and ADE20K [Zhou *et al.*, 2019] semantic segmentation. We further visualize fine-grained attention patterns captured by our structure-aware modeling. Performance gains on ImageNet-1K classification against complexity (measured by system-agnostic metrics such as FLOPs and Parameters) are shown in Fig. 1.

## 2 Related Work

**Token-based models:** Transformer architectures from language domain [Vaswani *et al.*, 2017; Devlin *et al.*, 2019] have been recently adopted to visual data in the seminal work ViT [Dosovitskiy *et al.*, 2021]. Even though attention mechanisms already existed in computer vision [Wang *et al.*, 2018; Zhao *et al.*, 2020], their true potential was realized when introduced with tokenization. Since then, a variety of token-based models have been introduced, some with the use of MLPs [Tolstikhin *et al.*, 2021; Touvron *et al.*, 2021a] or convolutions [Trockman and Kolter, 2022; Liu *et al.*, 2022]. DeiT [Touvron *et al.*, 2021b] introduces

an efficient training recipe, and [Caron *et al.*, 2021; Ransinghe *et al.*, 2022] use self-supervision. Swin Transformer [Liu *et al.*, 2021] introduces attention within shifted-windows, while downsampling progressively similar to [Heo *et al.*, 2021; Wang *et al.*, 2021; Fan *et al.*, 2021]. Another direction explores efficiency of such models [Zhai *et al.*, 2021; Bello, 2020; Graham *et al.*, 2021; Tang *et al.*, 2021; Yue *et al.*, 2021; Ryoo *et al.*, 2021].

**Token adoption in vision tasks:** Token-based models are already applied in most vision applications, including classification [Touvron *et al.*, 2021b; Liu *et al.*, 2021], object detection [Zhu *et al.*, 2020; Carion *et al.*, 2020], segmentation [Xie *et al.*, 2021; Duke *et al.*, 2021], image generation [Cao *et al.*, 2021; Esser *et al.*, 2021], video understanding [Nagrani *et al.*, 2021; Fan *et al.*, 2021; Arnab *et al.*, 2021; Dai *et al.*, 2022], dense prediction [Yang *et al.*, 2021a; Ranftl *et al.*, 2021], point clouds processing [Zhao *et al.*, 2021; Guo *et al.*, 2021] and reinforcement learning [Chen *et al.*, 2021a; Shang *et al.*, 2022].

**Structure with token-based models:** Some prior work in token-based models have explored structure, using hybrid architectures with convolutions [Xiao *et al.*, 2021; Peng *et al.*, 2021; d’Ascoli *et al.*, 2021]. A structure-based grouping method is proposed in T2T-ViT [Yuan *et al.*, 2021b]. With a complementary motivation to ours, TNT [Han *et al.*, 2021] and NesT [Zhang *et al.*, 2021] both consider a sub-token structure within tokens, but introduce additional tokens and become heavier with extra processing. [Yuan *et al.*, 2021a] has similarities with our channel mixing with token structure. Models such as ConvMixer [Trockman and Kolter, 2022], ConvNeXt [Liu *et al.*, 2022] and VAN [Guo *et al.*, 2022] also consider a convolutional design as ours (w/ Pointwise Conv and Depthwise Conv). However, they only consider structure *among* tokens, not structure *within* tokens. To our knowledge, this is the first work to preserve structure within tokens, without extra tokens or processing, i.e., with a minimal change in footprint.

<sup>1</sup>Information sharing based on either attention (MHSA), MLPs or convolutions is commonly referred to as *Mixing* in this paper.

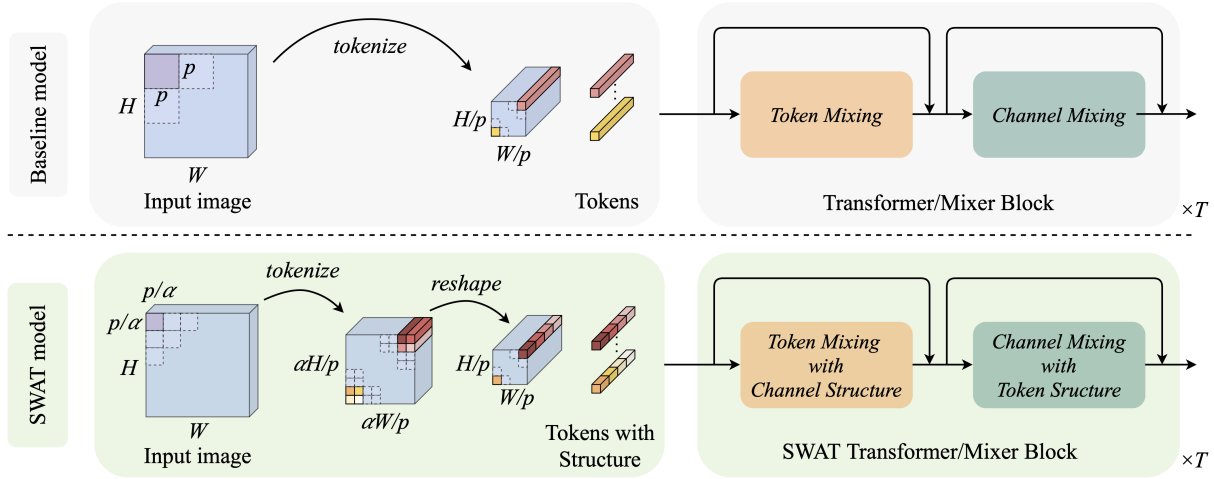


Figure 2: **SWAT Overview:** We show the architecture of SWAT (bottom) and a baseline model (top) in this figure. We propose two main contributions: (1) Structure-aware Tokenization and, (2) Structure-aware Mixing, which can be applied to common Transformer or Mixer architectures with minimal effort. Structure-aware tokenization preserves the spatial structure within a token, as channel segments. Simply put, first, we tokenize with a patch size  $(p/\alpha \times p/\alpha)$  instead of  $(p \times p)$ , resulting in  $\times \alpha^2$  more intermediate tokens. Next, we restructure  $\alpha \times \alpha$  neighboring tokens into one token (concatenating in channel dimension), which gives the same number of tokens (and the channel dimension) into Mixing operations, ensuring no additional cost in the downstream. However, now, this newly-preserved structure within tokens can be explored for better downstream processing, which was very limited previously. More on how we use the structure in Mixing is shown in Fig. 3 and Fig. 4.

