

# Unsupervised Domain Adaptation Based on Source-Guided Discrepancy\*

Seiichi Kuroki,<sup>†‡</sup> Nontawat Charoenphakdee,<sup>†‡</sup> Han Bao,<sup>†‡</sup>  
Junya Honda,<sup>†‡</sup> Issei Sato,<sup>†‡</sup> Masashi Sugiyama<sup>†‡</sup>

<sup>†</sup>The University of Tokyo, Tokyo, Japan

<sup>‡</sup>RIKEN, Tokyo, Japan

## Abstract

*Unsupervised domain adaptation* is the problem setting where data generating distributions in the source and target domains are different and labels in the target domain are unavailable. An important question in unsupervised domain adaptation is how to measure the difference between the source and target domains. Existing discrepancy measures for unsupervised domain adaptation either require high computation costs or have no theoretical guarantee. To mitigate these problems, this paper proposes a novel discrepancy measure called *source-guided discrepancy (S-disc)*, which exploits labels in the source domain unlike the existing ones. As a consequence, S-disc can be computed efficiently with a finite-sample convergence guarantee. In addition, it is shown that S-disc can provide a tighter generalization error bound than the one based on an existing discrepancy measure. Finally, experimental results demonstrate the advantages of S-disc over the existing discrepancy measures.

## Introduction

In the conventional supervised learning framework, we often assume that the training and test distributions are the same. However, this assumption may not hold in many practical applications such as natural language processing (Glorot, Bordes, and Bengio 2011), speech recognition (Sun et al. 2017), and computer vision (Saito, Ushiku, and Harada 2017). For instance, a personalized spam filter can be trained from the emails of all available users, but the training data may not represent the emails of the target user. Such scenarios can be formulated in the framework of *domain adaptation*, which has been studied extensively (Ben-David et al. 2007; Mansour, Mohri, and Rostamizadeh 2009a; Zhang, Zhang, and Ye 2012; Pan and Yang 2010; Sugiyama and Kawanabe 2012). An important challenge in domain adaptation is to find a classifier for a label-scarce target domain by exploiting a label-rich source domain. In particular, there are cases where we only have access to labeled data from the source domain and unlabeled data from the target domain, because annotating labels in the target domain is often time-consuming and expensive (Saito et al. 2018).

\*A longer version of this paper is available in Kuroki et al. (2018).

Copyright © 2019, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

This problem setting is called *unsupervised domain adaptation* (Ben-David et al. 2007), which is our interest in this paper.

Since domain adaptation cannot be performed effectively if the source and target domains are too different, an important topic to be addressed is how to measure the difference, or *discrepancy*, between the two domains. Many discrepancy measures have been used in previous studies such as the maximum mean discrepancy (Huang et al. 2007), Kullback-Leibler divergence (Sugiyama et al. 2008), Rényi divergence (Mansour, Mohri, and Rostamizadeh 2009b), and Wasserstein distance (Courty et al. 2017).

Apart from the discrepancy measures described above, Ben-David et al. (2007) proposed a discrepancy measure for binary classification that explicitly takes a hypothesis class into account. They showed that their discrepancy measure leads to a tighter generalization error bound than the  $L_1$  distance, which does not use information of the hypothesis class. Following this line of research, Mansour, Mohri, and Rostamizadeh (2009a) generalized the discrepancy measure of Ben-David et al. (2007) to arbitrary loss functions, and Germain et al. (2013) provided a PAC-Bayesian analysis based on the average disagreement of a set of hypotheses. These results provided a theoretical foundation for various applications of domain adaptation based on estimation of source-target discrepancy, including sentiment analysis (Glorot, Bordes, and Bengio 2011), source selection (Bhatt, Rajkumar, and Roy 2016), and multi-source domain adaptation (Zhao et al. 2017). However, their discrepancy measure considers the worst pair of hypotheses to bound the maximum gap in the loss between the domains, which may lead to a loose generalization error bound and expensive computation costs. Zhang, Zhang, and Ye (2012) and Mohri and Medina (2012) analyzed another discrepancy measure with a generalization error bound. Nevertheless, this discrepancy measure cannot be computed without labels in the target domain and therefore is not suitable for unsupervised domain adaptation.

To alleviate the limitations of the existing discrepancy measures for domain adaptation, we propose a novel discrepancy measure called *source-guided discrepancy (S-disc)*. By incorporating labels in the source domain, S-disc provides a tighter generalization error bound than the previously proposed discrepancy measure by Mansour, Mohri,

and Rostamizadeh (2009a). Furthermore, we provide an estimator of S-disc for binary classification that can be computed efficiently. We also establish consistency and analyze the convergence rate of the proposed S-disc estimator.

Our main contributions are as follows.

- We propose a novel discrepancy measure called *source-guided discrepancy (S-disc)*, which uses labels in the source domain to measure the difference between the two domains for unsupervised domain adaptation (Definition 2).
- We propose an efficient algorithm for estimating S-disc for the 0-1 loss (Algorithm 1).
- We show the consistency and elucidate the convergence rate of the estimator of S-disc (Theorem 4).
- We derive a generalization error bound in the target domain based on S-disc (Theorem 7), which is tighter than the existing bound provided by Mansour, Mohri, and Rostamizadeh (2009a). In addition, we derive a generalization error bound for a finite-sample case (Theorem 8).
- We demonstrate the effectiveness of S-disc for unsupervised domain adaptation in experiments.

## Problem Setting and Notation

In this section, we formulate the unsupervised domain adaptation problem. Let  $\mathcal{X}$  be the input space and  $\mathcal{Y}$  be the output space, which is  $\{+1, -1\}$  in binary classification. We define a domain as a pair  $(P_D, f_D)$ , where  $P_D$  is an input distribution on  $\mathcal{X}$  and  $f_D : \mathcal{X} \rightarrow \mathcal{Y}$  is a labeling function. In domain adaptation, we denote the source domain and the target domain as  $(P_S, f_S)$  and  $(P_T, f_T)$ , respectively. Unlike conventional supervised learning, we focus on the case where  $(P_T, f_T)$  differs from  $(P_S, f_S)$ . In unsupervised domain adaptation, we are given the following data:

- Unlabeled data in the target domain:  $\mathcal{T} = \{x_i^T\}_{i=1}^{n_T}$ .
- Labeled data in the source domain:  $\mathcal{S} = \{(x_j^S, y_j^S)\}_{j=1}^{n_S}$ .

For simplicity, we denote the input features from the source domain  $\{x_j^S\}_{j=1}^{n_S}$  as  $\mathcal{S}_{\mathcal{X}}$  and the empirical distribution corresponding to  $P_T$  (resp.  $P_S$ ) as  $\hat{P}_T$  (resp.  $\hat{P}_S$ ).

