

MBZUAI

Digital.Commons@MBZUAI

Computer Vision Faculty Publications

Scholarly Works

7-25-2022

Self-Distilled Vision Transformer for Domain Generalization

Maryam Sultana

Muzammal Naseer

Muhammad Haris Khan

Salman Khan

Fahad Shahbaz Khan

Follow this and additional works at: <https://dclibrary.mbzuai.ac.ae/cvfp>



Part of the [Artificial Intelligence and Robotics Commons](#)

Preprint: arXiv

Archived with thanks to arXiv

Preprint License: CC by 4.0

Uploaded 25 August 2022

Self-Distilled Vision Transformer for Domain Generalization

Maryam Sultana¹, Muzammal Naseer^{1,2}, Muhammad Haris Khan¹, Salman Khan^{1,2}, and Fahad Shahbaz Khan^{1,3}

¹ Mohamed Bin Zayed University of AI, UAE,

² Australian National University, AU,

³ Linköping University, Sweden,

maryam.sultana, muzammal.naseer, muhammad.haris, salman.khan, fahad.khan,
@mbzuai.ac.ae


Abstract. In recent past, several domain generalization (DG) methods have been proposed, showing encouraging performance, however, almost all of them build on convolutional neural networks (CNNs). There is little to no progress on studying the DG performance of vision transformers (ViTs), which are challenging the supremacy of CNNs on standard benchmarks, often built on i.i.d assumption. This renders the real-world deployment of ViTs doubtful. In this paper, we attempt to explore ViTs towards addressing the DG problem. Similar to CNNs, ViTs also struggle in out-of-distribution scenarios and the main culprit is overfitting to source domains. Inspired by the modular architecture of ViTs, we propose a simple DG approach for ViTs, coined as *self-distillation for ViTs*. It reduces the overfitting to source domains by easing the learning of input-output mapping problem through curating non-zero entropy supervisory signals for intermediate transformer blocks. Further, it does not introduce any new parameters and can be seamlessly plugged into the modular composition of different ViTs. We empirically demonstrate notable performance gains with different DG baselines and various ViT backbones in five challenging datasets. Moreover, we report favorable performance against recent state-of-the-art DG methods. Our code along with pre-trained models are publicly available at: <https://github.com/maryam089/SDViT>.

Keywords: Domain Generalization · Vision Transformers · Self Distillation.

1 Introduction

Since their inception, transformers have displayed remarkable performance in various natural language processing (NLP) tasks [1–3]. Owing to their success in NLP, recently, transformer design has been adopted for vision tasks [4]. Since then, we have been witnessing several vision transformer (ViT) models for image recognition [4], object detection [5, 6] and semantic segmentation [7, 8]. ViTs are intrinsically different in design compared to convolution neural networks (CNNs), since they lack explicit inductive biases such as, spatial connectivity and

Fig. 1: In-domain (validation) and out-of-domain (target) classification accuracy of ERM-CNN and ERM-ViT in four DG datasets. Similar to ERM-CNN, ERM-ViT also shows performance degradation in out-of-domain scenarios.



translation equivariance. They process (input) image as a sequence of patches that is enhanced via successive transformer blocks (comprised of self-attention mechanisms), thereby allowing the network to model relationship between any parts of the image. A useful consequence of such a processing is a wide receptive field that facilitates capturing the global context in contrast to a limited receptive field modelled in CNNs.

Many deep learning models are usually deployed in real-world scenarios where the test data is unknown in advance. When their predictions are used for decision making in safety-critical applications, such as medical diagnosis or self-driving car, an erroneous prediction can lead to dangerous consequences. This typically occurs because there is a distributional gap between the training and testing data. Hence, it is critical for deep learning models to provide reliable predictions that generalize across the different domains. Domain generalization (DG) is a problem setting in which data from multiple source domains is leveraged for training to generalize to a new (unseen) domain [9–18]. Existing DG methods aim to explicitly reduce domain gap in the feature space [9, 19, 20], learn well-transferable model parameters through meta-learning [21–24], propose different data augmentation techniques [15, 25–27], or leverage auxiliary tasks [12, 28]. Lately, Gulrajani and Lopez-Paz [13] show that a simple Empirical Risk Minimization (ERM) method obtains favourable performance against previous methods under a fair evaluation protocol termed as “Domainbed”. To our knowledge, almost all aforementioned DG approaches are based on CNNs, and there is little to no work on studying the DG performance of ViTs. So, in effect, despite ViTs demonstrating state-of-the-art performance on some standard benchmarks, often rooted in i.i.d assumption, their real-world deployment remain doubtful. To this end, we attempt to explore ViTs towards addressing DG problem.

We note that, ViT-based ERM (ERM-ViT), similar to its CNN-based counterpart (ERM-CNN), also suffers from performance degradation when facing out-of-distribution (OOD) target domain data (see Fig. 1). In the absence of any explicit overfitting prevention mechanism coupled with the one-hot encoded ground-truth supervision, which are essentially zero-entropy signals, as such it is challenging for a simple ERM-ViT model to obtain favourable OOD generalization. Under the hood, since the mapping problem is difficult, the model is prone to inadvertently exploiting non-generalizable, brittle features for making predictions. Fig. 2 visualizes attention maps from ERM-ViT on arbitrary images of four target domains in PACS dataset. ERM-ViT has the tendency to rely on




Fig. 2: Attention maps from ERM-ViT (top) and ERM-SDViT (bottom) corresponding to images of the target domains in PACS. ERM-ViT is prone to exploiting non-generalizable features e.g., background. Whereas ERM-SDViT is capable of learning cross-domain generalizable features e.g., object shape and its semantics.

non-object related features such as the background features, which are potentially non-transferable between the source and the target domains.

Inspired by the modular architecture of ViTs, we propose a light-weight plug-and-play DG approach for ViTs, namely *self-distillation for ViT (SDViT)*. It explicitly encourages the model towards learning generalizable, comprehensive features. ViTs process a sequence of input image patches repeatedly by multiple multi-headed self-attention layers, a.k.a transformer blocks [1]. These image patches are also known as *patch tokens*. A randomly initialized class token is usually appended to the set of image patches (tokens). This group is then passed through a sequence of transformer blocks followed by the passing of class token through a linear classifier to get final predictions. The class token can learn information that is useful while making final prediction. So, it can be extracted from the output of each transformer block and can be leveraged to get class-specific logits using the final classifier of the pretrained model [29]. Armed with this insight, we propose to transfer the so-called dark knowledge from the final classifier output to the intermediate blocks through developing a self-distillation strategy for ViT (sec. 3.2). It alleviates the overfitting to source domains by moderating the learning of input-output mapping problem via non-zero entropy supervision of intermediate blocks. We show that by improving the intermediate blocks, which are essentially multiple feature pathways, through soft supervision from final classifier facilitates the model towards learning cross-domain generalizable features (see Fig. 2). Our approach naturally fits into the modular and compositional architecture of different ViTs, and does not introduce any new parameters. As such it adds a minimal training overhead over the baseline. Extensive experiments have been conducted on five diverse datasets from DomainBed suite [13], including PACS, VLCS, OfficeHome, TerraIncognita, and DomainNet. We empirically show better performance across different DG baselines as well as different ViT backbones in all five datasets. Further, we demonstrate competitive performance against the recent state-of-the-art DG methods. With T2T-ViT-14 backbone, we obtain an (overall) average accuracy (five datasets) of 68.1%, thereby outperforming the existing best [30] by 1.3%.

