Gradient Estimation for Unseen Domain Risk Minimization with Pre-Trained Models

Byounggyu Lew[1,2†]   Donghyun Son[1,2†]   Buru Chang[1,3*]

[1]Hyperconnect   [2]Seoul National University   [3]Sogang University

{korts,happydh1}@snu.ac.kr   buru@sogang.ac.kr

Abstract

Domain generalization aims to build generalized models that perform well on unseen domains when only source domains are available for model optimization. Recent studies have shown that large-scale pre-trained models can enhance domain generalization by leveraging their generalization power. However, these pre-trained models lack target task-specific knowledge yet due to discrepancies between the pre-training objectives and the target task. Although the task-specific knowledge could be learned from source domains by fine-tuning, this hurts the generalization power of pre-trained models due to gradient bias toward the source domains. To alleviate this problem, we propose a new domain generalization method that estimates unobservable gradients that reduce potential risks in unseen domains using a large-scale pre-trained model. These estimated unobservable gradients allow the pre-trained model to learn task-specific knowledge further while preserving its generalization ability by relieving the gradient bias. Our experimental results show that our method outperforms baseline methods on DOMAINBED, a standard benchmark in domain generalization. We also provide extensive analyses to demonstrate that the pre-trained model can learn task-specific knowledge without sacrificing its generalization power.

1. Introduction

Many machine learning studies assume that training and test data are independent and identically distributed (*i.i.d*). However, this *i.i.d* assumption does not always hold in real-world scenarios where distribution shifts between training and test data occur frequently. Thus, traditional machine learning models often show poor performance on unseen domains shifted from source (training) domains [36, 44].

---

†Equal contribution.
*Corresponding author.
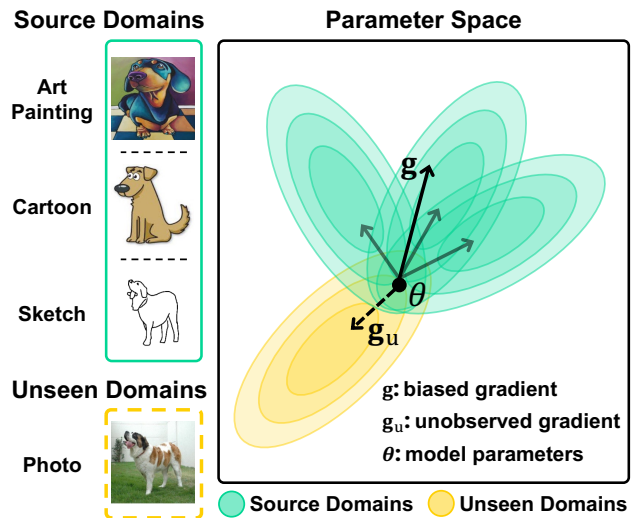This work was done while all authors were affiliated with Hyperconnect.

Figure 1: Model optimization is influenced by the gradient **g** biased toward the source domains, neglecting the unobservable gradient $\mathbf{g}_u$ that could minimize risks in the unseen domains. This can eventually result in a degradation of model performance on unseen domains.

To tackle this problem, *domain generalization* has attracted much attention recently.

The main goal of domain generalization is to build generalized models that also perform the target task (*e.g.,* classification) well on unseen domains (*e.g.,* realistic photo images) when only source domains (*e.g.,* cartoon or sketch images) are accessible during model optimization. Early domain generalization studies [31, 13, 24] have focused on learning domain-invariant representations across the source domains. However, [14] have recently shown that simple empirical risk minimization (ERM) [45] outperforms the previous methods on DOMAINBED, a benchmark for domain generalization, with pre-trained ResNet-50 [15]. Moreover, [54] provide empirical evidence that large-scale pre-trained models could enhance domain generalization by leveraging their generalization power.
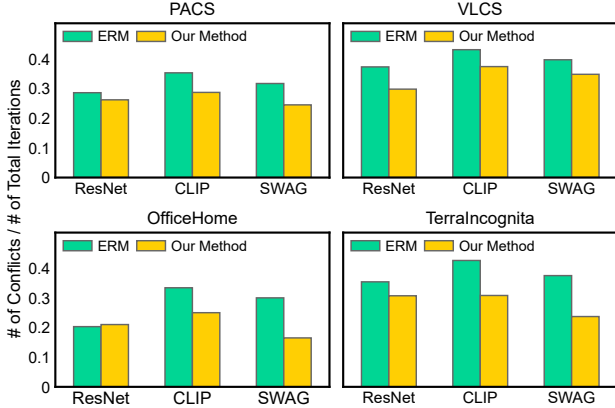
Figure 2: Gradient "*conflicts*" [53, 28] between $\mathbf{g}$ and $\mathbf{g}_u$ (*i.e.*, $\mathbf{g} \cdot \mathbf{g}_u < 0$) constantly occur throughout the whole fine-tuning iterations due to gradient bias. Our proposed method reduces the number of gradient conflicts by adding the estimated unobservable gradients $\tilde{\mathbf{g}}_u$ to the biased gradients $\mathbf{g}$. This observation indicates that gradient bias is relieved with the estimated gradients during model optimization. The more details are described in § 3.3.

Motivated by this, several studies have begun to leverage the generalization power of large-scale pre-trained models. [8] employ a pre-trained model for regularization, considering it as an approximation of the oracle model on any domain, and [26] utilize a frozen pre-trained model as a feature extractor. These studies have proven the usefulness of pre-trained models in domain generalization. However, the pre-trained models used in those studies cannot learn task-specific knowledge further since they are frozen during model optimization to preserve their generalization ability. To learn the task-specific knowledge, one can choose *fine-tuning* that updates all the parameters of pre-trained models by optimizing the models on the source domains. However, [21] demonstrate that fine-tuning distorts generalized representations of pre-trained models. Namely, fine-tuning hurts the generalization ability of pre-trained models.

