

Positional Label for Self-Supervised Vision Transformer

Zhemin Zhang¹, Xun Gong^{1,2,3*}

¹School of Computing and Artificial Intelligence, Southwest Jiaotong University, Chengdu, Sichuan, China

²Engineering Research Center of Sustainable Urban Intelligent Transportation, Ministry of Education, China

³Manufacturing Industry Chains Collaboration and Information Support Technology Key Laboratory of Sichuan Province, Chengdu, Sichuan, China

zheminzhang@my.swjtu.edu.cn, xgong@swjtu.edu.cn

Abstract

Positional encoding is important for vision transformer (ViT) to capture the spatial structure of the input image. General effectiveness has been proven in ViT. In our work we propose to train ViT to recognize the positional label of patches of the input image, this apparently simple task actually yields a meaningful self-supervisory task. Based on previous work on ViT positional encoding, we propose two positional labels dedicated to 2D images including absolute position and relative position. Our positional labels can be easily plugged into various current ViT variants. It can work in two ways: (a) As an auxiliary training target for vanilla ViT for better performance. (b) Combine the self-supervised ViT to provide a more powerful self-supervised signal for semantic feature learning. Experiments demonstrate that with the proposed self-supervised methods, ViT-B and Swin-B gain improvements of 1.20% (top-1 Acc) and 0.74% (top-1 Acc) on ImageNet, respectively, and 6.15% and 1.14% improvement on Mini-ImageNet. The code is publicly available at: <https://github.com/zhangzhemin/PositionalLabel>.

Introduction

Vision transformers (ViT) (Dosovitskiy et al. 2020; Zhao, Jia, and Koltun 2020) have recently emerged as an alternative to convolutional neural networks (CNNs) in computer vision. The design of ViT is inspired by the transformer of natural language processing (NLP), split an image into patches and provide the sequence of linear embeddings of these patches as an input to a Transformer (Touvron et al. 2021). Image patches are treated the same way as tokens (words) in NLP. The core of transformer is self-attention (Vaswani et al. 2017), which is able to model long-range dependencies in the data. However, self-attention is fully symmetric, it cannot capture the ordering of input tokens, which is undesirable for modeling structured data. Therefore, incorporating positional encoding is especially important for ViT.

There are mainly two classes of methods to add positional information for ViT. One is absolute and the other is relative. Absolute positional encoding (Gehring et al. 2017; Li et al. 2018): each absolute position of the input token sequence

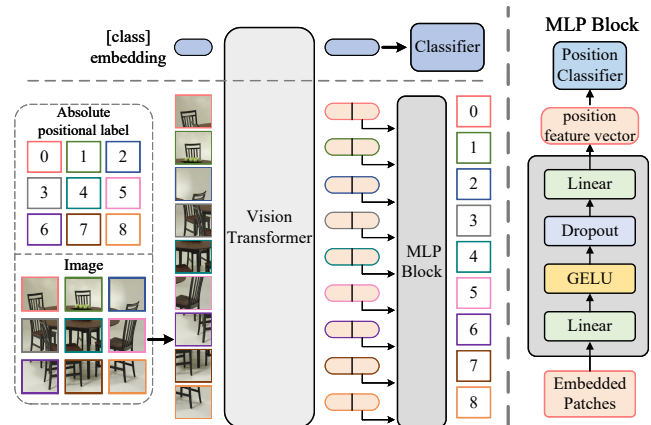


Figure 1: Our absolute positional label combined with ViT-B. We split an image into fixed-size patches, each corresponding to an absolute position, and feed the sequence of image patches to a standard ViT encoder. The full set of encoded patches output (intercept half) by the ViT is processed by a lightweight MLP block that outputs the positional feature vector corresponding to each image patch, which is finally fed into a positional classifier for training, the same as for vanilla classification. After training, the MLP block is discarded.

from 1 to maximum sequence length has a separate encoding vector. The encoding vector is then combined with the input tokens, usually element-wise add, to incorporate positional information into the ViT’s input tokens. On the other hand, relative positional encoding (Dai et al. 2019; Shaw, Uszkoreit, and Vaswani 2018) calculates the relative position between tokens based on the absolute position of each input token, and establishes a look-up table. Each relative position corresponds to a learnable parameter of the table to learn the pairwise relations of tokens. The parameters in the table interact with the self-attention weights. Relative positional encoding has been verified to be effective in ViT (Wu et al. 2021). Different from the previous explicit incorporation of positional information, we use the positional information as a supervision signal for ViT self-supervised training, making each encoded patch implicitly contains its positional information, as shown in Figure 1.

*Corresponding author.

