

# Generalization on Unseen Domains via Inference-time Label-Preserving Target Projections

Prashant Pandey<sup>1</sup>, Mrigank Raman<sup>\*1</sup>, Sumanth Varambally<sup>\*1</sup>, Prathosh AP<sup>1</sup>  
<sup>1</sup>IIT Delhi

{bsz178495, mt1170736, mt6170855, prathoshap}@iitd.ac.in

## Abstract

*Generalization of machine learning models trained on a set of source domains on unseen target domains with different statistics, is a challenging problem. While many approaches have been proposed to solve this problem, they only utilize source data during training but do not take advantage of the fact that a single target example is available at the time of inference. Motivated by this, we propose a method that effectively uses the target sample during inference beyond mere classification. Our method has three components - (i) A label-preserving feature or metric transformation on source data such that the source samples are clustered in accordance with their class irrespective of their domain (ii) A generative model trained on these features (iii) A label-preserving projection of the target point on the source-feature manifold during inference via solving an optimization problem on the input space of the generative model using the learned metric. Finally, the projected target is used in the classifier. Since the projected target feature comes from the source manifold and has the same label as the real target by design, the classifier is expected to perform better on it than the true target. We demonstrate that our method outperforms the state-of-the-art Domain Generalization methods on multiple datasets and tasks.*

## 1. Introduction

Domain shift refers to the existence of significant divergence between the distributions of the training and the test data [41]. This causes the machine learning models trained only on the training or the source data to perform poorly on the test or target data. A naive way of handling this problem is to fine-tune the model with new data which is often infeasible because of the difficulty in acquiring labelled data for every new target domain. The class of Domain Adaptation (DA) methods [42, 9, 14, 39, 26, 4, 31, 33, 34] tackle this problem by utilizing the (unlabeled) target data to minimize the domain shift; however they cannot be used when unlabeled

target data is unavailable.

Domain generalization (DG) [30, 21, 10, 22, 1, 23], on the other hand, views the problem from the following perspective: how to make a model trained on single or multiple source domains generalize on completely unseen target domains. These methods do so via (i) learning feature representations that are invariant to the data domains using methods such as adversarial learning [24, 25], (ii) simulating the domain shift while learning through meta-learning approaches [22, 1], and (iii) augmenting the source dataset with synthesized data from fictitious target domains [44, 48]. These methods have been shown to be effective in dealing with the problem of domain shift. However, most of the existing methods do not utilize the test sample from the target distribution available at the time of inference beyond mere classification. On the other hand, it is a common experience that when humans encounter an unseen object, they often relate it to a previously perceived similar object.

Motivated by this intuition, in this paper, we make the following contributions towards addressing the problem of DG: (a) Given samples from multiple source distributions, we propose to learn a source domain invariant representation that also preserves the class labels. (b) We propose to ‘project’ the target samples to the manifold of the source-data features before classification through an inference-time label-preserving optimization procedure over the input of a generative model (learned during training) that maps an arbitrary distribution (Normal) to the source-feature manifold. (c) We demonstrate through extensive experimentation that our method achieves new state-of-the-art performance on standard DG tasks while also outperforming other methods in terms of robustness and data efficiency.

## 2. Prior Work

**Meta-learning** : Meta-learning methods aim to improve model robustness against unseen domains by simulating domain shift during training. This is done by splitting the training set into a meta-train and meta-test set. [22] provide a general framework for meta-learning-based DG, where

\*Equal contribution

model parameters are updated to minimize loss over the meta-train and meta-test domains in a coordinated manner. [1] propose a pre-trained regularizer network which is used to regularize the learning objective of a domain-independent task network. [7] use a common feature extractor backbone network in conjunction with several domain-specific aggregation modules. An aggregation over these modules is performed during inference to predict the class label. [23] train separate feature extractors and classifiers on each of the source domains and minimize the loss on mismatched pairs of feature extractors and classifiers to improve model robustness. [15] utilise a probabilistic meta-learning model in which classifier parameters shared across domains are modeled as distributions. They also learn domain-invariant representations by optimizing a variational approximation to the information bottleneck. Since meta-learning methods are only trained on the simulated domain shifts, they might not always perform well on target domains that are not ‘covered’ in the simulated shifts.

**Data augmentation:** Augmenting the dataset with random transformations improves generalization [13]. Commonly used augmentation techniques include rotation, flipping, random cropping, random colour distortions, amongst others. [38] use gradients from a domain classifier to perturb images. However, these perturbations might not be reflective of practically observed domain shift. [48] aim to address this issue using an adversarial procedure to train a transformation network to produce an image translation that aims to generate novel domains while retaining class information. [47] generate images from pseudo-novel domains with an optimal transport based formulation while preserving semantic information with cycle-consistency and classification losses. [36] solve the problem of single-source DG by creating fictitious domains using Wasserstein Auto-Encoders in a meta-learning framework. While these generated domains differ significantly from the source domains, they potentially do not reflect practical domain differences.

**Domain-invariant representations:** Another common pervasive theme in domain generalization literature is transforming the source data into a lower-dimensional ‘feature’ space that is invariant to domains but retains the discriminative class information; these features are used for classification. [10] learn an auto-encoder to extract domain invariant features by reconstructing inter and cross domain images. [24] use adversarial auto-encoders to align the representations from all the source domains to a Laplacian prior using adversarial learning procedure. [6] employ episodic training to simulate domain shift while minimizing a global class-alignment loss and local sample-clustering objective to cluster points class-wise. [15] learn a kernel function that minimizes mean domain discrepancy and intra-class scatter while maximizing mean class discrepancy and multi-domain between-class scatter. [35] propose a low-rank de-

composition on the final classification layer to identifiably learn common and specific features across domains. [37] use domain-specific normalizations to learn representations that are domain-agnostic and semantically discriminative.

All of the above methods require domain labels, which might not always be viable. [5] aim to solve the problem of DG without domain labels by learning an auxiliary task of solving jigsaw puzzles. The idea is that features learned from such an auxiliary task will be invariant of the domains. [27] first assigns pseudo-labels inferred by clustering the domain discriminative features. They train a domain classifier against these pseudo-labels, which is further used to adversarially train a domain-invariant feature extractor. [29] use a semantic alignment loss as an additional regularizer while training the classifier for domain invariant feature learning. [16] iteratively locate dominant features activated on the training data using layer gradients, and learn useful features by self-challenging. [45] learn how to generalize across domains by simultaneously providing extrinsic supervision in the form of a metric learning task and intrinsic supervision in terms of a self-supervised auxiliary task. The most similar method to our own is [40] in that they also run an inference-time procedure. However, unlike our method, they use the test sample for updating model parameters.

