

Domain Generalization via Conditional Invariant Representations

Ya Li,[†] Mingming Gong,^{★*} Xinmei Tian,[†] Tongliang Liu,[‡] Dacheng Tao[‡]

[†]CAS Key Laboratory of Technology in Geo-Spatial Information Processing and Application Systems,
University of Science and Technology of China, China

[‡]UBTECH Sydney Artificial Intelligence Institute, SIT, FEIT, The University of Sydney, Australia

[★]Department of philosophy, Carnegie Mellon University

^{*}Department of Biomedical Informatics, University of Pittsburgh

muziyiye@mail.ustc.edu.cn, gongmingnju@gmail.com, xinmei@ustc.edu.cn,

tliang.liu@gmail.com, dacheng.tao@sydney.edu.au

Abstract

Domain generalization aims to apply knowledge gained from multiple labeled source domains to unseen target domains. The main difficulty comes from the dataset bias: training data and test data have different distributions, and the training set contains heterogeneous samples from different distributions. Let X denote the features, and Y be the class labels. Existing domain generalization methods address the dataset bias problem by learning a domain-invariant representation $h(X)$ that has the same marginal distribution $\mathbb{P}(h(X))$ across multiple source domains. The functional relationship encoded in $\mathbb{P}(Y|X)$ is usually assumed to be stable across domains such that $\mathbb{P}(Y|h(X))$ is also invariant. However, it is unclear whether this assumption holds in practical problems. In this paper, we consider the general situation where both $\mathbb{P}(X)$ and $\mathbb{P}(Y|X)$ can change across all domains. We propose to learn a feature representation which has domain-invariant class conditional distributions $\mathbb{P}(h(X)|Y)$. With the conditional invariant representation, the invariance of the joint distribution $\mathbb{P}(h(X), Y)$ can be guaranteed if the class prior $\mathbb{P}(Y)$ does not change across training and test domains. Extensive experiments on both synthetic and real data demonstrate the effectiveness of the proposed method.

Introduction

Recent years have witnessed a great success of supervised learning in various pattern recognition problems, such as image classification, object detection, and speech recognition. Standard supervised learning relies heavily on the *i.i.d.* data assumption; however, dataset-bias is unavoidable in many situations due to selection bias or mechanism changes. For example, this problem has been well recognized in the computer vision community (Torralba and Efros 2011; Khosla et al. 2012): the widely adopted vision datasets have their special properties and are not representative of the visual world. In medical diagnosis, the distribution of cell types varies from patient to patient, and we need to train a classifier on the data collected from previous patients that generalizes well to unseen patients (Blanchard, Lee, and Scott 2011;

Muandet, Balduzzi, and Schölkopf 2013). These problems are known as domain generalization, in which the training set consists of data from heterogeneous source domains, say patients, and the test data distribution is different from that of the training data.

To handle the distribution changes, many existing domain generalization methods aim to learn domain-invariant representations that have stable distributions across all source domains (Muandet, Balduzzi, and Schölkopf 2013; Erfani et al. 2016; Ghifary et al. 2017). The learned invariant representations are expected to generalize well to any unseen test set under the assumption that the changes of distribution across source and test domains are caused by some common factors whose effects are removed in the invariant representations. In computer vision, such factors could be illumination, camera viewpoints, and backgrounds. These methods have achieved good performance in computer vision (Ghifary et al. 2015; 2017) and medical diagnosis (Muandet, Balduzzi, and Schölkopf 2013).

However, existing methods that learn domain-invariant representations assume that only $\mathbb{P}(X)$ changes across domains while the conditional distribution $\mathbb{P}(Y|X)$ is rather stable. Thus, the conditional distribution $\mathbb{P}(Y|h(X))$ is also invariant, and the learning problem reduces to ensuring that the marginal distribution $\mathbb{P}(h(X))$ is invariant across domains. This assumption greatly simplifies the problem, but it is unclear whether this assumption holds in practical situations. According to some recent results in causal learning (Schölkopf et al. 2012; Janzing and Schölkopf 2010), $\mathbb{P}(Y|X)$ can be stable when $\mathbb{P}(X)$ changes in the situation where X is the cause for Y , i.e., the causal structure is $X \rightarrow Y$. This is because the mechanism that generates the cause, i.e., $\mathbb{P}(X)$, is not coupled with the mechanism that generates the effect from the cause, i.e., $\mathbb{P}(Y|X)$, and not vice versa. That is to say, if Y is the cause and X is the effect, $\mathbb{P}(X)$ often changes together with $\mathbb{P}(Y|X)$. In this situation, if $\mathbb{P}(X)$ changes, it is very likely that $\mathbb{P}(Y|X)$ also changes across domains, which violates the stability of $\mathbb{P}(Y|X)$ assumption. In practice, we have plenty of problems where the causal structure is $Y \rightarrow X$. For example,

, in face recognition, Y is person id, X is the feature, and θ is the viewpoint. Let us consider each viewpoint as a domain, then in each domain we have conditional distribution $P(X|Y, \theta = \theta_i)$. According to Bayes theorem, $P(Y|X, \theta = \theta_i) = P(X|Y, \theta = \theta_i)P(Y|\theta = \theta_i)/P(X|\theta = \theta_i)$, thus changes across domains. This conflicts with previous assumptions that $P(Y|X)$ keeps unchanged. There are also other examples, e.g. speaker recognition and person re-identification (Yang et al. 2017).

In this paper, we assume both $\mathbb{P}(X)$ and $\mathbb{P}(Y|X)$ change across domains. We aim to find a feature transformation $h(X)$ that has invariant class-conditional distribution $\mathbb{P}(h(X)|Y)$. To achieve so, we propose to minimize two regularization terms that enforce distribution invariance across source domains. The first term measures the variance of each class-conditional distribution across all source domains and then sums up the variances for all classes. The second term is the variance of class prior-normalized marginal distribution $\mathbb{P}_N(h(X))$, which measures the global distribution discrepancy. The normalization of class priors is introduced to remove the effects brought by possible changes in $\mathbb{P}(Y)$ across source domains. If the prior distribution $\mathbb{P}(Y)$ does not change across source domains, the second term reduces to the common technique used in existing domain-invariant representation learning methods (Muandet, Balduzzi, and Schölkopf 2013; Ghifary et al. 2017). To preserve the discriminative power of the learned representation, we also incorporate the intra-class and inter-class distances used in kernel Fisher discriminant analysis (FDA)(Mika et al. 1999).