We denote a loss function  $\ell: \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$ . For example, the 0-1 loss is given as  $\ell_{01}(y, y') = (1 - \text{sign}(yy'))/2$ . The expected loss for functions  $h, h' : \mathcal{X} \rightarrow \mathbb{R}$  and distribution  $P_D$  over  $\mathcal{X}$  is denoted by  $R_D^\ell(h, h') = \mathbb{E}_{x \sim P_D}[\ell(h(x), h'(x))]$ . We also denote an empirical risk as  $\hat{R}_D^\ell(h, h') = \mathbb{E}_{x \sim \hat{P}_D}[\ell(h(x), h'(x))]$ . In addition, we define the true risk minimizer and the empirical risk minimizer in a certain domain  $(P_D, f_D)$  in a hypothesis class  $\mathcal{H}$  as  $h_D^* = \arg \min_{h \in \mathcal{H}} R_D^\ell(h, f_D)$  and  $\hat{h}_D = \arg \min_{h \in \mathcal{H}} \hat{R}_D^\ell(h, f_D)$ , respectively. Here, note that the risk minimizer  $h_D^*$  is not necessarily equal to the labeling function  $f_D$  as we consider a restricted hypothesis class.

The goal in domain adaptation is to find a hypothesis  $h$  out of a hypothesis class  $\mathcal{H}$  so that it gives as small expected loss as possible in the target domain defined by

$$R_T^\ell(h, f_T) = \mathbb{E}_{x \sim P_T}[\ell(h(x), f_T(x))].$$

## Related Work

In unsupervised domain adaptation, it is essential to measure the difference between the source and target domains, because it might degrade the performance to use data collected from a source domain that is far from the target domain (Pan and Yang 2010). Since we cannot access labels from the target domain, it is impossible to measure the difference between the two domains on the basis of the *output space*. One way is to measure the difference between probability distributions in terms of the *input space*.

We will review such discrepancy measures in this section. First, we introduce the discrepancy measure proposed by Mansour, Mohri, and Rostamizadeh (2009a), which is defined as

$$\text{disc}_{\mathcal{H}}^\ell(P_T, P_S) = \sup_{h, h' \in \mathcal{H}} |R_T^\ell(h, h') - R_S^\ell(h, h')|. \quad (1)$$

We call this discrepancy measure  $\mathcal{X}$ -disc as it does not require labels from the source domain but input features. Note that  $\mathcal{X}$ -disc takes the hypothesis class  $\mathcal{H}$  into account.

Mansour, Mohri, and Rostamizadeh (2009a) showed that the following inequalities hold for any  $h, h' \in \mathcal{H}$ :

$$\begin{aligned} |R_T^\ell(h, h') - R_S^\ell(h, h')| &\leq \text{disc}_{\mathcal{H}}^\ell(P_T, P_S) \\ &\leq M \cdot L_1(P_T, P_S), \end{aligned}$$

where  $M > 0$  is a constant and  $L_1(\cdot, \cdot)$  is the  $L_1$  distance over distributions. Therefore, a tighter bound for the difference in the expected loss between the two domains can be obtained by considering a hypothesis class  $\mathcal{H}$ . However, a drawback of  $\mathcal{X}$ -disc is that it considers the worst pair of hypotheses as can be seen from (1). This may lead to a loose generalization error bound as we will show later, and also an intractable computation cost for empirical estimation. Ben-David et al. (2007) provided a computationally efficient proxy of  $\mathcal{X}$ -disc for the 0-1 loss defined as follows:

$$d_{\mathcal{H}}(P_T, P_S) = \sup_{h \in \mathcal{H}} \left| R_T^{\ell_{01}}(h, 1) - R_S^{\ell_{01}}(h, 1) \right|.$$

Although  $d_{\mathcal{H}}$  can be computed efficiently, there is no learning guarantee.

There is another variant of  $\mathcal{X}$ -disc called the generalized discrepancy (Cortes, Mohri, and Medina 2015). However, the theoretical analysis therein is only applicable to the regression task and not suitable for binary classification unlike the analysis for  $\mathcal{X}$ -disc.

Another discrepancy measure  $\mathcal{Y}$ -disc which also takes  $\mathcal{H}$  into account is as follows:

$$\mathcal{Y}\text{-disc}_{\mathcal{H}}^\ell(P_T, P_S) = \sup_{h \in \mathcal{H}} |R_T^\ell(h, f_T) - R_S^\ell(h, f_S)|.$$

While  $\mathcal{Y}$ -disc has been proposed to provide a tighter generalization error bound than  $\mathcal{X}$ -disc (Mohri and Medina 2012; Zhang, Zhang, and Ye 2012), it requires the labeling function  $f_T$  in the test domain which cannot be estimated in unsupervised domain adaptation in general.

## Source-guided Discrepancy (S-disc)

In this section, we propose a novel discrepancy measure called *source-guided discrepancy (S-disc)* to mitigate the

limitations of the existing discrepancy measures. Later, we will show that S-disc can provide a tighter generalization bound and can be computed efficiently compared with the existing measures.

We define S-disc as follows, which is obtained by fixing  $h'$  in the definition (1) of  $\mathcal{X}$ -disc to the true risk minimizer  $h_S^* = \arg \min_{h \in \mathcal{H}} R_S^\ell(h, f_S)$  in the source domain.

**Definition 1** (Source-guided discrepancy). Let  $\mathcal{H}$  be a hypothesis class and let  $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}_+$  be a loss function. S-disc between two distributions  $P_{D_1}$  and  $P_{D_2}$  is defined as

$$\zeta_{\mathcal{H}}^\ell(P_{D_1}, P_{D_2}) = \sup_{h \in \mathcal{H}} |R_{D_1}^\ell(h, h_S^*) - R_{D_2}^\ell(h, h_S^*)|. \quad (2)$$

S-disc satisfies the triangular inequality, i.e.,  $\zeta_{\mathcal{H}}^\ell(P_{D_1}, P_{D_2}) \leq \zeta_{\mathcal{H}}^\ell(P_{D_1}, P_{D_3}) + \zeta_{\mathcal{H}}^\ell(P_{D_3}, P_{D_2})$ , and is symmetric, i.e.,  $\zeta_{\mathcal{H}}^\ell(P_{D_1}, P_{D_2}) = \zeta_{\mathcal{H}}^\ell(P_{D_2}, P_{D_1})$ . However, in general, S-disc is not a distance as we may have  $\zeta_{\mathcal{H}}^\ell(P_{D_1}, P_{D_2}) = 0$  for some  $P_{D_1} \neq P_{D_2}$ .