In this paper, we seek to expand the applicability of positional information to serve as a supervised signal for self-supervised training of ViT. Based on the positional information, we propose a new ViT self-supervised loss function, namely position loss, which allows each token to implicitly contain its positional information. Specifically, we split an image into fixed-size patches and feed the sequence of image patches to a standard ViT. The full set of encoded patches output by the ViT is processed by a lightweight MLP block that outputs the positional label corresponding to each image patch. Based on previous research on ViT positional encoding, we propose two positional labels for ViT: absolute and relative positional labels. Section 3 illustrates how to apply these two positional labels to ViT and their differences.

Our position loss can work in two ways: **(a)** Combined with vanilla ViT (e.g., ViT-B (Dosovitskiy et al. 2020) and swin transformer (Liu et al. 2021)), as an auxiliary training task for the model. The ViT is trained under the joint supervision of the softmax loss and position loss, with a hyperparameter to balance the two supervision signals. We find that with joint supervision, the classification performance of ViT is significantly improved. **(b)** Combined with self-supervised ViT, e.g., MAE (He et al. 2021). MAE masks random patches from the input image and reconstructs the missing patches in the pixel space. Our position loss combined with MAE, allowing the model to not only reconstructs the pixels of the missing patches, but also outputs the positional information of the patches, providing a more powerful self-supervised signal for semantic feature learning. Our contributions can be summarized as:

- We propose two positional labels, absolute and relative positional labels, for ViT self-supervised learning. We also introduce efficient implementations of these two positional labels that can be easily plugged into self-attention layers.
- Our proposed positional label can be used as an alternative for positional encoding, which implicitly adds positional information to each image patch through ViT self-supervised training.
- The proposed positional labels can be combined with vanilla ViT and self-supervised ViT. Experiments show that, without adjusting any hyperparameters and settings, ViT-B and Swin-B obtained improvements of 1.20% (top-1 Acc) and 0.74% (top-1 Acc) on ImageNet, respectively. More improvements on the small dataset, ViT-B and Swin-B obtained improvements of 6.15% (top-1 Acc) and 1.14% (top-1 Acc) on Mini-ImageNet, respectively.

Related Work

Self-Supervised ViT

In pioneering works (Chen et al. 2020; He et al. 2021), training self-supervised Transformers for vision problems in general follows the masked auto-encoding paradigm in NLP (Devlin et al. 2018; Wang et al. 2019). iGPT (Chen et al. 2020) masks and reconstructs pixels. MAE (He et al. 2021) masks a high proportion of the input image patches

and reconstruct the missing patches. DINO (Caron et al. 2021) studies the importance of momentum encoder, multi-crop training (Caron et al. 2020), and the use of small patches with ViTs, to design a simple self-supervised approach that can be interpreted as a form of knowledge distillation with no labels. (Chen, Xie, and He 2021) focuses on training Transformers in the contrastive/Siamese paradigm, in which the loss is not defined for reconstructing the inputs. DILEMMA (Sameni, Jenni, and Favaro 2022) proposes a pseudo-task to train a ViT to detect which input token has been combined with an incorrect positional embedding to boost both shape and texture discriminability in models trained via self-supervised learning. In this work, we propose a novel self-supervised learning signal for ViT, named positional labels.

Positional Encoding

(a) Absolute Positional Encoding. Since transformer contains no recurrence and no convolution, in order for the model to make use of the order of the sequence, we need to inject some information about the position of the tokens. The original self-attention considers the absolute position, and add the absolute positional encodings to the input token embedding by element-wise addition (Vaswani et al. 2017). There are several choices of absolute positional encodings, such as Sinusoidal position (Vaswani et al. 2017) fixed encoding using sine and cosine functions sampling at different frequencies and (Gehring et al. 2017) the learnable encoding via training parameters.

(b) Relative Positional Encoding. Relative positional encoding is proposed firstly by (Shaw, Uszkoreit, and Vaswani 2018), extending the self-attention mechanism to efficiently consider representations of the relative positions, where relative positional encodings are added into the self-attention weight calculation. (Dai et al. 2019) proposed a novel positional encoding scheme with the prior of the sinusoid matrix and more learnable parameters, which not only enables capturing longer-term dependency, but also resolves the context fragmentation problem. (Wu et al. 2021) proposes new relative positional encoding methods dedicated to 2D images, called image RPE (iRPE). iRPE considers directional relative distance modeling as well as the interactions between queries and relative positional embeddings in self-attention mechanism. (Ramachandran et al. 2019) proposed 2D relative positional encoding that computes and concatenates separate encodings of each dimension. (Srinivas et al. 2021) proposes to incorporate the positional information in the convolutional feature map into ViT by concatenating the convolutional feature map with a set of feature maps produced via the self-attentive mechanism.