### 3. Proposed Method

#### 3.1. Problem Setting and Method Overview

Let  $\mathcal{X}$  and  $\mathcal{Y}$  respectively denote data and the label spaces. Let  $\mathcal{H}$  be the space of hypotheses where each hypothesis  $h$  in  $\mathcal{H}$  maps points from  $\mathcal{X}$  to a label in  $\mathcal{Y}$ . A domain is defined by the tuple  $(\mathcal{D}, g_{\mathcal{D}})$  where  $\mathcal{D}$  is a probability distribution over  $\mathcal{X}$  and  $g_{\mathcal{D}}$  where  $g_{\mathcal{D}} : \mathcal{X} \rightarrow \mathcal{Y}$  is a function that assigns ground-truth labels. It is generally assumed that the ground-truth labeling function  $g$  is same across all the domains. Domain Generalization is defined as the task where there are a total of  $N$  domains out of which  $|S|$  are source and  $|T|$  are target domains. The source and target domains are respectively denoted by  $\mathcal{D}_i^S, i \in [|S|]$  and  $\mathcal{D}_j^T, j \in [|T|]$ . The objective is to train a classifier on the source domains that predicts well on the target domain when the target samples are not available during training.

The motivation for our method comes from the following observation: DG methods that learn domain invariant representations do so only using the source data. Therefore, classifiers trained on such representations are not guaranteed to perform well on target data that is outside the source data manifolds. Hence, performance on the target data can be improved if the target sample is projected on to the manifold of the source features such that the ground-truth label is preserved, before classification. To this end, we propose a three-part procedure for domain generalization:

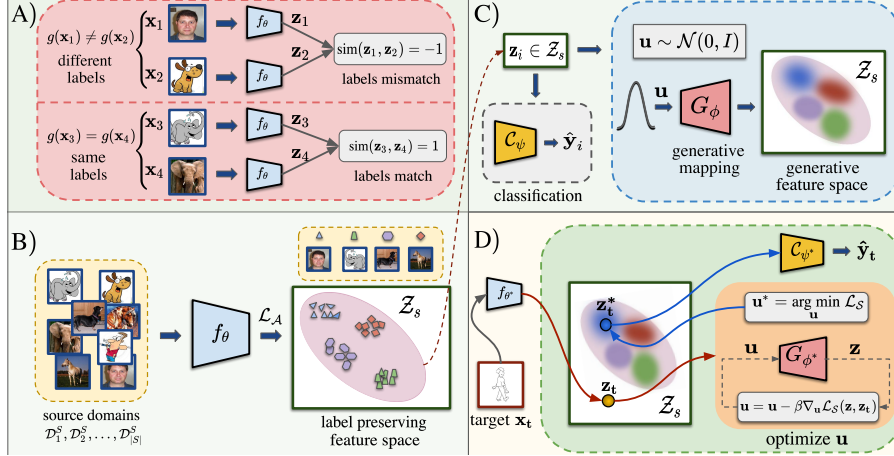


Figure 1: **A)** We design a function  $f$  (neural network  $f_\theta$ ) to learn a label-preserving metric that produces a similarity score of 1 when the ground truth labels (given by function  $g$ ) between a pair of images match and -1 otherwise. The function ‘sim’ refers to the cosine similarity function. **B)**  $f$  is implemented using a neural network  $f_\theta$ . During training, the examples from the source domains are utilized to create a source manifold  $\mathcal{Z}_s$  using loss  $\mathcal{L}_A$  such that the features on the manifold are implicitly *clustered* to preserve the labels of examples. **C)** A classifier  $\mathcal{C}_\psi$  and a generative model  $G_\phi$  are trained on the label-preserving features from manifold  $\mathcal{Z}_s$  such that  $G_\phi$  learns to map a Gaussian vector  $\mathbf{u}$  to a point on the manifold  $\mathcal{Z}_s$ . **D)** During inference,  $f_{\theta^*}$  projects target  $\mathbf{x}_t$  to a point  $\mathbf{z}_t$  on the label-preserving feature space. We propose an inference-time procedure to project the target feature to a point  $\mathbf{z}_t^*$  on the source manifold which is finally classified to predict its label  $\hat{\mathbf{y}}_t$ .  $\theta^*$ ,  $\psi^*$  and  $\phi^*$  indicate that the weights of their corresponding networks are fixed during inference.

1. **(Training):** Learn a label-preserving domain invariant representation using source data. We first transform the data from multiple source domains into a space where they are clustered according to class labels, irrespective of the domains and build a classifier on these features.
2. **(Training):** Learn to generate features from the domain invariant feature manifold created from the source data by constructing a generative model on it.
3. **(Inference):** Given a test target sample, project it on to the source-feature manifold in a label-preserving manner. This is done by solving an inference-time optimization problem on the input space of the aforementioned generative model. Finally, classify the projected target feature.

Note that the parameters of all the networks involved (Feature extractor, generative model and the classifier) are learned only using the source data and fixed during inference. Thus, our approach is well within the realm of DG despite a label-preserving optimization problem being solved for every target sample during inference. The overall procedure is depicted in Figure 1. In the subsequent sections, we describe all the aforementioned components in detail.

## 3.2. Label-Preserving Transformation

### 3.2.1 Domain invariant Features

The first step in our method is to learn a feature (metric) space such that the input images are clustered in accordance with their labels (classes) irrespective of their domains. We explicitly construct such a feature space  $\mathcal{F}$  with a function

$f : \mathcal{X} \rightarrow \mathcal{F}$  by solving the following optimization problem.

$$\arg \min_{\tilde{f}} \sum_{j=1}^N \sum_{i=1}^N (-1)^{\alpha(i,j)} \left\| \tilde{f}(\mathbf{x}_i) - \tilde{f}(\mathbf{x}_j) \right\|^2 \quad (1)$$

subject to  $\left\| \tilde{f}(\mathbf{x}_k) \right\| = 1 \quad \forall k \in [N]$

$$\alpha(i, j) = \begin{cases} 0 & g(\mathbf{x}_i) = g(\mathbf{x}_j) \\ 1 & \text{otherwise} \end{cases}$$

The function  $f$  is learned such that when a pair of source samples have the same ground truth labels, the norm of the difference between their representations under  $f$  is low irrespective of their domain membership and high when they belong to different classes. Under this formulation, the features in the space  $\mathcal{F}$  will be ‘clustered’ in accordance with their class labels, irrespective of the domains. In fact, it can be shown that the aforementioned  $f$  minimizes the  $\mathcal{H}$ -divergence [2] between any two pairs of domains.

