


Self-Promoted Supervision for Few-Shot Transformer

Bowen Dong¹, Pan Zhou², Shuicheng Yan², and Wangmeng Zuo^{1,3}

¹Harbin Institute of Technology ²National University of Singapore

³Peng Cheng Laboratory

{cndongsky,panzhou3,shuicheng.yan}@gmail.com, wzmzuo@hit.edu.cn

Abstract. The few-shot learning ability of vision transformers (ViTs) is rarely investigated though heavily desired. In this work, we empirically find that with the same few-shot learning frameworks, replacing the widely used CNN feature extractor with a ViT model often severely impairs few-shot classification performance. Moreover, our empirical study shows that in the absence of inductive bias, ViTs often learn the low-qualified token dependencies under few-shot learning regime where only a few labeled training data are available, which largely contributes to the above performance degradation. To alleviate this issue, we propose a simple yet effective few-shot training framework for ViTs, namely Self-promoted sUpervisiON (SUN). Specifically, besides the conventional global supervision for global semantic learning, SUN further pretrains the ViT on the few-shot learning dataset and then uses it to generate individual location-specific supervision for guiding each patch token. This location-specific supervision tells the ViT which patch tokens are similar or dissimilar and thus accelerates token dependency learning. Moreover, it models the local semantics in each patch token to improve the object grounding and recognition capability which helps learn generalizable patterns. To improve the quality of location-specific supervision, we further propose: 1) background patch filtration to filtrate background patches out and assign them into an extra background class; and 2) spatial-consistent augmentation to introduce sufficient diversity for data augmentation while keeping the accuracy of the generated local supervisions. Experimental results show that SUN using ViTs significantly surpasses other few-shot learning frameworks with ViTs and is the first one that achieves higher performance than those CNN state-of-the-arts. Our code is publicly available at <https://github.com/DongSky/few-shot-vit>.

Keywords: few-shot learning · location-specific supervision

1 Introduction

Vision transformers (ViTs) have achieved great success in the computer vision field, and even surpass corresponding state-of-the-art CNNs on many vision tasks, e.g. image classification [10,41,42,45,29,57], object detection [4,11] and segmentation [58,37]. One key factor contributing to ViTs' success is their powerful self-attention mechanism [43], which does not introduce any inductive bias




Fig. 1. Performance comparisons of different few-shot classification frameworks. Except the state-of-the-art RE-Net [19] which uses ResNet-12 [15] as the feature extractor, all methods use the same NesT transformer [57] as the feature extractor. With same ViT backbone, our SUN significantly surpasses other baselines.

and can better capture the long-range dependencies among local features in the data than the convolution mechanism in CNNs. This motivates us to investigate two important problems. First, we wonder whether ViTs can perform well under few-shot learning setting or not, which aims to recognize objects from novel categories with only a handful of labeled samples as reference. Second, if no, how to improve the few-shot learning ability of ViTs? The few-shot learning ability is heavily desired for machine intelligence, since in practice, many tasks actually have only a few labeled data due to data scarcity (e.g. disease data), high equipment cost to capture data and manual labeling cost (e.g. pixel-level labeling). In this work, we are particularly interested in few-shot classification [44,12,36] which is a benchmark task to evaluate few-shot learning capacity of a method.

For the first question, unfortunately, we empirically find that for representative few-shot learning frameworks, e.g. Meta-Baseline [6], replacing the CNN feature extractor by ViTs severely impairs few-shot classification performance. The most possible reason is the lack of inductive bias in ViTs—in absence of any prior inductive bias, ViTs needs a large amount of data to learn the dependencies among local features, i.e. patch tokens. Our analysis in Sec. 3 shows that introducing CNN-alike inductive bias can partly accelerate the token dependency learning in ViTs and thus partially mitigate this effect, which also accords with the observations in [22,56] under supervised learning. For example, in Fig. 1, the “Baseline” [5] and “Meta-Baseline” [6] use vanilla ViTs, while “CNN Distill” introduces CNN-alike inductive bias via distilling a pretrained CNN teacher network into the ViT and achieves higher performance. However, CNN-alike inductive bias cannot be inherently suitable for ViTs and cannot well enhance and accelerate token dependency learning in ViTs, because ViTs and CNNs have different network architectures, and also the inductive bias in CNNs often cannot well capture long-range dependencies. This naturally motivates us to consider how to improve the few-shot learning ability of ViTs.

We therefore propose an effective few-shot learning framework for ViTs, namely, *Self-promoted sUpervision* (SUN for short). SUN enables a ViT model

to generalize well on novel unseen categories with a few labeled training data. The key idea of SUN is to enhance and accelerate dependency learning among different tokens, i.e., patch tokens, in ViTs. This idea is intuitive, since if ViTs can learn the token dependencies fast and accurately, and thus naturally can learn from less labeled training data, which is consistent with few-shot learning scenarios. In particular, at the meta-training phase, SUN provides global supervision to global feature embedding and further employs individual location-specific supervision to guide each patch token. Here SUN first trains the ViT on the training data in few-shot learning, and uses it to generate patch-level pseudo labels as location-specific supervision. This is why we call our method “Self-Promoted Supervision” as it uses the same ViT to generate local supervision. To improve the quality of patch-level supervision, we further propose two techniques, namely 1) background patch filtration and 2) spatial-consistent augmentation. The former aims to alleviate effects of bad cases where background patches are wrongly assigned to a semantic class and have incorrect local supervision; the latter is to introduce sufficient diversity to data augmentation while keeping the accuracy of generated location-specific supervisions. Our location-specific dense supervision benefits ViT on few-shot learning tasks from two aspects. Firstly, considering the location-specific supervisions on all tokens are consistent, i.e. similar pseudo labels for similar local tokens, this dense supervision tells ViT which patch tokens are similar or dissimilar and thus can accelerate ViT to learn high-quality token dependencies. Secondly, the local semantics in each patch token are also well modeled to improve the object grounding and recognition capabilities of ViTs. Actually, as shown in [60,18], modeling semantics in local tokens can avoid learning skewed and non-generalizable patterns and thus substantially improve the model generalization performance which is also heavily desired in few-shot learning. So both aspects can improve the few-shot learning capacity of ViTs. Next, at the meta-tuning phase, similar to existing methods [6,51], SUN adapts the knowledge of ViT trained at the meta-training phase to new tasks by fine-tuning ViT on the corresponding training data.

Extensive experimental results demonstrate that with the same ViT architecture, our SUN framework significantly outperforms all baseline frameworks on the few-shot classification tasks, as shown in Fig. 1. Moreover, to the best of our knowledge, SUN with ViT backbones is the first work that applies ViT backbones to few-shot classification, and is also the first ViT method that achieves comparable even higher performance than state-of-the-art baselines using similar-sized CNN backbones, as illustrated in Fig. 1. This well shows the superiority of SUN, since CNN has high inductive bias and is actually much more suitable than ViT on few-shot learning problems. Additionally, our SUN also provides a simple yet solid baseline for few-shot classification.

