

# Token-Label Alignment for Vision Transformers

Han Xiao<sup>1,2,\*</sup> Wenzhao Zheng<sup>1,2,\*</sup> Zheng Zhu<sup>3</sup> Jie Zhou<sup>1,2</sup> Jiwen Lu<sup>1,2,†</sup>

<sup>1</sup>Beijing National Research Center for Information Science and Technology, China

<sup>2</sup>Department of Automation, Tsinghua University, China <sup>3</sup>PhiGent Robotics

{h-xiao20, zhengwz18}@mails.tsinghua.edu.cn; zhengzhu@ieee.org;

{jzhou, lujiwen}@tsinghua.edu.cn

## Abstract

*Data mixing strategies (e.g., CutMix) have shown the ability to greatly improve the performance of convolutional neural networks (CNNs). They mix two images as inputs for training and assign them with a mixed label with the same ratio. While they are shown effective for vision transformers (ViTs), we identify a token fluctuation phenomenon that has suppressed the potential of data mixing strategies. We empirically observe that the contributions of input tokens fluctuate as forward propagating, which might induce a different mixing ratio in the output tokens. The training target computed by the original data mixing strategy can thus be inaccurate, resulting in less effective training. To address this, we propose a token-label alignment (TL-Align) method to trace the correspondence between transformed tokens and the original tokens to maintain a label for each token. We reuse the computed attention at each layer for efficient token-label alignment, introducing only negligible additional training costs. Extensive experiments demonstrate that our method improves the performance of ViTs on image classification, semantic segmentation, objective detection, and transfer learning tasks. Code is available at: <https://github.com/Euphoria16/TL-Align>.*

## 1. Introduction

The recent developments of vision transformers (ViTs) have revolutionized the computer vision field and set new state-of-the-arts in various tasks, such as image classification [16, 43, 31, 10], object detection [5, 56, 12, 13], and semantic segmentation [28, 39, 54, 9]. The successful structure of alternative spatial mixing and channel mixing in ViTs also motivates the arising of high-performance MLP-like deep architectures [40, 41, 42] and promotes the evolution of better CNNs [15, 32, 17]. In addition to architec-

ture designs, an improved training strategy can also greatly boost the performance of a trained deep model [23, 44, 8, 7].

The training of modern deep architecture almost all adopts data mixing strategies for data augmentation [48, 45, 24, 47, 52, 53], which have been proven to consistently improve the generalization performance. They randomly mix two images and their labels with the same mixing ratio to produce mixed data. As the most commonly used data mixing strategy, CutMix [52] performs a copy-and-paste operation on the spatial domain to produce spatially mixed images. While data mixing strategies have been widely studied for CNNs [48, 45, 24], few works have explored their compatibilities with ViTs [7]. We find that self-attention in ViTs causes a fluctuation of the original spatial structure. Unlike the translation equivalence that ensures a global label consistency for CNNs, self-attention in ViTs undermines this global consistency and causes a misalignment between the token and label. This misalignment induces a different mixing ratio in the output tokens. The training targets computed by the original data mixing strategies can then be inaccurate, resulting in less effective training.

To address this, we propose a token-label alignment (TL-Align) method for ViTs to obtain a more accurate target for training. We present an overview of our method in Figure 1. We first assign a label to each input token in the mixed image according to the source of the token. We then trace the correspondence between the input tokens and the transformed tokens and align the labels accordingly. We assume that only the spatial self-attention and residual connection operation alter the presence of input tokens since channel MLP and layer normalization process each token independently. We reuse the computed attentions to linearly mix the labels of input tokens to obtain those of transformed tokens. The token-label alignment is performed iteratively to obtain a label for each output token. For class-token-based classification (e.g., ViT [16] and DeiT [43]), we directly use the aligned label for the output class token as the training target. For global-pooling-based classification (e.g., Swin [31]), we similarly average the labels of output tokens as the train-

\*Equal contribution.

†Corresponding author.

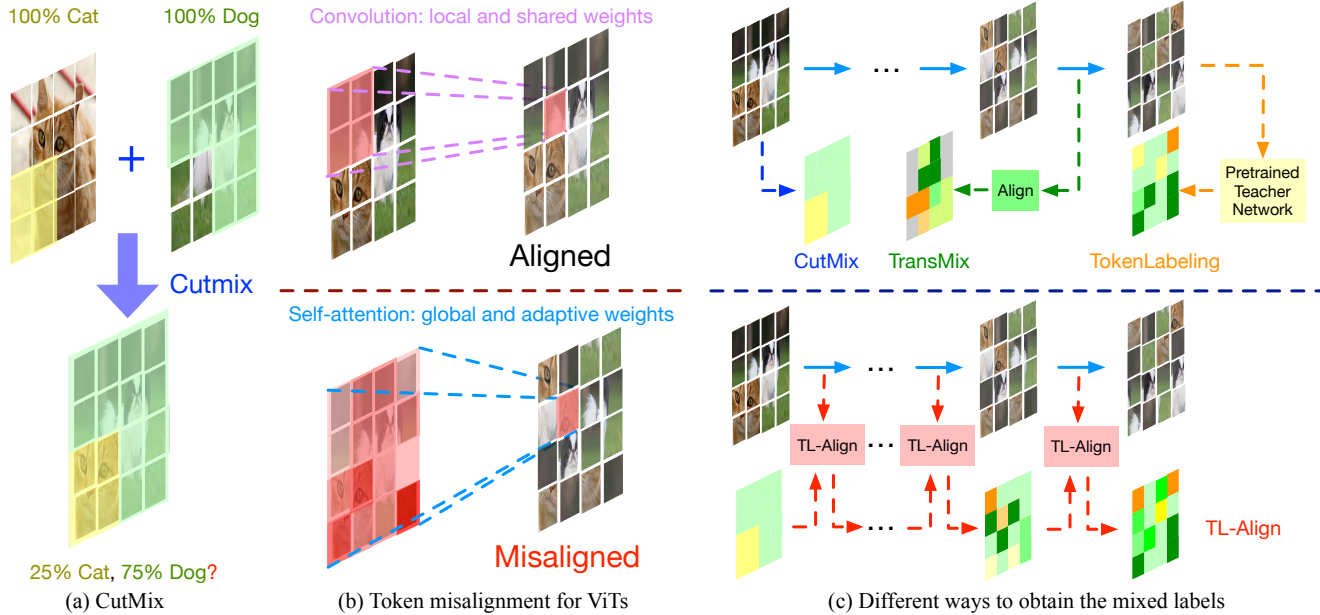


Figure 1. An overview of the proposed TL-Align. (a) CutMix-like methods [52] are widely used in model training, which spatially mix the tokens and their labels in the input space. (b) They are originally designed for CNNs and assume the processed tokens are spatially aligned with the input tokens. We show that it does not hold true for ViTs due to the global receptive field and adaptive weights. (c) Compared with existing methods, our method can effectively and efficiently align the tokens and labels without requiring a pretrained teacher network.

ing target. The proposed TL-Align is only used for training to improve performance and introduces no additional workload for inference. We apply the proposed TL-Align to various ViT variants with CutMix including plain ViTs (DeiT [43]) and hierarchical ViTs (Swin [31]). We observe a consistent performance boost across different models on ImageNet-1K [14]. Specifically, our TL-Align improves DeiT-S by 0.8% using the same training recipe. We evaluate the ImageNet-pretrained models on various downstream tasks including semantic segmentation, objection detection, and transfer learning. Experimental results also verify the robustness and generalization ability of our method.

