
Implicit Class-Conditioned Domain Alignment for Unsupervised Domain Adaptation

Xiang Jiang^{1,2} Qicheng Lao^{1,3} Stan Matwin^{2,4} Mohammad Havaei¹

Abstract

We present an approach for unsupervised domain adaptation—with a strong focus on practical considerations of within-domain class imbalance and between-domain class distribution shift—from a class-conditioned domain alignment perspective. Current methods for class-conditioned domain alignment aim to *explicitly* minimize a loss function based on pseudo-label estimations of the target domain. However, these methods suffer from pseudo-label bias in the form of error accumulation. We propose a method that removes the need for *explicit* optimization of model parameters from pseudo-labels. Instead, we present a sampling-based *implicit* alignment approach, where the sample selection is *implicitly* guided by the pseudo-labels. Theoretical analysis reveals the existence of a domain-discriminator shortcut in misaligned classes, which is addressed by the proposed approach to facilitate domain-adversarial learning. Empirical results and ablation studies confirm the effectiveness of the proposed approach, especially in the presence of within-domain class imbalance and between-domain class distribution shift.

1. Introduction

Supervised learning aims to extract statistical patterns from data by learning to approximate the conditional density $p(y|x)$. However, the generalization of the approximation is often sensitive to some dataset-specific factors. Dataset shift (Quionero-Candela et al., 2009) frequently arises from real-world applications and can manifest in many different ways, such as sample selection bias (Heckman, 1979; Torralba et al., 2011), class distribution shift (Webb & Ting, 2005), and covariate shift (Shimodaira, 2000). Unsupervised

Domain Adaptation (UDA) aims to address domain shift with access to labeled data in the source domain and unlabeled data in the target domain (Pan & Yang, 2009). The fundamental algorithmic issue is to infer domain-invariant representations.

While considerable progress has been made in UDA (Ganin et al., 2016), they tend to focus on marginal distribution matching in the feature space, and less emphasis is made on discovering label distributions. In real-world applications, it is very common to have *class imbalance* (Japkowicz & Stephen, 2002; Chawla, 2009) within each domain and *class distribution shift* (Tan et al., 2019) between different domains, necessitating the incorporation of label space distribution into adaptation. *Explicit* class-conditioned domain alignment (Xie et al., 2018; Pan et al., 2019; Liang et al., 2019a; Deng et al., 2019) has emerged as a key approach to promoting class-conditioned invariance by aligning prototypical representations of each class. While explicit alignment has the advantage of directly minimizing class-conditioned misalignment, it presents critical vulnerabilities to error accumulation (Chen et al., 2019a) and ill-calibrated probabilities (Guo et al., 2017) due to its dependence on *explicit* supervision from pseudo-labels provided by model predictions.

We propose *Implicit* Class-Conditioned Domain Alignment that removes the need for explicit pseudo-label based optimization. Instead, we use the pseudo-labels *implicitly* to *sample* class-conditioned data in a way that maximally aligns the joint distribution between features and labels. The primary advantage of the sampling-based implicit domain alignment is the ability to address within-domain class imbalance and between-domain class distribution shift, in addition to many other benefits such as applications in cost-sensitive learning. The proposed method is simple, effective, and is supported by theoretical analysis on the empirical estimations of domain divergence measures. It also overcomes limitations of explicit alignment by allowing the domain adaptation algorithm to discover class-conditioned domain-invariance in an unsupervised way without explicit supervision from pseudo-labels.

The contributions of this paper are as follows: (i) We propose implicit class-conditioned domain alignment to address the challenge of within-domain class imbalance and between-domain class distribution shift, which overcomes the limitation of error accumulation in explicit domain align-

¹Imagia, Canada ²Dalhousie University, Canada ³Mila, Université de Montréal, Canada ⁴Polish Academy of Sciences, Poland. Correspondence to: Xiang Jiang <xiang.jiang@dal.ca>.

ment; (ii) We provide theoretical analysis on the empirical domain divergence and reveal the existence of a shortcut function that interferes with domain-invariant learning, which is addressed by the proposed approach; (iii) We show that the proposed approach is orthogonal to the choice of domain adaptation algorithms and offers consistent improvements to two adversarial domain adaptation algorithms; (iv) We report state-of-the-art UDA performance under extreme within-domain class imbalance and between-domain class distribution shift, and competitive results on standard UDA tasks.