Compared to existing domain-invariant representation learning methods, our method does not require the assumption of stable $\mathbb{P}(Y|X)$ by exploiting the labels on the source domains which were overlooked in the previous methods. Especially, if the prior distribution $\mathbb{P}(Y)$ on the test sets is the same as that on the training set containing all source domains, our method is able to learn representations $h(X)$ that have invariant joint distribution $\mathbb{P}(h(X), Y)$ across all domains. We conduct a series of experiments on both synthetic and real data, and the results demonstrate the effectiveness of our method.

Related Work

Domain generalization has been widely applied in classification tasks (Xu et al. 2014; Duan et al. 2009; Muandet, Balduzzi, and Schölkopf 2013; Ghifary et al. 2017; 2015; Erfani et al. 2016). Compared with standard supervised learning, domain generalization methods aim to reduce data bias across different domains and improve the generalization of the learned model to unseen but related domains. For example, (Xu et al. 2014) assumed that positive samples within the same shared latent domain should have similar likelihood and proposed to exploit the low-rank structure from latent domains for domain generalization. (Muandet, Balduzzi, and Schölkopf 2013) proposed domain-invariant component analysis (DICA) through learning an invariant feature representation $h(X)$, in which the difference between marginal distributions $\mathbb{P}(h(X))$ is minimized. (Ghifary et al. 2017) proposed a unified framework called scatter

component analysis for domain adaptation and domain generalization. The scatter component analysis combines domain scatter (Muandet, Balduzzi, and Schölkopf 2013), kernel PCA (Schölkopf, Smola, and Müller 1998), and kernel FDA (Mika et al. 1999) in a single objective function. However, all these methods assume that the distribution between domains differs only in the marginal distribution $\mathbb{P}(X)$ while the conditional distribution $\mathbb{P}(Y|X)$ keeps stable or unchanged across domains. This assumption can simplify the problem of domain generalization, but it is easily violated in real-world applications.

Domain adaptation is a related problem which has been extensively studied in the literature (Baktashmotlagh et al. 2013; Huang et al. 2007; Pan et al. 2011; Long et al. 2017; Shao, Kit, and Fu 2014; Shao et al. 2016; Luo et al. 2017; Liu, Yang, and Tao 2017). Assuming that only $\mathbb{P}(X)$ changes, the distribution changes can be corrected by importance reweighting (Huang et al. 2007) or domain-invariant feature learning (Pan et al. 2011; Baktashmotlagh et al. 2013), using unlabeled data from source and target domains. Recently, several works attempted to work in the situation where both $\mathbb{P}(X)$ and $\mathbb{P}(Y|X)$ change across domains (Zhang et al. 2013; Gong et al. 2016; Long et al. 2017). (Zhang et al. 2013) and (Gong et al. 2016) proposed to consider the domain adaptation problem in the generalized target shift (GeTarS) scenario where the causal direction is $Y \rightarrow X$. In this scenario, both the change of distribution $\mathbb{P}(Y)$ and conditional distribution $\mathbb{P}(X|Y)$ are considered to reduce the data bias across domains. (Zhang et al. 2013) made an assumption that features from source domains can be transferred to the target domain by a location-scale transformation, which is restricted in real-world applications because of the presence of noises in features. (Gong et al. 2016) proposed to learn components whose conditional distribution $\mathbb{P}(h(X)|Y)$ is invariant across domains and estimate the target label distribution $\mathbb{P}^t(Y)$ through labeled source domain data and unlabeled target domain data. Since there are no labels in the target domain to match class-conditionals, the invariance of $\mathbb{P}(h(X)|Y)$ is achieved by minimizing the discrepancy of the marginal distribution $\mathbb{P}(h(X))$ under some untestable assumptions. (Long et al. 2017) proposed an iterative way to match the conditionals by using the predicted labels from previous iterations as pseudo labels. Different from the domain adaptation methods, domain generalization does not require unlabeled data from the target domains.

Conditional Invariant Domain Generalization

In this section, we first establish the basic notations of domains and formally introduce the definition of domain generalization. Then we give a detailed description of the proposed conditional invariant domain generalization (CIDG) method.

Problem Definition

Denote \mathcal{X} and \mathcal{Y} as the input feature and label spaces, respectively. A domain defined on $\mathcal{X} \times \mathcal{Y}$ can be represented by a joint probability distribution $\mathbb{P}(X, Y)$. For simplicity,

we denote the joint probability distribution $\mathbb{P}^s(X, Y)$ of the s -th source domain as \mathbb{P}^s . The domain \mathbb{P}^s is associated with a sample $D_s = \{x_i^s, y_i^s\}_{i=1}^{n^s}$, where $(x_i^s, y_i^s) \sim \mathbb{P}^s$ and n^s denotes the sample size of the domain \mathbb{P}^s . Then we can define domain generalization as follows.

Definition 1 (Domain Generalization). *Given multiple related source domains $\Omega = \{\mathbb{P}^1, \mathbb{P}^2, \dots, \mathbb{P}^m\}$ and each domain is associated with a sample $D_s = \{x_i^s, y_i^s\}_{i=1}^{n^s} \sim \mathbb{P}^s$, where $s = \{1, 2, \dots, m\}$. The goal of domain generalization is to learn a classification function $f : \mathcal{X} \rightarrow \mathcal{Y}$ from source domain datasets $\{D_s\}_{s=1}^m$ and apply it to an unseen but related target domain $\mathbb{P}^t(X, Y)$.*

Kernel Mean Embedding

Before introducing the proposed method, we briefly review the kernel mean embedding of distributions, which is an important mathematical tool to represent and compare distributions (Song, Fukumizu, and Gretton 2013; Sriperumbudur et al. 2010). Let \mathcal{H} denote a characteristic reproducing kernel Hilbert space (RKHS) on \mathcal{X} associated with a kernel $k(\cdot, \cdot) : \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, and ϕ be an associated mapping such that $\phi(x) \in \mathcal{H}$. Suppose we have two observations $x_1^s \in \mathcal{X}$ and $x_2^s \in \mathcal{X}$ from domain s , then we have $\langle \phi(x_1^s), \phi(x_2^s) \rangle = k(x_1^s, x_2^s)$. The kernel embedding of a distribution $\mathbb{P}(X)$ can be formulated as the following:

$$\mu_{\mathbb{P}_X} := E_{X \sim \mathbb{P}_X}[\phi(X)] = E_{X \sim \mathbb{P}_X}[k(X, \cdot)], \quad (1)$$

where \mathbb{P}_X denotes $\mathbb{P}(X)$ for simplicity. If a kernel is characteristic, then the mean embedding $\mu_{\mathbb{P}_X}$ is injective. All the information about the distribution can be preserved (Sriperumbudur et al. 2010). The kernel embedding cannot be computed directly and is usually estimated from observations. Given a sample $D = \{x_i\}_{i=1}^n$, where n is the sample size of the domain, and the kernel embedding can be empirically estimated as the following:

$$\hat{\mu}_{\mathbb{P}_X} = \frac{1}{n} \sum_{i=1}^n \phi(x_i) = \frac{1}{n} \sum_{i=1}^n k(x_i, \cdot). \quad (2)$$