## 2. Related Work

**Vision Transformer.** Transformers have been widely used in natural language processing and achieved great success on many language tasks. Recently, Vision Transformers (ViTs) have aroused extensive interest in computer vision due to their competitive performance compared with CNNs [16, 43, 31, 10]. Dosovitskiy et al. [16] firstly introduced transformers into the image classification task. They split the input image into non-overlapped patches and then feed them into the transformer encoders. Liu et al. [31] proposed a shifted windowing scheme to produce hierarchical feature maps suitable for dense prediction tasks. The great potential of vision transformer has motivated its adaptation to many challenging tasks including object detection [12, 56, 5], segmentation [9, 39], image enhancement [6, 27] and video understanding [33, 1].

Recently, some efforts have been devoted to producing better training targets to improve the performance of vision transformers [23, 44]. For example, DeiT [43] introduces a knowledge distillation procedure to reduce the training cost of ViTs and achieves a better accuracy/speed trade-off. TokenLabeling [23] employs a pretrained teacher annotator to predict a label for each token for dense knowledge distillation. Differently, we do not require a pretrained network to obtain the training targets. Our TL-Align maintains an aligned label for each token layer by layer and can be trained efficiently in an end-to-end manner.

**Data Mixing Strategy.** As an important type of data augmentation, data mixing strategies have demonstrated a consistent improvement in the generalization performance of CNNs. Zhang et al. [53] first proposed to combine a training pair to create augmented samples for model regularization. They perform linear interpolations on both the input images and associated targets. Following MixUp, CutMix [52] also utilizes the mixture of two input images but adopts a region copy-and-paste operation. Later methods including Puzzle Mix [24], SaliencyMix [45] and Attentive CutMix [48] leverage the salient regions for informative mixture generation. AlignMixUp [46] aligns two images geometrically in the feature space to improve the mixing representations. Recently, Yang et al. [51] proposed a RecursiveMix strategy which employs the historical input-prediction-label triplets for scale-invariant feature learning. Despite the better performance, a drawback of these methods is the heavily increased training cost due to the saliency extraction or historical information exploitation.

Most existing data mixing methods are originally designed for CNNs, and their effectiveness on ViTs has not been well explored. TransMix [7] uses the class attention map at the last layer to re-weight the mixing targets and assumes the output tokens to keep spatial correspondence with the input tokens. However, we identify a token fluctuation phenomenon for ViTs which may cause a mismatch between tokens and labels, leading to inaccurate label assignments in both the original CutMix and TransMix. To address this, we propose to align the label and token space by tracing their correspondence in a layerwise manner.

### 3. Proposed Approach

#### 3.1. Preliminaries

The convolution neural network (CNN) has been the dominant architecture for computer vision in the deep learning era, greatly improving the performance of many tasks. Its monopoly has been challenged by the recent emergence of vision transformers (ViTs), which first “patchify” each image into tokens and process them with alternating self-attention (SA) and multi-layer perceptron (MLP).

In addition to architecture design, training strategy also has a large effect on the model performance, especially the data augmentation strategy. Data mixing [48, 45, 24, 47, 52, 53] is an important set of data augmentation for the training of both CNNs and ViTs, as it significantly improves the generalization ability of models. As the most commonly used data mixing strategy, CutMix [52] aims to create virtual training samples from the given training samples  $(\mathbf{X}, y)$ , where  $\mathbf{X} \in \mathcal{R}^{H \times W \times C}$  denotes the input image and  $y$  is the corresponding label. CutMix randomly selects a local region from one input  $\mathbf{X}_1$  and uses it to replace the pixels in the same region of another input  $\mathbf{X}_2$  to generate a new sample  $\tilde{\mathbf{X}}$ . Similarly, the label  $\tilde{y}$  of  $\tilde{\mathbf{X}}$  is also the combination of the original labels  $y_1$  and  $y_2$ :

$$\begin{aligned} \tilde{\mathbf{X}} &= \mathbf{M} \odot \mathbf{X}_1 + (\mathbf{1} - \mathbf{M}) \odot \mathbf{X}_2 \\ \tilde{y} &= \lambda y_1 + (1 - \lambda) y_2 \end{aligned} \quad (1)$$

where  $\mathbf{M} \in \{0, 1\}^{H \times W}$  is a binary mask indicating the image each pixel belongs to,  $\mathbf{1}$  is an all-one matrix, and  $\odot$  is the element-wise multiplication.  $\lambda$  reflects the mixing ratio of two labels and is the proportion of pixels cropped from  $\mathbf{X}_1$  in the mixed image  $\tilde{\mathbf{X}}$ . For a cropped region  $[r_x, r_x + r_w] \times [r_y, r_y + r_h]$  from  $\mathbf{X}_1$ , we compute  $\lambda = \frac{r_w r_h}{WH}$  to obtain the initial mixed target  $\tilde{y}$ .

#### 3.2. The Token Fluctuation Phenomenon

CutMix is originally designed for CNNs and assumes the feature extraction process does not alter the mixing ratio. However, we discover that different from CNNs, self-attention in ViTs can lead to the fluctuation of some tokens.

The fluctuation further results in the mismatch between the token space and label space, which hinders the effective training of the network.

Formally, we use  $\mathbf{z}_i$  to denote a token of the image  $\mathbf{Z}$ , i.e.,  $\mathbf{z}_i$  is the transposed  $i$ -th column vector of  $\mathbf{Z}$ . We can then compute the  $i$ -th transformed token  $\hat{\mathbf{z}}_i$  after the spatial operation as  $\hat{\mathbf{z}}_i = \sum_{j=1}^N w_{i,j}^s \mathbf{z}_j$ , where  $w_{i,j}^s$  is the  $i, j$ -th element of the computed spatial mixing matrix  $\mathbf{w}^s(\mathbf{z})$ .

