

Learn to Preserve and Diversify: Parameter-Efficient Group with Orthogonal Regularization for Domain Generalization

Jiajun Hu^{1,2}, Jian Zhang^{1,2}, Lei Qi^{3,*}, Yinghuan Shi^{1,2,*}, Yang Gao^{1,2}

¹ State Key Laboratory for Novel Software Technology, Nanjing University, China

² National Institute of Healthcare Data Science, Nanjing University, China

³ School of Computer Science and Engineering, Key Lab of Computer Network and Information Integration (Ministry of Education), Southeast University, China

Abstract. Domain generalization (DG) aims to avoid the performance degradation of the model when the distribution shift between the limited training data and unseen test data occurs. Recently, foundation models with enormous parameters have been pre-trained with huge datasets, demonstrating strong generalization ability and showing promising direction for solving the DG problem. However, fully Fine-Tuning (FT) the foundation models results in unsatisfactory out-of-distribution accuracy due to the destroyed pre-trained generalized features. Recently, Parameter-Efficient Fine-Tuning (PEFT) alleviates the above problem by fine-tuning a small portion of the model parameters while keeping the rest frozen, which achieves better generalization performance compared to FT. Nevertheless, PEFT still suffers from the issue of overfitting to the training domains. To address the above issue, we propose Parameter-Efficient Group with Orthogonal regularization (PEGO) for vision transformers, which effectively preserves the generalization ability of the pre-trained network and learns more diverse knowledge compared with conventional PEFT. Specifically, we inject a group of trainable Low-Rank Adaptation (LoRA) modules into the pre-trained model and propose an orthogonal regularization loss to enhance the generalization ability of the model. Our framework achieves SOTA performance on five DG benchmarks, while only requiring training a small number of parameters without adding additional testing cost.

1 Introduction

Traditional machine learning algorithms assume that training data and test data come from independent and identical distributions [52]. However, a trained model suffers significant performance degradation when the distribution discrepancy (*a.k.a* domain gap) between the training and test data is large. To address this issue, Domain Generalization (DG) is proposed, which assumes that the model

* Corresponding Authors: Yinghuan Shi (syh@nju.edu.cn) and Lei Qi (qilei@seu.edu.cn).

trained with only source domains can generalize well in the unseen target domains. Previous DG works [1, 2, 23, 25, 45] mostly aim to facilitate model generalizability by either extracting domain-invariant features that are applicable across domains or augmenting training source domains with manually defined transformations, *e.g.*, style transfer [63, 75], Fourier transformation [26, 66], *etc.*

However, most DG works follow the same training strategy that directly fine-tunes a pre-trained model (*e.g.*, ResNet [27] pre-trained on ImageNet [18]) without considering the influence of the initial parameters on the ultimate generalization performance of finally-trained models. For example, we find that when using a randomly initialized ResNet-50 as the backbone, the average performance of Empirical Risk Minimization (ERM that trains a model by simply aggregating all source domains without any other techniques) on the PACS [42] dataset is only 35.8%. This result is much lower than the performance of ResNet-50 pre-trained on ImageNet, which is 84.2%. Particularly, on the domain *sketch* in PACS, where the distribution is far from ImageNet, fine-tuning the pre-trained ResNet-50 has a 55.3% performance improvement (24.0% \rightarrow 79.3%) than fine-tuning the randomly initialized ResNet-50. This indicates that *the generalizable knowledge from pre-trained models should be fully exploited for better generalization for downstream DG tasks.*

Recently, due to the flourishing development of deep learning, both the parameters and training data of deep models have largely increased and everyone can easily access these pre-trained large foundation models [8] based on the vision transformer architecture (ViT) [19]. Previous works [34, 40, 68] have shown that the vision transformer is more robust to unknown distributions than CNN [38], and CLIP [53] model trained on 400M image-text pairs has demonstrated strong zero-shot generalization ability. However, a critical question arises that *does directly fine-tuning stronger models lead to better results on DG tasks?* The answer is **NO**. Cha *et al.* [14] found that on the DomainBed [24] benchmark, the performance of fine-tuning ViT-B [19] pre-trained from CLIP [53] is lower than that of fine-tuning ResNet-50 pre-trained on ImageNet (ViT-B: 61.1% *vs.* ResNet-50: 64.2%). This is because direct fine-tuning distorts the generalizable features that originally reside in the pre-trained model [37]. Compared to the small models, a large number of parameters of the foundation models cause more severe overfitting issues when training with limited source domain data. Previous DG methods [1, 2, 23, 25, 45] mainly focus on how to extract domain-invariant features from limited source domains or perform data augmentation to generate more training data, ignoring how to preserve and exploit the generalization ability of the pre-trained models itself to improve the out-of-distribution generalization performance. Furthermore, with a huge number of parameters in foundation models, fine-tuning these parameters requires high training overhead in both GPU memory and training time, which significantly increases the difficulty of successfully fine-tuning a foundation model for users with limited resources.

To address the abovementioned two issues, Parameter-Efficient Fine-Tuning (PEFT) [28, 29, 32] has attracted significant interest in various language and visual tasks. Compared to full Fine-Tuning (FT), PEFT methods inject lightweight

trainable modules into the pre-trained model and freeze all the parameters of the pre-trained model. This approach reduces training overhead and achieves comparable or better performance than FT on downstream tasks [28]. Low-Rank Adaptation (LoRA) [29] is one of the most commonly employed PEFT implementations, which injects trainable rank decomposition matrices into every layer of the transformer [57]. Moreover, we discover that LoRA demonstrates substantial performance enhancement in addressing out-of-distribution tasks compared to FT and LoRA outperforms some conventional DG algorithms (In Sec. 4.3).

Despite the advantages of little computational overhead and overfitting alleviation, applying LoRA for DG bears two limitations. Firstly, although LoRA only injects little parameters into a foundation model that aims to alleviate the feature distortion problem, it still suffers from this problem since the features learned from the LoRA module may conflict with the feature of the pre-trained model, resulting in knowledge forgetting. Secondly, LoRA only employs a single low-rank module in each layer, which also easily overfits the training domain and cannot handle various unseen domains, further limiting its generalization performance.