Proposed Approach

The proposed conditional invariant domain generalization (CIDG) method aims to find a conditional invariant representation $h(X)$ (a linear transformation of the original features) to reduce the variance of the conditional distribution $\mathbb{P}(h(X)|Y)$ across source domains. Suppose we can learn a perfect conditional invariant representation $h(X)$, which satisfies $\mathbb{P}^{s=i}(h(X)|Y) = \mathbb{P}^{s=j}(h(X)|Y) = \mathbb{P}^t(h(X)|Y)$, $i, j \in \{1, 2, \dots, m\}$ and \mathbb{P}^t denotes the target domain. We can gather all the source domains to construct a new single domain with a joint distribution $\mathbb{P}^t(h(X)|Y)\mathbb{P}^{\text{new}}(Y)$. Therefore, under the condition $\mathbb{P}^{\text{new}}(Y) = \mathbb{P}^t(Y)$, the learned $h(X)$ has the invariant joint distribution across training and test domains. Contrarily, the previous method can only guarantee that $\mathbb{P}(h(X))$ is invariant, and whether $\mathbb{P}(Y|h(X))$ is invariant remains unknown. If $\mathbb{P}^{\text{new}}(Y)$ is different from $\mathbb{P}^t(Y)$, our method cannot guarantee the invariance of the joint distribution either. Nevertheless, our

method can at least guarantee invariant class-conditional distributions, which is still better than previous methods. This is because $\mathbb{P}(Y|h(X))$ is usually not very sensitive to the changes in the prior $\mathbb{P}(Y)$ if $h(X)$ is highly correlated with Y .

The learning of conditional invariant representations is achieved mainly through two regularization terms: total scatter of class-conditional distributions and scatter of class prior-normalized marginal distributions. The first term measures the variance of $\mathbb{P}(h(X)|Y)$ locally, while the second term measures the variance of $\mathbb{P}(h(X)|Y)$ globally. In addition to these two terms, we also incorporate several terms that measure the discriminative power of the representation $h(X)$ as done in the previous works. By minimizing the distribution variance across domains and maximizing the discriminative power in one objective function, we can obtain the conditional invariant representation which is predictable for the labels on unseen target domains.

Total scatter of class-conditional distributions Suppose we have m related domains $\{\mathbb{P}^1, \mathbb{P}^2, \dots, \mathbb{P}^m\}$ on $\mathcal{X} \times \mathcal{Y}$. The marginal distribution on \mathcal{X} of the s -th domain is denoted as \mathbb{P}_X^s . Suppose the class labels of each domain vary from 1 to C . For simplicity, the j -th class conditional distribution $\mathbb{P}^s(X|Y = j)$ of the s -th domain is denoted as \mathbb{P}_j^s . The total scatter of class-conditional distributions across domains can be formulated as:

$$\Psi \left(\left\{ \mu_{\mathbb{P}_1^s}, \mu_{\mathbb{P}_2^s}, \dots, \mu_{\mathbb{P}_C^s} \right\} \right) = \sum_{j=1}^C \frac{1}{m} \sum_{s=1}^m \|\mu_{\mathbb{P}_j^s} - \bar{\mu}_j\|_{\mathcal{H}}^2, \quad (3)$$

where $\bar{\mu}_j = \frac{1}{m} \sum_{s=1}^m \mu_{\mathbb{P}_j^s}$ and $\frac{1}{m} \sum_{s=1}^m \|\mu_{\mathbb{P}_j^s} - \bar{\mu}_j\|_{\mathcal{H}}^2$ is called the domain scatter (Ghifary et al. 2017) or distributional variance (Muandet, Balduzzi, and Schölkopf 2013). Instead of measuring the domain scatter w.r.t. the marginal distributions \mathbb{P}_X^s as done in previous works like (Ghifary et al. 2017), we measure the domain scatter w.r.t. each class-conditional distribution and then sum them together.

Before introducing the computation of the above scatter, we first give the formulation of the learned feature transformation. Denote the feature matrix $\mathbf{X} = [x_1, x_2, \dots, x_n]^T \in \mathbb{R}^{n \times d}$ as the data matrix of samples from m source domains, where d is the dimension of the feature space \mathcal{X} and $n = \sum_{s=1}^m n^s$. Define a set of functions $\Phi = [\phi(x_1), \phi(x_2), \dots, \phi(x_n)]^T$ related to the feature map $\phi : \mathbb{R}^d \rightarrow \mathcal{H}$. We aim to find a linear feature transformation \mathbf{W} transforming \mathcal{H} into a finite subspace $\mathcal{H} \rightarrow \mathbb{R}^q$, that is $h(x) = \mathbf{W}^T \phi(x)$. According to the kernel principal component analysis (KPCA) (Schölkopf, Smola, and Müller 1998), the linear transformation can be formulated as the linear combination of Φ , i.e., $\mathbf{W} = \Phi^T \mathbf{B}$, where $\mathbf{B} \in \mathbb{R}^{n \times q}$ is the coefficient matrix. By using this representation, we can avoid explicitly computing the feature map ϕ and use the kernel trick instead.

For simplicity, denote $\Psi \left(\{\mu_{\mathbb{P}_1^s}, \mu_{\mathbb{P}_2^s}, \dots, \mu_{\mathbb{P}_C^s}\} \right)$ as Ψ^{con} ,

$$\begin{aligned} \Psi^{con} &= \frac{1}{m} \sum_{s=1}^m \sum_{j=1}^C \|\mu_{\mathbb{P}_j^s} - \bar{\mu}_j\|_{\mathcal{H}}^2 \\ &= \frac{1}{m} \sum_{s=1}^m \sum_{j=1}^C \text{Tr} \left((\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)(\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)^\top \right) \\ &= \text{Tr} \left(\frac{1}{m} \sum_{s=1}^m \sum_{j=1}^C (\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)(\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)^\top \right), \end{aligned} \quad (4)$$

where $\text{Tr}(\cdot)$ is trace operator. To measure the distribution scatter of the distributions of $\mathbb{P}(h(X)|Y)$, we apply the linear feature transformation \mathbf{W} to the above scatter and obtain

$$\begin{aligned} \Psi_{\mathbf{B}}^{con} &= \text{Tr} \left(\frac{1}{m} \sum_{s=1}^m \sum_{j=1}^C \mathbf{B}^\top \Phi(\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)(\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)^\top \Phi^\top \mathbf{B} \right) \\ &= \text{Tr}(\mathbf{B}^\top \mathbf{H} \mathbf{B}), \end{aligned} \quad (5)$$