2 Related Work

Domain Generalization: A prevalent motivation of several existing DG methods is to learn the underlying domain-invariant representations from the available source data. The seminal work of Vapnik *et al.* [31] introduced Empirical Risk Minimization (ERM), which minimizes the sum of squared errors across domains. Following this route, we observe several variants aimed at uncovering the domain-invariant features via matching distributions across domains. For instance, Muandet *et al.* [9] employed maximum mean discrepancy (MMD) constraint, Ghifary *et al.* [10] proposed a multi-task autoencoder, and Yang *et al.* [32] used canonical correlation analysis (CCA). Arjovsky *et al.* [33] proposed the learning of invariant predictors across various source domains. A few methods used low rank regularization to extract invariant features e.g., [11, 34]. Meta-learning based methods have also been used as regularizers. Li *et al.* [24] switched domain-specific feature extractors and classifiers across domains via episodic training. Balaji *et al.* [23] learned a regularization function in an episodic training paradigm. Furthermore, some DG methods masked features via ranking gradients [14], utilized auxiliary tasks [12, 28], employed domain-specific masks [35], and exploited domain-specific normalizations [36]. A few DG approaches proposed contrastive semantic alignment and self-supervised contrastive formulations [17, 22, 37]. Another class of DG methods employ various data augmentation techniques to improve the diversity of source domains. Shankar *et al.* [25] proposed Crossgrad training, Volpi *et al.* [26] imposed wasserstein constraint in semantic space, Zhou *et al.* [27] learned a generator to generate new examples, and Khan *et al.* [15] estimated class-conditional covariance matrices for generating novel source features. Recently, Gulrajani *et al.* [13] demonstrated that, under a fair evaluation protocol, a simple empirical risk minimization (ERM) method can achieve state-of-the-art DG performance. Cha *et al.* [30] proposed stochastic weight averaging in a dense manner to achieve flatter minima for DG. We note that, all aforementioned DG methods are based on CNN architecture, however, little to no attention has been paid towards investigating the DG performance of ViTs. To this end, we choose to study the performance of ViTs under domain generalization with ERM as a simple, but strong DG baseline.

Vision transformers: operate in a hierarchical manner by processing input images as a sequence of non-overlapping patches via the self attention mechanism. Recently, we have seen some ViT-based methods for image classification [4, 38, 39], object detection [5, 40], and semantic segmentation [41, 42]. Dosovitskiy *et al.* [4] proposed the first fully functional ViT model for image classification. Despite its promising performance, its adoption remained limited because it requires large-scale datasets for model training and huge computation resources. Towards improving data efficiency in ViTs, Touvron *et al.* [39] developed Data-efficient image Transformer (DeiT); it attains competitive results against the CNN by training only on ImageNet and without leveraging external data. Similarly, Yuan *et al.* [43] proposed Tokens-To-Token Vision Transformer (T2T-ViT) strategy. It progressively structurizes the patch tokens in a way that the local structure represented by surrounding tokens can be modeled while reducing the tokens

length. Furthermore, Wu *et al.* [38] proposed a hybrid approach, namely Convolutional Vision Transformer (CvT), by combining the strengths of CNNs and ViTs aimed at improving the performance and robustness of ViTs, while maintaining computational and memory efficiency. Recently, Zhang *et al.* [44] studied the performance of ViTs under distribution shifts and proposed a generalization-enhanced vision transformer from the outlook of self-supervised learning and information theory. They concluded that by scaling the capacity of ViTs the out-of-distribution (OOD) generalization performance can be enhanced, mostly under the domain adaptation settings. On the other hand, we show that it is possible to improve the OOD generalization performance of ViTs without introducing any new parameters under the established DG protocols [13].

Knowledge Distillation: was initially designed for model compression and aims at matching the output of a teacher model to a student model whom capacity is smaller than the teacher model [45]. Zhang *et al.* [46] partitioned a CNN model into several blocks, and the knowledge from the full (deeper part) of the model is squeezed into the shallow parts. Yun *et al.* [47] proposed a self-distillation approach based on penalizing the predictive distributions between similar data samples. In particular, it distills the predictive distribution between different samples of the same label during training. Towards addressing DG problem, Wang *et al.* [48] proposed a teacher-student distillation strategy, based on CNNs, and a gradient filter as an efficient regularization term. In contrast, we propose a new self-distillation strategy to enhance the DG capabilities of ViTs. It prevents introducing any new parameters via seamlessly exploiting the modular architecture of ViTs.

3 Proposed Approach

In this paper, we aim to explore ViTs towards tackling domain generalization problem. We observe that a simple, but competitive DG baseline (ERM) built on ViT displays notable performance decay in a typical DG setting (Fig. 1). Towards this end, we propose a simple plug-and-play DG approach for ERM-ViT, termed as self-distillation for ViTs, that explicitly facilitates the model towards exploiting cross-domain transferable features (Fig. 2). The proposed self-distillation approach reduces the overfitting to the source domains by easing the mapping problem via non-zero entropy supervision of multiple feature pathways. Furthermore, it seamlessly fits into the modular architecture of ViTs and avoids introducing any new parameters.

3.1 Preliminaries

Problem Settings: In traditional domain generalization (DG) setting [13], we assume the availability of data from a set of training (source) domains $\mathcal{D} = \{\mathcal{D}_k\}_{k=1}^K$. Where \mathcal{D}_k denotes a distribution over the input space \mathcal{X} and K is the total number of training domains. From a domain k , we sample J training datapoints that comprise of input x and label y as pairs $(x_j^k \in \mathcal{X}, y_j^k \in \mathcal{Y})_{j=1}^J$.

Besides a set of training (source) domains, we also assume a set of target domains $\{\mathcal{T}\}_{t=1}^T$, where T is the total number of target domains and is typically set to 1. The goal in DG is to learn a mapping $\mathcal{F}_\theta : \mathcal{X} \rightarrow \mathcal{Y}$ that provides accurate predictions on data from an unseen target domain \mathcal{T}_t .

Empirical risk minimization (ERM) for DG: Assume a loss function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y}$ which can quantify the prediction error, such as standard Cross Entropy (CE) for image recognition task. A simple DG baseline accumulates the data from multiple source domains \mathcal{D} and searches for a predictor minimizing the following empirical risk [31]: $\frac{1}{N} \sum_{i=1}^N \mathcal{L}(\mathcal{F}_\theta(x_k^j, y_k^j))$. Where $N = K \times J$ is the total number of data points from all source domains. Recently, Gulrajani and Lopez-Paz [13] demonstrated that this simple ERM based DG baseline shows competitive results or even performs better than many previous state-of-the-art DG methods under a fair evaluation protocol.

ViT based ERM: While exploring ViTs for DG, we observe that ViT-based ERM (ERM-ViT) shows notable performance drop, similar to their CNN-based ERM counterpart (Fig. 1). This is likely due to the lack of any explicit overfitting mechanism and the supervision from one-hot encoded ground truth labels. It renders the overall learning of the mapping problem, from input space to label space, rather difficult. As a result, the model is more prone to exploiting non-generalizable, brittle features, such as the specific background of a domain (Fig. 2 & 6). To tackle this problem, in the next section, we propose a new self-distillation technique for improving the DG capabilities of ViTs. The core idea is to ease the mapping problem by generating non-zero entropy supervision for multiple feature pathways in ViTs. This enables the model towards utilizing more generalizable features (Fig. 2 & 6), that are mostly shared across source and target domains.

3.2 Self-distilled Vision Transformer for Domain Generalization

ViTs have modular architectures: Assume the model \mathcal{F} is composed of n intermediate layers and the final classifier h such as $\mathcal{F} = (f_1 \circ f_2 \circ f_3 \circ \dots \circ f_n) \circ h$, where f_i represents an intermediate block or layer. In the case of ViT (*e.g* DeiT-Small [39]), f_i is based on self-attention transformer block and such network design is monolithic as any transformer block produces equi-dimensional features that is $\mathbb{R}^{m \times d}$, where m represent the number of input features or tokens and each token has d dimensions. The monolithic design approach of ViT allows a self-ensemble behavior [29], where the output of each block can be processed by the final classifier h to create an intermediate classifier⁴.

$$\mathcal{F}_i = f_i \circ h \quad (1)$$

Our goal is to induce the so-called dark knowledge, non-zero entropy supervisory signals, from the final classifier to these sub-models, manifesting multiple feature pathways.

⁴ For non-monolithic ViT designs and CNNs, where feature dimension changes across the layers, the intermediate classifier can be obtained via $\mathcal{F}_i = (f_i \circ g_i) \circ h$, where g projects the output of f_i to the same dimension as h .




Fig. 3: Proposed self-distillation in ViTs for domain generalization (ERM-SDViT). ViTs build upon a modular and a hierarchical architecture, where a model is comprised of n intermediate blocks/layers (f_i) and a final classifier h . The ‘Selector’ chooses a random block from the range of intermediate blocks and makes a prediction after passing its classification token through the final classifier. This way the dark knowledge, as non-zero entropy signals, is distilled from the final classification token to the intermediate class tokens during training.

Self-distillation in ViTs: As discussed earlier, ViTs can be easily dissected into number of sub-models due to their monolithic architectural design as each transformer block produces a classification token that can be processed by the final classification head (Eqn. 1) to produce a class-specific score. Each sub-model represents a discriminative feature pathway through the network. We believe that inducing dark knowledge from final classifier output to these sub-models via soft supervision during training can enhance the overall model’s capability towards learning object semantics.