With the assumption of the linear information integration, we define the contribution of an original token  $\mathbf{z}_i$  to a mixed token  $\hat{\mathbf{z}}_j$  as  $c(\mathbf{z}_i, \hat{\mathbf{z}}_j) = \frac{|w_{i,j}^s|}{\sum_{k=1}^N |w_{k,j}^s|}$ , where  $|\cdot|$  denotes the absolute value. We can then compute the presence of a token  $\mathbf{z}_i$  in all the mixed image tokens as:

$$p(\mathbf{z}_i) = \sum_{j=1}^N c(\mathbf{z}_i, \hat{\mathbf{z}}_j) = \sum_{j=1}^N \frac{|w_{i,j}^s|}{\sum_{k=1}^N |w_{k,j}^s|}. \quad (2)$$

For non-strided depth-wise convolution, each token is multiplied by each element in the convolutional kernel due to the translation invariance. We can thus obtain:

$$\sum_{l=1}^N |w_{i,l}^s| = \sum_{j=1}^N |w_{k,j}^s| = \sum_{k=1, l=1}^M |K_{k,l}|, \quad \forall i, j \in \mathbf{P}_{NE}, \quad (3)$$

where  $\mathbf{P}_{NE}$  denotes the set of positions that are not at the edge of the image,  $K_{k,l}$  denotes the value of the  $k, l$ -th position of the convolution kernel  $\mathbf{K}$  and  $M$  is the kernel size. We can infer that  $p(\mathbf{z}_i) = 1, \forall i \in \mathbf{P}_{NE}$ , i.e., the effect of all the internal tokens does not change during the convolution process. However, for self-attention in ViTs, (3) does not hold due to the non-existence of translation invariance. The fluctuation of  $p(\mathbf{z})$  is further amplified by the input dependency of the spatial mixing matrix  $\mathbf{w}^s(\mathbf{z})$  induced by self-attention. As an extreme case, we may obtain  $p(\mathbf{z}) \sim 0$  for certain tokens. The fluctuation of tokens will alter the proportion of mixing (i.e.,  $\lambda$ ) and the network might even completely ignore one of the mixed images. The actual label of the processed tokens can then deviate from the mixed label computed by (1), resulting in less effective training.

#### 3.3. Token-Label Alignment

Each token in ViTs interacts with other tokens using the self-attention mechanism. The input-dependent weights empower ViTs with more flexibility but also result in a mismatch between the processed token and the initial token. To address this, we propose a token-label alignment (TL-Align) method to trace the correspondence between the input and transformed tokens to obtain the aligned labels for the resulting representations, as illustrated in Figure 2.

Specifically, ViTs first split the mixed input  $\tilde{\mathbf{X}}$  after CutMix (1) to a sequence of  $N$  non-overlapped patches and then flatten them to obtain the original image tokens

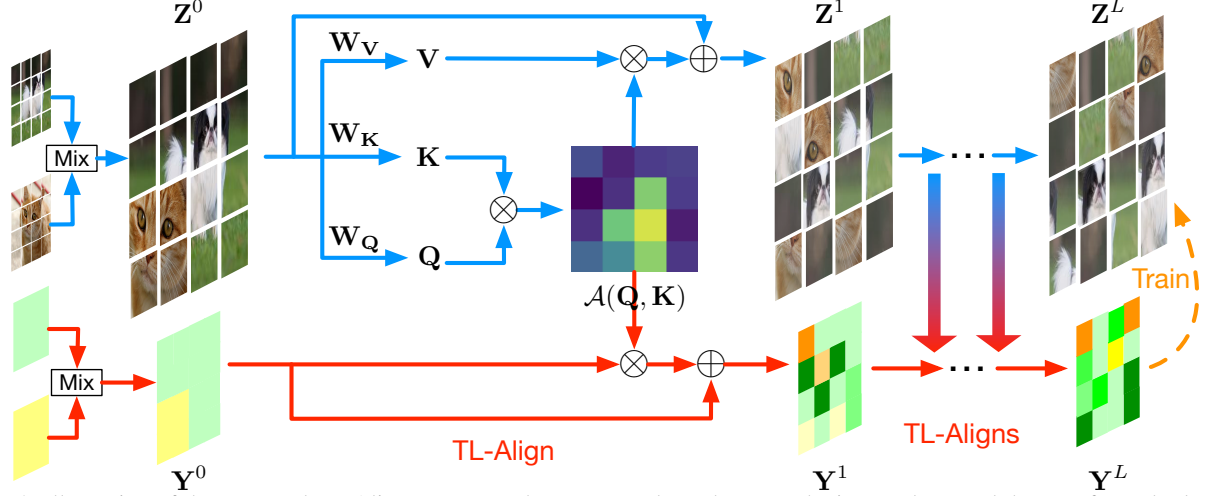


Figure 2. Illustration of the proposed TL-Align. We trace the correspondence between the input tokens and the transformed tokens and align the labels accordingly. We reuse the computed attentions to linearly mix the labels of input tokens to obtain those of transformed tokens. The token-label alignment is performed iteratively to obtain a label for each output token.

$\{\tilde{\mathbf{x}}_1, \tilde{\mathbf{x}}_2, \dots, \tilde{\mathbf{x}}_N\}$ . We then project them into a proper dimension and add positional embeddings:

$$\mathbf{Z}^0 = [\tilde{\mathbf{z}}_{cls}; \tilde{\mathbf{x}}_1 \cdot \mathbf{E}; \tilde{\mathbf{x}}_2 \cdot \mathbf{E}; \dots; \tilde{\mathbf{x}}_N \cdot \mathbf{E}] + \mathbf{E}_{pos}, \quad (4)$$

where  $\tilde{\mathbf{z}}_{cls} \in \mathcal{R}^{1 \times d}$  denotes the class token,  $N$  is the number of tokens,  $\mathbf{E}$  represents the patch projector, and  $\mathbf{E}_{pos} \in \mathcal{R}^{(N+1) \times d}$  is the position embeddings. Note that we adopt the process of the original transformer architecture [16] as an example without loss of generality. Other models may omit the class token and use a relative positional embedding instead, which does not affect the utility of the proposed TL-Align method.

We first assign each token  $\mathbf{z}_i \in \mathcal{R}^{1 \times d}$  with a label embedding  $\mathbf{y}_i \in \mathcal{R}^{1 \times C}$ :

$$\mathbf{Y}^0 = [\tilde{\mathbf{y}}_{cls}^0; \tilde{\mathbf{y}}_1^0; \tilde{\mathbf{y}}_2^0; \dots; \tilde{\mathbf{y}}_N^0], \quad (5)$$

where the sum of elements in each  $\mathbf{y}_i$  equals 1 (i.e.,  $\sum_{j=1}^C y_{i,j} = 1$ ) and  $y_{i,j}$  indicates how much the  $i$ -th token belong to the  $j$ -th class. We initialize the label embedding following the conventional data mixing paradigm. For example, when using CutMix to mix two images  $\mathbf{X}_1$  and  $\mathbf{X}_2$  from the  $j$ -th class and the  $k$ -th class with a mixing ratio of  $\lambda$ , we set  $\tilde{y}_{cls,j} = \lambda$  and  $\tilde{y}_{cls,k} = 1 - \lambda$  for the class token. For each patch token, we set  $\tilde{y}_{i,j} = 1$  if it comes from  $\mathbf{X}_1$  and  $\tilde{y}_{i,k} = 1$  if it comes from  $\mathbf{X}_2$ . If a patch token contains both the mixed images, we use the mixing ratio within this patch as the label. For MixUp, we can simply set all label embeddings  $\{\tilde{\mathbf{y}}_i\}$  with  $\tilde{y}_{i,j} = \lambda$  and  $\tilde{y}_{i,k} = 1 - \lambda$ .