where \mathbf{H} is:

$$\mathbf{H} = \sum_{s=1}^m \frac{1}{m} \sum_{j=1}^C \Phi(\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)(\mu_{\mathbb{P}_j^s} - \bar{\mu}_j)^\top \Phi^\top, \quad (6)$$

in which $\mu_{\mathbb{P}_j^s}$ and $\bar{\mu}_j$ can be computed according to the empirical estimation shown in equation (2). Denote $x_{k \sim j}^s$ as the k -th sample belonging to the j -th class in the s -th domain, where $s \in \{1, 2, \dots, m\}$ and $j \in \{1, 2, \dots, C\}$. Let n_j^s denote the sample size of the j -th class from the s -th domain, we have:

$$\hat{\mu}_{\mathbb{P}_j^s} = \frac{1}{n_j^s} \sum_{k=1}^{n_j^s} \phi(x_{k \sim j}^s), \quad \hat{\bar{\mu}}_j = \frac{1}{m} \sum_{s=1}^m \hat{\mu}_{\mathbb{P}_j^s}, \quad (7)$$

where $k \sim j$ denotes the indices of examples in the j -th class.

Scatter of class prior-normalized marginal distributions

The scatter of each class-conditional distribution is estimated locally using the samples from that class. When the number of examples in each class is small, optimizing (5) can easily overfit the data. To further improve the estimation accuracy, we propose another regularization term which measures the scatter of class-prior normalized marginal distributions. The new regularization term is able to measure the global distance between all class-conditionals. In the s -th domain, the marginal distribution is defined as

$$\mathbb{P}^s(X) = \sum_{j=1}^C \mathbb{P}^s(X|Y=j) \mathbb{P}^s(Y=j). \quad (8)$$

If the class prior distribution $\mathbb{P}(Y)$ does not change across domains, and we can also find a feature representation that has an invariant class-conditional $\mathbb{P}(h(X)|Y)$ across source domains, we can say that $\mathbb{P}(h(X))$ is also domain-invariant,

but not vice versa. Nevertheless, searching for a representation that reduces the discrepancy between the marginal distributions can to some extent reduce the discrepancy of class conditional distributions, though the original purpose was to match marginal distributions only (Muandet, Balduzzi, and Schölkopf 2013; Ghifary et al. 2017). However, if the class prior changes across source domains, the above statements are no longer true. That is to say, even if the class conditionals are domain-invariant, the marginal distribution are not invariant because of the changes in $\mathbb{P}(Y)$. To mitigate this issue, we propose to match the class-prior normalized marginal distribution, which is defined as follows:

$$\mathbb{P}_N^s(X) = \sum_{j=1}^C \mathbb{P}^s(X|Y=j) \frac{1}{C} \quad (9)$$

It can be seen that the class-prior normalized marginal distribution enforces the same prior probability for each class. Therefore, the changes in the prior distribution across source domains are adjusted, which guarantees that the prior-normalized marginal distribution is domain-invariant when the class conditionals are invariant. By embedding the class prior-normalized marginal distribution into a Hilbert space, the scatter of the normalized marginal distribution across domains can be formulated as:

$$\Psi^{prior} = \frac{1}{m} \sum_{s=1}^m \|\bar{\mu}_N - \mu_{\mathbb{P}_N^s}\|_{\mathcal{H}}^2, \quad (10)$$

where $\mu_{\mathbb{P}_N^s} = E_{x \sim \mathbb{P}_N^s}[\phi(x)]$, and \mathbb{P}_N^s is the prior-normalized marginal distribution of the s -th domain. $\bar{\mu}_N = \frac{1}{m} \sum_{s=1}^m \mu_{\mathbb{P}_N^s}$ is the kernel mean of the class prior-normalized marginal distribution \mathbb{P}_N of all domains. To learn the domain-invariant representation, we apply the linear feature transformation \mathbf{W} to the above scatter, resulting in:

$$\begin{aligned} \Psi_{\mathbf{B}}^{prior} &= \text{Tr} \left(\frac{1}{m} \sum_{s=1}^m \mathbf{B}^\top \Phi(\bar{\mu}_N - \mu_{\mathbb{P}_N^s})(\bar{\mu}_N - \mu_{\mathbb{P}_N^s})^\top \Phi^\top \mathbf{B} \right) \\ &= \text{Tr}(\mathbf{B}^\top \mathbf{L} \mathbf{B}), \end{aligned} \quad (11)$$

where \mathbf{L} can be formulated as follows:

$$\mathbf{L} = \frac{1}{m} \sum_{s=1}^m \Phi(\bar{\mu}_N - \mu_{\mathbb{P}_N^s})(\bar{\mu}_N - \mu_{\mathbb{P}_N^s})^\top \Phi^\top. \quad (12)$$

$\mu_{\mathbb{P}_N^s}$ in (12) can be empirically estimated from the observations as:

$$\hat{\mu}_{\mathbb{P}_N^s} = \frac{1}{C} \sum_{j=1}^C \frac{1}{n_j^s} \sum_{k=1}^{n_j^s} \phi(x_{k \sim j}^s). \quad (13)$$

Note that if n_j^s are identical for all j , that is the classes are balanced, the class prior-normalized marginal distribution reduces to the empirical estimate of the original marginal distribution $\hat{\mu}_{\mathbb{P}^s} = \frac{1}{n^s} \sum_{k=1}^{n^s} \phi(x_k^s)$ adopted in (Muandet, Balduzzi, and Schölkopf 2013; Ghifary et al. 2017).

Preserving Discriminative Power In addition to the above proposed two domain-invariance regularization terms, we also consider extra terms to preserve the discriminativeness of the learned representation. There have been plenty of works in supervised dimension reduction in the *i.i.d.* case, and kernel Fisher discriminant analysis (Mika et al. 1999) is a representative method which has been used in domain generalization (Ghifary et al. 2017). Because the focus of our method is to better learn the domain-invariant representations, we incorporate kernel Fisher discriminant analysis for fair comparison to existing methods. Specifically, the examples with the same label should be similar and the examples with different labels should be well separated. These two constraints can be formulated as two regularization terms: within-class scatter and between-class scatter, which are briefly described as follows.

Between-class scatter:

$$\Psi_B^{between} = Tr(\mathbf{B}^\top \mathbf{P} \mathbf{B}), \quad (14)$$

where matrix \mathbf{P} can be computed as:

$$e\mathbf{q} : \mathbf{P} \mathbf{P} = \sum_{j=1}^C n_j \Phi(\mu_j - \bar{\mu}_b)(\mu_j - \bar{\mu}_b)^\top \Phi^\top, \quad (15)$$

and $n_j = \sum_{s=1}^m n_j^s$ denotes the number of examples in the j -th class from all domains. Note that μ_j and $\bar{\mu}_b$ can be empirically estimated as $\hat{\mu}_j = \frac{1}{n_j} \sum_{s=1}^m \sum_{k=1}^{n_j^s} \phi(x_{k \sim j}^s)$ and $\hat{\bar{\mu}}_b = \frac{1}{n} \sum_{j=1}^C n_j \hat{\mu}_j$.