**Proposition 1.** *The label-preserving transformation  $f$  defined in Eq. 1, reduces the  $\mathcal{H}$ -divergence between any two pair of domains on which it is learned (Proof in Appendix).*

In summary, the proposed feature transformation merges multiple source domains into a single feature domain such that images that have the same labels cluster into a group in the feature space.

### 3.2.2 Learning the $f$ -function

We propose to learn  $f$  by parameterizing it with a deep neural network,  $f_\theta$ . It is easy to see that the objective func-

tion in Eq. 1 reduces to an optimization of cosine-similarity between the pair of samples  $f_\theta(\mathbf{x}_i)$  and  $f_\theta(\mathbf{x}_j)$ . Suppose  $\mathbf{z}_i = f_\theta(\mathbf{x}_i)$  and  $\mathbf{z}_j = f_\theta(\mathbf{x}_j)$  represent the feature vectors of inputs, the cosine-similarity  $s_{i,j}$  between  $\mathbf{z}_i$  and  $\mathbf{z}_j$  is given by,

$$s_{i,j} = \frac{\mathbf{z}_i \cdot \mathbf{z}_j}{\|\mathbf{z}_i\| \|\mathbf{z}_j\|} \quad (2)$$

Note that the optimization problem in Eq. 1 seeks  $s_{i,j}$  to be high when the labels are same and low when they are different (denoted by the  $\alpha(i,j)$  term in Eq. 1). Thus, we first translate  $s_{i,j}$  into logits for a sigmoid activation and use binary cross entropy on the generated probabilities. However, since  $-1 \leq s_{i,j} \leq 1$ , we scale it with a small positive constant  $\tau$  (typically 0.1) to widen the range of the generated logits. Mathematically, we can write  $\mathbf{p}_{i,j} = \text{sigmoid}(s_{i,j}/\tau)$ . Under this formulation, one can treat  $\mathbf{p}_{i,j}$  as a similarity score which should be 1 if  $(\mathbf{x}_i, \mathbf{x}_j)$  have the same label ( $\alpha(i,j) = 0$ ) and 0 otherwise. Thus, we finally use a binary cross entropy loss  $\mathcal{L}_{\mathcal{A}}$  between  $\mathbf{p}_{i,j}$  and  $1 - \alpha(i,j)$ , to train the  $f_\theta$  network.

### 3.3. Inference-time Target Projections

In a DG setting, the feature transformation mentioned in the previous sections is learned on the source domains. Let  $\mathcal{Z}_s$  denote the manifold created by learning such features using the source data. The classifier  $\mathcal{C}_\psi$  is trained on points from the source data feature manifold  $\mathcal{Z}_s$ . Because of the domain shift, the feature  $f(\mathbf{x}_t)$  corresponding to a test target point  $\mathbf{x}_t$  might not fall on the source data feature manifold  $\mathcal{Z}_s$ . This causes the classifier to fail on the target feature  $f(\mathbf{x}_t)$ . To address this issue, we propose to project or ‘push’ the target feature onto  $\mathcal{Z}_s$  while preserving the ground truth label of  $\mathbf{x}_t$ , so that the classifier can better discern the class label of  $\mathbf{x}_t$ . We propose to accomplish such a label-preserving projection by solving an inference-time optimization problem on the input space of a generative model trained on  $\mathcal{Z}_s$ . For convenience, we define a function  $\tilde{g} : \mathcal{F} \rightarrow \mathcal{Y}$  such that  $\tilde{g}(f(x)) = g(x)$

#### 3.3.1 Generating the Source-Feature Manifold $\mathcal{Z}_s$

Once the label-preserving source-feature manifold  $\mathcal{Z}_s$  is obtained through  $f_\theta$ , we build a generative model on it. That is, a transformation from samples from an arbitrary distribution, such as Normal distribution, to the source-data feature manifold  $\mathcal{Z}_s$  is learned. We choose two state-of-the-art neural generative models for this purpose: (a) Variational Auto-Encoder (VAE) [18]: In this setting, a VAE is trained using the source-data features  $\mathbf{z}$ , by encoding them to produce the latent space  $\mathbf{u} \sim \mathcal{N}(0, I)$ . A decoder  $G_\phi$  reconstructs (generates) the source-feature manifold  $\mathcal{Z}_s$  by minimizing a regularized norm-based loss. (b) Generative Adversarial Networks (GAN) [11]: Here, a GAN is trained with a generator network  $G_\phi$  that maps an arbitrary latent space

$\mathbf{u} \sim \mathcal{N}(0, I)$  to the source-feature manifold  $\mathcal{Z}_s$ . Note that these generative models are trained on the source-features alone and fixed during the inference procedure. We denote the trained generative model by  $G_{\phi^*}$ .

#### 3.3.2 Label-preserving Projections

The final component of our method is to project the target features on to the source-feature manifold during inference. It is to be noted that the transformation  $f_\theta$  is constructed such that when a pair of samples have zero distance in that space, they have same ground truth label. That is, if  $\|f_\theta(\mathbf{x}_1) - f_\theta(\mathbf{x}_2)\| = 0$ , then  $g(\mathbf{x}_1) = g(\mathbf{x}_2)$ . We exploit this property and solve a (per-sample) optimization procedure on the input space of the generative model  $G_{\phi^*}(\mathbf{u})$  to obtain the target-feature projection.