Unlike the previous explicit combination of positional information, we use the positional information as a supervised signal for ViT self-supervised training, such that each encoded patch implicitly contains its positional information. Section 4.2 discusses the relationship between our positional label and conventional positional encoding and demonstrates that they can be used jointly.

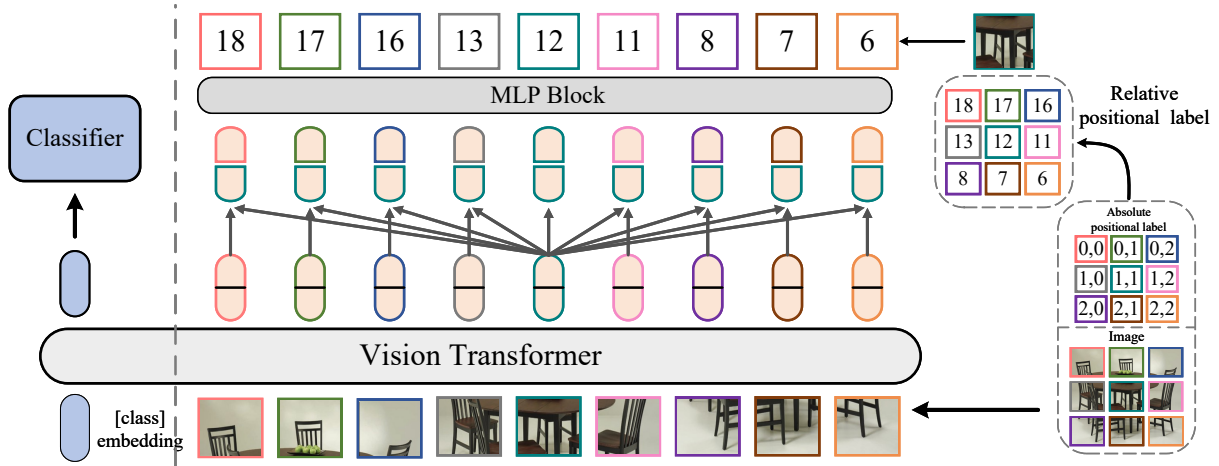


Figure 2: Our relative positional label combined with ViT-B. For the brevity of the image, here we take the relative positional label of the central image patch of the input image as an example, and the rest of the image patches are calculated in the same way. We combine the center encoded patch with other encoded patches pairwise, input a lightweight MLP block, and output the relative positional labels between the center image patch and other image patches. We use the method in Swin-ViT to convert 2D absolute positional labels to 1D relative positional labels

Method

Absolute Positional Label

Here we take ViT-B (Dosovitskiy et al. 2020) as an example, other ViT variants (e.g., Swin-B (Liu et al. 2021)) are calculated in the same way. An overview of the absolute position as an auxiliary training task for classification is depicted in Figure 1.

We split an image $x \in \mathbb{R}^{H \times W \times C}$ into fixed-size patches $x_m \in \mathbb{R}^{N \times (m^2 \cdot C)}$, where (H, W) is the resolution of the original image, C is the number of channels, (m, m) is the resolution of each image patch, and $N = HW/m^2$ is the number of generated patches. x_{class} is the classification token, whose state at the output of the ViT serves as the image representation z_{class} (Eq.2). In ViT-B, as an alternative to raw image patches, the input sequence can be formed from feature maps of a CNN (He et al. 2016), the patch embedding projection $E \in \mathbb{R}^{(m^2 \cdot C) \times D}$ is applied to patches extracted from a CNN feature map. The linear projection E maps x to D dimensions. We refer to the output y of this projection as the patch embeddings:

$$[y_{class}; y_1; y_2; \dots; y_N] = [x_{class}; x_m^1 E; x_m^2 E; \dots; x_m^N E] \quad (1)$$

Feed the sequence of patch embeddings to a standard ViT encoder:

$$[z_{class}; z_1; z_2; \dots; z_N] = \text{ViT}([y_{class}; y_1; y_2; \dots; y_N]) \quad (2)$$

where z_{class} is the image representation, which is input to the classifier to calculate the classification loss. θ denotes the trainable parameters of ViT. $z_i = (z_1, \dots, z_N)$ is the representation of the image patch $x_m^i = (x_m^1, \dots, x_m^N)$. We intercept half of z_i , named $\frac{1}{2}z_i$. Then use a trainable lightweight MLP block (shown in Figure 1) to project

$\frac{1}{2}z_i = (\frac{1}{2}z_1, \dots, \frac{1}{2}z_N)$ as the d -dimensional positional feature vector:

$$[p_1; p_2; \dots; p_N] = \text{MLP}\left(\left[\frac{1}{2}z_1; \frac{1}{2}z_2; \dots; \frac{1}{2}z_N\right]\right) \quad (3)$$