2. Preliminaries

We follow the notations by (Ben-David et al., 2010) and define a domain as an ordered pair consisting of a distribution \mathcal{D} on the input space \mathcal{X} , and a labeling function $f: \mathcal{X} \rightarrow \mathcal{Y}$ that maps \mathcal{X} to the label space \mathcal{Y} . The source and target domains are denoted by $\langle \mathcal{D}_S, f_S \rangle$ and $\langle \mathcal{D}_T, f_T \rangle$, respectively.

In unsupervised domain adaptation, the model is trained on labeled data from the source domain, together with unlabeled data from the target domain. The goal is to obtain a model $h \in \mathcal{H}$ which learns domain-invariant representations while simultaneously minimizing the classification error on \mathcal{D}_S .

Adversarial training is the prevailing approach for domain adaptation (Ganin et al., 2016). It formulates a minimax problem where the maximizer maximizes the estimation of the domain divergence between the empirical samples, and the minimizer minimizes the sum of the source error and the domain divergence estimation obtained from the maximizer.

While matching the marginal distribution is a good step towards domain-invariant learning, it is still susceptible to the problem of conditional distribution mismatching. Prototype-based class-conditioned domain alignment (Luo et al., 2017; Xie et al., 2018; Chen et al., 2019a; Pan et al., 2019; Liang et al., 2019a;b) is designed to address this problem. We refer to this group of methods as *explicit* class-conditioned domain alignment. The explicit alignment is achieved by incorporating an auxiliary loss that minimizes the Euclidean distance of the class-conditioned prototypical representations \mathbf{c}_j between the source and target domains. The class-conditioned prototype \mathbf{c}_j is the average representation for all examples in a domain with class label j .

The main limitation of explicit class-conditioned domain alignment is in its reliance on explicit optimization of model parameters based on pseudo-labels. This learning procedure is vulnerable to error accumulation (Chen et al., 2019a) as mistakes in the pseudo-label predictions can gradually accumulate leading to poor local minima in EM-style training. Furthermore, the pseudo-labels are likely to suffer from ill-calibrated probabilities (Guo et al., 2017), especially for deep learning methods, which exacerbate the critical problem of error accumulation with misleadingly confident mistakes.

3. Method

We begin with theoretical motivations of implicit alignment by decomposing the empirical domain divergence measure into class-aligned and class-misaligned divergence, and show that misaligned divergence is detrimental to domain adaptation. We then present the proposed implicit domain alignment framework that addresses class-misalignment.

3.1. Theoretical Motivations

The $\mathcal{H}\Delta\mathcal{H}$ divergence between two domains is defined as

$$d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{D}_S, \mathcal{D}_T) = 2 \sup_{h, h' \in \mathcal{H}} |\mathbb{E}_{\mathcal{D}_T}[h \neq h'] - \mathbb{E}_{\mathcal{D}_S}[h \neq h']|, \quad (1)$$

where \mathcal{H} denotes some hypothesis space, and $h \neq h'$ is the abbreviation for $h(x) \neq h'(x)$. (Ben-David et al., 2010) theorized that the target domain error $\epsilon_T(h)$ is bounded by the error of the source domain $\epsilon_S(h)$ and the empirical domain divergence $\hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{U}_S, \mathcal{U}_T)$ where $\mathcal{U}_S, \mathcal{U}_T$ are unlabeled empirical samples drawn from $\mathcal{D}_S, \mathcal{D}_T$.

In deep learning, minibatch-based optimization limits the amount of data available at each training step. This necessitates the analysis of the empirical estimations of $d_{\mathcal{H}\Delta\mathcal{H}}$ at the minibatch level, so as to shed light on the learning dynamics.