Let  $\mathbf{z}_t = f_\theta(\mathbf{x}_t)$  denote the feature vector corresponding to a test target sample  $\mathbf{x}_t$ . Our goal is to find the projected target feature in the source-feature manifold  $\mathbf{z}_t^* \in \mathcal{Z}_s$  that has the same ground-truth label as that of  $\mathbf{z}_t$ . By construction of  $f_\theta$  (Eq. 1), the cosine distance between  $\mathbf{z}_t$  and  $\mathbf{z}_t^*$  should be low if their ground truth labels are to match. Based on this, we devise the following optimization problem on the input space of  $G_{\phi^*}$  to find  $\mathbf{z}_t^*$ :

$$\mathcal{L}_S = \left[ 1 - \frac{\mathbf{z}_t \cdot G_{\phi^*}(\mathbf{u})}{\|\mathbf{z}_t\| \|G_{\phi^*}(\mathbf{u})\|} \right] \quad (3)$$

$$\mathbf{u}^* = \arg \min_{\mathbf{u}} \mathcal{L}_S \quad (4)$$

$$\mathbf{z}_t^* = G_{\phi^*}(\mathbf{u}^*) \quad (5)$$

The objective function in Eq. 4 seeks to find the projected target feature  $\mathbf{z}_t^*$  (via  $\mathbf{u}^*$  from Eq. 5) whose cosine distance is least from the true target feature  $\mathbf{z}_t$ . The implicit assumption of our method is that when the target example is projected on to the source manifold, minimizing distance is equivalent to preserving labels.

#### 3.3.3 Analysis and Implementation

In this section, we analyze the performance of the classifier on account of it using the projected target instead of the real target features. We start by upper bounding the expected value of the misclassification when the projected target is used in the classifier, in the below proposition.

**Proposition 2.** *The expected misclassification rate obtained with a classifier  $h$  when the projected target is used instead of the true target, obeys the following upper-bound:*

$$\mathbb{E}_{(\mathcal{D}^T, \mathcal{D}^{T^*})} |\tilde{g}(\mathbf{z}_t) - h(\mathbf{z}_t^*)| \leq \underbrace{\mathbb{E}_{\mathcal{D}^{T^*}} |\tilde{g}(\mathbf{z}_t^*) - h(\mathbf{z}_t^*)|}_i + \underbrace{\mathbb{E}_{(\mathcal{D}^T, \mathcal{D}^{T^*})} |\tilde{g}(\mathbf{z}_t) - \tilde{g}(\mathbf{z}_t^*)|}_{ii} \quad (6)$$

where  $\mathcal{D}$  and  $\mathcal{D}^{T^*}$  respectively denote the true and the projected target distributions respectively (Proof in Appendix).



The term (i) in Eq. 6 is the misclassification error of  $h$  on the projected target and term (ii) is the difference between the ground truth labels of the true and the projected targets. Given that our overall objective is to minimize the LHS of Eq. 6, the optimization procedure in the previous section aims to minimize term (ii) while term (i) is expected to be less since the projected target is expected to lie on the source-feature manifold  $\mathcal{Z}_s$  on which the classifier is trained.

In the implementation, during inference, we optimize the objective in Eq. 4 (reducing term (ii) in Eq. 6) by gradient descent. A discussion on the choice of stopping criteria can be found in Section 4.7. The training and the inference procedures are detailed in Algorithm 1 and shown in Figure 1.

---

### Algorithm 1: Inference-time Target Projections

---

#### Training

**Input:** Batch size  $N$ , learning rate  $\eta$ , source data  $\{(\mathbf{x}_k, \mathbf{y}_k)\}$ ;

**Result:** Trained  $f_{\theta^*}$ , generative model  $G_{\phi^*}$ , classifier  $\mathcal{C}_{\psi^*}$

```

for sampled minibatch  $\{(\mathbf{x}_k, \mathbf{y}_k)\}_{k=1}^N$  do
  for all  $i \in \{1, \dots, N\}$  and  $j \in \{1, \dots, N\}$  do
     $(\mathbf{z}_i, \mathbf{z}_j) \leftarrow (f_{\theta}(\mathbf{x}_i), f_{\theta}(\mathbf{x}_j))$ 
     $\mathbf{s}_{i,j} \leftarrow \frac{\mathbf{z}_i \cdot \mathbf{z}_j}{\|\mathbf{z}_i\| \|\mathbf{z}_j\|}$ 
     $\mathbf{y}_{i,j} = \delta_{\mathbf{y}_i, \mathbf{y}_j}$ 
     $\mathbf{p}_{i,j} \leftarrow \text{sigmoid}(\mathbf{s}_{i,j} / \tau)$ 
  end
   $\mathcal{L}_{\mathcal{A}} \leftarrow \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \text{BCELoss}(\mathbf{p}_{i,j}, \mathbf{y}_{i,j})$ 
   $\theta \leftarrow \theta - \eta \nabla_{\theta} \mathcal{L}_{\mathcal{A}}$ 

```

**end**

Train  $G_{\phi}$  and  $\mathcal{C}_{\psi}$  on  $\{(f_{\theta}(\mathbf{x}_k), \mathbf{y}_k)\}$ .

---

#### Inference

**Input:** Target image  $\mathbf{x}_t$ , trained network  $f_{\theta^*}$ , generative model

$G_{\phi^*}$ , classifier  $\mathcal{C}_{\psi^*}$ , iteration rate  $\beta$ ;

**Result:** Target label  $\hat{\mathbf{y}}_t$

$\mathbf{z}_t \leftarrow f_{\theta^*}(\mathbf{x}_t)$ ;

Sample  $\mathbf{u}$  from  $\mathcal{N}(0, I)$ ;

Initialize  $U$  and  $L$  as empty lists

**for** all  $i \in \{1, \dots, M\}$  **do**

$\mathbf{z} \leftarrow G_{\phi^*}(\mathbf{u})$

$\mathcal{L}_{\mathcal{S}} \leftarrow 1 - \frac{\mathbf{z} \cdot \mathbf{z}_t}{\|\mathbf{z}\| \|\mathbf{z}_t\|}$

$\mathbf{u} \leftarrow \mathbf{u} - \beta \nabla_{\mathbf{u}} \mathcal{L}_{\mathcal{S}}$