### 3 Spatial Structure Within and Among Tokens

In SWAT family of models, we explore the benefits of preserving spatial structure not only among tokens, but within tokens as well. To do this with a general framework, we consider all token-based models (eg: ViTs [Dosovitskiy *et al.*, 2021], Mixers [Tolstikhin *et al.*, 2021]) as a unified architecture, which consists of two main components: (1) *Tokenization*, for converting image patches into tokens, and, (2) *Mixing*, for sharing information within and among tokens. Mixing can mean either the use of Multi-Layer Perceptron (MLP), Multi-Headed Self-Attention (MHSA) or convolution for information sharing. In this framework, we suggest improvements to both Tokenization and Mixing. When these components are adopted *together* in a network, it can preserve and utilize the spatial structure. Namely, we introduce *Structure-aware Tokenization* and *Structure-aware Mixing*, which we describe below in detail.

#### 3.1 Structure-aware Tokenization

Here, we propose to preserve the spatial structure within tokens, not imposing any additional burden on downstream processing. The idea is to keep spatial information within tokens separated as its channel segments, so that the ‘mixing’ component can later take advantage of it. In general, image patches are converted into tokens by sliding a large convolutional kernel with a stride (eg: a  $16 \times 16$  kernel with a stride of 16), which extracts a set of tokens. In such a setting, all the spatial information within a patch is directly fused into the channels of the corresponding token, losing the explicit structure in the process. In our method, we replace this direct fusion, retaining structural information within tokens.

More concretely, let us consider an input image of size

$H \times W \times 3$ , and a baseline tokenizer which converts image patches into tokens by extracting non-overlapping patches of size  $p \times p$ . This is usually implemented as a convolutional layer with  $C$  kernels of size  $(p \times p)$ , applied at a stride of  $p$ . The output here will be an  $H/p \times W/p$  2D structure of tokens, which is reshaped to create a sequence of  $HW/p^2$  tokens of embedding dimension  $C$  (refer Fig. 2 top). Even though these tokens are processed downstream as a sequence, they can be reshaped back into the original 2D structure of  $H/p \times W/p$  whenever necessary. It has been observed that the tokens preserve this structure (*among* tokens) through skip connections and positional encodings [Caron *et al.*, 2021; Naseer *et al.*, 2021], even after a series of Mixing blocks. However, the structure within a  $p \times p$  patch is irreversibly lost, i.e., although each token is a linear abstraction of  $p \times p$  pixels, remapping the token back to its original  $p \times p$  shape in subsequent layers is not directly feasible.

In contrast, the proposed tokenizer in SWAT retains the structure *within* a token (refer Fig. 2 bottom). We do this by first having  $C/\alpha^2$  convolutional kernels of size  $(p/\alpha \times p/\alpha)$  (where  $\alpha > 1$ ) and applying it with a stride of  $p/\alpha$ . The resulting intermediate set of tokens will have a 2D structure of  $\alpha H/p \times \alpha W/p$  and a dimension of  $C/\alpha^2$ . Next, such  $\alpha \times \alpha$  neighboring tokens are reshaped into a single token (concatenating in the channel dimension), creating the same number of tokens  $HW/p^2$  as the baseline, with the embedding dimension of  $C$ . By doing so, we now have an  $\alpha \times \alpha$  2D structure within each token— as its channel segments, which can be preserved throughout downstream processing, by the same principles: skip connections and (optional) positional embeddings. Note that the SWAT tokenizer will have a fewer parameters, in fact,  $3Cp^2/\alpha^4$ , compared to that of the baseline

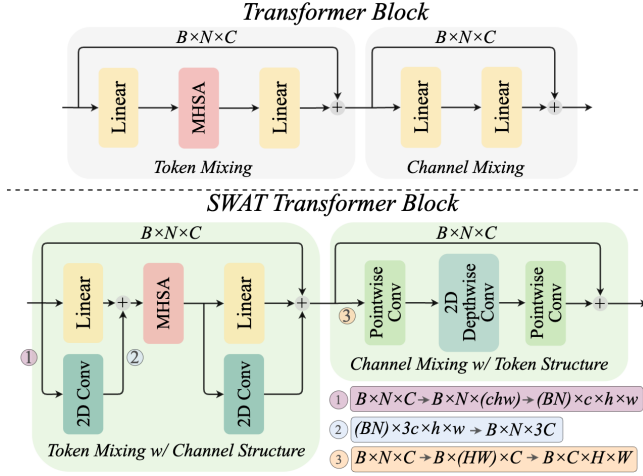


Figure 3: **SWAT Transformer Block**: We benefit from 2D channel structure (within tokens) in token mixing, and 2D token structure (among tokens) in channel mixing: (1) We insert a 2D Conv in-parallel to each Linear projection in attention (MHSA) block, applied on a reshaped input. It explicitly considers a structured local neighborhood *within* a token during token mixing. (2) In channel mixing, we first replace Linear layers with Pointwise Conv as a design decision for implementation simplicity. Next, we insert a 2D Depthwise Conv to consider a structured local neighborhood *among* tokens. Key tensor reshape operations are highlighted.

( $3Cp^2$ ), which can impair the learning capacity. To avoid this in practice, we use a bottleneck structure of multiple layers instead of a single convolution layer (still having the same downsampling factor of  $1/\alpha$  as the baseline), which will enable the tokenizer to have an equivalent capacity, while introducing structure within tokens.

### 3.2 Structure-aware Mixing

To make use of the structured tokens (w/ spatial structure both within and among) generated by the SWAT tokenizer, we propose *Structure-aware Mixing*. The idea is straightforward: when we have such a 2D structure, the corresponding elements (either tokens or channels) will have the notion of neighboring elements in the 2D space, which gives an inductive bias that we can benefit from. Our approach uses this locality in a form of 2D convolutions, mixing information in a local region of elements, in addition to the usual global information sharing in Transformer/Mixer models. We present this idea in two parts: (1) *Token Mixing with Channel Structure* and, (2) *Channel Mixing with Token Structure*.

