

Adaptive Methods for Real-World Domain Generalization

Abhimanyu Dubey*
MIT
dubeya@mit.edu

Vignesh Ramanathan
Facebook AI
vigneshr@fb.com

Alex Pentland
MIT
pentland@mit.edu

Dhruv Mahajan
Facebook AI
dhruvm@fb.com

Abstract

Invariant approaches have been remarkably successful in tackling the problem of domain generalization, where the objective is to perform inference on data distributions different from those used in training. In our work, we investigate whether it is possible to leverage domain information from the unseen test samples themselves. We propose a domain-adaptive approach consisting of two steps: a) we first learn a discriminative domain embedding from unsupervised training examples, and b) use this domain embedding as supplementary information to build a domain-adaptive model, that takes both the input as well as its domain into account while making predictions. For unseen domains, our method simply uses few unlabelled test examples to construct the domain embedding. This enables adaptive classification on any unseen domain. Our approach achieves state-of-the-art performance on various domain generalization benchmarks. In addition, we introduce the first real-world, large-scale domain generalization benchmark, Geo-YFCC, containing 1.1M samples over 40 training, 7 validation and 15 test domains, orders of magnitude larger than prior work. We show that the existing approaches either do not scale to this dataset or underperform compared to the simple baseline of training a model on the union of data from all training domains. In contrast, our approach achieves a significant 1% improvement.

1. Introduction

Domain generalization refers to the problem of learning a classifier from a heterogeneous collection of distinct training domains that can generalize to new unseen test domains [7]. Among the various formalizations proposed for the problem [19, 6], the most effective is domain-invariant learning [4, 15]. It learns feature representations invariant to the underlying domain, providing a universal classifier that can generalize to new domains [26, 26, 33, 25, 16].

It can be demonstrated that domain-invariant classifiers




Figure 1: A visual for *adaptive* domain generalization. Here we denote by \hat{f}_i the optimal classifier in a class \mathcal{F} for domain i (out of 8). For test domain D_T , the optimal *universal* classifier \hat{f}_* may be far from the optimal classifier for D_T .

minimize the *average risk* across domains [33]. However, this may not guarantee good performance for any *specific* test domain, particularly if the distribution of domains has high variance, as visualized heuristically in Figure 1. In this case, the optimal classifier for a target domain can lie far from the optimal *universal* classifier. In our work, we propose an *adaptive* classifier that can be adapted to any new domain using very few unlabelled samples without any further training. Unlike invariant approaches, this requires a few *unsupervised* samples from any domain while *testing*¹. However, this requirement is trivial, since this set is available by definition in all practical settings.

Our approach consists of two steps. We first embed each domain into a vector space using very few unlabelled samples from the domain. Next, we leverage these *domain embeddings* as supplementary signals to learn a *domain-adaptive* classifier. During testing, the classifier is supplied the corresponding embedding obtained from test samples. Our contributions can be summarized as follows.

1. We adapt low-shot *prototypical learning* [39] to construct *domain embeddings* from unlabelled samples of each domain. We also provide an algorithm to use these embeddings as supplementary signals to train *adaptive* classifiers.
2. We justify our design choices with novel theoretic

*Work done while visiting Facebook AI.

¹This does not correspond to unsupervised domain adaptation, where unsupervised samples are assumed to be present during *training*.

cal results based on the framework of kernel mean embeddings [18]. Furthermore, we leverage the theory of risk minimization under product kernel spaces [7] to derive generalization bounds on the average risk for adaptive classifiers.

3. We introduce the first large-scale, real-world domain generalization benchmark, dubbed **Geo-YFCC**, which contains 40 training, 7 validation and 15 test domains, and an overall 1.1M samples. Geo-YFCC is constructed from the popular YFCC100M [42] dataset partitioned by geographical tags, and exhibits several characteristics of *domain shift* [23], *label shift* [3], and long-tailed class distributions [49, 45, 9] common in real-world classification. Because of its scale and diversity, this benchmark is substantially more challenging than the incumbent benchmarks.

4. On existing benchmarks and notably, two large-scale benchmarks, we demonstrate the strength of fine-tuning (ERM) with *adaptive* classification over existing approaches. In addition to the effectiveness of adaptive classification, this suggests, along with the claims of Gulrajani and Lopez-Paz [19] that rigorous model selection and large benchmarks are essential in understanding domain generalization, where naive ERM can also be a powerful baseline.

2. Related Work

Our paper is inspired by several areas of research, and we enumerate the connections with each area sequentially.

Domain Generalization. The problem of domain generalization was first studied as a variant of multi-task learning in Blanchard *et al.* [7]. Many *domain-invariant* approaches [41, 25, 26], have been developed using the popular domain adaptation algorithm introduced in [15], where the authors learned invariant representations via adversarial training. This has been followed by alternative formulations for invariant learning such as MMD minimization [26], correlation alignment [41], and class-conditional adversarial learning [25]. Other approaches include meta learning [24], invariant risk minimization [2], distributionally robust optimization [38], mixup [46, 47, 44], and causal matching [32]. *Adversarial training* with improvements [15, 35] has also been used to learn invariant representations. Complementary to these approaches, we focus instead on learning *adaptive* classifiers that are specialized to each target domain². Our approach also does not use any unsupervised data from unseen domains during training, as is done in the problem of unsupervised domain adaptation [17, 4].

Domain-Adaptive Learning. The idea of using *kernel mean embeddings* (KME) for adaptive domain generalization was proposed in the work of Blanchard *et al.* [7]. Kernel mean embeddings have also been used for personal-

ized learning in both multi-task [11] and multi-agent learning [13] bandit problems. A rigorous treatment of *domain-adaptive* generalization in the context of KME approaches is provided in Deshmukh *et al.* [12]. Our approach complements this theoretical line of work, by providing an efficient algorithm for classification, that requires only a few unsupervised samples for competitive performance. We also extend the results of [7] to a larger class of functions.

Large-Scale Learning. We propose domain generalization benchmarks based on both the ImageNet LSVRC12 [10] and YFCC100M [42] datasets, which have been instrumental in accelerating research for image classification. We consider the challenges encountered specifically for large-scale computer vision, which is plagued by issues such as long-tailed data distributions [27], and large training times [1]. It is important to note that in large-scale settings, it is difficult to perform extensive hyper-parameter tuning and optimal model selection owing to large training times.