$(U[i], L[i]) \leftarrow (\mathbf{u}, \mathcal{L}_{\mathcal{S}})$

**end**

Smoothen  $L$  by window-averaging

$\mathbf{u}^* \leftarrow U[\arg \max_i \delta^2 L]$

$\hat{\mathbf{y}}_t \leftarrow \mathcal{C}_{\psi^*}(G_{\phi^*}(\mathbf{u}^*))$

---

## 4. Experiments and Results

We have considered four standard DG datasets - PACS [21], VLCS [8], Office-Home [43] and Digits-DG [48] to demonstrate the efficacy of our method. All these datasets contain four domains out of which three are used as sources and the other as a target in a leave-one-out strategy. We use a VAE as the generative model  $G_{\phi}$  in all our main results owing to its stability of training vis-a-vis a GAN. The per-

formance metric is the classification accuracy and we compare against a baseline Deep All method: classifier trained on the combined source domains without employing any DG techniques. For each target domain, we have reported our average and standard deviation for five independent runs of the model. We also compare our method with the existing DG methods and report the results, dataset wise. We report the standard deviation as 0 for models which have not reported them. For each dataset, we use the validation set for selecting hyperparameters if it is available. Otherwise, we split the data from source domains and use the smaller set for hyperparameter selection. We also use data augmentation for regularizing the network  $f_{\theta}$ . For further details on the datasets, machine configuration and choice of hyperparameters, please refer to the Appendix.

Method	Art.	Cartoon	Sketch	Photo	Avg.
AlexNet					
Deep All	65.96±0.2	69.50±0.2	59.89±0.3	89.45±0.3	71.20
Jigen [5]	67.63±0.0	71.71±0.0	65.18±0.0	89.00±0.0	73.38
MMLD [27]	69.27±0.0	72.83±0.0	66.44±0.0	88.98±0.0	74.38
MASF [6]	70.35±0.3	72.46±0.2	67.33±0.1	90.68±0.1	75.21
EISNet [45]	70.38±0.4	71.59±1.3	70.25±1.4	91.20±0.0	75.86
RSC [16]	71.62±0.0	75.11±0.0	66.62±0.0	90.88±0.0	76.05
Ours	<b>72.67±0.5</b>	<b>76.51±0.3</b>	<b>73.09±0.2</b>	<b>92.01±0.3</b>	<b>78.57</b>
ResNet-18					
Deep All	77.65±0.2	75.36±0.3	69.08±0.2	95.12±0.1	79.30
MMLD [27]	81.28±0.0	77.16±0.0	72.29±0.0	96.09±0.0	81.83
EISNet [45]	81.89±0.9	76.44±0.3	74.33±1.4	95.93±0.1	82.15
L2A-OT [47]	83.30±0.0	78.20±0.0	73.60±0.0	96.20±0.0	82.80
DSON [37]	84.67±0.0	77.65±0.0	<b>82.23±0.0</b>	95.87±0.0	85.11
RSC [16]	83.43±0.8	80.31±1.8	80.85±1.2	95.99±0.3	85.15
Ours	<b>86.39±0.3</b>	<b>81.26±0.2</b>	81.79±0.1	<b>97.15±0.4</b>	<b>86.65</b>
ResNet-50					
Deep All	81.31±0.3	78.54±0.4	69.76±0.4	94.97±0.1	81.15
MASF [6]	82.89±0.2	80.49±0.2	72.29±0.2	95.01±0.1	82.67
EISNet [45]	86.64±1.4	81.53±0.6	78.07±1.4	97.11±0.4	85.84
DSON [37]	87.04±0.0	80.62±0.0	82.90±0.0	95.99±0.0	86.64
RSC [16]	87.89±0.0	82.16±0.0	83.85±0.0	97.92±0.0	87.83
Ours	<b>90.25±0.4</b>	<b>85.19±0.2</b>	<b>86.20±0.5</b>	<b>98.97±0.1</b>	<b>90.15</b>

Table 1: Comparison of performance between different models on PACS [21] dataset with AlexNet, ResNet-18 and ResNet-50 as backbones for the  $f_{\theta}$  network.

### 4.1. Multi-source Domain Generalization

**PACS:** The PACS dataset consists of images from Photo, Art Painting, Cartoon and Sketch domains. We follow the experimental protocol defined in [21]. We use ResNet-50, ResNet-18 and AlexNet as backbones for the feature extractor network  $f_{\theta}$  and train them on source domains. To perform target projection, we learn to sample from the features produced by  $f_{\theta}$  by training a VAE on the latent space ( $\mathbf{u}$ ). We achieve state-of-the-art results with all three choices of backbones as shown in Table 1.

**VLCS:** VLCS comprises of the VOC2007(Pascal), LabelMe, Caltech and Sun domains, all of which contain photos. We follow the same experimental setup as men-

Method	Caltech	LabelMe	Pascal	Sun	Avg.
Deep All	96.45±0.1	60.03±0.5	70.41±0.4	62.63±0.3	72.38
Jigen [5]	96.93±0.0	60.90±0.0	70.62±0.0	64.30±0.0	73.19
MMLD [27]	96.66±0.0	58.77±0.0	71.96±0.0	68.13±0.0	73.88
MASF [6]	94.78±0.2	64.90±0.1	69.14±0.2	67.64±0.1	74.11
EISNet [45]	97.33±0.4	63.49±0.8	69.83±0.5	68.02±0.8	74.67
RSC [16]	97.61±0.0	61.86±0.0	73.93±0.0	68.32±0.0	75.43
Ours	<b>98.12±0.1</b>	<b>66.80±0.3</b>	<b>74.77±0.4</b>	<b>70.43±0.1</b>	<b>77.53</b>

Table 2: Comparison of performance between different models using AlexNet backbone on VLCS [8] dataset.

tioned in [27] where we train on three source domains with 70% data from each and test on all the examples from the fourth target domain. We utilize similar setup for  $f_\theta$  and  $G_\phi$  as in the PACS dataset with AlexNet as the backbone. We achieve SOTA results on VLCS as evident from Table 2. We emphasize that unlike PACS dataset where the domains differ in image styles, VLCS consists of domains that contain only photos. Thus, we demonstrate that our method generalizes well even when the source domains are not diverse.