We perform TL-Align in a layer-wise manner and compute the aligned labels based on the operation on the tokens. Formally, ViTs use self-attention to perform spatial mixing

of the input tokens  $\mathbf{Z}$ :

$$\begin{aligned} \mathbf{Q} &= \mathbf{Z} \cdot \mathbf{W}_Q, \mathbf{K} = \mathbf{Z} \cdot \mathbf{W}_K, \mathbf{V} = \mathbf{Z} \cdot \mathbf{W}_V, \\ \mathcal{A}(\mathbf{Q}, \mathbf{K}) &= \text{Softmax}(\mathbf{Q} \cdot \mathbf{K}^T / \sqrt{d}), \\ \hat{\mathbf{Z}} &= \text{SA}(\mathbf{Z}) = \mathcal{A}(\mathbf{Q}, \mathbf{K}) \cdot \mathbf{V}. \end{aligned} \quad (6)$$

To align the labels, we update the label embeddings  $\mathbf{Y}$  using the same attention matrix  $\mathcal{A}(\mathbf{Q}, \mathbf{K})$ :

$$\hat{\mathbf{Y}} = \mathcal{A}(\mathbf{Q}, \mathbf{K}) \cdot \mathbf{Y}. \quad (7)$$

ViTs usually adopt multi-head self-attention (MSA) to perform multiple self-attentions parallelly:

$$\hat{\mathbf{Z}} = \text{MSA}(\mathbf{Z}) = [\text{SA}_1(\mathbf{Z}); \text{SA}_2(\mathbf{Z}); \dots; \text{SA}_H(\mathbf{Z})] \cdot \mathbf{w}_h, \quad (8)$$

where  $H$  is the number of heads and  $\mathbf{w}_h \in \mathcal{R}^{d \times d}$ . We then adapt our label alignment to MSA by simply taking the average of all the attention matrices for alignment:

$$\hat{\mathbf{Y}} = \text{TL-Align-S}(\mathbf{Z}, \mathbf{Y}) := \frac{1}{H} \sum_{i=1}^H \mathcal{A}_i(\mathbf{Q}, \mathbf{K}) \cdot \mathbf{Y}, \quad (9)$$

where  $\mathcal{A}_i$  is the attention matrix corresponding to the  $i$ -th head  $\text{SA}_i$ .

Each transformer block  $l$  processes the tokens by both spatial and channel mixing:

$$\begin{aligned} \hat{\mathbf{Z}}^{l-1} &= \text{MSA}(\text{LN}(\mathbf{Z}^{l-1})), \mathbf{Z}^{l-1} = \hat{\mathbf{Z}}^{l-1} + \mathbf{Z}^{l-1}, \\ \hat{\mathbf{Z}}^l &= \text{MLP}(\text{LN}(\mathbf{Z}^{l-1})), \mathbf{Z}^l = \hat{\mathbf{Z}}^l + \mathbf{Z}^{l-1}, \end{aligned} \quad (10)$$

where MLP and LN denote the MLP module and layer normalization [2], respectively. Our TL-Align then aligns the label embeddings in a similar manner:

$$\begin{aligned} \hat{\mathbf{Y}}^{l-1} &= \text{TL-Align-S}(\mathbf{Y}^{l-1}), \mathbf{Y}^{l-1} = \text{Norm}(\hat{\mathbf{Y}}^{l-1} + \mathbf{Y}^{l-1}), \\ \hat{\mathbf{Y}}^l &= \mathbf{Y}^{l-1}, \mathbf{Y}^l = \text{Norm}(\hat{\mathbf{Y}}^l + \mathbf{Y}^{l-1}), \end{aligned} \quad (11)$$

where Norm denotes the normalization operation. We implement Norm by a simple average.

Hierarchical vision transformers such as Swin [31] further introduce a patch aggregation operation to merge multiple patches. They usually concatenate multiple tokens across the channels to reduce the spatial resolution. Instead of concatenation, we simply add the label embeddings of the merged tokens followed by normalization as the aligned labels. The proposed TL-Align can be generalized to different architectures composed of spatial mixing, channel mixing, point-wise transformation, residual connection, and spatial aggregation. We provide detailed illustrations of alignment with these operations in the supplementary.

We synchronously align the labels with the processed tokens layer by layer and obtain the aligned tokens  $\mathbf{Z}^L$  and labels  $\mathbf{Y}^L$ . The final representation of the image  $\mathbf{z}$  is either the class token  $\mathbf{z}_{cls}^L$  [16, 43] or the average pooling of all the spatial tokens  $\frac{1}{N} \sum_{i=1}^N \mathbf{z}_i^L$  [31]. The aligned label  $\mathbf{y}_{align}$  for the image is then  $\mathbf{y}_{cls}^L$  or  $\frac{1}{N} \sum_{i=1}^N \mathbf{y}_i^L$  depending on the specific model. We then adopt the aligned label  $\mathbf{y}_{align}$  to train the network and can adapt to different loss functions and training schemes:

$$J = J(\mathbf{z}, \text{stop-gradient}(\mathbf{y}_{align})). \quad (12)$$

We do not back-propagate through the aligned label as they only serve as a more accurate target.

Our TL-Align serves as a plug-and-play module on various vision transformers while only introducing negligible training costs. We adjust the label of each token adaptively during the layer-by-layer propagation and preserve alignment between tokens and labels throughout the forward process. TL-Align is only used during training and introduces no additional computation cost when inference.

## 4. Experiments

In this section, we conducted extensive experiments to evaluate the proposed TL-Align method. We demonstrate the improvement of TL-Align on various vision transformers and compare it with state-of-the-art training strategies concerning the accuracy, network complexity, and training speed. We also examine the transferability on downstream tasks including semantic segmentation, object detection, and transfer learning. We further provide in-depth analysis to evaluate the effectiveness of TL-Align.

### 4.1. ImageNet Classification

**Implementation Details.** We first evaluate TL-Align on ImageNet [38] for image classification. We implement our method based on PyTorch [35] and the timm library [50]. We conduct experiments on various transformer architectures: three variants of DeiT [43] (DeiT-T, DeiT-S, and DeiT-B), two variants of PVT [49] (PVT-T, PVT-S) and

Table 1. **Results on ImageNet classification task.** We compare the parameters, FLOPs, and accuracy of different vision transformer backbones without and with our TL-Align.