The probability that p_i belongs to y_i by SoftMax is:

$$P_{p_i, y_i} = \frac{\exp(w_{y_i}^T p_i)}{\sum_{j=1}^c \exp(w_j^T p_i)} \quad (4)$$

where $[w_1, \dots, w_c] \in \mathbb{R}^{d \times c}$ is the weights in the positional classifier, c denotes the number of classes (In ViT-B, $c = 196$), and d denotes the dimension of the feature. Each x_m^i corresponds to an absolute positional label $y_i = (y_1, \dots, y_N)$, as shown in Figure 1 ($y_1 = (0)$, \dots , $y_9 = (8)$). Absolute positional label trains the model by minimizing cross-entropy:

$$L_p = -\frac{1}{n} \sum_{i=1}^n \log(P_{p_i, y_i}) \quad (5)$$

In the conventional classification task, we adopt the joint supervision of softmax loss and position loss to train the ViTs for discriminative feature learning. The formulation is given in Eq.6.

$$L = L_s + \lambda L_p \quad (6)$$

where the scalar λ is used for balancing the two loss functions.

Relative Positional Label

Besides the absolute position of each input image patch, we also consider the relative positional relationship between patches. Relative position methods encode the relative distance between input patches and learn the pairwise relationships of patches. We encode the relative position between

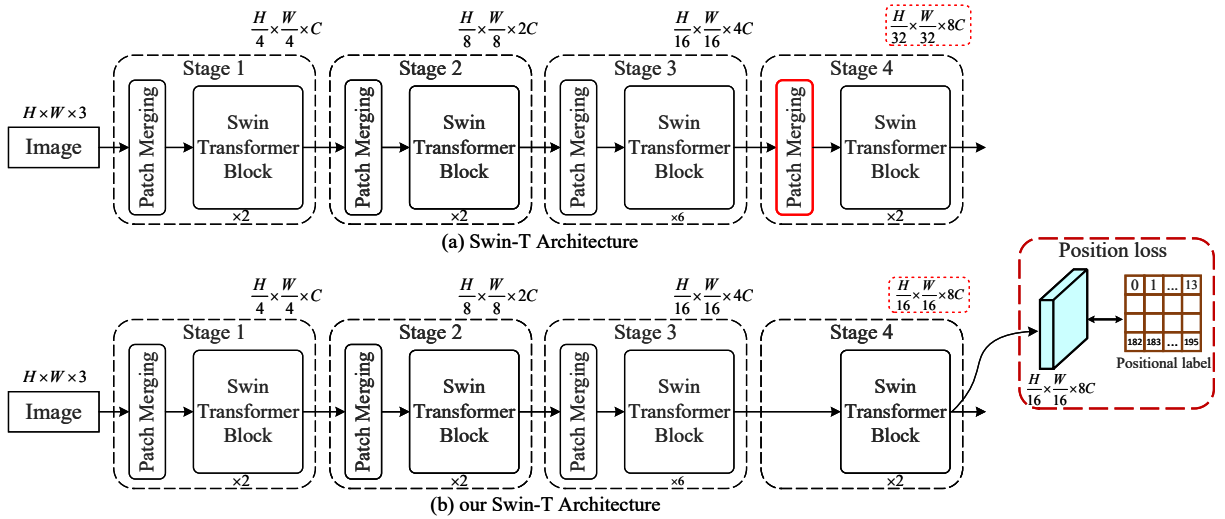


Figure 3: (a) The architecture of the original Swin Transformer (Swin-T); (b) The architecture of our modified Swin-T. This illustration was inspired by Swin-ViT figure 3(a).

the input image patches x_m^i and x_m^j into a vector p_{ij}^r . To establish the pairwise positional relationship between the output image patch embeddings, we intercept half of z_i , and combine them in pairs, as shown in Figure 2:

$$z_{ij} = \text{concat} \left(\frac{1}{2}z_i, \frac{1}{2}z_j \right) \quad (7)$$

where the *concat* operation concatenates the two embeddings. Same as the absolute positional label, feed z_{ij} to a lightweight MLP block:

$$p_{ij}^r = \text{MLP} (z_{ij}) \quad (8)$$

where p_{ij}^r denotes the relative positional feature vector between x_m^i and x_m^j output by the model. We use the method in Swin-ViT (Liu et al. 2021) to convert 2D absolute positional labels to 1D relative position labels y_{ij} . Similar to Eq.4 and 5, the optimization goal of the relative positional label is:

$$L_p^r = -\frac{1}{n^2} \sum_{i=1}^n \sum_{j=1}^n \log \left(P_{p_{ij}^r, y_{ij}} \right) \quad (9)$$

In the conventional classification task, the relative position loss is trained in the same way as Eq.6.