Method	Artistic	Clipart	Product	Real-World	Avg.
Deep All	52.06±0.5	46.12±0.3	70.45±0.2	72.45±0.2	60.27
D-SAM [7]	58.03±0.0	44.37±0.0	69.22±0.0	71.45±0.0	60.77
Jigen [5]	53.04±0.0	47.51±0.0	71.47±0.0	72.79±0.0	61.20
MMD-AAE [24]	56.50±0.0	47.30±0.0	72.10±0.0	74.80±0.0	62.70
DSON [37]	59.37±0.0	45.70±0.0	71.84±0.0	74.68±0.0	62.90
RSC [16]	58.42±0.0	47.90±0.0	71.63±0.0	74.54±0.0	63.12
L2A-OT [47]	60.60±0.0	50.10±0.0	74.80±0.0	77.00±0.0	65.60
Ours	<b>62.63±0.2</b>	<b>55.79±0.3</b>	<b>76.86±0.1</b>	<b>78.98±0.1</b>	<b>68.56</b>

Table 3: Comparison of performance between different models using ResNet-18 backbone on Office-Home [43].

Method	MNIST	MNIST-M	SVHN	SYN	Avg.
Deep All	95.24±0.1	58.36±0.6	62.12±0.5	78.94±0.3	73.66
Jigen [5]	96.50±0.0	61.40±0.0	63.70±0.0	74.00±0.0	73.90
CCSA [29]	95.20±0.0	58.20±0.0	65.50±0.0	79.10±0.0	74.50
MMD-AAE [24]	96.50±0.0	58.40±0.0	65.00±0.0	78.40±0.0	74.60
CrossGrad [38]	96.70±0.0	61.10±0.0	65.30±0.0	80.20±0.0	75.80
L2A-OT [47]	96.70±0.0	63.90±0.0	68.60±0.0	83.20±0.0	78.10
Ours	<b>97.99±0.1</b>	<b>66.52±0.4</b>	<b>71.31±0.3</b>	<b>85.40±0.5</b>	<b>80.30</b>

Table 4: Comparison of performance between different models on Digits-DG [48] dataset.

**Office-Home:** Office-Home contains images from 4 domains namely Artistic, Clipart, Product and Real-World. We follow the experimental protocol as outlined in [7]. We utilize a ResNet-18 as backbone with two additional fully-connected layers. The reconstruction error in VAE is minimized using both L1 and L2 loss functions. We achieve SOTA results on all the domains using a ResNet-18 backbone as shown in Table 3. It should be noted that Clipart is a difficult domain to generalize to as it is dissimilar to other domains. We achieve 5.6% improvement over the nearest competitor on the Clipart domain, showcasing the merit of our method in generalizing to dissimilar target domains.

**Digits-DG:** Digits-DG is the task of digit recognition us-

ing MNIST [20], MNIST-M [9], SVHN [32] and SYN [9] domains that differ drastically in font style and background. We follow the experimental setup of [48] and use their architecture for the feature extractor  $f_\theta$ . The  $G_\phi$  network is implemented in a similar way as it is implemented for the PACS dataset with ResNet-18 as the backbone. Table 4 shows the performance of our method on Digits-DG dataset compared against various SOTA methods.

## 4.2. Pairwise $\mathcal{H}$ -divergence

We examine the effectiveness of the proposed method in projecting target features onto the source-feature manifold by computing a proxy measure for  $\mathcal{H}$ -divergence called the  $\mathcal{A}$ -distance [3], that measures the divergence between two domains. We compute the  $\mathcal{A}$ -distance between features from the Sketch domain (target domain) and each of the three source domains (Photo, Cartoon and Art), obtained in two ways - first, from a ResNet-18 model trained on the source domains (Deep All), and second, the features obtained after performing target projections using our method. We compare the  $\mathcal{A}$ -distance obtained from these two methods in Figure 2a. It is observed that the source-target divergence is reduced substantially compared to Deep All, suggesting the effectiveness of our projection scheme in bringing the target points onto the source manifold.

## 4.3. Ablation Studies

To quantify the importance of the two components (i.e. the networks  $f_\theta$  and  $G_\phi$ ), we perform ablative studies using the Office-Home and Digits-DG datasets. For the Office-Home dataset, the Deep All model is trained with ResNet-18 as the backbone while the backbone proposed by [48] is used for the Deep All model on the Digits-DG dataset. We classify target domain features extracted from the feature extractor network without the target projection procedure as shown in Figure 2b. We observe a marginal improvement over Deep All by classifying on the domain-invariant features learnt by  $f_\theta$  without using the  $G_\phi$  network. We attribute this to the explicit class separability imposed by the  $f_\theta$  network in the feature space  $\mathcal{F}$ . When we exclude the  $f_\theta$  network by training the generative model on Deep All features and perform the projection procedure on these features, we still observe a sizable performance boost over merely classifying with the Deep All model. This shows the effectiveness of the target projection procedure as a sampling strategy. The best performance is achieved with all components working together as seen in Figure 2b.

## 4.4. Sampling strategy

To quantify the effect of the specific generative model used for the  $G_\phi$  network, we report performance with three samplers on VLCS dataset, namely: (a) VAE, (b) 1-Nearest Neighbor (1-NN) of the target sample with the source samples using the similarity metric as in Eq. 2, (c) GAN, and show results in Table 5. It is seen that sampling using con-




Figure 2: **(left to right)** (a) Plot showing  $\mathcal{A}$ -divergence between source and projected target features. (b) Ablation on Office-Home dataset highlighting the relative importance of different components employed in our method. (c) Relative performance of different methods trained on 20%, 40%, 60% and 80% of PACS dataset. For each setting, we report average performance with all examples from the target domains using leave-one-out strategy. (d) Relative performance of DG methods trained on a single domain (Sketch (S)) and tested on Photo (P), Cartoon (C) and Art (A) domains.

tinuous generative models (VAE and GAN) offer better performance compared to 1-NN sampling. On domains like LabelMe and Sun, they offer improvement of about 5.7% and 7.0% respectively over the 1-NN method of sampling. This is attributed to the fact that generative models can approximate the true data distribution by learning from empirical data and offer infinite sampling while 1-NN search is restricted to the existing training points only.

Method	Caltech	LabelMe	Pascal	Sun	Avg.
Deep All	96.45	60.03	70.41	62.63	72.38
1-NN	96.51	61.44	71.82	63.46	73.31
Ours ( $G_\phi = \text{GAN}$ )	97.89	<b>67.18</b>	74.59	70.28	77.48
Ours ( $G_\phi = \text{VAE}$ )	<b>98.12</b>	66.80	<b>74.77</b>	<b>70.43</b>	<b>77.53</b>

Table 5: Performance of our method on VLCS dataset with different sampling strategies.