S-disc has the following three advantages over existing discrepancy measures. First, the computation cost of S-disc is low since we do not need to consider a pair of hypotheses unlike  $\mathcal{X}$ -disc as we can see from (1) and (2). Second, S-disc leads to a tighter bound as discussed below. Mansour, Mohri, and Rostamizadeh (2009a) derived a generalization error bound by bounding the difference  $|R_T^\ell(h, h_S^*) - R_S^\ell(h, h_S^*)|$  with  $\mathcal{X}$ -disc. On the other hand, for any  $h \in \mathcal{H}$ , it is easy to see that from the definition of S-disc

$$\begin{aligned} |R_T^\ell(h, h_S^*) - R_S^\ell(h, h_S^*)| &\leq \zeta_{\mathcal{H}}^\ell(P_T, P_S) \\ &\leq \text{disc}_{\mathcal{H}}^\ell(P_T, P_S), \end{aligned} \quad (3)$$

which implies that S-disc gives a tighter generalization error bound than  $\mathcal{X}$ -disc. Third, since  $h_S^*$  is estimated by labeled data from the source domain, we can compute S-disc without labels from the target domain unlike  $\mathcal{Y}$ -disc. For these reasons, S-disc is more suitable for unsupervised domain adaptation than the existing discrepancy measures. A comparison of S-disc with the existing discrepancy measures is summarized in Table 1.

## S-disc Estimation for the 0-1 Loss

Here we consider the task of binary classification with the 0-1 loss, where the output space  $\mathcal{Y}$  is  $\{+1, -1\}$ . The following theorem states that S-disc estimation can be reduced to a cost-sensitive classification problem. We consider a symmetric hypothesis class  $\mathcal{H}$ , which is closed under negation, i.e., for any  $h \in \mathcal{H}$ ,  $-h$  is contained in  $\mathcal{H}$ .

**Theorem 2.** For the 0-1 loss and a symmetric hypothesis class  $\mathcal{H}$ , the following equality holds:

$$\zeta_{\mathcal{H}}^{\ell_{01}}(\hat{P}_T, \hat{P}_S) = 1 - \min_{h \in \mathcal{H}} J_{\ell_{01}}(h),$$

where  $J_\ell(h)$  is defined as

$$\begin{aligned} J_\ell(h) &= \frac{1}{n_S} \sum_{j=1}^{n_S} \ell(h(x_j^S), h_S^*(x_j^S)) \\ &\quad + \frac{1}{n_T} \sum_{i=1}^{n_T} \ell(h(x_i^T), -h_S^*(x_i^T)). \end{aligned}$$

---

## Algorithm 1: S-disc Estimation for the 0-1 Loss

---

**Input:** labeled source data  $\mathcal{S}$ , unlabeled target data  $\mathcal{T}$ , surrogate loss  $\ell_{\text{sur}}$ , hypothesis class  $\mathcal{H}$ .

**Output:**  $\zeta_{\mathcal{H}}^\ell(\hat{P}_T, \hat{P}_S)$ .

**Source learning:**

Learn a classifier  $\hat{h}_S$  using labeled source data  $\mathcal{S}_{\mathcal{X}}$ .

**Pseudo labeling:**

- $\tilde{\mathcal{S}} = \{(x, \text{sign} \circ \hat{h}_S(x)) \mid x \in \mathcal{S}_{\mathcal{X}}\}$ ,

- $\tilde{\mathcal{T}} = \{(x, -\text{sign} \circ \hat{h}_S(x)) \mid x \in \mathcal{T}\}$ .

**Cost sensitive learning from pseudo labeled data**

Learn another classifier  $h'' \in \mathcal{H}$  using  $\tilde{\mathcal{S}}$  and  $\tilde{\mathcal{T}}$  to minimize the surrogate cost-sensitive risk  $J_{\ell_{\text{sur}}}$ .

**return**  $\zeta_{\mathcal{H}}^\ell(\hat{P}_T, \hat{P}_S) = 1 - J_{\ell_{01}}(h'')$ .

---

The proof of Theorem 2 is given in Kuroki et al. (2018). Note that the minimization of  $J_{\ell_{01}}(h)$  corresponds to empirical risk minimization for  $\mathcal{S}^* = \{(x, h_S^*(x)) \mid x \in \mathcal{S}_{\mathcal{X}}\}$  and  $\mathcal{T}^* = \{(x, -h_S^*(x)) \mid x \in \mathcal{T}\}$ . Therefore, Theorem 2 naturally suggests a three-step algorithm illustrated in Algorithm 1, where  $h_S^*$  is replaced with  $\hat{h}_S$ . This allows us to compute S-disc by minimizer  $h''$  of the cost-sensitive risk  $J_{\ell_{01}}(h)$ , which weights the pseudo-labeled data  $\tilde{\mathcal{S}}$  and  $\tilde{\mathcal{T}}$  with costs  $1/n_S$  and  $1/n_T$ , respectively.

Since minimization of the 0-1 loss is computationally hard (Ben-David, Eiron, and Long 2003; Feldman et al. 2012), we use a surrogate loss such as the hinge loss  $\ell_{\text{hinge}}(y, y') = \max(0, 1 - yy')$  (Bartlett, Jordan, and McAuliffe 2006). Note that the 0-1 loss is used for calculating S-disc in the final step, i.e.,  $1 - J_{\ell_{01}}(h'')$ , while a surrogate loss is only used to train a classifier. Hinge loss minimization with the sequential minimal optimization (SMO) algorithm requires  $O((n_T + n_S)^3)$  for the entire algorithm (Platt 1999). This low computation cost has a big advantage over  $\mathcal{X}$ -disc (see Table 1).

## Theoretical Analysis

In this section, we show that S-disc can be estimated from finite data with the consistency guarantee. After that, we show that our proposed S-disc is useful for deriving a tighter generalization error bound than  $\mathcal{X}$ -disc. To derive theoretical results, the *Rademacher complexity* is used, which captures the complexity of a set of functions by measuring the capability of a hypothesis class to correlate with the random noise.

