

Towards Accurate and Robust Domain Adaptation under Noisy Environments

Zhongyi Han¹, Xian-Jin Gui², Chaoran Cui^{3*} and Yilong Yin^{1*}

¹School of Software, Shandong University, Jinan 250101. China

²National Key Laboratory for Novel Software Technology, Nanjing University, Nanjing 210023, China

³School of Computer Science and Technology, Shandong University of Finance and Economics
hanzhongyicn@gmail.com, guixj@lamda.nju.edu.cn, crcui@sdufe.edu.cn, ylyin@sdu.edu.cn

Abstract

In non-stationary environments, learning machines usually confront the domain adaptation scenario where the data distribution does change over time. Previous domain adaptation works have achieved great success in theory and practice. However, they always lose robustness in noisy environments where the labels and features of examples from the source domain become corrupted. In this paper, we report our attempt towards achieving accurate noise-robust domain adaptation. We first give a theoretical analysis that reveals how harmful noises influence unsupervised domain adaptation. To eliminate the effect of label noise, we propose an offline curriculum learning for minimizing a newly-defined empirical source risk. To reduce the impact of feature noise, we propose a proxy distribution based margin discrepancy. We seamlessly transform our methods into an adversarial network that performs efficient joint optimization for them, successfully mitigating the negative influence from both data corruption and distribution shift. A series of empirical studies show that our algorithm remarkably outperforms state of the art, over 10% accuracy improvements in some domain adaptation tasks under noisy environments.

1 Introduction

Conventional learning theories assume that the learning machines are under a static environment where we draw training and test examples from an identical distribution. If the training and test distributions substantially differ, the learning machines would probably lose generalization ability [Valiant, 1984]. However, we expect our learning machines can perform well in crucial non-stationary environments where the test domain is similar yet distinct from the training domain. Unsupervised domain adaptation is the key machine learning topic to deal with this problem [Ben-David *et al.*, 2007].

Prominent theoretical advance and effective algorithm have been achieved in domain adaptation. Ben-David *et al.* [2007; 2010a] propose a pioneering $\mathcal{H}\Delta\mathcal{H}$ distance and give rigorous generalization bounds to measure the distribution shift.

A series of seminal studies then extend the $\mathcal{H}\Delta\mathcal{H}$ distance to discrepancy distance [Mansour *et al.*, 2009], to domain disagreement distance [Germain *et al.*, 2013], to generalized discrepancy distance [Cortes *et al.*, 2019], or to margin disparity discrepancy [Zhang *et al.*, 2019], etc. With significant theoretical findings, domain adaptation algorithms have made significant advances. Previous studies have explored various techniques concerning statistic moment matching-based algorithms [Pan *et al.*, 2010; Long *et al.*, 2017], gradual transition-based algorithms [Gopalan *et al.*, 2011], and pseudo labeling-based algorithms [Sener *et al.*, 2016; Saito *et al.*, 2017]. More interestingly, adversarial learning-based algorithms introducing a domain discriminator for minimizing distribution discrepancy, yields state of the art performance on many visual tasks [Ganin and Lempitsky, 2015; Tzeng *et al.*, 2017; Long *et al.*, 2018].

While previous works have achieved significant successes, they easily fail in realistic scenarios. The reason is that they implicitly assume an ideal learning environment where the source data are noise-free, which is difficult to hold in practice. *Domain adaptation in noisy environments* relaxes the assumption of clean data in standard domain adaptation to more realistic scenarios, where labels and features of examples from source domain are noisy. However, this more general yet challenging problem is under-explored so far. There is only one pioneering work [Shu *et al.*, 2019] that proposes an online curriculum learning algorithm to handle this problem. Although this work makes some progress, how to conduct theoretical analysis and construct a robust algorithm for unsupervised domain adaptation under noisy environments is still an open problem.

In this paper, we give a theoretical analysis of unsupervised domain adaptation when trained with noisy labels and features. Our theoretical analysis reveals that label noise worsens the expected risk of the source distribution; thus, we define a conditional-weighted empirical risk and propose an offline curriculum learning, which also provides some effective remedial measures for the online curriculum learning [Shu *et al.*, 2019]. Meanwhile, our theoretical analysis reveals that feature noise mainly aggravates the distribution discrepancy between source and target distribution; thus, we propose a novel proxy margin discrepancy by introducing a proxy distribution to provide an optimal solution for distribution discrepancy minimization. We finally transform our methods into an adversarial network that performs dynamic joint optimization between them. A

*Co-corresponding author

series of empirical studies on synthetic and real datasets show that our algorithm remarkably outperforms previous methods.

2 Domain Adaptation in Noisy Environments

2.1 Learning Set-Up

In domain adaptation, a sample of m labeled training examples $\{(x_i, y_i)\}_{i=1}^m$ is drawn according to a *source distribution* Q defined on $\mathcal{X} \times \mathcal{Y}$, where \mathcal{X} is the feature set and \mathcal{Y} is the label set. \mathcal{Y} is $\{1, \dots, K\}$ in multi-class classification. Meanwhile, a sample of n unlabeled test examples $\{(x_i)\}_{i=1}^n$ is drawn according to a *target distribution* P that somewhat differs from Q . Assume that the distributions Q and P should not be dissimilar substantially [Ben-David *et al.*, 2010b].

For domain adaptation in noisy environments, we relax the assumption of clean data in standard domain adaptation to that source distribution may be corrupted with feature noise and label noise independently. The source distribution is corrupted into a *noisy source distribution* Q_n . We denote by \tilde{x} the noisy features, and p_f the probability of corrupting one feature with harmful noises e , *i.e.*, $p(\tilde{x}_i = x_i + e) = p_f$ and $p(\tilde{x}_i = x_i) = 1 - p_f$. We denote by \tilde{y} the noise labels corrupted with a noise transition matrix $\mathcal{T} \in \mathbb{R}^{K \times K}$ where $\mathcal{T}_{ij} = p(\tilde{y} = j | y = i)$ denotes the probability of labeling an i -th class example as j .