## 4.5. Low resource settings

In this section, we demonstrate the efficacy of our method in low-resource settings. We show that our method generalizes well when used in scarce resource settings and single domain DG problems, and can also easily be extended to supervised domain adaptation with access to scant labelled target samples. Our method has a two-fold advantage that makes it especially data-efficient: (a) since  $f_\theta$  is trained on pairs of images, it can learn effectively even on small datasets (b) the generative model  $G_\phi$  enables infinite sampling during the target projection procedure.

### 4.5.1 Scarce resource setting

We compare our method against Deep All and two state-of-the-art methods, RSC [16] and Jigen [5]. For each source domain in PACS we train on  $\{20\%, 40\%, 60\%, 80\%\}$  data from each domain and test on the entire target domain. From Figure 2c, it is evident that our method outperforms all the baselines at every considered resource setting.




Figure 3: **(left)** (a) Effect of deviation from the optimal number of iterations of each target example during inference on PACS. **(right)** (b) Plot of iterations vs loss for projection of a target example from Sketch domain on source manifold ( $\mathcal{Z}_s$ ) created by Photo, Art and Cartoon domains.

### 4.5.2 Supervised Domain Adaptation

This setting has also been referred to as Few-shot Domain Adaptation in [36] and [28]. In this setting, in addition to the source data, we assume that we have access to a limited number of labelled samples ( $|\mathcal{T}|$ ) from the target domain at train time. We train the feature extractor network  $f_\theta$ , generative model  $G_\phi$  and the classifier  $\mathcal{C}_\psi$  on the source data and fine tune on the target samples. We compare our method against M-ADA [36] and FADA [28] on the Digits dataset. The results are presented in Table 6. We outperform our nearest competitors by a significant margin, thus highlighting the adaptability of our method to this use-case.

### 4.5.3 Single Source Domain Generalization

In single source DG, we only have access to a single domain during training and aim to generalize to all other unseen domains. We train on the Sketch domain of the PACS dataset and test on the other three domains (i.e. Photo, Art Painting and Cartoon). We compare our method against Jigen [5] and RSC [16]. The results are presented in Figure 2d. We also examine the individual effects of each of the components  $f_\theta$  and  $G_\phi$  of our method in this scenario by examining the per-

Method	$ T $	U→M	M→S	S→M	Avg.
FADA [28]	7	91.50	47.00	87.20	75.23
CCSA [29]	10	95.71	37.63	94.57	75.97
M-ADA [36]	0	71.19	36.61	60.14	55.98
	7	92.33	56.33	89.90	79.52
	10	93.67	57.16	91.81	80.88
Ours	0	74.52	42.96	64.12	60.53
	7	93.81	58.92	92.02	81.58
	10	<b>96.10</b>	<b>60.07</b>	<b>95.33</b>	<b>83.83</b>

Table 6: Comparison of few-shot domain adaptation performance between different models on Digits dataset (MNIST [20] (M), USPS [17] (U) and SVHN [32] (S)).

Method	S → P	S → C	S → A
Deep All	29.88	32.47	30.56
Ours (w/o $G_\phi$ )	33.76	37.94	36.02
Ours (w/o $f_\theta$ )	50.39	66.82	44.94
Ours ( $f_\theta + G_\phi$ )	53.82	70.33	50.61

Table 7: Ablation on Single Source DG with Sketch (S) as source and Photo (P), Cartoon (C) and Art (A) as unseen target domains.

formance without each of them (Table 7). We observe that without  $G_\phi$ , the performance difference between Deep All and our method is substantially lower than the improvement obtained by performing target projections on Deep All features, highlighting the effectiveness of the projection procedure. The best performance is obtained when both are used together, since the projection procedure effectively utilizes the label-preserving metric defined by  $f_\theta$ .

#### 4.6. Robust Domain Generalization

We examine the robustness of our method against different types of corruptions. We benchmark against the CIFAR-10-C dataset [12], which consists of images with 19 corruptions types applied at five levels of severities on the test set of CIFAR-10. We follow the protocol detailed in [36] and train our model on the CIFAR-10 dataset using a Wide Residual Network (WRN) backbone [46]. The results are presented severity level-wise in Table 8 and by type of corruption in Figure 4a. Similar to 4.5.3, we observe the effectiveness of  $G_\phi$  in generalizing to corruptions.

Method	Level 1	Level 2	Level 3	Level 4	Level 5
ERM [19]	87.8±0.1	81.5±0.2	75.5±0.4	68.2±0.6	56.1±0.8
M-ADA [36]	90.5±0.3	86.8±0.4	82.5±0.6	76.4±0.9	65.6±1.2
Ours	<b>93.6±0.2</b>	<b>89.2±0.4</b>	<b>85.3±0.1</b>	<b>79.0±0.3</b>	<b>68.2±0.6</b>

Table 8: Comparison of performance on CIFAR-10-C with 5 different levels of severity. Accuracy averaged over all 19 corruption levels

#### 4.7. Determining the stopping criteria

The inference-time iterative optimization process discussed in 3.3.3 needs a stopping criteria, since stopping




Figure 4: (left) (a) Performance comparison between M-ADA and our method on 10 out of the 19 corruption types at severity level 5 from the CIFAR-10-C dataset. (right) (b) Plot showing relative importance of each component (right) of our method on the CIFAR-10-C dataset. Accuracy averaged over all 19 corruption levels.

too early would not guarantee the label preservation, while stopping too late may take the projected target to a low-probability region on the source manifold.

We address this issue heuristically: stop the iteration process at the “elbow-point” (maxima of the second derivative) of the loss curve as a function of the number of iterations (denoted by  $n^*$ ). This choice is inspired by the observation that the elbow-point reflects the point of diminishing returns; for a constant iteration rate ( $\beta$ ), the loss decreases at a slower rate beyond this point.

As empirical evidence, we vary the stopping point for each target sample by a fixed number of iterations  $\epsilon$  around the maxima  $n^*$  of the second derivative values of the sample and calculate the accuracy on the target projections so obtained. The plot of accuracy vs  $\epsilon$  is shown in Figure 3a. We observe that the highest accuracy is obtained around  $\epsilon = 0$  which indicates the correctness of the stopping criteria  $n^*$ . For negative values of  $\epsilon$ , the accuracy drop can be explained by inadequate label preservation, while the projected target leaves the source manifold for positive values of  $\epsilon$ .