Random sub-model distillation: The number of sub-models within a given ViT depends on the number of Transformer blocks (see Fig. 3) and distilling the knowledge to all of the sub-models at once pose optimization difficulties during online training. Therefore, we introduce a simple technique that randomly samples one sub-model based on Eqn. 1 from all the possible set of sub-models (see Fig. 3). In this manner, our approach trains all sub-models but knowledge is transferred to only a single sub-model at any step of the training. This strategy eases the optimization and lead to better domain generalization.

Impact on internal representations: In Fig. 4, we plot the block-wise accuracy of baseline (ERM-ViT) and our method (ERM-SDViT). Random sub-model distillation improves the accuracy of all blocks, in particular, the improvement is more pronounced for the earlier blocks. Besides later blocks, it also encourages earlier blocks to bank on transferable representations, yet discriminative representations. Since these earlier blocks manifest multiple discriminative feature pathways, we believe that they better facilitate the overall model towards capturing the semantics of the object class.

Training Objective: For a given input x , the prediction error from the final classification token of the ViT is computed using cross-entropy loss in comparison




Fig. 4: Block-wise accuracy of baseline (ERM-ViT [39]) and our method (ERM-SDViT), featuring random sub-model distillation for DG. Results are reported on four challenging target domains from PACS dataset.

with one-hot encoded ground-truths as follows.

$$\mathcal{L}_{\text{CE}}(\mathcal{F}(x), y) = - \sum_{j=1}^n y_j \log(\mathcal{F}(x)_j), \quad (2)$$

where n is the output dimension of the final classifier. We randomly sample an intermediate block as shown in Fig. 3 and produce logits from the classification token of a sub-model by applying Eqn.1. We then compute the difference between the final and randomly sampled intermediate classification token by comparing the KL divergence between their logit distributions as follows:

$$\mathcal{L}_{\text{KL}}(\mathcal{F}(x) \parallel \mathcal{F}_i(x)) = \sum_{j=1}^n \sigma(\mathcal{F}(x)/\tau)_j \log \frac{\sigma(\mathcal{F}(x)/\tau)_j}{\sigma(\mathcal{F}_i(x)/\tau)_j}, \quad (3)$$

where σ denotes the softmax operation and τ represent temperature used to rescale the logits [45]. The model is optimized by minimizing the overall loss based on Eqns. 2 and 3 and given as follows:

$$\mathcal{L} = \mathcal{L}_{\text{CE}} + \lambda \mathcal{L}_{\text{KL}}, \quad (4)$$

where λ balances the contribution of \mathcal{L}_{KL} towards the overall loss \mathcal{L} .

4 Experiments

Datasets: Following the work of Gulrajani and Lopez-Paz [13], we rigorously evaluate the effectiveness of our proposed method and draw comparisons with the

existing state-of-the-art on five benchmark datasets including PACS [11], VLCS [49], OfficeHome [50], TerraIncognita [51] and DomainNet [52]. PACS [11] contains four domains $d \in \{\text{Art, Cartoons, Photos, Sketches}\}$, 7 classes and a total of 9,991 images. VLCS [49] comprises of four domains as well $d \in \{\text{Caltech101, LabelMe, SUN09, VOC2007}\}$, 5 classes and offers 10,729 images. OfficeHome [50] also contains four domains $d \in \{\text{Art, Clipart, Product, Real}\}$, 65 classes and a total of 15,588 images. TerraIncognita [51]: has four camera-trap domains $d \in \{\text{L100, L38, L43, L46}\}$, 10 classes and offers 24,778 wild photographs. DomainNet [52] contains six domains $d \in \{\text{Clipart, Infograph, Painting, Quickdraw, Real, Sketch}\}$, 345 classes and 586,575 images.

Implementation and training/testing details: To allow fair comparisons, we follow the training and evaluation protocol of Gulrajani and Lopez-Paz [13]. We use the training domain validation protocol for model selection. After partitioning each training domain data into the training and validation subsets (80%/20%), the validation data from each training domain are pooled to obtain an overall validation set. The model that maximizes the accuracy on this overall validation set is considered as the best model which is then evaluated on the target domain to report classification (top-1) accuracy. For all our ViT-based methods, including the proposed approach, we use AdamW [53] optimizer and use the default hyperparameters (HPs) of ERM from [13]⁵, including the batch size of 32, the learning rate of $5e-05$, and the weight decay of 0.0. Note that, only the values of our method-specific HPs, λ and τ , are sought via grid search in the ranges $\{0.1, 0.2, 0.5\}$ and $\{3.0, 5.0\}$, respectively, using the validation set. We report accuracy for each target domain and their average where a model is trained/validated on training domains and evaluated on an (unseen) target domain. Each accuracy on target domain is an average over three different trials with different train-validation splits.

Evaluation with different ViT backbones: We establish the generalizability of our proposed approach by reporting results with three different ViT backbones, namely **DeiT** [39], **CvT** [38], and **T2T-ViT** [43]. **DeiT** is a data-efficient image transformer and was trained on 1.2 million ImageNet examples. We use the DeiT-Small model since it has 22M parameters and can be regarded as a ResNet-50 counterpart containing 23.5M parameters. Note that, we utilize the DeiT-Small model without the distillation token and a student-teacher formulation. **CvT** introduces convolutions into ViT to improve the accuracy and efficiency. The introduction of convolution allows ViT to handle shift, scale and distortion invariance while retaining the important features of ViTs such as the dynamic attention and global context. We use CvT-21 that contains 32M parameters in our baselines and proposed self-distilled ViTs. **T2T-ViT** relies on a progressive tokenization mechanism to aggregate neighboring Tokens to one Token; it can encode the local structure information of surrounding tokens and reduce the length of tokens iteratively. Further, it builds on an efficient ViT backbone by borrowing architecture designs from CNN. We use T2T-ViT-14 model, containing 21.5M

⁵ They are default parameters in the pre-defined ranges [13] for random HP search.

Table 1: Comparison with the several (17) existing state-of-the-art DG methods on five different benchmarks, including PACS, VLCS, OfficeHome, TerraIncognita and DomainNet. Best results are in bold and the second best are underlined.