**Definition 3** (Rademacher complexity (Bartlett and Mendelson 2002)). Let  $\mathcal{H}$  be a set of real-valued functions defined over a set  $\mathcal{X}$ . Given a sample  $(x_1, \dots, x_m) \in \mathcal{X}^m$  independently and identically drawn from a distribution  $\mu$ , the Rademacher complexity of  $\mathcal{H}$  is defined as

$$\mathfrak{R}_{\mu, m}(\mathcal{H}) = \mathbb{E}_{x_1, \dots, x_m} \mathbb{E}_{\sigma} \left[ \sup_{h \in \mathcal{H}} \frac{1}{m} \sum_{i=1}^m \sigma_i h(x_i) \right],$$

where the inner expectation is taken over  $\sigma = (\sigma_1, \dots, \sigma_m)$

Table 1: Comparison of S-disc with the existing discrepancy measures for unsupervised domain adaptation in binary classification. We assume the hypothesis class  $\mathcal{H}$  satisfies (4). We consider the hinge loss for S-disc,  $\mathcal{X}$ -disc, and  $d_{\mathcal{H}}$ . The computational complexity of S-disc and  $d_{\mathcal{H}}$  are based on the empirical hinge loss minimization, which is solved with the kernel support vector machine by the SMO algorithm (Platt 1999). The computational complexity of  $\mathcal{X}$ -disc is that based on SDP relaxation by the ellipsoid method (Bubeck 2015) given in Kuroki et al. (2018).  $\mathcal{Y}$ -disc is not computable with unlabeled target data.

	S-disc (proposed)	$d_{\mathcal{H}}$	$\mathcal{X}$ -disc	$\mathcal{Y}$ -disc
taking $\mathcal{H}$ into account	✓	✓	✓	✓
convergence rate	$\mathcal{O}_p\left(n_{\mathcal{T}}^{-\frac{1}{2}} + n_{\mathcal{S}}^{-\frac{1}{2}}\right)$	$\mathcal{O}_p\left(n_{\mathcal{T}}^{-\frac{1}{2}} + n_{\mathcal{S}}^{-\frac{1}{2}}\right)$	$\mathcal{O}_p\left(n_{\mathcal{T}}^{-\frac{1}{2}} + n_{\mathcal{S}}^{-\frac{1}{2}}\right)$	N/A
target generalization error bound	✓	N/A	✓	✓
computational complexity	$O\left((n_{\mathcal{T}} + n_{\mathcal{S}})^3\right)$	$O\left((n_{\mathcal{T}} + n_{\mathcal{S}})^3\right)$	$O\left((n_{\mathcal{T}} + n_{\mathcal{S}} + d)^8\right)$	N/A

which are mutually independent uniform random variables taking values in  $\{+1, -1\}$ .

Hereafter, we use the following notation:

- $\mathcal{H} \otimes \mathcal{H} := \{x \mapsto h(x) \cdot h'(x) \mid h, h' \in \mathcal{H}\}$ ,
- $\ell \circ (\mathcal{H} \otimes \mathcal{H}) := \{x \mapsto \ell(h(x), h'(x)) \mid h, h' \in \mathcal{H}\}$ .

### Consistency of S-disc

In this section, we show the estimator  $\varsigma_{\mathcal{H}}^{\ell}(\widehat{P}_{\mathcal{T}}, \widehat{P}_{\mathcal{S}})$  converges to the true S-disc  $\varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}})$  as the numbers of samples  $n_{\mathcal{T}}$  and  $n_{\mathcal{S}}$  increase. The following theorem gives the deviation of the empirical S-disc estimator, which is a general result that does not depend on the specific choice of the loss  $\ell$  and the hypothesis class  $\mathcal{H}$ .

**Theorem 4.** Assume the loss function  $\ell$  is bounded from above by  $M > 0$ . For any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} & \left| \varsigma_{\mathcal{H}}^{\ell}(\widehat{P}_{\mathcal{T}}, \widehat{P}_{\mathcal{S}}) - \varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}}) \right| \\ & \leq 2\mathfrak{R}_{P_{\mathcal{T}}, n_{\mathcal{T}}}(\ell \circ (\mathcal{H} \otimes \mathcal{H})) + 2\mathfrak{R}_{P_{\mathcal{S}}, n_{\mathcal{S}}}(\ell \circ (\mathcal{H} \otimes \mathcal{H})) \\ & \quad + M \sqrt{\frac{\log \frac{4}{\delta}}{2n_{\mathcal{T}}}} + M \sqrt{\frac{\log \frac{4}{\delta}}{2n_{\mathcal{S}}}}. \end{aligned}$$

The proof of this theorem is given in Kuroki et al. (2018).

Theorem 4 guarantees the consistency of  $\varsigma_{\mathcal{H}}^{\ell}(\widehat{P}_{\mathcal{T}}, \widehat{P}_{\mathcal{S}})$  under the condition that  $\mathfrak{R}_{P_{\mathcal{T}}, n_{\mathcal{T}}}(\ell \circ (\mathcal{H} \otimes \mathcal{H}))$  and  $\mathfrak{R}_{P_{\mathcal{S}}, n_{\mathcal{S}}}(\ell \circ (\mathcal{H} \otimes \mathcal{H}))$  are well-controlled.

To derive a specific convergence rate, we consider an assumption given by

$$\mathfrak{R}_{P_{\mathcal{T}}, n_{\mathcal{T}}}(\mathcal{H} \otimes \mathcal{H}) \leq \frac{\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}}}{\sqrt{n_{\mathcal{T}}}}, \quad \mathfrak{R}_{P_{\mathcal{S}}, n_{\mathcal{S}}}(\mathcal{H} \otimes \mathcal{H}) \leq \frac{\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}}}{\sqrt{n_{\mathcal{S}}}}, \quad (4)$$

for some constant  $\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}} > 0$  depending only on the hypothesis class  $\mathcal{H}$ . Lemma 5 below shows that this assumption is naturally satisfied in the linear-in-parameter model.

**Lemma 5.** Let  $\mathcal{H}$  be the linear-in-parameter model class, i.e.,  $\mathcal{H} := \{x \mapsto w^{\top} \phi(x) \mid w \in \mathbb{R}^d, \|w\|_2 \leq \Lambda\}$  for fixed basis functions  $\phi : \mathcal{X} \rightarrow \mathbb{R}^d$  satisfying  $\|\phi\|_{\infty} \leq D_{\phi}$ . Then, for any distribution  $\mu$  over  $\mathcal{X}$  and  $m \in \mathbb{N}$ ,

$$\mathfrak{R}_{\mu, m}(\mathcal{H} \otimes \mathcal{H}) \leq \frac{\Lambda^2 D_{\phi}^2}{\sqrt{m}}.$$

The proof of Lemma 5 is given in Kuroki et al. (2018). Subsequently, we derive the convergence rate bound for the 0-1 loss.

**Corollary 6.** When we consider  $\ell = \ell_{01}$ , it holds for any  $\delta \in (0, 1)$  that, with probability at least  $1 - \delta$ ,

$$\begin{aligned} & \left| \varsigma_{\mathcal{H}}^{\ell}(\widehat{P}_{\mathcal{T}}, \widehat{P}_{\mathcal{S}}) - \varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}}) \right| \\ & \leq \frac{\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}}}{\sqrt{n_{\mathcal{T}}}} + \frac{\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}}}{\sqrt{n_{\mathcal{S}}}} + \sqrt{\frac{\log \frac{4}{\delta}}{2n_{\mathcal{T}}}} + \sqrt{\frac{\log \frac{4}{\delta}}{2n_{\mathcal{S}}}} \end{aligned}$$

under the assumption in (4).