Coupling Positional Label and MAE

The above two positional labels can not only be used as auxiliary supervision signals for classification tasks, but also can be used for self-supervised training of ViT, which can be combined with various current ViT self-supervised methods to obtain more powerful self-supervised signals, e.g., MAE (He et al. 2021).

MAE masks a large random subset of image patches (e.g., 75%). A small subset of visible patches is then fed into the encoder. Mask tokens are introduced after the encoder, and the full set of encoded patches and mask tokens is processed by a small decoder that reconstructs the original image in




Figure 4: Our absolute positional label is combined with MAE. The output of the encoder is not only input to the decoder for MAE reconstruction of the image, but also to our positional label self-supervised MLP to output the absolute position of visible patches. This illustration was inspired by MAE.

pixels. This section illustrates our positional label combined with MAE, as shown in Figure 4. The output of the encoder is not only input to the decoder for MAE reconstruction of the image, but also to our positional label self-supervised MLP to output the absolute position of visible patches. Different from the original MAE, we do not add the encoder’s positional encoding during pre-training, and only use our positional labels for self-supervised training to fuse the positional information. Add the encoder’s positional encoding and use the proposed positional label, when fine-tuning the model.

Coupling Positional Label and Hierarchical ViT

Hierarchical ViT (e.g., Swin-ViT) to produce a hierarchical representation, the number of tokens is reduced by patch merging layers as the network gets deeper. The "Stage 3" and "Stage 4" of the original Swin-T, with output resolutions of $\frac{H}{16} \times \frac{W}{16}$ and $\frac{H}{32} \times \frac{W}{32}$, respectively, as shown in Figure 3(a). Taking the input image resolution of 224×224 as an example, the sequence length of the encoded patches output by the original Swin-T is $\frac{224}{32} \times \frac{224}{32} = 49$, which means that the number of classes of the positional label classification task $c = 49$, much smaller than the number of classes of ViT-B ($c = 196$), which we found experimentally that this affects the performance of positional label. In order to adapt to the positional label, we slightly modified the architecture of Swin-ViT and removed the last patch merging layer, so that the output resolution of Swin-T's "Stage 4" is $\frac{H}{16} \times \frac{W}{16}$, keeping the number of classes of the positional label classification task $c = 196$, as shown in Figure 3(b). We found that this will not affect the performance of Swin-ViT, but it will increase the amount of calculation, which is also the limitation of the positional label in hierarchical ViT, and we will explore the solution to this problem in the next work. In the following experiments, we use the modified architecture for hierarchical ViT.

Experiments

We conduct experiments on ImageNet-1K (Deng et al. 2009) image classification, Caltech-256 (Griffin, Holub, and Perona 2007) and Mini-ImageNet (Krizhevsky, Sutskever, and Hinton 2012) small datasets image classification. In the following, we first conduct ablation experiments to demonstrate that positional encoding can be combined with the proposed positional labels, and then combine the proposed positional labels with the latest ViT architecture to demonstrate the effectiveness of the positional labels. Finally, we combine the proposed positional labels with MAE to demonstrate the potential of positional labels in self-supervised ViT training.

Experiment Settings

Dataset. For image classification, we benchmark the proposed positional label on the ImageNet-1K, which contains 1.28M training images and 50K validation images from 1,000 classes. To explore the performance of positional label on small datasets, we also conducted experiments on Caltech-256 and Mini-ImageNet. Caltech-256 has 257 classes with more than 80 images in each class. Mini-ImageNet contains a total of 60,000 images from 100 classes.

Implementation details. This setting mostly follows (Liu et al. 2021). We use the PyTorch toolbox (Paszke et al. 2019) to implement all our experiments. We employ an AdamW (Kingma and Ba 2014) optimizer for 300 epochs using a cosine decay learning rate scheduler and 20 epochs of linear warm-up. A batch size of 256, an initial learning rate of 0.001, and a weight decay of 0.05 are used. ViT-B/16 uses an image size 384 and others use 224. We include most of

Method	Top-1 acc.	Top-5 acc.
ViT-B + PE (Baseline)	77.91	92.48
ViT-B	76.34	90.83
ViT-B + APL	78.19	92.87
ViT-B + RPL	78.00	92.40
ViT-B + PE + RPL	79.07	93.67
ViT-B + PE + APL	79.11	93.73

Table 1: Top-1 and Top-5 accuracies (%) for ViT-B model, on the ImageNet dataset. PE: positional encoding term of ViT; APL: the proposed absolute positional label; RPL: the proposed relative positional label.