To address the above limitations of LoRA, we propose a novel **Parameter Efficient Group with Orthogonal** regularization (**PEGO**) framework to fully exploit the potential of the pre-trained foundation models to solve the DG problem. First, we **Preserve** the generalization ability of the pre-trained model learned from the large-scale pre-training by imposing an orthogonal regularization loss between the pre-trained weights and the weights of LoRA layers. In this way, we can effectively minimize the distortion of the pre-trained generalized features. Second, we employ a group of LoRA modules for each layer to **Diversify** feature representations during training. With the learned abundant features, the model can better handle various unseen domains during the test. To further encourage diversity, the orthogonal constraints are also added between the weights of these LoRA modules. We summarize the contributions of this work as follows:

1. We propose a novel PEFT framework named PEGO that can effectively alleviate the overfitting issue with little computational overhead.
2. We design an orthogonal regularization loss to facilitate knowledge preservation of the pre-trained model.
3. We design the LoRA group with diversity constraints to learn diverse features to handle various unseen domains.
4. On five DomianBed benchmarks, PEGO achieves state-of-the-art performance compared to previous DG algorithms. Moreover, our method outperforms other PFET methods and the methods exploiting pre-trained models.

2 Related Work

2.1 Domain Generalization

Domain generalization (DG) aims to ensure that a model trained on source domains can generalize well to any unseen target domains. There are various

methods for DG problem, mainly including data augmentation [51,62,66,72,73,75], meta-learning [5,20,41,61,70], ensemble learning [3,13,54,74], self-supervised learning [12,59], adversarial learning [23,43,45], causal learning [1,2,47], test time adaptation [15,31,69,71], *etc.*

Most of the previous DG works choose small models (*e.g.*, AlexNet [36] and ResNet [27]) as the pre-trained backbone. Different from them, several methods exploit large pre-trained models for out-of-distribution generalization. MIRO [14] proposes mutual information regularization by assuming the pre-trained model as the oracle model. GESTUR [39] designs a gradient estimation to reduce potential risks in unseen domains utilizing a large pre-trained model. These two methods require significant training costs due to optimizing all the parameters of the pre-trained model. Moreover, there are some recent works [17,30,48,56] that utilize the text information from vision-language models to enhance the generalization ability of the fine-tuned model, but these methods rely on the jointly trained text encoder and visual encoder.

2.2 Parameter-Efficient Fine-Tuning

Parameter-Efficient Fine-Tuning (PEFT) [28] is first proposed to fine-tune large pre-trained transformers in natural language processing tasks, effectively reducing the computational and storage cost. PEFT methods inject some lightweight modules into the foundation model and only optimize a small portion of the model parameters to achieve a similar or higher performance compared with FT on downstream tasks. PEFT methods are also applied to deal with visual tasks. For example, VPT [32] introduces trainable prompt tokens in the input space for ViT [19]. One of the most influential PEFT methods is LoRA [29], which inserts trainable low-rank decomposition matrices into the transformer block while freezing the pre-trained model weights. It does not introduce additional inference latency and works with any neural network with dense layers.

2.3 Orthogonal Regularization

Orthogonality in neural networks [6,10,11,65] has been widely studied to improve training stability, training speed, and the performance of the model. For example, Xie *et al.* [65] utilize the orthonormality among different filter banks to mitigate the gradient vanishing on training deep convolutional neural networks (CNN). Bansal *et al.* [6] develop novel orthogonality regularizations, which achieve better accuracy, and faster and more stable convergences on several CNNs. In the DG literature, there are some works [4,9,16,33] applying orthogonal loss to disentangle the domain-invariant representations and domain-specific representations, but these methods usually require designing the additional network to generate two decoupled representations. Orthogonality is also extensively used in continual learning to prevent catastrophic forgetting of past tasks [22,55,60,67]. The most relevant to our work is O-LoRA [60], which utilizes orthogonal subspace learning for continual learning in language models. Different from their task, this paper aims to solve the DG problem when fine-tuning the foundation model and

we additionally consider the relationship between the pre-trained weights and the injected LoRA layers.

3 Methods

3.1 Problem Formulation

For the classification task in the DG setting, the training data D_{tr} usually contains n_d domains $\{D_i\}_{i=1}^{n_d}$ ($n_d > 1$). The test data D_{te} is not accessible during training but shares the same category space as D_{tr} . We define that x_i is the i -th sample in D_{tr} with the category label y_i . The distribution of training data P_{tr} is the joint distribution over image space and label space in D_{tr} , and there is a distribution gap between P_{tr} and the distribution of test data P_{te} .

The entire model includes a feature extractor $F(\cdot; \theta)$ parameterized by θ and a classifier $H(\cdot; \psi)$ parameterized by ψ . Previous DG works usually choose a pre-trained model and fine-tune it on the distribution of training data P_{tr} . The goal of DG is that F and H trained on D_{tr} can generalize well on D_{te} .

3.2 Revisit Previous Methods

Fine-Tuning. Fine-Tuning (FT, updating all the parameters of the model) is the most common training manner in previous DG works. The optimization objective of FT can be defined as the following formula:

$$\min_{\theta, \psi} \mathcal{L}_{cls} = \mathcal{L}_{CE}(H(F(x; \theta); \psi), y), \quad (1)$$

where $(x, y) \sim P_{tr}$ and \mathcal{L}_{CE} is the cross-entropy loss. Gulrajani *et al.* [24] claim that ERM based on FT has competitive performance compared to most DG algorithms. However, Kumar *et al.* [37] argue that FT distorts pre-trained features, which leads to poor out-of-distribution performance. In Sec. 4.3, we also find that the performances of ERM are much lower than those of other methods when choosing a foundation model as the backbone.

Parameter-Efficient Fine-Tuning. Compared with FT, Parameter-Efficient Fine-Tuning (PEFT) only updates a small number of parameters. For any PEFT method, the network with PEFT modules can be defined as $G(\cdot; \Phi)$ parameterized by Φ , where the trainable parameters are $\phi \subseteq \Phi$ (the parameter number of ϕ is much smaller than that of Φ). The general optimization objective of any PEFT method can be formulated as follows:

$$\min_{\phi, \psi} \mathcal{L}_{cls} = \mathcal{L}_{CE}(H(G(x; \Phi); \psi), y). \quad (2)$$

Due to not updating any parameters of the foundation model, PEFT effectively inherits the generalization ability of the foundation model. However, in the DG setting, the trainable parameters in PEFT still suffer from overfitting to the source domains, resulting in a performance decrease on the unseen test domain. At the same time, the injected PEFT modules also partially distort the generalized features produced by the foundation model during training.