Definition 3.1. Let $\mathcal{B}_S, \mathcal{B}_T$ be minibatches from \mathcal{U}_S and \mathcal{U}_T , respectively, where $\mathcal{B}_S \subseteq \mathcal{U}_S, \mathcal{B}_T \subseteq \mathcal{U}_T$, and $|\mathcal{B}_S| = |\mathcal{B}_T|$. The empirical estimation of $d_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{B}_S, \mathcal{B}_T)$ over the minibatches $\mathcal{B}_S, \mathcal{B}_T$ is defined as

$$\hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{B}_S, \mathcal{B}_T) = \sup_{h, h' \in \mathcal{H}} \left| \sum_{\mathcal{B}_T} [h \neq h'] - \sum_{\mathcal{B}_S} [h \neq h'] \right|. \quad (2)$$

Theorem 3.2 (The decomposition of $\hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{B}_S, \mathcal{B}_T)$). Let \mathcal{H} be a hypothesis space and \mathcal{Y} be the label space of the classification task where $\mathcal{B}_S, \mathcal{B}_T$ are minibatches drawn from $\mathcal{U}_S, \mathcal{U}_T$, respectively, and Y_S, Y_T are the label set of $\mathcal{B}_S, \mathcal{B}_T$. We define three disjoint sets on the label space: the shared labels $Y_C := Y_S \cap Y_T$, and the domain-specific labels $\bar{Y}_S := Y_S - Y_C$, and $\bar{Y}_T := Y_T - Y_C$. We also define the following disjoint sets on the input space where $\mathcal{B}_S^C := \{x \in \mathcal{B}_S | y \in Y_C\}$, $\mathcal{B}_S^{\bar{C}} := \{x \in \mathcal{B}_S | y \notin Y_C\}$, $\mathcal{B}_T^C := \{x \in \mathcal{B}_T | y \in Y_C\}$, $\mathcal{B}_T^{\bar{C}} := \{x \in \mathcal{B}_T | y \notin Y_C\}$. The empirical $\hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{B}_S, \mathcal{B}_T)$ divergence can be decomposed into class aligned divergence and class-misaligned divergence:

$$\hat{d}_{\mathcal{H}\Delta\mathcal{H}}(\mathcal{B}_S, \mathcal{B}_T) = \sup_{h, h' \in \mathcal{H}} \left| \xi^C(h, h') + \xi^{\bar{C}}(h, h') \right|, \quad (3)$$

where

$$\xi^C(h, h') = \sum_{\mathcal{B}_T^C} \mathbb{1}[h \neq h'] - \sum_{\mathcal{B}_S^C} \mathbb{1}[h \neq h'], \quad (4)$$

$$\xi^{\bar{C}}(h, h') = \sum_{\mathcal{B}_T^{\bar{C}}} \mathbb{1}[h \neq h'] - \sum_{\mathcal{B}_S^{\bar{C}}} \mathbb{1}[h \neq h']. \quad (5)$$




Figure 1. Illustration of the domain discriminator shortcut. The domain discriminator aims to distinguish between different domains (red and blue), where the decision boundary is represented by dashed lines. But misaligned samples create a shortcut where the domain labels can be directly determined by the misaligned class labels (3 and 6). The decision boundary of the resulting shortcut is independent of the covariate that causes the domain difference, which does not contribute to adversarial domain-invariant learning.

The proof is provided in supplementary materials.

Remark 3.3 (The domain discriminator shortcut). *Let the ordered triple (x, y_c, y_d) denote data sample x , and its associating class label y_c and domain label y_d , respectively, where $x \in \mathcal{B}$, $y_c \in Y$ and $y_d \in \{0, 1\}$. Let f_c be a classifier that maps x to a class label y_c . Let f_d be a domain discriminator that maps x to a binary domain label y_d . For the empirical class-misaligned divergence $\xi^{\bar{C}}(h, h')$ with sample $x \in \mathcal{B}_S^{\bar{C}} \cup \mathcal{B}_T^{\bar{C}}$, there exists a domain discriminator shortcut function*

$$f_d(x) = \begin{cases} 1 & f_c(x) \in \bar{Y}_S \\ 0 & f_c(x) \in \bar{Y}_T, \end{cases} \quad (6)$$

such that the domain label can be solely determined by the domain-specific class labels. This shortcut interferes with adversarial domain adaptation because the model could bypass the optimization for domain-invariant representations, but rather optimize for a shortcut function that is independent of the covariate contributing to the domain difference.

Figure 1 illustrates a toy example where the source and target domains are aligned for class 4 but misaligned between classes 3 and 6 as a result of random sampling in the minibatch construction. The domain discriminator aims to predict domain labels based on their domain information, i.e., red and blue. However, due to the class shortcut for the misaligned samples (3 and 6), the domain discriminator could infer domain labels based on class information directly (digits 3 and 6), without the need to learn domain-specific information. This problem of class-misalignment is especially pronounced under extreme within-domain class imbalance and between-domain class distribution shift, where a simple random sample is more likely to fail in providing good coverage of the label space.