Model	Image Size	Params	FLOPs	Top-1(%)	Top-5(%)
DeiT-T	224 <sup>2</sup>	5.7M	1.6G	72.2	91.3
+TL-Align				<b>73.2</b>	<b>91.7</b>
PVT-T	224 <sup>2</sup>	13.2M	1.9G	75.1	92.4
+TL-Align				<b>75.5</b>	<b>93.0</b>
DeiT-S	224 <sup>2</sup>	22M	4.6G	79.8	95.0
+TL-Align				<b>80.6</b>	95.0
PVT-S	224 <sup>2</sup>	24.5M	3.8G	79.8	95.0
+TL-Align				<b>80.4</b>	<b>95.5</b>
Swin-T	224 <sup>2</sup>	29M	4.5G	81.2	95.5
+TL-Align				<b>81.4</b>	<b>95.7</b>
Swin-S	224 <sup>2</sup>	50M	8.8G	83.0	96.3
+TL-Align				<b>83.4</b>	<b>96.5</b>
DeiT-B	224 <sup>2</sup>	86M	17.5G	81.8	95.5
+TL-Align				<b>82.3</b>	<b>95.8</b>
Swin-B	224 <sup>2</sup>	88M	15.4G	83.5	96.4
+TL-Align				<b>83.7</b>	<b>96.5</b>

Table 2. **Comparison of our TL-Align with other training strategies on ImageNet.**

Method	Params	Trainig Time	Acc.(%)
Vanilla	22M	1×	76.4
CutMix [52]	22M	1×	79.8
Puzzle-Mix [24]	22M	2.24×	79.8
F-Mix [18]	22M	1.08×	77.4
ResizeMix [36]	22M	1.03×	78.6
SaliencyMix [45]	22M	1.26×	79.2
Attentive-CutMix [48]	46M	1.31×	77.5
TransMix [7]	22M	1.16×	80.1
TokenMix [30]	22M	1.65×	80.3
CutMix + TL-Align	22M	1.24×	<b>80.6</b>

Table 3. **Results on semantic segmentation ADE20K.**

Backbone	Params	FLOPs	mIoU	mIoU (MS)	mAcc
DeiT-S	58M	1032G	43.8	45.1	55.2
+TL-Align			<b>44.5</b>	<b>45.7</b>	<b>55.5</b>
Swin-T	60M	945G	44.4	45.8	55.6
+TL-Align			<b>44.7</b>	<b>46.5</b>	<b>56.4</b>
Swin-S	81M	1038G	47.6	49.5	58.8
+TL-Align			<b>48.0</b>	<b>49.7</b>	<b>59.5</b>
Swin-B	121M	1188G	48.1	49.7	59.1
+TL-Align			<b>48.3</b>	<b>50.1</b>	<b>59.7</b>

three variants of Swin Transformer [31] (Swin-T, Swin-S, and Swin-B). For tiny and small models, we train from scratch for 300 epochs following the same training recipe using CutMix as DeiT [43], PVT [49] and Swin [31]. We keep all the data augmentation policies and all hyperparameter settings unchanged for fair comparisons. The only modification is that we replace the mixing targets in Cut-

Table 4. Experimental results on object detection and instance segmentation on COCO.

Backbone	Params	FLOPs	Schedule	AP <sup>box</sup>	AP <sup>box</sup> <sub>50</sub>	AP <sup>box</sup> <sub>75</sub>	AP <sup>mask</sup>	AP <sup>mask</sup> <sub>50</sub>	AP <sup>mask</sup> <sub>75</sub>
Swin-T +TL-Align	86M	745G	3x	50.4 <b>50.5</b>	69.2 <b>69.4</b>	54.7 <b>54.9</b>	43.7 <b>43.8</b>	66.6 66.6	47.3 47.3
Swin-S +TL-Align	107M	838G	3x	51.9 <b>52.2</b>	70.7 <b>71.1</b>	56.3 <b>56.7</b>	45.0 <b>45.2</b>	68.2 <b>68.4</b>	48.8 <b>49.1</b>
Swin-B +TL-Align	145M	982G	3x	51.9 <b>52.3</b>	70.5 <b>71.2</b>	56.4 <b>56.9</b>	45.0 <b>45.3</b>	68.1 <b>68.7</b>	48.9 <b>49.1</b>

Table 5. Results on different transfer learning datasets.

Model	Params	FLOPs	C-10	C-100	Flowers	Cars
ResNet50	26M	4.1G	-	-	96.2	90.0
ViT-B/16	86M	55.4G	98.1	87.1	89.5	-
ViT-L/16	307M	190.7G	97.9	86.4	89.7	-
DeiT-T	5.7M	1.6G	97.6	85.7	97.1	90.1
+TL-Align	5.7M	1.6G	<b>97.8</b>	<b>86.4</b>	<b>97.9</b>	<b>90.7</b>
DeiT-S	22M	4.6G	97.9	90.2	98.1	91.4
+TL-Align	22M	4.6G	<b>98.8</b>	<b>90.4</b>	<b>98.3</b>	<b>91.8</b>
DeiT-B	86M	17.5G	99.1	90.8	98.4	92.1
+TL-Align	86M	17.5G	99.1	90.5	<b>98.6</b>	<b>93.0</b>




Figure 3. The Root Mean Square Error (RMSE) between original CutMix targets and labels obtained by T-L Align.

MiX with the labels obtained by our TL-Align. For base models (i.e., DeiT-B and Swin-B), we finetune the official pre-trained models for 40 epochs with a constant learning rate of  $1e-5$  and a weight decay of  $1e-8$ .

**Performance on Different Architectures.** As shown in Table 1, TL-Align steadily improves the performance of different vision transformer architectures. Specifically, TL-Align boosts the top-1 accuracy of DeiT-T, DeiT-S, and DeiT-B by 1.0%, 0.8%, and 0.5%, respectively, in a parameter-free manner. Moreover, our method is generalizable and can be directly applied to hierarchical vision transformers like Swin. It is worth noting that most existing methods need either architecture modifications (adding a class token in [7]) or extra computations (saliency map extraction in [45]) when applied to Swin. In contrast, our TL-Align method can be used as a plug-and-play module and achieves consistent improvement on variants of Swin.

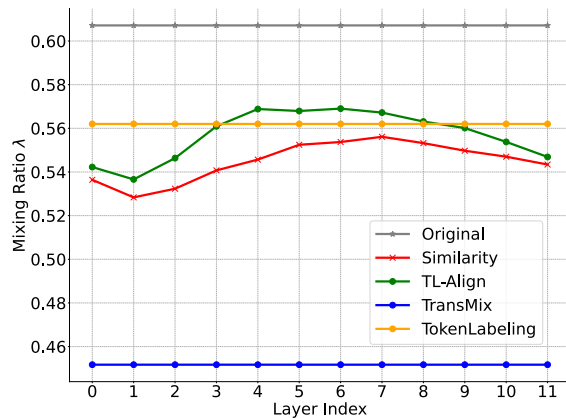


Figure 4. Visualization of mixing ratio  $\lambda$  of fluctuating tokens from different layers. We compare the results of TL-Align with CutMix, token similarity, TransMix, and TokenLabeling.