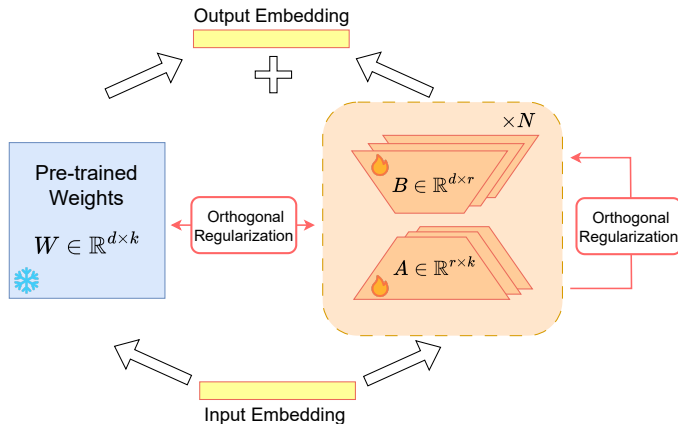


Fig. 1: Illustration of our method: **Parameter Efficient Group with Orthogonal regularization (PEGO)**. Different from previous DG work updating all the parameters of the pre-trained model, we freeze the parameters of the model and inject a group of trainable parameter-efficient modules into it. Moreover, we apply an orthogonal regularization loss between the pre-trained weights and the LoRA modules to preserve the generalization ability of the pre-trained model (*Learn to Preserve*) and employ another orthogonal regularization loss on different LoRA modules within the group to encourage them to learn diverse knowledge during training (*Learn to Diversify*).

3.3 Parameter-Efficient Group with Orthogonal Regularization

To address the issue of FT and PEFT in DG, we propose **Parameter Efficient Group with Orthogonal regularization (PEGO)** based on LoRA [29], which is a classic PEFT method for pre-trained transformer. In this section, we first review the key technologies of LoRA and then introduce the details of our method.

LoRA. Hu *et al.* [29] propose **Low-Rank Adaptation (LoRA)**, utilizing the hypothesis that the pre-trained weights have a low intrinsic rank. Specifically, for a linear layer in the pre-trained model, which is parameterized by $W \in \mathbb{R}^{d \times k}$, there is a low-rank decomposition $W + \Delta W = W + BA$, where $B \in \mathbb{R}^{d \times r}$, $A \in \mathbb{R}^{r \times k}$ and $r \ll \min(d, k)$. During training, LoRA only updates injected low-rank weights A and B while keeping W frozen. We define the input feature of the linear layer as z_{in} , and the output feature of the linear layer can be calculated through the following forward process:

$$z_{out} = (W + \Delta W)z_{in} = (W + BA)z_{in} = Wz_{in} + BAz_{in}. \quad (3)$$

In the original paper, the LoRA layers are only injected into the W^q and W^v , which are the query and value projection matrices of the self-attention modules. After training, we can directly compute $W_{final} = W + BA$ as the final weight. Therefore, there is no additional inference latency in LoRA.

PEGO. To further improve the out-of-distribution performance of LoRA, we propose the idea of learning to **Preserve** and **Diversify**. The former aims to preserve the generalization performance of the pre-trained model during fine-tuning, while the latter aims to learn more diverse knowledge from source domains. Specifically, we propose to inject a group of LoRA modules into the pre-trained network and apply an orthogonal regularization loss to achieve the above two goals simultaneously. Fig. 1 provides an illustration of our method.

- **Learn to Preserve**

Inspired by the orthogonal gradient updating [22, 55, 67] used to prevent catastrophic forgetting, we strive to constrain the gradient subspace of fine-tuning to be orthogonal to the gradient subspace of large-scale pre-training task. This enables the model to learn useful information in the source domains while preserving the generalization ability of the pre-trained model. However, we may not be able to access the pre-training dataset and it is not practical to calculate the full gradient of the large-scale pre-training task. Following Wang *et al.* [60] who regard the weights of the LoRA layer as the gradient subspace of a certain task, we propose to consider the original pre-trained weights as the gradient subspace of the pre-training task. Similarly, the fine-tuned LoRA weights can be considered as the gradient subspace updated on the source domains. With these two gradient subspaces, we propose an orthogonal regularization loss to constrain the pre-trained weights W orthogonal to the weights of injected module BA . Specifically, the loss can be formulated as follows:

$$\mathcal{L}_{preserve} = \|W^T(BA)\|_1, \quad (4)$$

where $\|\cdot\|_1$ is $L1$ norm. As the same as LoRA, we only fine-tune the low-rank weight matrices A and B while keeping the rest parameters frozen during training.

Furthermore, we analyze the above loss from the perspective of the feature-level. According to the forward process in Eq. (3), we define the output feature of the pre-trained weight $z_{init} = Wz_{in}$ and the output feature of the LoRA layer $z_{new} = BAz_{in}$. While the loss restricts BA orthogonal to W , it indirectly constrains z_{init} orthogonal to z_{new} . This can be demonstrated by the following transpose transformation:

$$z_{init}^T z_{new} = (Wz_{in})^T BAz_{in} = z_{in}^T (W^T BA) z_{in}. \quad (5)$$

Since the features generated by fine-tuning are encouraged to be orthogonal to pre-trained features, the generalization ability of the pre-trained model is preserved well. At the same time, during the implementation, we discovered that optimizing weights to be orthogonal requires fewer computational resources and results in better performance compared to optimizing features to be orthogonal.

- **Learn to Diversify**

In the DG literature, increasing the diversity of the training trajectories is used to improve the generalization performance of the model. For example, Zhang *et al.* [70] propose a multi-view algorithm for employing multiple optimization trajectories; Arpit *et al.* [3] perform the ensemble of multiple independently trained models. However, these methods require additional multi-step

gradient updates or training multiple models. Benefiting from the lightweight and easily scalable characteristics of LoRA, we propose to introduce multiple LoRA modules and apply orthogonal regularization to facilitate the model learning diverse knowledge.

Different from the original LoRA where the pre-trained weight W comes with only one trainable low-rank matrices BA , we employ a parameter-efficient group of LoRA layers $g = \{A_i, B_i\}_{i=1}^N$ in our framework, where N is the number of LoRA layers. Moreover, we adopt a pairwise orthogonal regularization loss to enhance the diversity of knowledge learned by each LoRA layer. Specifically, for a certain LoRA layer $g_i = \{A_i, B_i\}$, its weight matrix $B_i A_i$ is encouraged to be orthogonal to the weight matrix of other LoRA layers $\{B_j A_j\}_{j \neq i}^N$. Formally, the loss that aims to learn to diversify is defined as follows:

$$\mathcal{L}_{diversify} = \sum_{i=1}^N \sum_{j=i+1}^N \|(B_i A_i)^T (B_j A_j)\|_1. \quad (6)$$

The above loss promotes the orthogonality between the weights of different LoRA modules, resulting in the output features of different LoRA modules are also encouraged to orthogonal. Our method learns more diverse optimization trajectories compared to the original LoRA and only increases little training cost.