3. Approach

We assume that each *domain* $D \in \mathcal{D}$ is a probability distribution over $\mathcal{X} \times \mathcal{Y}$, i.e., the product of the input space \mathcal{X} and output space \mathcal{Y} , and there exists a *mother* distribution \mathfrak{P} which is a measure over the space of domains \mathcal{D} . A *training domain* $\hat{D}(n)$ is obtained by first sampling a domain D from \mathfrak{P} , and then sampling n points (X and Y) from $\mathcal{X} \times \mathcal{Y}$ following D . A training set can then be constructed by taking the union of N such domains $(\hat{D}_i(n))_{i=1}^N$. Correspondingly, any test domain $\hat{D}_T(n_T)$ is also constructed by first drawing a sample $D_T \sim \mathcal{D}$, then drawing n_T samples from D_T , and discarding the labels.

Consider a family of functions \mathcal{F} . For each $D \in \mathcal{D}$, the optimal classifier f_D in \mathcal{F} (for some loss function ℓ) can be given by $f_D = \arg \min_{f \in \mathcal{F}} \mathbb{E}_{(\mathbf{x}, y) \sim D} [\ell(f(\mathbf{x}), y)]$. We denote the optimal empirical risk minimization (ERM) classifier for each domain \hat{D}_i in the training set domain $D_i(n)$ as $\hat{f}_i = \arg \min_{f \in \mathcal{F}} \mathbb{E}_{(\mathbf{x}, y) \sim D(n)} [\ell(f(\mathbf{x}), y)]$. The *universal* expected risk minimizer (i.e., classifier minimizing overall risk over $D \sim \mathfrak{P}$) can then be given by f_* such that,

$$f_* = \arg \min_{f \in \mathcal{F}} \mathbb{E}_{D \sim \mathfrak{P}} \mathbb{E}_{\mathbf{x}, y \sim D} [\ell(f(\mathbf{x}), y)]. \quad (1)$$

Typically, the optimal classifier for D_T within \mathcal{F} may be far from the (non-adaptive) universal classifier \hat{f}_* , a distance which we call the “adaptivity gap”. A visual heuristic representation of the above intuition is provided in Figure 1. Ideally, we would like the classifier to be chosen closer to \hat{f}_T such that it obtains lower error. Note that the *adaptivity gap* may be large even when the training set contains samples close to \hat{f}_T (in the figure, this refers to functions \hat{f}_7 and \hat{f}_8), since the *universal* classifier minimizes overall risk.

Our approach reduces this discrepancy by learning an *adaptive* classifier, i.e., a function $F : \mathcal{X} \times \mathcal{D} \rightarrow \mathcal{Y}$ that

²Note that our algorithm is not opposed to invariant learning, we merely focus on learning *distribution-adaptive* classifiers, that can be combined with invariant approaches with ease, see Section 5.1.

takes *both* the point \mathbf{x} and the domain D of \mathbf{x} into consideration. Our goal is to make sure that for any test domain D_T , the resulting classifier $f_T = F(\cdot, D_T)$ from F lies close to \hat{f}_T . We assume that a small set of unsupervised data is available from D_T to characterize $F(\cdot, D_T)$ from F^3 . We follow a two step approach:

(A) **Computing domain embeddings.** We first learn a function $\mu : D \rightarrow \mathbb{R}^{d_D}$, that maps any domain $D \in \mathcal{D}$ (including empirical domains) to a finite vector (domain embedding). Ideally, we want this embedding to be learned from a few samples belonging to the domain. To achieve this, we leverage kernel mean embeddings [34], a formalization that can represent probability distributions as the average of some feature vectors Φ_D in a suitable Hilbert space. μ is obtained as the feature $\Phi_D(\mathbf{x})$ averaged over all $\mathbf{x} \in D$.

(B) **ERM using augmented inputs.** With a suitable domain embedding μ , we learn a neural network F directly over the *augmented input* $(\mathbf{x}, \mu(\hat{D}))$, where \mathbf{x} is a point from the domain \hat{D} , using regular ERM (i.e., minimizing cross-entropy loss from predictions, see Algorithm 2).

Inference. For an unseen domain \hat{D}_T , we first compute the domain embedding $\mu(\hat{D}_T)$ using unsupervised samples (or the test set when none are present) from \hat{D}_T and the learned function Φ from (A). We then provide the computed embedding as an additional input to the model F from (B) to get the *adaptive* model $F(\cdot, \mu(\hat{D}_T))$.

3.1. Prototypical Domain Embeddings

Kernel Mean Embeddings. An elegant approach to embed distributions into vector spaces is the nonparametric method of *kernel mean embeddings* (KME) [34]. The fundamental idea behind KMEs is to consider each probability distribution $D \in \mathcal{D}$ as a member of a reproducing kernel Hilbert space (RKHS) $\mathcal{H}(\mathcal{X})$ with some kernel k and feature $\Phi_D : \mathbb{R}^d \rightarrow \mathbb{R}^{d_D}$. Then, the KME of any distribution $D \in \mathcal{D}$ can be given as $\mu(D) = \int_{\mathcal{X}} \Phi_D(\mathbf{x}) dD(\mathbf{x})$.

The strength of this approach is that, if a suitable Φ_D is chosen, then it is possible to learn a function over the space of distributions \mathcal{D} via the average feature embedding of samples drawn from D . Hence, we propose to learn both a domain embedding μ using features Φ_D and the target function $F(\cdot, \mu(\cdot))$, such that when any domain D is presented, we can provide an *adapted* function $F(\cdot, \mu(D))$. We learn a neural network $\Phi_D : \mathcal{X} \rightarrow \mathbb{R}^{d_D}$, parameterized by weights θ , and then compute the n -sample *empirical* KME μ of each training domain \hat{D} as:

$$\mu(\hat{D}) = \frac{1}{n} \sum_{\mathbf{x} \in \hat{D}} \Phi_D(\mathbf{x}; \theta). \quad (2)$$

³This assumption is trivially satisfied by the test set itself. Note that we do not access D_T during training and hence we do not perform any *transductive learning*, unlike unsupervised domain adaptation.

Low-Shot Domain Prototypes. To be useful in classification, we require μ to satisfy two central criteria:

(A) **Expressivity.** μ , and consequently, Φ_D , must be expressive, i.e., for two domains D and D' , $\|\mu(D) - \mu(D')\|_2$ must be large if they are very distinct (i.e., domain shift), and small if they are similar. For example, if Φ_D is constant, the KME will provide no additional information about D , and we can expect the *expected risk* itself to be large in this case, equivalent to non-adaptive classification.