**Comparison with Other Training Strategies.** We also compare our method with the state-of-the-art training strategies for data mixing on DeiT-S, as shown in Table 2. Specifically, we train the DeiT-S model while only disabling CutMix as the baseline method, which is denoted as Vanilla. Moreover, since TransMix [7] and TokenMix [30] report the EMA accuracy with different hyperparameters, we reproduce it under the same training recipe [43] for a fair comparison. We see that TL-Align shows better performance than the other mixup variants while maintaining the number of parameters and training speed. Puzzle-Mix obtains the same classification accuracy as CutMix but results in a much lower training speed as it relies on an extra model to get the optimal solution. SaliencyMix and Attentive-CutMix lead to performance degeneration when built upon DeiT-S backbone. Notably, our method also achieves higher top-1 accuracy than ViT-targeted TransMix and TokenMix. While the teacher annotator in TokenMix introduces heavy training costs, TL-Align obtains accurate token-label alignment in an efficient manner. Moreover, TransMix has limitations in handling token fluctuation and cannot be applied to ViTs without a class token, such as Swin.

## 4.2. Downstream Tasks

**Semantic Segmentation.** We evaluate our TL-Align on ADE20K dataset [55] for semantic segmentation. We adopt DeiT-S and three variants of Swin Transformer as backbones equipped with UpperNet for segmentation. As pre-

Table 6. **Comparison results of model generalization ability and robustness.** We evaluate them on various out-of-distribution/corrupted datasets and against adversarial attacks.  $\uparrow$  denotes higher is better, and  $\downarrow$  denotes lower is better.

Model	Params	FLOPs	ImageNet		Generalization	Robustness			
			Top-1 $\uparrow$	Top-5 $\uparrow$	IN-V2 $\uparrow$	IN-A $\uparrow$	IN-C $\downarrow$	IN-R $\uparrow$	AutoAttack $\uparrow$
DeiT-T	5.7M	1.6G	72.2	91.3	60.4	7.7	69.1	34.1	3.9
+TL-Align	5.7M	1.6G	<b>73.2</b>	<b>91.7</b>	<b>61.4</b>	6.1	<b>68.0</b>	<b>34.6</b>	<b>4.4</b>
DeiT-S	22M	4.6G	79.8	95.0	68.5	18.9	54.7	42.5	6.9
+TL-Align	22M	4.6G	<b>80.6</b>	95.0	<b>68.9</b>	<b>19.2</b>	<b>53.2</b>	<b>43.2</b>	<b>7.5</b>
DeiT-B	86M	17.5G	81.8	95.5	70.5	27.9	48.5	45.3	-
+TL-Align	86M	17.5G	<b>82.3</b>	<b>95.8</b>	<b>70.9</b>	<b>29.0</b>	<b>47.1</b>	44.4	-

Table 7. **Ablation of using different data mixing strategies.**

MixUp	CutMix	Random	Block-wise	Top-1 (%)	+TL-Align Top-1 (%)
$\times$	$\times$	$\times$	$\times$	76.4	-
$\times$	$\checkmark$	$\times$	$\times$	79.8	80.6
$\checkmark$	$\checkmark$	$\times$	$\times$	79.8	80.2
$\times$	$\times$	$\checkmark$	$\times$	79.7	80.2
$\times$	$\times$	$\times$	$\checkmark$	80.0	80.3

sented in Table 3, TL-Align improves the segmentation performance on both DeiT and Swin at different model scales.

**Object Detection and Instance Segmentation.** We also examine the performance of TL-Align on object detection and instance segmentation on the COCO 2017 dataset [29]. We apply our TL-Align to Swin [31] due to the advantage of the hierarchical representations on object detection tasks. We adopt the Cascade Mask-RCNN [4] framework and use the training strategy of 3x schedule. As shown in Table 4, we observe consistent improvements on all variants of Swin Transformer. This demonstrates the advantages to learn token-level meaningful features suitable for dense prediction tasks.

**Transfer Learning.** We further evaluate the transferred classification performance of TL-Align on CIFAR-10 [26], CIFAR-100 [26], Flowers [34] and Cars [25]. We use pre-trained models on ImageNet and finetune them on these datasets following existing works [43]. We compare the performance with and without TL-Align on three variants of DeiT [43], as shown in Table 5. TL-Align obtains significant performance gains for all variants on the four datasets.

### 4.3. Performance Analysis and Visualization

**Effectiveness of Token-Label Alignment.** We first quantize the difference between the original targets and aligned labels and investigate its correlation with the model. Specifically, we compute the Root Mean Square Error (RMSE) between the original targets and labels obtained by our TL-Align. As shown in Figure 3, the RMSE decreases when enlarging the model size. This indicates that larger models demonstrates less token fluctuation since the self-correction ability is also enhanced as the model capacity scales up. Moreover, the RMSE for Swin Transformer

Table 8. **Ablation of different TL-Align operations.**

Alignment	Top-1 Acc.(%)	$\Delta$ (%)
None (DeiT-S baseline)	79.8	-
TL-Align-S (Layer 12)	80.1	+0.3
TL-Align-S (Layer 2,4,6,8)	80.2	+0.4
Normalization Disabled	80.3	+0.5
Default (TL-Align)	<b>80.6</b>	<b>+0.8</b>

tends to be lower compared with DeiT of similar model size. This is due to the local-window self-attention in Swin which preserves more local information. These observations are consistent with our experimental results: the improvements on small models and DeiT-like backbones tend to be more significant as they encounter more token fluctuation.

**Visualization of the Layer-wise Mixing Ratio of Fluctuated Tokens.** To investigate the effectiveness of TL-Align, we compute a similarity-based ‘‘ground-truth’’ mixing ratio for each layer. Specifically, we compute the similarities of tokens between the mixed and unmixed images and use them as the label of each token. We compare them with the mixing ratios produced by TL-Align, CutMix [52], TransMix [7], and TokenLabeling [23]. As shown in Figure 4, the similarity-based mixing ratio changes at each layer, resulting from token fluctuation. However, CutMix, TransMix, and TokenLabeling assume the output tokens keep spatial correspondence with the input tokens and compute a fixed mixing ratio. TL-Align assigns dynamic labels to tokens using layer-wise alignment, which is more accurate compared with other methods.