In this paper, we interpret the above issue in terms of *gradient bias* during model optimization. As shown in Figure 1, the gradient of naive fine-tuning is biased toward the source domains because it is computed by only the source domains, disregarding unseen domains. Although this biased gradient reduces empirical risks in the source domains with the learning of task-specific knowledge, it probably increases risks in the unseen domains. We argue that such gradient bias would be relieved if unobservable gradients that lower the risks in the unseen domains are obtainable.

To this end, we propose a new domain generalization method named *GESTUR*, which estimates the unobservable gradients with a large-scale pre-trained model. GESTUR consists of two key components: a task expert (TE) and a generalization expert (GE), which are initialized with a large-scale pre-trained model. Based on ERM where gradients tend to be biased to the source domains, TE learns task-specific knowledge from source domains directly to transfer the knowledge to GE. Meanwhile, GE learns the task-specific knowledge from TE indirectly via exponential moving average (EMA) while preserving the generalization ability of a large-scale pre-trained model. Still, the gradient bias of TE might impair the generalization ability of GE. To mitigate this, GE is utilized to estimate unobservable gradients that minimize risks in unseen domains for TE based on the assumption that large-scale pre-trained models could act as a loose approximation of the oracle model of unseen domains (§ 2). As shown in Figure 2, the biased gradient of TE is relieved by simply adding the estimated unobservable gradients to the biased gradients, improving domain generalization performance (§ 3). Extensive experiments and analyses demonstrate that GESTUR outperforms baseline methods by learning the task-specific knowledge appropriately from source domains while preserving the generalization ability of large-scale pre-trained models.

Contributions: (1) We propose a simple yet effective domain generalization method that learns task-specific knowledge while preserving the generalization ability of large-scale pre-trained models. Based on the two experts, TE and GE, our proposed method estimates unobservable gradients that reduce potential risks in unseen domains to relieve gradient bias toward source domains. (2) We conduct extensive experiments to show the effectiveness of our proposed method in domain generalization. By providing careful analyses, we demonstrate that the unobservable gradients could be estimated with a large-scale pre-trained model, and it relieves the gradient bias. We also demonstrate that our proposed method learns task-specific knowledge without sacrificing the generalization ability of the large-scale pre-trained model.

2. Methodology

2.1. Preliminaries

Problem formulation. Let $\mathcal{D}_s$ and $\mathcal{D}_u$ be sets of source domains and unseen domains, respectively. Each domain $\mathcal{D}$ contains the total number of $n_{\mathcal{D}}$ data samples, $\{(x_i, y_i)\}_{i=1}^{n_{\mathcal{D}}} \sim \mathcal{D}$, where each data sample (x_i, y_i) consists of an input x_i and its target label y_i. The $n_{\mathcal{D}}$ data samples are *i.i.d* over some probability distribution. The main goal of domain generalization is to build a model θ that performs well on the unseen domains $\mathcal{D}_u$ when the source domains $\mathcal{D}_s$ are only available:

$$\min_{\theta} \mathbb{E}_{\mathcal{D} \sim \mathcal{D}_u} \mathbb{E}_{(x,y) \sim \mathcal{D}} [\ell((x, y); \theta)], \quad (1)$$

where $\ell((x,y);\theta)$ is the loss function defined for the model θ on the data sample (x,y). Note that this study focuses on solving classification tasks. Hence, we denote the model in detail as $\theta = \{\theta^f; \theta^c\}$ consisting of its feature extractor θ^f and classifier θ^c.

Motivation. With success in many downstream tasks, it has become a convention to initialize the feature extractor θ^f with a large-scale pre-trained model. Although pre-trained models provide better feature representations than randomly initialized parameters, they do not fully equip task-specific knowledge yet. It is because there is a discrepancy between the pre-training objective and the target task. For example, CLIP [37] is pre-trained to match web-crawled image-caption pairs, whereas the target task is to classify data into seven classes (*e.g.,* horse and dog), in the case of PACS [23]. Therefore, many studies have adopted fine-tuning that updates all the parameters of the feature extractor θ^f to learn the task-specific knowledge by optimizing the model on source domains $\mathcal{D}_s$. However, [21] observe that fine-tuning impairs generalization ability of pre-trained models during the learning of task-specific knowledge.

We try to interpret this issue at the gradient level. Based on ERM [45], the gradient $\mathbf{g}$ of fine-tuning is computed for the model θ on the source domains $\mathcal{D}_s$, as follows:

$$\mathbf{g} = \nabla_\theta \mathbb{E}_{(x,y)\sim\mathcal{B}}[\ell((x,y);\theta)], \tag{2}$$

where $\mathcal{B}$ is a mini-batch sampled from the source domains $\mathcal{D}_s$. The gradient $\mathbf{g}$ is influenced by only the source domains $\mathcal{D}_s$ because the unseen domains $\mathcal{D}_u$ are not accessible. Namely, the gradient is biased toward the source domains. We presume that this gradient bias degrades generalization performance in the unseen domains.

2.2. GESTUR: <u>G</u>rdient <u>Est</u>imation for <u>U</u>nseen Domain <u>R</u>isk Minimization with Pre-Trained Models

We hypothesize that the gradient bias mentioned above could be relieved if the unobservable gradient $\mathbf{g}_u$ minimizing risks in the unseen domains is obtainable. To achieve this, we borrow the assumption of [8] that large-scale pre-trained models are the approximation of the oracle model θ^* which is optimally generalized for any domain $\mathcal{D}$. Since the unobservable gradient $\mathbf{g}_u$ cannot be computed directly from the unseen domains $\mathcal{D}_u$, we consider the direction from the current model θ to the oracle model θ^* as the unobservable gradient $\mathbf{g}_u$. However, the oracle model is inaccessible in practice. Hence, we estimate the unobservable gradient using a large-scale pre-trained model as the approximation of the oracle model.