We denote by $L : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}$ a loss function defined over pairs of labels. For multi-class classification, we denote by $f : \mathcal{X} \rightarrow \mathbb{R}^K$ a scoring function, which induces a labeling function $h_f : \mathcal{X} \rightarrow \mathcal{Y}$ where $h_f : x \rightarrow \arg \max_{y \in \mathcal{Y}} [f(x)]_y$. For any distribution Q on $\mathcal{X} \times \mathcal{Y}$ and any labeling function $h \in \mathcal{H}$, we denote $\epsilon_Q(h) = \mathbb{E}_{(x,y) \sim Q} L(h(x), y)$ the expected risk. Our objective is to select a hypothesis f out of a hypothesis set \mathcal{F} with a small expected risk $\epsilon_P(h_f)$ on the target distribution.

2.2 Theoretical Analysis

For in-depth analysis, we derive an upper bound of expected target risk in noisy environments according to the triangle inequality. Full proofs will be provided in the longer version.

Proposition 2.1. For any hypothesis $h \in \mathcal{H}$, the bound of target expected risk $\epsilon_P(h_f)$ in noisy environments is given by

$$\epsilon_P(h) \leq \epsilon_{Q_n}(h) + |\epsilon_{Q_n}(h, h^*) - \epsilon_P(h, h^*)| + \lambda, \quad (1)$$

where $\lambda = \epsilon_{Q_n}(h^*) + \epsilon_P(h^*)$ is the ideal combined error of h^* :

$$h^* = \arg \min_{h \in \mathcal{H}} \epsilon_P(h) + \epsilon_{Q_n}(h). \quad (2)$$

Proof Sketch: According to the triangle inequality, we have

$$\begin{aligned} \epsilon_P(h) &\leq \epsilon_P(h^*) + \epsilon_P(h, h^*) \\ &= \epsilon_P(h^*) + \epsilon_{Q_n}(h, h^*) + \epsilon_P(h, h^*) - \epsilon_{Q_n}(h, h^*) \\ &\leq \epsilon_P(h^*) + \epsilon_{Q_n}(h, h^*) + |\epsilon_P(h, h^*) - \epsilon_{Q_n}(h, h^*)| \\ &\leq \epsilon_{Q_n}(h) + |\epsilon_{Q_n}(h, h^*) - \epsilon_P(h, h^*)| \\ &\quad + [\epsilon_{Q_n}(h^*) + \epsilon_P(h^*)] \end{aligned} \quad (3)$$

This bound is informative if the three terms are small enough; however, noises do worsen all of them as follows.

The first term, expected risk of the noisy source distribution, is given by $\epsilon_{Q_n}(h) = \mathbb{E}_{(\tilde{x}, \tilde{y}) \sim Q_n} L(h(\tilde{x}), \tilde{y})$. Label noise

worsens the generalization ability of h in the first term, and the more label noises, the worse generalization ability. However, it is uncertain how feature noise behaves. Ilyas *et al.* [2019] concluded that deep neural networks (DNN) trained with adversarial perturbations might have better generalization and transfer-ability. Our empirical studies also find little difference between reserving and eliminating feature-noise examples when minimizing empirical source risk using DNNs.

The second term represents the measure of distribution discrepancy between the noisy source distribution and target distribution. Feature noise is dependent of it, but label noise is not. The following theorem reveals the dependence.

Theorem 2.1. Assume the source distribution Q be corrupted by feature noise into a new distribution R while target distribution P is unaffected. Assume feature noise is harmful and independent of the source and target distributions. Feature noise will enlarge the distribution discrepancy (*disc*) between source and target distribution, *i.e.* $\text{disc}(Q, P) \leq \text{disc}(R, P)$.

Proof Sketch: Given distributions Q_X and P_X over feature space \mathcal{X} , assuming \mathcal{H} denotes a symmetric hypothesis class, the \mathcal{H} -divergence [Ben-David *et al.*, 2007] is derived into

$$d_{\mathcal{H}}(Q, P) = 2 \left(1 - \inf_{h \in \mathcal{H}} \left(\Pr_{x \sim Q_X} [h(x) = 0] + \Pr_{x \sim P_X} [h(x) = 1] \right) \right). \quad (4)$$

When injecting harmful feature noise into source examples, the infimum would become lower; thus $d_{\mathcal{H}}(Q, P) \leq d_{\mathcal{H}}(R, P)$.

The third term λ is assumed to be small enough [Ben-David *et al.*, 2010b]; however, if feature-noise and label-noise are heavy, the ideal hypothesis error of target distribution will be unbounded, and this bound would be uninformative.

In summary, the above analysis suggests two critical aspects of achieving robust domain adaptation. Firstly, we should achieve a *robust empirical risk minimization* over the noisy source distribution by mitigating the negative influence of label noise. Secondly, we should concern a *robust distribution discrepancy* by reducing the impact of feature noise.

3 The Proposed Methods

Based on our analysis, we present the offline curriculum learning for the robust empirical risk minimization (Sec. 3.1), and then introduce the robust proxy margin discrepancy (Sec. 3.2).

3.1 Offline Curriculum Learning

To improve the empirical source risk, a natural idea is to eliminate label noise examples to recover a clean distribution. The pioneering work [Shu *et al.*, 2019] adopts an online curriculum learning that defines a loss value threshold γ to eliminate noisy examples in each training epoch $t \in T$ by

$$\min_f \frac{1}{m} \sum_{i=1}^m \mathbf{w}_i^t L(f(\tilde{x}_i; \theta_f^t), \tilde{y}_i), \text{ where} \quad (5)$$

$$\mathbf{w}_i^t = \mathbf{1}(\ell_i^t \leq \gamma), \quad i = 1, \dots, m, \quad t = 1, \dots, T, \quad (6)$$

where \mathbf{w}_i^t indicates whether or not to keep the i -th example, and θ_f^t is the parameter of hypothesis f in the t -th epoch. $\ell_i^t = L(f(\tilde{x}_i; \theta_f^t), \tilde{y}_i)$ minimized by stochastic gradient descent (SGD). Symbol $\mathbf{1}$ is the indicator function. γ is set

according to cross-validation. The reason behind its success is the *small loss* criterion that reveals clean examples have smaller loss values than noisy examples [Han *et al.*, 2018].