Method	Top-1 acc.	Top-5 acc.
Swin-T + PE (Baseline)	81.32	95.64
Swin-T	80.14	94.93
Swin-T + APL	80.89	95.17
Swin-T + RPL	81.15	95.30
Swin-T + PE + RPL	81.93	95.67
Swin-T + PE + APL	81.51	95.83

Table 2: Top-1 and Top-5 accuracies (%) for Swin-T model, on the ImageNet dataset.

the augmentation and regularization strategies of (Liu et al. 2021) in training.

Ablation Studies

Positional encoding. Incorporating explicit representations of positional information is known to be especially important for Transformers, since the model is otherwise entirely invariant to sequence ordering, which is undesirable for modeling structured data. Experiments show that our proposed positional label can completely replace positional encoding to add positional information to the ViT model. Surprisingly, we found that the proposed positional label can be used in combination with conventional positional encoding, the model does not learn an identity and can achieve the best performance gain. We experiment on ViT-B and Swin-T, as shown in Table 1 and Table 2.

Hyperparameter λ . The hyperparameter λ balances the classification loss and position loss. It is essential to our




Figure 5: Illustration of position loss with different λ , using the ViT-B architecture.




Figure 6: Data augmentation. Our positional label works with different image augmentations.

model. So we conduct an experiment to investigate the sensitiveness of this parameter. In our experiments, we vary λ from 0 to 1.25 to learn different models. The Top-1 accuracies of these models on the ImageNet-1K dataset are shown in Figure 5. Properly choosing the value of λ can improve the accuracy of the models. We also observe that the performance of the models remains largely stable across a wide range of λ . All experiments in this paper set $\lambda = 0.5$.

Data augmentation. Because the data augmentation affects our positional labels, as shown in Figure 6, the image uses different data augmentations, and the bird head corresponds to different positional labels. Left: positional label is 2. Middle: positional label is 0. Right: positional label is 8. Table 3 studies the influence of data augmentation on our positional label. Using random horizontal flipping for our positional label improves the performance, but combining random horizontal flipping and random vertical flipping has little effect on the results, suggesting that positional label does not require large augmentation to regularize training.

Image Classification on the ImageNet-1K

We combine absolute and relative positional labels with various ViT variants, each ViT variant follows the approach of adding positional encoding in its paper. The experimental results show that the proposed positional label can significantly improve the performance of ViT, as shown in Table 4: +1.20% for ViT-B with the absolute positional label (79.11%) over ViT-B (77.91%), and +0.74% for Swin-B with the relative positional label (84.09%) over Swin-B (83.35%). The experimental results show that the performance improvement of positional labels for the full-

Method	Top-1 acc.	Top-5 acc.
crop + RHF (Baseline)	77.91	92.48
crop + APL	78.54	93.13
crop + RPL	78.26	92.95
crop + RHF + APL	79.11	93.73
crop + RHF + RPL	79.07	93.67
crop + RHF + RVF + APL	79.09	93.69
crop + RHF + RVF + RPL	79.10	92.87

Table 3: Data augmentation. Our positional label works with minimal augmentation, using the ViT-B architecture. crop: random resized crop; RHF: random horizontal flipping; RVF: random vertical flipping.

Method	Top-1 acc.	Top-5 acc.
ViT-B (Dosovitskiy et al. 2020)	77.91	92.48
DeiT-B (Touvron et al. 2021)	81.87	93.92
Swin-B (Liu et al. 2021)	83.35	96.38
NesT-B (Zhang et al. 2022)	83.67	96.16
ViT-B + APL	79.11	93.73
DeiT-B + APL	82.49	94.34
Swin-B + APL	83.76	96.83
NesT-B + APL	84.13	96.85
ViT-B + RPL	79.07	93.67
DeiT-B + RPL	82.85	94.21
Swin-B + RPL	84.09	96.77
NesT-B + RPL	83.93	96.81

Table 4: Top-1 and Top-5 accuracies (%) of different ViT backbones, on the ImageNet dataset. Each ViT variant follows the approach of adding positional encoding in its paper.

attention transformer (e.g., ViT-B) is higher than that of the local-attention transformer (e.g., Swin-B), which also indicates that the local-attention transformer can process limited data more efficiently.

In the classification task, we use a classification loss function combined with our positional labels for joint training. We observe the classification loss changing curve of ViT and find an interesting phenomenon that the classification loss of ViT with positional labels is smaller (as shown in Figure 7), which indicates that our positional labels can benefit the convergence of ViT.

Image Classification on Caltech-256 and Mini-ImageNet

We show the performance of ViT on small datasets in Table 5 and Table 6. It is known that ViTs usually perform poorly on such tasks as they typically require large datasets to be trained on. The models that perform well on large-scale ImageNet do not necessary work perform on small-scale Mini-ImageNet and Caltech-256, e.g., ViT-B has top-1 accuracy of 58.28% and Swin-B has top-1 accuracy of 67.39% on the Mini-ImageNet, which suggests that ViTs are more challenging to train with less data. Our proposed positional label can accelerate the convergence of ViT and significantly im-




Figure 7: During training, the classification loss of ViT-B with the positional label and ViT-B without the positional label.