Note that we aim to estimate the unobservable gradient $\mathbf{g}_u$ for the unseen domains $\mathcal{D}_u$ to alleviate gradient bias,

---

Algorithm 1 GESTUR

1: **Input:** task expert θ_{TE}, generalization expert θ_{GE}, gradient scale factor λ, and moving average coefficient m.
2: **Init:** initialize the feature extractors θ_{TE}^f and θ_{GE}^f with a pre-trained model θ_0 and randomly initialize the classifiers θ_{TE}^c and θ_{GE}^c.
3: **Output:** the updated generalization expert θ_{GE}
4: **for** sampled mini-batch $\mathcal{B}$ from the source domains $\mathcal{D}_s$ **do**
5:      $\mathbf{g} = \nabla_\theta \mathbb{E}_{(x,y)\sim\mathcal{B}}[\ell((x,y);\theta_{\text{TE}})]$
6:      $\tilde{\mathbf{g}}_u^f = \theta_{\text{GE}}^f - \theta_{\text{TE}}^f$
7:      $\tilde{\mathbf{g}}_u^f = \lambda\|\mathbf{g}^f\|_2 \cdot \dfrac{\tilde{\mathbf{g}}_u^f}{\|\tilde{\mathbf{g}}_u^f\|_2}$
8:      $\mathbf{g}^f = (\mathbf{g}^f + \tilde{\mathbf{g}}_u^f)/2$
9:      update θ_{TE}^f with $\mathbf{g}^f$ and update θ_{TE}^c with $\mathbf{g}^c$
10:      update $\theta_{\text{GE}} = m\theta_{\text{GE}} + (1-m)\theta_{\text{TE}}$
11: **end for**

---

so the above assumption needs to be more elaborate due to the following reasons. First, we intend to design the unobservable gradient for the unseen domains $\mathcal{D}_u$ only rather than any domain $\mathcal{D}$. Second, pre-trained models do not have task-specific knowledge yet, as described in § 2.1. Therefore, we slightly modify the assumption as follows: pre-trained models are the *loose* approximation of the oracle model θ_u^* of *the unseen domains* $\mathcal{D}_u$, and they could get *closer* to the oracle model by learning task-specific knowledge. Based on this assumption, we propose a simple yet effective domain generalization method, GESTUR, which estimates the unobservable gradient $\mathbf{g}_u$ for unseen domain risk minimization with a large-scale pre-trained model.

Task expert and generalization expert. GESTUR consists of two classification models: a task expert (TE, θ_{TE}) and a generalization expert (GE, θ_{GE}), which are complementary to each other. Their feature extractors are both initialized with a large-scale pre-trained model θ_0, respectively. TE learns task-specific knowledge from the source domains $\mathcal{D}_s$ directly to transfer the knowledge to GE. Meanwhile, GE also learns task-specific knowledge from TE via EMA, but it aims to preserve the generalization ability of the pre-trained model deliberately. Here, the gradient bias of TE might hurt the generalization ability of GE. To relieve the gradient bias, GE is used to estimate the unobservable gradient $\mathbf{g}_u$ as the loose approximation of the oracle model θ_u^* for the unseen domains $\mathcal{D}_u$. By adding the estimated unobservable gradient to the gradient of TE, the gradient bias could be relieved. Our proposed GESTUR is summarized in Algorithm 1.

Gradient estimation. Using Equation 2, the gradient $\mathbf{g}$ for TE is computed as $\mathbf{g} = \nabla_\theta \mathbb{E}_{(x,y)\sim\mathcal{B}}[\ell((x,y);\theta_{\text{TE}})]$ while learning task-specific knowledge. The gradient $\mathbf{g}$ is biased toward the source domains $\mathcal{D}_s$. A gradient that minimizes

risks in the unseen domains could relieve the gradient bias, but it is unobservable. When we have access to the oracle model θ_u^* of the unseen domains $\mathcal{D}_u$, we can direct the current model to head to the oracle model instead of empirically calculating the unobservable gradient from the unseen domains. Hence, we treat the direction from the current model θ_{TE} to the oracle model θ_u^* as the unobservable gradient $\mathbf{g}_u$:

$$\mathbf{g}_u = \theta_u^* - \theta_{\text{TE}}. \qquad (3)$$

In fact, it is infeasible to access the oracle model. Thus, we estimate the unobservable gradient using GE that approximates the oracle model loosely, as follows:

$$\tilde{\mathbf{g}}_u = \theta_{\text{GE}} - \theta_{\text{TE}}. \qquad (4)$$

This estimated gradient $\tilde{\mathbf{g}}_u$ is used to relieve the gradient bias during the parameter optimization.

Parameter optimization. We want to emphasize again that GESTUR leverages the generalization power of large-scale pre-trained models to relieve gradient bias which distorts the generalized feature representations of the feature extractor θ^f. Hence, we limit the scope of usage of the estimated unobservable gradient $\tilde{\mathbf{g}}_u$ only to the feature extractor θ^f, not the classifier θ^c.

For TE, the estimated gradient $\tilde{\mathbf{g}}_u^f$ for the feature extractor θ_{TE}^f is added to the biased gradient $\mathbf{g}^f$ for the same feature extractor, as follows:

$$\mathbf{g}^f = \frac{1}{2}\Big(\mathbf{g}^f + \lambda\|\mathbf{g}^f\|_2 \cdot \frac{\tilde{\mathbf{g}}_u^f}{\|\tilde{\mathbf{g}}_u^f\|_2}\Big), \qquad (5)$$

where λ is a gradient scale factor that controls the influence of the normalized $\tilde{\mathbf{g}}_u^f$. The feature extractor θ_{TE}^f is updated with the gradient $\mathbf{g}^f$ adjusted by $\tilde{\mathbf{g}}_u^f$. On the other hand, the classifier θ_{TE}^c of TE is updated with its original gradient $\mathbf{g}^c$.

As our assumption, GE can get closer to the oracle model θ^* by learning task-specific knowledge. However, the generalization ability of GE decreases when we optimize GE on the source domains $\mathcal{D}_s$ directly to learn the task-specific knowledge. Therefore, we inject the learned task-specific knowledge of TE into GE delicately via EMA:

$$\theta_{\text{GE}} = m\theta_{\text{GE}} + (1-m)\theta_{\text{TE}}, \qquad (6)$$

where m is the moving average coefficient. By encouraging the parameters of GE to change slowly, EMA is helpful in preserving the generalization ability of GE. Since the goal of domain generalization is to build a model that minimizes the risk of the unseen domains, we choose GE, designed to approximate the oracle model of the domains, as our final model θ.