While online curriculum learning is intuitive, there are still four serious issues. *Firstly*, when source data are class-imbalanced, the magnitude of loss values will exist substantial differences across different classes; online curriculum learning will cause biased estimation. *Secondly*, it is in the expectation that correct examples will have smaller losses than incorrect ones in every epoch. However, the SGD algorithm makes the loss value of each example fluctuate in different epochs, which would cause the online curriculum learning unstable and unreliable to some extent. *Thirdly*, setting a fixed threshold of γ by cross-validation, is unpractical in domain adaptation. We cannot access the annotated data of the target domain. It is also unreasonable because it will lead DNNs to overfit on label-noise examples when the loss values of label-noise examples drop gradually with the increasing of the epoch. *Finally*, online curriculum learning treats feature and label noise examples indiscriminately and removing all of them in the whole course. However, eliminating feature-noise examples is needless when minimizing empirical source risk according to our analysis because lots of valuable features would be wasted in such a way. To resolve the above issues, we define the robust conditional-weighted empirical risk.

Definition 3.1 (Conditional-Weighted Empirical Risk). Denote by T the epoch number to filter out noisy examples, and denote by γ_k the loss value threshold of class $k \in K$, we define the conditional-weighted empirical risk over noisy source distribution for any classifier $f \in \mathcal{F}$ by

$$\epsilon_{\hat{Q}_n}(f) = \frac{1}{m} \sum_{i=1}^m w(\tilde{x}_i, \tilde{y}_i) L(f(\tilde{x}_i), \tilde{y}_i), \text{ where } \quad (7)$$

$$w(\tilde{x}_i, \tilde{y}_i) = \mathbb{1}(\bar{\ell}_i \leq \gamma_{\tilde{y}_i}), \quad i = 1, \dots, m, \text{ and } \quad (8)$$

$$\bar{\ell}_i = \frac{1}{T} \sum_{t=1}^T L(f(\tilde{x}_i; \theta_f^t), \tilde{y}_i). \quad (9)$$

To optimize this risk, we propose the offline curriculum learning, which consists of two steps. In the first step, we construct the early training curriculum to filter out noisy examples. In short, we first learn a classifier f and collect examples' loss values in T epochs. We then average each example's loss value (in Eq. (9)) and rank them by class in ascending order. We use the loss value of the $(m_k \times p_k)$ -th example as the loss value threshold γ_k of class k . m_k is the example number of class k , and $p_k = 1 - r_k$ in which r_k denotes the label-noise rate of class k . In practice, we find it is safer to set $p_k = \max\{1 - 1.2r_k, 0.8(1 - r_k)\}$. Finally, we conduct an offline exam by comparing each example's averaged loss value of $\bar{\ell}_i$ with $\gamma_{\tilde{y}_i}$ to decide each example's weight $w(\tilde{x}_i, \tilde{y}_i)$.

In the second step, we minimize Eq. (7) using the trusted examples ($w(\tilde{x}_i, \tilde{y}_i) = 1$) and use them to do domain adaptation. We choose cross-entropy loss, which can be optimized by SGD efficiently. Denote by σ the softmax function, *i.e.*, for $z \in \mathbb{R}^K$, $\sigma_k(z) \triangleq \frac{e^{z_k}}{\sum_{k=1}^K e^{z_k}}$, Eq. (7) is optimized by

$$\min_f -\mathbb{E}_{(\tilde{x}, \tilde{y}) \sim \hat{Q}_n} w(\tilde{x}_i, \tilde{y}_i) \log[\sigma_{\tilde{y}}(f(\tilde{x}))]. \quad (10)$$

Offline curriculum learning (OCL) provides point-to-point remedial measures for online curriculum learning. Firstly, OCL ranks and selects small-loss examples class by class, fully considering the class-conditional information and successfully achieving an unbiased estimation. Secondly, OCL examines the average loss value of each example along the whole curriculum process for comparison, avoiding the randomness of loss values in different epochs. Thirdly, OCL gets an efficient threshold of γ by leveraging the estimated label-noise rate, which can be estimated by existing methods [Patrini *et al.*, 2017]. We argue that we cannot expect to obtain a reliable classifier behind learning guarantees without known rough noise information. Finally, OCL only mitigates the negative influence of label noise but keeps valuable information on feature-noise examples to get better generalization. One might wonder how to separate label-noise and feature-noise examples. Our empirical studies point out that in the training process of DNNs with SGD, as for examples with the same noisy label, the ones with feature noises only will have smaller losses than the ones with false labels. Furthermore, the selected trusted source data by OCL directly supports shrinking the ideal combined hypothesis error of λ ; such establishes a valid learning guarantee for the expected target risk.

3.2 Proxy Margin Discrepancy

Recall that the measure of distribution discrepancy appeared as a critical term in the target error bound; thus, we especially propose the proxy margin discrepancy to mitigate feature noise.

Theorem 2.1 implies that a better guarantee would hold if we could select, instead of Q_n , other *proxy distribution* Q' with a smaller discrepancy $disc(Q', P)$. We thus introduce the proxy discrepancy using the Q' to give an optimal solution for distribution discrepancy minimization, by shrinking the discrepancy between source, proxy, and target distributions.

Definition 3.2 (Proxy Discrepancy, PD). Given a proxy distribution Q' , a noisy source distribution Q_n , and a target distribution P , for any measure of distribution discrepancy $disc$, we define the Proxy Discrepancy and its empirical version by

$$\begin{aligned} d_\Lambda(Q_n, P) &= disc(Q_n, Q') + disc(Q', P), \\ d_\Lambda(\hat{Q}_n, \hat{P}) &= disc(\hat{Q}_n, \hat{Q}') + disc(\hat{Q}', \hat{P}). \end{aligned} \quad (11)$$

In practice, we need to select an efficient $disc$. In the seminal works, Zhang *et al.* [2019] propose the novel margin disparity discrepancy (MDD) that embeds margin loss into a disparity discrepancy with informative generalization bound.

