

Localized Adversarial Domain Generalization

Wei Zhu^{1,2} * Le Lu³ Jing Xiao⁴ Mei Han¹ Jiebo Luo² Adam P. Harrison¹

¹ PAII Inc. ² University of Rochester ³ Alibaba DAMO Academy ⁴ PingAn Insurance Group

{zwviews, tiger.lelu, jiebo.luo, adam.p.harrison}@gmail.com

xiaojing661@pingan.com.cn, hanmei613@piai-labs.com

Abstract

Deep learning methods can struggle to handle domain shifts not seen in training data, which can cause them to not generalize well to unseen domains. This has led to research attention on domain generalization (DG), which aims to the model’s generalization ability to out-of-distribution. Adversarial domain generalization is a popular approach to DG, but conventional approaches (1) struggle to sufficiently align features so that local neighborhoods are mixed across domains; and (2) can suffer from feature space over collapse which can threaten generalization performance. To address these limitations, we propose localized adversarial domain generalization with space compactness maintenance (LADG) which constitutes two major contributions. First, we propose an adversarial localized classifier as the domain discriminator, along with a principled primary branch. This constructs a min-max game whereby the aim of the featurizer is to produce locally mixed domains. Second, we propose to use a coding-rate loss to alleviate feature space over collapse. We conduct comprehensive experiments on the Wilds DG benchmark to validate our approach, where LADG outperforms leading competitors on most datasets.

1. Introduction

Deep neural networks can suffer from poor generalization performance on out-of-distribution (OOD) data from unseen domains, *i.e.*, from *domain shift*. For this reason, techniques to improve generalization abilities have gained increasing attention over the past decade, *e.g.* domain generalization (DG) and domain adaptation (DA). DA requires access to data from testing domains during training. In contrast, DG aims to construct a generalized model by exposing the training process to multiple training domains without exposure to OOD data from testing domains. As such, DG can reflect challenges in many real-life applications.

*Work was done while Wei Zhu interned at PAII Inc.

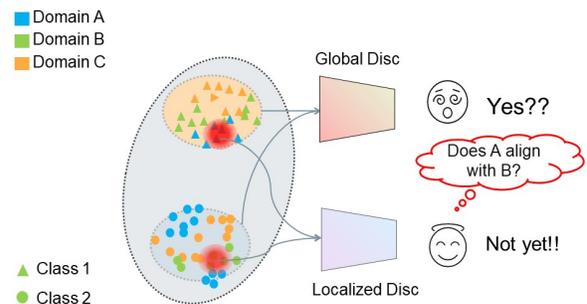


Figure 1. Localized discriminator can align two domains in a more fine-grained way.

Different DG methods have been proposed, including empirical risk minimization (ERM)-based methods [40, 52, 59], meta-learning based methods [28], domain-invariant representation based methods [16, 29, 33, 44], invariant risk minimization based methods [3, 54], and gradient agreement methods [39, 42]. Among them, domain invariant representation based methods, specifically adversarial domain generalization (ADG) methods, seem to be the most popular according to recent literature [19, 43]. ADG methods are inspired by generative adversarial networks [18] and learn a common feature space by adversarial learning for training domains [16]—the common feature space is expected to help generalization to unseen domains [16]. Although intuitively reasonable and technically sound, most of these methods show little performance gain over the baseline ERM in practice as indicated by recent benchmarks [19, 25].

In this work, we argue that two issues limit the performance impact of current ADG methods. First, we find that ERM feature representations are surprisingly already roughly aligned, with features clustered class-wisely regardless of their domain labels. The discrepancy between domains can be observed at local regions where local neighborhoods are not mixed across domains. Since ADG methods operate under the assumption of significant domain shift, this less obvious domain-level discrepancy challenges

conventional ADG methods. Second, we measure the compactness of the feature space using three different metrics and found that the featurizer may trivially fool the discriminator by over collapsing the feature space. The collapsed feature space can cause overfitting [50].

To address these limitations, we propose localized adversarial domain generalization with space compactness maintenance (LADG). As shown in Fig. 1, LADG incorporates a localized domain classifier [6] with adversarial learning. Since the predictions of local classifiers are made according to samples around a target sample, a local domain classifier can thus be used to describe the quality of alignment of local regions. We adopt label propagation in this paper as it is effective and differentiable [32, 53]. We also outline how to integrate localized classification properly within a generator loss that encourages neighborhood mixing. To alleviate feature space over collapse, LADG measures and maintains the compactness of the feature space using a differentiable coding ratio [50].

We summarize our contributions as follows:

1. We argue that ERM can already roughly align training domains with a linear primary task predictor, which undermines the assumption of most existing ADG methods. We also observe that applying ADG to the feature space can lead to feature space over collapse.
2. We propose localized adversarial domain generalization (LADG) to alleviate these two limitations. LADG adopts a principled label propagation as the domain discriminator to allow a fine-grained domain alignment and penalizes the space collapse caused by adversarial learning.
3. We conduct extensive experiments on benchmark datasets to verify our observations and show the effectiveness of the proposed method.

2. Related Work

DG is an active research topic with varied approaches. For instance, some methods enhance the generalization ability of ERM directly, *e.g.*, using mixup-based methods to generate novel data composed of mixtures of domains [47, 48, 52, 59]. Other methods try to improve the performance over the worst groups or domains [21, 40]. Recently, self-supervised learning has also been applied to DG, which is inspired by its success on general representation learning [15, 24, 36]. Nonetheless, domain-invariant learning has attracted increasing attention for DG, and can be categorized as either representation invariant [16, 17, 29, 33, 41, 44, 54], predictor invariant [1, 3, 11, 26], and gradient invariant methods [38, 39, 42].

We focus on domain invariant representation learning, specifically adversarial domain generalization (ADG) [2,

16, 29, 33, 35, 58] in this paper. Domain adversarial neural network (DANN) adopts a gradient reversal layer and updates the featurizer to fool the domain discriminator by generating domain-invariant representations [16]. Conditional DANN (CDANN)’s discriminator takes the primary task label into consideration when distinguishing samples from different domains [30, 33]. MMD-AAE adopts maximum mean discrepancy (MMD) to align domains, matching the latent representation to a prior distribution by adversarial learning [29]. Matsuura *et al.* extend DANN [35] with unknown domain labels. Zhou *et al.* adopt a domain transformation network to augment the training data to different training domains for adversarial learning [58]. Sicilia *et al.* propose to interpolate the feature space by using the gradient of the domain discriminator [43]. All these works operate under the assumption of significant domain shifts in the feature space between training domains. However, we empirically observe that the domain-level discrepancy is not as obvious as expected under standard benchmark domain generalization settings. As such, LADG is developed to align the domains in a more fine-grained and local way. Moreover, we propose a loss to prevent space collapse that could potentially benefit all ADG methods and also other fields using adversarial learning, *e.g.*, domain adaptation [56] and adversarial debiasing [23, 51].