Within-class scatter:

$$\Psi_B^{within} = Tr(\mathbf{B}^\top \mathbf{Q} \mathbf{B}), \quad (16)$$

where the matrix \mathbf{Q} can be computed as:

$$\mathbf{Q} = \sum_{j=1}^C \sum_{s=1}^m \sum_{k=1}^{n_j^s} \Phi(\phi(x_{k \sim j}^s) - \mu_j)(\phi(x_{k \sim j}^s) - \mu_j)^\top \Phi^\top. \quad (17)$$

Objective Function and Optimization In this subsection, we first formulate our objective function with the above regularization terms and then find the solutions by maximizing the objective function.

The proposed CIDG aims to learn an invariant feature transformation by solving the following optimization problem:

$$\arg \max_B \frac{\Psi_B^{between}}{\Psi_B^{con} + \Psi_B^{prior} + \Psi_B^{within}}. \quad (18)$$

The numerator enforces the distance between features in different classes to be large. The denominator aims to learn a conditional invariant feature representation and reduce the distance between features in the same class simultaneously.

Replace the scatters with equation, (5), (11), (14), (16) and introduce several trade-off parameters γ, α , the above objective function can be reformulated as follows:

$$\arg \max_B \frac{Tr(\mathbf{B}^\top \mathbf{P} \mathbf{B})}{Tr(\mathbf{B}^\top (\gamma \mathbf{H} + \alpha \mathbf{L} + \mathbf{Q}) \mathbf{B})}, \quad (19)$$

where $0 < \gamma, 0 < \alpha$ are trade-off parameters, which need to be selected according to the validation set.

Note that the above objective function is invariant when rescaling $\mathbf{B} \rightarrow \eta \mathbf{B}$, where η is a constant. Consequently, (19) can be reformulated as the following constrained optimization problem:

$$\begin{aligned} \arg \max_B \quad & Tr(\mathbf{B}^\top \mathbf{P} \mathbf{B}) \\ \text{s.t.} \quad & Tr(\mathbf{B}^\top (\gamma \mathbf{H} + \alpha \mathbf{L} + \mathbf{Q}) \mathbf{B}) = 1, \end{aligned} \quad (20)$$

which yields Lagrangian:

$$\begin{aligned} L(\mathbf{B}) = & Tr(\mathbf{B}^\top \mathbf{P} \mathbf{B}) \\ & - Tr((\mathbf{B}^\top (\gamma \mathbf{H} + \alpha \mathbf{L} + \mathbf{Q}) \mathbf{B} - \mathbf{I}_q) \Gamma), \end{aligned} \quad (21)$$

where \mathbf{I}_q is an identity matrix of dimension q and $\Gamma = \text{diag}(\lambda_1, \lambda_2, \dots, \lambda_q)$ is a diagonal matrix with the Lagrange multipliers aligned in the diagonal. Solving (21) by setting the derivative w.r.t. \mathbf{B} to be zero, we arrive at a standard eigenvalue decomposition problem:

$$\mathbf{P} \mathbf{B} = (\gamma \mathbf{H} + \alpha \mathbf{L} + \mathbf{Q}) \mathbf{B} \Gamma. \quad (22)$$

In practice, the term $(\gamma \mathbf{H} + \alpha \mathbf{L} + \mathbf{Q})$ is added by a small constant $\epsilon \mathbf{I}$ to get a more stable solution, becoming $(\gamma \mathbf{H} + \alpha \mathbf{L} + \mathbf{Q} + \epsilon \mathbf{I})$. We summarize the algorithm of our CIDG in Algorithm 1.

Algorithm 1 Conditional invariant domain generalization

Require: m source domains with datasets $S_D = \{D_s = \{x_i^s, y_i^s\}_{i=1}^{n^s}, s = \{1, 2, \dots, m\}\}$, trade-off parameters γ, α .

Ensure: Invariant feature transformation \mathbf{B}^* and corresponding eigenvalues Γ^*

- 1: Construct kernel matrix \mathbf{K} from data samples of all domains, $\mathbf{K}(i, j) = k(x_i, x_j), \forall x_i, x_j \in S_D$, and construct matrices $\mathbf{H}, \mathbf{L}, \mathbf{P}, \mathbf{Q}$ from equations (14), (5), (11), (16).
 - 2: Centering the kernel matrix $\mathbf{K} \leftarrow \mathbf{K} - \mathbf{1}_n \mathbf{K} - \mathbf{K} \mathbf{1}_n + \mathbf{1}_n \mathbf{K} \mathbf{1}_n$, where $n = \sum_{i=1}^m n^s$ and $\mathbf{1}_n \in \mathbb{R}^{n \times n}$ denotes a matrix with all entries equal to $\frac{1}{n}$.
 - 3: Solve the equation (22) to get the optimal feature transformation matrix \mathbf{B}^* and the corresponding eigenvalues Γ^* with the first q leading eigenvalues.
 - 4: When given a target domain with a set of data $D_t = \{x_i^t, y_i^t\}_{i=1}^{n^t}$, construct a kernel matrix \mathbf{K}^t with samples from source domains and samples from the target domain, $\mathbf{K}^t(i, j) = k(x_i, x_j), \forall x_i \in S_D, x_j \in D_t$. Then we apply the centering operation to $\mathbf{K}^t \leftarrow \mathbf{K}^t - \mathbf{1}_n \mathbf{K}^t - \mathbf{K}^t \mathbf{1}_{n^t} + \mathbf{1}_n \mathbf{K}^t \mathbf{1}_{n^t}$, where $\mathbf{1}_{n^t} \in \mathbb{R}^{n^t \times n^t}$ denotes a matrix with all entries equal to $\frac{1}{n^t}$.
 - 5: The learned feature matrix of the target domain can be computed as $\mathbf{X}^* = (\mathbf{K}^t)^\top \mathbf{B}^* (\Gamma^*)^{-\frac{1}{2}}$.
-

Experiments

In this section, we conduct experiments on one synthetic data and two real-world image classification datasets to

domain index	domain 1			domain 2			domain 3		
class index	1	2	3	1	2	3	1	2	3
x	(1,0.3)	(2, 0.3)	(3, 0.3)	(3.5, 0.3)	(4.5, 0.3)	(5.5, 0.3)	(8, 0.3)	(9.5, 0.3)	(10, 0.3)
y	(2,0.3)	(1, 0.3)	(2, 0.3)	(2.5, 0.3)	(1.5, 0.3)	(2.5, 0.3)	(2.5, 0.3)	(1.5, 0.3)	(2.5, 0.3)
# samples	30	20	30	20	60	40	40	40	40

Table 1: Details of the generated distributions of three domains.