Definition 3.3 (Margin Disparity Discrepancy, MDD). Given a hypothesis set \mathcal{F} , and a specific classifier $f \in \mathcal{F}$, the MDD and its empirical version induced by $f' \in \mathcal{F}$ are defined by

$$\begin{aligned} d_{f, \mathcal{F}}^{(\rho)}(Q, P) &= \sup_{f' \in \mathcal{F}} \left(disp_P^{(\rho)}(f', f) - disp_Q^{(\rho)}(f', f) \right), \\ d_{f, \mathcal{F}}^{(\rho)}(\hat{Q}, \hat{P}) &= \sup_{f' \in \mathcal{F}} \left(disp_{\hat{P}}^{(\rho)}(f', f) - disp_{\hat{Q}}^{(\rho)}(f', f) \right), \end{aligned} \quad (12)$$

where $disp^{(\rho)}(f, f')$ is the margin disparity defined below. Let $\rho_{f'}$ denote the margin of a hypothesis f' and Φ_ρ denote ρ -margin loss, the margin disparity is defined by

$$disp_Q^{(\rho)}(f', f) = \mathbb{E}_{x \sim Q} \Phi_\rho(\rho_{f'}(x, h_f(x))). \quad (13)$$

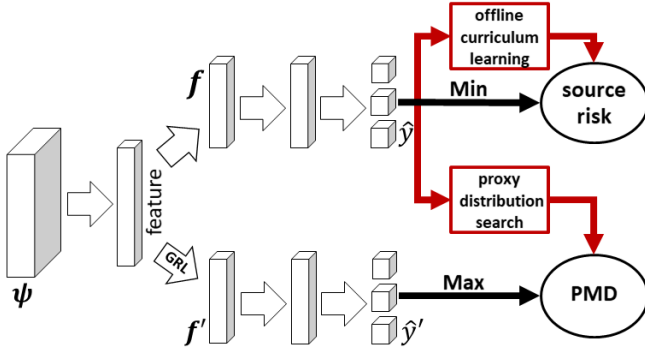


Figure 1: The adversarial network for robust domain adaptation.

Definition 3.4 (Proxy Margin Discrepancy, PMD). With the definition of MDD and a proxy distribution Q' , we define the Proxy Margin Discrepancy and its empirical version by

$$\begin{aligned} d_{\Lambda}^{(\rho)}(Q_n, P) &= d_{f, \mathcal{F}}^{(\rho)}(Q_n, Q') + d_{f, \mathcal{F}}^{(\rho)}(Q', P), \\ d_{\Lambda}^{(\rho)}(\hat{Q}_n, \hat{P}) &= d_{f, \mathcal{F}}^{(\rho)}(\hat{Q}_n, \hat{Q}') + d_{f, \mathcal{F}}^{(\rho)}(\hat{Q}', \hat{P}). \end{aligned} \quad (14)$$

To reveal the capacity of proxy distribution, we derive a generalization bound for PMD based on the bound of MDD.

Theorem 3.1. Based on the definition of Rademacher Complexity [Mansour *et al.*, 2009], for any $\delta > 0$, with probability $1 - 2\delta$ over a sample of size u drawn from Q' , a sample of size m drawn from Q , and a sample of size n drawn from P , the following holds simultaneously for any scoring function f :

$$\begin{aligned} & |d_{\Lambda}^{(\rho)}(\hat{Q}_n, \hat{P}) - d_{\Lambda}^{(\rho)}(Q_n, P)| \\ & \leq \frac{2K}{\rho} \mathfrak{R}_{m, Q_n}(\Pi_H^{\mathcal{F}}) + \frac{4K}{\rho} \mathfrak{R}_{u, Q'}(\Pi_H^{\mathcal{F}}) + \frac{2K}{\rho} \mathfrak{R}_{n, P}(\Pi_H^{\mathcal{F}}) \\ & \quad + \sqrt{\frac{\log \frac{2}{\delta}}{2m}} + 2\sqrt{\frac{\log \frac{2}{\delta}}{2u}} + \sqrt{\frac{\log \frac{2}{\delta}}{2n}}. \end{aligned} \quad (15)$$

where \mathfrak{R} denote Rademacher complexity, $\Pi_H^{\mathcal{F}}$ is defined as $\{x \rightarrow f(x, h(x)) | h \in H, f \in \mathcal{F}\}$, where h is induced by f .

To achieve effective optimization of PMD, we adopt a deep adversarial learning network using cross-entropy loss instead of the margin loss of Φ_{ρ} according to [Zhang *et al.*, 2019]. As shown in Fig. 1, it consists of a representation function ψ and two classifiers f and f' with coincide to Definition 3.3. Theorem 3.1 implies searching for an empirical proxy distribution with large size of u is crucial. To achieve that, we assume that \hat{Q}' is a better distribution that can be derived from noisy source distribution \hat{Q}_n , and let classifier f search for it by

$$\hat{Q}' = \arg \min_{\hat{Q}' \in \mathcal{Q}} -\mathbb{E}_{(\tilde{x}, y) \sim \hat{Q}'} \log[\sigma_y(f(\psi(\tilde{x})))] \quad (16)$$

where \mathcal{Q} denotes the set of distributions with support of trusted source examples by $\text{supp}(\mathbf{w}\hat{Q}_n)$. When the size of \hat{Q}' is zero, optimizing Eq. (16) is meaningless. Thus we limit the size of \hat{Q}' by $|\hat{Q}'| \geq \tau|\mathbf{w}\hat{Q}_n|$ where $\tau \in [0, 1]$ controls the ratio

Algorithm 1: Robust Domain Adaptation Algorithm.

input : parameters $\theta_{\psi}, \theta_f, \theta_{f'}$, learning rate η , fixed p_k , fixed τ , max epoch T , max iteration N_{max}
output : $\theta_{\psi}, \theta_f, \theta_{f'}$

- 1 /* stage 1: filter out label-noise examples */
- 2 **initialize** parameters θ_f
- 3 **for** $t = 1, 2, \dots, T$ **do**
- 4 **collect** loss values $[\ell_i^t]_{i=1}^m$ of source examples \hat{Q}_n
- 5 **update** $\theta_f = \theta_f - \eta \nabla \ell(f, \psi(\hat{Q}_n))$
- 6 **end**
- 7 **average** each example' loss value $\bar{\ell}_i = \frac{1}{T} \sum_{t=1}^T \ell_i^t$
- 8 **rank** average loss values by class in ascending order
- 9 **assign** $w = 1$ to the top $m_k \times p_k$ examples of class k
- 10 **obtain** trusted data $\hat{Q}(w = 1)$, untrusted $\hat{Q}''(w = 0)$
- 11 /* stage 2: conduct domain adaptation */
- 12 **initialize** parameters $\theta_{\psi}, \theta_f, \theta_{f'}, \tau' = 1$
- 13 **for** $n = 1, 2, \dots, N_{max}$ **do**
- 14 **fetch** batch \bar{Q} from \hat{Q} , \bar{P} from \hat{P}
- 15 **update** $\theta_f = \theta_f - \eta \nabla \ell(f, \psi(\bar{Q}'))$
- 16 **update** $\theta_{\psi} = \theta_{\psi} - \eta \nabla \ell(f, \psi(\bar{Q}))$
- 17 **obtain** $\bar{Q}' = \arg \min_{\bar{Q}': |\bar{Q}'| \geq \tau' |\bar{Q}|} \ell(f, \psi(\bar{Q}'))$
- 18 **update** $\theta_{f'} = \theta_{f'} - \eta \nabla d_{\Lambda}^{(\rho)}(\psi(\bar{Q}'), \psi(\bar{P}))$
- 19 **update** $\theta_{\psi} = \theta_{\psi} - \eta \nabla d_{\Lambda}^{(\rho)}(\psi(\bar{Q}'), \psi(\bar{P}))$
- 20 **update** $\tau' = \min\{\frac{n}{N_{max}}, \tau\}$
- 21 **end**