3. Adversarial Domain Generalization

We focus on the problem of domain generalization. In this setting, the training set, $\mathcal{D}_{tr} = \{x_i, y_i, e_i\}_i^n$, is composed of data from S different domains, *i.e.*, $\mathcal{D}_{tr} = \{\mathcal{D}_s\}_s^S$, where $\mathcal{D}_s = \{x_i, y_i, s\}_{i=1}^{n_s}$ selects the samples from the s -th domain, x_i is the training data, y_i denotes the ground-truth label for the primary task, and the categorical value, e_i , denotes the domain label. We note that the primary task label y_i can take different forms for different tasks, *e.g.* a categorical value for classification and a continuous value for regression. We adopt a deep neural network to handle the problem, which contains a featurizer ϕ and a predictor w . The features extracted by ϕ are denoted as $H = \{H_s\}_s^S$, where $H_s = \{\mathbf{h}_i\}_{i=1}^{n_s} = \{\phi(x_i)\}_{i=1}^{n_s}$ selects the learned features for s -th domain. If the primary task is classification, we can also use $H^y = \{\phi(x_i) \mid \forall y_i = y\}_{i=1}^{n_y}$ to denote the features extracted from samples of the y -th class.

It is reasonable and common to implement the predictor, ϕ , using a single *linear layer* following standard benchmarks [19]. We will follow the same practice, which will allow us to directly measure the impact of our contributions compared to alternative approaches. The objective of domain generalization is to obtain a model that achieves superior performance on testing data from unseen domains.

A highly popular approach is adversarial domain generalization (ADG) [16, 33, 35, 43]. These are inspired by gen-

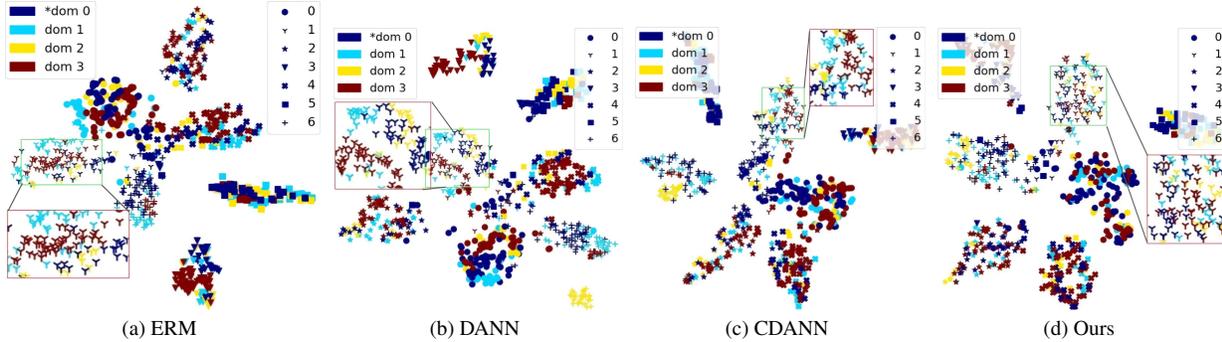


Figure 2. Visualization of learned representation on PACS dataset. * denotes testing domain. Different shape(color) represents different class(domain).

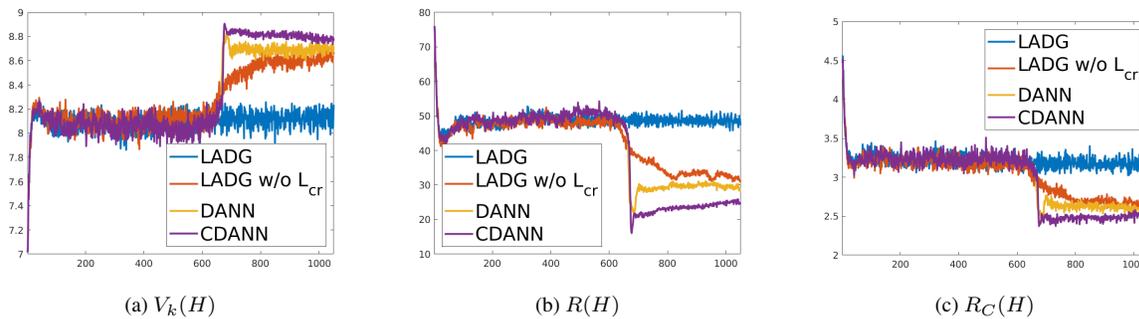


Figure 3. The feature space will collapse with adversarial domain alignment. We pretrain the model for 600 steps with ERM and then ADG methods are applied. Our method trained with \mathcal{L}_{cr} can alleviate space collapse.

erative adversarial networks [18] and handle the problem by learning a domain-invariant representation with the aid of a domain discriminator, denoted as η . We use the seminal DANN [16] work as an exemplar to illustrate the approach. Given the training data, DANN trains the featurizer ϕ and the classifier w to minimize the primary task loss:

$$\mathcal{L}_t(w, \phi) = \frac{1}{n} \sum_{x_i, y_i \in \mathcal{D}_{tr}} \ell(w \circ \phi(x_i), y_i), \quad (1)$$

where $\ell(\cdot)$ is the primary-task loss function. The featurizer representations likely overfit to the training domains and would struggle under the presence of domain shifts. This can lead to serious performance drops when testing the trained model on unseen domains. ADG methods aim to learn domain-invariant features to improve the generalization ability by using the domain discriminator, η , to distinguish samples from different domains using the extracted features h_i by minimizing the domain classification loss:

$$\mathcal{L}_{dom}(\eta) = -\frac{1}{n} \sum_{x_i, e_i \in \mathcal{D}_{tr}} \log(p_i), \quad (2)$$

where p_i is the pseudo-probability for the ground-truth domain and the dependence on x_i is implied. DANN uses the

cross-entropy loss to train the domain classifier η . To obtain domain-invariant representations, the feature extractor ϕ is adversarially trained to maximize \mathcal{L}_{dom} so that the extracted features H are indistinguishable to the domain classifier. We summarize the training objective as follows:

$$\begin{aligned} \mathcal{L}_{gen} &:= \min_{w, \phi} \mathcal{L}_t(w, \phi) - \lambda \mathcal{L}_{dom}(\eta) \\ \mathcal{L}_{disc} &:= \min_{\eta} \mathcal{L}_{dom}(\eta), \end{aligned} \quad (3)$$

where λ is used to balance the adversarial domain classification loss and the primary task loss.