- **Final Object**

We combine $\mathcal{L}_{preserve}$ and $\mathcal{L}_{diversify}$ as the optimization objective of orthogonal regularization, which takes the following form:

$$\mathcal{L}_O(W) = \sum_{i=1}^N \|W^T (B_i A_i)\|_1 + \sum_{i=1}^N \sum_{j=i+1}^N \|(B_i A_i)^T (B_j A_j)\|_1. \quad (7)$$

In line with LoRA, we only apply the above loss to W^q and W^v in each block of the pre-trained transformer. During training, we only fine-tune the LoRA group and the classification head while keeping the rest parameters frozen. The final orthogonal regularization loss can be given by:

$$\mathcal{L}_{OR} = \sum_{b=1}^B (\mathcal{L}_O(W_b^q) + \mathcal{L}_O(W_b^v)), \quad (8)$$

where B is the number of blocks in ViT, W_b^q and W_b^v indicates the query and value projection matrices in the b -th block. Finally, we combine \mathcal{L}_{cls} with \mathcal{L}_{OR} as the final optimization object of the model:

$$\mathcal{L}_{final} = \mathcal{L}_{cls} + \alpha \mathcal{L}_{OR}, \quad (9)$$

where α is the balancing coefficient of two losses. When the model is deployed to the test environment after training, we merge the group of LoRA layers to the pre-trained weight:

$$W_{final} = W + \sum_{i=1}^N B_i A_i. \quad (10)$$

As the same as LoRA, there is no additional testing latency in our method.

Table 1: Performance comparison with DG methods. Leave-one-domain-out accuracy (%) on five DomainBed benchmarks. In addition to the results of our method, other results come from Lew *et al.* [39]. OH, TI and DN indicate OfficeHome, TerraIncognita, and DomainNet, respectively (similarly hereinafter).

Algorithm	PACS	VLCS	OH	TI	DN	Avg
ERM (FT)	83.4±0.5	75.9±1.3	66.4±0.5	35.3±0.8	44.4±0.6	61.1
SWAD [13]	91.3±0.1	79.4±0.4	76.9±0.1	45.4±0.5	51.7±0.8	68.9
SMA [3]	92.1±0.2	79.7±0.2	78.1±0.1	48.3±0.7	55.9±0.2	70.8
MIRO [14]	95.6±0.8	82.2±0.3	82.5±0.1	54.3±0.4	54.0±0.3	73.7
GESTUR [39]	96.0±0.0	82.8±0.1	84.2±0.1	55.7±0.2	58.9±0.1	75.5
Ours	96.5±0.1	83.2±0.3	84.2±0.1	57.3±0.3	59.3±0.1	76.1

4 Experiments

4.1 DataSets

We use five common datasets in DomainBed [24] evaluation benchmarks to verify the effectiveness of our method: **1. PACS** [42] includes 9,991 images, 7 categories, and 4 domains. The domain shift between each domain is large (*e.g.*, *photo* and *sketch*). **2. VLCS** [21] includes 10,729 images, 5 categories, and 4 domains. The domain shift between each domain is mainly from different viewpoints. **3. OfficeHome** [58] includes 15,500 images, 65 categories, and 4 domains. It contains more categories and smaller domain shift than PACS. **4. TerraIncognita** [7] includes 24,788 images, 10 categories, and 4 domains. The images are taken in four different wild locations and it’s a challenging dataset. **5. DomainNet** [50] includes 586,575 images, 345 categories, and 6 domains. Its number of images and categories far exceeds the above datasets.

4.2 Implementation Details

Different from DomainBed [24] using ResNet-50 [27] as the backbone, we utilize ViT-B/16 [19] pre-trained by CLIP [53] as the default foundation model. During training, we employ the Adam [35] optimizer with a learning rate of 5e-4 for 5000 iterations, except for the DomainNet, which requires 15000 iterations for convergence. The 32 images in each source domain construct a whole batch. Our data augmentation includes random horizontal flip, color jittering, and random graying in DomainBed. We set the rank of LoRA r to 4 and the balancing coefficient α to 1e-3. For the numbers of LoRA layers N in each group, we define the hyperparameter search space as $N \in \{2, 4, 6\}$. More details about the evaluation protocol and hyperparameters search can be found in Supplementary.

4.3 Main Results

Comparison with DG Methods. We first compare PEGO with the baseline ERM (*i.e.*, FT) and state-of-the-art (SOTA) DG methods, including: SWAD [13],

Table 2: Performance comparison with PEFT methods. Leave-one-domain-out accuracy (%) on five DomainBed benchmarks.

Algorithm	PACS	VLCS	OH	TI	DN	Avg
ERM (FT)	83.4±0.5	75.9±1.3	66.4±0.5	35.3±0.8	44.4±0.6	61.1
Adapter [28]	92.0±0.5	79.8±0.4	72.9±0.4	44.4±0.8	56.2±0.1	69.1
LoRA [29]	96.0±0.1	82.7±0.0	83.4±0.1	54.8±0.6	58.1±0.1	75.0
VPT [32]	96.2±0.3	82.9±0.3	83.4±0.3	54.2±0.7	58.9±0.1	75.1
Ours	96.5±0.1	83.2±0.3	84.2±0.1	57.3±0.3	59.3±0.1	76.1

SMA [3], MIRO [14] and GESTUR [39]. SWAD and SMA are both ensemble methods that show significant improvement compared to ERM on DomainBed using ResNet-50 as the backbone; MIRO and GESTUR both aim to preserve and exploit the generalization ability of the pre-trained network.

The results of the performance comparison are shown in Tab. 1. Compared to ERM, all DG approaches have significantly improved, indicating much potential for further enhancement when using foundation models for the DG problem. Furthermore, PEGO outperforms all previous methods on all five datasets, demonstrating the superiority of our method. Especially on the challenging TerraIncognita dataset, PEGO achieves a remarkable improvement which is more than 1.6% compared with other methods (55.7% \rightarrow 57.3%).

Comparison with PEFT Methods. In this subsection, we make a performance comparison between PEGO and several PEFT methods, containing Adapter [28] and LoRA [29] which are widely utilized in language tasks, VPT [32] which is designed specifically for vision transformer.

As shown in Tab. 2, the performances of all PEFT methods on five DomainBed benchmarks are significantly higher than that of FT. This indicates that full fine-tuning considerably harms the generalization performance of the pre-trained model, while PEFT can effectively address this issue. Moreover, our method achieves state-of-the-art performance on five datasets and 1.1% improvement of average performances compared to LoRA which is the basis of our method (75.0% \rightarrow 76.1%).