between \hat{Q}' and \hat{Q}_n . Since the early stage of DNNs is unstable, we design an incremental learning to gradually increase the size of \hat{Q}' . We let $\tau' = \min\{\frac{n}{N_{max}}, \tau\}$, instead of τ . Here n, N_{max} are the current and maximum iterations, respectively.

Meanwhile, since the optimization of PMD is a min-max game, we let the representation function ψ to minimize PMD and the classifier f' to maximize it. To avoid the problem of vanishing gradients, on the target domain, we use a modified cross-entropy loss: $\text{disp}_P(f', f) = \log[1 - \sigma_{f(\psi(x))}(f'(\psi(x)))]$. The final optimization of PMD is by

$$\begin{aligned} & \min_{\psi} \max_{f' \in \mathcal{F}} \alpha \mathbb{E}_{\tilde{x} \sim \psi(\hat{Q}_n)} \log[\sigma_{f(\psi(\tilde{x}))}(f'(\psi(\tilde{x})))] \\ & \quad - \mathbb{E}_{x \sim \psi(\hat{Q}')} \log[\sigma_{f(\psi(x))}(f'(\psi(x)))], \end{aligned} \quad (17)$$

and,

$$\begin{aligned} & \min_{\psi} \max_{f' \in \mathcal{F}} \alpha \mathbb{E}_{x \sim \psi(\hat{Q}')} \log[\sigma_{f(\psi(x))}(f'(\psi(x)))] \\ & \quad + \mathbb{E}_{x \sim \psi(\hat{P})} \log[1 - \sigma_{f(\psi(x))}(f'(\psi(x)))], \end{aligned} \quad (18)$$

where $\alpha \triangleq \exp \rho$ is designed to attain the margin ρ .

In summary, we give a novel perspective when facing feature noises in domain adaptation. Briefly speaking, PMD injects a proxy distribution to bridge noisy source distribution and target distribution, successfully alleviating the negative impact of feature noise. The theory-induced adversarial network performs proxy distribution search and disentangles the min-max game simultaneously, successfully learning domain-invariant representation in noise environments. The optimiza-

| Method | Label Corruption | | | | | | | Feature Corruption | | | | | | | Mixed Corruption | | | | | | |
|-------------|------------------|-------------|-------------|-------------|-------------|-------------|-------------|--------------------|-------------|-------------|-------------|-------------|-------------|-------------|------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | A→W | W→A | A→D | D→A | W→D | D→W | Avg | A→W | W→A | A→D | D→A | W→D | D→W | Avg | A→W | W→A | A→D | D→A | W→D | D→W | Avg |
| ResNet | 47.2 | 33.0 | 47.1 | 31.0 | 68.0 | 58.8 | 47.5 | 70.2 | 55.1 | 73.0 | 55.0 | 94.5 | 87.2 | 72.5 | 58.8 | 39.1 | 69.3 | 37.7 | 75.2 | 75.5 | 59.3 |
| SPL | 72.6 | 50.0 | 75.3 | 38.9 | 83.3 | 64.6 | 64.1 | 75.8 | 59.7 | 75.7 | 56.7 | 93.9 | 87.8 | 74.9 | 77.3 | 57.5 | 78.4 | 47.5 | 93.4 | 83.5 | 72.9 |
| MentorNet | 74.4 | 54.2 | 75.0 | 43.2 | 85.9 | 70.6 | 67.2 | 76.0 | 60.3 | 75.5 | 59.1 | 93.4 | 89.9 | 75.7 | 76.8 | 59.5 | 78.2 | 52.3 | 94.4 | 89.0 | 75.0 |
| DAN | 63.2 | 39.0 | 58.0 | 36.7 | 71.6 | 61.6 | 55.0 | 73.9 | 60.2 | 72.2 | 59.6 | 92.5 | 88.0 | 74.4 | 64.4 | 45.1 | 71.2 | 44.7 | 79.3 | 78.3 | 63.8 |
| RTN | 64.6 | 56.2 | 76.1 | 49.0 | 82.7 | 71.7 | 66.7 | 81.0 | 64.6 | 81.3 | 62.3 | 95.2 | 91.0 | 79.2 | 76.7 | 56.9 | 84.1 | 56.4 | 93.0 | 86.7 | 75.6 |
| DANN | 61.2 | 46.2 | 57.4 | 42.4 | 74.5 | 62.0 | 57.3 | 71.3 | 54.1 | 69.0 | 54.1 | 84.5 | 84.6 | 69.6 | 69.7 | 50.0 | 69.5 | 49.1 | 80.1 | 79.7 | 66.4 |
| ADDA | 61.5 | 49.2 | 61.2 | 45.5 | 74.7 | 65.1 | 59.5 | 76.8 | 62.0 | 79.8 | 60.1 | 93.7 | 89.3 | 77.0 | 69.7 | 54.5 | 72.4 | 56.0 | 87.5 | 85.5 | 70.9 |
| MDD | 74.7 | 55.1 | 76.7 | 54.3 | 89.2 | 81.6 | 71.9 | 92.9 | 66.8 | 88.0 | 70.9 | 99.8 | 96.6 | 85.8 | 88.7 | 63.1 | 81.9 | 68.5 | 94.6 | 89.3 | 81.0 |
| TCL | 82.0 | 65.7 | 83.3 | 60.5 | 90.8 | 77.2 | 76.6 | 84.9 | 62.3 | 83.7 | 64.0 | 93.4 | 91.3 | 79.9 | 87.4 | 64.6 | 83.1 | 62.2 | 99.0 | 92.7 | 81.5 |
| Ours | 89.7 | 67.2 | 92.0 | 65.5 | 96.0 | 92.7 | 83.6 | 95.1 | 68.4 | 89.4 | 72.4 | 99.8 | 97.8 | 87.2 | 93.1 | 69.5 | 92.0 | 71.5 | 99.0 | 93.1 | 86.4 |