3.1. Limitations

Although ADG methods have achieved progress, they do not show significant improvements over empirical risk minimization (ERM) according to recent benchmarks [19]. We argue that two issues limit the performance of current ADG approaches: incomplete alignment and feature space over collapse. To illustrate this, we use three existing methods as exemplars: ERM, DANN [16], and CDANN [33]. DANN learns to match the feature distributions across domains $P(H^s)$ where $H^s = \phi(X^s)$, while CDANN aims to match the class-conditioned feature distributions. We

provide observations and analysis using experiments conducted on the PACS dataset [27], which is a commonly used image classification dataset for domain generalization that comprises 7 different classes from four domains $S \in \{art, cartoons, photo, sketches\}$. Additional analyses on other datasets can be found in the supplementary materials.

3.1.1 Incomplete Alignment

We train ADG methods under the benchmark setting with a linear predictor w and use tSNE [46] to visualize the learned representations. The results are shown in Fig. 2. ADG methods operate under the assumption that the latent representation of samples from different training domains learned by ERM would be located at different sub-spaces, *i.e.*, domain shift, and learning to align the data distributions across different domains would lead to a common representation space, *i.e.*, a domain-invariant representation. However, observations from the visualization results of Fig. 2a question this assumption. First, when using ERM, the features for different domains are surprisingly already roughly aligned and the domain shift is not as obvious as expected. Instead, samples from the same class are grouped together regardless of their domain labels. Second, the features from same domains often form localized clusters and the clusters will be spread over the whole space. The localized clusters are also cross-distributed across domains, which leads to localized domain-level statistical difference across domains. This suggests that ERM already incompletely aligns features, *i.e.*, in a gross but non-local manner.

Since features are already at least partly aligned, the impact of additional measures that aim to align the domain-level distributions $P(H^s)$, *e.g.*, DANN [16] and MMD [29], can be limited. In principle, ADG approaches could more completely align distributions across domains, *i.e.*, produce local mixing without the localized clusters of Fig. 2. Yet, in practice it is extremely difficult to train classification-based discriminators to learn a rich and complete enough representation of the domain distributions that actually describes the degree of local neighborhood-level mixing. As Fig. 2b demonstrates, DANN features still exhibit the localized clusters of ERM.

An intuitive way to alleviate this problem is to align the domain distributions conditioned on the class label $P(H_y^s)$ as adopted by CDANN [30, 33]. However, the conditioned domain distributions are still not well mixed (see Fig. 2c) because of the intra-class intra-domain heterogeneity. Additionally, the class-conditioned domain alignment may partially dismiss the fact that the domain information contained in samples from different classes could mutually benefit the domain discrimination across classes. Finally, it is nontrivial to apply class-conditioned methods to general domain

adaptation problems, *e.g.*, regression or object detection.

To mitigate the problem, we introduce localized adversarial domain generalization (Section 4.1), which aligns the domain features at all local regions by incorporating a localized classifier with adversarial learning [6].

3.1.2 Feature Space Collapse

Another concern for ADG is feature space over collapse. An ideal result for ADG is to make features to be domain-invariant while maintaining the compactness of the feature space and keeping the primary task loss \mathcal{L}_t small. When updating the feature extractor ϕ to fool the domain discriminator by maximizing the domain classification loss \mathcal{L}_{dom} , a trivial way is to collapse the whole space or class-wisely collapse the feature space so that the \mathcal{L}_{dom} could be maximized. To verify our assumption that the feature space tends to be over compact with ADG, we adopt three different strategies to measure the compactness of the learned feature space $H \in \mathcal{R}^{n \times d}$:

1. We calculate the average degree of the k-nearest neighbors (KNN) graph. That is, we average the sum of cosine similarity between the i -th sample and its K-nearest neighbors:

$$V_k(H) = \frac{1}{N} \sum_{i=1}^N \sum_{j \in \mathcal{N}_i} \frac{\mathbf{h}_i^\top \mathbf{h}_j}{\|\mathbf{h}_i\| \|\mathbf{h}_j\|}, \quad (4)$$

where the set \mathcal{N}_i represents the neighbors of the i -th sample. $S_k(H)$ can be used to monitor the local density;

2. We calculate coding rate following [50]:

$$R(H) = \frac{1}{2} \log \det \left(I + \frac{d}{n\epsilon^2} \mathbf{H}^\top \mathbf{H} \right), \quad (5)$$

where ϵ denotes the precision and we use $\mathbf{H} = [\mathbf{h}_1 \dots \mathbf{h}_n]^\top$ to represent a set of features in matrix form. $R(H)$ indicates the number of bits needed to encode the data H up to a precision ϵ , and thus can be seen as a measurement of the compactness of the feature space. Please refer to [34, 50] for details. We normalize H to eliminate scale effects when calculating $R(H)$;

3. We average over the class-wise coding rate for the learned features as [50]

$$R_C(H) = \sum_{y=1}^C \frac{N_y}{N} R(H^y, \epsilon). \quad (6)$$

$R_C(H)$ can be used to measure the compactness for the class-wise feature space [50].

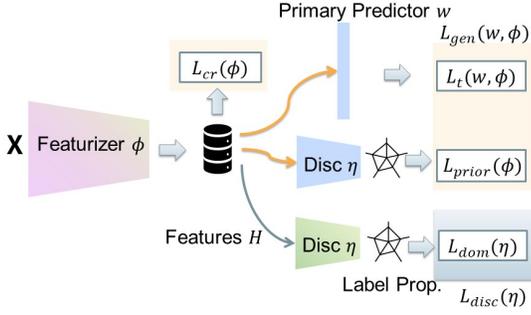


Figure 4. Block diagram of our method.

For simplicity, we calculate all measurements with samples in the minibatch, and the results are shown in Fig. 3. According to our experiments, when we train the model with ERM, all feature space compactness measurements remain stable. In contrast, we observe a sharp decrease of all measurements when we apply ADG. This implies that the features extracted by ϕ tend to collapse to trivially maximize the loss of the domain discriminator. The collapse often leads to overfitting and poses a threat to the generalization performance for unseen domains [50]. While the feature space is expected to shrink with ADG by removing features correlated with domain information, the space is at high risk for over collapse. As shown in Fig. 3c, we alleviate the over collapse by proposing a loss to encourage to maintain the compactness of the feature space.