Algorithm	Backbone	# Params	VLCS	PACS	OfficeHome	TerraInc	DomainNet	Average
ERM [13]	ResNet-50	23.5M	77.5 \pm 0.4	85.5 \pm 0.2	66.5 \pm 0.3	46.1 \pm 1.8	40.9 \pm 0.1	63.3
IRM [33]	ResNet-50	23.5M	78.5 \pm 0.5	83.5 \pm 0.8	64.3 \pm 2.2	47.6 \pm 0.8	33.9 \pm 2.8	61.5
GroupDRO [54]	ResNet-50	23.5M	76.7 \pm 0.6	84.4 \pm 0.8	66.0 \pm 0.7	43.2 \pm 1.1	33.3 \pm 0.2	60.7
Mixup [55]	ResNet-50	23.5M	77.4 \pm 0.6	84.6 \pm 0.6	68.1 \pm 0.3	47.9 \pm 0.8	39.2 \pm 0.1	63.4
MLDG [21]	ResNet-50	23.5M	77.2 \pm 0.4	84.9 \pm 1.0	66.8 \pm 0.6	47.7 \pm 0.9	41.2 \pm 0.1	63.5
CORAL [56]	ResNet-50	23.5M	78.8 \pm 0.6	86.2 \pm 0.3	68.7 \pm 0.3	47.6 \pm 1.0	41.5 \pm 0.1	64.5
MMD [20]	ResNet-50	23.5M	77.5 \pm 0.9	84.6 \pm 0.5	66.3 \pm 0.1	42.2 \pm 1.6	23.4 \pm 9.5	58.8
DANN [19]	ResNet-50	23.5M	78.6 \pm 0.4	83.6 \pm 0.4	65.9 \pm 0.6	46.7 \pm 0.5	38.3 \pm 0.1	62.6
CDANN [57]	ResNet-50	23.5M	77.5 \pm 0.1	82.6 \pm 0.9	65.8 \pm 1.3	45.8 \pm 1.6	38.3 \pm 0.3	62.0
MTL [58]	ResNet-50	23.5M	77.2 \pm 0.4	84.6 \pm 0.5	66.4 \pm 0.5	45.6 \pm 1.2	40.6 \pm 0.1	62.8
SagNet [16]	ResNet-50	23.5M	77.8 \pm 0.5	86.3 \pm 0.2	68.1 \pm 0.1	48.6 \pm 1.0	40.3 \pm 0.1	64.2
ARM [59]	ResNet-50	23.5M	77.6 \pm 0.3	85.1 \pm 0.4	64.8 \pm 0.3	45.5 \pm 0.3	35.5 \pm 0.2	61.7
VREx [60]	ResNet-50	23.5M	78.3 \pm 0.2	84.9 \pm 0.6	66.4 \pm 0.6	46.4 \pm 0.6	33.6 \pm 2.9	61.9
RSC [14]	ResNet-50	23.5M	77.1 \pm 0.5	85.2 \pm 0.9	65.5 \pm 0.9	46.6 \pm 1.0	38.9 \pm 0.5	62.6
SelfReg [17]	ResNet-50	23.5M	77.5 \pm 0.0	86.5 \pm 0.3	69.4 \pm 0.2	51.0 \pm 0.4	44.6 \pm 0.1	65.8
mDSDI [18]	ResNet-50	23.5M	79.0 \pm 0.3	86.2 \pm 0.2	69.2 \pm 0.4	48.1 \pm 1.4	42.8 \pm 0.1	65.0
SWAD [30]	ResNet-50	23.5M	79.1 \pm 0.1	88.1 \pm 0.1	70.6 \pm 0.2	50.0 \pm 0.3	46.5 \pm 0.1	66.8
ERM-ViT [39]	DeiT-Small	22M	78.3 \pm 0.5	84.9 \pm 0.9	71.4 \pm 0.1	43.4 \pm 0.5	45.5 \pm 0.0	64.7
ERM-SDViT	DeiT-Small	22M	78.9 \pm 0.4	86.3 \pm 0.2	71.5 \pm 0.2	44.3 \pm 1.0	45.8 \pm 0.0	65.4
ERM-SDViT + T3A	DeiT-Small	22M	<u>81.6 \pm 0.1</u>	86.7 \pm 0.2	72.5 \pm 0.3	44.9 \pm 0.4	47.4 \pm 0.1	66.6
ERM-ViT [38]	CvT-21	32M	79.0 \pm 0.3	86.9 \pm 0.3	75.5 \pm 0.0	48.7 \pm 0.4	50.4 \pm 0.1	68.1
ERM-SDViT	CvT-21	32M	79.2 \pm 0.4	<u>88.3 \pm 0.2</u>	<u>75.6 \pm 0.2</u>	49.7 \pm 1.4	<u>50.4 \pm 0.0</u>	68.6
ERM-SDViT + T3A	CvT-21	22M	81.9 \pm 0.4	88.9 \pm 0.5	77.0 \pm 0.2	51.4 \pm 0.7	52.0 \pm 0.0	70.2
ERM-ViT [43]	T2T-ViT-14	21.5M	78.9 \pm 0.3	86.8 \pm 0.4	73.7 \pm 0.2	48.1 \pm 0.2	48.1 \pm 0.1	67.1
ERM-SDViT	T2T-ViT-14	21.5M	79.5 \pm 0.8	88.0 \pm 0.7	74.2 \pm 0.3	50.6 \pm 0.8	48.2 \pm 0.2	68.1
ERM-SDViT + T3A	T2T-ViT-14	21.5M	81.2 \pm 0.3	87.8 \pm 0.6	75.5 \pm 0.2	50.5 \pm 0.6	50.2 \pm 0.1	<u>69.0</u>

parameters, which is approx. equivalent to the ResNet-50 model, comprising 23.5M parameters.

4.1 Comparison with the state-of-the-art

We compare our approach to several (in particular, 17) existing state-of-the-art algorithms for DG (see Table 1) listed in Domainbed suite [13]. Specifically, we include the following DG algorithms: Empirical Risk Minimization (ERM) [13], Invariant Risk Minimization (IRM) [33], Group Distributionally Robust Optimization (GroupDRO) [54], Inter-domain Mixup (Mixup) [55], Meta-Learning for Domain Generalization (MLDG) [21], Deep CORrelation ALignment (CORAL) [56], Maximum Mean Discrepancy (MMD) [20], Domain Adversarial Neural Networks (DANN) [19], Class-conditional DANN (CDANN) [57], Marginal Transfer Learning (MTL) [58], Style-Agnostic Networks (SagNet) [16], Adaptive Risk Minimization (ARM) [59], Variance Risk Extrapolation (VREx) [60], Representation Self Challenging (RSC) [14], Self-supervised contrastive regularization (SelfReg) [17], meta-Domain Specific-Domain Invariant (mDSDI) [18], and Stochastic Weight Averaging Densely (SWAD) [30].

VLCS and OfficeHome: In VLCS, our approach (ERM-SDViT) records the best classification accuracy of 79.5% with T2T-ViT-14 backbone outperforming the baseline (ERM-ViT) and DG SOTA algorithms. Similarly, in OfficeHome, our method outperforms all other methods under all three ViT backbones. In particular, it displays the best accuracy of 75.6% with CvT-21 backbone.




Fig. 5: t-SNE visualization of features from different blocks (9 & 12) in baseline and our approach. Left: Features are colored corresponding to their class labels (classes: 7, PACS dataset). Right: Features are colored corresponding to their domain labels. Our approach has performed well for instance, in class wise t-SNE in block 9, the features of class 0 and 1 (highlighted in red circle) are well separated as compared to ERM-ViT baseline. Similarly in class 0 and 4 in the final 12th block features of our ERM-SDViT approach are also separated clearly. While in domain-wise t-SNE, similar pattern is observed, as source and target domain features are more overlapped with each other and clearly separated as well.

PACS, DomainNet and TerraInc: In PACS our approach delivers the top accuracy of 88.3% and in DomainNet it achieves an accuracy of 50.4% with CvT-21 backbone. In TerraIncognita, our method achieves a competitive accuracy of 50.6% with T2T-ViT-14 backbone against the top performing method of SelfReg [17]. In the overall average accuracy over five datasets, our method outperforms the existing state-of-the-art in DG with CvT-21 and T2T-ViT-14 backbones. Moreover, it provides notable gains over the baseline (ERM-ViT) under the three recent ViT backbone architectures.

4.2 Ablation Study and Analysis

In this section, we report ablation studies and different analysis to understand our proposed approach. In all experiments, unless stated otherwise, the baseline method is ERM-ViT with DeiT-Small backbone [39].

With a recent DG baseline: We show that our proposed approach is also effective in further improving the performance of a strong DG baseline namely T3A [61] (see Table 1 and 2). T3A computes a pseudo-prototype representation for each class using unlabeled data augmented by the base classifier trained in




Fig. 6: Comparison of attention maps between the baseline ERM-ViT and our proposed ERM-SDViT (backbone: DeiT-Small) on four target domains of PACS dataset. Note that the attention maps are computed at the final block of ViT models.

Table 2: Our proposed approach is also effective in further improving the performance of a strong DG baseline namely T3A [61]. Results are reported with different ERM-ViT backbone architectures: DeiT-Small, CvT-21, and T2T-ViT-14.

Model	Backbone	# of Params	Art	Cartoons	Photos	Sketches	Average
ERM	ResNet-50	23.5M	81.3 ± 0.6	80.9 ± 0.3	96.3 ± 0.6	78.0 ± 1.6	84.1 ± 0.4
ERM-ViT	DeiT-Small	22M	87.4 ± 1.2	81.5 ± 0.8	98.1 ± 0.1	72.6 ± 3.3	84.9 ± 0.9
ERM-ViT + T3A	DeiT-Small	22M	88.1 ± 1.5	81.8 ± 0.8	98.3 ± 0.1	73.8 ± 2.7	85.5 ± 0.7
ERM-SDViT	DeiT-Small	22M	87.6 ± 0.3	82.4 ± 0.4	98.0 ± 0.3	77.2 ± 1.0	86.3 ± 0.2
ERM-SDViT + T3A	DeiT-Small	22M	88.2 ± 0.4	82.9 ± 0.5	98.3 ± 0.1	77.2 ± 0.9	86.7 ± 0.2
ERM-ViT	CvT-21	32M	89.0 ± 1.0	84.8 ± 0.6	98.8 ± 0.2	78.6 ± 0.3	87.8 ± 0.1
ERM-ViT + T3A	CvT-21	32M	90.1 ± 0.7	85.3 ± 0.6	99.0 ± 0.1	79.6 ± 0.4	88.5 ± 0.1
ERM-SDViT	CvT-21	32M	90.8 ± 0.1	84.1 ± 0.5	98.3 ± 0.2	80.0 ± 1.3	88.3 ± 0.2
ERM-SDViT + T3A	CvT-21	32M	91.2 ± 0.8	83.5 ± 0.2	98.3 ± 0.1	82.5 ± 1.5	88.9 ± 0.5
ERM-ViT	T2T-ViT-14	21.5M	89.6 ± 0.9	81.0 ± 0.9	98.9 ± 0.2	77.6 ± 2.6	86.8 ± 0.4
ERM-ViT + T3A	T2T-ViT-14	21.5M	90.7 ± 1.0	82.4 ± 0.6	99.1 ± 0.1	78.5 ± 2.2	87.7 ± 0.4
ERM-SDViT	T2T-ViT-14	21.5M	90.2 ± 1.2	82.7 ± 0.7	98.6 ± 0.2	80.5 ± 2.2	88.0 ± 0.7
ERM-SDViT + T3A	T2T-ViT-14	21.5M	89.2 ± 1.8	84.0 ± 0.1	98.7 ± 0.1	79.3 ± 1.1	87.8 ± 0.6