2 Related Work

Vision Transformer. Unlike CNNs that use convolutions to capture local information in images, ViTs [10] model long-range interaction explicitly, and have

shown great potential for vision tasks. However, to achieve comparable performance with CNN of a similar model size trained on ImageNet [8], ViTs need huge-scale datasets to train [8,38], greatly increasing computational cost. To alleviate this issue and further improve performance, many works [41,18,53,49,29,57] [47,52,40,23,26,45,50,54] propose effective solutions. However, these ViTs still need to train on large-scale datasets, e.g. ImageNet [8] or JFT-300M [38], and often fail on small datasets [57,22,56]. Recently, a few works [22,3] also propose ViTs for small datasets. For instance, Li *et al.* [22] initialized ViT from a well-trained CNN to borrow inductive bias in CNN to ViT. Cao *et al.* [3] first used instance discriminative loss to pretrain ViT and then fine-tuned it on target datasets. In comparison, we propose self-promoted supervision to provide dense local supervisions, enhancing token dependency learning and alleviating the data-hungry issue. Moreover, we are interested in few-shot learning problems that focus on initializing ViTs with small training data from base classes and generalizing on novel unseen categories with a few labeled data.

Few-Shot Classification. Few-shot classification methods can be coarsely divided into three groups: 1) *optimization-based methods* [12,35,20,61,62] for fast adaptation to new tasks; 2) *memory-based methods* [33,31,14] for storing important training samples; and 3) *metric-based methods* [36,6,44,51] which learn the similarity metrics among samples. This work follows the pipeline of metric-based methods for the simplicity and high accuracy. It is worth noting that some recent works [9,17,51,59,27] also integrate transformer layers with metric-based methods for few shot learning. Different from previous works [9,17,51,59,27] which leverage transformer layers as few-shot classifiers, our work focuses on improving the few-shot classification accuracy of ViT backbones.


3 Empirical Study of ViTs for Few-Shot Classification

Here we first introduce the few-shot classification task. Given a labeled base dataset \mathbb{D}_{base} which contains base classes C_{base} , this task aims to learn a meta-learner f (feature extractor) such that f can be fast adapted to unseen classes C_{novel} (i.e. $C_{\text{base}} \cap C_{\text{novel}} = \emptyset$) in which each class has a few labeled training samples. This task is a benchmark to evaluate few-shot learning capacity of a method, since it needs to well learn the knowledge in the base data \mathbb{D}_{base} and then fast adapt the learnt knowledge to the new data C_{novel} via only a few training data. Next, we investigate the ViTs’ performance on this important task and also analyze the potential reasons for their performance.

Performance Investigation of ViTs. We investigate the few-shot learning performance of various ViTs, including LV-ViT [18] (single stage ViT), Le-ViT [13] (multi-stage ViT), Visformer [7] (CNN-enhanced ViT), Swin [29] and NesT [57] (locality-enhanced ViT). For a fair comparison, we scale the depth and width of these ViTs such that their model sizes are similar to ResNet-12 [15] ($\sim 12.5\text{M}$ parameters) which is the most commonly used architecture and achieves (nearly) state-of-the-art performance on few-shot classification tasks.

Table 1. Classification accuracy (%) of the Meta-Baseline few-shot learning framework using ViTs and ResNet-12 as a feature extractor on *miniImageNet*.

	Backbone	Params	Meta-Training Phase		Meta-Tuning Phase	
			5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
CNN	ResNet-12	12.5M	60.00±0.44	80.55±0.31	64.53±0.45	81.41±0.31
	LV-ViT	13.5M	43.08±0.38	59.03±0.39	43.07±0.39	58.87±0.39
	Le-ViT	12.6M	38.89±0.39	53.51±0.37	40.65±0.40	54.23±0.38
ViT	Swin	12.3M	48.26±0.42	65.72±0.38	54.63±0.45	70.60±0.38
	Visformer	12.4M	43.55±0.38	60.49±0.39	47.61±0.43	63.00±0.39
	NesT	12.6M	49.23±0.43	66.57±0.39	54.57±0.46	69.85±0.38

**Fig. 2.** Accuracy curve comparison of different feature extractors in Meta-Baseline.

Here we choose a simple yet effective Meta-Baseline [6] as the few-shot learning framework to test the performance of the above five ViTs. Specifically, in the meta-training phase, Meta-Baseline performs conventional supervised classification; in meta-tuning phase, given a training task, it computes the class prototypes by averaging the sample features from the same class on the support set, and then minimizes the cosine distance between the sample feature and the corresponding class prototypes. For inference, similar to meta-tuning phase, given a test task τ , it first calculates the class prototypes on the support set of τ , and assigns a class label to a test sample according to cosine distance between the test sample feature and the class prototypes. See more details in Sec. 4.2.

Table 1 summarizes the experimental results on *miniImageNet* [44]. Here, “Meta-Training Phase” means we test ViTs after meta-training phase in Meta-Baseline, while “Meta-Tuning Phase” denotes the test performance of ViTs after meta-tuning phase. Frustratingly, one can observe that even with the same Meta-Baseline framework, all ViTs perform much worse than ResNet-12 after meta-training phase. Moreover, after meta-tuning phase, all ViTs still suffer from much worse few-shot learning performance. These observations inspire us to investigate what happens during the meta-training period of ViTs and why ViTs perform much worse than CNNs.

Analysis. To analyze performance of ViTs in the meta-training phase, Fig. 2 plots accuracy curves on training and validation sets in base dataset \mathbb{D}_{base} , and reports the 5-way 1-shot accuracy on the training and test sets of the novel classes. According to Fig. 2(a) and 2(b), all ViTs converge well on training and validation sets in the base dataset \mathbb{D}_{base} . Fig. 2(c) shows that ViTs also converge well on the training data of novel categories during the whole meta-training phase, while Fig. 2(d) indicates that ViTs generalize poorly to the test

Table 2. Empirical analysis of ViTs on two datasets. (Top) shows the effect of inductive bias (IB) to few-shot accuracy of ViTs; (Middle) studies whether CNN-alike inductive bias (CaIB) can enhance token dependency learning; and (Bottom) analyzes local attention (LA) in enhancing token dependency learning.

Method	ViT	Var	<i>miniImageNet</i>		<i>tieredImageNet</i>	
			5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
(Top)		IB				
Meta-Baseline [6]	NesT		54.57±0.46	69.85±0.38	63.73±0.47	79.33±0.38
[6]+CNN [46]	NesT	✓	57.91±0.48	73.31±0.38	63.46±0.51	80.13±0.38
(Middle)		CaIB				
Meta-Baseline [6]	NesT		54.57±0.46	69.85±0.38	63.73±0.47	79.33±0.38
[6]+CNN Distill	NesT	✓	55.79±0.45	71.81±0.37	64.48±0.50	80.43±0.37
(Bottom)		LA				
Meta-Baseline [6]	LeViT		38.89±0.39	53.51±0.37	55.75±0.49	71.99±0.39
Meta-Baseline [6]	Swin	✓	54.63±0.45	70.60±0.38	62.68±0.50	78.52±0.38
Meta-Baseline [6]	NesT	✓	54.57±0.46	69.85±0.38	63.73±0.47	79.33±0.38

data of novel categories. For instance, though NesT [57] achieves $\sim 52\%$ accuracy in the first 30 meta-training epochs, its accuracy drops to 49.2% rapidly around the 30th meta-training epochs. In contrast, ResNet-12 maintains nearly 60% 1-shot accuracy in the last 100 meta-training epochs. These observations show that compared to CNN feature extractors, ViTs often suffer from generalization issues on novel categories though they enjoy good generalization ability on \mathbb{D}_{base} .