4. Our Method

We introduce localized adversarial domain generalization with space compactness maintenance (LADG) to deal with the above two limitations. Fig. 4 outlines our approach. LADG applies a localized classifier to examine every local region to achieve instance-wise domain alignment. Additionally, LADG adopts a coding rate inspired loss to prevent space collapse.

4.1. Local Classification

To deal with the problem of incomplete alignment, LADG adopts a local classifier to ADG. Local classifiers predict the label for each sample based on the local regions of interests around the target sample [6]. The general idea here is that if a discriminator can correctly predict a sample’s domain from its neighbors’ domains, then local neighborhoods are not well mixed. Fooling such a discriminator can lead to more complete distributional alignment and better domain generalization performance. Typical local classification methods include KNN, decision trees, random forests [8], and localized logistic regression [20]. LADG instead adopts label propagation, which is effective, differentiable, and can be seamlessly incorporated with deep

neural networks. Label propagation has been successfully applied to different problems, including few-shot learning [32], node classifiers [53], metric learning [31], debiasing [60], semi-supervised learning [22], and domain adaptation [10, 55], but is not well studied for domain generalization.

Algorithm 1: Localized ADG

Input: the training set: \mathcal{D}_{tr} , $\rho = 0.2$, $\xi = 0.99$

Output: the trained weight ϕ , w

while the maximal iterations are not reached **do**

$\{x_i, y_i, e_i\}_{i=1}^n \sim \mathcal{D}_{tr}$ // sample minibatch

$\{\mathbf{h}_i\}_{i=1}^n = \{\phi(x_i)\}_{i=1}^n$ // extract features

/* Update Domain Disc η */

Obtain domain prediction p_{ij} by Alg. 2

Update η by minimizing \mathcal{L}_{disc} (13)

/* Update ϕ and w */

Obtain p_{ij} by Alg. 2 with updated η

$\mathcal{L}_t(w, \phi) = \frac{1}{n} \sum_{i=1}^n \ell(w \circ \phi(x_i), y_i)$ (1)

$\mathcal{L}_{prior}(\phi) = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^S q_j \log(p_{ij})$ (11)

$\mathcal{L}_{cr}(\phi) = \frac{1}{\rho} \log \cosh(\rho(R(H) - \bar{R}(H)))$ (12)

Update ϕ and w by minimizing \mathcal{L}_{gen} (13)

$\bar{R}(H) \leftarrow \xi \bar{R}(H) + (1 - \xi)R(H)$

end

Algorithm 2: Label Propagation

Input: Extracted features $\{\mathbf{h}_i\}_{i=1}^n$, domain label $\{e_i\}_{i=1}^n$

Output: Domain prediction: p_{ij}

$\mathbf{g}_i = \eta(\mathbf{h}_i)$ // Get domain features

Construct similarity graph \mathbf{A} by (7)

$\mathbf{S} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$

Obtain converged prediction \mathbf{R}^* by (9)

$p_{ij} = \frac{\exp(r_{ij}^*)}{\sum_{j=1}^S \exp(r_{ij}^*)}$

LADG trains a discriminator that adversarially projects feature representations into a new space, *i.e.*, $\mathbf{g}_i = \eta(\mathbf{h}_i)$. We can use a standard MLP for the discriminator. Alternatively, other architectures are possible, *e.g.*, GatedGCN [9] which allows the discriminator to exploit a sample’s neighbors when projecting the features. This can give it additional information and a greater ability to discriminate, thereby further challenging the featurizer. Both options work well, with the latter providing a slight edge in our experiments. An affinity matrix, $\mathbf{A} = \{a_{ij}\}$, is then constructed based on the transformed features and their K-nearest neighbors:

$$a_{ij} = \begin{cases} \exp\left(\frac{\tau \mathbf{g}_i^\top \mathbf{g}_j}{2\|\mathbf{g}_i\| \|\mathbf{g}_j\|}\right) & \text{if } j \in \mathcal{N}_i \\ 0 & \text{otherwise} \end{cases}, \quad (7)$$

where τ is a scale factor. Because it is intractable to construct an affinity matrix across an entire dataset, we in-

dataset	task	input	# samples	# domains
iWildCam	Cls.	image	203,029	323
Camelyon17	Cls.	image	455,954	5
PovertyMap	Reg.	image	19,669	23x2
FMoW	Cls.	image	523,846	16x5
CivilComments	Cls.	text	448,000	16
Amazon	Cls.	text	539,502	2586

Table 1. Wilds datasets statistics.

stead construct one for each mini-batch. A Laplacian normalized similarity graph can then be obtained as $\mathbf{S} = \mathbf{D}^{-1/2} \mathbf{A} \mathbf{D}^{-1/2}$ [13], where the degree matrix $\mathbf{D} = \{d_{ij}\}$ is a diagonal matrix defined as $d_{ii} = \sum_j a_{ij}$. We use the constructed graph to propagate the domain labels [32, 53, 57], here represented as an $n \times S$ matrix, $\mathbf{E} = \{e_{ij}\}$, where each row is a one-hot vector of the categorical domain label:

$$\mathbf{R}^{t+1} = \alpha \mathbf{S} \mathbf{R}^t + (1 - \alpha) \mathbf{E}, \quad (8)$$

where \mathbf{R}^t are the propagated soft domain labels for the t -th propagation timestep, and $\alpha = 0.8$ is the restart probability. We initialize $\mathbf{R}^0 \leftarrow \mathbf{E}$. The converged propagation results can be calculated as [57]

$$\mathbf{R}^* = (\mathbf{I} - \alpha \mathbf{S})^{-1} \mathbf{E}. \quad (9)$$

We normalize the converged results into domain-prediction pseudo-probabilities:

$$p_{ij} = \frac{\exp(r_{ij}^*)}{\sum_{j=1}^S \exp(r_{ij}^*)}. \quad (10)$$

LADG updates the *discriminator* by minimizing the domain classification loss L_{dom} in (2) with the p_{ij} defined in (10). We conduct the label propagation with samples in the mini-batch.