3. Experiments

3.1. Experimental setup

Datasets. We conduct experiments using five popular domain generalization benchmark datasets: PACS [23] (4 domains & 7 classes), VLCS [12] (4 domains & 5 classes), OfficeHome [46] (4 domains & 65 classes), TerraIncognita [4] (4 domains & 10 classes), and DomainNet [34] (6 domains & 345 classes).

Pre-trained models. We employ three pre-trained models of different sizes to verify that the proposed method performs well with various pre-trained models generally: ResNet-50 [15] pre-trained on ImageNet [9] (RN50), ViT-B/16 [10] with CLIP [37] (CLIP), and RegNetY-16GF [38] with SWAG [42] (SWAG).

Evaluation protocol. We adopt the experimental protocol of DOMAINBED, which enforces fair and realistic evaluations (*e.g.,* same model selection criterion) across competitors. We divide the data from each domain into 80% and 20% splits and follow *training-domain validation set* strategy for the model selection and the hyperparameter search in every experiment. We also repeat every experiment three times to reduce the randomness in dataset splits and parameter initialization, similar to [14], and report the mean and standard error of the experimental results.

Implementation details. Our implementation is built on the codebase of [8]. We use Adam optimizer [19] for parameter optimization. GESTUR has two hyperparameters, the gradient scale factor (λ) and the moving average coefficient (m). In every experiment, we search the optimal λ from $\{0.01, 0.05, 0.1, 0.5\}$ and fix m as 0.999. Other hyperparameters such as learning rate, weight decay, and dropout are searched in the same way as [8]. We explain more details of implementation in Appendix A.

Baselines. We exhaustively compare our proposed method with various baseline methods in the experiment. For simplicity, we report only the experimental results of baseline methods that show the higher performance than ERM [45], the simplest baseline method. We describe the baseline methods and report the full version of the results in Appendix B.1.

3.2. Main results

Results on RN50. The first part of Table 1 shows the experimental results where RN50 is used to initialize the feature extractor. GESTUR achieves the best performance for all the datasets except DomainNet. In detail, our method outperforms ERM by an average of 3.2%p. Furthermore,

Table 1: Evaluation results (%) on the five datasets with the three different pre-trained models. The best performance is in bold. GESTUR outperforms baseline methods in all experiments across all the types of the pre-trained models.

Method	PACS	VLCS	OfficeHome	TerraInc	DomainNet	Avg.
Using ResNet-50 pre-trained on ImageNet.						
ERM	84.2 ±0.1	77.3 ±0.1	67.6 ±0.2	47.8 ±0.6	44.0 ±0.1	64.2
SagNet	86.3 ±0.2	77.8 ±0.5	68.1 ±0.1	48.6 ±1.0	40.3 ±0.1	64.2
SelfReg	85.6 ±0.4	77.8 ±0.9	67.9 ±0.7	47.0 ±0.3	42.8 ±0.0	64.2
CORAL	86.2 ±0.3	78.8 ±0.6	68.7 ±0.3	47.6 ±1.0	41.5 ±0.1	64.5
mDSDI	86.2 ±0.2	79.0 ±0.3	69.2 ±0.4	48.1 ±1.4	42.8 ±0.1	65.1
GVRT	85.1 ±0.3	79.0 ±0.2	70.1 ±0.1	48.0 ±1.4	44.1 ±0.1	65.2
MIRO	85.4 ±0.4	79.0 ±0.0	70.5 ±0.4	50.4 ±1.1	44.3 ±0.2	65.9
SMA	87.5 ±0.2	78.2 ±0.2	70.6 ±0.1	50.3 ±0.5	46.0 ±0.1	66.5
SWAD	**88.1** ±0.1	79.1 ±0.1	70.6 ±0.2	50.0 ±0.3	**46.5** ±0.1	66.9
GESTUR	88.0 ±0.2	**80.1** ±0.2	**71.1** ±0.0	**51.3** ±0.2	46.3 ±0.1	**67.4**
Using ViT-B/16 with CLIP.						
ERM	83.4 ±0.5	75.9 ±1.3	66.4 ±0.5	35.3 ±0.8	44.4 ±0.6	61.1
SWAD	91.3 ±0.1	79.4 ±0.4	76.9 ±0.1	45.4 ±0.5	51.7 ±0.8	68.9
SMA	92.1 ±0.2	79.7 ±0.2	78.1 ±0.1	48.3 ±0.7	55.9 ±0.2	70.8
MIRO	95.6 ±0.8	82.2 ±0.3	82.5 ±0.1	54.3 ±0.4	54.0 ±0.3	73.7
GESTUR	**96.0** ±0.0	**82.8** ±0.1	**84.2** ±0.1	**55.7** ±0.2	**58.9** ±0.1	**75.5**
Using RegNetY-16GF with SWAG.						
ERM	89.6 ±0.4	78.6 ±0.3	71.9 ±0.6	51.4 ±1.8	48.5 ±0.6	68.0
SWAD	94.7 ±0.2	79.7 ±0.2	80.0 ±0.1	57.9 ±0.7	53.6 ±0.6	73.2
MIRO	**97.4** ±0.2	79.9 ±0.6	80.4 ±0.2	58.9 ±1.3	53.8 ±0.1	74.1
SMA	95.5 ±0.0	80.7 ±0.1	82.0 ±0.0	59.7 ±0.0	60.0 ±0.0	75.6
GESTUR	96.9 ±0.1	**83.5** ±0.1	**83.1** ±0.0	**61.1** ±0.4	**60.1** ±0.0	**76.9**

our method improves the runner-up by: 1.0%*p* in VLCS, 0.5%*p* in OfficeHome, and 1.3%*p* in TerraIncognita. Especially, the proposed method outperforms the state-of-the-art method (SWAD [7]) by an average of 0.5%*p*.