the source domains in an online manner. A test image is classified based on its distance to the pseudo-prototype representation. Our proposed approach with T3A (ERM-SDViT+T3A) consistently improves the performance of the baseline (ERM-ViT+T3A) with all three ViT backbones.

Feature visualizations: Fig. 5 (left) visualizes the class-wise feature representations of different blocks using t-SNE in baseline and our approach. Compared to baseline, our approach facilitates learning more discriminative features. As a result, the intra-class features are compactly clustered while the inter-class features are far apart. Likewise, fig. 5 (right) visualizes the same features, however, they are colored based on their domain labels. Compared to baseline, our method promotes a greater overlap between the features of source and target domains.

On different block selection techniques: We show the impact on performance with different ways of selecting blocks in our self-distillation method (see Table 3). First, we restrict the random sampling of blocks to earlier blocks i.e.

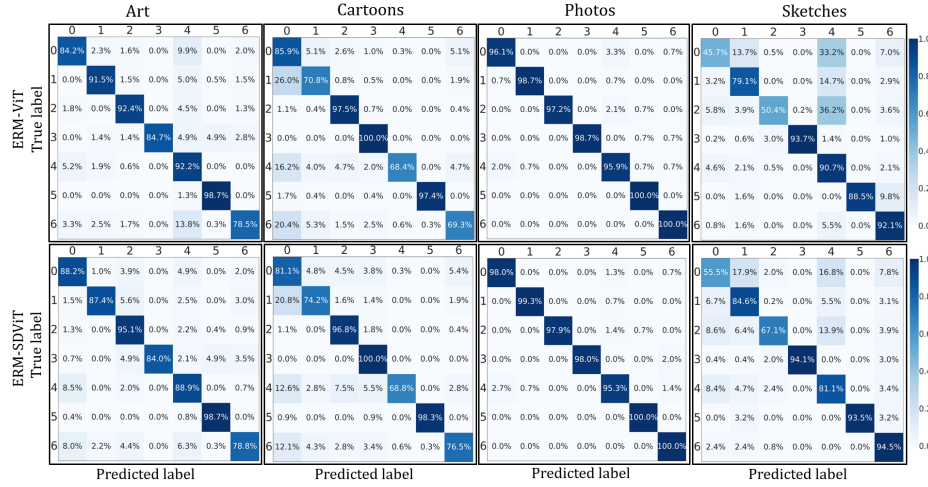


Fig. 7: Confusion matrices of the baseline and our method on PACS dataset. The classes in the Figure are ‘0’:Dog, ‘1’: Elephant, ‘2’:Giraffe, ‘3’:Guitar, ‘4’:Horse, ‘5’:House, and ‘6’:Person.

Table 3: Impact on performance with different ways of selecting transformer blocks in our self-distillation method.

Model	Art	Cartoons	Photos	Sketches	Average
ERM-ViT	87.4 ± 1.2	81.5 ± 0.8	98.1 ± 0.1	72.6 ± 3.3	84.9 ± 0.9
ERM-SDViT[0-5]	87.3 ± 0.2	82.1 ± 0.4	98.3 ± 0.1	76.6 ± 2.1	86.1 ± 0.5
ERM-SDViT[6-11]	86.8 ± 0.8	81.4 ± 0.3	98.6 ± 0.1	74.3 ± 1.7	85.3 ± 0.3
ERM-ViT[self.dist all blocks]	87.8 ± 1.9	82.2 ± 0.7	97.9 ± 0.1	75.0 ± 1.1	85.7 ± 0.3
ERM-SDViT[0-11] Ours	87.6 ± 0.3	82.4 ± 0.4	98.0 ± 0.3	77.2 ± 1.0	86.3 ± 0.2

from block 0 to block 5. Second, we limit the random sampling of blocks to later blocks i.e. from block 6 to block 11. Finally, we do not perform any sampling in any range and rather include all the blocks [46]. Compared to all aforementioned block selection techniques, our proposal of random sampling from the full range of blocks shows the best (overall) average accuracy of 86.3%. Sampling from earlier blocks seems beneficial as compared to the later blocks. When earlier blocks, with relatively longer feature pathways to the final block, are provided with soft target labels, there is potentially a greater room for exploring cross-domain generalizable features.

Confusion matrices: Fig. 7 visualizes the confusion matrices for the baseline and our method on PACS dataset. Compared to baseline, our method produces less number of false positives, particularly in ‘Sketches’ as target domain.

What kind of features our DG approach facilitates?: We attempt to visualize the features that are used by the baseline and our method to make predictions through visualizing attention maps (see Fig. 6). In all target domains,

Table 4: Training overhead, computed as relative % increase in training time (hrs.), introduced by our method on top of the baseline.

Model	Art	Cartoons	Photos	Sketches
ERM-ViT	0.266	0.270	0.278	0.267
ERM-SDViT	0.279	0.276	0.279	0.278
Rel.overhead	4.88	2.22	0.35	4.11

Table 5: Performance under various types of domain shifts. We categorize five DG datasets, including PACS, VLCS, OfficeHome (OH), TerraIncognita (Terra), and DomainNet, into these four different domain shift categories based on the type of shift(s) exhibited by them. Each entry in the Table is the accuracy (%) averaged over the datasets belonging to a certain domain shift category.

Methods	Shift Type			
	Background Shifts (VLCS,Terra)	Corruption Shifts (Terra)	Texture Shifts (PACS,DomainNet)	Style Shifts (OH,PACS,DomainNet)
ERM-ViT	60.0	43.4	65.2	67.2
ERM-SDViT	61.6	44.3	66.0	67.8

we observe that our method mostly attends to features that mainly capture the semantics and the shape of the object class. On the contrary, the baseline has a greater tendency to attend background features and thus focus less on the object class features. Our self-distillation approach encourages the intermediate blocks to bank on more robust and generalizable features via providing supervision with a soft non-zero entropy signal, thereby reducing the overfitting of the model to the source domains.

Training overhead: Table 4 reports training overhead, computed as relative % increase in training time (hrs.), introduced by our method on top of the baseline. Our method adds very little training overhead over the baseline.

Performance under different domain shifts: We benchmark the performance under various types of domain shifts, including background shifts, corruption shifts, texture shifts, and style shifts. For instance, background shifts do not affect pixel, shape, texture and structures in foreground object. Whereas style shifts typically depict variance at different stages of concepts, including texture, shape and object part [44]. To this end, we classify five DG datasets, including PACS, VLCS, OfficeHome, TerraIncognita, and DomainNet, into these four different domain shift categories based on the type of shift(s) exhibited by them. Table 5 reports the results by ERM-ViT (baseline) and ERM-SDViT (ours) under four different kinds of domain shifts. We observe that ERM-SDViT outperforms ERM-ViT across the whole spectrum of domain shifts. It shows that our proposed self-distillation strategy allows the learning of cross-domain generalizable features that are effective under the full spectrum of domain shifts.

5 Conclusion

Towards exploring ViTs for the DG problem, we propose self-distilled ViTs for tackling this critical issue. Our Self-distillation strategy in ViTs is a simple and a light-weight approach to alleviate the problem of overfitting to source domains.