Analogously, we could follow standard ADG practices and simply train the featurizer to produce domain-invariant representations by maximizing the same L_{dom} in (3). However, we found that this will not only make the training process unstable, but also likely make the domains trivially aligned when the number of training domains $S \geq 3$. To see this, assume four domains as $\{A, B, C, D\}$. We could trivially update the featurizer ϕ to fool the discriminator by aligning A with B , and C with D , without even aligning the A and C domains together. Instead we need a loss that directly formulates our end goal of local neighborhood mixing. To do this, we calculate a prior distribution of domain labels for each minibatch, $q_j = \frac{1}{n} \sum_{i=1}^n e_{ij}$. The featurizer is then trained to match each sample’s propagated domain pseudo-probability with the prior distribution:

$$\mathcal{L}_{prior}(\phi) = -\frac{1}{n} \sum_{i=1}^n \sum_{j=1}^S q_j \log(p_{ij}), \quad (11)$$

thus encouraging each local neighborhood to be well mixed with all domains. We empirically validate the effectiveness of (11) in the experiments.

4.2. Alleviate Space Over Collapse

The featurizer may trivially minimize the feature space to make the domains less distinguishable, *i.e.*, space collapse, which will dramatically degrade the generalization performance for testing domains. To alleviate this problem, we penalize feature space over collapse by encouraging to maintain the compactness for the feature space $R(H)$:

$$\mathcal{L}_{cr}(\phi) = \frac{1}{\rho} \log \cosh(\rho(R(H) - \bar{R}(H))). \quad (12)$$

We elaborate on several points for (12). First, we calculate the coding rate $R(H)$ within each minibatch for tractability. Second, $\bar{R}(H)$ is the moving average of $R(H)$ across batches, and is updated as $\bar{R}(H) \leftarrow \xi \bar{R}(H) + (1 - \xi) R(H)$, where $\xi = 0.99$. $\bar{R}(H)$ can be initialized during pretraining. We adopt a Log-Cosh loss instead of L_1 , which makes the loss effectively zero beyond a certain tolerance value, allowing the network to focus on more egregious violations. A hyperparameter ρ further smoothes the loss curve, which we set to 0.2. Lastly, we use the overall coding rate $R(H)$ instead of the class-wise one $R_y(H)$ since we observe that stable $R(H)$ directly leads to stable $R_y(H)$ and it is nontrivial to calculate the class-wise coding rate $R_y(H)$ for general primary tasks.

4.3. Overall Training

We summarize the overall training objective of LADG as

$$\begin{aligned} \mathcal{L}_{gen} &:= \min_{w, \phi} \mathcal{L}_t(w, \phi) + \lambda \mathcal{L}_{prior}(\phi) + \gamma \mathcal{L}_{cr}(\phi) \\ \mathcal{L}_{dom} &:= \min_{\eta} \mathcal{L}_{adv}(\eta), \end{aligned} \quad (13)$$

where λ and γ are used to balance different terms. As the domain discriminator will be discarded during inference, our method introduces no extra parameters and computational cost compared with ERM for testing. We summarize the training process of LADG in Algorithm 1. Because we calculate affinities across mini-batches, the sampling strategy plays an important role. We randomly select K domains and populate the mini-batch with an equal number of samples from each domain.

5. Experiments

As has been noted, robustness to synthetic distributional shifts do not necessarily translate to robustness to real-world shifts [14, 45]. For this reason, we conduct experiments primarily on the Wilds benchmark [25], which is

Methods	iWildCam F1	Camelyon17 avg acc	PovertyMap wg r	FMoW wg acc	CivilComments wg acc	Amazon 10% acc
ERM	31.0(1.3)	70.3(6.4)	0.45(0.06)	32.3(1.25)	56.0(3.6)	53.8 (0.8)
CORAL	<u>32.8(0.1)</u>	59.5(7.7)	0.44(0.06)	31.7(1.24)	65.6(1.3)	52.9(0.8)
IRM	15.1(4.9)	64.2(8.1)	0.43 (0.07)	30.0(1.37)	66.3 (2.1)	52.4 (0.8)
Group DRO	23.9 (2.1)	68.4 (7.3)	0.39 (0.06)	30.8 (0.81)	<u>70.0(2.0)</u>	53.3 (0.0)
FISH	22.0 (1.8)	74.7 (7.1)	0.30 (0.01)	34.6 (0.18)	73.1* (1.2)	<u>53.3*</u> (0.0)
DANN	27.9(1.1)	65.6(6.7)	<u>0.45(0.06)</u>	30.6(0.15)	66.0(1.0)	<u>53.3(0.0)</u>
CDANN	<u>32.7(1.5)</u>	66.3(6.1)	n/a	33.1(0.11)	69.5(1.5)	<u>53.3(0.0)</u>
LADG	33.1 (0.6)	76.5 (7.7)	0.48 (0.07)	<u>33.5</u> (2.08)	66.9(1.3)	<u>53.3</u> (0.0)

Table 2. Out-of-distribution test results on Wilds. wg means worst group performance. * FISH adopts different backbone for CivilComments and Amazon [42]. Best results are highlighted. DANN and CDANN are implemented by ourselves.

Camelyon17	val acc	test acc*
ERM	84.9(3.1)	70.3(6.4)
CORAL	86.2 (1.4)	59.5(7.7)
IRM	86.2 (1.4)	64.2(8.1)
Group DRO	85.5(2.2)	68.4(7.3)
FISH	83.9(1.2)	74.7(7.1)
DANN	85.6(1.1)	65.6(6.7)
CDANN	86.0(1.3)	66.3(6.1)
LADG	85.4(0.8)	76.5 (7.7)

Table 3. Results on Camelyon17. Validation results are OOD. Test acc is the main summary metric.

iWildCam	ID F1	ID acc	OOD F1*	OOD acc
ERM	47.0 (1.4)	75.7 (0.3)	31.0(1.3)	71.6(2.5)
CORAL	43.5(3.5)	73.7(0.4)	<u>32.8(0.1)</u>	<u>73.3(4.3)</u>
IRM	22.4(7.7)	59.9(8.1)	15.1(4.9)	<u>59.8(3.7)</u>
Group DRO	37.5(1.7)	71.6(2.7)	23.9(2.1)	72.7(2.0)
FISH	40.3(0.6)	<u>73.8(0.1)</u>	22.0(1.8)	64.7(2.6)
DANN	42.3(0.5)	73.5(0.1)	27.9(1.1)	71.8(1.3)
CDANN	<u>46.7(0.7)</u>	75.0(0.2)	32.7(1.5)	72.9(1.5)
LADG	46.4(1.2)	73.7(0.8)	33.1 (0.6)	74.4 (2.7)

Table 4. Results on iWildCam. All metrics are from the test set. OOD F1 is the main summary metric.