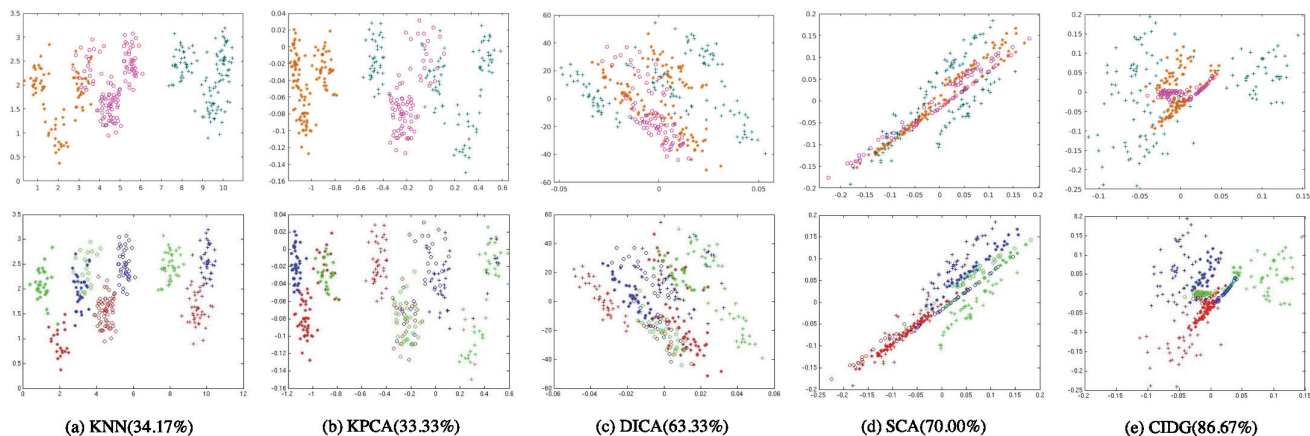


Figure 1: Performance comparison between different methods. The figures in the first row visualize the samples according to three different domains (yellow, magenta, cyan). The figures in the second row visualize the samples of three classes (green, red, blue) in different domains (star, circle, cross). Note that the left two domains (yellow, magenta) are source domains and the right one (cyan) is target domain.

demonstrate the effectiveness of our conditional invariant domain generalization (CIDG) method. The synthetic data are two dimensional, which facilitate the comparison of the performance of different methods through the visualization of the data distribution. The two real-world image classification datasets are the VLCS and Office+Caltech datasets, which are widely used datasets to evaluate the performance of domain generalization and domain adaptation (Ghifary et al. 2017; Gong et al. 2016; Khosla et al. 2012). We compare our CIDG with several state-of-the-art domain generalization methods, which are summarized below.

- K-nearest neighbors (KNN) using the original features, which serves as the baseline method.
- Kernel principal component analysis (KPCA) (Schölkopf, Smola, and Müller 1998) which finds the dominant components of the original features. KNN is applied for classification on the KPCA features.
- Undo-Bias (Khosla et al. 2012), which is a multi-task learning method aims to reduce the data bias. Because undo-bias is a binary classification algorithm, we use the one-vs-rest strategy for multi-class classification.
- Domain invariant component analysis (DICA) (Muandet, Balduzzi, and Schölkopf 2013), which is a domain generalization method learns an domain-invariant feature representation in terms of marginal distributions. We use KNN to do classification on the learned feature representation.

- Scatter component analysis (SCA) (Ghifary et al. 2017), which is a another method that learns domain-invariant features in terms of marginal distributions. The method incorporates discriminative terms and domain scatter terms into a unified framework.

Note that we have also conducted experiments using kernel finsher discriminant analysis (FDA), however, it performs worse than KPCA. Consequently, we do not report the results of KLDA in this paper.

Synthetic Dataset

In this section, we randomly generate two dimensional examples for source domains and target domain from different Gaussian distributions $\mathcal{N}(\mu, \sigma)$, where μ is the mean and σ is the standard deviation. The values of mean μ and standard deviation σ pairs (μ, σ) of different classes in three domains are shown in Table . We consider the first two domains as source domains and the third one as a target domain. The first row of Figure visualizes the samples from three different domains corresponding to three different colors (yellow, magenta, cyan), and the domains are domain 1, domain 2 and domain 3 from left to right. The second row of Figure shows that each domain has three clusters (green, red, blue) corresponding to three different classes and the domains are represented by different shapes (star, circle, cross). The first column illustrates the raw feature distributions.

We compare our CIDG with KNN, KPCA, DICA, and SCA to evaluate the distributions of the learned feature rep-

Source	Target	INN	KPCA	DICA	Undo-bias	SCA	CIDG
L,C,S	V	53.27 ± 1.52	58.62 ± 1.44	58.29 ± 1.51	57.73 ± 1.02	57.48 ± 1.78	65.65 ± 0.52
V,C,S	L	50.35 ± 0.94	53.80 ± 1.78	50.35 ± 1.45	58.16 ± 2.13	52.07 ± 0.86	60.43 ± 1.57
V,L,S	C	76.82 ± 1.56	85.84 ± 1.64	73.32 ± 4.13	82.18 ± 1.77	70.39 ± 1.42	91.12 ± 1.62
V,C,L	S	51.78 ± 2.07	53.23 ± 0.62	54.97 ± 0.61	55.02 ± 2.53	54.46 ± 2.71	60.85 ± 1.05
C,S	V,L	52.44 ± 1.87	55.74 ± 1.01	53.76 ± 0.96	56.83 ± 0.67	56.05 ± 0.98	59.25 ± 1.21
C,L	V,S	45.04 ± 2.49	45.13 ± 3.01	44.81 ± 1.62	52.16 ± 0.80	48.97 ± 1.04	54.04 ± 0.91
C,V	L,S	47.09 ± 2.49	55.79 ± 1.57	49.81 ± 1.40	59.00 ± 2.49	53.47 ± 0.71	61.61 ± 0.67
L,S	V,C	57.09 ± 1.43	58.50 ± 3.84	44.09 ± 0.58	51.16 ± 3.52	49.98 ± 1.84	55.65 ± 3.57
L,V	S,C	59.21 ± 1.84	63.88 ± 0.36	61.22 ± 0.95	64.26 ± 2.77	66.68 ± 1.09	70.89 ± 1.31
V,S	L,C	58.39 ± 0.78	64.56 ± 0.99	60.68 ± 1.36	68.58 ± 1.62	63.29 ± 1.34	70.44 ± 1.43

Table 2: Performance comparison between different methods with respect to accuracy (%) on VLCS dataset.