(B) **Consistency.** Φ must have little dependence on both the choice, and the number of samples used to construct it. This is desirable as we only have access to the empirical distribution \hat{D} corresponding to any domain D . Consistency can be achieved, if we learn Φ_D such that for any sample \mathbf{x} from domain D , $\Phi_D(\mathbf{x})$ is strongly clustered around $\mu(D)$.

We show that *prototypical networks* [39] originally proposed for the task of few-shot learning, can be adapted to construct domain embeddings which satisfy both criteria. We train the network based on the algorithm from Snell *et al.* [39], but using domain identities as the labels instead of class labels. The network is trained such that embeddings of all points from a domain are tightly clustered and far away from embeddings corresponding to points of other domains. In Section 3.3, we demonstrate consistency by deriving a concentration bound for the clustering induced by the network. Expressivity is guaranteed by the loss function used to train the network, which forces domain embeddings to be discriminative, as we show next.

Our prototypical network (embedding function Φ_D) is trained with SGD [8]. At each iteration, we first sample a subset of N_t domains from the training domains. Next, we sample two sets of points from each domain. The first set is used to obtain the embedding for the domain, by running the points through the network Φ_D and averaging (Eq. 2). This results in an approximate embedding $\mu(\hat{D})$ for each domain D . Every point \mathbf{x} in the second set is given a probability of belonging to a domain D :

$$p_{\theta}(\mathbf{x} \in D) = \frac{\exp\left(-\|\mu(\hat{D}) - \Phi_D(\mathbf{x})\|_2^2\right)}{\sum_{i=1}^{N_t} \exp\left(-\|\mu(\hat{D}_i) - \Phi_D(\mathbf{x})\|_2^2\right)}. \quad (3)$$

Learning proceeds by minimizing the negative log-likelihood $J(\theta) = -\log p_{\theta}(\mathbf{x} \in D_{\mathbf{x}})$ for the correct domain $D_{\mathbf{x}}$ from which \mathbf{x} was sampled. This ensures that probability is high only if a point is closer to its own domain embedding and farther away from other domains, making the domain embedding *expressive*. Once training is complete, we obtain the final model parameters θ^* . For each domain $\hat{D}_i(n)$, the *domain prototype* is constructed as:

$$\mu(\hat{D}_i(n)) := \frac{1}{n} \sum_{j=1}^n \Phi_D(\mathbf{x}; \theta^*). \quad (4)$$

Algorithm 1 Prototypical Training Pseudocode

PROTOTYPE TRAINING

Input. N domains $(D_i)_{i=1}^N$ with n samples each. N_t : #domains sampled each batch, N_s : #support examples, N_q : #query examples per batch.
Output. Embedding neural network $\Phi_{\mathcal{D}}: \mathcal{X} \rightarrow \mathbb{R}^{d_{\mathcal{D}}}$.

Initialize. weights θ of $\Phi_{\mathcal{D}}$ randomly.
for Round $t = 1$ to T **do**
 $D_t \leftarrow \text{RANDOMLY SAMPLE}(N, N_t)$. // sample N_t domains
 for Domain d in D_t **do**
 $S_d \leftarrow \text{RANDOMLY SAMPLE}(D_d, N_s)$. // sample support points
 $S_q \leftarrow \text{RANDOMLY SAMPLE}(D_d, N_q)$. // sample query points
 $\hat{\mu}_d \leftarrow \frac{1}{N_s} \sum_{\mathbf{x} \in S_d} \Phi_{\mathcal{D}}(\mathbf{x}; \theta)$. // compute prototype
 end for
 $J_{\theta}(t) \leftarrow 0$.
 for Domain d in D_t **do**
 for $\mathbf{x} \in S_q$ **do**
 Update Loss $J_{\theta}(t)$ based on Equation 3.
 end for
 end for
 $\theta \leftarrow \text{SGD STEP}(J(t), \theta)$. // gradient descent step
end for

PROTOTYPE COMPUTATION ($\Phi_{\mathcal{D}}, S$)

Input. Trained prototype network $\Phi_{\mathcal{D}}$ with weights θ^* .
Set S of points to create prototype.
Output. Prototype $\mu(S)$.
return $\mu(S) \leftarrow \frac{1}{|S|} \sum_{\mathbf{x} \in S} \Phi_{\mathcal{D}}(\mathbf{x}, \theta^*)$.

Prototype training and construction is completely unsupervised as long as we are provided the inputs partitioned by domains. We repeat *only* Equation 4 for each test domain as well (no retraining). See Algorithm 1 for details.

3.2. ERM on Augmented Inputs

Once we have the domain prototype for each training domain from the previous step, we create the *augmented* training set by appending the corresponding prototype to each input. We obtain the training set $\tilde{S}_{\text{tm}} = \cup_{i=1}^N \tilde{D}_i$, where,

$$\tilde{D}_i = \left(\mathbf{x}_{ij}, \mu(\hat{D}_i), y_{ij} \right)_{j=1}^n. \quad (5)$$

The second step is to then train the *domain-adaptive* classifier F over \tilde{S}_{tm} using regular ERM training. F consists of two neural networks (Algorithm 2). The first network F_{ft} with weights ω_{ft} takes \mathbf{x} (image) as input and outputs a feature vector $\Phi_{\mathcal{X}}(\mathbf{x})$. We then concatenate the image feature $\Phi_{\mathcal{X}}$ with the domain embedding $\mu(D)$ and pass through the second network F_{mlp} which predicts the class label y . We refer to this approach as DA-ERM. In Section 5.1, we present a few variants that combine the complementary nature of adaptive and invariant approaches.

When testing, we first compute the domain prototypes for a test domain \hat{D}_T (Equation 4) using a small number of unlabeled examples⁴, and then produce predictions via $F(\cdot, \hat{D}_T)$. Alternatively, it is possible to decouple the prototype construction and inference procedures in the test phase

⁴We can also use a small sample of the total test set to create $\mu(\hat{D}_T)$.

Algorithm 2 Adaptive Training Pseudocode

ADAPTIVE TRAINING

Input. Prototype network $\Phi_{\mathcal{D}}$ with weights θ^* , Domains $(D_i)_{i=1}^N$.
points N_p to sample from each domain to create prototype.
Output. F_{ft} with weights ω_{ft} , F_{mlp} with weights ω_{mlp} .