#### Token Mixing with Channel Structure

Token Mixing happens in different ways in Transformers [Dosovitskiy *et al.*, 2021; Touvron *et al.*, 2021b] and Mixers [Tolstikhin *et al.*, 2021; Touvron *et al.*, 2021a]. In Transformers, each token attends to every other token pairwise and dynamically (w/ input-dependent weights). In an attention block, a MHSA layer is sandwiched between two Linear projection layers. Here, by design, token mixing (i.e., information sharing among tokens) happens while also mixing channels. These Linear layers may reshuffle channels and

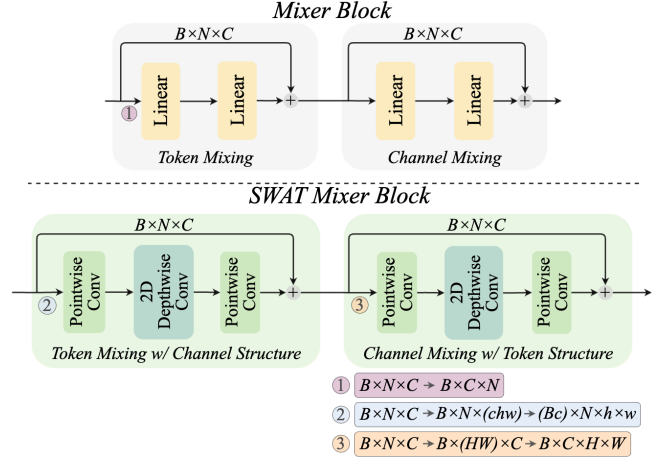


Figure 4: **SWAT Mixer Block**: In both token mixing and channel mixing, we first replace Linear layers with Pointwise Conv, applied on a reshaped input (eg: a Linear layer on a  $(B \times N \times C)$  shaped tensor equals to a Pointwise Conv on  $(B \times C \times N)$ , in a PyTorch-like channel-first Conv implementation). This is a design decision for implementation simplicity, which makes no change in how an input is processed. Now, we can easily explore the 2D channel structure (*within* tokens) in token mixing and, 2D token structure (*among* tokens) in channel mixing, by inserting a 2D Depthwise Conv. Key tensor reshape operations are highlighted.

waste our newly-introduced structure within tokens, as there is not even a skip-connection to save it. In contrast, in Mixers, token mixing is done with static relations (w/ learned weights), while not reshuffling channels. Simply put, tokens are mixed channel-wise, without damaging the structure within tokens. Therefore, we follow different designs in Transformers and Mixers to introduce our *token mixing with channel structure*.

**Transformers:** We insert a 2D Conv in-parallel<sup>2</sup> to the Linear layers before and after MHSA, to explore the channel structure (structure within tokens). See Fig. 3 bottom-left. After SWAT tokenizer, the channel dimension  $C$  has an internal structure of  $c \times h \times w$  (as in Fig. 2, with usual notation), which we use to reshape the input as,

$$B \times N \times C \rightarrow B \times N \times (chw) \rightarrow (BN) \times c \times h \times w.$$

Here,  $B$  represents batch,  $N$ , num. of tokens and  $C$ , embedding dimension. When a 2D Conv is applied on this tensor<sup>3</sup>, it can mix channel information similar to a Linear layer, but also considering the inductive bias of channel structure.

**Mixers:** In Mixers, we first replace the two Linear layers in token mixing with Pointwise  $1 \times 1$  Conv. See Fig. 4 bottom-left. We do this just to simplify the implementation, w/o changing the underlying operation (i.e., Linear =  $1 \times 1$  Conv). Applying a Linear layer on a tensor of shape  $B \times C \times N$  is the same as applying a Pointwise Conv on a tensor  $B \times N \times C$  (again, we consider a PyTorch-like implementation of channel-first Conv

<sup>2</sup>Why in-parallel? To retain a capacity (params) similar to the baseline. Refer to Appendix for more details.

<sup>3</sup>Here we consider a PyTorch-like channel-first implementation of convolution (eg: 2D Conv has an input shape of  $B \times C \times H \times W$ ).

and channel-last Linear). Now, we can conveniently consider the 2D structure in channels (*within* tokens). Next, we insert a 2D Depthwise Conv in-between the Pointwise Conv layers, applied on a reshaped input as,

$$B \times N \times C \rightarrow B \times N \times (chw) \rightarrow (Bc) \times N \times h \times w.$$

Altogether, this token mixing block now considers the channel structure (i.e., structure *within* tokens).

### Channel Mixing with Token Structure

Channel Mixing operation is the same for both Transformers and Mixers. In a baseline, two Linear layers are applied on an input tensor shaped as  $B \times N \times C$  to mix channel information. In SWAT, we wish to do this while considering the token structure. Hence, we replace the two Linear layers with the same sandwich block: 2D Depthwise Conv in-between two Pointwise Conv, applied on an input reshaped as,

$$B \times N \times C \rightarrow B \times (HW) \times C \rightarrow B \times C \times H \times W.$$

See Fig. 3 or Fig. 4 bottom-right. This channel mixing block now considers token structure (i.e., structure *among* tokens).

Specific hyperparameter settings and ablations related to (1) newly-introduced structure within tokens, and (2) level of structure-awareness in mixing, are included in Appendix. When experimenting with pyramid architectures (eg: Swin), we need to explicitly preserve structure when downsampling, and how we do this is also described in Appendix.

## 4 Experiments

In this section, we evaluate our family of models, SWAT on image classification and semantic segmentation. We use Imagenet-1K [Deng *et al.*, 2009] and ADE20K [Zhou *et al.*, 2019] as benchmarks to compare against common Transformer/Mixer/Conv architectures such as DeiT [Touvron *et al.*, 2021b], Swin [Liu *et al.*, 2021], MLP-Mixer [Tolstikhin *et al.*, 2021], ResMLP [Touvron *et al.*, 2021a] and VAN [Guo *et al.*, 2022]. In our ablations, we further evaluate the benefits of preserving structure.

### 4.1 ImageNet Classification

ImageNet-1K [Deng *et al.*, 2009] is a commonly-used classification benchmark, with 1.2M training images and 50K validation images, annotated with 1000 categories. For all our models, we report Top-1 (%) accuracy on single-crop evaluation with complexity metrics such as Parameters and FLOPs. We train all our models for 300 epochs on inputs of  $224 \times 224$  using the `timm` [Wightman, 2019] library. We use the original hyperparameters for all backbones, without further tuning. All models are trained with Mixed Precision.