Results on `CLIP` and `SWAG`. In the second and third parts of Table 1, we show the experimental results where the larger pre-trained models, `CLIP` and `SWAG`, are used to initialize the feature extractor, respectively. In summary, GESTUR achieves the best performance in all the datasets. In detail, our method outperforms MIRO that also considers large-scale pre-trained models as the approximation of the oracle model by 1.8%*p* and 2.8%*p* on `CLIP` and `SWAG`, respectively. From this, we verify the validity of our assumption that large-scale pre-trained models are the *loose* approximation of the oracle model and they can get closer to the oracle model by learning task-specific knowledge further. Our method successfully leverages the generalization ability of pre-trained models compared to other baseline methods with our assumption. Interestingly, we observe that the performance gap between the proposed method and ERM increases as the size of the pre-trained model increases.

3.3. Comparison with ERM in terms of gradient bias

Setup. As described in § 1, we suspect that the gradient bias degrades the domain generalization performance. We further conduct analysis to check how much gradient bias occurs during the fine-tuning and how much gradient bias is alleviated by our proposed method. To quantify the gradient bias, we borrow the concept of *gradient conflict* [53, 28]: there is a conflict between two gradients g_i and g_j if $g_i \cdot g_j < 0$. For every iteration, we first sample two mini-batches from both source domains and an unseen domain, respectively. We then compute losses of the mini-batches, and calculate gradients g and g_u from the losses, respectively. Finally, we count the number of iterations where the gradient conflict ($g \cdot g_u < 0$) occurs, for ERM and GESTUR. Here, we update the model using only the gradient g since unseen domains are inaccessible in practice.

Results. As shown in Table 2, GESTUR reduces the gradient conflicts of ERM by around 11.5%, 21%, and 28.2% for the pre-trained models, respectively. From this, we verify that our proposed method relieves gradient bias by esti-

Table 2: The percentage (%) of gradient conflicts between $\mathbf{g}$ and $\mathbf{g}_u$ ($\mathbf{g} \cdot \mathbf{g}_u < 0$) to the whole training iterations. GESTUR reduces the total number of the gradients conflicts with the estimated unobservable gradients $\tilde{\mathbf{g}}_u$.

Method	PACS	VLCS	OH	TI	Avg.
Using ResNet-50 pre-trained on ImageNet.					
ERM	28.6	37.3	**20.3**	35.4	30.4
GESTUR	**26.2**	**29.8**	21.0	**30.7**	**26.9**
Using ViT-B/16 with CLIP.					
ERM	35.3	43.1	33.4	42.6	38.6
GESTUR	**28.7**	**37.4**	**25.0**	**30.8**	**30.5**
Using RegNetY-16GF with SWAG.					
ERM	31.7	39.7	30.0	37.5	34.7
GESTUR	**24.5**	**34.8**	**16.5**	**23.7**	**24.9**

Table 3: Evaluation results (%) on PACS with the three different pre-trained models varying λ. The larger scale factor λ improves the generalization performance when larger pre-trained models CLIP and SWAG are given because these larger models better approximate the oracle model than the small one, RN50.

Pre-trained model	Dataset (size)	λ			
		0.01	0.05	0.1	0.5
RN50	ImageNet (1.3M)	**88.0**	86.0	82.1	73.4
CLIP	CLIP (400M)	94.8	96.0	**96.2**	96.0
SWAG	Instagram (3.6B)	96.3	96.9	97.6	**97.9**

mating unobservable gradients with the pre-trained model. We observe that gradient conflicts occur more often in GESTUR than ERM on only the experimental setup (Office-Home w/ RN50), which is consistent with the performance in Table 5 where ERM outperforms GESTUR w/ TE. This observation indicates that the domain generalization performance is affected by gradient bias represented as the gradient conflicts in this analysis. Additional analysis on the similarity of the true and estimated unobservable gradients is provided in Appendix C.3.

3.4. Relationship between λ and the size of the pre-trained model

Setup. Our proposed GESTUR controls the scale of the estimated unobservable gradients that reduce risks in unseen domains using the gradient scale factor λ. To verify the effect of the scale factor, we observe the performance change varying the scale factor.

Results. In Table 3, RN50 achieves the best performance with $\lambda = 0.01$. On the other hand, the larger pre-trained models, CLIP and SWAG achieve the best performance with the relatively larger $\lambda = 0.1$ and $\lambda = 0.5$, respectively. We summarize more results on other datasets (*i.e.,* VLCS, OfficeHome, and TerraIncognita) in Appendix B.2, and they show the similar pattern as in PACS.

Intuitively, the larger pre-trained models act as a better approximation of the oracle model than the small one because they are likely to encounter various domains from the huge web-crawled datasets during pre-training. They help to estimate unobservable gradients more accurately. The larger gradient scale factor, gradients $\mathbf{g}$ of TE is more affected by the estimated unobservable gradients $\tilde{\mathbf{g}}_u$ while optimizing the model on source domains. From this, we

can conclude that the larger scale factor improves the generalization performance when larger pre-trained models are given.

3.5. Task-specific knowledge learned by the generalization expert

Setup. We conduct additional experiments to show that the feature extractor θ_{GE}^f of GE learns task-specific knowledge successfully. Linear probing that updates parameters of only the classifier while freezing those of the feature extractor is common practice for assessing representation quality. We assume that the more task-specific knowledge the feature extractor learns, the better linear probing performance it exhibits in unseen domains targeting the same task. In detail, we first train GE on source domains and then evaluate linear probing performance on an unseen domain with the trained feature extractor of GE. We compare it with the case that a *frozen* pre-trained model is used as the feature extractor. For linear probing, we simply train a logistic regression classifier on the output feature representations of each feature extractor using the unseen domain only. Note that, in this analysis, we use CLIP and SWAG which are pre-trained with objectives significantly different from the target task to demonstrate the effectiveness of the newly learned task-specific knowledge clearly.