Source	Target	INN	KPCA	DICA	Undo-bias	SCA	CIDG
W,D,C	A	87.65 ± 2.46	90.92 ± 1.03	80.34 ± 2.65	89.56 ± 1.55	89.97 ± 1.85	93.24 ± 0.71
A,W,D	C	67.00 ± 0.67	74.23 ± 1.34	64.55 ± 2.85	82.27 ± 1.49	77.90 ± 1.28	85.07 ± 0.93
A,W,C	D	97.36 ± 1.92	94.34 ± 1.19	93.21 ± 1.92	95.28 ± 2.45	93.21 ± 3.50	97.36 ± 0.92
A,C,D	W	82.11 ± 0.67	88.84 ± 2.17	69.68 ± 3.22	90.18 ± 2.10	81.26 ± 3.15	90.53 ± 2.66
A,C	D,W	60.95 ± 1.31	75.81 ± 2.94	60.41 ± 1.94	80.24 ± 2.21	76.89 ± 0.99	83.65 ± 2.24
D,W	A,C	60.47 ± 0.99	65.75 ± 1.74	43.02 ± 3.24	74.14 ± 3.45	69.53 ± 1.87	65.91 ± 1.42
A,W	C,D	71.11 ± 0.81	76.26 ± 1.13	69.29 ± 1.77	81.77 ± 1.77	78.99 ± 1.54	83.89 ± 2.97
A,D	C,W	60.95 ± 1.31	75.81 ± 2.94	68.49 ± 2.88	81.23 ± 2.17	75.84 ± 1.66	84.66 ± 3.27
C,W	A,D	89.08 ± 2.26	91.45 ± 1.27	83.01 ± 2.42	91.73 ± 0.67	90.46 ± 1.72	93.41 ± 0.92
C,D	A,W	86.19 ± 1.58	90.36 ± 1.26	79.69 ± 1.11	90.67 ± 1.87	88.61 ± 0.38	91.70 ± 1.35

Table 3: Performance comparison between different methods with respect to accuracy (%) on office+caltech dataset.

resentation across domains. Since Undo-Bias is a SVM-based method that does not need to explicitly learn a feature representation, we do not compare the results with Undo-Bias on synthetic data. We use the RBF kernel for all the methods involving computation of kernel matrices. In all experiments, domain 1 and domain 2 are used as source domains and domain 3 is used as the unseen target domain. From the results in Figure , we can see that the proposed CIDG achieves the best accuracy of 86.67%. KPCA almost has no improvement over the baseline KNN method on the synthetic dataset. DICA can cluster one class (blue) well but performs badly for the other two classes. SCA can learn better feature distribution but the blue class and the green class are mixed in the learned representation. Additionally, the samples in the same class lie in a line rather than reside in a clear cluster. Our CIDG can learn more robust feature representations and the learned features in the same class are distributed in a well-shaped cluster.

VLCS Dataset

VLCS is an image classification dataset widely used for evaluating the performance of domain generalization. This dataset contains images from four different sub-datasets corresponding to four domains: PASCAL VOC2007 (V) (Everingham et al. 2010), LabelMe (L) (Russell et al. 2008), Caltech-101 (C) (Griffin, Holub, and Perona 2007), and SUN09 (S) (Choi et al. 2010). Five shared classes (bird, car, chair, dog and person) are selected from these four datasets.

The images are preprocessed by subtracting the mean values and cropped on the central 224×224 region out of the 256×256 resized images. Then the preprocessed images are fed into the DeCAF network and extracted the 4096 dimensional DeCAF₆ features (Donahue et al. 2014). We randomly select 70% of the data as training set from each domain and repeat the random selection five times. The mean classification accuracy and standard deviation of the five random selection are given for each method. All parameters are selected through validation, in which 30% of the training data is selected as validation set. All kernel methods use a RBF kernel and the learned features are classified using KNN except for Undo-Bias. The results are shown in Table .

From the results in Table , we can see that our conditional invariant domain generalization (CIDG) performs the best on 9 of the 10 domain generalization tasks. KPCA performs the best when L,S are source domains and V,C are target domains. Note that almost all the domain generalization methods outperform the INN on raw features. However, some methods on several domain tasks perform even worse than INN on raw features. This is mainly because that features of real world images are complicated and noisy. The learned features are not discriminative when generalized to target domains.

Office+Caltech Dataset

The Office+Caltech image dataset consists of ten overlapping categories between the Office dataset and the Caltech-256 dataset (C). Because the Office dataset contains three sub-datasets: AMAZON (A), DSLR (D), and WEBCAM (W), we have four different domains in total. Similarly, We randomly select 70% of the data as training set from each domain and repeat the random selection five times. The mean classification accuracy and standard deviation of the five random selection are reported for each method. The feature extraction is the same as that used for the VLCS dataset except we use the CAFFE network (Jia et al. 2014) instead of the DeCAF network. The other settings are the same as those in experiments on the VLCS dataset.

From the results in Table , we can find that the proposed CIDG achieves the best performance on 9 of the 10 domain generalization tasks. This further validates that enforcing conditional invariance is more reasonable than enforcing only marginal invariance. Note that Undo-bias is a SVM-based method. It is possibly the main reason why it outperforms CIDG when using D,W as source domains and A,C as target domains.

Conclusion

In this paper, we have proposed a conditional invariant domain generalization approach considering the situation that both $\mathbb{P}(X)$ and $\mathbb{P}(Y|X)$ change across domains. Different from previous works which assume that only $\mathbb{P}(X)$ changes, our proposed method can learn representations that have invariant joint distribution $\mathbb{P}(h(X), Y)$ across domains if the prior distribution $\mathbb{P}(Y)$ does not change between the source domains and the target domains. Two regularization terms that enforce class-conditional distribution invariance across domains are proposed and validated on both synthetic and real datasets.

Acknowledgments

This work was supported by National Key Research and Development Program of China 2017YFB1002203, NSFC No.61572451, No.61390514, and No. 61632019, Youth Innovation Promotion Association CAS CX2100060016, Fok Ying Tung Education Foundation WF2100060004, and Australian Research Council Projects FL-170100117, DP-180103424, DP-140102164, LP-150100671.

References

Baktashmotlagh, M.; Harandi, M.; Lovell, B.; and Salzmann, M. 2013. Unsupervised domain adaptation by domain invariant projection. In *Computer Vision (ICCV), 2013 IEEE International Conference on*, 769–776.