**SWAT is generally-applicable and scalable:** In Table 1, we present the performance of SWAT with the two main types of token-based models: those using attention (MHSA) such as DeiT [Touvron *et al.*, 2021b] and Swin [Liu *et al.*, 2021], or those using MLPs such as Mixer [Tolstikhin *et al.*, 2021]. In both model families, SWAT consistently outperforms the baselines across different model scales, verifying that our Structure-aware Tokenization and Structure-aware Mixing can

Model	Model scale	Top-1 (%)	Params. (M)	FLOPs (G)
DeiT (Touvron <i>et al.</i> )	Ti	72.2	5.7	1.3
	S	79.8	22.1	4.6
	B/32	75.5	88.2	4.3
SWAT <sub>DeiT</sub> (ours)	Ti	<b>(+3.5)</b> 75.7	5.8	1.4
	S	<b>(+0.7)</b> 80.5	22.3	4.9
	B/32	<b>(+0.7)</b> 76.2	86.3	4.5
Mixer (Tolstikhin <i>et al.</i> )	Ti	68.3	5.1	1.0
	S	75.7	18.5	3.8
	B/32	75.5	60.3	3.2
SWAT <sub>Mixer</sub> (ours)	Ti	<b>(+4.0)</b> 72.3	5.1	1.0
	S	<b>(+2.2)</b> 77.9	18.6	3.8
	B/32	<b>(+1.6)</b> 77.1	58.4	3.3
Swin (Liu <i>et al.</i> )	Ti	81.3	28.3	4.5
	S	83.0	49.6	8.7
SWAT <sub>Swin</sub> (ours)	Ti	<b>(+0.4)</b> 81.7	27.1	4.7
	S	<b>(+0.3)</b> 83.3	48.9	9.1

Table 1: **SWAT is generally-applicable and scalable.** We compare SWAT with DeiT [Touvron *et al.*, 2021b], MLP-Mixer [Tolstikhin *et al.*, 2021] and Swin [Liu *et al.*, 2021] on ImageNet-1K. We report the performance in Tiny, Small and Base/32 (i.e., patch size of  $32 \times 32$ ) configurations. SWAT models consistently outperform their counterparts with minimal change in parameters or computations. All models are trained for 300 epochs at  $224 \times 224$  resolution. Performance improvement is in **bold**.

be applied in both cases. Specifically, we consider Tiny, Small and Base/32 (i.e., patch size of  $32 \times 32$ ) model scales, with varying range of parameters and computations. These are standard models reported in previous work. We implement our tokenizer and replace Transformer/Mixing blocks with ours in each configuration (eg: DeiT-Ti  $\rightarrow$  SWAT<sub>DeiT</sub>-Ti). In all configurations, SWAT models show consistent improvements. In SWAT<sub>DeiT</sub>, Tiny version achieves the highest gain of +3.5%, with +0.7% in Small and Base/32. In SWAT<sub>Mixer</sub>, all Tiny (+4.0%), Small (+2.2%) and Base/32 (+1.6%) versions show a considerable improvement over baselines. SWAT<sub>Swin</sub> shows +0.4% w/ Tiny and +0.3% w/ Small models. Overall, SWAT models have minimal (or no) increment in parameters or computations. The performance vs. complexity graphs are shown in Fig. 1.

**SWAT is competitive with SOTA:** In Table 2, we implement SWAT with multiple families of token-based models, either Transformer/Mixer/Convolutional, including DeiT [Touvron *et al.*, 2021b], Swin [Liu *et al.*, 2021], Mixer [Tolstikhin *et al.*, 2021], ResMLP [Touvron *et al.*, 2021a] and VAN [Guo *et al.*, 2022]. We report the performance in mid-sized (14-30M parameters) standard configurations. We use the same hyperparameter settings and training recipes as the corresponding original baselines. We observe consistent gains in SWAT family of models: +2.2% in SWAT<sub>Mixer</sub> and +1.2% in SWAT<sub>ResMLP</sub>, +0.7% in SWAT<sub>DeiT</sub>, +0.4% in SWAT<sub>Swin</sub> and +0.6% in SWAT<sub>VAN</sub>, with minimal change in parameters and computations compared to baselines. This further shows that SWAT can be generally-adopted to any token-based architecture with minimal effort and cost.

Model	Type	Top-1 (%)	Params. (M)	FLOPs (G)	
ResNet (He <i>et al.</i> )	CNN	78.8	25.6	4.1	
ResNeXt* (Xie <i>et al.</i> )		77.6	25.0	4.3	
EfficientNet* (Tan and Le)		82.6	19.3	4.4	
RegNetY* (Radosavovic <i>et al.</i> )		79.4	20.6	4.0	
ConvMixer (Trockman and Kolter)		80.2	21.1	-	
ConvNeXt (Liu <i>et al.</i> )		82.1	29.0	4.5	
Mixer (Tolstikhin <i>et al.</i> )	MLP	75.7	18.5	3.8	
SWAT <sub>Mixer</sub> (ours)		(+2.2) 77.9	18.6	3.8	
gMLP (Touvron <i>et al.</i> )		79.6	20.0	4.5	
ResMLP* (Touvron <i>et al.</i> )		76.6	15.4	3.0	
SWAT <sub>ResMLP</sub> * (ours)		(+1.2) 77.8	15.6	3.1	
PoolFormer (Yu <i>et al.</i> )		80.3	21.4	3.6	
CycleMLP (Chen <i>et al.</i> )		81.6	27.0	3.9	
DeiT (Touvron <i>et al.</i> )		Attention	79.8	22.1	4.6
SWAT <sub>DeiT</sub> (ours)			(+0.7) 80.5	22.3	4.9
T2T-ViT (Yuan <i>et al.</i> )			81.5	21.5	4.8
TNT (Han <i>et al.</i> )	81.5		23.8	5.2	
NesT (Zhang <i>et al.</i> )	81.5		17.0	5.8	
PVT (Wang <i>et al.</i> )	79.8		24.5	3.8	
Twins (Chu <i>et al.</i> )	81.7		24.0	2.8	
Focal (Yang <i>et al.</i> )	82.2		29.1	4.9	
Swin (Liu <i>et al.</i> )	81.3		28.3	4.5	
SWAT <sub>Swin</sub> (ours)	(+0.4) 81.7		27.1	4.7	
ConViT (d'Ascoli <i>et al.</i> )	Hybrid	81.3	27.0	5.4	
CvT (Wu <i>et al.</i> )		81.6	20.0	4.5	
Conformer (Peng <i>et al.</i> )		81.3	23.5	5.2	
CeiT (Yuan <i>et al.</i> )		82.0	24.2	4.8	
MobileFormer* (Chen <i>et al.</i> )		79.3	14.0	0.5	
VAN (Guo <i>et al.</i> )		82.8	26.6	5.0	
SWAT <sub>VAN</sub> (ours)		(+0.6) 83.4	27.0	5.8	