Next, we empirically analyze why ViTs perform worse than CNNs on few-shot classification. The most possible explanation is that the lack of inductive bias leads to the performance degeneration. To illustrate this point, we introduce inductive bias (in CNN) into ViT through three ways. **a)** For each stage, we use a ViT branch and a CNN branch to independently extract image features, and combine their features for fusion. See implementation details in the suppl.. In this way, this mechanism inherits the inductive bias from the CNN branch. Table 2 (Top) shows that by introducing this CNN inductive bias, the new NesT (i.e. “[6]+CNN”) largely surpasses the vanilla NesT (“Meta-Baseline [6]”). **b)** We train a CNN model on \mathbb{D}_{base} , and use it to teach ViT via knowledge distillation [16]. This method, i.e. “[6]+CNN Distill” in Table 2 (Middle), can well introduce the CNN-alike inductive bias into ViT to enhance dependency learning, and improves the vanilla “Meta-Baseline [6]” by a significant margin. To analyze the quality of token dependency, we visualize the attention maps of the last block of vanilla ViT and CNN-distilled ViT in Fig. 4(a). One can find that CNN-distilled ViT can learn better token dependency than vanilla ViT, since the former captures almost all semantic tokens while the later one only captures a few semantic patches. **c)** Local attention enjoys better few-shot learning performance. Table 2 (Bottom) shows that with the same Meta-Baseline [6] framework, Swin and NesT achieve much higher accuracy than LeViT, since local attention in Swin and NesT introduces (CNN-alike) inductive bias and can enhance token dependency learning than the global attention in LeViT. All these results show that inductive bias benefits few-shot classification.




Fig. 3. Visualization of attention maps from vanilla ViTs with different training epochs.




Fig. 4. (a) Attention maps from vanilla ViT and CNN-distilled ViT. (b) training accuracy curves between ViT and CNN-distilled ViT. (c&d) Attention score distance between ViT and CNN-distilled ViT at the middle layer (middle) and last layer (right).

Moreover, we also observe that the convergence speed is not equivalent to the quality of token dependency learning. In comparison to vanilla ViT, CNN-distilled ViT can learn higher-qualified token dependency as shown in Fig. 4(a), but has slower convergence speed and lower training accuracy as illustrated in Fig. 4(b). So one can conclude that the convergence speed actually does not reflect the quality of token dependency learning. Besides, for vanilla ViT, its training accuracy increases rapidly in the first 30 epochs but becomes saturated in the following epochs in Fig. 4(b), while its attention maps at the last block still evolve in Fig. 3. Indeed, Fig. 4(c) (and 4(d)) indicates that attention score distance descends stably for all training epochs. So these results all show that convergence speed does not well reflect the quality of token dependency learning.

4 Self-Promoted Supervision for Few-Shot Classification

From the observations and analysis in Sec. 3, one knows that 1) directly replacing the CNN feature extractor with ViT in existing few-shot learning frameworks often leads to severe performance degradation; and 2) the few-shot learning frameworks that boost token dependency learning to remedy lack of inductive bias can improve performance of ViT on few-shot learning tasks. Accordingly, here we propose a Self-promoted sUpervision (SUN) framework which aims to improve the few-shot learning ability of ViTs via enhancing and accelerating token dependency learning in ViTs via a dense location-specific supervision. Similar to conventional few-shot learning frameworks, e.g. Meta-Baseline, our SUN framework also consists of two phases, i.e., meta-training and meta-tuning which are illustrated in Fig. 5. In the following, we will introduce these phases.

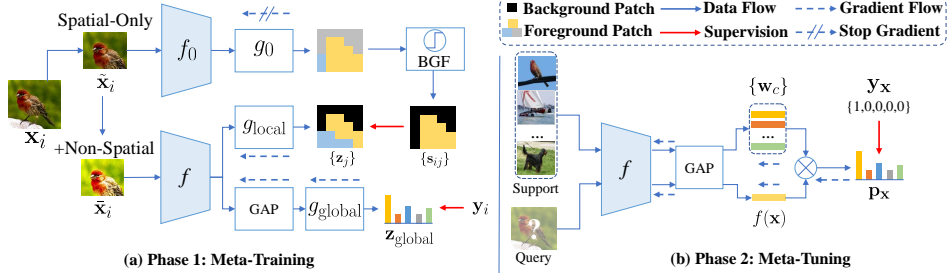


Fig. 5. Training pipeline of ViTs with self-promoted supervision (SUN). For meta-training phase (a), given an image \mathbf{x}_i , we use spatial-consistent augmentation to generate two crops $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{x}}_i$. Given $\tilde{\mathbf{x}}_i$, teacher (ViT f_0 & classifier g_0) with background filtration (BGF) generates its location-specific supervision $\{s_{ij}\}$. Given $\tilde{\mathbf{x}}_i$, the meta-learner f extracts its token features, and then global classifier g_{global} and local classifier g_{local} respectively predict the global semantic label \tilde{y}_i of $\tilde{\mathbf{x}}_i$ and the semantic labels $\{\tilde{s}_{ij}\}$ of all patches. Finally, SUN uses ground-truth labels y_i and the location-specific label $\{s_{ij}\}$ to supervise \tilde{y}_i and $\{\tilde{s}_{ij}\}$ for optimizing f and $g_{global/local}$. During meta-tuning phase (b), SUN adopts an existing few-shot method (e.g., Meta-Baseline [6]) to fine-tune f and then predicts the label for each query \mathbf{x} with a support set S .

4.1 Meta-Training

For meta-training phase in Fig. 5(a), similar to Meta-Baseline, our target is to learn a meta learner f mentioned in Sec. 3 such that f is able to fast adapt itself to new classes with a few training data. Here f is actually a feature extractor and is a ViT model in this work. The main idea of SUN is a dense location-specific supervision which aims to enhance and accelerate token dependency learning in ViT and thus boost the learning efficiency on the limited training data.