Results. As shown in Table 4, GE outperforms *frozen* in all benchmark datasets except the one case where the two models reach the near 99% performance. This shows that GE is learning task-specific knowledge further during training, which makes it a better approximation of the oracle model. The result supports our claim that pre-trained models are not fully equipped with target task-specific knowledge, and injecting the knowledge to the models further increases performance.

Table 4: Linear probing performance (%) with the two different feature extractors: *frozen* pre-trained model θ_0 and the feature extractor θ_{GE}^f of GE. These results demonstrate that GE successfully learns the task-specific knowledge while preserving the generalization power of pre-trained models.

Model	PACS	VLCS	OfficeHome	TerraInc	Avg.
Using ViT-B/16 with CLIP.					
frozen	98.5 ±0.1	88.5 ±0.2	89.3 ±0.1	83.4 ±0.2	89.9
GE	**98.7** ±0.1	**90.0** ±0.6	**89.4** ±0.3	**88.3** ±0.1	**91.6**
Using RegNetY-16GF with SWAG.					
frozen	**98.9** ±0.1	87.1 ±0.2	89.6 ±0.0	89.3 ±0.0	91.2
GE	98.7 ±0.1	**88.8** ±0.2	**90.1** ±0.2	**90.2** ±0.1	**92.0**

3.6. Comparison between the task expert and the generalization expert

Setup. GESTUR consists of two essential components: the task expert (TE) and the generalization expert (GE). In this paper, we use GE as the final model based on the assumption that GE is set as the approximation of the oracle model of unseen domains. Nevertheless, TE is also designed to preserve the generalization ability of pre-trained models since it also uses the estimated unobservable gradient in every update to relieve its gradient bias. Therefore, we compare the performance of ERM, GESTUR w/ GE, and its variant GESTUR w/ TE based on the hyperparameters searched in § 3.2 to show that they preserve the generalization ability of large-scale pre-trained models.

Results. As shown in Table 5, GESTUR w/ GE achieves the best performance in all experiments. Also, GESTUR w/ TE outperforms ERM by averages of 9.8%p and 4.5%p when using CLIP and SWAG, respectively. The performance of GESTUR w/ TE is higher when the larger pre-trained models are given, similar to the observation in § 3.2. These observations demonstrate that GESTUR w/ TE could preserve the generalization ability of the pre-trained models with the estimated unobservable gradient, *i.e.,* the gradient bias of TE is relieved. Moreover, GESTUR w/ GE shows a higher performance than GESTUR w/ TE, which indicates that EMA ensures the model preserves generalization ability during the learning of task-specific knowledge stable. From these, we reaffirm the justification for our choice of GE as the final model.

3.7. Performance on source domains

Setup. Domain generalization aims to improve the generalization performance on unseen domains shifted from source domains. Thus, domain generalization studies often do not consider situations where the source domains

Table 5: Evaluation results (%) across the four datasets using the three different pre-trained models. We separate the cases where GESTUR uses TE and GE as the final model, respectively. GESTUR performs best when GE is used as the final model. From this, we show that EMA helps the model to preserve generalization ability during the learning of task-specific knowledge.

Method	PACS	VLCS	OfficeHome	TerraInc	Avg.
Using ResNet-50 pre-trained on ImageNet.					
ERM	84.2	77.3	<u>67.6</u>	<u>47.8</u>	<u>69.2</u>
GESTUR w/ TE	<u>84.9</u>	<u>79.2</u>	66.3	45.6	69.0
GESTUR w/ GE	**88.0**	**80.1**	**71.1**	**51.3**	**72.6**
Using ViT-B/16 with CLIP.					
ERM	83.4	75.9	66.4	35.3	65.3
GESTUR w/ TE	<u>90.7</u>	<u>82.4</u>	<u>76.9</u>	<u>50.4</u>	<u>75.1</u>
GESTUR w/ GE	**96.0**	**82.8**	**84.2**	**55.7**	**79.7**
Using RegNetY-16GF with SWAG.					
ERM	89.6	78.6	71.9	51.4	72.9
GESTUR w/ TE	<u>94.8</u>	<u>82.5</u>	<u>77.7</u>	<u>54.7</u>	<u>77.4</u>
GESTUR w/ GE	**96.9**	**83.5**	**83.1**	**61.1**	**81.2**

Table 6: Evaluation results (%) on the four datasets with `RN50`. Notably, we compute the average performance across the source domains rather than in the unseen target domain. Our experiments show that GESTUR w/ TE is worse than ERM but GESTUR w/ GE is better. These results demonstrate that our two-expert architecture is robust even when the testing domains are exactly the same as the training domains.

Method	PACS	VLCS	OfficeHome	TerraInc	Avg.
ERM	97.4 ±0.2	86.7 ±0.1	82.9 ±0.3	**92.2** ±0.1	89.8
GESTUR w/ TE	97.1 ±0.1	86.9 ±0.2	81.7 ±0.2	89.4 ±0.1	88.8
GESTUR w/ GE	**98.2** ±0.1	**87.4** ±0.2	**84.8** ±0.3	91.4 ±0.1	**90.5**

and the target domains are similar. To verify whether estimated unobservable gradients are useful when the unseen domains are similar to the source domains, we report the performance on the training-domain validation set, simulating the situations when the testing domains are exactly the same as the training domains.

Results. The evaluation results are summarized in Table 6. GESTUR w/ TE shows worse performance than ERM, indicating that the estimated unobservable gradients act as noisy gradients. Namely, gradients biased toward the source domains are more helpful than estimated unobservable gradients when the source domains and the target domains are similar. Nonetheless, GESTUR w/ GE performs better than ERM, demonstrating that our two-expert architecture is robust to various situations even when source domains and unseen domains are similar or not.

4. Related Work

4.1. Domain generalization

Domain alignment. Domain alignment is to learn domain-invariant feature representations by removing domain-specific knowledge in the representations. Adversarial training is widely adopted to learn domain invariant features through a min-max game between a feature extractor and a domain discriminator [13, 25, 29, 59]. On the other hand, several studies [31, 43, 24] aim to minimize feature divergence across source domains. Recently, contrastive learning-based algorithms [18, 52] have been proposed to minimize distances between feature representations of samples in the same class, regardless their domains.