Table 2: **SWAT is competitive with SOTA.** We report experiments on ImageNet-1K with different families of token-based models in mid-sized configurations (14-30M params.). We implement SWAT with DeiT [Touvron *et al.*, 2021b], Swin [Liu *et al.*, 2021], MLP-Mixer [Tolstikhin *et al.*, 2021], ResMLP [Touvron *et al.*, 2021a] and VAN [Guo *et al.*, 2022] baselines, and train with original hyperparameter settings. SWAT outperforms all baselines consistently with minimal change in complexity, showing competitive performance with SOTA models. In general, models are trained for 300 epochs at 224 × 224 resolution (exceptions denoted with \* are discussed in appendix). Performance improvement is in **bold**.

**SWAT shows more fine-grained attention patterns:** In Fig. 5, we visualize token attention values in Tiny configurations of DeiT [Touvron *et al.*, 2021b] and SWAT<sub>DeiT</sub>. We use the code from DINO [Caron *et al.*, 2021] paper as a base. However, in our models, since we do not use a class token, we cannot visualize the attention on a single token as in [Caron *et al.*, 2021]. Instead, we show the attention maps of the final layer of each model, averaged across tokens. We consider larger image size (1024 × 1024) compared to training (224 × 224) to get higher resolution visualizations. We use the same patch size of 16 and interpolate positional encodings accordingly. We can see clear differences between the attention in DeiT [Touvron *et al.*, 2021b] and SWAT<sub>DeiT</sub>. In SWAT, we have more contrastive attention which resembles

Model	Structure-aware			Top-1 (%)	Params. (M)	FLOPs (G)
	Tokenize	Tk. Mix.	Ch. Mix.			
DeiT				73.3	5.72	1.25
			✓	74.6	5.96	1.30
		✓		73.7	5.72	1.30
	✓			73.0	5.58	1.30
	✓	✓		74.5	5.59	1.35
SWAT <sub>DeiT</sub>	✓	✓	✓	<b>75.7</b>	5.83	1.40
Mixer				68.3	5.07	0.97
			✓	70.8	5.28	1.01
		✓		68.9	5.08	0.97
	✓			67.9	4.88	0.94
	✓	✓		70.2	4.88	0.95
SWAT <sub>Mixer</sub>	✓	✓	✓	<b>72.3</b>	5.10	0.99

Table 3: **Ablations on Structure with DeiT-Ti [Touvron *et al.*, 2021b] and Mixer-Ti [Tolstikhin *et al.*, 2021] on ImageNet-1K.** We report the gains from (1) Structure-aware Tokenization, (2) Token Mixing with Channel Structure, and (3) Channel Mixing with Token Structure. Structure-aware inputs and Structure-aware Mixing gives consistent improvements, as shown in **bold**. One key observation shows that, our Tokenization and our Token Mixing should *always* be coupled together.

fine-grained structures (eg: boundaries in object segments), since we preserve such structure within tokens. In contrast, DeiT attention is smoothed-out and subtle. Also, the attention weights in the SWAT model are less-noisy. It is worth noting that we use the same resolution (i.e., same number of tokens) in both cases.

We include a detailed analysis of model throughput (im/s) in SWAT models and their baselines at inference, in the Appendix. We consider FLOPs and parameters as our metrics of complexity, as they are system-agnostic and reproducible.

### 4.2 Ablations on ImageNet

In this section, We present ablations on Tiny versions of SWAT<sub>DeiT</sub> and SWAT<sub>Mixer</sub>. Specifically, in Table 3, we focus on Structure-aware Tokenization, Token Mixing with Channel Structure and Channel Mixing with Token Structure.

**Structure-aware Tokenization:** We compare different settings with SWAT tokenizer. Bottom line is that *Structure-aware Tokenization should always be coupled with the Structure-aware Token Mixing*. It makes sense: if we prepare tokens with structure and not take advantage of it during mixing, it does not really have a benefit and the reduced capacity (due to our Tokenization) may even drop the performance. In DeiT [Touvron *et al.*, 2021b], we see such performance drop of -0.3% when we do not use the structure (*within tokens*) explicitly. In Mixer [Tolstikhin *et al.*, 2021], this drop is -0.4%. In both cases, when we specifically make use of the newly-introduced structure *within tokens*, we see consistent gains (+1.5% in DeiT and +2.3% in Mixer) over our Tokenization-only versions.

**Channel Structure (*within tokens*) in Token Mixing:** Here, we intend to consider a local neighborhood within tokens. Even if such a structure is not present (i.e., not having our Tokenization), models can benefit slightly: +0.4% in DeiT