**Evaluation of Robustness and Generalization.** We further conduct experiments to validate the generalization and robustness of TL-Align, as shown in Table 6. We employ four corrupted and out-of-distribution datasets for robustness evaluation, including ImageNet-A [22], ImageNet-C [21], and ImageNet-R [20]. We also adopt AutoAttack [11] to evaluate the adversarial robustness on the ImageNet validation set. Due to memory limitation, we do not experiment with DeiT-B on AutoAttack. We use mean Corruption Error (mCE) for ImageNet-C and Top-1 Accuracy for others as the evaluation metric. For generalization evaluation, we adopt the ImageNet-V2 dataset [37]. We see that TL-Align improves both robustness and general-

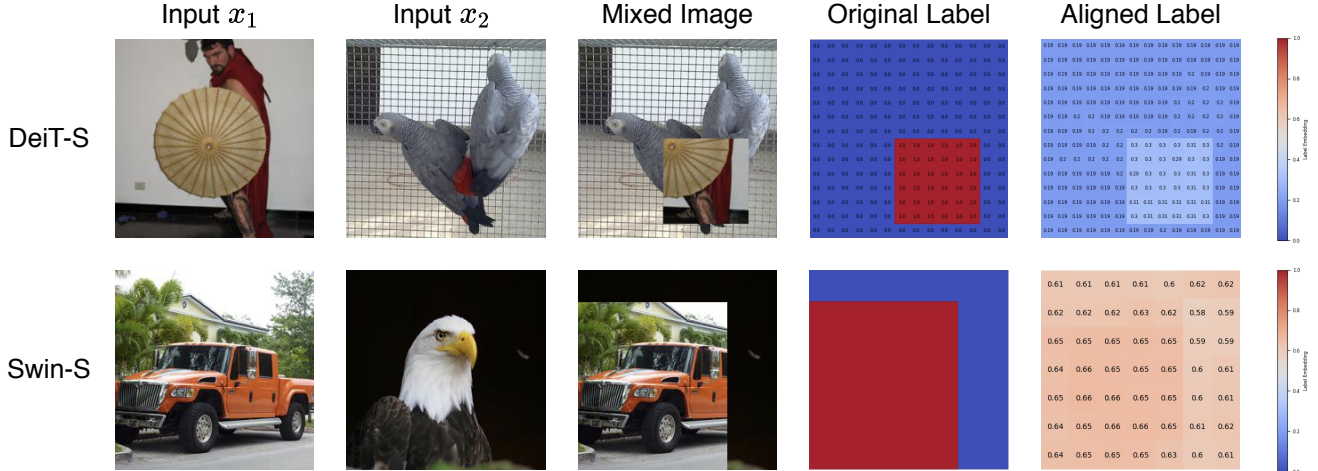


Figure 5. **The visualization of results on DeiT-S and Swin-S.** We visualize the input images, the mixed image, the original label embedding, and the label embedding after token-label alignment.

ization, showing the superiority of adopting TL-Align for pre-training.

#### Ablation Study on different Data Mixing Strategies.

Due to the efficiency of the proposed layer-wise alignment, TL-Align can be directly applied to a wide range of data mixing strategies. We adopt MixUp, CutMix, a random mixing strategy and a block-wise mixing strategy to evaluate the generalizability of TL-Align. The random mixing and block-wise mixing strategies are inspired by MAE [19] and BEiT [3] and we replace the masking operation with image mixing on patch-level and block-level (both of size  $16 \times 16$ ) respectively. The comparison results of training DeiT-S with and without our approach is demonstrated in Table 7. Specifically, TL-Align improves CutMix by 0.8%, MixUp+CutMix by 0.4%, random mixing by 0.5% and block-wise mixing by 0.3% respectively, further verifying the generalizability of the proposed TL-Align.

**Ablation Study on Different Label Alignment Operations.** Our TL-Align aligns the labels with tokens transformed by spatial self-attention and residual connection layer-by-layer. To investigate the effect of reusing attention maps and normalization, we conduct an ablation study regarding different alignment operations on DeiT-S. We try aligning the labels only by using the attention map of Layer 12, which is equivalent to TransMix [7]. We also test the performance of applying alignment to several middle transformer layers and disabling normalization. As presented in Table 8, incomplete alignment at a part of layers marginally boosts the performance as it cannot well handle the token fluctuation issue. Disabling normalization leads to 0.3% accuracy drop due to the inaccurate alignment at the presence of residual connections. This demonstrates the significance of the token-label alignment by attention utilization and normalization in a layer-wise manner.

**Visualizations of Aligned Labels.** We visualize the labels obtained by TL-Align on DeiT-S [43] and Swin-S [31]

as shown in Figure 5. Specifically, the aligned label embedding is obtained after the final transformer block for both DeiT-S and Swin-S. The value of the label embedding represents the probability of the belonged class of the corresponding token. We use red to denote larger probabilities towards the first image and blue for the second image. We observe that the aligned labels can deviate from the original labels and result in different mixing ratios for training. Therefore, using the original ratio as the training target may produce false training signals and lead to inferior performance. We see that TL-Align can correct the labels when the images are mixed with uninformative tokens. More visualization results are included in the supplementary.

## 5. Conclusion

In this paper, we have presented a token-label alignment method for training better vision transformers. As important subsets of data augmentation methods, data mixing strategies can generally improve the performance of both CNNs and ViTs. We identify a token fading issue for ViTs and address it by tracing the correspondence between transformed tokens and the original tokens to obtain a label for each output token to obtain more accurate training signals. Experimental results have demonstrated that our TL-Align can consistently improve the performance of various ViT models. The generalization performance of TL-Align to other architectures such as MLP-like models remains unknown and is a promising future direction to explore.

## Acknowledgements

This work was supported in part by the National Key Research and Development Program of China under Grant 2022ZD0160102 and in part by the National Natural Science Foundation of China under Grant 62125603.