Table 1: Accuracy (%) on **Office-31** with 40% corruption of Label, Feature, and Both.

| Method | Bing-Caltech | Office-Home | | | | | | | | | | | | | |
|-------------|--------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|-------------|--|
| | B→C | Ar→Cl | Ar→Pr | Ar→Rw | Cl→Ar | Cl→Pr | Cl→Rw | Pr→Ar | Pr→Cl | Pr→Rw | Rw→Ar | Rw→Cl | Rw→Pr | Avg | |
| ResNet | 74.4 | 27.1 | 50.7 | 61.7 | 41.1 | 53.8 | 56.3 | 40.9 | 28.0 | 61.8 | 51.3 | 33.0 | 65.9 | 47.6 | |
| SPL | 75.3 | 32.4 | 56.0 | 67.4 | 41.9 | 55.3 | 57.2 | 47.9 | 32.9 | 69.3 | 60.0 | 36.2 | 70.4 | 52.2 | |
| MentorNet | 75.6 | 34.5 | 57.1 | 66.7 | 43.3 | 56.1 | 57.6 | 48.5 | 34.0 | 70.2 | 59.8 | 37.2 | 70.4 | 53.0 | |
| DAN | 75.0 | 31.2 | 52.3 | 61.2 | 41.2 | 53.1 | 54.6 | 40.7 | 30.3 | 61.5 | 51.7 | 36.7 | 67.4 | 48.5 | |
| RTN | 75.8 | 29.3 | 57.8 | 66.3 | 44.0 | 58.6 | 58.3 | 46.0 | 30.1 | 67.5 | 56.3 | 32.2 | 69.9 | 51.4 | |
| DANN | 72.3 | 32.9 | 50.6 | 60.1 | 38.6 | 49.2 | 50.6 | 39.9 | 32.6 | 60.4 | 50.5 | 38.4 | 67.4 | 47.6 | |
| ADDA | 74.7 | 32.6 | 52.0 | 60.6 | 42.6 | 53.5 | 54.3 | 43.0 | 31.6 | 63.1 | 52.7 | 37.7 | 67.5 | 49.3 | |
| MDD | 78.9 | 44.6 | 62.4 | 68.8 | 46.7 | 58.9 | 60.8 | 45.5 | 39.5 | 65.2 | 59.8 | 47.1 | 72.9 | 56.0 | |
| TCL | 79.0 | 38.8 | 62.1 | 69.4 | 46.5 | 58.5 | 59.8 | 51.3 | 39.9 | 72.3 | 63.4 | 43.5 | 74.0 | 56.6 | |
| Ours | 81.7 | 50.8 | 68.7 | 72.3 | 55.6 | 67.4 | 67.9 | 57.8 | 50.5 | 74.6 | 69.5 | 57.7 | 80.2 | 64.4 | |

Table 2: Accuracy (%) on real dataset **Bing-Caltech** and **Office-Home** with 40% Mixed Corruption.

| Method | Office-31 40% Mixed Corruption | | | | | | |
|-------------|--------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | A→W | W→A | A→D | D→A | W→D | D→W | Avg |
| del-OCL | 90.4 | 62.1 | 84.3 | 68.1 | 96.2 | 89.7 | 81.8 |
| del-PMD | 92.2 | 68.8 | 89.2 | 69.7 | 98.8 | 92.7 | 85.2 |
| Ours | 93.1 | 69.5 | 92.0 | 71.5 | 99.0 | 93.1 | 86.4 |

Table 3: Ablation study by deleting OCL or PMD.

| Method | Office-31 40% Label Corruption | | | | | | |
|-------------|--------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | A→W | W→A | A→D | D→A | W→D | D→W | Avg |
| Online | 83.1 | 55.4 | 79.1 | 49.9 | 88.6 | 83.8 | 73.3 |
| Offline | 90.1 | 59.8 | 91.3 | 52.7 | 85.7 | 75.3 | 75.9 |
| Ours | 89.7 | 67.2 | 92.0 | 65.5 | 96.0 | 92.7 | 83.6 |

Table 4: An analysis of our offline curriculum learning.

tion process uses incremental learning to control the size of proxy distribution, providing a stable solution for PMD.

3.3 Robust Joint Optimization

Finally, we unify the learning with the robust conditional-weighted empirical risk (\mathcal{E}_w) and proxy margin discrepancy in a joint min-max problem. We state the final problem as

$$\begin{aligned} \min_{f, \psi} \mathcal{E}_w(\psi(\hat{Q}_n)) + \beta d_{\Lambda}^{(\rho)}(\psi(\hat{Q}_n), \psi(\hat{P})), \\ \max_{f'} d_{\Lambda}^{(\rho)}(\psi(\hat{Q}_n), \psi(\hat{P})), \end{aligned} \quad (19)$$

where β is the trade-off coefficient between source error and discrepancy. In practice, we split the optimization process of Eq. (19) into two stages: filter out label-noise examples and conduct domain adaptation, as shown in Algorithm 1.

| Method | Office-31 40% Feature Corruption | | | | | | |
|----------|----------------------------------|-------------|-------------|-------------|-------------|-------------|-------------|
| | A→W | W→A | A→D | D→A | W→D | D→W | Avg |
| TCL | 84.9 | 62.3 | 83.7 | 64.0 | 93.4 | 91.3 | 79.9 |
| Ours-del | 90.1 | 71.5 | 89.9 | 69.5 | 98.6 | 92.7 | 85.4 |
| Ours-res | 95.1 | 68.4 | 89.4 | 72.4 | 99.8 | 97.8 | 87.2 |

Table 5: An analysis of the feature-noise effect on source risk.

| Method | Wider Feature Noise Levels on A→W | | | | | |
|-------------|-----------------------------------|-------------|-------------|-------------|-------------|-------------|
| | 0% | 20% | 40% | 60% | 80% | Avg |
| DANN | 85.5 | 73.7 | 71.3 | 60.4 | 53.2 | 57.4 |
| TCL | 87.5 | 85.9 | 84.9 | 61.9 | 31.2 | 58.6 |
| Ours | 95.9 | 95.7 | 95.1 | 87.2 | 80.4 | 75.7 |

Table 6: An analysis of different discrepancies on feature noises.