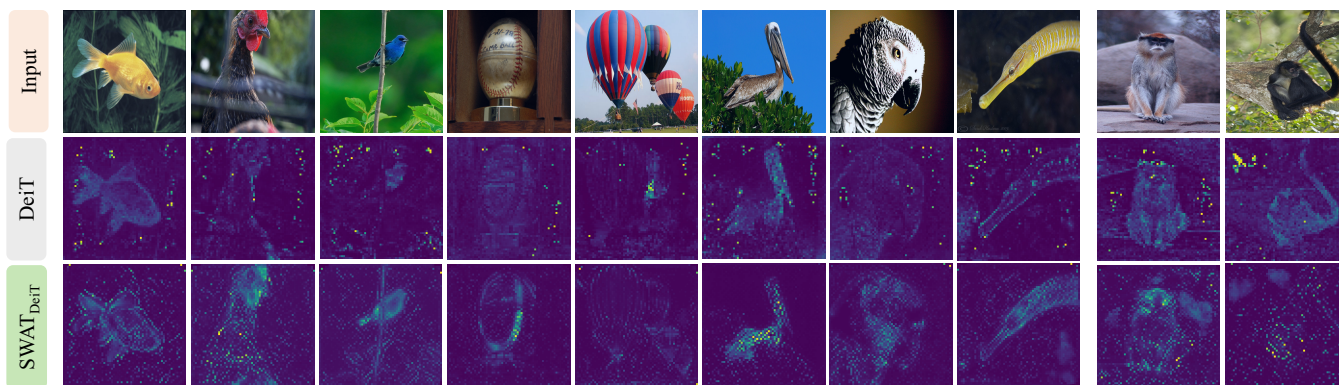


Figure 5: **Visualization of token attention** in DeiT-Ti [Touvron *et al.*, 2021b] and SWAT<sub>DeiT</sub>-Ti. We use the code from DINO [Caron *et al.*, 2021] for visualization. However since we do not use class tokens as in DINO, we show the attention averaged across tokens. SWAT, as it preserves structure even within tokens, shows more contrastive and fine-grained attention maps compared to DeiT (even though we consider same number of tokens, i.e., resolution, in both). Note the better-visible boundaries and segments in SWAT, compared to smoothed-out variations in DeiT. Some cases where SWAT fails to capture fine details are also shown (to the right). Best viewed in color and zoomed-in.

[Touvron *et al.*, 2021b] and +0.6% in Mixer [Tolstikhin *et al.*, 2021]. This is due to the inductive bias of replacing Linear layers with Conv. However, the true potential of this comes when a structure *within* tokens is explicitly available, where we see a +1.2% improvement in DeiT [Touvron *et al.*, 2021b] and +1.9% in Mixer [Tolstikhin *et al.*, 2021].

#### Token Structure (among tokens) in Channel Mixing:

Here, we consider a local neighborhood among tokens. In DeiT [Touvron *et al.*, 2021b], we see +1.3% boost, and in Mixer [Tolstikhin *et al.*, 2021], a +2.5% boost in performance.

### 4.3 Semantic Segmentation

ADE20K [Zhou *et al.*, 2019] benchmark contains annotations for semantic segmentation across 150 categories. It comes with 25K annotated images in total, with 20K training, 2K validation and 3K testing. We report mIoU for our models in multi-scale testing (i.e., [0.5, 0.75, 1.0, 1.25, 1.5, 1.75]× the training resolution) similar to previous work [Liu *et al.*, 2021], along with complexity metrics such as parameters, FLOPs (for input size of  $512 \times 2048$  similar to [Liu *et al.*, 2021]) and frame-rate. We follow a similar training recipe to Swin [Liu *et al.*, 2021]. Our backbones are pretrained on ImageNet-1K [Deng *et al.*, 2009] for 300 epochs at  $224 \times 224$ , before re-training with a decoder for segmentation at  $512 \times 512$ . We use UperNet [Xiao *et al.*, 2018] as our decoder within `mmsegmentation` [OpenMMLab, 2020] framework. We use the original hyperparameter settings as the baseline.

**Results:** In Table 4, we show the performance of SWAT<sub>Swin</sub> and SWAT<sub>VAN</sub> backbones when used with the UperNet [Xiao *et al.*, 2018] head for semantic segmentation on ADE20K, and compare with similar-sized baselines. SWAT<sub>Swin</sub> gives +0.7 mIoU and SWAT<sub>VAN</sub> gives +0.6 mIoU improvement over the respective baselines, when trained under the same settings. However, FPS of the SWAT<sub>Swin</sub> based model is slightly lower, due to extra convolutions introduced in SWAT. In the Appendix, we include segmentation masks generated by SWAT<sub>Swin</sub> and Swin [Liu *et al.*, 2021] backbones, which qualitatively show this improvement.

Method	Backbone	mIoU	Params. (M)	FLOPs (G)	FPS
DANet (Fu <i>et al.</i> )	Resnet-101 (He <i>et al.</i> )	45.2	69	1119	15.2
DLab.v3+ (Chen <i>et al.</i> )		44.1	63	1021	16.0
ACNet (Fu <i>et al.</i> )		45.9	-	-	-
DNL (Yin <i>et al.</i> )		46.0	69	1249	14.8
OCRNet (Yuan <i>et al.</i> )		45.3	56	923	19.3
UperNet (Xiao <i>et al.</i> )		44.9	86	1029	20.1
UperNet (Xiao <i>et al.</i> )	DeiT-S (Touvron <i>et al.</i> )	44.0	52	1099	16.2
	Swin-Ti (Liu <i>et al.</i> )	45.8	60	945	18.5
	SWAT <sub>Swin</sub> -Ti (ours)	<b>46.5</b>	59	950	16.9
	VAN-B2 (Guo <i>et al.</i> )	50.1	57	948	-
	SWAT <sub>VAN</sub> -B2 (ours)	<b>50.7</b>	55	952	-

Table 4: **SWAT for semantic segmentation** on ADE20K [Zhou *et al.*, 2019] dataset. We report results in the same setting as Swin [Liu *et al.*, 2021] using `mmsegmentation` [OpenMMLab, 2020] framework. FPS is measured on a single V100 GPU. SWAT outperforms respective baselines, but is slightly slower due to extra convolutions.

## 5 Conclusion

In this work, we present the merits of preserving spatial structure, both *within* and *among* Tokens, in common Transformer/Mixer/Convolutional token-based architectures. Our two key contributions are: (1) *Structure-aware Tokenization* and (2) *Structure-aware Mixing*, which can be adopted in different families of models with minimal effort. The resulting family of models, SWAT, outperforms the corresponding baselines and shows competitive performance with SOTA models on multiple benchmarks, with minimal change in parameters and computations. We hope that SWAT will open-up new ways of making use of spatial structure as an inductive bias in token-based models.

## Acknowledgements

This work was supported by the National Science Foundation (IIS-2104404 and CNS-2104416). We thank the Robotics Lab at SBU for helpful discussions.