// compute prototypes for training domains
for Domain d from 1 to N **do**
 $S_d \leftarrow \text{RANDOMLY SAMPLE}(D_d, N_p)$.
 $\mu(\hat{D}_d) \leftarrow \text{PROTOTYPE COMPUTATION}(\Phi_{\mathcal{D}}, S_d)$.
end for
Create augmented dataset $(\tilde{D}_i)_{i=1}^n$ (Equation 5).
for SGD round $t = 1$ to T **do**
 Sample batch (\mathbf{x}, μ, y) from augmented dataset.
 $\Phi_{\mathcal{X}}(\mathbf{x}) \leftarrow F_{\text{ft}}(\mathbf{x}; \omega_{\text{ft}})$. // compute image features
 $\Phi(\mathbf{x}, \mu) \leftarrow \text{CONCAT}(\Phi_{\mathcal{X}}(\mathbf{x}), \mu)$. // concatenate features
 $\hat{y} \leftarrow F_{\text{mlp}}(\Phi(\mathbf{x}, \mu); \omega_{\text{mlp}})$. // compute predictions
 $J_{\omega}(t) \leftarrow \text{CROSSENTROPY}(\hat{y}, y)$. // compute loss
 $\omega_{\text{ft}}, \omega_{\text{mlp}} \leftarrow \text{SGD STEP}(J_{\omega}(t), \omega_{\text{ft}}, \omega_{\text{mlp}})$. // gradient descent
end for

ADAPTIVE INFERENCE

Input. Trained networks $\Phi_{\mathcal{D}}, F_{\text{ft}}, F_{\text{mlp}}$.
A set S of N points for prototype from domain D .
Output. Adaptive classifier for domain D .
 $\mu \leftarrow \text{PROTOTYPE COMPUTATION}(\Phi_{\mathcal{D}}, S)$.
return $F(\mathbf{x}) = F_{\text{mlp}}(\text{CONCAT}(F_{\text{ft}}(\mathbf{x}), \mu))$.

(since we can use unsupervised samples obtained *a priori* to construct prototypes). Note that once the prototype $\mu(\hat{D}_T)$ is computed, we only need the adapted classifier $F(\cdot, \mu(\hat{D}_T))$ for inference, not the prototypical network.

3.3. Theoretical Guarantees

Here we provide an abridged summary of our theoretical contributions, and defer the complete results and proofs to the appendix. We first demonstrate that expressivity and consistency are theoretically motivated using the framework of kernel mean embeddings [34]. Our first result is a concentration bound on the n -point approximation error of μ .

Theorem 1. (*μ approximation, informal*) For any domain D and feature Φ such that $\mathbb{E}_{\mathbf{x} \sim D}[\|\Phi(\mathbf{x})\|_2^2] \leq \sigma^2$, let the corresponding empirical dataset be $\hat{D}(n)$, let $\mu(D)$ be the true mean embedding and $\mu(\hat{D})$ be its approximation. Then with high probability, $\|\mu(D) - \mu(\hat{D})\|_{\infty} \lesssim \tilde{O}(\sigma/\sqrt{n})$ ⁵.

This suggests that for low error, the variance σ must be small for each domain $D \in \mathcal{D}$, a property achieved by prototypes [39]. Next, we provide a generalization bound that controls the gap between the training error $\hat{L}_{n,N}(F)$ and test error $L(F)$ (RHS of Eqn. 1) for F that is *adaptive*, i.e., a function of $\mathcal{X} \times \mathcal{D}$ and lies in the Hilbert space determined by some kernel κ . We build on the framework of [7] and extend their results to a more general class of kernel functions.

Theorem 2. (*error bound, informal*) Let F be an adaptive classifier defined over $\mathcal{X} \times \mathcal{D}$ that lies in an RKHS \mathcal{H}_{κ} such

⁵The \tilde{O} notation hides polylogarithmic factors and failure probability.

that $\|f\|_{\mathcal{H}_\kappa} \leq R$. Then, we have, with high probability that $|L(F) - \widehat{L}_{n,N}(F)| \lesssim \widetilde{O}\left(R \cdot \sigma((\log N)/n)^{\frac{1}{2}} + RN^{-\frac{1}{2}}\right)$.

Theorem 2 quantifies the error in terms of the feature variance σ , “norm” of function R and the number of samples n . In contrast to the results presented in prior work [7, 33, 21], we provide a tighter dependence on n that incorporates variance σ as well. See Appendix for more details.

4. Large Scale Benchmarks

There has been significant interest recently in domain generalization for vision problems [19]. However, we still lack large-scale benchmarks for rigorous evaluation. The largest dataset currently in use consists of only 7 domains and 5 classes [36], much smaller compared to real-world settings. Moreover, the model selection method of choice for this benchmark is *leave-one-domain-out* cross validation, which is infeasible in most large-scale applications with many domains. To address these issues, we introduce two large-scale domain generalization benchmarks.

4.1. Geo-YFCC

YFCC100M [42] is a large-scale dataset consisting of 100M images sourced from Flickr, along with tags submitted by users. Geotags (latitude and longitude) are also available for many of these images. We construct a dataset where each domain corresponds to a specific country.

Construction. We use geotags to partition images based on their country of origin. For the label space \mathcal{Y} , we consider the 4K categories from ImageNet-5K [10] not present in ILSVRC12 [37]. These categories are selected in order to eliminate biased prior knowledge from pre-training on ILSVRC12, which is the de-facto choice for large-scale pre-training. For each of the 4K labels, we select the corresponding images from YFCC100M based on a simple keyword-filtering of image tags. This provides us 1,261 categories with at least 1 image present. Furthermore, each category is present in at least 5 countries. We group images by their country of origin and only retain countries that have at least 10K images. For any domain with more than 20K images, we randomly sub-sample to limit it to 20K images. Therefore, each *domain* (i.e., country) has anywhere between 10K-20K images, giving us a total of 1,147,059 images from 1,261 categories across 62 countries (domains), and each image is associated with a class label and country (domain). We refer to this resulting dataset as **Geo-YFCC**. To the best of our knowledge, the scale of this dataset in the number of images, labels and domains is orders of magnitude more than prior datasets.