*Proof.* It simply follows from Theorem 4, the assumption in (4), and the fact  $\mathfrak{R}_{\mu, k}(\ell_{01} \circ (\mathcal{H} \otimes \mathcal{H})) = \frac{1}{2} \mathfrak{R}_{\mu, k}(\mathcal{H} \otimes \mathcal{H})$  for any distribution  $\mu$  and  $k > 0$  (Mohri, Rostamizadeh, and Talwalkar 2012, Lemma 3.1).  $\square$

From this corollary, we see that the empirical S-disc has the consistency with convergence rate  $\mathcal{O}_p(n_{\mathcal{T}}^{-1/2} + n_{\mathcal{S}}^{-1/2})$  under a mild condition.

### Generalization Error Bound

In the previous section, we showed that S-disc  $\varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}})$  can be estimated by  $\varsigma_{\mathcal{H}}^{\ell}(\widehat{P}_{\mathcal{T}}, \widehat{P}_{\mathcal{S}})$ . In this section, we give two bounds on the generalization error  $R_{\mathcal{T}}^{\ell}(h, f_{\mathcal{T}})$  for the target domain in terms of  $\varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}})$  or  $\varsigma_{\mathcal{H}}^{\ell}(\widehat{P}_{\mathcal{T}}, \widehat{P}_{\mathcal{S}})$ .

The first bound shows the relationship between target risk  $R_{\mathcal{T}}^{\ell}(h, f_{\mathcal{T}})$  and source risk  $R_{\mathcal{S}}^{\ell}(h, h_{\mathcal{S}}^*)$ .

**Theorem 7.** Assume that  $\ell$  obeys the triangular inequality, i.e.,  $\ell(u, v) \leq \ell(u, w) + \ell(w, v)$ ,  $\forall u, v, w \in \mathbb{R}$ , such as the 0-1 loss. Then, for any hypothesis  $h \in \mathcal{H}$ ,

$$\begin{aligned} & R_{\mathcal{T}}^{\ell}(h, f_{\mathcal{T}}) - R_{\mathcal{T}}^{\ell}(h_{\mathcal{T}}^*, f_{\mathcal{T}}) \\ & \leq R_{\mathcal{S}}^{\ell}(h, h_{\mathcal{S}}^*) + R_{\mathcal{T}}^{\ell}(h_{\mathcal{S}}^*, h_{\mathcal{T}}^*) + \varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}}). \end{aligned} \quad (5)$$

*Proof.* Since

$$R_{\mathcal{T}}^{\ell}(h, f_{\mathcal{T}}) \leq R_{\mathcal{T}}^{\ell}(h, h_{\mathcal{S}}^*) + R_{\mathcal{T}}^{\ell}(h_{\mathcal{S}}^*, h_{\mathcal{T}}^*) + R_{\mathcal{T}}^{\ell}(h_{\mathcal{T}}^*, f_{\mathcal{T}})$$

holds from the triangular inequality, we have

$$\begin{aligned} & R_{\mathcal{T}}^{\ell}(h, f_{\mathcal{T}}) - R_{\mathcal{T}}^{\ell}(h_{\mathcal{T}}^*, f_{\mathcal{T}}) \\ & \leq R_{\mathcal{T}}^{\ell}(h, h_{\mathcal{S}}^*) + R_{\mathcal{T}}^{\ell}(h_{\mathcal{S}}^*, h_{\mathcal{T}}^*) \\ & \leq R_{\mathcal{S}}^{\ell}(h, h_{\mathcal{S}}^*) + R_{\mathcal{T}}^{\ell}(h_{\mathcal{S}}^*, h_{\mathcal{T}}^*) + \varsigma_{\mathcal{H}}^{\ell}(P_{\mathcal{T}}, P_{\mathcal{S}}), \end{aligned}$$

where the last inequality follows from the definition of S-disc.  $\square$

The LHS of (5) represents the regret arising from the use of hypothesis  $h$  instead of  $h_T^*$  in the target domain. Theorem 7 shows that the regret  $R_T^\ell(h, f_T) - R_T^\ell(h_T^*, f_T)$  is bounded by three terms: (i) the expected loss with respect to  $h_S^*$  in the source domain, (ii) the difference between  $h_T^*$  and  $h_S^*$  in the target domain, and (iii) S-disc between  $P_T$  and  $P_S$ . Note that if the source and target domains are sufficiently close, we can expect the second term  $R_T^\ell(h_S^*, h_T^*)$  and the third term  $\zeta_{\mathcal{H}}^\ell(P_T, P_S)$  to be small. This fact indicates that for an appropriate source domain, minimization of estimation error  $R_S^\ell(h, h_S^*)$  in the source domain leads to a better generalization in the target domain.

We can see an advantage of the generalization error bound based on S-disc through comparison with the bound based on  $\mathcal{X}$ -disc (Mansour, Mohri, and Rostamizadeh 2009a, Theorem 8) given by

$$\begin{aligned} R_T^\ell(h, f_T) - R_T^\ell(h_T^*, f_T) \\ \leq R_S^\ell(h, h_S^*) + R_T^\ell(h_T^*, h_S^*) + \text{disc}_{\mathcal{H}}^\ell(P_T, P_S). \end{aligned} \quad (6)$$

The upper bound (6) using  $\mathcal{X}$ -disc has the same form as the upper bound (5) except for the term  $\zeta_{\mathcal{H}}^\ell(P_T, P_S)$ . Since S-disc is never larger than  $\mathcal{X}$ -disc (see the inequality (3)), S-disc gives a tighter bound than  $\mathcal{X}$ -disc.

The following theorem shows the generalization error bound for the finite-sample case.

**Theorem 8.** When we consider  $\ell = \ell_{01}$ , for any  $h \in \mathcal{H}$  and  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} R_T^\ell(h, f_T) - R_T^\ell(h_T^*, f_T) \\ \leq \widehat{R}_S^\ell(h, h_S^*) + R_T^\ell(h_S^*, h_T^*) + \zeta_{\mathcal{H}}^\ell(\widehat{P}_T, \widehat{P}_S) \\ + \frac{\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}}}{\sqrt{n_T}} + \frac{\mathfrak{C}_{\mathcal{H} \otimes \mathcal{H}}}{\sqrt{n_S}} + \sqrt{\frac{\log \frac{5}{\delta}}{2n_T}} + 2\sqrt{\frac{\log \frac{5}{\delta}}{2n_S}} \end{aligned}$$

under the assumption in (4).