## References

- [Arnab *et al.*, 2021] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. ViViT: A Video Vision Transformer. In *ICCV*, October 2021.
- [Bello, 2020] Irwan Bello. LambdaNetworks: Modeling long-range Interactions without Attention. In *ICLR*, 2020.
- [Cao *et al.*, 2021] Jiezhong Cao, Yawei Li, Kai Zhang, and Luc Van Gool. Video Super-Resolution Transformer. *arXiv:2106.06847*, 2021.
- [Carion *et al.*, 2020] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-End Object Detection with Transformers. In *ECCV*. Springer, 2020.
- [Caron *et al.*, 2021] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging Properties in Self-Supervised Vision Transformers. In *ICCV*, October 2021.
- [Chen *et al.*, 2018] Liang-Chieh Chen, Yukun Zhu, George Papandreou, Florian Schroff, and Hartwig Adam. Encoder-Decoder with Atrous Separable Convolution for Semantic Image Segmentation. In *ECCV*, 2018.
- [Chen *et al.*, 2021a] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision Transformer: Reinforcement Learning via Sequence Modeling. *URL Workshop in ICML*, 2021.
- [Chen *et al.*, 2021b] Shoufa Chen, Enze Xie, Chongjian Ge, Ding Liang, and Ping Luo. CycleMLP: A MLP-like Architecture for Dense Prediction. *arXiv:2107.10224*, 2021.
- [Chen *et al.*, 2022] Yinpeng Chen, Xiyang Dai, Dongdong Chen, Mengchen Liu, Xiaoyi Dong, Lu Yuan, and Zicheng Liu. Mobileformer: Bridging mobilenet and transformer. In *CVPR*, 2022.
- [Chu *et al.*, 2021] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. *NeurIPS*, 34, 2021.
- [Dai *et al.*, 2021] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. UP-DETR: Unsupervised Pre-training for Object Detection with Transformers. In *CVPR*, 2021.
- [Dai *et al.*, 2022] Rui Dai, Srijan Das, Kumara Kahatapitiya, Michael S Ryoo, and Francois Bremond. MS-TCT: Multi-Scale Temporal ConvTransformer for Action Detection. In *CVPR*, 2022.
- [Dauphin *et al.*, 2017] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. Language Modeling with Gated Convolutional Networks. In *ICML*. PMLR, 2017.
- [Deng *et al.*, 2009] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*. IEEE, 2009.
- [Devlin *et al.*, 2019] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. BERT: Pre-training of deep bidirectional transformers for language understanding. Association for Computational Linguistics, June 2019.
- [Dosovitskiy *et al.*, 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. 2021.
- [Duke *et al.*, 2021] Brendan Duke, Abdalla Ahmed, Christian Wolf, Parham Aarabi, and Graham W Taylor. SSTVOS: Sparse Spatiotemporal Transformers for Video Object Segmentation. In *CVPR*, 2021.
- [d’Ascoli *et al.*, 2021] Stéphane d’Ascoli, Hugo Touvron, Matthew L Leavitt, Ari S Morcos, Giulio Biroli, and Levent Sagun. ConViT: Improving Vision Transformers with Soft Convolutional Inductive Biases. In *ICML*. PMLR, 2021.
- [Esser *et al.*, 2021] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming Transformers for High-Resolution Image Synthesis. In *CVPR*, 2021.
- [Fan *et al.*, 2021] Haoqi Fan, Bo Xiong, Karttikeya Mangalam, Yanghao Li, Zhicheng Yan, Jitendra Malik, and Christoph Feichtenhofer. Multiscale Vision Transformers. In *ICCV*, October 2021.
- [Fu *et al.*, 2019a] Jun Fu, Jing Liu, Haijie Tian, Yong Li, Yongjun Bao, Zhiwei Fang, and Hanqing Lu. Dual Attention Network for Scene Segmentation. In *CVPR*, 2019.
- [Fu *et al.*, 2019b] Jun Fu, Jing Liu, Yuhang Wang, Yong Li, Yongjun Bao, Jinhui Tang, and Hanqing Lu. Adaptive Context Network for Scene Parsing. In *ICCV*, 2019.
- [Graham *et al.*, 2021] Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Hervé Jégou, and Matthijs Douze. LeViT: A Vision Transformer in ConvNet’s Clothing for Faster Inference. In *ICCV*, October 2021.
- [Guo *et al.*, 2021] Meng-Hao Guo, Jun-Xiong Cai, Zheng-Ning Liu, Tai-Jiang Mu, Ralph R Martin, and Shi-Min Hu. PCT: Point Cloud Transformer. *Computational Visual Media*, 7(2), 2021.
- [Guo *et al.*, 2022] Meng-Hao Guo, Cheng-Ze Lu, Zheng-Ning Liu, Ming-Ming Cheng, and Shi-Min Hu. Visual Attention Network. *arXiv:2202.09741*, 2022.
- [Han *et al.*, 2021] Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in Transformer. In *NeurIPS*, 2021.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *CVPR*, 2016.
- [Heo *et al.*, 2021] Byeongho Heo, Sangdoon Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking Spatial Dimensions of Vision Transformers. In *ICCV*, October 2021.
- [LeCun *et al.*, 2015] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep Learning. *Nature*, 521(7553), 2015.
- [Liu *et al.*, 2021] 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*, October 2021.
- [Liu *et al.*, 2022] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A ConvNet for the 2020s. In *CVPR*, 2022.
- [Nagrani *et al.*, 2021] Arsha Nagrani, Shan Yang, Anurag Arnab, Aren Jansen, Cordelia Schmid, and Chen Sun. Attention Bottlenecks for Multimodal Fusion. *NeurIPS*, 34, 2021.
- [Naseer *et al.*, 2021] Muzammal Naseer, Kanchana Ranasinghe, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Intriguing Properties of Vision Transformers. In *NeurIPS*, 2021.