Train-Test Split. We randomly partition the data in to 45 training, 7 validation and 15 test domains (by country). For

each domain, we sample 3K points to create a per-domain test set and use the remaining points for training and validation. Since any image may have multiple labels, we convert it to a single-label dataset by replicating each such image and associating each separate copy for each label⁶, expanding the total image set to 1,809,832. We plan on releasing the complete partitions and annotations upon publication.

Analysis: Geo-YFCC exhibits several properties of real-world data, such as long-tailed label distributions [29], considerable *covariate shift* [40] and *label shift* [3] across domains. Figure 2A shows the label distributions from 4 sample domains, where the categories are ordered based on the frequency of labels from the “USA” domain. We see that each domain exhibits long tails and labels shifts (i.e., the marginal distribution of labels is different in each domain [28]). For example, “Cambodia” has “Temple” as the most frequent category while it is a low-frequency concept for “USA”. Figure 2B shows sample images for some categories from different domains, highlighting substantial *covariate shift* (i.e., shift in the marginal distribution of \mathcal{X}).

4.2. Long-Tailed ImageNet (LT-ImageNet)

We construct another dataset from the ImageNet-5K [10] benchmark by artificially introducing label shifts and long-tailed label distributions. The controlled settings in this dataset allow us to perform various ablations and understand our algorithm (Section 5.3). We subsample 500 categories from the 5K categories as the *base set* (excluding ILSVRC classes as earlier), from which we partition samples from each class into distinct train, val and test subsets. Each training domain is constructed as follows. We first select K head classes randomly from the base set. We then randomly sample A points from train samples of each head class and Af , $f < 1$ points from the remaining $500 - K$ tail classes to construct the train split of that domain, in order to mimic long-tailed phenomena. We use $A = 50$ and 300 to construct val and test splits for the domain from the corresponding splits of each class. We repeat this process N times to construct N training domains. Our overall training set has a non-uniform representation of classes in each domain, along with variation in long-tailed effects. We refer to the resulting dataset as **LT-ImageNet**.

To construct the val and test set of domains, we sample 10 domains with classes restricted to the set of all classes present in the training set (since we are not tackling zero-shot domain generalization), to ensure that the training data has at least 1 sample from each class (across all domains). We believe that these two large-scale datasets will help the community in properly understanding the effectiveness of existing and future domain generalization algorithms.

⁶This does not create any overlap between the train and test data as this duplication is done after the train-test split. As each image can belong only to one country, this ensures that no images are shared across domains.

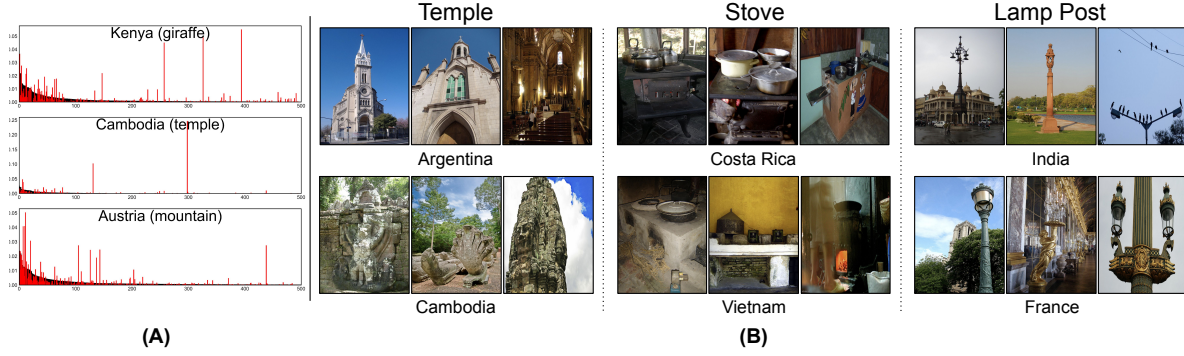


Figure 2: (A, left) Relative frequencies of categories from 3 different domains overlaid on the those of the USA, with most frequent category in parentheses. (B, right) Samples from 3 classes from different domains, highlighting covariate shift.

5. Experiments

We compare our algorithm with various recent algorithms on domain generalization problems both at the small-scale (i.e., fewer domains and classes per domain, and significant hyper-parameter tuning is possible) and at the real-world (large-scale) level. All benchmark algorithms are implemented from the DOMAINBED [19] suite with the hyper-parameter value ranges suggested therein.

Domain Embeddings (Section 3.1). We use a ResNet-50 [20] neural network pre-trained on ILSVRC12 [37] and introduce an additional fully-connected layer of dimension d_D on top of the pool5 layer. The output of this layer is used to construct the domain embedding (Equation 4). Unless otherwise stated, $d_D = 1024$, and the number of domains N_t used in N_t -way classification during each training step is set to 4. See Appendix for optimal hyper-parameters. To construct the domain embedding (Equation 4), we use at most 200 random points from the test set itself for small scale benchmarks in Section 5.1. For large-scale benchmarks, we leverage a subset of the samples from training domains and use a separate set of held-out points for test domains. For LT-ImageNet and Geo-YFCC, 300 and 5K points are used respectively.

DA-ERM (Section 3.2). The first neural network, F_{ft} , is again a ResNet-50 [20] network pre-trained on ILSVRC12 [37] with an additional fully-connected layer of dimension 1024 added to the pool5 layer. This layer is then concatenated with the d_D dimensional domain embedding and passed through F_{mlp} . Unless otherwise stated, the MLP has two layers with hidden dimension $d_{mlp} = 1024$, and output dimension set to number of classes. We use the standard cross-entropy loss for training. See the Appendix for optimal training hyper-parameters for each dataset.

5.1. Small-Scale Benchmarks

We conduct small-scale benchmark comparisons carefully following the DOMAINBED suite [19], which con-

tains 7 datasets in total. This paper focuses specifically on real-world domain generalization, and hence we select the 5 larger datasets (VLCS [14], PACS [23], Office-Home [43], Domain-Net [36] and Terra Incognita [5]) for our experiments, and discard the MNIST-based [22] datasets.

In this setting we select hyper-parameters by *leave one domain out* cross validation, the de facto technique for model selection for domain generalization, wherein we run N trials of each hyper-parameter setting, setting one domain for testing and the rest for training, and selecting the set of hyper-parameters that maximize validation accuracy over the training domains, averaged over all N trials. Performance under alternative techniques for model selection as well as the training details and hyper-parameter ranges for different algorithms can be found in the Appendix.