## References

- [1] Anurag Arnab, Mostafa Dehghani, Georg Heigold, Chen Sun, Mario Lučić, and Cordelia Schmid. Vivit: A video vision transformer. In *ICCV*, pages 6836–6846, 2021. 2
- [2] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E Hinton. Layer normalization. *arXiv*, abs/1607.06450, 2016. 4
- [3] Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021. 8
- [4] Zhaowei Cai and Nuno Vasconcelos. Cascade r-cnn: Delving into high quality object detection. In *CVPR*, pages 6154–6162, 2018. 7
- [5] Nicolas Carion, Francisco Massa, Gabriel Synnaeve, Nicolas Usunier, Alexander Kirillov, and Sergey Zagoruyko. End-to-end object detection with transformers. In *ECCV*, pages 213–229, 2020. 1, 2
- [6] Hanqing Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *CVPR*, pages 12299–12310, 2021. 2
- [7] Jie-Neng Chen, Shuyang Sun, Ju He, Philip Torr, Alan Yuille, and Song Bai. Transmix: Attend to mix for vision transformers. *arXiv preprint arXiv:2111.09833*, 2021. 1, 3, 5, 6, 7, 8
- [8] Tianlong Chen, Zhenyu Zhang, Yu Cheng, Ahmed Awadallah, and Zhangyang Wang. The principle of diversity: Training stronger vision transformers calls for reducing all levels of redundancy. In *CVPR*, 2022. 1
- [9] Bowen Cheng, Alexander G Schwing, and Alexander Kirillov. Per-pixel classification is not all you need for semantic segmentation. In *NeurIPS*, 2021. 1, 2
- [10] Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. In *NeurIPS*, 2021. 1, 2
- [11] Francesco Croce and Matthias Hein. Reliable evaluation of adversarial robustness with an ensemble of diverse parameter-free attacks. In *International conference on machine learning*, pages 2206–2216, 2020. 7
- [12] Xiyang Dai, Yinpeng Chen, Jianwei Yang, Pengchuan Zhang, Lu Yuan, and Lei Zhang. Dynamic detr: End-to-end object detection with dynamic attention. In *ICCV*, pages 2988–2997, 2021. 1, 2
- [13] Zhigang Dai, Bolun Cai, Yugeng Lin, and Junying Chen. Up-detr: Unsupervised pre-training for object detection with transformers. In *CVPR*, pages 1601–1610, 2021. 1
- [14] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, pages 248–255. Ieee, 2009. 2
- [15] 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. In *CVPR*, 2022. 1
- [16] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 1, 2, 4, 5
- [17] 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. 1
- [18] Ethan Harris, Antonia Marcu, Matthew Painter, Mahesan Niranjan, Adam Prügel-Bennett, and Jonathon Hare. Fmix: Enhancing mixed sample data augmentation. *arXiv preprint arXiv:2002.12047*, 2020. 5
- [19] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 16000–16009, 2022. 8
- [20] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8340–8349, 2021. 7
- [21] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *arXiv preprint arXiv:1903.12261*, 2019. 7
- [22] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15262–15271, 2021. 7
- [23] Zi-Hang Jiang, Qibin Hou, Li Yuan, Daquan Zhou, Yujun Shi, Xiaojie Jin, Anran Wang, and Jiashi Feng. All tokens matter: Token labeling for training better vision transformers. In *NeurIPS*, 2021. 1, 2, 7
- [24] Jang-Hyun Kim, Wonho Choo, and Hyun Oh Song. Puzzle mix: Exploiting saliency and local statistics for optimal mixup. In *ICML*, pages 5275–5285, 2020. 1, 2, 3, 5
- [25] Jonathan Krause, Michael Stark, Jia Deng, and Li Fei-Fei. 3d object representations for fine-grained categorization. In *ICCVW*, pages 554–561, 2013. 7
- [26] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 7
- [27] Wenbo Li, Xin Lu, Jiangbo Lu, Xiangyu Zhang, and Jiaya Jia. On efficient transformer and image pre-training for low-level vision. *arXiv preprint arXiv:2112.10175*, 2021. 2
- [28] Xiaoxiao Li, Ziwei Liu, Ping Luo, Chen Change Loy, and Xiaoou Tang. Not all pixels are equal: Difficulty-aware semantic segmentation via deep layer cascade. In *CVPR*, pages 3193–3202, 2017. 1
- [29] 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*, pages 740–755, 2014. 7
- [30] Jihao Liu, Boxiao Liu, Hang Zhou, Hongsheng Li, and Yu Liu. Tokenmix: Rethinking image mixing for data augmentation in vision transformers. In *ECCV*, pages 455–471, 2022. 5, 6

- [31] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 1, 2, 5, 7, 8
- [32] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *arXiv preprint arXiv:2201.03545*, 2022. 1
- [33] Ze Liu, Jia Ning, Yue Cao, Yixuan Wei, Zheng Zhang, Stephen Lin, and Han Hu. Video swin transformer. *arXiv preprint arXiv:2106.13230*, 2021. 2
- [34] Maria-Elena Nilsback and Andrew Zisserman. Automated flower classification over a large number of classes. In *Indian Conference on Computer Vision, Graphics and Image Processing*, 2008. 7
- [35] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *NIPS*, pages 8026–8037, 2019. 5
- [36] Jie Qin, Jiemin Fang, Qian Zhang, Wenyu Liu, Xingang Wang, and Xinggang Wang. Resizemix: Mixing data with preserved object information and true labels. *arXiv preprint arXiv:2012.11101*, 2020. 5
- [37] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400, 2019. 7
- [38] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015. 5
- [39] Robin Strudel, Ricardo Garcia, Ivan Laptev, and Cordelia Schmid. Segmenter: Transformer for semantic segmentation. In *ICCV*, 2021. 1, 2
- [40] Yehui Tang, Kai Han, Jianyuan Guo, Chang Xu, Yanxi Li, Chao Xu, and Yunhe Wang. An image patch is a wave: Quantum inspired vision mlp. In *CVPR*, 2022. 1
- [41] 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*, 34, 2021. 1
- [42] 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*, 2021. 1
- [43] Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *ICML*, pages 10347–10357, 2021. 1, 2, 5, 6, 7, 8
- [44] Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit. *arXiv preprint arXiv:2204.07118*, 2022. 1, 2
- [45] AFM Uddin, Mst Monira, Wheemyung Shin, TaeChoong Chung, Sung-Ho Bae, et al. Saliencymix: A saliency guided data augmentation strategy for better regularization. *arXiv preprint arXiv:2006.01791*, 2020. 1, 2, 3, 5, 6
- [46] Shashanka Venkataramanan, Ewa Kijak, Laurent Amsaleg, and Yannis Avrithis. Alignmixup: Improving representations by interpolating aligned features. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19174–19183, 2022. 2
- [47] Vikas Verma, Alex Lamb, Christopher Beckham, Amir Najafi, Ioannis Mitliagkas, David Lopez-Paz, and Yoshua Bengio. Manifold mixup: Better representations by interpolating hidden states. In *ICML*, pages 6438–6447, 2019. 1, 3
- [48] Devesh Walawalkar, Zhiqiang Shen, Zechun Liu, and Marios Savvides. Attentive cutmix: An enhanced data augmentation approach for deep learning based image classification. *arXiv preprint arXiv:2003.13048*, 2020. 1, 2, 3, 5
- [49] 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. 5
- [50] Ross Wightman. Pytorch image models. <https://github.com/rwightman/pytorch-image-models>, 2019. 5
- [51] Lingfeng Yang, Xiang Li, Borui Zhao, Renjie Song, and Jian Yang. Recursivemix: Mixed learning with history. *arXiv preprint arXiv:2203.06844*, 2022. 2
- [52] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *ICCV*, pages 6023–6032, 2019. 1, 2, 3, 5, 7
- [53] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. In *ICLR*, 2018. 1, 2, 3
- [54] Sixiao Zheng, Jiachen Lu, Hengshuang Zhao, Xiatian Zhu, Zekun Luo, Yabiao Wang, Yanwei Fu, Jianfeng Feng, Tao Xiang, Philip HS Torr, et al. Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In *CVPR*, pages 6881–6890, 2021. 1
- [55] Bolei Zhou, Hang Zhao, Xavier Puig, Tete Xiao, Sanja Fidler, Adela Barriuso, and Antonio Torralba. Semantic understanding of scenes through the ade20k dataset. *IJCV*, 127:302–321, 2019. 6
- [56] 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. 1, 2