composed of real-world datasets representing natural distributional shifts from many different modalities and applications, *e.g.*, biomedical imagery, satellite imagery, and text. Importantly, all Wilds datasets reflect distributional shifts that produce significant performance drops between in-distribution (ID) and out-of-distribution (OOD) data [25]. We follow the experimental settings of FISH [42], a recent and leading approach, and test on the Poverty [49], Camelyon17 [4], FMoW [12], CivilComments [7], iWildCam [5], and Amazon [37] Wilds datasets. *Note CivilComments is*

PovertyMap	val r	test r	val wg r	test wg r*
ERM	0.80(0.04)	0.78(0.04)	0.51(0.06)	0.45(0.06)
CORAL	0.80(0.04)	0.78(0.05)	0.52(0.06)	0.44(0.06)
IRM	0.81(0.03)	0.77(0.05)	0.53(0.05)	0.43(0.07)
Group DRO	0.78(0.05)	0.75(0.07)	0.46(0.04)	0.39(0.06)
FISH	0.82 (0.0)	0.80 (0.02)	0.47(0.01)	0.30(0.01)
DANN	0.80(0.04)	<u>0.79(0.03)</u>	0.52(0.05)	<u>0.45(0.06)</u>
CDANN	n/a	n/a	n/a	n/a
LADG	0.80(0.03)	<u>0.79</u> (0.04)	0.51(0.07)	0.48 (0.07)

Table 5. Results on PovertyMap. CDANN is not applicable for regression. r denotes Pearson correlation coefficient. All validation metrics are OOD. wg means worst-group performance and test wg r coefficient is the main summary metric.

FMoW	val acc	test acc	val wr acc	test wr acc*
ERM	59.5 (0.37)	53.0 (0.55)	48.9(0.62)	32.3(1.25)
CORAL	56.9(0.25)	50.5(0.36)	47.1(0.43)	31.7(1.24)
IRM	57.4(0.37)	50.8(0.13)	47.5(1.57)	30.0(1.37)
Group DRO	<u>58.8</u> (0.19)	<u>52.1</u> (0.5)	46.5(0.25)	30.8(0.81)
FISH	57.8(0.15)	51.8(0.32)	49.5(2.34)	34.6 (0.18)
DANN	57.1(0.73)	50.2(0.36)	<u>50.2</u> (1.01)	30.6(0.15)
CDANN	56.8(0.11)	51.0(0.36)	50.0(0.31)	33.1(0.11)
LADG	57.2(0.15)	51.2(0.24)	51.0 (0.47)	<u>33.5</u> (2.08)

Table 6. Results on FMoW. All validation results are OOD. wr means worst-region performance and test wr acc is the main summary metric.

designed for sub-population shifts (training and testing domains overlap), so it is not a DG dataset, but we include it for completeness. We summarize the statistics of the used datasets in Table 1, and we refer the readers to [25] for details on these datasets. We report the Wilds leaderboard metrics, which always include OOD test, but also may include ID test or OOD validation metrics. We also provide the results on domainbed including the PACS dataset in supplementary material [19].

Civ.Comm.	val avg acc	val wg acc	test avg acc	test wg acc*
ERM	92.3 (0.2)	50.5(1.9)	92.2 (0.1)	56.0(3.6)
CORAL	88.9(0.6)	64.7(1.4)	88.7(0.5)	65.6(1.3)
IRM	89.0(0.7)	65.9(2.8)	88.8(0.7)	66.3(2.1)
Group DRO	90.1(0.4)	67.7(1.8)	89.9(0.5)	70.0(2.0)
FISH*	89.7*(0.3)	72.0* (1.0)	89.5*(0.2)	73.1* (1.2)
DANN	89.1(0.5)	68.8(1.1)	88.9(0.5)	66.0(1.0)
CDANN	89.6(0.2)	69.4(0.5)	89.3(0.2)	69.5(1.5)
LADG	88.8(0.6)	66.8(2.1)	88.6(0.6)	66.9(1.3)

Table 7. Results on CivilComments. * FISH uses a different backbone [42]. For CivilComments, the training and test domains overlap and there is only sub-population shift. wg means worst-group performance and test wg acc is the main summary metric.

Amazon	val acc	test acc	val 10% acc	test 10% acc
ERM	72.7 (0.1)	71.9 (0.1)	55.2 (0.7)	53.8 (0.8)
CORAL	72.0(0.3)	71.1(0.3)	54.7(0.1)	53.3(0.1)
IRM	71.5(0.3)	70.5(0.3)	54.2(0.8)	52.4(0.8)
Group DRO	70.7(0.6)	70.0(0.6)	54.7(0.0)	53.3(0.0)
FISH*	72.5*(0.0)	71.7*(0.1)	54.7*(0.0)	53.3*(0.0)
DANN	72.1(0.0)	71.4(0.1)	54.7(0.0)	53.3(0.0)
CDANN	72.3(0.2)	71.2(0.2)	54.7(0.0)	53.3(0.0)
LADG	71.9(0.1)	71.1(0.2)	54.7(0.0)	53.3(0.0)

Table 8. Results on Amazon. * FISH uses a different backbone [42]. Validation metrics are OOD and test 10% acc is the main summary metric.

Experimental Settings: We follow the settings and evaluation metrics of Wilds benchmark [25]. We compare our methods with ERM [25], Coral [44], Group DRO [40], IRM [3], and Fish [42], which we retrieve from the Wilds benchmark leaderboard [25]. We implement DANN [16] and CDANN [33] ourselves. We adopt the same featurizer ϕ , predictor w , and model selection strategies used in the Wilds benchmark. We provide all hyperparameters in the supplementary material. Following the Wilds specifications [25], for iWildCam, FMoW, and Amazon, we average the results over three different runs, and for Camelyon17 (CivilComment), we run the experiments ten (five) times. For PovertyMap, we average the results over five different folds.

Results: We summarize the results in Table 2 for the main summary metrics, and detailed results for each dataset are provided in Tables 3–8. From Table 2, LADG achieves state-of-the-art results on 3 datasets and the second-best result on FMoW. Notably, LADG performance is much more stable, with the relative (and rank-order) performance of some methods, *e.g.* FISH, varying considerably from dataset to dataset. Despite not being designed to address sub-population shifts, LADG still exhibits good performance on CivilComments, outperforming ERM. All methods perform comparably for Amazon (likely because its huge numbers of domains make it difficult to beat ERM),

	Camelyon17	PovertyMap	
	avg acc	r	wg r
DANN	65.6(6.7)	0.79(0.03)	0.45(0.06)
DANN w/ \mathcal{L}_{cr}	73.8(7.9)	0.78(0.04)	0.47(0.06)
LADG w/o \mathcal{L}_{cr}	68.8(5.4)	0.79(0.05)	0.43(0.07)
LADG w/o \mathcal{L}_{prior}	72.8(7.3)	0.80 (0.07)	0.46(0.13)
LADG w/o GNN	75.7(6.3)	0.79(0.05)	0.49(0.03)
LADG	76.5(7.7)	0.79(0.04)	0.48(0.07)

Table 9. Ablation studies on Wilds.

but LADG matches the best DG methods’ performances. These results demonstrate that for true DG dataset, LADG (1) exhibits leading performance; and (2) its performance is consistently high.