Domain Mixup. Since the small-scale datasets have typically 4 and at most 7 domains in total (1 of which is used as testing), the resulting training set (of domains) is not representative enough to learn a sufficiently discriminatory domain prototype. To alleviate this issue, we introduce *mixup* [48] in the prototype training step, where, for each batch of points from training domains, we construct *synthetic* domains by averaging points belonging to different domains and providing that as additional samples (belonging to a new domain). For example, during training, if we have a batch of points from each of the N_t training domains, we will create additional N_t synthetic batches where each synthetic batch is created by randomly averaging two batches. The averaging ratio for each synthetic domain is randomly chosen uniformly between (0.2, 0.8).

Results. Results are summarized in Table 1, where each experiment is averaged over 8 independent trials (for our algorithm, this includes retraining of the domain prototypical network). DA-ERM provides an average improvement of **1.2%** in accuracy across all 5 datasets over the ERM baseline that simply combines data from all the training domains to train a model. The improvement is more sub-

Algorithm	VLCS [14]	PACS [23]	OffHome [43]	DNet [36]	TerraInc [5]	Average
ERM [19]	76.7 \pm 0.9	83.2 \pm 0.7	67.2 \pm 0.5	41.1 \pm 0.8	46.2 \pm 0.3	62.9 \pm 0.6
IRM [2]	76.9 \pm 0.5	82.8 \pm 0.5	67.0 \pm 0.4	35.7 \pm 1.9	43.8 \pm 1.0	61.2 \pm 0.8
DRO [38]	77.3 \pm 0.2	83.3 \pm 0.3	66.8 \pm 0.1	33.0 \pm 0.5	42.0 \pm 0.6	60.4 \pm 0.4
Mixup [44]	78.2 \pm 0.6	83.9 \pm 0.8	68.3 \pm 0.4	40.2 \pm 0.4	46.2 \pm 0.5	63.3 \pm 0.7
MLDG [24]	76.8 \pm 0.4	82.6 \pm 0.6	67.7 \pm 0.3	42.2 \pm 0.6	46.0 \pm 0.4	63.0 \pm 0.4
CORAL [41]	77.3 \pm 0.3	83.3 \pm 0.5	68.6 \pm 0.1	42.1 \pm 0.7	47.7 \pm 0.3	63.8 \pm 0.3
MMD [25]	76.1 \pm 0.7	83.1 \pm 0.7	67.3 \pm 0.2	38.7 \pm 0.5	45.8 \pm 0.6	62.2 \pm 0.6
C-DANN [26]	73.8 \pm 1.1	81.4 \pm 1.3	64.2 \pm 0.5	39.5 \pm 0.2	40.9 \pm 0.7	60.1 \pm 0.9
DA-ERM	78.0 \pm 0.2	84.1 \pm 0.5	67.9 \pm 0.4	43.6 \pm 0.3	47.3 \pm 0.5	64.1 \pm 0.8
DA-MMD	78.6 \pm 0.5	84.4 \pm 0.3	68.2 \pm 0.1	43.3 \pm 0.5	47.9 \pm 0.4	64.5 \pm 0.5
DA-CORAL	78.5 \pm 0.4	84.5 \pm 0.7	68.9 \pm 0.4	43.9 \pm 0.3	48.1 \pm 0.3	64.7 \pm 0.6

Table 1: Small-Scale Benchmark Comparisons on DOMAINBED.

Algorithm	LT-ImageNet		Geo-YFCC	
	Train	Test	Train	Test
	Top-1 / 5	Top-1 / 5	Top-1 / 5	Top-1 / 5
ERM [19]	70.9 / 90.3	53.7 / 77.5	28.5 / 56.3	22.4 / 48.2
CORAL [41]	71.1 / 89.3	53.5 / 76.2	25.5 / 51.2	21.8 / 46.4
MMD [25]	70.9 / 88.0	52.8 / 76.4	25.4 / 50.9	21.8 / 46.2
DA-ERM	73.4 / 91.8	56.1 / 80.5	28.3 / 55.8	23.4 / 49.1

Table 2: Large-Scale Comparisons.

stantial in Domain-Net [36], that has a larger number of domains. This can be attributed to the increased diversity in the training data, which allows us to construct embeddings that cover the space of domains better. In comparison to other domain-invariant approaches, DA-ERM either performs better or competitively with the best approach on each individual dataset.

To demonstrate the complimentary value of *adaptive* learning to domain-invariant methods, we modify the second step of our algorithm and consider two variants of the DA-ERM approach: DA-MMD and DA-CORAL, by replacing the cross-entropy in DA-ERM with the loss terms in domain-invariant approaches MMD [25] and CORAL [41] respectively⁷ (see the appendix for details). Compared to CORAL, DA-CORAL provides a significant improvement of 0.9%, and DA-MMD achieves an impressive improvement of **2.3%** (62.2% vs. 64.5%) over MMD.

5.2. Large-Scale Benchmarks

We now present results on two large-scale benchmarks proposed in Section 4. We compare our approach with ERM, CORAL [41] (the best performing baseline in Section 5.1), and, MMD [25]. We were unable to replicate several algorithms at scale, such as DANN [15], Conditional-DANN [30] and MMLD [35]: these algorithms rely on adversarial training that requires careful hyper-parameter optimization, which we were unable to accomplish based on hyper-parameters suggested in DOMAINBED [19].

LT-ImageNet. We set $N=50$, $K=100$, $A=350$ and $f=0.1$ to construct the LT-ImageNet (50, 100, 350, 0.1) dataset and report top-1 and top-5 accuracies on the test splits of train and test domains in Table 2 (left). We observe that DA-ERM improves top-1 performance by 2.3% (71.1% vs. 73.4%) on train domains and 2.4% (53.7% vs 56.1%) on test domains. We show the effect of varying the dataset settings (N , K , A and f) in Section 5.3.