4 Experiments

We evaluate our algorithm on three datasets against state of the art methods. The code is at <https://github.com/zhyhan/RDA>.

4.1 Setup

Office-Home is a more challenging domain adaptation dataset consisting of 15,599 images with 65 unbalanced classes. It consists of four more diverse domains: **Artistic images**, **Clip Art**, **Product images**, and **Real-world images**.

Following the protocol in the pioneering work [Shu *et al.*, 2019], we create corrupted counterparts on the above two clean datasets as follows. We make three types of corruption: label corruption, feature corruption, and mixed corruption on source domains. Label corruption uniformly changes the label of each image into a random class with probability p_{noise} . Feature corruption refers to each image corrupted by Gaussian blur and salt-and-pepper noise with probability p_{noise} . Mixed corruption refers to each image corrupted by label corruption and




Figure 2: An illustration of loss distributions of the three types noises.




Figure 3: An analysis of ten methods in various mixed noise levels.

feature corruption with probability $p_{noise}/2$ independently. While past literature mainly studied label corruption, we also study mixed corruption with distribution shifts.

Bing-Caltech is a real noisy dataset created with Bing and Caltech-256 datasets with 256 classes. The Bing dataset consists of rich mixed noises because it was collected by retrieving from the Bing image search engine using the category labels of Caltech-256. We naturally set Bing as a source domain, while Caltech-256 as the target domain.

We compare our designed Robust Domain Adaptation (RDA) algorithm with state of the art methods: **ResNet-50** [He *et al.*, 2016], Self-Paced Learning (SPL) [Kumar *et al.*, 2010], **MentorNet** [Jiang *et al.*, 2018], Deep Adaptation Network (DAN) [Long *et al.*, 2015], Residual Transfer Network (RTN) [Long *et al.*, 2016], Domain Adversarial Network (DANN) [Ganin and Lempitsky, 2015], Adversarial Discriminative Domain Adaptation (ADDA) [Tzeng *et al.*, 2017], Margin Disparity Discrepancy based algorithm (MDD) [Zhang *et al.*, 2019], and Transferable Curriculum Learning (TCL) [Shu *et al.*, 2019]. Note that TCL is the pioneering work with the same setup of datasets as ours.

We implement our algorithm in Pytorch. We use **ResNet-50** as the representation function with parameters pre-trained from ImageNet. The main classifier and adversarial classifier are both 2-layer neural networks. We set the early training epoch T to 30. α is set to 3 and β is set to 0.1 according

to [Zhang *et al.*, 2019]. We use mini-batch SGD with the Nesterov momentum 0.9. The initial learning rate of the classifiers f and f' is 0.001, which is ten times than of the representation function ψ . The set of r_k and τ depends on noise rates.

4.2 Results

Table 1 reports the results on Office-31 under 40% corruption of label, feature, and both. Our algorithm significantly outperforms existing domain adaptation methods and noise label methods (MentorNet, SPL) on almost all tasks. Note that the standard domain adaptation methods (DANN, MDD, etc.) suffer from over-fitting under noises while our algorithm performs stably positive adaptation, which demonstrates its robustness and generalization ability. While the pioneering work TCL uses additional entropy minimization to enhance its performance, our algorithm outperforms it by a large margin. Table 2 reports the results on the real dataset Bing-Caltech and the synthetic dataset Office-Home under 40% mixed corruption, where we make a remarkable performance boost.

4.3 Analysis

Ablation Study. Table 3 reports the results of our ablation study. Abandoning either the Offline Curriculum Learning (**del-OCL**) or the Proxy Margin Discrepancy (**del-PMD**) results in a significant decline, which demonstrates the efficacy and robustness of both methods in noisy environments.

Offline Curriculum Learning. We dissect the strengths of offline curriculum learning. Table 4 reports the results of three curriculum modes: online, offline, ours (Offline Curriculum Learning) on Office-31 under 40% label corruption. The online mode performs worse than the offline mode, which performs worse than us, proving the necessity of considering the averaged loss and class priors. Furthermore, we also justify our claim that we need to treat label-noise and feature-noise examples separately. As shown in Table 5, our method gets better performance than TCL that treats them indiscriminately. Table 5 also reports that reserving feature-noise examples (**ours-res**) generally gets better results than deleting them (**ours-del** and **TCL**) when optimizing empirical source risk.

Proxy Margin Discrepancy. While Table 1 has demonstrated the strengths of our proxy margin discrepancy of mitigating the negative influence from feature-noise examples, we provide a broader spectrum for more in-depth analysis. Table 6

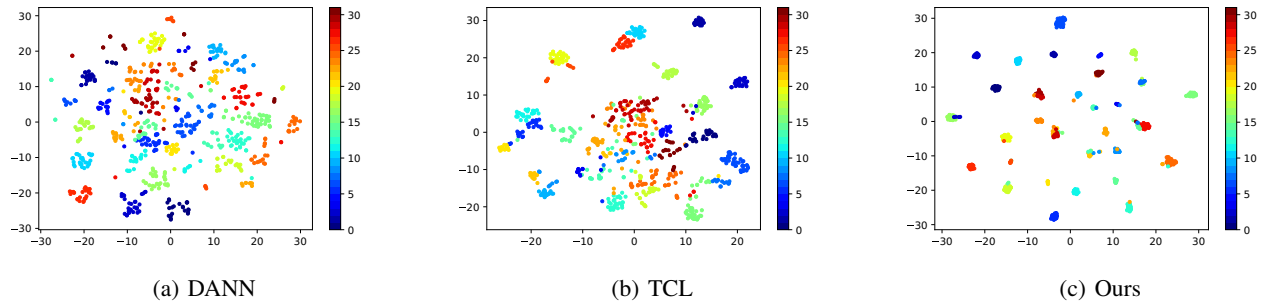


Figure 4: The t-SNE visualization of class features on target domain.

reports the results of representative methods on wider feature-noise levels. Our method maintains stable performance while DANN and TCL drop rapidly with the increase of noise levels.