The proof of this theorem is given in Kuroki et al. (2018).

Theorem 8 tells us that when  $n_S, n_T \rightarrow \infty$  the following three terms are dominating in the bound of the regret in the target domain  $R_T^\ell(h, f_T) - R_T^\ell(h_T^*, f_T)$ : (i) the empirical loss with respect to  $h_S^*$  in the source domain, (ii) the difference between  $h_T^*$  and  $h_S^*$  in the target domain, and (iii) S-disc between the two empirical distributions  $\zeta_{\mathcal{H}}^\ell(\widehat{P}_T, \widehat{P}_S)$ . Therefore, if  $h_T^*$  is sufficiently close to  $h_S^*$ , selecting a good source in terms of S-disc allows us to achieve good target generalization.

## Comparison with Existing Discrepancy Measures

In the previous section, we showed the consistency of the estimator for S-disc and derived a generalization error bound of S-disc tighter than  $\mathcal{X}$ -disc. In this section, we first compare these theoretical guarantees of S-disc with those for the existing ones in more detail. We next discuss their computation cost. This is also an important aspect when we apply these discrepancy measures to sentiment analysis (Bhatt, Rajkumar, and Roy 2016), adversarial learning (Zhao et al.

2017), and computer vision (Saito et al. 2018) for source selection or reweighting of the source data. In fact, in these applications the discrepancy  $d_{\mathcal{H}}$  instead of  $\mathcal{X}$ -disc is used for ease of computation even though  $d_{\mathcal{H}}$  has no theoretical guarantee on the generalization error. The results of this section are summarized in Table 1.

## Convergence Rates of Discrepancy Estimators

Here we discuss the consistency and convergence rates of the estimators of discrepancy measures.

The empirical estimator of  $d_{\mathcal{H}}$  is consistent, and its convergence rate is  $\mathcal{O}_p(((\log n_T)/n_T)^{1/2} + ((\log n_S)/n_S)^{1/2})$  (Ben-David et al. 2010, Lemma 1). This rate is slower than the rate  $\mathcal{O}_p(n_T^{-1/2} + n_S^{-1/2})$  for S-disc with appropriately controlled Rademacher complexities  $\mathfrak{R}_{P_T, n_T}(\mathcal{H})$  and  $\mathfrak{R}_{P_S, n_S}(\mathcal{H})$  of the hypothesis class, such as the linear-in-parameter model (Mohri, Rostamizadeh, and Talwalkar 2012, Theorem 4.3). Here recall that no generalization error bound is known in terms of  $d_{\mathcal{H}}$  even though  $d_{\mathcal{H}}$  itself is consistently estimated.

On the other hand, the empirical estimator of  $\mathcal{X}$ -disc is shown to be consistent (Mansour, Mohri, and Rostamizadeh 2009a, Corollary 7) and its convergence rate is  $\mathcal{O}_p(n_T^{-1/2} + n_S^{-1/2})$  in the case that the loss function is  $\ell_q$  loss, i.e.,  $\ell_q(y, y') = |y - y'|^q$ . Thus, the derived rate is the same as S-disc whereas the requirement on the loss function is more restrictive than the one for S-disc in Theorem 8.

Note that the above difference of the theoretical guarantees does not come from the inherent difference of these estimators. This is because we adopted the analysis based on the Rademacher complexity, which has not been well studied in the context of unsupervised domain adaptation. This is a distribution-dependent complexity measure and less pessimistic compared with the VC-dimension used in the previous work (Ben-David et al. 2010). In fact, the known guarantees on  $d_{\mathcal{H}}$  and  $\mathcal{X}$ -disc can be improved to Propositions 9 and 10 given below.

**Proposition 9.** For any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| d_{\mathcal{H}}(\widehat{P}_T, \widehat{P}_S) - d_{\mathcal{H}}(P_T, P_S) \right| \\ \leq 2\mathfrak{R}_{P_T, n_T}(\mathcal{H}) + 2\mathfrak{R}_{P_S, n_S}(\mathcal{H}) + \sqrt{\frac{2 \log \frac{4}{\delta}}{n_T}} + \sqrt{\frac{2 \log \frac{4}{\delta}}{n_S}}. \end{aligned}$$

**Proposition 10.** Assume the loss function  $\ell$  is upper bounded by  $M > 0$ . For any  $\delta \in (0, 1)$ , with probability at least  $1 - \delta$ ,

$$\begin{aligned} \left| \text{disc}_{\mathcal{H}}^\ell(\widehat{P}_T, \widehat{P}_S) - \text{disc}_{\mathcal{H}}^\ell(P_T, P_S) \right| \\ \leq 2\mathfrak{R}_{P_T, n_T}(\ell \circ (\mathcal{H} \otimes \mathcal{H})) + 2\mathfrak{R}_{P_S, n_S}(\ell \circ (\mathcal{H} \otimes \mathcal{H})) \\ + M\sqrt{\frac{\log \frac{4}{\delta}}{2n_T}} + M\sqrt{\frac{\log \frac{4}{\delta}}{2n_S}}. \end{aligned}$$

The proofs of these propositions are quite similar to that of Theorem 4 and omitted.




Figure 1: 2D plots of three domains.

In summary, under the mild assumptions, the convergence rates of the empirical estimators of  $d_{\mathcal{H}}$  and  $\mathcal{X}$ -disc are  $\mathcal{O}_p(n_T^{-1/2} + n_S^{-1/2})$  as well as S-disc.

### Computational Complexity

Computation of  $d_{\mathcal{H}}$  can be done by the empirical risk minimization (Ben-David et al. 2010, Lemma 2). The original form is given with the 0-1 loss, which can be efficiently minimized with a surrogate loss. When the hinge loss is applied, the minimization can be carried out with the computation cost  $\mathcal{O}((n_T + n_S)^3)$  by the SMO algorithm (Platt 1999).

On the other hand, no efficient algorithm is given for the computation of  $\mathcal{X}$ -disc in the classification setting.<sup>1</sup> For a fair comparison, we give a relatively efficient algorithm to compute  $\mathcal{X}$ -disc (1) in the classification setting with the hinge loss, based on *semidefinite relaxation*. Unfortunately, the computational complexity of the relaxed algorithm is still  $\mathcal{O}((n_T + n_S + d)^8)$ , which is prohibitive compared with the computation of S-disc and  $d_{\mathcal{H}}$ .

### Experiments

In this section, we provide the experimental results that illustrate the failure of existing discrepancy measures and the advantage of using S-disc. We illustrate the advantage of S-disc in terms of computation time, the empirical convergence, and the performance on the source selection task.