Ablation Study: In this section, we conduct ablation studies on Camelyon17 and Poverty Map to validate the different components of the proposed methods. The experiments follow the exact same settings as Wilds benchmark. The results are shown in Table 9. DANN w/ \mathcal{L}_{cr} equips DANN with \mathcal{L}_{cr} loss. LADG w/o \mathcal{L}_{prior} replaces prior matching loss with \mathcal{L}_{adv} and LADG w/o GNN replace GNN with MLP. The experiments show that the coding ratio loss \mathcal{L}_{cr} is important for our method, and LADG outperforms its direct counterpart DANN either with or without \mathcal{L}_{cr} loss.

6. Limitations and Future Work

Our analysis and experiments are conducted under the benchmark settings with a linear predictor w [19, 25], and we do not experimentally verify the performance of our method under the settings with non-linear predictors. We expect incompletely aligned distributions with unmixed local neighborhoods to also be a problem when using non-linear predictors and that LADG’s innovations would equally apply. We leave this for future work. Another limitation is that our method treats each domain as equally important. In practice, we may over-sample the domains that are deviated from others to improve the worst group performance for better generalization ability [40]. Finally, improving performance when domains numbers are huge is another important avenue for future work.

7. Conclusions

We proposed LADG, which addresses two limitations of existing ADG approaches: (1) incomplete distributional alignment with unmixed neighborhoods and (2) feature space over collapse. LADG incorporates an adversarial localized classifier to align domains within every local region. Moreover, LADG penalizes space over collapse by encouraging to maintain the compactness of the feature space via a differentiable coding-rate formulation. Experiments on leading benchmarks validate LADG’s effectiveness and consistent performance across datasets.

References

- [1] Faruk Ahmed, Yoshua Bengio, Harm van Seijen, and Aaron Courville. Systematic generalisation with group invariant predictions. In *International Conference on Learning Representations*, 2020. [2](#)
- [2] Isabela Albuquerque, João Monteiro, Mohammad Darvishi, Tiago H Falk, and Ioannis Mitliagkas. Adversarial target-invariant representation learning for domain generalization. 2020. [2](#)
- [3] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. [1](#), [2](#), [8](#)
- [4] Peter Bandi, Oscar Geessink, Quirine Manson, Marcory Van Dijk, Maschenka Balkenhol, Meyke Hermesen, Babak Ehteshami Bejnordi, Byungjae Lee, Kyunghyun Paeng, Aoxiao Zhong, et al. From detection of individual metastases to classification of lymph node status at the patient level: the camelyon17 challenge. *IEEE transactions on medical imaging*, 38(2):550–560, 2018. [7](#)
- [5] Sara Beery, Arushi Agarwal, Elijah Cole, and Vighnesh Birodkar. The iwildcam 2021 competition dataset. *arXiv preprint arXiv:2105.03494*, 2021. [7](#)
- [6] Bernd Bischl, Julia Schiffner, and Claus Weihs. Benchmarking local classification methods. *Computational Statistics*, 28(6):2599–2619, 2013. [2](#), [4](#), [5](#)
- [7] Daniel Borkan, Lucas Dixon, Jeffrey Sorensen, Nithum Thain, and Lucy Vasserman. Nuanced metrics for measuring unintended bias with real data for text classification. In *Companion proceedings of the 2019 world wide web conference*, pages 491–500, 2019. [7](#)
- [8] Leo Breiman. Random forests. *Machine learning*, 45(1):5–32, 2001. [5](#)
- [9] Xavier Bresson and Thomas Laurent. Residual gated graph convnets. *CoRR*, abs/1711.07553, 2017. [5](#)
- [10] Tianle Cai, Ruiqi Gao, Jason D Lee, and Qi Lei. A theory of label propagation for subpopulation shift. *arXiv preprint arXiv:2102.11203*, 2021. [5](#)
- [11] Shiyu Chang, Yang Zhang, Mo Yu, and Tommi Jaakkola. Invariant rationalization. In *International Conference on Machine Learning*, pages 1448–1458. PMLR, 2020. [2](#)
- [12] Gordon Christie, Neil Fendley, James Wilson, and Ryan Mukherjee. Functional map of the world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6172–6180, 2018. [7](#)
- [13] Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997. [6](#)
- [14] Josip Djolonga, Jessica Yung, Michael Tschannen, Rob Romijnders, Lucas Beyer, Alexander Kolesnikov, Joan Puigcerver, Matthias Minderer, Alexander D’Amour, Dan Moldovan, Sylvain Gelly, Neil Houlsby, Xiaohua Zhai, and Mario Lucic. On Robustness and Transferability of Convolutional Neural Networks. In *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 16453–16463, Nashville, TN, USA, June 2021. IEEE. [6](#)
- [15] Qi Dou, Daniel Coelho de Castro, Konstantinos Kamnitsas, and Ben Glocker. Domain generalization via model-agnostic learning of semantic features. *Advances in Neural Information Processing Systems*, 32:6450–6461, 2019. [2](#)
- [16] 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](#), [2](#), [3](#), [4](#), [8](#)
- [17] Mingming Gong, Kun Zhang, Tongliang Liu, Dacheng Tao, Clark Glymour, and Bernhard Schölkopf. Domain adaptation with conditional transferable components. In *International conference on machine learning*, pages 2839–2848. PMLR, 2016. [2](#)
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014. [1](#), [3](#)
- [19] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020. [1](#), [2](#), [3](#), [7](#), [8](#)
- [20] David J Hand and Veronica Vinciotti. Local versus global models for classification problems: fitting models where it matters. *The American Statistician*, 57(2):124–131, 2003. [5](#)
- [21] Zeyi Huang, Haohan Wang, Eric P Xing, and Dong Huang. Self-challenging improves cross-domain generalization. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pages 124–140. Springer, 2020. [2](#)
- [22] Ahmet Iscen, Giorgos Tolias, Yannis Avrithis, and Ondrej Chum. Label propagation for deep semi-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5070–5079, 2019. [5](#)
- [23] Byungju Kim, Hyunwoo Kim, Kyungsu Kim, Sungjin Kim, and Junmo Kim. Learning not to learn: Training deep neural networks with biased data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9012–9020, 2019. [2](#)
- [24] 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. [2](#)
- [25] Pang Wei Koh, Shiori Sagawa, Sang Michael Xie, Marvin Zhang, Akshay Balsubramani, Weihua Hu, Michihiro Yasunaga, Richard Lanus Phillips, Irena Gao, Tony Lee, et al. Wilds: A benchmark of in-the-wild distribution shifts. In *International Conference on Machine Learning*, pages 5637–5664. PMLR, 2021. [1](#), [6](#), [7](#), [8](#)
- [26] 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. [2](#)
- [27] 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. [4](#)