Geo-YFCC. Table 2 (right) shows comparisons on the test

⁷Auxiliary loss terms minimizing the differences between the domains is applied to the d_{mlp} dimensional hidden layer in the MLP F_{mlp} .

splits for train and test domains of the Geo-YFCC dataset. DA-ERM provides an absolute improvement of 1% in comparison to benchmark techniques on test domains. It must be noted that this problem is substantially more challenging, as we have 1,261 classes with significant domain differences between the countries (Section 4). We also observe that all other domain-invariant baselines were unable to obtain improvements over the simple ERM baseline. Figure 3A shows t-SNE [31] visualizations of prototypical embeddings for each domain, grouped by continent. Countries belonging to the same continent are mapped closer together, expressing the semantics captured by μ . Additionally, we note that our algorithm simultaneously improves performance on test domains while maintaining performance on training domains, unlike other techniques.

5.3. Ablations

5.3.1 Dataset Ablations

We consider the LT-ImageNet dataset with default settings mentioned in Section 5.2, unless otherwise stated.

Effect of diversity in training data. We first vary the number of available training domains (N), while sampling the same number of points from each domain. From Figure 3B, we observe that compared to other algorithms, the improvements in domain-aware training are larger, which we attribute to diversity in prototype training. Next, we repeat the same experiment, however, we subsample points from each domain such that the overall number of training points are the same, i.e., NA , remains constant. We see in Figure 3C that DA-ERM performs much better than other algorithms with increasing number of domains. Note that once $N > 100$, a slight dip is observed in our algorithm. This can be understood intuitively from Theorem 2, as it suggests such an inflection point when the total number of points (nN) is fixed due to the interaction of the two terms of the RHS. As the number of domains N increases, the first term decreases, while the second term increases.

Effect of Tail Index. In Figure 3D, we demonstrate the

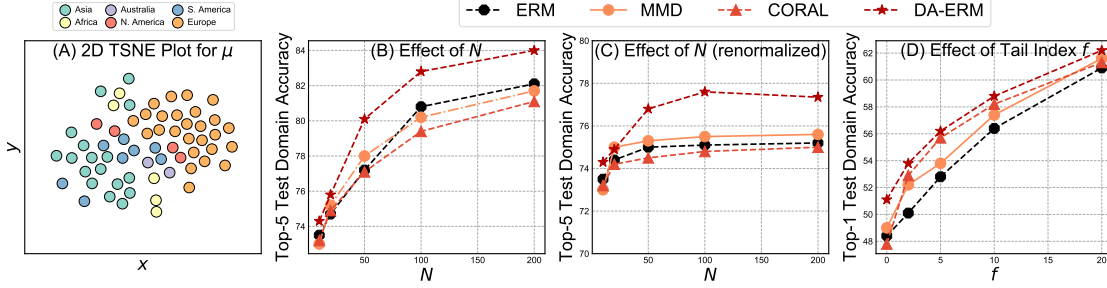


Figure 3: Ablation experiments: Fig. A (left) is a t-SNE visualization of domain embeddings. Figs. B (center left) and C (center right) show the effect of changing the number of domains, where the total number of training points increases linearly with the number of domains in (B) and is constant in (C). Fig. D (right) demonstrates the effect of the tail index f .

effect of varying f : the parameter f controls the extent of imbalance present in each domain, by adjusting the fraction of points available in the *tail classes* compared to the head classes. We see that when the dataset is severely imbalanced (i.e., $f < 5\%$), domain-adaptive training considerably outperforms existing techniques, with a top-1 performance gap of about **3.5%** on average compared to naive ERM. As f increases and the imbalance becomes smaller, our method performs closer to ERM itself, with an average improvement of only **1.8%**. Since most real-world datasets are extremely long-tailed, we expect this to be a valuable feature of domain-adaptive learning in practice.

5.3.2 Algorithm Ablations

We perform several ablation studies on the domain embedding algorithm itself to verify the two properties of *consistency* and *expressivity* defined in Section 3.1.

Consistency. We varied the number of points used to construct the prototype on Geo-YFCC and observed that the performance declines only for fewer than 50 points. For other values up to 2K, performance remains virtually identical (see Appendix). This is desirable as in many settings we do not have access to many samples from new domains.

Expressivity. We measure the ability of the prototype to efficiently embed different domains distinctly within the embedding space. We evaluate the following two alternate approaches to construct domain embedding: (a) *Mean*: we select the ERM baseline model and simply compute the average `pool5` features for each domain; (b) *Softmax*: we train the domain embedding network via cross-entropy with domain IDs as labels instead of the prototypical loss. The results are summarized in Table 3. For reference, we also report results of naive ERM (*None*) and supplying incorrect embeddings (i.e., the embedding of some other random domain) during evaluation of test domains (*Random*). We observe that using random features degrades performance significantly, indicating that domain embeddings indeed play

Dataset	Top-1 Accuracy on Test Domains				
	None	Random	Mean	Softmax	Prototype
LT - ImageNet	76.9	75.3	77.9	78.8	80.5
Geo - YFCC	22.4	21.3	23.0	23.5	23.4

Table 3: Ablation of various domain embedding algorithms.

an important role in the final model. Mean features, while useful (i.e., perform better than naive ERM), are not as expressive as prototypical learning. Softmax features show mixed results by performing worse than Prototype on LT-ImageNet and slightly better on Geo-YFCC. We additionally study the effect of number of domains sampled per round in prototypical training, and observe no significant effect for values in the range $[3, 8]$.

6. Conclusion

The problem of domain generalization lies at the heart of many machine learning applications. Research on the problem has largely been restricted to small-scale datasets that are constructed from artificially diverse sources that are not representative of real-world settings. For example, it is impractical to train a model on sketches if the target domain contains photographs: it is much easier in practice to curate a dataset of photographs instead. With this paper *work*, our goal is to provide an algorithm and evaluation procedure that is closer to real-world applications, and involves experimentation at scale. We presented an *adaptive* domain generalization framework, and conducted the first set of experiments at scale for this problem. Our results suggest that we need significant effort to scale to real-world problem settings, and with our contributions we expect to initiate research in this new direction.

Acknowledgements. We would like to sincerely thank Deepti Ghadiyaram, Filip Radenovic and Cheng-Yang Fu for their helpful feedback on our paper.