Method	Top-1 acc.	Top-5 acc.
ViT-B (Dosovitskiy et al. 2020)	58.28	79.57
DeiT-B (Touvron et al. 2021)	63.67	83.92
Swin-B (Liu et al. 2021)	67.39	86.88
NesT-B (Zhang et al. 2022)	67.43	86.75
ViT-B + APL	64.43	83.73
DeiT-B + APL	66.49	85.34
Swin-B + APL	68.91	87.83
NesT-B + APL	68.73	87.95
ViT-B + RPL	63.97	83.06
DeiT-B + RPL	66.85	85.21
Swin-B + RPL	69.11	88.02
NesT-B + RPL	69.56	88.67

Table 5: Top-1 and Top-5 accuracies (%) of various ViT variants, on the Mini-ImageNet dataset.

Method	Top-1 acc.	Top-5 acc.
ViT-B (Dosovitskiy et al. 2020)	37.57	56.83
DeiT-B (Touvron et al. 2021)	40.56	61.24
Swin-B (Liu et al. 2021)	46.67	67.22
NesT-B (Zhang et al. 2022)	45.54	66.35
ViT-B + APL	40.73	60.63
DeiT-B + APL	41.79	62.86
Swin-B + APL	48.91	68.33
NesT-B + APL	49.77	68.84
ViT-B + RPL	40.51	61.07
DeiT-B + RPL	41.28	61.56
Swin-B + RPL	48.80	68.09
NesT-B + RPL	48.97	67.99

Table 6: Top-1 and Top-5 accuracies (%) of various ViT variants, on Caltech-256 dataset.

prove the performance of ViT on small datasets like Mini-ImageNet and Caltech-256. On Mini-ImageNet, the top-1 accuracy of ViT-B and Swin-B are improved by 6.15% and 1.72%, respectively. On Caltech-256, the top-1 accuracy of ViT-B and Swin-B are improved by 3.16% and 2.24%, respectively.

Coupling Positional Label and MAE

We do self-supervised pre-training on the ImageNet-1K training set. Then we do supervised training to evaluate the representations with end-to-end fine-tuning. We report top-1 validation accuracy. This experiment uses ViT-B as the backbone, and the experimental setting follows the default setting of MAE (He et al. 2021). Table 7 is a comparison between ViT-B trained from scratch vs. fine-tuned from MAE with

Method	Top-1 acc.
ViT-B (Dosovitskiy et al. 2020)	77.91
ViT-B + MAE (He et al. 2021)	79.54
ViT-B + MAE + APL	80.40
ViT-B + MAE + RPL	80.21

Table 7: Experimental results of positional label combined with MAE on the ImageNet-1K. Pre-trained for 400 epochs.




Figure 8: The position loss with different masking ratios.

our positional label.

MAE finds that masking a high proportion of the input image, e.g., 75%, yields a nontrivial and meaningful self-supervisory task. Random sampling with a high masking ratio largely eliminates redundancy, thus creating a task that cannot be easily solved by extrapolation from visible neighboring patches. We find that this also applies to our proposed positional label, where the encoder uses a small subset of visible patches to reconstruct the overall positional information of the image, yielding a more powerful self-supervised signal. As shown in Figure 8, the training difficulty of the model increases as the masking ratio increases.

Under this design, MAE and our positional label can achieve a win-win scenario: MAE reconstructs clear masked patch pixels that help the positional loss capture the overall positional information of the image, while the positional loss also incorporates the positional information into the MAE reconstruction, allowing the model to consider the overall structure of the image during the reconstruction process.

Conclusions

In this paper, we propose two positional labels for self-supervised training of ViT. The extensive experiments show that our method can be combined with various ViT variants, bringing significant improvements on classification tasks. Our methods could be easily plugged into self-attention layers. In addition, our positional labels can also be used for ViT fully self-supervised methods to provide a powerful self-supervised signal.

Acknowledgments

We acknowledge with thanks the discussion with Jinyi Wu and Shijie Wu from Southwest Jiaotong University, as well as the many others who have. This work is partially supported by the National Natural Science Foundation of China (61876158), Sichuan Science and Technology Program (2023YFG0267), Fundamental Research Funds for the Central Universities (2682021ZTPY030, 2682022KJ045).