It provides soft supervision to the intermediate blocks of ViTs to strengthen their internal representations, thereby moderating the learning of input-output mapping problem. Moreover, it avoids introduction of any new parameters and can be easily plugged into various ViT architectures, owing to their monolithic and modular architecture. Experimental results on five datasets with different DG baselines and ViT backbones, including comparisons with the recent SOTA, validate the effectiveness of our approach for ViTs tackling DG problem.

References

1. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)
2. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
3. Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J.D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al.: Language models are few-shot learners. *Advances in neural information processing systems* **33** (2020) 1877–1901
4. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* (2020)
5. Carion, N., Massa, F., Synnaeve, G., Usunier, N., Kirillov, A., Zagoruyko, S.: End-to-end object detection with transformers. In: *European conference on computer vision*, Springer (2020) 213–229
6. Zhu, X., Su, W., Lu, L., Li, B., Wang, X., Dai, J.: Deformable detr: Deformable transformers for end-to-end object detection. *arXiv preprint arXiv:2010.04159* (2020)
7. Zheng, S., Lu, J., Zhao, H., Zhu, X., Luo, Z., Wang, Y., Fu, Y., Feng, J., Xiang, T., Torr, P.H., et al.: Rethinking semantic segmentation from a sequence-to-sequence perspective with transformers. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. (2021) 6881–6890
8. 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*. (2021) 568–578
9. Muandet, K., Balduzzi, D., Schölkopf, B.: Domain generalization via invariant feature representation. In: *International Conference on Machine Learning*. (2013) 10–18
10. Ghifary, M., Bastiaan Kleijn, W., Zhang, M., Balduzzi, D.: Domain generalization for object recognition with multi-task autoencoders. In: *Proceedings of the IEEE international conference on computer vision*. (2015) 2551–2559
11. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: *Proceedings of the IEEE international conference on computer vision*. (2017) 5542–5550
12. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2019) 2229–2238

13. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. *ArXiv abs/2007.01434* (2021)
14. Huang, Z., Wang, H., Xing, E.P., Huang, D.: Self-challenging improves cross-domain generalization. (2020)
15. Khan, M.H., Zaidi, T., Khan, S., Khan, F.S.: Mode-guided feature augmentation for domain generalization. (2021)
16. Nam, H., Lee, H., Park, J., Yoon, W., Yoo, D.: Reducing domain gap by reducing style bias. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. (2021) 8690–8699
17. Kim, D., Yoo, Y., Park, S., Kim, J., Lee, J.: Selfreg: Self-supervised contrastive regularization for domain generalization. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. (2021) 9619–9628
18. Bui, M.H., Tran, T., Tran, A., Phung, D.: Exploiting domain-specific features to enhance domain generalization. *Advances in Neural Information Processing Systems* **34** (2021)
19. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *The Journal of Machine Learning Research* **17** (2016) 2096–2030
20. Li, H., Pan, S.J., Wang, S., Kot, A.C.: Domain generalization with adversarial feature learning. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. (2018) 5400–5409
21. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Learning to generalize: Meta-learning for domain generalization. In: *Thirty-Second AAAI Conference on Artificial Intelligence*. (2018)
22. Dou, Q., de Castro, D.C., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. In: *Advances in Neural Information Processing Systems*. (2019) 6450–6461
23. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: *Advances in Neural Information Processing Systems*. (2018) 998–1008
24. Li, D., Zhang, J., Yang, Y., Liu, C., Song, Y.Z., Hospedales, T.M.: Episodic training for domain generalization. In *ICCV* (2019)
25. Shankar, S., Piratla, V., Chakrabarti, S., Chaudhuri, S., Jyothi, P., Sarawagi, S.: Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745* (2018)
26. Volpi, R., Namkoong, H., Sener, O., Duchi, J.C., Murino, V., Savarese, S.: Generalizing to unseen domains via adversarial data augmentation. In: *Advances in Neural Information Processing Systems*. (2018) 5334–5344
27. Zhou, K., Yang, Y., Hospedales, T., Xiang, T.: Learning to generate novel domains for domain generalization. (2020)
28. Wang, S., Yu, L., Li, C., Fu, C.W., Heng, P.A.: Learning from extrinsic and intrinsic supervisions for domain generalization. (2020)
29. Naseer, M., Ranasinghe, K., Khan, S., Khan, F.S., Porikli, F.: On improving adversarial transferability of vision transformers. *arXiv preprint arXiv:2106.04169* (2021)
30. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: Swad: Domain generalization by seeking flat minima. *Advances in Neural Information Processing Systems* **34** (2021)
31. Vapnik, V.: *The nature of statistical learning theory*. Springer science & business media (1999)

32. YANG, P.Y., Gao, W.: Multi-view discriminant transfer learning. (2013)
33. Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant risk minimization. arXiv preprint arXiv:1907.02893 (2019)
34. Xu, Z., Li, W., Niu, L., Xu, D.: Exploiting low-rank structure from latent domains for domain generalization. In: European Conference on Computer Vision, Springer (2014) 628–643
35. Chattopadhyay, P., Balaji, Y., Hoffman, J.: Learning to balance specificity and invariance for in and out of domain generalization. (2020)
36. Seo, S., Suh, Y., Kim, D., Kim, G., Han, J., Han, B.: Learning to optimize domain specific normalization for domain generalization. (2020)
37. Motiian, S., Piccirilli, M., Adjero, D.A., Doretto, G.: Unified deep supervised domain adaptation and generalization. In: Proceedings of the IEEE International Conference on Computer Vision. (2017) 5715–5725
38. 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. (2021) 22–31
39. Touvron, H., Cord, M., Douze, M., Massa, F., Sablayrolles, A., Jégou, H.: Training data-efficient image transformers & distillation through attention. In: International Conference on Machine Learning, PMLR (2021) 10347–10357
40. Dai, X., Chen, Y., Yang, J., Zhang, P., Yuan, L., Zhang, L.: Dynamic detr: End-to-end object detection with dynamic attention. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). (2021) 2988–2997
41. Strudel, R., Garcia, R., Laptev, I., Schmid, C.: Segmenter: Transformer for semantic segmentation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. (2021) 7262–7272
42. Lu, Z., He, S., Zhu, X., Zhang, L., Song, Y.Z., Xiang, T.: Simpler is better: Few-shot semantic segmentation with classifier weight transformer. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. (2021) 8741–8750
43. 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. (2021) 558–567
44. Zhang, C., Zhang, M., Zhang, S., Jin, D., Zhou, Q., Cai, Z., Zhao, H., Yi, S., Liu, X., Liu, Z.: Delving deep into the generalization of vision transformers under distribution shifts. arXiv preprint arXiv:2106.07617 (2021)
45. Hinton, G., Vinyals, O., Dean, J.: Distilling the knowledge in a neural network (2015). arXiv preprint arXiv:1503.02531 **2** (2015)
46. Zhang, L., Song, J., Gao, A., Chen, J., Bao, C., Ma, K.: Be your own teacher: Improve the performance of convolutional neural networks via self distillation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. (2019) 3713–3722
47. Yun, S., Park, J., Lee, K., Shin, J.: Regularizing class-wise predictions via self-knowledge distillation. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. (2020) 13876–13885
48. Wang, Y., Li, H., Chau, L.p., Kot, A.C.: Embracing the dark knowledge: Domain generalization using regularized knowledge distillation. In: Proceedings of the 29th ACM International Conference on Multimedia. (2021) 2595–2604
49. Fang, C., Xu, Y., Rockmore, D.N.: Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In: Proceedings of the IEEE International Conference on Computer Vision. (2013) 1657–1664

50. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. (2017) 5018–5027
51. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: *Proceedings of the European conference on computer vision (ECCV)*. (2018) 456–473
52. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *Proceedings of the IEEE International Conference on Computer Vision*. (2019) 1406–1415
53. Loshchilov, I., Hutter, F.: Fixing weight decay regularization in adam. (2018)
54. Sagawa, S., Koh, P.W., Hashimoto, T.B., Liang, P.: Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731* (2019)
55. Yan, S., Song, H., Li, N., Zou, L., Ren, L.: Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677* (2020)
56. Sun, B., Saenko, K.: Deep coral: Correlation alignment for deep domain adaptation. In: *European conference on computer vision*, Springer (2016) 443–450
57. Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., Tao, D.: Deep domain generalization via conditional invariant adversarial networks. In: *Proceedings of the European Conference on Computer Vision (ECCV)*. (2018) 624–639
58. Blanchard, G., Deshmukh, A.A., Dogan, U., Lee, G., Scott, C.: Domain generalization by marginal transfer learning. *arXiv preprint arXiv:1711.07910* (2017)
59. Zhang, M., Marklund, H., Dhawan, N., Gupta, A., Levine, S., Finn, C.: Adaptive risk minimization: Learning to adapt to domain shift. *Advances in Neural Information Processing Systems* **34** (2021)
60. Krueger, D., Caballero, E., Jacobsen, J.H., Zhang, A., Binas, J., Zhang, D., Le Priol, R., Courville, A.: Out-of-distribution generalization via risk extrapolation (rex). In: *International Conference on Machine Learning*, PMLR (2021) 5815–5826
61. Iwasawa, Y., Matsuo, Y.: Test-time classifier adjustment module for model-agnostic domain generalization. *Advances in Neural Information Processing Systems* **34** (2021)

Appendix

More t-SNE feature visualizations: Fig. 8 (left) visualizes the class-wise feature representations of different blocks using t-SNE in baseline (ERM-ViT) and our model (ERM-SDViT) for Caltech101 target domain in the VLCS dataset. In comparison to baseline, our method facilitates improved learning of discriminative features and hence reduces the intra-class variance while increasing the inter-class variance in the feature space. Similarly, Fig. 8 (right) visualizes the same features, however, on the basis of source and target domain labels. Compared to baseline, our method promotes a greater overlap between the features of source and target domain features.

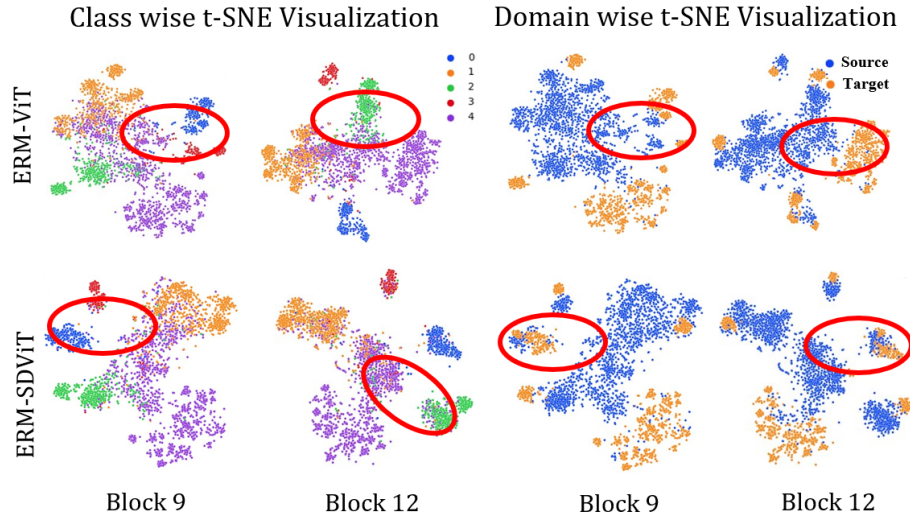


Fig. 8: t-SNE visualization of features from different blocks (9 & 12) in baseline (ERM-ViT) and our approach for Caltech101 target domain in VLCS dataset. Left: Features are colored corresponding to their class labels (classes: 5). Right: Features are colored corresponding to their domain labels. Our approach (ERM-SDViT) improves class-wise discrimination. For instance, in class-wise t-SNE in block 9, the features of class 0 and 3 (highlighted in red circle) are well-separated as compared to the baseline (ERM-ViT). Similarly, in block 12, the features of class 2 and 4 are clearly distinguishable. In domain-wise t-SNE, for our approach, source and target domain features show greater overlap with each other.

Confusion matrices on other DG dataset: Fig. 9 visualizes the confusion matrices for the baseline and our method on VLCS dataset. In comparison to baseline, our method is capable of reducing false positives in all four target domains.

Attention visualizations on other DG datasets: We also visualize attention maps from different images of four datasets, including VLCS, OfficeHome,





Fig. 9: Confusion matrices of the baseline and our method on VLCS dataset. The classes in the figure are ‘0’:Bird, ‘1’:Car, ‘2’:Chair, ‘3’:Dog, and ‘4’:Person.

Fig. 10: Comparison of attention maps between the baseline (ERM-ViT) and our proposed method (ERM-SDViT) on four target domains of TerraIncognita dataset. The ViT backbone is DeiT-Small.



TerraIncognita and DomainNet in Figs. 10, 11 and 12. It can be observed that in all target domains of the four datasets, our method mostly relies on features corresponding to the foreground object’s semantics rather than the background information. However, the baseline approach (ERM-ViT) mostly capitalizes more on the background features and pays less attention to the features belonging to foreground object. For instance, in Fig. 10, target domain: Location_46 of the TerraIncognita dataset, our method is capable of focussing on the foreground object (a dog), which occupies a small fraction of the overall image. However, the baseline model is prone to attending more to the background features, which are prevalent in the image. Note that the attention maps are computed at the final block of ViT models.

Recognition accuracy on target domains of other DG datasets: Tables 6 and 7 compares target domain-wise recognition accuracy on VLCS, OfficeHome, TerraIncognita, and DomainNet datasets of our method with the baseline utilizing three ViT backbones and a DG baseline (T3A [61]).

Table 6: Comparison of target domain-wise classification accuracy on **VLCS**, **OfficeHome**, and **TerraIncognita** datasets. Results are reported of our method with the baseline using three different ViT backbones, including DeiT-Small [39], CvT-21 [38], and T2T-ViT-14 [43], and a DG baseline (T3A [61]).