Specifically, to train the meta-learner f on the base dataset \mathbb{D}_{base} , SUN uses global supervision, i.e., the ground-truth label, to guide the global average token of all patch tokens in f for learning the global semantics of the whole image. More importantly, SUN further employs individual location-specific supervision to supervise each patch token in f . To generate patch-level pseudo labels as location-specific supervision, SUN first conducts the supervised classification pretraining to optimize a teacher model f_g consisting of a ViT f_0 with the same architecture as f and a classifier g_0 on the base dataset \mathbb{D}_{base} . This is reasonable, since teacher f_0 is trained on \mathbb{D}_{base} and its predicted patch-level pseudo labels are semantically consistent with the ground-truth label of the whole image, i.e., similar pseudo labels for similar local and class tokens. This is why our method is so called *self-promoted* since f_0 has the same architecture as f . Formally, for each training image \mathbf{x}_i in the base dataset \mathbb{D}_{base} , SUN first calculates the per-token classification confidence score $\hat{\mathbf{s}}_i$ as follows:

$$\hat{\mathbf{s}}_i = [\hat{s}_{i1}, \hat{s}_{i2}, \dots, \hat{s}_{iK}] = f_g(\mathbf{x}_i) = [g(\mathbf{z}_1), g(\mathbf{z}_2), \dots, g(\mathbf{z}_K)] \in \mathbb{R}^{c \times K}, \quad (1)$$

where $\mathbf{z} = f(\mathbf{x}_i) = [\mathbf{z}_{\text{cls}}, \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_K]$ denotes the class and patch tokens of image \mathbf{x}_i , and $\hat{\mathbf{s}}_{ij}$ is the pseudo label of the j -th local patch \mathbf{x}_{ij} in sample \mathbf{x}_i . Accordingly, for each patch \mathbf{x}_{ij} , the position with relative high confidence in $\hat{\mathbf{s}}_{ij}$ indicates that this patch contains the semantics of corresponding categories.

To fully exploit the power of local supervision, we further propose background filtration (BGF) technique which targets at classifying background patches into a new unique class and improving quality of patch-level supervision. This technique is necessary, since being trained on \mathbb{D}_{base} which has no background class, the teacher f_0 always wrongly assigns background patches into a base (semantic) class instead of background class and provides inaccurate location-specific supervision. To tackle this issue, we filtrate the local patches with very low confidence score as background patches and classify them into a new unique class. Specifically, given a batch of m patches with confidence score $\{\hat{\mathbf{s}}_{ij}\}$, we first select the maximum score for each patch, and then sort the m scores of patches by an ascending order. Next, we view the top $p\%$ patches with the lowest scores as background patches. Meanwhile, we add a new unique class, i.e., the background class, into the base classes C_{base} . Accordingly, we increase one dimension for the generated pseudo label $\hat{\mathbf{s}}_i \in \mathbb{R}^c$ to obtain $\mathbf{s}_{ij} \in \mathbb{R}^{c+1}$. For background patches, the last positions of their pseudo label $\mathbf{s}_{ij} \in \mathbb{R}^{c+1}$ are one and remaining positions are zero. For non-background patches, we only set the last positions of their pseudo label as zero and do not change other values. In this way, given one image \mathbf{x}_i with ground-truth label \mathbf{y}_i , we define its overall training loss:

$$\mathcal{L}_{\text{SUN}} = H(g_{\text{global}}(\mathbf{z}_{\text{global}}), \mathbf{y}_i) + \lambda \sum_{j=1}^K H(g_{\text{local}}(\mathbf{z}_j), \mathbf{s}_{ij}), \quad (2)$$

where $\mathbf{z} = f(\mathbf{x}_i) = [\mathbf{z}_{\text{cls}}, \mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_K]$ denotes the class and patch tokens of image \mathbf{x}_i , and $\mathbf{z}_{\text{global}}$ is a global average pooling of all patch tokens $\{\mathbf{z}_j\}_{j=1}^K$. Here H denotes the cross entropy loss, and g_{global} and g_{local} are two trainable classifiers for global semantic classification and local patch classification respectively. For λ , we set it as 0.5 in all experiments for simplicity.

Now we discuss the two benefits of the dense supervision on few-shot learning tasks. Firstly, all location-specific supervisions are generated by the teacher f_0 and thus guarantees similar pseudo labels for similar local tokens. This actually tells ViT which patch tokens are similar or dissimilar and thus can accelerate token dependency learning. Secondly, in contrast to the global supervision on the whole image, the location-specific supervisions are at a fine-grained level, namely the patch level, and thus help ViTs to easily discover the target objects and improve the recognition accuracy. This point is also consistent with some previous literature. For instance, works [60,18] demonstrate that modeling semantics in local tokens can avoid learning skewed and non-generalizable patterns and thus substantially improve the model generalization performance. Therefore, both aspects are heavily desired in few-shot learning problems and can increase the few-shot learning capacity of ViTs.

To further improve the robustness of local supervision from teacher f_0 while keeping sufficient data diversity, we propose a ‘‘Spatial-Consistent Augmentation’’ (SCA) strategy to improve generalization. SCA consists of a *spatial-only*

augmentation and a *non-spatial* augmentation. In spatial-only augmentation, we only introduce spatial transformation (e.g., random crop and resize, flip and rotation) to augment the input images \mathbf{x}_i and obtain $\tilde{\mathbf{x}}_i$. For non-spatial augmentation, it only leverages non-spatial augmentations, e.g. color jitter, on $\tilde{\mathbf{x}}_i$ to obtain $\bar{\mathbf{x}}_i$. During the meta-training phase, we feed $\tilde{\mathbf{x}}_i$ into the teacher f_0 to generate \mathbf{s}_{ij} used in Eqn. (2), and feed $\bar{\mathbf{x}}_i$ into the target meta-learner f . In this way, the samples used to train meta-learner f are of high diversity, while still enjoying very accurate location-specific supervisions \mathbf{s}_{ij} since the teacher f_0 uses a weak augmentation $\tilde{\mathbf{x}}_i$ to generate the location-specific supervisions. This also helps improve the generalization ability of ViTs.

4.2 Meta-Tuning

For meta-tuning phase in Fig. 5(b), our target is to finetune the meta-learner f via training it on multiple “ N -way K -shot” tasks $\{\tau\}$ sampled from the base dataset \mathbb{D}_{base} , such that f can be adapted to a new task which contains unseen classes C_{novel} (i.e. $C_{\text{novel}} \notin \mathbb{D}_{\text{base}}$) with only a few labeled training samples. To this end, without loss of generality, we follow a simple yet effective meta-tuning method in Meta Baseline [6]. Besides, we also investigate different meta-tuning methods, such as FEAT [51] and DeepEMD [55], in the suppl.. With the same FEAT and DeepEMD, our SUN framework still shows superiority over other few-shot learning frameworks. See more details in the suppl.. For completeness, we introduce the meta-tuning in Meta-Baseline in the following. Specifically, given a task τ with support set \mathbf{S} , SUN calculates the classification prototype \mathbf{w}_k of class k via $\mathbf{w}_k = \sum_{\mathbf{x} \in \mathbf{S}_k} \text{GAP}(f(\mathbf{x})) / |\mathbf{S}_k|$, where \mathbf{S}_k denotes the support samples from class c and GAP means global average pooling operation. Then for each query image \mathbf{x} , meta-learner f calculates the classification confidence score of the k -th class:

$$\mathbf{p}_k = \frac{\exp(\gamma \cdot \cos(\text{GAP}(f(\mathbf{x})), \mathbf{w}_k))}{\sum_{k'} \exp(\gamma \cdot \cos(\text{GAP}(f(\mathbf{x})), \mathbf{w}_{k'}))}, \quad (3)$$