3.2. Implicit Class-Conditioned Domain Alignment

Having identified the domain discriminator shortcut in class misaligned empirical samples, we now propose framework that aligns the two domains from a sampling perspective.

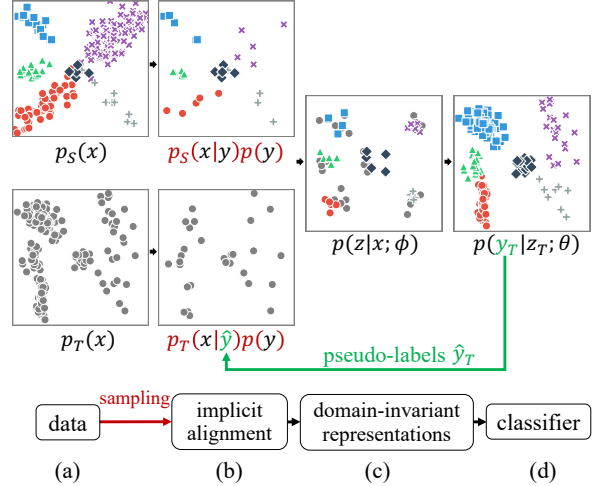


Figure 2. The proposed framework. (a) We aim to align the source domain $p_S(x)$, colored by classes, with unlabeled target domain $p_T(x)$. (b) For $p_S(x)$, we sample $x \sim p_S(x|y)p(y)$ based on the alignment distribution $p(y)$. For $p_T(x)$, we sample a class aligned minibatch $x \sim p_T(x|\hat{y})p(y)$ using identical $p(y)$, with the help of pseudo-labels \hat{y}_T . (c) The adversarial training aims to acquire domain-invariant representations z from the feature extractor parameterized by ϕ . (d) The classifier predicts class labels from z .

Figure 2 depicts the proposed implicit class-conditioned domain alignment framework. We aim to align $p_S(x)$ and $p_T(x)$ at the input and label space jointly with the factorization $p(x, y) = p(x|y)p(y)$ while ensuring that the sampled classes are aligned between the two domains. The alignment distribution $p(y)$ is pre-specified, e.g., uniform distribution, to ensure samples are aligned in the shared label space in spite of different empirical label distributions of the two domains. This implicit alignment procedure minimizes the class-misaligned divergence $\xi^{\bar{C}}(h, h')$, providing a more reliable empirical estimation of domain divergence. For the unlabeled target domain, we use the model predictions to sample class-conditioned data from $p_T(x|\hat{y})$ to approximate $p_T(x|y)$.

3.2.1. CLASS-ALIGNED SAMPLING STRATEGY

Algorithm 1 presents the proposed sampling procedure that selects class-aligned examples for minibatch training. It is a type of stratified sampling where the dataset is partitioned into mutually exclusive subgroups to reflect the label information in a class-aligned manner.

First, we predict pseudo-labels of the target domain using the classifier $f_c(\cdot; \theta)$ parameterized by θ . The pseudo-labels will be later used in class-conditioned sampling. Second, we sample a set Y from the label space \mathcal{Y} where $p(y)$ defines the probability with which we pick the classes to align so as to ensure the empirical samples of the source and target domains share the same Y . This in turn minimizes the class-misaligned divergence $\xi^{\bar{C}}(h, h')$. Third, for each class $y_i \in Y$, we sample class-conditioned examples for the source