Comparison with Methods Exploiting Pre-trained Models. There are some works in other fields utilizing pre-trained models to improve the generalization ability of the fine-tuned model. Following Cha *et al.* [14], we select several classic methods to make a performance comparison with PEGO. Specifically, L²-SP [44] employs L^2 penalty between the pre-trained model and fine-tuned model during training; LP-FT [37] proposes a strategy of first linear probing and then full fine-tuning; LwF [46] constrains the outputs of fine-tuned model for old tasks to be similar to that of pre-trained network; WiSE-FT [64] ensembles the weights of the zero-shot and fine-tuned networks to realize robust fine-tuning.

We report the results of all methods on four DomainBed benchmarks in Tab. 3. Our method achieves the best performance on PACS, OfficeHome, and TerraIncognita, except for VLCS, where it is only 0.7% lower than WiSE-FT.

Table 3: Performance comparison with methods exploiting pre-trained models. Leave-one-domain-out accuracy (%) on four DomainBed benchmarks. The results of WiSE-FT come from Lew *et al.* [39] and we report the rest results. The best and second-best accuracy are **bolded** and underlined, respectively.

Algorithm	PACS	VLCS	OH	TI	Avg
ERM (FT)	83.4±0.5	75.9±1.3	66.4±0.5	35.3±0.8	65.3
L ² -SP [44]	92.2±0.7	81.0±0.2	68.2±0.5	39.4±1.6	70.2
LP-FT [37]	94.2±0.7	77.5±0.4	72.0±0.4	39.0±1.5	70.7
LwF [46]	93.6±0.6	81.9±0.4	80.7±0.4	39.4±0.6	73.9
WiSE-FT [64]	<u>94.5±0.0</u>	83.9±0.3	<u>83.9±0.2</u>	<u>47.5±1.2</u>	<u>77.5</u>
Ours	96.5±0.1	<u>83.2±0.3</u>	84.2±0.1	57.3±0.3	80.3

Table 4: Ablation study on two orthogonal regularization losses. Leave-one-domain-out accuracy (%) on PACS and OfficeHome.

	$\mathcal{L}_{preserve}$	$\mathcal{L}_{diversify}$	PACS	OH
PEGO	✓	✓	96.55±0.11	84.21±0.10
	✓		96.14±0.25	82.95±0.02
		✓	96.37±0.20	83.76±0.08
			95.34±0.30	82.85±0.07
LoRA [29]			95.96±0.12	83.41±0.11

Furthermore, PEGO outperforms the previous best method by 2.8% on the average accuracy of four benchmarks (77.5% \rightarrow 80.3%). This is mainly because PEGO has over 9.8% improvement on the TerraIncognita dataset (47.5% \rightarrow 57.3%), and its standard error is significantly smaller. The above results indicate that although previous methods achieve some improvement compared with FT and aim to preserve the generalization ability of pre-trained models, there is still a significant overfitting phenomenon when using these methods. Our method effectively alleviates the above issue.

5 Further Analysis

5.1 Ablation Study

Effectiveness of Two Orthogonal Losses. To verify the improvement of the model’s generalization performance by employing our proposed losses, we conduct ablation experiments about $\mathcal{L}_{preserve}$ and $\mathcal{L}_{diversify}$ on the PACS and OfficeHome benchmarks. As shown in Tab. 4, when $\mathcal{L}_{preserve}$ and $\mathcal{L}_{diversify}$ are applied simultaneously, the model achieves the best performance (blue row). We notice that the performance of injecting a group of LoRA layers into the pre-trained model without applying regularization loss (4th row) is worse than the performance of the original LoRA. This indicates that increasing the number of

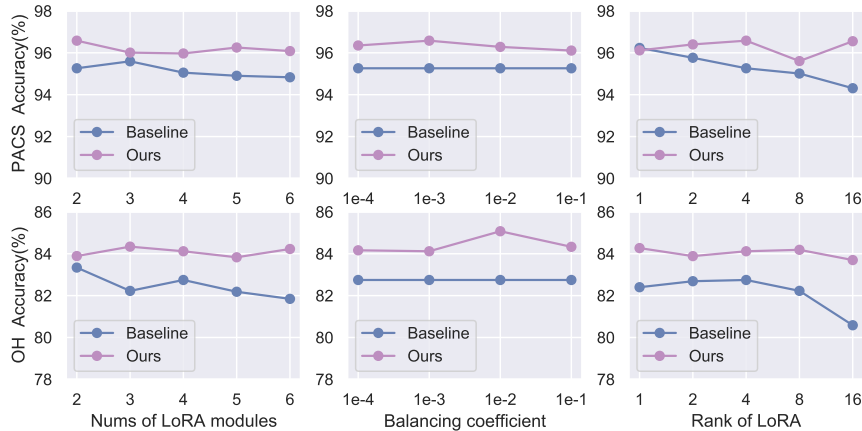


Fig. 2: Leave-one-domain-out accuracy (%) on PACS and OfficeHome when choosing different numbers of LoRA modules N , balancing coefficient α and rank of LoRA r . Baseline (blue line) indicates injecting a group of LoRA layers into the pre-trained model without applying $\mathcal{L}_{preserve}$ and $\mathcal{L}_{diversify}$ (*i.e.*, balancing coefficient is zero).

Table 5: Performance comparison with Zero-shot CLIP. Leave-one-domain-out accuracy (%) on four DomainBed benchmarks. The results of Zero-shot come from Lew *et al.* [39]. The best and second-best accuracy are **bolded** and underlined, respectively.

Algorithm	PACS	VLCS	OH	TI	Avg
ERM (FT)	83.4±0.5	75.9±1.3	66.4±0.5	35.3±0.8	65.3
Zero-shot [53]	96.8±0.0	<u>81.7±0.3</u>	<u>83.0±0.3</u>	31.3±0.2	<u>73.2</u>
Ours	<u>96.5±0.1</u>	83.2±0.3	84.2±0.1	57.3±0.3	80.3

training parameters without regularization cannot improve the model’s generalization performance, while our losses can effectively utilize more parameters.


Effects of Numbers of LoRA Modules. Intuitively, the more LoRA modules in our method, the higher the probability of learning diverse knowledge. However, excessive modules lead to complicating loss optimization and increasing training overhead. The first column of Fig. 2 shows the performances of PEGO and Baseline when choosing different numbers of modules. PEGO achieves higher accuracy and is more stable than the Baseline.

Effects of Balancing Coefficient. The second column of Fig. 2 shows the performances of Baseline and PEGO with different balancing coefficients. In a wide range of coefficients from $1e-4$ to $1e-1$, PEGO outperforms Baseline (balancing coefficient is zero), demonstrating the effectiveness of our proposed loss.

Effects of Rank of LoRA. As shown in the third column of Fig. 2, when the rank of the LoRA module is too high (greater than 8), the accuracy of Baseline significantly decreases, while the accuracy of our method remains stable.