#### Illustration

We illustrate the failure of  $d_{\mathcal{H}}$ , which is the well-known proxy for  $\mathcal{X}$ -disc. We compared S-disc with  $d_{\mathcal{H}}$  in the toy experiment.

We generated 200 data points per class for each of two

<sup>1</sup>A computation algorithm of  $\mathcal{X}$ -disc for the 0-1 loss is given only in the one-dimensional case (Mansour, Mohri, and Rosztamzadeh 2009a, Section 5.2).

sources  $S_1$  and  $S_2$  and target domain as follows:

$$P_{S_i}(X|Y = j) = \mathcal{N}(\mu_i^{(j)}, I) \quad (i = 1, 2),$$

$$P_T(X|Y = j) = \mathcal{N}(\mu_T^{(j)}, I),$$

where

$$\mu_1^{(0)} = (-5, -5), \quad \mu_2^{(0)} = (0, 3), \quad \mu_T^{(0)} = (-5, -3),$$

$$\mu_1^{(1)} = (5, -5), \quad \mu_2^{(1)} = (2, -3), \quad \mu_T^{(1)} = (5, -3).$$

In this experiment, we used the support vector machine with a linear kernel<sup>2</sup>. For these data, we obtained the following results:

$$\zeta_{\mathcal{H}}^{\ell}(\widehat{P}_T, \widehat{P}_{S_1}) = 0.27, \quad \zeta_{\mathcal{H}}^{\ell}(\widehat{P}_T, \widehat{P}_{S_2}) = 0.49,$$

$$d_{\mathcal{H}}(\widehat{P}_T, \widehat{P}_{S_1}) = 0.69, \quad d_{\mathcal{H}}(\widehat{P}_T, \widehat{P}_{S_2}) = 0.49.$$

These values indicate that while  $d_{\mathcal{H}}$  regards  $S_2$  as the better source, S-disc regards  $S_1$  as the better source for target domain.

In this example, the loss calculated on the target domain of the classifier trained on  $S_1$  is 0.0 and the loss of the classifier trained on  $S_2$  is 0.49. This implies that S-disc is the better discrepancy measure to measure the quality of source domains for a better generalization in a given target domain.

Intuitively,  $d_{\mathcal{H}}$  is a heuristic measure based on a classifier separating the source and target domains without considering labels in the source domain. Once the supports of the input distributions are not highly overlapped between the source and target domains,  $d_{\mathcal{H}}$  may immediately regard that domains are totally different from each other even if the risk minimizers are highly similar between these domains, which resulted in the failure of  $d_{\mathcal{H}}$ . On the other hand, S-disc can prevent such a problem by taking labels from the source domain into account as illustrated in the toy experiments and justified in Theorems 7 and 8.

### Comparison of Computation Time

We compared the computation time to estimate S-disc,  $d_{\mathcal{H}}$ , and  $\mathcal{X}$ -disc. We used 2-dimensional 200 examples for both source and target domains. Each domain consists of 100 positive examples and 100 negative examples. For the computation of  $\mathcal{X}$ -disc, we used the relaxed algorithm<sup>3</sup> shown in Kuroki et al. (2018). For both  $d_{\mathcal{H}}$  and S-disc, we used the support vector machine with a linear kernel for the computation of S-disc and  $d_{\mathcal{H}}$ . The simulation was run on 2.8GHz Intel<sup>®</sup> Core i7. The results shown in Figure 2 demonstrate that the computation time of  $\mathcal{X}$ -disc is prohibitive while both S-disc and  $d_{\mathcal{H}}$  are feasible to compute as suggested in Table 1.

### Empirical Convergence

We compared the empirical convergence of S-disc and  $d_{\mathcal{H}}$  based on logistic regression implemented with scikit-learn with default parameters (Pedregosa et al. 2011). Note that

<sup>2</sup>We used scikit-learn (Pedregosa et al. 2011) to implement it.

<sup>3</sup>We used picos (<https://picos-api.gitlab.io/picos/>) for the SDP solver with cvxopt (<https://cvxopt.org/>) backend.




Figure 2: Comparison of computation time in the log scale.




Figure 3: Empirical convergence of  $d_{\mathcal{H}}$  and S-disc.

$\mathcal{X}$ -disc is not used due to computational intractability. Here, we used MNIST (LeCun, Cortes, and Burges 2010) dataset, and considered a binary classification task to separate odd and even digits. We defined the source and target domains as follows:

- Source domain  $S_1$ : MNIST,
- Source domain  $S_2$ : MNIST from zero to seven,
- Target domain: MNIST.

The number of examples was ranged from  $\{1000, 2000, \dots, 20000\}$  for each domain. Note that while examples from the source domain  $S_1$  are drawn from the same distribution of the target domain, the sample of the source domain  $S_2$  is affected by sample selection bias (Cortes et al. 2008). Figure 3 shows the empirical convergence of the estimator of both discrepancy measures. It is observed that both discrepancy measures indicate that  $S_1$  is a better source than  $S_2$  for target domain. However, since the true value of the discrepancy of  $S_1$  from the target domain is supposed to be zero, we can observe that the  $d_{\mathcal{H}}$  will be converged much more slowly than S-disc or converged to a non-zero value.

### Source Selection

We compared the performance in the source selection task between S-disc and  $d_{\mathcal{H}}$ . We used the source domains and the target domain as follows:

- Clean source domains: Five grayscale MNIST-M (Ganin et al. 2016),




Figure 4: Source selection performance with varying noise rates. Blue dotted line denotes the maximum score (five).

- Noisy source domains: Five grayscale MNIST-M corrupted by Gaussian random noise,
- Target domain: MNIST.

The MNIST-M dataset is known to be useful for the domain adaptation task when the target domain is MNIST. The task of each domain is to classify between even and odd digits and logistic regression with default parameters was used for computing S-disc and  $d_{\mathcal{H}}$ . The objective of this source selection task is to correctly rank five clean source domains over noisy domains. We ranked each source by computing discrepancy values between the target domain and each source and rank them in ascending order. The score was calculated by counting how many clean sources are ranked in the first five sources. We varied the number of examples from  $\{200, 400, \dots, 4000\}$  for each domain. The Gaussian noise with standard deviation  $\epsilon = 30, 40, 50$  were added and clipped to force the value to be between  $0 - 255$ . For each number of examples per class, the experiments were conducted 15 times and the average score was used.

Figure 4 shows the performance of each discrepancy measure with different noise rates. As the number of examples increased, S-disc achieved a better performance. In contrast,  $d_{\mathcal{H}}$  cannot distinguish between noisy and clean source domains. In fact,  $d_{\mathcal{H}}$  always returned one, which indicates that MNIST-M is unrelated to MNIST. Unlike the previous experiment on empirical convergence, the difference between the two domains may be harder to identify since the source domains and the target domain are not exactly the same (MNIST vs MNIST-M). As a result, our experiment demonstrates the failure of  $d_{\mathcal{H}}$  in the source selection task and suggests the advantage of S-disc in this task.