Dataset			VLCS				
Model	Backbone	# of Params	Caltech101	LableMe	SUN09	VOC2007	Average
ERM	ResNet-50	23.5M	98.1 \pm 0.4	64.1 \pm 0.5	70.7 \pm 0.9	74.8 \pm 2.4	76.9 \pm 0.6
ERM-ViT	DeiT-Small	22M	96.7 \pm 0.8	65.2 \pm 1.0	73.9 \pm 0.3	77.4 \pm 0.3	78.3 \pm 0.5
ERM-SDViT	DeiT-Small	22M	96.8 \pm 0.5	64.2 \pm 0.8	76.2 \pm 0.4	78.5 \pm 0.4	78.9 \pm 0.4
ERM-SDViT + T3A	DeiT-Small	22M	98.9 \pm 0.2	65.9 \pm 0.3	79.8 \pm 0.4	81.9 \pm 0.4	81.6 \pm 0.1
ERM-ViT	CvT-21	32M	97.3 \pm 0.5	65.2 \pm 0.9	76.6 \pm 1.1	76.9 \pm 0.3	79.0 \pm 0.3
ERM-SDViT	CvT-21	32M	96.5 \pm 0.7	63.3 \pm 0.4	78.1 \pm 0.2	78.9 \pm 0.8	79.2 \pm 0.4
ERM-SDViT + T3A	CvT-21	32M	98.4 \pm 0.3	66.8 \pm 0.5	80.1 \pm 1.0	80.6 \pm 0.7	81.9 \pm 0.4
ERM-ViT	T2T-ViT-14	21.5M	96.5 \pm 0.5	64.5 \pm 0.1	76.4 \pm 0.4	78.2 \pm 1.0	78.9 \pm 0.3
ERM-SDViT	T2T-ViT-14	21.5M	96.9 \pm 0.4	64.0 \pm 0.5	76.7 \pm 1.4	80.4 \pm 1.3	79.5 \pm 0.8
ERM-SDViT + T3A	T2T-ViT-14	21.5M	98.6 \pm 0.3	66.5 \pm 0.7	78.2 \pm 0.5	81.7 \pm 0.9	81.2 \pm 0.3
Dataset			OfficeHome				
Model	Backbone	# of Params	Art	Clipart	Product	Real World	Average
ERM	ResNet-50	23.5M	58.8 \pm 1.0	51.3 \pm 0.4	73.7 \pm 0.4	74.7 \pm 0.6	64.6 \pm 0.2
ERM-ViT	DeiT-Small	22M	67.6 \pm 0.3	57.0 \pm 0.6	79.4 \pm 0.1	81.6 \pm 0.4	71.4 \pm 0.1
ERM-SDViT	DeiT-Small	22M	68.3 \pm 0.8	56.3 \pm 0.2	79.5 \pm 0.3	81.8 \pm 0.1	71.5 \pm 0.2
ERM-SDViT + T3A	DeiT-Small	22M	69.1 \pm 1.0	57.9 \pm 0.4	80.7 \pm 0.0	82.3 \pm 0.1	72.5 \pm 0.3
ERM-ViT	CvT-21	32M	74.4 \pm 0.2	59.8 \pm 0.5	83.5 \pm 0.4	84.1 \pm 0.2	75.5 \pm 0.0
ERM-SDViT	CvT-21	32M	73.8 \pm 0.6	60.7 \pm 0.9	83.0 \pm 0.3	85.0 \pm 0.3	75.6 \pm 0.2
ERM-SDViT + T3A	CvT-21	32M	75.2 \pm 0.7	62.7 \pm 0.8	84.2 \pm 0.6	86.1 \pm 0.0	77.0 \pm 0.2
ERM-ViT	T2T-ViT-14	21.5M	70.2 \pm 0.5	59.0 \pm 0.6	81.9 \pm 0.3	83.6 \pm 0.6	73.7 \pm 0.2
ERM-SDViT	T2T-ViT-14	21.5M	71.1 \pm 0.5	59.2 \pm 0.3	82.8 \pm 0.4	83.5 \pm 0.3	74.2 \pm 0.3
ERM-SDViT + T3A	T2T-ViT-14	21.5M	70.8 \pm 0.4	61.9 \pm 0.7	84.1 \pm 0.2	85.0 \pm 0.3	75.5 \pm 0.2
Dataset			TerraIncognita				
Model	Backbone	# of Params	location_38	location_43	location_46	location_100	Average
ERM	ResNet-50	23.5M	56.3 \pm 1.1	36.8 \pm 4.6	52.6 \pm 0.4	35.2 \pm 1.7	45.2 \pm 1.2
ERM-ViT	DeiT-Small	22M	50.2 \pm 1.4	30.6 \pm 0.9	53.2 \pm 0.2	39.6 \pm 1.0	43.4 \pm 0.5
ERM-SDViT	DeiT-Small	22M	55.9 \pm 1.7	31.7 \pm 2.6	52.2 \pm 0.3	37.4 \pm 0.6	44.3 \pm 1.0
ERM-SDViT + T3A	DeiT-Small	22M	53.8 \pm 1.2	36.2 \pm 1.0	51.1 \pm 1.0	38.5 \pm 1.3	44.9 \pm 0.4
ERM-ViT	CvT-21	32M	51.4 \pm 1.8	40.1 \pm 1.7	57.6 \pm 1.0	45.7 \pm 0.6	48.7 \pm 0.4
ERM-SDViT	CvT-21	32M	53.6 \pm 3.3	42.7 \pm 1.6	58.2 \pm 1.0	44.5 \pm 1.8	49.7 \pm 1.4
ERM-SDViT + T3A	CvT-21	32M	58.1 \pm 0.7	46.2 \pm 0.3	57.0 \pm 1.0	44.1 \pm 2.2	51.4 \pm 0.7
ERM-ViT	T2T-ViT-14	21.5M	52.5 \pm 1.7	43.0 \pm 1.3	53.7 \pm 1.1	43.0 \pm 1.6	48.1 \pm 0.2
ERM-SDViT	T2T-ViT-14	21.5M	57.2 \pm 2.9	45.4 \pm 2.4	57.7 \pm 0.8	41.9 \pm 0.4	50.6 \pm 0.8
ERM-SDViT + T3A	T2T-ViT-14	21.5M	59.3 \pm 1.2	48.2 \pm 1.0	53.1 \pm 0.9	41.5 \pm 0.2	50.5 \pm 0.6

Table 7: Comparison of target domain-wise classification accuracy on **DomainNet** dataset. Results are reported of our method with the baseline using three different ViT backbones, including DeiT-Small [39], CvT-21 [38], and T2T-ViT-14 [43], and a DG baseline (T3A [61]).

Dataset	DomainNet								
Model	Backbone	# of Params	Clipart	Infograph	Painting	Quickdraw	Real	Sketch	Average
ERM	ResNet-50	23.5M	57.6 \pm 0.6	18.5 \pm 0.3	45.9 \pm 0.7	11.6 \pm 0.1	59.5 \pm 0.3	48.6 \pm 0.3	40.3 \pm 0.1
ERM-ViT	DeiT-Small	22M	62.9 \pm 0.2	23.3 \pm 0.1	53.1 \pm 0.2	15.7 \pm 0.1	65.7 \pm 0.1	52.4 \pm 0.2	45.5 \pm 0.0
ERM-SDViT	DeiT-Small	22M	63.4 \pm 0.1	22.9 \pm 0.0	53.7 \pm 0.1	15.0 \pm 0.4	67.4 \pm 0.1	52.6 \pm 0.2	45.8 \pm 0.0
ERM-SDViT + T3A	DeiT-Small	22M	64.3 \pm 0.2	23.7 \pm 0.0	54.2 \pm 0.3	19.7 \pm 0.4	69.6 \pm 0.1	53.2 \pm 0.2	47.4 \pm 0.1
ERM-ViT	CvT-21	32M	69.0 \pm 0.2	27.2 \pm 0.2	58.4 \pm 0.2	17.1 \pm 0.3	71.6 \pm 0.1	59.2 \pm 0.3	50.4 \pm 0.1
ERM-SDViT	CvT-21	32M	68.9 \pm 0.1	26.7 \pm 0.3	58.0 \pm 0.1	17.3 \pm 0.1	71.9 \pm 0.0	59.5 \pm 0.3	50.4 \pm 0.0
ERM-SDViT + T3A	CvT-21	32M	69.7 \pm 0.1	27.6 \pm 0.2	58.7 \pm 0.1	23.0 \pm 0.1	73.6 \pm 0.2	59.6 \pm 0.1	52.0 \pm 0.0
ERM-ViT	T2T-ViT-14	21.5M	67.0 \pm 0.3	25.2 \pm 0.2	55.3 \pm 0.3	15.3 \pm 0.2	70.3 \pm 0.1	55.9 \pm 0.2	48.1 \pm 0.1
ERM-SDViT	T2T-ViT-14	21.5M	67.6 \pm 0.2	25.0 \pm 0.2	55.8 \pm 0.4	15.2 \pm 0.3	70.0 \pm 0.1	55.9 \pm 0.1	48.2 \pm 0.2
ERM-SDViT + T3A	T2T-ViT-14	21.5M	68.2 \pm 0.1	25.8 \pm 0.2	56.7 \pm 0.3	20.7 \pm 0.2	72.4 \pm 0.1	57.0 \pm 0.2	50.2 \pm 0.1




Fig. 11: Comparison of attention maps between the baseline (ERM-ViT) and our proposed method (ERM-SDViT) on four target domains of VLCS and OfficeHome datasets. The ViT backbone is DeiT-Small.




Fig. 12: Comparison of attention maps between the baseline (ERM-ViT) and our proposed method (ERM-SDViT) on six target domains of DomainNet datasets. The ViT backbone is DeiT-Small.

Training overhead on target domains of other DG datasets: Table 8 and 9 reports training overhead, computed as relative % increase in training time (hrs.) on TerraIncognita and DomainNet datasets. The numbers report the training time increase introduced by our method on top of the baseline. The results shows that in both large-scale DG benchmark datasets i.e. TerraIncognita (24K images) and DomainNet (500K images), our model (ERM-SDViT) is not exceeding more than 20% relative overhead training time. Note that this training time could differ with GPU utilization. Results are reported with DeiT-Small (22M params.) backbone on Nvidia RTX A6000 GPU.

Table 8: Training overhead, computed as relative % increase in training time (hrs.), introduced by our method on top of the baseline.

Dataset:	TerraIncognita			
Model	Location_38	Location_43	Location_46	Location_100
ERM-ViT	0.268	0.268	0.270	0.268
ERM-SDViT	0.276	0.282	0.282	0.302
Rel.overhead	2.975	5.068	4.447	12.620

Table 9: Training overhead, computed as relative % increase in training time (hrs.), introduced by our method on top of the baseline.

Dataset:	DomainNet					
Model	Clipart	Infograph	Painting	Quickdraw	Real	Sketch
ERM-ViT	0.418	0.423	0.422	0.430	0.436	0.430
ERM-SDViT	0.482	0.444	0.510	0.469	0.446	0.460
Rel.overhead	15.376	5.124	20.928	9.079	2.463	7.089