Table 6: Leave-one-domain-out accuracy (%) on four DomainBed benchmarks when using ViT-L/14 pre-trained by CLIP as the backbone.

Algorithm	PACS	VLCS	OH	TI	Avg
ERM	88.0±4.1	77.5±0.6	53.0±3.2	43.3±0.7	65.5
LoRA [29]	98.1±0.0	83.7±0.3	87.9±0.1	52.7±0.8	80.6
PEGO	98.0±0.1	83.7±0.2	88.6±0.1	57.2±0.5	81.9

**Fig. 3:** The visualization of the feature space (before the classifier) extracted by FT model, pre-trained model, LoRA model, and our model when training the PACS dataset and the test domain is *art painting*.

5.2 Comparison with the Zero-shot Baseline

In Sec. 4.3, we choose the CLIP [53] pre-trained model as the default backbone. CLIP learns representations on 400 million image-text pairs and has demonstrated strong zero-shot ability on plenty of visual datasets.

Tab. 5 shows the performance comparison between our method and Zero-shot CLIP on four DomainBed benchmarks. In addition to being slightly lower than Zero-shot on the simple benchmark PACS by 0.3%, PEGO achieves the best performances on the other three benchmarks. It outperforms zero-shot by 7.1% in average performance (73.2% \rightarrow 80.3%). Besides, we notice that the accuracy of Zero-shot on TerraIncognita is worse than ERM and our method surpasses Zero-shot by 26% (31.3% \rightarrow 57.3%). This result is consistent with the finding of Cho *et al.* [17]. Although CLIP can leverage text information to achieve zero-shot without training data, we argue that source domain data is crucial for enhancing the generalization ability of the model and the key factor is whether robust fine-tuning can be accomplished.

5.3 Experiment Using ViT-L as the Backbone

In Sec. 4.3, we choose ViT-B/16 [19] pre-trained by CLIP [53] as the backbone for all the experiments. To verify the effectiveness of our method on larger models, we conduct the experiment using ViT-L/14 pre-trained by CLIP as the backbone. Tab. 6 provides the performances of ERM, LoRA [29] and PEGO on four DomainBed benchmarks. Both LoRA and PEGO outperform ERM significantly on all the benchmarks and achieve similar accuracy on PACS and VLCS. However, on the other two benchmarks, PEGO has a significant improvement compared to LoRA, especially on TerraIncognita (52.7% \rightarrow 57.2%).

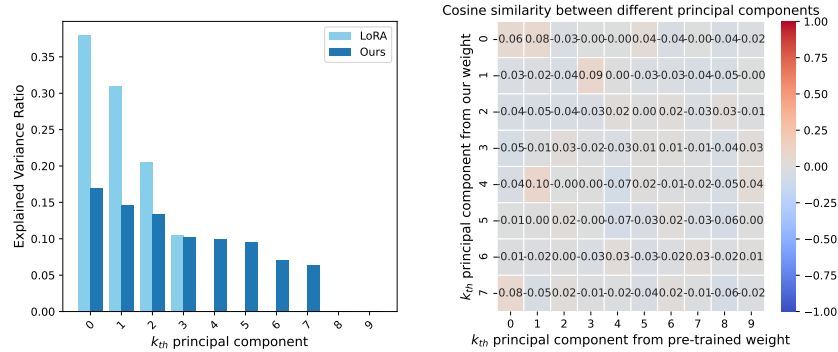


Fig. 4: **Left:** Explained Variance Ratio of the top-10 PCs in LoRA weight and the top-10 PCs in PEGO weight. **Right:** Cosine similarity between the top-10 PCs of pre-trained weight and the top-8 PCs of PEGO weight.

5.4 Visualization of Feature Space

To understand whether our method can “*Learn to Preserve*”, we visualize the difference in feature space between our method and the pre-trained model with PCA [49] and compare with full fine-tuning (FT) and LoRA. As shown in Fig. 3, fine-tuning all layers largely distorts the feature distribution and LoRA can partially alleviate it. Instead, our method successfully preserves pre-trained features by further using our orthogonal loss $\mathcal{L}_{preserve}$.

5.5 Principal Component Analysis on Weights

To confirm that our method achieves orthogonal regularization, we experiment by decomposing the model weights to get their principal components (PCs). Specifically, to verify the effect of “*Learn to Preserve*”, we expect the learned PCs to be orthogonal to the PCs of pre-trained weights, and to verify “*Learn to Diversify*”, we expect the learned PCs to be more than the original LoRA. As shown in Fig. 4, the weight of our model has more PCs (8 *vs.* 4) than LoRA and also exhibits orthogonal (zero cosine similarity) to the PCs of the pre-trained weights, validating the effectiveness of our proposed two losses.

6 Conclusion

In this paper, we address the problem of using foundation models in DG from a novel perspective of Learning to Preserve and Diversify. Specifically, we propose Parameter-Efficient Group with Orthogonal regularization (PEGO), which effectively preserves the generalization ability of pre-trained models and learns diverse knowledge. We conduct comparative experiments and ablation experiments to demonstrate the effectiveness and stability of PEGO. Our simple method can be applied to any neural network architecture with linear layers and is training-friendly without additional testing costs.

Acknowledgements

The work is supported by the NSFC Project (62222604, 62206052, 62192783), Jiangsu Natural Science Foundation Project (BK20210224), China Postdoctoral Science Foundation (2024M750424), the Fundamental Research Funds for the Central Universities (020214380120), and the State Key Laboratory Funds for Key Project (ZZKT2024A14).