Data augmentation. Many studies have employed data augmentation techniques to improve domain generalization performance. For example, [14] apply simple data augmentation techniques as a default setup in DOMAINBED and some studies [48, 50, 51] utilize Mixup [55]. Recently, a few works [58, 32, 17] focus on image style, based on the idea that domain gap is closely related to image style. On the other side, some works on single domain generalization introduce adversarial data augmentation [47, 11, 35] to generate hard samples adversarially while assuring their reliability.

Gradient-based. Recently, several studies utilize gradients to build generalized models, especially by aligning gradients from different domains. [28] exploit gradient agreement for gradient surgery, based on the hypothesis that conflicting gradients contain domain-specific information. [41] propose a training method that maximizes inner product between source domain gradients to match optimization paths across domains. Similarly, [39] try to match domain-level Hessian to align loss landscapes across domains. As another line of work, [16] introduce the self-challenging algorithm that iteratively masks dominant features, which are selected by the scale of the gradients.

Meta-learning-based. Since simulating domain shift by dividing source domains into meta-train and meta-test domains was first introduced in MLDG [22], several approaches have been proposed in a similar setting. For example, [3] propose to learn a regularizer for classifier weights and [56] bring the idea of Reptile [33] to MLDG to further increase performance with a multi-view framework. On the other hand, [57] employ meta-learning to adaptively predict model parameters from a batch of inputs.

Others. Some of the works bring concepts of causality [27], optimize the worst-case performance [40, 20], uti-lize text labels [30], or average model weights from different epochs [7, 2].

Our work differs from aforementioned approaches in that we mainly concentrate on effectively utilizing large-scale pre-trained models.

4.2. Domain generalization with pre-trained models

Recently, [14] empirically show that simple ERM [45] outperforms most of early methods with pre-trained ResNet-50 [15]. [54] show that using large-scale models pre-trained on massive datasets improves out-of-distribution performance. [21] find that fine-tuning distorts pre-trained features and propose the *linear-probing then fine-tuning* to mitigate the feature distortion. [49] find that linearly interpolating the zero-shot and fine-tuned parameters of a pre-trained model improves performance in both source and unseen domains. Although GESTUR's EMA (Equation 6) looks similar to their interpolation, GESTUR updates the pre-trained model to inject task-specific knowledge. [26] propose a method to efficiently leverage a pool of large-scale pre-trained models through specialty-aware ensemble learning. [8] propose MIRO, a regularization method that targets to minimize mutual information with pre-trained models which approximate the oracle model. In this work, we share similar motivation with MIRO in that we initially approximate the oracle model with a large-scale pre-trained model. However, we iteratively inject task-specific knowledge into the approximation of the oracle model, resulting in a better approximation.

5. Conclusion and Future Work

In this paper, we propose a new domain generalization method that learns task-specific knowledge while preserving the generalization ability of large-scale pre-trained models. We point out that gradient bias toward source domains hurts the generalization ability of pre-trained models during fine-tuning. To alleviate the gradient bias, our proposed method estimates unobservable gradients that minimize risk in unseen domains based on two key components: a task expert and a generalization expert. Experimental results on DOMAINBED show that our proposed method outperforms baseline methods in domain generalization. Through extensive analyses, we also demonstrate that the estimated unobservable gradients effectively reduce gradient bias, thereby helping to learn task-specific knowledge without hurting the generalization power of large-scale pre-trained models.

Although we verify the effectiveness of our proposed method, it heavily relies on the capability of pre-trained models. When unseen domains that pre-trained models did not encounter are given (*e.g.* ResNet trained on ImageNet does not see medical images), the pre-trained models might not act as an approximation of the oracle model of the domains. We will address this issue in future work.

References

[1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. 12

[2] Devansh Arpit, Huan Wang, Yingbo Zhou, and Caiming Xiong. Ensemble of averages: Improving model selection and boosting performance in domain generalization. In Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho, editors, *Advances in Neural Information Processing Systems*, 2022. 8, 12

[3] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. *Advances in neural information processing systems*, 31, 2018. 8

[4] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European conference on computer vision (ECCV)*, pages 456–473, 2018. 4, 12

[5] Gilles Blanchard, Aniket Anand Deshmukh, Ürun Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *The Journal of Machine Learning Research*, 22(1):46–100, 2021. 12

[6] Manh-Ha Bui, Toan Tran, Anh Tran, and Dinh Phung. Exploiting domain-specific features to enhance domain generalization. *Advances in Neural Information Processing Systems*, 34:21189–21201, 2021. 12

[7] Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. Swad: Domain generalization by seeking flat minima. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. 5, 8, 12, 13, 14

[8] Junbum Cha, Kyungjae Lee, Sungrae Park, and Sanghyuk Chun. Domain generalization by mutual-information regularization with pre-trained models. *arXiv preprint arXiv:2203.10789*, 2022. 2, 3, 4, 8, 12, 13, 14

[9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. 4, 12

[10] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 4, 12

[11] Xinjie Fan, Qifei Wang, Junjie Ke, Feng Yang, Boqing Gong, and Mingyuan Zhou. Adversarially adaptive normalization for single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8208–8217, 2021. 8

[12] Chen Fang, Ye Xu, and Daniel N Rockmore. 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*, pages 1657–1664, 2013. 4, 12

[13] Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016. 1, 8, 12

[14] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. In *International Conference on Learning Representations*, 2021. 1, 4, 8, 13

[15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 4, 8, 12

[16] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *European Conference on Computer Vision*, pages 124–140. Springer, 2020. 8, 12

[17] Juwon Kang, Sohyun Lee, Namyup Kim, and Suha Kwak. Style neophile: Constantly seeking novel styles for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7130–7140, 2022. 8

[18] Daehee Kim, Youngjun Yoo, Seunghyun Park, Jinkyu Kim, and Jaekoo Lee. Selfreg: Self-supervised contrastive regularization for domain generalization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9619–9628, 2021. 8, 12