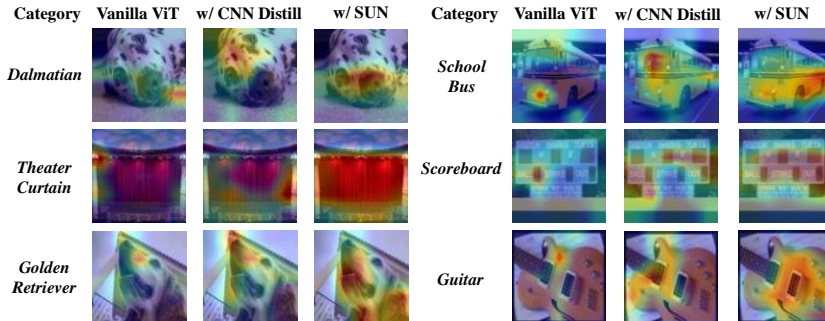
where \cos denotes cosine similarity and γ is a temperature parameter. Finally, it minimizes the cross-entropy loss $\mathcal{L}_{\text{few-shot}} = H(\mathbf{p}_{\mathbf{x}}, \mathbf{y}_{\mathbf{x}})$ to fine-tune meta-learner f on the various sampled tasks $\{\tau\}$, where $\mathbf{p}_{\mathbf{x}} = [\mathbf{p}_1, \dots, \mathbf{p}_c]$ is the prediction and $\mathbf{y}_{\mathbf{x}}$ is ground-truth label of \mathbf{x} . After this meta-tuning, given a new test task τ' with support set \mathbf{S}' , we follow the above step to compute its classification prototypes, and then use Eqn. (3) to predict the labels of the test samples.

5 Experiments

Following [5,6,63], we evaluate SUN with various ViTs on three widely used few-shot benchmarks, i.e., *miniImageNet* [44], *tieredImageNet* [34] and *CIFAR-FS* [2] whose details are deferred to the supplement. For fairness, for each benchmark, we follow [5,6,63] and report the average accuracy with 95% confidence interval on 2,000 tasks under 5-way 1-shot and 5-way 5-shot classification settings. More training details can be found in the suppl..

Table 3. Comparison between SUN and Meta Baseline on *miniImageNet*.

Backbone	Params	Meta-Baseline [5]		SUN (Ours)	
		5-way 1-shot	5-way 5-shot	5-way 1-shot	5-way 5-shot
LV-ViT	12.6M	43.08±0.38	59.03±0.39	59.00±0.44	75.47±0.34
Swin	12.6M	54.63±0.45	70.60±0.38	64.94±0.46	80.40±0.32
Visformer	12.5M	47.61±0.43	63.00±0.39	67.80±0.45	83.25±0.29
NesT	12.8M	54.57±0.46	69.85±0.38	66.54±0.45	82.09±0.30

**Fig. 6.** Visualization of attention maps from vanilla ViT, CNN-distilled ViT and ViT with SUN. ViT with SUN performs better, thus obtains better token dependency.

5.1 Comparison on Different ViTs

We evaluate our SUN on four different ViTs, i.e., LV-ViT, Swin, Visformer and NesT, which cover most of existing ViT types. Table 3 shows that on *miniImageNet*, SUN significantly surpasses Meta-Baseline [6] on the four ViTs. Specifically, it makes 15.9%, 10.3%, 20.2%, and 12.0% improvement over Meta-Baseline under 5-way 1-shot setting. These results demonstrate the superior generalization ability of our SUN on novel categories, and also its good compatibility with different kinds of ViTs. Besides, among all ViTs, Visformer achieves the best performance, since it uses many CNN modules and can better handle the few-shot learning task. In the following Sec. 5.2 and 5.4, we use the second best NesT as our ViT feature extractor, since NesT does not involve CNN modules and only introduce modules to focus on local features and thus better reveals few-shot learning ability of conventional ViT architecture. And in Sec. 5.3, we use both NesT and Visformer to make comparisons with the state-of-the-art CNN-based few-shot classification methods.

SUN v.s. ViT in dependency learning. Following the analysis in Sec. 3, we visualize the attention maps of vanilla ViT, CNN-distilled ViT and ViT with SUN to compare the quality of token dependency learning. Fig. 6 demonstrate the comparison result. ViT with SUN can capture more semantic tokens than vanilla ViT and CNN-distilled ViT on various categories. The result indicates that ViT with SUN can learn better token dependency during the training.

Table 4. Comparison of different ViT based few-shot classification frameworks on *miniImageNet*. “AC” means introducing CNN modules in ViT backbone to fuse features. “Distill” is knowledge distillation. “FT” is using meta-tuning to finetune meta learner (see Sec. 4.2). Except [22], all methods use NesT as ViT feature extractor.

Method	Backbone	AC	Distill	FT	<i>miniImageNet</i>	
					5-way 1-shot	5-way 5-shot
Meta-Baseline [6]	<i>ResNet-12</i>			✓	64.53±0.45	81.41±0.31
Baseline [5]	<i>ViT</i>				49.23±0.43	66.57±0.39
Meta-Baseline [6]	<i>ViT</i>			✓	54.57±0.46	69.85±0.38
Re-Parameter [22]	<i>ViT</i>	✓		✓	46.59±0.44	62.56±0.41
[6]+CNN-Distill [46]	<i>ViT</i>	✓	✓	✓	55.79±0.45	71.81±0.37
Semiformer [46]	<i>CNN+ViT</i>	✓	✓		56.62±0.46	72.91±0.39
[6]+Semiformer [46]	<i>CNN+ViT</i>	✓	✓	✓	57.91±0.48	73.31±0.38
[6]+DrLoc [28]	<i>ViT</i>			✓	57.85±0.46	74.03±0.35
BML [63]	<i>ViT</i>		✓		59.35±0.45	76.00±0.35
SUN (Ours)	<i>ViT</i>		✓	✓	66.54±0.45	82.09±0.30

5.2 Comparison among Different Few-shot Learning Frameworks

In Table 4, we compare SUN with other few-shot learning frameworks on the same ViT architecture. In Table 4, the method “[6]+CNN-Distill” pretrains a CNN to distill a ViT; “Semiformer” [46] introduces CNN modules into ViT; “Re-Parameter” [22] initializes a ViT by a trained CNN; “[6]+DrLoc” predicts the relative positions between local token pairs; “BML” [63] introduces mutual learning based few-shot learning framework. All these methods propose their own techniques to improve few-shot learning performance of ViT via introducing CNN-alike inductive bias or constructing tasks to better learn token dependency. Except [22] that uses its CNN-transformer-mixed architecture, all others adopt NesT [57] as a ViT feature extractor or use ResNet-12 as the CNN feature extractor. Table 4 shows that our SUN significantly surpasses all these methods in terms of accuracy on novel categories. Specifically, our SUN outperforms the runner-up by 7.19% and 6.09% under 5-way 1-shot and 5-shot settings, respectively. These results show that even on the same ViT, SUN achieves much better few-shot accuracy on novel categories than other few-shot learning frameworks.