References

1. Ahuja, K., Caballero, E., Zhang, D., Gagnon-Audet, J.C., Bengio, Y., Mitliagkas, I., Rish, I.: Invariance principle meets information bottleneck for out-of-distribution generalization. In: Conference on Neural Information Processing Systems (NeurIPS). pp. 3438–3450 (2021)
2. Arjovsky, M., Bottou, L., Gulrajani, I., Lopez-Paz, D.: Invariant risk minimization. arXiv:1907.02893 (2019)
3. Arpit, D., Wang, H., Zhou, Y., Xiong, C.: Ensemble of averages: Improving model selection and boosting performance in domain generalization. In: Conference on Neural Information Processing Systems (NeurIPS). pp. 8265–8277 (2022)
4. Bai, H., Sun, R., Hong, L., Zhou, F., Ye, N., Ye, H.J., Chan, S.H.G., Li, Z.: DecAug: Out-of-distribution generalization via decomposed feature representation and semantic augmentation. In: AAAI Conference on Artificial Intelligence (AAAI). pp. 6705–6713 (2021)
5. Balaji, Y., Sankaranarayanan, S., Chellappa, R.: Metareg: Towards domain generalization using meta-regularization. In: Conference on Neural Information Processing Systems (NeurIPS) (2018)
6. Bansal, N., Chen, X., Wang, Z.: Can we gain more from orthogonality regularizations in training deep networks? In: Conference on Neural Information Processing Systems (NeurIPS) (2018)
7. Beery, S., Van Horn, G., Perona, P.: Recognition in terra incognita. In: European Conference on Computer Vision (ECCV). pp. 456–473 (2018)
8. Bommasani, R., Hudson, D.A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M.S., Bohg, J., Bosselut, A., Brunskill, E., et al.: On the opportunities and risks of foundation models. arXiv:2108.07258 (2021)
9. Bousmalis, K., Trigeorgis, G., Silberman, N., Krishnan, D., Erhan, D.: Domain separation networks. In: Conference on Neural Information Processing Systems (NeurIPS) (2016)
10. Brock, A., Donahue, J., Simonyan, K.: Large scale gan training for high fidelity natural image synthesis. In: International Conference on Learning Representations (ICLR) (2019)
11. Brock, A., Lim, T., Ritchie, J., Weston, N.: Neural photo editing with introspective adversarial networks. In: International Conference on Learning Representations (ICLR) (2017)
12. Carlucci, F.M., D’Innocente, A., Bucci, S., Caputo, B., Tommasi, T.: Domain generalization by solving jigsaw puzzles. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2229–2238 (2019)
13. Cha, J., Chun, S., Lee, K., Cho, H.C., Park, S., Lee, Y., Park, S.: SWAD: Domain generalization by seeking flat minima. In: Conference on Neural Information Processing Systems (NeurIPS). pp. 22405–22418 (2021)

14. Cha, J., Lee, K., Park, S., Chun, S.: Domain generalization by mutual-information regularization with pre-trained models. In: European Conference on Computer Vision (ECCV). pp. 440–457 (2022)
15. Chen, L., Zhang, Y., Song, Y., Shan, Y., Liu, L.: Improved test-time adaptation for domain generalization. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 24172–24182 (2023)
16. Chen, Y., Wang, Y., Pan, Y., Yao, T., Tian, X., Mei, T.: A style and semantic memory mechanism for domain generalization. In: International Conference on Computer Vision (ICCV). pp. 9164–9173 (2021)
17. Cho, J., Nam, G., Kim, S., Yang, H., Kwak, S.: Promptstyler: Prompt-driven style generation for source-free domain generalization. In: International Conference on Computer Vision (ICCV). pp. 15702–15712 (2023)
18. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 248–255 (2009)
19. 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. In: International Conference on Learning Representations (ICLR) (2021)
20. Dou, Q., Coelho de Castro, D., Kamnitsas, K., Glocker, B.: Domain generalization via model-agnostic learning of semantic features. In: Conference on Neural Information Processing Systems (NeurIPS) (2019)
21. Fang, C., Xu, Y., Rockmore, D.N.: Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In: International Conference on Computer Vision (ICCV). pp. 1657–1664 (2013)
22. Farajtabar, M., Azizan, N., Mott, A., Li, A.: Orthogonal gradient descent for continual learning. In: International Conference on Artificial Intelligence and Statistics (AISTATS). pp. 3762–3773 (2020)
23. Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., Lempitsky, V.: Domain-adversarial training of neural networks. *Journal of Machine Learning Research (JMLR)* pp. 2096–2030 (2016)
24. Gulrajani, I., Lopez-Paz, D.: In search of lost domain generalization. In: International Conference on Learning Representations (ICLR) (2021)
25. Guo, J., Qi, L., Shi, Y.: Domaindrop: Suppressing domain-sensitive channels for domain generalization. In: International Conference on Computer Vision (ICCV). pp. 19114–19124 (2023)
26. Guo, J., Wang, N., Qi, L., Shi, Y.: ALOFT: A lightweight mlp-like architecture with dynamic low-frequency transform for domain generalization. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 24132–24141 (2023)
27. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 770–778 (2016)
28. Houlisby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., Gelly, S.: Parameter-efficient transfer learning for nlp. In: International Conference on Machine Learning (ICML). pp. 2790–2799 (2019)
29. Hu, E.J., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al.: LoRA: Low-rank adaptation of large language models. In: International Conference on Learning Representations (ICLR) (2022)

30. Huang, Z., Zhou, A., Ling, Z., Cai, M., Wang, H., Lee, Y.J.: A sentence speaks a thousand images: Domain generalization through distilling clip with language guidance. In: International Conference on Computer Vision (ICCV). pp. 11685–11695 (2023)
31. Iwasawa, Y., Matsuo, Y.: Test-time classifier adjustment module for model-agnostic domain generalization. In: Conference on Neural Information Processing Systems (NeurIPS). pp. 2427–2440 (2021)
32. Jia, M., Tang, L., Chen, B.C., Cardie, C., Belongie, S., Hariharan, B., Lim, S.N.: Visual prompt tuning. In: European Conference on Computer Vision (ECCV). pp. 709–727 (2022)
33. Jo, S.Y., Yoon, S.W.: Poem: Polarization of embeddings for domain-invariant representations. AAAI Conference on Artificial Intelligence (AAAI) pp. 8150–8158 (2023)
34. Kim, D., Wang, K., Sclaroff, S., Saenko, K.: A broad study of pre-training for domain generalization and adaptation. In: European Conference on Computer Vision (ECCV). pp. 621–638 (2022)
35. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: International Conference on Learning Representations (ICLR) (2015)
36. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Conference on Neural Information Processing Systems (NeurIPS) (2012)
37. Kumar, A., Raghunathan, A., Jones, R.M., Ma, T., Liang, P.: Fine-tuning can distort pretrained features and underperform out-of-distribution. In: International Conference on Learning Representations (ICLR) (2022)
38. LeCun, Y., Bottou, L., Bengio, Y., Haffner, P.: Gradient-based learning applied to document recognition. Proceedings of the IEEE pp. 2278–2324 (1998)
39. Lew, B., Son, D., Chang, B.: Gradient estimation for unseen domain risk minimization with pre-trained models. In: International Conference on Computer Vision Workshops (ICCVW). pp. 4436–4446 (2023)
40. Li, B., Shen, Y., Yang, J., Wang, Y., Ren, J., Che, T., Zhang, J., Liu, Z.: Sparse mixture-of-experts are domain generalizable learners. In: International Conference on Learning Representations (ICLR) (2023)
41. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.: Learning to generalize: Meta-learning for domain generalization. In: AAAI Conference on Artificial Intelligence (AAAI). pp. 3490–3497 (2018)
42. Li, D., Yang, Y., Song, Y.Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: International Conference on Computer Vision (ICCV). pp. 5542–5550 (2017)
43. Li, H., Pan, S.J., Wang, S., Kot, A.C.: Domain generalization with adversarial feature learning. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 5400–5409 (2018)
44. LI, X., Grandvalet, Y., Davoine, F.: Explicit inductive bias for transfer learning with convolutional networks. In: International Conference on Machine Learning (ICML). pp. 2825–2834 (2018)
45. Li, Y., Tian, X., Gong, M., Liu, Y., Liu, T., Zhang, K., Tao, D.: Deep domain generalization via conditional invariant adversarial networks. In: European Conference on Computer Vision (ECCV). pp. 624–639 (2018)
46. Li, Z., Hoiem, D.: Learning without forgetting. IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2017)
47. Mahajan, D., Tople, S., Sharma, A.: Domain generalization using causal matching. In: International Conference on Machine Learning (ICML). pp. 7313–7324 (2021)