Algorithm 1 The proposed implicit alignment training

```

1: Input: dataset  $S = \{(x_i, y_i)\}_{i=1}^N, T = \{x_i\}_{i=1}^M$ ,
2:   label space  $\mathcal{Y}$ , label alignment distribution  $p(y)$ ,
3:   classifier  $f_c(\cdot; \theta)$ 
4: while converged do
5:   # predict pseudo-labels for T
6:    $\hat{T} \leftarrow \{(x_i, \hat{y}_i)\}_{i=1}^M$  where  $x_i \in T$  and  $\hat{y}_i = f_c(x_i; \theta)$ 
7:   # sample N unique classes in the label space
8:    $Y \leftarrow$  draw  $N$  samples in  $\mathcal{Y}$  from  $p(y)$ 
9:   # sample K examples conditioned on each  $y_j \in Y$ 
10:  for  $y_j$  in  $Y$  do
11:     $(X'_S, Y'_S) \leftarrow$  draw  $K$  samples in  $S$  from
       $p_S(x|y=y_j)$ 
12:     $X'_T \leftarrow$  draw  $K$  samples in  $\hat{T}$  from  $p_T(x|\hat{y}=y_j)$ 
13:  end for
14:  # domain adaptation training on this minibatch
15:  train minibatch  $(X'_S, Y'_S, X'_T)$ 
16: end while
    
```

and target domains, respectively, and store them in (X'_S, Y'_S) and X'_T . This is equivalent to performing a table lookup to select a subset \mathcal{B}_i where all examples belong to class y_i , followed by random sampling in \mathcal{B}_i . We use pseudo-labels to sample the target domain due to the lack of ground-truth labels. Once we obtained the class-aligned minibatch, we use it to train unsupervised domain adaptation algorithm and repeat this process until the model converges.

This algorithm addresses class imbalance within each domain as well as class distribution shift between different domains by specifying the sampling strategy $p(y)$ in the label space. We use uniform sampling $p(y)$ for all experiments in this paper, and leave more advanced specifications and their applications to cost-sensitive domain adaptation as future work.

3.2.2. INTEGRATING IMPLICIT ALIGNMENT INTO CLASSIFIER-BASED DOMAIN DISCREPANCY MEASURE

Section 3.2.1 describes the implicit alignment algorithm from a sampling perspective, where we sample minibatches in a way that maximizes class alignment implicitly. This sampling strategy is independent of the choice of domain divergence measures. In this section, we show how to integrate the sampling approach into Margin Disparity Discrepancy (MDD) (Zhang et al., 2019b)—a state-of-the-art classifier-based domain discrepancy measure—to further facilitate implicit alignment. MDD is defined as

$$d_{f, \mathcal{F}}(S, T) = \sup_{f' \in \mathcal{F}} \left(\text{disp}_{\mathcal{D}_T}(f', f) - \text{disp}_{\mathcal{D}_S}(f', f) \right), \quad (7)$$

where f and f' are two independent scoring functions that predict class probabilities, and $\text{disp}(f', f)$ is a disparity measure between the scores provided by the classifiers f' and f . The domain divergence is to estimate the discrepancy

between the disparity measures of the two domains.

Following notations in Theorem 3.2, we define the empirical MDD on class-misaligned samples as

$$\hat{d}_{f, \mathcal{F}}(\mathcal{B}_S^{\bar{C}}, \mathcal{B}_T^{\bar{C}}) = \sup_{f' \in \mathcal{F}} \left(\sum_{\mathcal{B}_T^{\bar{C}}} \text{disp}(f', f) - \sum_{\mathcal{B}_S^{\bar{C}}} \text{disp}(f', f) \right). \quad (8)$$

Because $\mathcal{B}_S^{\bar{C}}$ and $\mathcal{B}_T^{\bar{C}}$ are disjoint in the label space, there exists a shortcut solution

$$\text{disp}(f'(x), f(x)) = \begin{cases} 0 & f_c(x) \in \bar{Y}_S \\ 1 & f_c(x) \in \bar{Y}_T, \end{cases} \quad (9)$$

which maximizes the divergence estimation of Eq. (8). Although class-aligned sampling can mitigate this problem, it is difficult to fully eliminate the impact of misalignment due to imperfect pseudo-labels. To further eliminate the detrimental impact of class-misalignment, we introduce a masking scheme on the scoring functions f and f' defined as

$$\begin{aligned} & \hat{d}_{f, \mathcal{F}}(\mathcal{B}_S, \mathcal{B}_T) \\ &= \sup_{f' \in \mathcal{F}} \left(\sum_{\mathcal{B}_T} \text{disp}(f' \odot \omega, f \odot \omega) - \sum_{\mathcal{B}_S} \text{disp}(f' \odot \omega, f \odot \omega) \right), \end{aligned} \quad (10)$$

where $f \odot \omega$ denotes element-wise multiplication between the output of f and ω . The alignment mask ω is a binary vector that denotes whether the i -th class is present in the sampled classes Y (i.e., the classes that we intend to align in the current minibatch). By doing so, we simultaneously align the source and target domains (i) in the input space and (ii) in the functional approximations of the domain divergence by masking the scoring functions f and f' .