5.3 Comparison with State-of-The-Arts

Here we compare SUN with state-of-the-arts (SoTAs), including CNN based methods and ViT based one, on *miniImageNet* [44], *tieredImageNet* [34] and CIFAR-FS [2]. Table 5 reports the evaluation results. Without introducing extremely complex few-shot learning methods like [48,19], our SUN achieves comparable performance with SoTA on *miniImageNet* [44], and sets new SoTAs on *tieredImageNet* [34] and CIFAR-FS [2]. Specifically, on *tieredImageNet* under 5-way 1-shot and 5-shot settings, our SUN (Visformer) respectively obtains 72.99% and 86.74%, and respectively improves $\sim 0.8\%$ and $\sim 0.2\%$ over the SoTA TPMN [48]. On CIFAR-FS dataset, our SUN (NesT) obtains 78.17% and 88.98% in terms of 1-shot accuracy and 5-shot accuracy, which significantly outper-

Table 5. Comparison with SoTA few-shot learning methods under 5-way few-shot classification setting. The results of the best 2 methods are in bold font.

Method	Classifier Params	<i>miniImageNet</i>		<i>tieredImageNet</i>		CIFAR-FS	
		1-shot	5-shot	1-shot	5-shot	1-shot	5-shot
ResNet-12/18 as feature extractor							
MetaOptNet [20]	0	64.09±0.62	80.00±0.45	65.81±0.74	81.75±0.53	72.00±0.70	84.20±0.50
DeepEMD [55]	0	65.91±0.82	82.41±0.56	71.16±0.80	86.03±0.58	46.47±0.70	63.22±0.71
FEAT [51]	1.05M	66.78±0.20	82.05±0.14	70.80±0.23	84.79±0.16	-	-
TADAM [32]	1.23M	58.50±0.30	76.70±0.30	-	-	-	-
Rethink-Distill [39]	225K	64.82±0.60	82.14±0.43	71.52±0.69	86.03±0.49	73.90±0.80	86.90±0.50
DC [24]	224K	61.26±0.20	79.01±0.13	-	-	-	-
CloserLook++ [5]	131K	51.87±0.77	75.68±0.63	-	-	-	-
Meta-Baseline [6]	0	63.17±0.23	79.26±0.17	68.62±0.27	83.29±0.18	-	-
Neg-Cosine [25]	131K	63.85±0.81	81.57±0.56	-	-	-	-
AFHN [21]	359K	62.38±0.72	78.16±0.56	-	-	68.32±0.93	81.45±0.87
Centroid [1]	10K	59.88±0.67	80.35±0.73	69.29±0.56	85.97±0.49	-	-
RE-Net [19]	430K	67.60±0.44	82.58±0.30	71.61±0.51	85.28±0.35	74.51±0.46	86.60±0.32
TPMN [48]	16M	67.64±0.63	83.44±0.43	72.24±0.70	86.55±0.63	75.50±0.90	87.20±0.60
NesT ViT as feature extractor							
CloserLook++ [5]	180K	49.23±0.43	66.57±0.39	59.13±0.46	77.88±0.39	63.89±0.49	80.43±0.37
Meta-Baseline [6]	0	54.57±0.46	69.85±0.38	63.73±0.47	79.33±0.38	68.05±0.48	81.53±0.36
BML [63]	180K	59.35±0.45	76.00±0.35	66.98±0.50	83.75±0.34	67.51±0.48	82.17±0.36
SUN	0	66.54±0.45	82.09±0.30	72.93±0.50	86.70±0.33	78.17±0.46	88.98±0.33
Visformer ViT as feature extractor							
SUN	0	67.80±0.45	83.25±0.30	72.99±0.50	86.74±0.33	78.37±0.46	88.84±0.32

forms all the state-of-the-art methods by at least 2.5% in terms of 1-shot accuracy. Meanwhile, our SUN (Visformer) also obtains 67.80% 1-shot accuracy on *miniImageNet test* set and surpasses the SoTA CNN-based few-shot classification methods by $\sim 1.2\%$. These results well demonstrate the effectiveness of ViTs with our SUN in few-shot learning.

5.4 Ablation Study

Here we show the key ablation studies, more details can be found in the suppl.. **Effect of each Phase in SUN.** We investigate the effect of each training phase in SUN. For fairness, we always use the same ViT feature extractor [57] as meta-learner f for few-shot classification. Table 6 shows that on *miniImageNet* dataset, compared to the baseline (denoted by “Base” in Table 6), the teacher ViT f_g of meta-training phase significantly outperforms the baseline by 13.2% and 12.9% in terms of 1-shot and 5-shot accuracy. By introducing the location-specific supervision, type (b) obtains $\sim 0.7\%$ improvement for 1-shot accuracy. Then after introducing SCA and BGF, type (d) can further lead to a gain of $\sim 1.8\%$ 1-shot accuracy as well as $\sim 1.0\%$ 5-shot accuracy. A possible reason is that the SCA improves the quality of location-specific supervision from the pre-trained transformer while keeping the strong augmentation ability for the target transformer in the meta-training phase. Meanwhile, BGF eliminates the negative effect from mislabeling background patches and improves the feature representation quality. These two techniques significantly improve location-specific supervision and improve the generalization ability of ViT feature extractors.

We also investigate meta-tuning phase. As shown in type (e) in Table 6, the meta-learner f after meta-training phase can be naturally adopted to the

Table 6. Ablation study of SUN on *miniImageNet*. “Local” means location-specific supervision. SCA is spatial-consistent augmentation. BGF means background filtration. And f_g means the teacher ViT model in SUN meta-training phase.

Type	Meta-Training				Meta-Tuning	<i>miniImageNet</i>	
	f_g	Local	SCA	BGF	Meta-Baseline [6]	5-way 1-shot	5-way 5-shot
Base						49.23±0.43	66.57±0.39
(a)	✓					62.40±0.44	79.45±0.32
(b)	✓	✓				63.06±0.45	79.91±0.32
(c)	✓	✓	✓			64.50±0.45	80.59±0.31
(d)	✓	✓	✓	✓		64.84±0.45	80.96±0.32
(e)	✓	✓	✓	✓	✓	66.54±0.45	82.09±0.30

Table 7. Ablation of meta-tuning, where “SUN-*” means ours and others are SoTAs.

Methods	MetaBaseline[6]	FEAT[51]	DeepEMD[55]	COSOC[30]	SUN-M	SUN-F	SUN-D
5-way 1-shot	64.53±0.45	66.78±0.20	68.77±0.29	69.28±0.49	67.80±0.45	66.90±0.44	69.56±0.44
5-way 5-shot	81.41±0.31	82.05±0.14	84.13±0.53	85.16±0.42	83.25±0.30	82.63±0.30	85.38±0.49

existing few-shot learning methods and obtain further improvement. Specifically, our SUN (denoted by (e)) also surpasses f after meta-training phase by 1.7% and 1.1% in terms of 1-shot and 5-shot accuracy.