References

Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised Learning of Visual Features by Contrasting Cluster Assignments. In Larochelle,

- H.; Ranzato, M.; Hadsell, R.; Balcan, M.; and Lin, H., eds., *Advances in Neural Information Processing Systems*, volume 33, 9912–9924. Curran Associates, Inc.
- Caron, M.; Touvron, H.; Misra, I.; Jégou, H.; Mairal, J.; Bojanowski, P.; and Joulin, A. 2021. Emerging Properties in Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9650–9660.
- Chen, M.; Radford, A.; Child, R.; Wu, J.; Jun, H.; Luan, D.; and Sutskever, I. 2020. Generative Pretraining From Pixels. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 1691–1703. PMLR.
- Chen, X.; Xie, S.; and He, K. 2021. An Empirical Study of Training Self-Supervised Vision Transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 9640–9649.
- Dai, Z.; Yang, Z.; Yang, Y.; Carbonell, J. G.; Le, Q. V.; and Salakhutdinov, R. 2019. Transformer-XL: Attentive Language Models Beyond a Fixed-Length Context. *CoRR*, abs/1901.02860.
- Deng, J.; Dong, W.; Socher, R.; Li, L.-J.; Li, K.; and Fei-Fei, L. 2009. ImageNet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, 248–255.
- Devlin, J.; Chang, M.; Lee, K.; and Toutanova, K. 2018. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. *CoRR*, abs/1810.04805.
- Dosovitskiy, A.; Beyer, L.; Kolesnikov, A.; Weissenborn, D.; Zhai, X.; Unterthiner, T.; Dehghani, M.; Minderer, M.; Heigold, G.; Gelly, S.; Uszkoreit, J.; and Houshy, N. 2020. An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale. *CoRR*, abs/2010.11929.
- Gehring, J.; Auli, M.; Grangier, D.; Yarats, D.; and Dauphin, Y. N. 2017. Convolutional Sequence to Sequence Learning. In Precup, D.; and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 1243–1252. PMLR.
- Griffin, G.; Holub, A.; and Perona, P. 2007. Caltech-256 object category dataset. *Technical Report UCB/CSD-04-1366*, California Institute of Technology.
- He, K.; Chen, X.; Xie, S.; Li, Y.; Dollár, P.; and Girshick, R. B. 2021. Masked Autoencoders Are Scalable Vision Learners. *CoRR*, abs/2111.06377.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kingma, D. P.; and Ba, J. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25: 1097–1105.
- Li, C.; Zhang, Z.; Lee, W. S.; and Lee, G. H. 2018. Convolutional Sequence to Sequence Model for Human Dynamics. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Liu, Z.; Lin, Y.; Cao, Y.; Hu, H.; Wei, Y.; Zhang, Z.; Lin, S.; and Guo, B. 2021. Swin Transformer: Hierarchical Vision Transformer Using Shifted Windows. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10012–10022.
- Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32: 8026–8037.
- Ramachandran, P.; Parmar, N.; Vaswani, A.; Bello, I.; Levskaya, A.; and Shlens, J. 2019. Stand-Alone Self-Attention in Vision Models. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Sameni, S.; Jenni, S.; and Favaro, P. 2022. DILEMMA: Self-Supervised Shape and Texture Learning with Transformers. *arXiv preprint arXiv:2204.04788*.
- Shaw, P.; Uszkoreit, J.; and Vaswani, A. 2018. Self-Attention with Relative Position Representations. *CoRR*, abs/1803.02155.
- Srinivas, A.; Lin, T.-Y.; Parmar, N.; Shlens, J.; Abbeel, P.; and Vaswani, A. 2021. Bottleneck Transformers for Visual Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16519–16529.
- Touvron, H.; Cord, M.; Douze, M.; Massa, F.; Sablayrolles, A.; and Jégou, H. 2021. Training data-efficient image transformers-and; distillation through attention. In Meila, M.; and Zhang, T., eds., *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, 10347–10357. PMLR.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Wang, A.; Pruksachatkun, Y.; Nangia, N.; Singh, A.; Michael, J.; Hill, F.; Levy, O.; and Bowman, S. 2019. SuperGLUE: A Stickier Benchmark for General-Purpose Language Understanding Systems. In Wallach, H.; Larochelle, H.; Beygelzimer, A.; d'Alché-Buc, F.; Fox, E.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Wu, K.; Peng, H.; Chen, M.; Fu, J.; and Chao, H. 2021. Rethinking and Improving Relative Position Encoding for Vision Transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 10033–10041.

Zhang, Z.; Zhang, H.; Zhao, L.; Chen, T.; Arik, S. O.; and Pfister, T. 2022. Nested hierarchical transformer: Towards accurate, data-efficient and interpretable visual understanding. In *AAAI Conference on Artificial Intelligence (AAAI)*, volume 2022.

Zhao, H.; Jia, J.; and Koltun, V. 2020. Exploring Self-Attention for Image Recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.