Blanchard, G.; Lee, G.; and Scott, C. 2011. Generalizing from several related classification tasks to a new unlabeled sample. In *Advances in neural information processing systems*, 2178–2186.

Choi, M. J.; Lim, J. J.; Torralba, A.; and Willsky, A. S. 2010. Exploiting hierarchical context on a large database of object categories. In *Computer vision and pattern recognition (CVPR), 2010 IEEE conference on*, 129–136. IEEE.

Donahue, J.; Jia, Y.; Vinyals, O.; Hoffman, J.; Zhang, N.; Tzeng, E.; and Darrell, T. 2014. Decaf: A deep convolutional activation feature for generic visual recognition. In *International conference on machine learning*, 647–655.

Duan, L.; Tsang, I. W.; Xu, D.; and Chua, T.-S. 2009. Domain adaptation from multiple sources via auxiliary classifiers. In *Proceedings of the 26th Annual International Conference on Machine Learning*, 289–296. ACM.

Erfani, S. M.; Baktashmotlagh, M.; Moshtaghi, M.; Nguyen, V.; Leckie, C.; Bailey, J.; and Ramamohanarao, K. 2016. Robust domain generalisation by enforcing distribution invariance. In *IJCAI*, 1455–1461.

Everingham, M.; Van Gool, L.; Williams, C. K.; Winn, J.; and Zisserman, A. 2010. The pascal visual object classes (voc) challenge. *International journal of computer vision* 88(2):303–338.

Ghifary, M.; Bastiaan Kleijn, W.; Zhang, M.; and Balduzzi, D. 2015. Domain generalization for object recognition with multi-task autoencoders. In *Proceedings of the IEEE international conference on computer vision*, 2551–2559.

Ghifary, M.; Balduzzi, D.; Kleijn, W. B.; and Zhang, M. 2017. Scatter component analysis: A unified framework for domain adaptation and domain generalization. *IEEE transactions on pattern analysis and machine intelligence* 39(7):1414–1430.

Gong, M.; Zhang, K.; Liu, T.; Tao, D.; Glymour, C.; and Schölkopf, B. 2016. Domain adaptation with conditional transferable components. In *International Conference on Machine Learning*, 2839–2848.

Griffin, G.; Holub, A.; and Perona, P. 2007. Caltech-256 object category dataset.

Huang, J.; Smola, A.; Gretton, A.; Borgwardt, K.; and Schölkopf, B. 2007. Correcting sample selection bias by unlabeled data. In *NIPS 19*, 601–608.

Janzing, D., and Schölkopf, B. 2010. Causal inference using the algorithmic markov condition. *IEEE Transactions on Information Theory* 56(10):5168–5194.

Jia, Y.; Shelhamer, E.; Donahue, J.; Karayev, S.; Long, J.; Girshick, R.; Guadarrama, S.; and Darrell, T. 2014. Caffe: Convolutional architecture for fast feature embedding. In *Proceedings of the 22nd ACM international conference on Multimedia*, 675–678. ACM.

Khosla, A.; Zhou, T.; Malisiewicz, T.; Efros, A. A.; and Torralba, A. 2012. Undoing the damage of dataset bias. In *European Conference on Computer Vision*, 158–171. Springer.

Liu, T.; Yang, Q.; and Tao, D. 2017. Understanding how feature structure transfers in transfer learning. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence*, 2365–2371.

Long, M.; Zhu, H.; Wang, J.; and Jordan, M. I. 2017. Deep transfer learning with joint adaptation networks. In Precup, D., and Teh, Y. W., eds., *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, 2208–2217. International Convention Centre, Sydney, Australia: PMLR.

- Luo, Y.; Wen, Y.; Liu, T.; and Tao, D. 2017. General heterogeneous transfer distance metric learning via knowledge fragments transfer.
- Mika, S.; Ratsch, G.; Weston, J.; Schölkopf, B.; and Müllers, K.-R. 1999. Fisher discriminant analysis with kernels. In *Neural Networks for Signal Processing IX, 1999. Proceedings of the 1999 IEEE Signal Processing Society Workshop.*, 41–48. IEEE.
- Muandet, K.; Balduzzi, D.; and Schölkopf, B. 2013. Domain generalization via invariant feature representation. In *Proceedings of the 30th International Conference on Machine Learning (ICML-13)*, 10–18.
- Pan, S. J.; Tsang, I. W.; Kwok, J. T.; and Yang, Q. 2011. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks* 22:199–220.
- Russell, B. C.; Torralba, A.; Murphy, K. P.; and Freeman, W. T. 2008. Labelme: a database and web-based tool for image annotation. *International journal of computer vision* 77(1):157–173.
- Schölkopf, B.; Janzing, D.; Peters, J.; Sgouritsa, E.; Zhang, K.; and Mooij, J. 2012. On causal and anticausal learning. *arXiv preprint arXiv:1206.6471*.
- Schölkopf, B.; Smola, A.; and Müller, K.-R. 1998. Non-linear component analysis as a kernel eigenvalue problem. *Neural computation* 10(5):1299–1319.
- Shao, M.; Ding, Z.; Zhao, H.; and Fu, Y. 2016. Spectral bisection tree guided deep adaptive exemplar autoencoder for unsupervised domain adaptation. In *AAAI, 2023–2029*.
- Shao, M.; Kit, D.; and Fu, Y. 2014. Generalized transfer subspace learning through low-rank constraint. *International Journal of Computer Vision* 109(1-2):74–93.
- Song, L.; Fukumizu, K.; and Gretton, A. 2013. Kernel embeddings of conditional distributions: A unified kernel framework for nonparametric inference in graphical models. *IEEE Signal Processing Magazine* 30(4):98–111.
- Sriperumbudur, B. K.; Gretton, A.; Fukumizu, K.; Schölkopf, B.; and Lanckriet, G. R. 2010. Hilbert space embeddings and metrics on probability measures. *Journal of Machine Learning Research* 11(Apr):1517–1561.
- Torralba, A., and Efros, A. A. 2011. Unbiased look at dataset bias. In *Computer Vision and Pattern Recognition (CVPR), 2011 IEEE Conference on*, 1521–1528. IEEE.
- Xu, Z.; Li, W.; Niu, L.; and Xu, D. 2014. Exploiting low-rank structure from latent domains for domain generalization. In *European Conference on Computer Vision*, 628–643. Springer.
- Yang, X.; Wang, M.; Hong, R.; Tian, Q.; and Rui, Y. 2017. Enhancing person re-identification in a self-trained subspace. *arXiv preprint arXiv:1704.06020*.
- Zhang, K.; Schölkopf, B.; Muandet, K.; and Wang, Z. 2013. Domain adaptation under target and conditional shift. In *International Conference on Machine Learning*, 819–827.