Meta-Tuning Methods. We also evaluate our SUN with different meta-tuning methods. To empirically analyze our SUN, we fix Visformer [7] as ViT backbone and choose Meta-Baseline [6], FEAT [51] and a dense prediction method DeepEMD [55]. As shown in Table 7, all SUN-M/F/D improve the accuracy and perform better than corresponding CNN-based baselines. Moreover, SUN-D achieves 69.56% 5-way 1-shot accuracy, even slightly surpasses the SoTA COSOC [30].

6 Conclusion

In this work, we first empirically showed that replacing the CNN feature extractor as a ViT model leads to severe performance degradation on few-shot learning tasks, and the slow token dependency learning on limited training data largely contributes to this performance degradation. Then for the first time, we proposed a simple yet effective few-shot training framework for ViTs, i.e. Self-promoted sUpervision (SUN), to resolve this issue. By firstly pretraining the ViT on the few-shot learning dataset, SUN adopts it to generate individual location-specific supervision for 1) guiding the ViT on which patch tokens are similar or dissimilar and thus accelerating token dependency learning; 2) improving the object grounding and recognition ability. Experimental results on few-shot classification tasks demonstrate the superiority of SUN in accuracy over state-of-the-arts.

Acknowledgement This work was supported in part by the National Key R&D Program of China under Grant No. 2021ZD0112100, and the Major Key Project of PCL under Grant No. PCL2021A12.

References

1. Afrasiyabi, A., Lalonde, J.F., Gagné, C.: Associative alignment for few-shot image classification. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J.M. (eds.) *Computer Vision – ECCV 2020*. pp. 18–35. Springer International Publishing, Cham (2020)
2. Bertinetto, L., Henriques, J.F., Torr, P., Vedaldi, A.: Meta-learning with differentiable closed-form solvers. In: *International Conference on Learning Representations (2019)*, <https://openreview.net/forum?id=HyxnZh0ct7>
3. Cao, Y.H., Yu, H., Wu, J.: Training vision transformers with only 2040 images (2022)
4. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: Vedaldi, A., Bischof, H., Brox, T., Frahm, J. (eds.) *Computer Vision - ECCV 2020 - 16th European Conference, Glasgow, UK, August 23-28, 2020, Proceedings, Part I. Lecture Notes in Computer Science*, vol. 12346, pp. 213–229. Springer (2020). https://doi.org/10.1007/978-3-030-58452-8_13, https://doi.org/10.1007/978-3-030-58452-8_13
5. Chen, W.Y., Liu, Y.C., Kira, Z., Wang, Y.C., Huang, J.B.: A closer look at few-shot classification. In: *International Conference on Learning Representations (2019)*
6. Chen, Y., Liu, Z., Xu, H., Darrell, T., Wang, X.: Meta-baseline: Exploring simple meta-learning for few-shot learning. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 9062–9071 (October 2021)
7. Chen, Z., Xie, L., Niu, J., Liu, X., Wei, L., Tian, Q.: Visformer: The vision-friendly transformer. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 589–598 (October 2021)
8. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *2009 IEEE conference on computer vision and pattern recognition*. pp. 248–255. Ieee (2009)
9. Doersch, C., Gupta, A., Zisserman, A.: Crosstransformers: spatially-aware few-shot transfer. In: Larochelle, H., Ranzato, M., Hadsell, R., Balcan, M.F., Lin, H. (eds.) *Advances in Neural Information Processing Systems*. vol. 33, pp. 21981–21993. Curran Associates, Inc. (2020), <https://proceedings.neurips.cc/paper/2020/file/fa28c6cdf8dd6f41a657c3d7caa5c709-Paper.pdf>
10. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Deghani, M., Minderer, M., Heigold, G., Gelly, S., Uszkoreit, J., Houlsby, N.: An image is worth 16x16 words: Transformers for image recognition at scale. In: *International Conference on Learning Representations (2021)*
11. Fang, Y., Liao, B., Wang, X., Fang, J., Qi, J., Wu, R., Niu, J., Liu, W.: You only look at one sequence: Rethinking transformer in vision through object detection. *Advances in Neural Information Processing Systems 34 pre-proceedings (2021)*
12. Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: *Proceedings of the 34th International Conference on Machine Learning (2017)*
13. Graham, B., El-Nouby, A., Touvron, H., Stock, P., Joulin, A., Jégou, H., Douze, M.: Levit: A vision transformer in convnet’s clothing for faster inference. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 12259–12269 (October 2021)
14. Graves, A., Wayne, G., Danihelka, I.: Neural turing machines. *arXiv preprint arXiv:1410.5401 (2014)*
15. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 770–778 (2016)

16. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network. In: NIPS Deep Learning and Representation Learning Workshop (2015), <http://arxiv.org/abs/1503.02531>
17. Hou, R., Chang, H., MA, B., Shan, S., Chen, X.: Cross attention network for few-shot classification. In: Wallach, H., Larochelle, H., Beygelzimer, A., d'Alché-Buc, F., Fox, E., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 32. Curran Associates, Inc. (2019), <https://proceedings.neurips.cc/paper/2019/file/01894d6f048493d2cacde3c579c315a3-Paper.pdf>
18. Jiang, Z., Hou, Q., Yuan, L., Daquan, Z., Shi, Y., Jin, X., Wang, A., Feng, J.: All tokens matter: Token labeling for training better vision transformers. In: *Advances in Neural Information Processing Systems* (2021)
19. Kang, D., Kwon, H., Min, J., Cho, M.: Relational embedding for few-shot classification. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 8822–8833 (October 2021)
20. Lee, K., Maji, S., Ravichandran, A., Soatto, S.: Meta-learning with differentiable convex optimization. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* pp. 10649–10657 (2019)
21. Li, K., Zhang, Y., Li, K., Fu, Y.: Adversarial feature hallucination networks for few-shot learning. In: *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
22. Li, S., Chen, X., He, D., Hsieh, C.J.: Can vision transformers perform convolution? (2021)
23. Li, Y., Zhang, K., Cao, J., Timofte, R., Van Gool, L.: Localvit: Bringing locality to vision transformers. *arXiv preprint arXiv:2104.05707* (2021)
24. Lifchitz, Y., Avrithis, Y., Picard, S., Bursuc, A.: Dense classification and implanting for few-shot learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2019)
25. Liu, B., Cao, Y., Lin, Y., Li, Q., Zhang, Z., Long, M., Hu, H.: Negative margin matters: Understanding margin in few-shot classification. In: *European Conference on Computer Vision*. pp. 438–455 (2020)
26. Liu, H., Dai, Z., So, D., Le, Q.: Pay attention to mlps. *Advances in Neural Information Processing Systems* **34** (2021)
27. Liu, L., Hamilton, W.L., Long, G., Jiang, J., Larochelle, H.: A universal representation transformer layer for few-shot image classification. In: *International Conference on Learning Representations* (2021), <https://openreview.net/forum?id=04cII6MumYV>
28. Liu, Y., Sangineto, E., Bi, W., Sebe, N., Lepri, B., Nadai, M.D.: Efficient training of visual transformers with small datasets. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems* (2021), <https://openreview.net/forum?id=SCN8UaetXx>
29. Liu, Z., Lin, Y., Cao, Y., Hu, H., Wei, Y., Zhang, Z., Lin, S., Guo, B.: Swin transformer: Hierarchical vision transformer using shifted windows. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 10012–10022 (October 2021)
30. Luo, X., Wei, L., Wen, L., Yang, J., Xie, L., Xu, Z., Tian, Q.: Rectifying the short-cut learning of background for few-shot learning. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) *Advances in Neural Information Processing Systems* (2021), <https://openreview.net/forum?id=N1i6BJzouX4>
31. Mishra, N., Rohaninejad, M., Chen, X., Abbeel, P.: A simple neural attentive meta-learner. In: *International Conference on Learning Representations* (2018), <https://openreview.net/forum?id=B1DmUzWAW>