Noise Levels. Fig. 3 reports the results in various mixed noise levels on $A \rightarrow W$ task. Our method outperforms all the compared methods at each level, which demonstrates that our method is more robust under severe noisy environments. In particular, our methods achieve the best result when the noise level is 0%, which proves that our method can also fit into the standard domain adaptation scenarios.

Loss Distribution. Fig. 2 exhibits the loss distributions of the three types of noises of 40% on the Office-31 Amazon dataset. Fig. 2 (a, b) verifies that correct examples have smaller losses than incorrect examples. Fig. 2 (c) verifies that feature-noise examples generally have smaller losses than label-noise examples, which fits our intuition that label noise is more harmful than feature noise.

Feature Visualization. Fig. 4 illustrates the t-SNE embeddings [Donahue *et al.*, 2014] of the learned representations by DANN, TCL, and ours, on 40% mixed corruption of $A \rightarrow W$ task. While the learned features of DANN, and TCL are mixed up, ours is more discriminative in every class on the target domain, which verifies that our method can learn domain-invariant representations in noise environments.

5 Conclusion

We presented a new analysis of domain adaptation in noisy environments, an under-explored but more realistic scenario. We also proposed an offline curriculum learning and a new proxy discrepancy tailored to label and feature noises, respectively. Our methods are more robust for achieving real-world applications in noisy non-stationary environments. The theory-induced algorithm yields the state of the art results in all tasks.

Acknowledgements

This work is supported by the National Natural Science Foundation of China (61876098, 61573219, 61701281), the National Key R&D Program of China (2018YFC0830100, 2018YFC0830102), and the Fostering Project of Dominant Discipline and Talent Team of Shandong Province Higher Education Institutions.

References

- [Ben-David *et al.*, 2007] Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. In *Advances in neural information processing systems*, pages 137–144, 2007.
- [Ben-David *et al.*, 2010a] 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.
- [Ben-David *et al.*, 2010b] Shai Ben-David, Tyler Lu, Teresa Luu, and Dávid Pál. Impossibility theorems for domain adaptation. In *International Conference on Artificial Intelligence and Statistics*, pages 129–136, 2010.
- [Cortes *et al.*, 2019] Corinna Cortes, Mehryar Mohri, and Andrés Muñoz Medina. Adaptation based on generalized discrepancy. *The Journal of Machine Learning Research*, 20(1):1–30, 2019.
- [Donahue *et al.*, 2014] Jeff Donahue, Yangqing Jia, Oriol Vinyals, Judy Hoffman, Ning Zhang, Eric Tzeng, and Trevor Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *Proceedings of the 31th International Conference on Machine Learning*, pages 647–655, 2014.
- [Ganin and Lempitsky, 2015] Yaroslav Ganin and Victor S. Lempitsky. Unsupervised domain adaptation by backpropagation. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 1180–1189, 2015.
- [Germain *et al.*, 2013] Pascal Germain, Amaury Habrard, François Laviolette, and Emilie Morvant. A pac-bayesian approach for domain adaptation with specialization to linear classifiers. In *Proceedings of the 30th International Conference on Machine Learning*, pages 738–746, 2013.
- [Gopalan *et al.*, 2011] Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *2011 international conference on computer vision*, pages 999–1006. IEEE, 2011.
- [Han *et al.*, 2018] Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi

- Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In *Advances in neural information processing systems*, pages 8527–8537, 2018.
- [He *et al.*, 2016] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [Ilyas *et al.*, 2019] Andrew Ilyas, Shibani Santurkar, Dimitris Tsipras, Logan Engstrom, Brandon Tran, and Aleksander Madry. Adversarial examples are not bugs, they are features. In *Advances in Neural Information Processing Systems*, pages 125–136, 2019.
- [Jiang *et al.*, 2018] Lu Jiang, Zhengyuan Zhou, Thomas Leung, Li-Jia Li, and Li Fei-Fei. Mentornet: Learning data-driven curriculum for very deep neural networks on corrupted labels. In *Proceedings of the 35th International Conference on Machine Learning*, pages 2309–2318, 2018.
- [Kumar *et al.*, 2010] M Pawan Kumar, Benjamin Packer, and Daphne Koller. Self-paced learning for latent variable models. In *Advances in Neural Information Processing Systems*, pages 1189–1197, 2010.
- [Long *et al.*, 2015] Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *Proceedings of the 32nd International Conference on Machine Learning*, pages 97–105, 2015.
- [Long *et al.*, 2016] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Unsupervised domain adaptation with residual transfer networks. In *Advances in Neural Information Processing Systems*, pages 136–144, 2016.
- [Long *et al.*, 2017] Mingsheng Long, Han Zhu, Jianmin Wang, and Michael I Jordan. Deep transfer learning with joint adaptation networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2208–2217, 2017.
- [Long *et al.*, 2018] 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.
- [Mansour *et al.*, 2009] Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. In *22nd Conference on Learning Theory, COLT*, 2009.
- [Pan *et al.*, 2010] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 22(2):199–210, 2010.
- [Patrini *et al.*, 2017] Giorgio Patrini, Alessandro Rozza, Aditya Krishna Menon, Richard Nock, and Lizhen Qu. Making deep neural networks robust to label noise: A loss correction approach. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 2233–2241, 2017.
- [Saito *et al.*, 2017] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning*, pages 2988–2997, 2017.
- [Sener *et al.*, 2016] Ozan Sener, Hyun Oh Song, Ashutosh Saxena, and Silvio Savarese. Learning transferrable representations for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems*, pages 2110–2118, 2016.
- [Shu *et al.*, 2019] Yang Shu, Zhangjie Cao, Mingsheng Long, and Jianmin Wang. Transferable curriculum for weakly-supervised domain adaptation. In *Thirty-Third AAAI Conference on Artificial Intelligence*, pages 4951–4958, 2019.
- [Tzeng *et al.*, 2017] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*, pages 7167–7176, 2017.
- [Valiant, 1984] Leslie G Valiant. A theory of the learnable. In *Proceedings of the sixteenth annual ACM symposium on Theory of computing*, pages 436–445. ACM, 1984.
- [Zhang *et al.*, 2019] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael I. Jordan. Bridging theory and algorithm for domain adaptation. In *Proceedings of the 36th International Conference on Machine Learning*, pages 7404–7413, 2019.