- [28] 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. 1
- [29] 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, 2, 4
- [30] 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. 2, 4
- [31] Bin Liu, Zhirong Wu, Han Hu, and Stephen Lin. Deep metric transfer for label propagation with limited annotated data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, pages 0–0, 2019. 5
- [32] Yanbin Liu, Juho Lee, Minseop Park, Saehoon Kim, Eunho Yang, Sung Ju Hwang, and Yi Yang. Learning to propagate labels: Transductive propagation network for few-shot learning. *arXiv preprint arXiv:1805.10002*, 2018. 2, 5, 6
- [33] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *arXiv preprint arXiv:1705.10667*, 2017. 1, 2, 3, 4, 8
- [34] Yi Ma, Harm Derksen, Wei Hong, and John Wright. Segmentation of multivariate mixed data via lossy data coding and compression. *IEEE transactions on pattern analysis and machine intelligence*, 29(9):1546–1562, 2007. 4
- [35] 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. 2
- [36] 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. 2
- [37] Jianmo Ni, Jiacheng Li, and Julian McAuley. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 188–197, 2019. 7
- [38] Giambattista Parascandolo, Alexander Neitz, Antonio Orvieto, Luigi Gresele, and Bernhard Schölkopf. Learning explanations that are hard to vary. *arXiv preprint arXiv:2009.00329*, 2020. 2
- [39] Alexandre Rame, Corentin Dancette, and Matthieu Cord. Fishr: Invariant gradient variances for out-of-distribution generalization. *arXiv preprint arXiv:2109.02934*, 2021. 1, 2
- [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. 1, 2, 8
- [41] Shiv Shankar, Vihari Piratla, Soumen Chakrabarti, Siddhartha Chaudhuri, Preethi Jyothi, and Sunita Sarawagi. Generalizing across domains via cross-gradient training. *arXiv preprint arXiv:1804.10745*, 2018. 2
- [42] Yuge Shi, Jeffrey Seely, Philip HS Torr, N Siddharth, Awni Hannun, Nicolas Usunier, and Gabriel Synnaeve. Gradient matching for domain generalization. *arXiv preprint arXiv:2104.09937*, 2021. 1, 2, 7, 8
- [43] Anthony Sicilia, Xingchen Zhao, and Seong Jae Hwang. Domain adversarial neural networks for domain generalization: When it works and how to improve. *arXiv preprint arXiv:2102.03924*, 2021. 1, 2
- [44] Baochen Sun, Jiashi Feng, and Kate Saenko. Correlation alignment for unsupervised domain adaptation. In *Domain Adaptation in Computer Vision Applications*, pages 153–171. Springer, 2017. 1, 2, 8
- [45] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring Robustness to Natural Distribution Shifts in Image Classification. *arXiv:2007.00644 [cs, stat]*, Sept. 2020. arXiv: 2007.00644. 6
- [46] Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008. 4
- [47] 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. 2
- [48] Yuan Wu, Diana Inkpen, and Ahmed El-Roby. Dual mixup regularized learning for adversarial domain adaptation. In *European Conference on Computer Vision*, pages 540–555. Springer, 2020. 2
- [49] Christopher Yeh, Anthony Perez, Anne Driscoll, George Azhari, Zhongyi Tang, David Lobell, Stefano Ermon, and Marshall Burke. Using publicly available satellite imagery and deep learning to understand economic well-being in africa. *Nature communications*, 11(1):1–11, 2020. 7
- [50] Yaodong Yu, Kwan Ho Ryan Chan, Chong You, Chaobing Song, and Yi Ma. Learning diverse and discriminative representations via the principle of maximal coding rate reduction. *Advances in Neural Information Processing Systems*, 33, 2020. 2, 4, 5
- [51] Brian Hu Zhang, Blake Lemoine, and Margaret Mitchell. Mitigating unwanted biases with adversarial learning. In *Proceedings of the 2018 AAAI/ACM Conference on AI, Ethics, and Society*, pages 335–340, 2018. 2
- [52] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 1, 2
- [53] Kai Zhang, Yaokang Zhu, Jun Wang, and Jie Zhang. Adaptive structural fingerprints for graph attention networks. In *International Conference on Learning Representations*, 2019. 2, 5, 6
- [54] Marvin Zhang, Henrik Marklund, Nikita Dhawan, Abhishek Gupta, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group distribution shift. *arXiv preprint arXiv:2007.02931*, 2020. 1, 2

- [55] Yabin Zhang, Bin Deng, Kui Jia, and Lei Zhang. Label propagation with augmented anchors: A simple semi-supervised learning baseline for unsupervised domain adaptation. In *European Conference on Computer Vision*, pages 781–797. Springer, 2020. [5](#)
- [56] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *International Conference on Machine Learning*, pages 7404–7413. PMLR, 2019. [2](#)
- [57] Dengyong Zhou, Olivier Bousquet, Thomas Lal, Jason Weston, and Bernhard Schölkopf. Learning with local and global consistency. *Advances in neural information processing systems*, 16, 2003. [6](#)
- [58] Kaiyang Zhou, Yongxin Yang, Timothy Hospedales, and Tao Xiang. Deep domain-adversarial image generation for domain generalisation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 13025–13032, 2020. [2](#)
- [59] Kaiyang Zhou, Yongxin Yang, Yu Qiao, and Tao Xiang. Mixstyle neural networks for domain generalization and adaptation. *arXiv preprint arXiv:2107.02053*, 2021. [1](#), [2](#)
- [60] Wei Zhu, Haitian Zheng, Haofu Liao, Weijian Li, and Jiebo Luo. Learning bias-invariant representation by cross-sample mutual information minimization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15002–15012, 2021. [5](#)