## 5. Conclusion

We propose a novel Domain Generalization technique where the source domains are utilized to learn a domain invariant label-preserving metric space. During inference, every target sample is projected onto this space so that the classifier trained on the source features can generalize well on the projected target sample. We have demonstrated that this method yields SOTA results on Multi Source, Single Source and Robust Domain Generalization settings. In addition, the data-efficiency of the method makes it suitable to work well in Low Resource settings. Future iterations of work could attempt to extend this method to Domain Generalization for Segmentation and Zero-Shot Learning.



## References

- [1] Yogesh Balaji, Swami Sankaranarayanan, and Rama Chellappa. Metareg: Towards domain generalization using meta-regularization. In *Advances in Neural Information Processing Systems*, pages 998–1008, 2018.
- [2] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- [3] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Advances in neural information processing systems*, pages 137–144, 2007.
- [4] Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3722–3731, 2017.
- [5] Fabio M Carlucci, Antonio D’Innocente, Silvia Bucci, Barbara Caputo, and Tatiana Tommasi. Domain generalization by solving jigsaw puzzles. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2229–2238, 2019.
- [6] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. In *Advances in Neural Information Processing Systems*, pages 6450–6461, 2019.
- [7] Antonio D’Innocente and Barbara Caputo. Domain generalization with domain-specific aggregation modules. In *German Conference on Pattern Recognition*, pages 187–198. Springer, 2018.
- [8] 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 (ICCV)*, December 2013.
- [9] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pages 1180–1189, 2015.
- [10] Muhammad Ghifary, W Bastiaan Kleijn, Mengjie Zhang, and David Balduzzi. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, pages 2551–2559, 2015.
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [12] Dan Hendrycks and Thomas Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. In *International Conference on Learning Representations*, 2019.
- [13] Alex Hernández-García and Peter König. Further advantages of data augmentation on convolutional neural networks. In *International Conference on Artificial Neural Networks*, pages 95–103. Springer, 2018.
- [14] Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pages 1989–1998, 2018.
- [15] Shoubo Hu, Kun Zhang, Zhitang Chen, and Laiwan Chan. Domain generalization via multidomain discriminant analysis. In *Uncertainty in artificial intelligence: proceedings of the... conference. Conference on Uncertainty in Artificial Intelligence*, volume 35. NIH Public Access, 2019.
- [16] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. *arXiv preprint arXiv:2007.02454*, 2020.
- [17] Jonathan J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- [18] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [19] Vladimir Koltchinskii. *Oracle Inequalities in Empirical Risk Minimization and Sparse Recovery Problems: Ecole d’Été de Probabilités de Saint-Flour XXXVIII-2008*, volume 2033. Springer Science & Business Media, 2011.
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [21] 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.
- [22] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- [23] Da Li, Jianshu Zhang, Yongxin Yang, Cong Liu, Yi-Zhe Song, and Timothy M Hospedales. Episodic training for domain generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1446–1455, 2019.
- [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.
- [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.
- [26] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in neural information processing systems*, pages 136–144, 2016.
- [27] Toshihiko Matsuura and Tatsuya Harada. Domain generalization using a mixture of multiple latent domains. In *AAAI*, pages 11749–11756, 2020.
- [28] Saeid Motiian, Quinn Jones, Seyed Iranmanesh, and Gianfranco Doretto. Few-shot adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 6670–6680, 2017.

- [29] Saeid Motiian, Marco Piccirilli, Donald A Adjeroh, and Gianfranco Doretto. Unified deep supervised domain adaptation and generalization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5715–5725, 2017.
- [30] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18, 2013.
- [31] Zak Murez, Soheil Kolouri, David Kriegman, Ravi Ramamoorthi, and Kyungnam Kim. Image to image translation for domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4500–4509, 2018.
- [32] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bisacco, Bo Wu, and Andrew Y Ng. Reading digits in natural images with unsupervised feature learning. 2011.
- [33] Pau Panareda Busto and Juergen Gall. Open set domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 754–763, 2017.
- [34] Prashant Pandey, Aayush Kumar Tyagi, Sameer Ambekar, and Prathosh. Skin segmentation from nir images using unsupervised domain adaptation through generative latent search. *arXiv preprint arXiv:2006.08696*, 2020.
- [35] Vihari Piratla, Praneeth Netrapalli, and Sunita Sarawagi. Efficient domain generalization via common-specific low-rank decomposition. *arXiv preprint arXiv:2003.12815*, 2020.
- [36] 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.
- [37] Seonguk Seo, Yumin Suh, Dongwan Kim, Jongwoo Han, and Bohyung Han. Learning to optimize domain specific normalization for domain generalization. *arXiv preprint arXiv:1907.04275*, 2019.
- [38] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. In *International Conference on Learning Representations*, 2018.
- [39] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [40] Yu Sun, Xiaolong Wang, Zhuang Liu, John Miller, Alexei Efros, and Moritz Hardt. Test-time training with self-supervision for generalization under distribution shifts. In *International Conference on Machine Learning*, pages 9229–9248. PMLR, 2020.
- [41] Antonio Torralba and Alexei A Efros. Unbiased look at dataset bias. In *CVPR 2011*, pages 1521–1528. IEEE, 2011.
- [42] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7167–7176, 2017.
- [43] Hemant 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.
- [44] Riccardo Volpi, Hongseok Namkoong, Ozan Sener, John C Duchi, Vittorio Murino, and Silvio Savarese. Generalizing to unseen domains via adversarial data augmentation. In *Advances in neural information processing systems*, pages 5334–5344, 2018.
- [45] Shujun Wang, Lequan Yu, Caizi Li, Chi-Wing Fu, and Pheng-Ann Heng. Learning from extrinsic and intrinsic supervisions for domain generalization. *arXiv preprint arXiv:2007.09316*, 2020.
- [46] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. In Edwin R. Hancock Richard C. Wilson and William A. P. Smith, editors, *Proceedings of the British Machine Vision Conference (BMVC)*, pages 87.1–87.12. BMVA Press, September 2016.
- [47] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Learning to generate novel domains for domain generalization. *arXiv preprint arXiv:2007.03304*, 2020.
- [48] Kaiyang Zhou, Yongxin Yang, Timothy M Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *AAAI*, pages 13025–13032, 2020.