4. Experiments

4.1. Setup

Datasets. We evaluate on Office-31, Office-Home and VisDA2017. Office-31 (Saenko et al., 2010) has three domains (Amazon, DSLR and Webcam) with 31 classes. We use three versions of Office-Home (Venkateswara et al., 2017) that contains four domains (Art, Clip Art, Product, and Real-world) with 65 classes: (i) “standard”: the standard Office-Home dataset. (ii) “balanced” (Tan et al., 2019): a subset of the standard dataset where each class has the same number of examples. (iii) “RS-UT”: Reversely-unbalanced Source (RS) and Unbalanced-Target (UT) distribution (Tan et al., 2019) where both domains are imbalanced, but the majority class in the source domain is the minority class in the target domain. VisDA2017 (synthetic→real) (Peng et al., 2017) is a large-scale dataset with 12 classes and more than 200k images.

Model architecture. We use ResNet-50 (He et al., 2016) pre-trained from ImageNet (Russakovsky et al., 2015) as

Table 1. Per-class average accuracy on Office-Home dataset with RS-UT label shifts (ResNet-50).

Methods	Rw→Pr	Rw→Cl	Pr→Rw	Pr→Cl	Cl→Rw	Cl→Pr	Avg
Source Only [†]	69.77	38.35	67.31	35.84	53.31	52.27	52.81
BSP (Chen et al., 2019c) [†]	72.80	23.82	66.19	20.05	32.59	30.36	40.97
PADA (Cao et al., 2018) [†]	60.77	32.28	57.09	26.76	40.71	38.34	42.66
BBSE (Lipton et al., 2018) [†]	61.10	33.27	62.66	31.15	39.70	38.08	44.33
MCD (Saito et al., 2018) [†]	66.03	33.17	62.95	29.99	44.47	39.01	45.94
DAN (Long et al., 2015) [†]	69.35	40.84	66.93	34.66	53.55	52.09	52.90
F-DANN (Wu et al., 2019) [†]	68.56	40.57	67.32	37.33	55.84	53.67	53.88
JAN (Long et al., 2017) [†]	67.20	43.60	68.87	39.21	57.98	48.57	54.24
DANN (Ganin et al., 2016) [†]	71.62	46.51	68.40	38.07	58.83	58.05	56.91
MDD (random sampler)	71.21	44.78	69.31	42.56	52.10	52.70	55.44
MDD (source-balanced sampler)	76.06	47.38	71.56	40.03	57.46	58.54	58.50
COAL (Tan et al., 2019) ^{†,‡}	73.65	42.58	73.26	40.61	59.22	57.33	58.40
MDD+Explicit Alignment (basic) [‡]	69.52	44.70	69.59	40.27	53.02	53.39	55.08
MDD+Explicit Alignment (moving avg.) [‡]	71.37	45.26	69.69	40.28	52.92	52.69	55.37
MDD+Explicit Alignment (curriculum) [‡]	70.02	45.48	69.71	40.86	53.26	52.99	55.39
MDD+Implicit Alignment	76.08	50.04	74.21	45.38	61.15	63.15	61.67

[†] Source: Data of these baseline methods are cited from (Tan et al., 2019).

[‡] Methods using explicit class-conditioned domain alignment.

the backbone, and use hyper-parameters from (Zhang et al., 2019b) for MDD-based domain discrepancy measure. The batch size is 31 for Office-31 and 50 for Office-Home.

Baselines. Our main explicit alignment baselines are COAL (Tan et al., 2019), PACET (Liang et al., 2019b) and MCS (Liang et al., 2019a), state-of-the-art explicit alignment methods based on domain discriminator discrepancy. As our domain discrepancy measure is MDD, we re-implement various MDD-based explicit alignment for fair comparison.

Computational efficiency. We only update pseudo-labels periodically, i.e., every 20 steps, instead of at every training step. We show in the supplementary materials that our method does not require more frequent pseudo-label updates.