32. Oreshkin, B., Rodríguez López, P., Lacoste, A.: Tadam: Task dependent adaptive metric for improved few-shot learning. In: Bengio, S., Wallach, H., Larochelle, H., Grauman, K., Cesa-Bianchi, N., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*. vol. 31. Curran Associates, Inc. (2018), <https://proceedings.neurips.cc/paper/2018/file/66808e327dc79d135ba18e051673d906-Paper.pdf>
33. Ravi, S., Larochelle, H.: Optimization as a model for few-shot learning. In: *ICLR* (2017)
34. Ren, M., Ravi, S., Triantafillou, E., Snell, J., Swersky, K., Tenenbaum, J.B., Larochelle, H., Zemel, R.S.: Meta-learning for semi-supervised few-shot classification. In: *International Conference on Learning Representations* (2018), <https://openreview.net/forum?id=HJcSzz-CZ>
35. Rusu, A.A., Rao, D., Sygnowski, J., Vinyals, O., Pascanu, R., Osindero, S., Hadsell, R.: Meta-learning with latent embedding optimization. In: *International Conference on Learning Representations* (2019), <https://openreview.net/forum?id=BJgklhAcK7>
36. Snell, J., Swersky, K., Zemel, R.: Prototypical networks for few-shot learning. In: *Proceedings of the 31st International Conference on Neural Information Processing Systems*. pp. 4080–4090 (2017)
37. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 7262–7272 (October 2021)
38. Sun, C., Shrivastava, A., Singh, S., Gupta, A.: Revisiting unreasonable effectiveness of data in deep learning era. In: *Proceedings of the IEEE international conference on computer vision*. pp. 843–852 (2017)
39. Tian, Y., Wang, Y., Krishnan, D., Tenenbaum, J.B., Isola, P.: Rethinking few-shot image classification: a good embedding is all you need? *ECCV* (2020)
40. Tolstikhin, I.O., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Steiner, A., Keysers, D., Uszkoreit, J., et al.: Mlp-mixer: An all-mlp architecture for vision. *Advances in Neural Information Processing Systems* **34** (2021)
41. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers and distillation through attention. In: *International Conference on Machine Learning*. vol. 139, pp. 10347–10357 (July 2021)
42. Touvron, H., Cord, M., Sablayrolles, A., Synnaeve, G., Jégou, H.: Going deeper with image transformers. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 32–42 (October 2021)
43. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L., Polosukhin, I.: Attention is all you need. In: *Advances in neural information processing systems*. pp. 5998–6008 (2017)
44. Vinyals, O., Blundell, C., Lillicrap, T., Wierstra, D., et al.: Matching networks for one shot learning. In: *Advances in Neural Information Processing Systems*. pp. 3630–3638 (2016)
45. Wang, W., Xie, E., Li, X., Fan, D.P., Song, K., Liang, D., Lu, T., Luo, P., Shao, L.: Pyramid vision transformer: A versatile backbone for dense prediction without convolutions. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 568–578 (October 2021)
46. Weng, Z., Yang, X., Li, A., Wu, Z., Jiang, Y.G.: Semi-supervised vision transformers (2021)

47. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 22–31 (October 2021)
48. Wu, J., Zhang, T., Zhang, Y., Wu, F.: Task-aware part mining network for few-shot learning. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8433–8442 (October 2021)
49. Xiao, T., Dollar, P., Singh, M., Mintun, E., Darrell, T., Girshick, R.: Early convolutions help transformers see better. In: Beygelzimer, A., Dauphin, Y., Liang, P., Vaughan, J.W. (eds.) Advances in Neural Information Processing Systems (2021)
50. Xu, W., Xu, Y., Chang, T., Tu, Z.: Co-scale conv-attentional image transformers. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 9981–9990 (October 2021)
51. Ye, H.J., Hu, H., Zhan, D.C., Sha, F.: Few-shot learning via embedding adaptation with set-to-set functions. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 8808–8817 (2020)
52. Yu, W., Luo, M., Zhou, P., Si, C., Zhou, Y., Wang, X., Feng, J., Yan, S.: Metaformer is actually what you need for vision. arXiv preprint arXiv:2111.11418 (2021)
53. Yuan, L., Chen, Y., Wang, T., Yu, W., Shi, Y., Jiang, Z.H., Tay, F.E., Feng, J., Yan, S.: Tokens-to-token vit: Training vision transformers from scratch on imagenet. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 558–567 (October 2021)
54. Yuan, Y., Fu, R., Huang, L., Lin, W., Zhang, C., Chen, X., Wang, J.: Hrformer: High-resolution transformer for dense prediction. NeurIPS (2021)
55. Zhang, C., Cai, Y., Lin, G., Shen, C.: Deepemd: Few-shot image classification with differentiable earth mover’s distance and structured classifiers. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (June 2020)
56. Zhang, H., Duan, J., Xue, M., Song, J., Sun, L., Song, M.: Bootstrapping vits: Towards liberating vision transformers from pre-training (2021)
57. Zhang, Z., Zhang, H., Zhao, L., Chen, T., , Arik, S.O., Pfister, T.: Nested hierarchical transformer: Towards accurate, data-efficient and interpretable visual understanding. In: AAAI Conference on Artificial Intelligence (AAAI) (2022)
58. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., Zhang, L.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6881–6890 (June 2021)
59. Zhmoginov, A., Sandler, M., Vladymyrov, M.: Hypertransformer: Model generation for supervised and semi-supervised few-shot learning (2022)
60. Zhong, Z., Zheng, L., Kang, G., Li, S., Yang, Y.: Random erasing data augmentation. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 13001–13008 (2020)
61. Zhou, P., Yuan, X., Xu, H., Yan, S., Feng, J.: Efficient meta learning via minibatch proximal update. In: Neural Information Processing Systems (2019)
62. Zhou, P., Zou, Y., Yuan, X.T., Feng, J., Xiong, C., Hoi, S.: Task similarity aware meta learning: Theory-inspired improvement on maml. In: Uncertainty in Artificial Intelligence. pp. 23–33. PMLR (2021)
63. Zhou, Z., Qiu, X., Xie, J., Wu, J., Zhang, C.: Binocular mutual learning for improving few-shot classification. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 8402–8411 (October 2021)