- [OpenMMLab, 2020] OpenMMLab. MMSegmentation: OpenMMLab Semantic Segmentation Toolbox and Benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. Accessed: 2023-01-18.
- [Peng *et al.*, 2021] 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*, October 2021.
- [Radosavovic *et al.*, 2020] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *CVPR*, 2020.
- [Ranasinghe *et al.*, 2022] Kanchana Ranasinghe, Muzammal Naseer, Salman Khan, Fahad Shahbaz Khan, and Michael S Ryoo. Self-supervised Video Transformer. In *CVPR*, 2022.
- [Ranftl *et al.*, 2021] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision Transformers for Dense Prediction. In *ICCV*, 2021.
- [Ryoo *et al.*, 2021] Michael S Ryoo, AJ Piergiovanni, Anurag Arnab, Mostafa Dehghani, and Anelia Angelova. TokenLearner: Adaptive Space-Time Tokenization for Videos. In *NeurIPS*, 2021.
- [Shang *et al.*, 2022] Jinghuan Shang, Kumara Kahatapitiya, Xiang Li, and Michael S Ryoo. StARformer: Transformer with State-Action-Reward Representations. In *ECCV*, 2022.
- [Tan and Le, 2019] Mingxing Tan and Quoc Le. EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks. In *ICML*. PMLR, 2019.
- [Tang *et al.*, 2021] Chuanxin Tang, Yucheng Zhao, Guangting Wang, Chong Luo, Wenxuan Xie, and Wenjun Zeng. Sparse MLP for Image Recognition: Is Self-Attention Really Necessary? *arXiv:2109.05422*, 2021.
- [Tolstikhin *et al.*, 2021] Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Peter Steiner, Daniel Keysers, Jakob Uszkoreit, et al. MLP-Mixer: An All-MLP Architecture for Vision. In *NeurIPS*, 2021.
- [Touvron *et al.*, 2021a] Hugo Touvron, Piotr Bojanowski, Mathilde Caron, Matthieu Cord, Alaaeldin El-Nouby, Edouard Grave, Gauthier Izacard, Armand Joulin, Gabriel Synnaeve, Jakob Verbeek, et al. ResMLP: Feedforward networks for image classification with data-efficient training. *arXiv:2105.03404*, 2021.
- [Touvron *et al.*, 2021b] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*. PMLR, 2021.
- [Trockman and Kolter, 2022] Asher Trockman and J Zico Kolter. Patches are all you need? *arXiv:2201.09792*, 2022.
- [Vaswani *et al.*, 2017] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention Is All You Need. In *NeurIPS*, 2017.
- [Wang *et al.*, 2018] Xiaolong Wang, Ross Girshick, Abhinav Gupta, and Kaiming He. Non-local Neural Networks. In *CVPR*, 2018.
- [Wang *et al.*, 2021] 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*, October 2021.
- [Wightman, 2019] Ross Wightman. Pytorch Image Models. <https://github.com/rwightman/pytorch-image-models>, 2019. Accessed: 2023-01-18.
- [Wu *et al.*, 2021] Haiping Wu, Bin Xiao, Noel Codella, Mengchen Liu, Xiyang Dai, Lu Yuan, and Lei Zhang. CvT: Introducing Convolutions to Vision Transformers. In *ICCV*, October 2021.
- [Xiao *et al.*, 2018] Tete Xiao, Yingcheng Liu, Bolei Zhou, Yuning Jiang, and Jian Sun. Unified Perceptual Parsing for Scene Understanding. In *ECCV*, 2018.
- [Xiao *et al.*, 2021] Tete Xiao, Piotr Dollar, Mannat Singh, Eric Mintun, Trevor Darrell, and Ross Girshick. Early Convolutions Help Transformers See Better. In *NeurIPS*, 2021.
- [Xie *et al.*, 2017] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *CVPR*, 2017.
- [Xie *et al.*, 2021] Enze Xie, Wenhai Wang, Zhiding Yu, Anima Anandkumar, Jose M Alvarez, and Ping Luo. SegFormer: Simple and Efficient Design for Semantic Segmentation with Transformers. In *NeurIPS*, 2021.
- [Yang *et al.*, 2021a] Guanglei Yang, Hao Tang, Mingli Ding, Nicu Sebe, and Elisa Ricci. Transformer-Based Attention Networks for Continuous Pixel-Wise Prediction. In *ICCV*, 2021.
- [Yang *et al.*, 2021b] 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. In *NeurIPS*, 2021.
- [Yin *et al.*, 2020] Minghao Yin, Zhuliang Yao, Yue Cao, Xiu Li, Zheng Zhang, Stephen Lin, and Han Hu. Disentangled Non-Local Neural Networks. In *ECCV*. Springer, 2020.
- [Yu *et al.*, 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. In *CVPR*, 2022.
- [Yuan *et al.*, 2021a] Kun Yuan, Shaopeng Guo, Ziwei Liu, Aojun Zhou, Fengwei Yu, and Wei Wu. Incorporating Convolution Designs Into Visual Transformers. In *ICCV*, October 2021.
- [Yuan *et al.*, 2021b] Li Yuan, Yunpeng Chen, Tao Wang, Weihao Yu, Yujun Shi, Zi-Hang Jiang, Francis E.H. Tay, Jiashi Feng, and Shuicheng Yan. Tokens-to-Token ViT: Training Vision Transformers From Scratch on ImageNet. In *ICCV*, October 2021.
- [Yue *et al.*, 2021] Xiaoyu Yue, Shuyang Sun, Zhanghui Kuang, Meng Wei, Philip H.S. Torr, Wayne Zhang, and Dahua Lin. Vision Transformer With Progressive Sampling. In *ICCV*, October 2021.
- [Zhai *et al.*, 2021] Shuangfei Zhai, Walter Talbott, Nitish Srivastava, Chen Huang, Hanlin Goh, Ruixiang Zhang, and Josh Susskind. An Attention Free Transformer. *arXiv:2105.14103*, 2021.
- [Zhang *et al.*, 2021] Zizhao Zhang, Han Zhang, Long Zhao, Ting Chen, and Tomas Pfister. Aggregating Nested Transformers. *arXiv:2105.12723*, 2021.
- [Zhao *et al.*, 2020] Hengshuang Zhao, Jiaya Jia, and Vladlen Koltun. Exploring Self-attention for Image Recognition. In *CVPR*, 2020.
- [Zhao *et al.*, 2021] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point Transformer. In *ICCV*, 2021.
- [Zhou *et al.*, 2019] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic Understanding of Scenes Through the ADE20K Dataset. *IJCV*, 127(3), 2019.
- [Zhu *et al.*, 2020] Xizhou Zhu, Weijie Su, Lewei Lu, Bin Li, Xiaogang Wang, and Jifeng Dai. Deformable DETR: Deformable Transformers for End-to-End Object Detection. In *ICLR*, 2020.