4.2. Evaluating Extreme Class Distribution Shift

We use Office-Home (RS-UT), described in Figure 3 (a), to evaluate the performance of different methods under extreme within-domain class imbalance and between-domain class distribution shift where the majority classes in the source domain are minority classes in the target domain. Table 1 presents the per-class average accuracy on Office-Home (RS-UT). Our main baseline is the explicit alignment method “covariate and label shift co-alignment” (COAL) designed to address data imbalance and class distribution shift. Our proposed implicit domain alignment works the best.

4.2.1. THE IMPACT OF CLASS DISTRIBUTION SHIFT

Many baseline methods suffer from class distribution shift, and their performances degrade to “Source Only” training as they do not take into account within-domain class imbalance and between-domain class distribution shift. For MDD-based methods, after we apply balanced sampling for the source domain, the per-class average accuracy improved from 55.44% to 58.50%, which indicates balanced sampling is helpful for

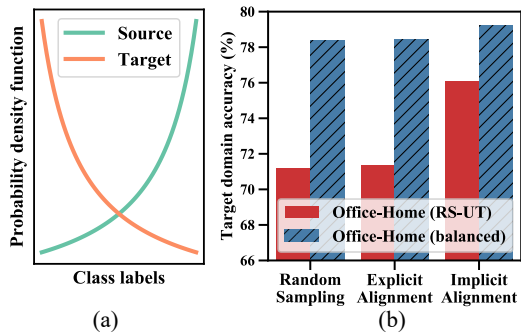


Figure 3. (a) Source and target class distribution of Office-Home (RS-UT). (b) Accuracy comparison between Office-Home (RS-UT) and Office-Home (balanced) for Rw→Pr.

class distribution shift, despite only in the source domain.

4.2.2. THE EFFECTIVENESS OF IMPLICIT ALIGNMENT

The effectiveness of implicit alignment is demonstrated through the comparison between “MDD+Implicit Alignment” and “MDD (source-balanced sampler)”. Both methods use the same sampling procedure for the source. The only difference is that implicit alignment aligns the two domains by sampling aligned classes in the target domain, whereas “source-balanced sampler” only takes random samples from the target domain. Table 1 shows that implicit alignment performs better than “source-balanced sampler” because it is better-aligned, which confirms the effectiveness of implicit alignment. Besides, the proposed method also outperforms MDD-based explicit alignment, which validates the effectiveness of implicit alignment over the explicit alignment.

Figure 3 (b) compares the baseline, implicit and explicit alignments on Office-Home (balanced) and Office-Home (RS-UT). We observe that implicit alignment performs the best on both datasets. More importantly, implicit alignment is more robust to class distribution shift which greatly out-performs other methods under RS-UT

Table 2. Accuracy (%) on Office-31 (standard) for unsupervised domain adaptation (ResNet-50). We repeated each experiment 5 times with different random seeds and report the average and the standard error of the accuracy.

Method	A → W	D → W	W → D	A → D	D → A	W → A	Avg
Source only	68.4±0.2	96.7±0.1	99.3±0.1	68.9±0.2	62.5±0.3	60.7±0.3	76.1
DAN (Long et al., 2015)	80.5±0.4	97.1±0.2	99.6±0.1	78.6±0.2	63.6±0.3	62.8±0.2	80.4
DANN (Ganin et al., 2016)	82.0±0.4	96.9±0.2	99.1±0.1	79.7±0.4	68.2±0.4	67.4±0.5	82.2
ADDA (Tzeng et al., 2017)	86.2±0.5	96.2±0.3	98.4±0.3	77.8±0.3	69.5±0.4	68.9±0.5	82.9
JAN (Long et al., 2017)	85.4±0.3	97.4±0.2	99.8±0.2	84.7±0.3	68.6±0.3	70.0±0.4	84.3
MADA (Pei et al., 2018)	90.0 ± 0.1	97.4±0.1	99.6±0.1	87.8±0.2	70.3±0.3	66.4±0.3	85.2
GTA (Sankaranarayanan et al., 2018)	89.5±0.5	97.9±0.3	99.8±0.4	87.7±0.5	72.8±0.3	71.4±0.4	86.5
MCD (Saito et al., 2018)	88.6±0.2	98.5±0.1	100.0±0.0	92.2±0.2	69.5±0.1	69.7±0.3	86.5
CDAN (Long et al., 2018)	94.1±0.1	98.6±0.1	100.0±0.0	92.9±0.2	71.0±0.3	69.3±0.3	87.7
MDD (Zhang et al., 2019b)	94.5±0.3	98.4±0.1	100.0±0.0	93.5±0.2	74.6±0.3	72.2±0.1	88.9
PACET (Liang et al., 2019b) [‡]	90.8	97.6	99.8	90.8	73.5	73.6	87.4
CAT (Deng et al., 2019) [‡]	94.4±0.1	98.0±0.2	100.0±0.0	90.8±1.8	72.2±0.2	70.2±0.1	87.6
MDD (source-balanced sampler)	90.4±0.4	98.7±0.1	99.9±0.1	90.4±0.2	75.0±0.5	73.7±0.9	88.0
MDD+Explicit Alignment [‡]	92.3±0.1	98.2±0.1	99.8±0.0	92.3±0.3	74.6±0.2	72.9±0.7	88.4
MDD+Implicit Alignment	90.3±0.2	98.7±0.1	99.8±0.0	92.1±0.5	75.3±0.2	74.9±0.3	88.8