References

- [1] Sami Abu-El-Haija, Nisarg Kothari, Joonseok Lee, Paul Natsev, George Toderici, Balakrishnan Varadarajan, and Sudheendra Vijayanarasimhan. Youtube-8m: A large-scale video classification benchmark. *arXiv preprint arXiv:1609.08675*, 2016. [2](#)
- [2] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019. [2](#), [7](#)
- [3] Kamyar Azizzadenesheli, Anqi Liu, Fanny Yang, and Animashree Anandkumar. Regularized learning for domain adaptation under label shifts. *arXiv preprint arXiv:1903.09734*, 2019. [2](#), [5](#)
- [4] Mahsa Baktashmotlagh, Mehrtash T Harandi, Brian C Lovell, and Mathieu Salzmann. Unsupervised domain adaptation by domain invariant projection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 769–776, 2013. [1](#), [2](#)
- [5] Sara Beery, Grant Van Horn, and Pietro Perona. Recognition in terra incognita. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 456–473, 2018. [6](#), [7](#)
- [6] 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. [1](#)
- [7] Gilles Blanchard, Gyemin Lee, and Clayton Scott. Generalizing from several related classification tasks to a new unlabeled sample. In *Advances in neural information processing systems*, pages 2178–2186, 2011. [1](#), [2](#), [4](#), [5](#)
- [8] Léon Bottou. Large-scale machine learning with stochastic gradient descent. In *Proceedings of COMPSTAT'2010*, pages 177–186. Springer, 2010. [3](#)
- [9] Dong Cao, Xiangyu Zhu, Xingyu Huang, Jianzhu Guo, and Zhen Lei. Domain balancing: Face recognition on long-tailed domains. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5671–5679, 2020. [2](#)
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. [2](#), [5](#)
- [11] Aniket Anand Deshmukh, Urun Dogan, and Clay Scott. Multi-task learning for contextual bandits. In *Advances in neural information processing systems*, pages 4848–4856, 2017. [2](#)
- [12] Aniket Anand Deshmukh, Yunwen Lei, Srinagesh Sharma, Urun Dogan, James W Cutler, and Clayton Scott. A generalization error bound for multi-class domain generalization. *arXiv preprint arXiv:1905.10392*, 2019. [2](#)
- [13] Abhimanyu Dubey and Alex Pentland. Kernel methods for cooperative multi-agent contextual bandits. *arXiv preprint arXiv:2008.06220*, 2020. [2](#)
- [14] Chen Fang, Ye Xu, and Daniel N Rockmore. Unbiased metric learning: On the utilization of multiple datasets and web images for softening bias. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1657–1664, 2013. [6](#), [7](#)
- [15] 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](#), [7](#)
- [16] 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. [1](#)
- [17] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2066–2073. IEEE, 2012. [2](#)
- [18] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. A kernel two-sample test. *The Journal of Machine Learning Research*, 13(1):723–773, 2012. [2](#)
- [19] Ishaan Gulrajani and David Lopez-Paz. In search of lost domain generalization. *arXiv preprint arXiv:2007.01434*, 2020. [1](#), [2](#), [5](#), [6](#), [7](#)
- [20] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. [6](#)
- [21] Shoubo Hu, Kun Zhang, Zhitang Chen, and Laiwan Chan. Domain generalization via multidomain discriminant analysis. In *Uncertainty in Artificial Intelligence*, pages 292–302. PMLR, 2020. [5](#)
- [22] Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>. [6](#)
- [23] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Deeper, broader and artier domain generalization. In *Proceedings of the IEEE international conference on computer vision*, pages 5542–5550, 2017. [2](#), [6](#), [7](#)
- [24] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M Hospedales. Learning to generalize: Meta-learning for domain generalization. *arXiv preprint arXiv:1710.03463*, 2017. [2](#), [7](#)
- [25] Ya Li, Mingming Gong, Xinmei Tian, Tongliang Liu, and Dacheng Tao. Domain generalization via conditional invariant representation. *arXiv preprint arXiv:1807.08479*, 2018. [1](#), [2](#), [7](#)
- [26] 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. [1](#), [2](#), [7](#)
- [27] Dahua Lin. Distribution-balanced loss for multi-label classification in long-tailed datasets. 2020. [2](#)
- [28] Zachary C Lipton, Yu-Xiang Wang, and Alex Smola. Detecting and correcting for label shift with black box predictors. *arXiv preprint arXiv:1802.03916*, 2018. [5](#)
- [29] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed

- recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019. 5
- [30] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pages 1640–1650, 2018. 7
- [31] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(Nov):2579–2605, 2008. 7
- [32] Divyat Mahajan, Shruti Tople, and Amit Sharma. Domain generalization using causal matching. *arXiv preprint arXiv:2006.07500*, 2020. 2
- [33] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning*, pages 10–18, 2013. 1, 5
- [34] Krikamol Muandet, Kenji Fukumizu, Bharath Sriperumbudur, and Bernhard Schölkopf. Kernel mean embedding of distributions: A review and beyond. *arXiv preprint arXiv:1605.09522*, 2016. 3, 4
- [35] Zhongyi Pei, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Multi-adversarial domain adaptation. *arXiv preprint arXiv:1809.02176*, 2018. 2, 7
- [36] Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 1406–1415, 2019. 5, 6, 7
- [37] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115(3):211–252, 2015. 5, 6
- [38] 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. 2, 7
- [39] Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087, 2017. 1, 3, 4
- [40] Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(May):985–1005, 2007. 5
- [41] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European conference on computer vision*, pages 443–450. Springer, 2016. 2, 7
- [42] Bart Thomee, David A Shamma, Gerald Friedland, Benjamin Elizalde, Karl Ni, Douglas Poland, Damian Borth, and Li-Jia Li. Yfcc100m: The new data in multimedia research. *Communications of the ACM*, 59(2):64–73, 2016. 2, 5
- [43] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *(IEEE) Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. 6, 7
- [44] 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, 7
- [45] Liuyu Xiang, Guiguang Ding, and Jungong Han. Learning from multiple experts: Self-paced knowledge distillation for long-tailed classification. In *European Conference on Computer Vision*, pages 247–263. Springer, 2020. 2
- [46] Minghao Xu, Jian Zhang, Bingbing Ni, Teng Li, Chengjie Wang, Qi Tian, and Wenjun Zhang. Adversarial domain adaptation with domain mixup. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 6502–6509, 2020. 2
- [47] Shen Yan, Huan Song, Nanxiang Li, Lincan Zou, and Liu Ren. Improve unsupervised domain adaptation with mixup training. *arXiv preprint arXiv:2001.00677*, 2020. 2
- [48] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 6
- [49] Xiao Zhang, Zhiyuan Fang, Yandong Wen, Zhifeng Li, and Yu Qiao. Range loss for deep face recognition with long-tailed training data. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5409–5418, 2017. 2