48. Mao, X., Chen, Y., Jia, X., Zhang, R., Xue, H., Li, Z.: Context-aware robust fine-tuning. *International Journal of Computer Vision (IJCV)* pp. 1–16 (2023)
49. Pearson, K.: LIII. On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin philosophical magazine and journal of science* pp. 559–572 (1901)
50. Peng, X., Bai, Q., Xia, X., Huang, Z., Saenko, K., Wang, B.: Moment matching for multi-source domain adaptation. In: *International Conference on Computer Vision (ICCV)*. pp. 1406–1415 (2019)
51. Qi, L., Yang, H., Shi, Y., Geng, X.: Normaug: Normalization-guided augmentation for domain generalization. *IEEE Transactions on Image Processing (TIP)* pp. 1419–1431 (2024)
52. Quiñero-Candela, J., Sugiyama, M., Schwaighofer, A., Lawrence, N.D.: *Dataset shift in machine learning*. Mit Press (2008)
53. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: *International Conference on Machine Learning (ICML)*. pp. 8748–8763 (2021)
54. Rame, A., Kirchmeyer, M., Rahier, T., Rakotomamonjy, A., Gallinari, P., Cord, M.: Diverse weight averaging for out-of-distribution generalization. In: *Conference on Neural Information Processing Systems (NeurIPS)*. pp. 10821–10836 (2022)
55. Saha, G., Garg, I., Roy, K.: Gradient projection memory for continual learning. In: *International Conference on Learning Representations (ICLR)* (2021)
56. Shu, Y., Guo, X., Wu, J., Wang, X., Wang, J., Long, M.: CLIPood: Generalizing clip to out-of-distributions. In: *International Conference on Machine Learning (ICML)* (2023)
57. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. In: *Conference on Neural Information Processing Systems (NeurIPS)* (2017)
58. Venkateswara, H., Eusebio, J., Chakraborty, S., Panchanathan, S.: Deep hashing network for unsupervised domain adaptation. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5018–5027 (2017)
59. Wang, S., Yu, L., Li, C., Fu, C.W., Heng, P.A.: Learning from extrinsic and intrinsic supervisions for domain generalization. In: *European Conference on Computer Vision (ECCV)*. pp. 159–176 (2020)
60. Wang, X., Chen, T., Ge, Q., Xia, H., Bao, R., Zheng, R., Zhang, Q., Gui, T., Huang, X.J.: Orthogonal subspace learning for language model continual learning. In: *Conference on Empirical Methods in Natural Language Processing (EMNLP)*. pp. 10658–10671 (2023)
61. Wang, X., Zhang, J., Qi, L., Shi, Y.: Generalizable decision boundaries: Dualistic meta-learning for open set domain generalization. In: *International Conference on Computer Vision (ICCV)*. pp. 11564–11573 (2023)
62. Wang, Y., Qi, L., Shi, Y., Gao, Y.: Feature-based style randomization for domain generalization. *IEEE Transactions on Circuits and Systems for Video Technology (TCSVT)* pp. 5495–5509 (2022)
63. Wang, Z., Luo, Y., Qiu, R., Huang, Z., Baktashmotlagh, M.: Learning to diversify for single domain generalization. In: *International Conference on Computer Vision (ICCV)*. pp. 834–843 (2021)
64. Wortsman, M., Ilharco, G., Kim, J.W., Li, M., Kornblith, S., Roelofs, R., Lopes, R.G., Hajishirzi, H., Farhadi, A., Namkoong, H., et al.: Robust fine-tuning of zero-shot models. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 7959–7971 (2022)

65. Xie, D., Xiong, J., Pu, S.: All you need is beyond a good init: Exploring better solution for training extremely deep convolutional neural networks with orthonormality and modulation. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 6176–6185 (2017)
66. Xu, Q., Zhang, R., Zhang, Y., Wang, Y., Tian, Q.: A fourier-based framework for domain generalization. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 14383–14392 (2021)
67. Zeng, G., Chen, Y., Cui, B., Yu, S.: Continual learning of context-dependent processing in neural networks. *Nature Machine Intelligence (MNI)* pp. 364–372 (2019)
68. Zhang, C., Zhang, M., Zhang, S., Jin, D., Zhou, Q., Cai, Z., Zhao, H., Liu, X., Liu, Z.: Delving deep into the generalization of vision transformers under distribution shifts. In: IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7277–7286 (2022)
69. Zhang, J., Qi, L., Shi, Y., Gao, Y.: Generalizable model-agnostic semantic segmentation via target-specific normalization. *Pattern Recognition (PR)* p. 108292 (2022)
70. Zhang, J., Qi, L., Shi, Y., Gao, Y.: MVDG: A unified multi-view framework for domain generalization. In: European Conference on Computer Vision (ECCV). pp. 161–177 (2022)
71. Zhang, J., Qi, L., Shi, Y., Gao, Y.: Domainadaptor: A novel approach to test-time adaptation. In: International Conference on Computer Vision (ICCV). pp. 18971–18981 (2023)
72. Zhou, K., Yang, Y., Hospedales, T., Xiang, T.: Deep domain-adversarial image generation for domain generalisation. In: AAAI Conference on Artificial Intelligence (AAAI). pp. 13025–13032 (2020)
73. Zhou, K., Yang, Y., Hospedales, T., Xiang, T.: Learning to generate novel domains for domain generalization. In: European Conference on Computer Vision (ECCV). pp. 561–578 (2020)
74. Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Domain adaptive ensemble learning. *IEEE Transactions on Image Processing (TIP)* pp. 8008–8018 (2021)
75. Zhou, K., Yang, Y., Qiao, Y., Xiang, T.: Domain generalization with mixstyle. In: International Conference on Learning Representations (ICLR) (2021)