[19] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR (Poster)*, 2015. 4

[20] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghuai Zhang, Remi Le Priol, and Aaron Courville. Out-of-distribution generalization via risk extrapolation (rex). In *International Conference on Machine Learning*, pages 5815–5826. PMLR, 2021. 8, 12

[21] Ananya Kumar, Aditi Raghunathan, Robbie Matthew Jones, Tengyu Ma, and Percy Liang. Fine-tuning can distort pretrained features and underperform out-of-distribution. In *International Conference on Learning Representations*, 2022. 2, 3, 8

[22] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 2018. 8, 12

[23] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. 3, 4, 12

[24] Haoliang Li, Sinno Jialin Pan, Shiqi Wang, and Alex C Kot. Domain generalization with adversarial feature learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5400–5409, 2018. 1, 8, 12

[25] Ya Li, Xinmei Tian, Mingming Gong, Yajing Liu, Tongliang Liu, Kun Zhang, and Dacheng Tao. Deep domain generalization via conditional invariant adversarial networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 624–639, 2018. 8, 12

[26] Ziyue Li, Kan Ren, Xinyang Jiang, Bo Li, Haipeng Zhang, and Dongsheng Li. Domain generalization using pretrained models without fine-tuning. *arXiv preprint arXiv:2203.04600*, 2022. 2, 8

[27] Fangrui Lv, Jian Liang, Shuang Li, Bin Zang, Chi Harold Liu, Ziteng Wang, and Di Liu. Causality inspired representation learning for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8046–8056, 2022. 8

[28] Lucas Mansilla, Rodrigo Echeveste, Diego H Milone, and Enzo Ferrante. Domain generalization via gradient surgery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6630–6638, 2021. 2, 5, 8

[29] Toshihiko Matsuura and Tatsuya Harada. Domain generalization using a mixture of multiple latent domains. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 11749–11756, 2020. 8

[30] Seonwoo Min, Nokyung Park, Siwon Kim, Seunghyun Park, and Jinkyu Kim. Grounding visual representations with texts for domain generalization. *arXiv preprint arXiv:2207.10285*, 2022. 8, 12

[31] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18. PMLR, 2013. 1, 8

[32] Hyeonseob Nam, HyunJae Lee, Jongchan Park, Wonjun Yoon, and Donggeun Yoo. Reducing domain gap by reducing style bias. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8690–8699, 2021. 8, 12

[33] Alex Nichol, Joshua Achiam, and John Schulman. On first-order meta-learning algorithms. *arXiv preprint arXiv:1803.02999*, 2018. 8

[34] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1406–1415, 2019. 4, 12

[35] Fengchun Qiao, Long Zhao, and Xi Peng. Learning to learn single domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12556–12565, 2020. 8

[36] Joaquin Quinonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. Mit Press, 2008. 1

[37] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, pages 8748–8763. PMLR, 2021. 3, 4, 12

[38] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10428–10436, 2020. 4, 12

[39] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. In *International Conference on Machine Learning*, pages 18347–18377. PMLR, 2022. 8, 12

[40] Shiori Sagawa, Pang Wei Koh, Tatsunori B Hashimoto, and Percy Liang. Distributionally robust neural networks for group shifts: On the importance of regularization for worst-case generalization. *arXiv preprint arXiv:1911.08731*, 2019. 8, 12

[41] Yuge Shi, Jeffrey Seely, Philip Torr, Siddharth N, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. In *International Conference on Learning Representations*, 2022. 8, 12

[42] Mannat Singh, Laura Gustafson, Aaron Adcock, Vinicius de Freitas Reis, Bugra Gedik, Raj Prateek Kosaraju, Dhruv Mahajan, Ross Girshick, Piotr Dollár, and Laurens van der Maaten. Revisiting weakly supervised pre-training of visual perception models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 804–814, June 2022. 4, 12

[43] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pages 443–450. Springer, 2016. 8, 12

[44] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1521–1528. IEEE, 2011. 1

[45] Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 1999. 1, 3, 4, 8, 12

[46] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5018–5027, 2017. 4, 12

[47] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. *Advances in neural information processing systems*, 31, 2018. 8

[48] Yufei Wang, Haoliang Li, and Alex C Kot. Heterogeneous domain generalization via domain mixup. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 3622–3626. IEEE, 2020. 8, 12

[49] Mitchell Wortsman, Gabriel Ilharco, Jong Wook Kim, Mike Li, Simon Kornblith, Rebecca Roelofs, Raphael Gontijo Lopes, Hannaneh Hajishirzi, Ali Farhadi, Hongseok Namkoong, and Ludwig Schmidt. Robust fine-tuning of zero-shot models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7959–7971, June 2022. 8, 12, 14

[50] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6502–6509, 2020. 8, 12

[51] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677*, 2020. 8, 12

[52] Xufeng Yao, Yang Bai, Xinyun Zhang, Yuechen Zhang, Qi Sun, Ran Chen, Ruiyu Li, and Bei Yu. Pcl: Proxy-based contrastive learning for domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7097–7107, 2022. 8

[53] Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020. 2, 5

[54] Yaodong Yu, Heinrich Jiang, Dara Bahri, Hossein Mobahi, Seungyeon Kim, Ankit Singh Rawat, Andreas Veit, and Yi Ma. An empirical study of pre-trained vision models on out-of-distribution generalization. In *NeurIPS 2021 Workshop on Distribution Shifts: Connecting Methods and Applications*, 2021. 1, 8

[55] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 8

[56] Jian Zhang, Lei Qi, Yinghuan Shi, and Yang Gao. Mvdg: A unified multi-view framework for domain generalization, 2021. 8

[57] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: Learning to adapt to domain shift. *Advances in Neural Information Processing Systems*, 34:23664–23678, 2021. 8, 12

[58] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Domain generalization with mixstyle. In *International Conference on Learning Representations*, 2021. 8, 12

[59] Wei Zhu, Le Lu, Jing Xiao, Mei Han, Jiebo Luo, and Adam P Harrison. Localized adversarial domain generalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7108–7118, 2022. 8