### Conclusion

We proposed a novel discrepancy measure for unsupervised domain adaptation called source-guided discrepancy (S-disc). We provided a computationally efficient algorithm for the estimation of S-disc with respect to the 0-1 loss, and also established the consistency of the estimator with the convergence rate. Moreover, we derived a generalization



bound based on S-disc, which is tighter than that of  $\mathcal{X}$ -disc. Finally, we demonstrated the failure of existing discrepancy measures and the advantages of S-disc through experiments.

## Acknowledgements

We thank Ikko Yamane, Futoshi Futami, and Kento Nozawa for the useful discussion. NC was supported by MEXT scholarship. JH was supported by KAKENHI 17H00757, IS was supported by JST CREST Grant Number JPMJCR17A1, and MS was supported by KAKENHI 17H01760.

## References

- Bartlett, P. L., and Mendelson, S. 2002. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research* 3:463–482.
- Bartlett, P. L.; Jordan, M. I.; and McAuliffe, J. D. 2006. Convexity, classification, and risk bounds. *Journal of the American Statistical Association* 101(473):138–156.
- Ben-David, S.; Blitzer, J.; Crammer, K.; and Pereira, F. 2007. Analysis of representations for domain adaptation. In *NIPS*, 137–144.
- Ben-David, S.; Blitzer, J.; Crammer, K.; Kulesza, A.; Pereira, F.; and Vaughan, J. W. 2010. A theory of learning from different domains. *Machine Learning* 79(1-2):151–175.
- Ben-David, S.; Eiron, N.; and Long, P. M. 2003. On the difficulty of approximately maximizing agreements. *Journal of Computer and System Sciences* 66(3):496–514.
- Bhatt, H. S.; Rajkumar, A.; and Roy, S. 2016. Multi-source iterative adaptation for cross-domain classification. In *IJ-CAI*, 3691–3697.
- Bubeck, S. 2015. *Convex Optimization: Algorithms and Complexity*, volume 8. Now Publishers, Inc.
- Cortes, C.; Mohri, M.; Riley, M.; and Rostamizadeh, A. 2008. Sample selection bias correction theory. In *ALT*, 38–53.
- Cortes, C.; Mohri, M.; and Medina, A. M. 2015. Adaptation algorithm and theory based on generalized discrepancy. In *SIGKDD*, 169–178.
- Courty, N.; Flamary, R.; Habrard, A.; and Rakotomamonjy, A. 2017. Joint distribution optimal transportation for domain adaptation. In *NIPS*, 3730–3739.
- Feldman, V.; Guruswami, V.; Raghavendra, P.; and Wu, Y. 2012. Agnostic learning of monomials by halfspaces is hard. *SIAM Journal on Computing* 41(6):1558–1590.
- Ganin, Y.; Ustinova, E.; Ajakan, H.; Germain, P.; Larochelle, H.; Laviolette, F.; Marchand, M.; and Lempitsky, V. 2016. Domain-adversarial training of neural networks. *Journal of Machine Learning Research* 17(1):2096–2030.
- Germain, P.; Habrard, A.; Laviolette, F.; and Morvant, E. 2013. A PAC-Bayesian approach for domain adaptation with specialization to linear classifiers. In *ICML*, 738–746.
- Glorot, X.; Bordes, A.; and Bengio, Y. 2011. Domain adaptation for large-scale sentiment classification: A deep learning approach. In *ICML*, 513–520.
- Huang, J.; Gretton, A.; Borgwardt, K. M.; Schölkopf, B.; and Smola, A. J. 2007. Correcting sample selection bias by unlabeled data. In *NIPS*, 601–608.
- Kuroki, S.; Charoenphakdee, N.; Bao, H.; Honda, J.; Sato, I.; and Sugiyama, M. 2018. Unsupervised domain adaptation based on source-guided discrepancy. *arXiv:1809.03839*.
- LeCun, Y.; Cortes, C.; and Burges, C. 2010. MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist/>.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009a. Domain adaptation: Learning bounds and algorithms. In *COLT*.
- Mansour, Y.; Mohri, M.; and Rostamizadeh, A. 2009b. Multiple source adaptation and the Rényi divergence. In *UAI*, 367–374.
- Mohri, M., and Medina, A. M. 2012. New analysis and algorithm for learning with drifting distributions. In *ALT*, 124–138.
- Mohri, M.; Rostamizadeh, A.; and Talwalkar, A. 2012. *Foundations of Machine Learning*. MIT Press.
- Pan, S. J., and Yang, Q. 2010. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering* 22(10):1345–1359.
- Pedregosa, F.; Varoquaux, G.; Gramfort, A.; Michel, V.; Thirion, B.; Grisel, O.; Blondel, M.; Prettenhofer, P.; Weiss, R.; Dubourg, V.; Vanderplas, J.; Passos, A.; Cournapeau, D.; Brucher, M.; Perrot, M.; and Duchesnay, É. 2011. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research* 12:2825–2830.
- Platt, J. C. 1999. Sequential minimal optimization: A fast algorithm for training support vector machines. In *Advances in Kernel Methods: Support Vector Learning*. MIT Press. 185–208.
- Saito, K.; Watanabe, K.; Ushiku, Y.; and Harada, T. 2018. Maximum classifier discrepancy for unsupervised domain adaptation. In *CVPR*, 3723–3732.
- Saito, K.; Ushiku, Y.; and Harada, T. 2017. Asymmetric tri-training for unsupervised domain adaptation. In *ICML*, 2988–2997.
- Sugiyama, M., and Kawanabe, M. 2012. *Machine Learning in Non-Stationary Environments: Introduction to Covariate Shift Adaptation*. MIT Press.
- Sugiyama, M.; Suzuki, T.; Nakajima, S.; Kashima, H.; von Büna, P.; and Kawanabe, M. 2008. Direct importance estimation for covariate shift adaptation. *Annals of the Institute of Statistical Mathematics* 60(4):699–746.
- Sun, S.; Zhang, B.; Xie, L.; and Zhang, Y. 2017. An unsupervised deep domain adaptation approach for robust speech recognition. *Neurocomputing* 257:79–87.
- Zhang, C.; Zhang, L.; and Ye, J. 2012. Generalization bounds for domain adaptation. In *NIPS*, 3320–3328.
- Zhao, H.; Zhang, S.; Wu, G.; Costeira, J. P.; Moura, J. M. F.; and Gordon, G. J. 2017. Multiple source domain adaptation with adversarial training of neural networks. *arXiv:1705.09684*.