[‡] Methods using explicit class-conditioned domain alignment.

Table 3. Accuracy (%) on Office-Home (standard) for unsupervised domain adaptation (ResNet-50).

Method	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
Source only	34.9	50.0	58.0	37.4	41.9	46.2	38.5	31.2	60.4	53.9	41.2	59.9	46.1
DAN (Long et al., 2015)	43.6	57.0	67.9	45.8	56.5	60.4	44.0	43.6	67.7	63.1	51.5	74.3	56.3
DANN (Ganin et al., 2016)	45.6	59.3	70.1	47.0	58.5	60.9	46.1	43.7	68.5	63.2	51.8	76.8	57.6
JAN (Long et al., 2017)	45.9	61.2	68.9	50.4	59.7	61.0	45.8	43.4	70.3	63.9	52.4	76.8	58.3
CDAN (Long et al., 2018)	50.7	70.6	76.0	57.6	70.0	70.0	57.4	50.9	77.3	70.9	56.7	81.6	65.8
BSP (Chen et al., 2019c)	52.0	68.6	76.1	58.0	70.3	70.2	58.6	50.2	77.6	72.2	59.3	81.9	66.3
MDD (Zhang et al., 2019b)	54.9	73.7	77.8	60.0	71.4	71.8	61.2	53.6	78.1	72.5	60.2	82.3	68.1
MCS (Liang et al., 2019a) [‡]	55.9	73.8	79.0	57.5	69.9	71.3	58.4	50.3	78.2	65.9	53.2	82.2	66.3
MDD+Explicit Alignment [‡]	54.3	74.6	77.6	60.7	71.9	71.4	62.1	52.4	76.9	71.1	57.6	81.3	67.7
MDD (source-balanced sampler)	55.3	75.0	79.1	62.3	70.1	73.2	63.5	53.2	78.7	70.4	56.2	82.0	68.3
MDD+Implicit Alignment	56.2	77.9	79.2	64.4	73.1	74.4	64.2	54.2	79.9	71.2	58.1	83.1	69.5

[‡] Methods using explicit class-conditioned domain alignment.

distribution shift and has a smaller performance drop from the balanced version of Office-Home.

4.3. Evaluating Standard Domain Adaptation Datasets

Table 2 and Table 3 summarize the results for the standard Office-31 and Office-Home datasets which have a small degree of class imbalance. Our method outperforms the baselines in 3 out of 6 domain pairs for Office-31, and 10 out of 12 domain pairs for Office-Home (standard). The proposed implicit alignment exhibits larger performance gains on the Office-Home dataset because the dataset is more difficult for domain adaptation, and it has 65 classes compared with the 31 classes in Office-31. We also report state-of-the-art results for VisDA in Table 4.

Similar to findings in Section 4.2, we observe source-balanced sampling is helpful when comparing “MDD (source-balanced sampler)” with the MDD standard baseline, even without extreme class distribution shift.

The proposed method outperforms the state-of-the-art

explicit alignment methods—PACET and MCS—across all domain pairs. We find it ineffective to incorporate prototype-based explicit alignment into MDD. This is in contrast with domain-discriminator-based adversarial learning, where explicit alignment is shown to improve domain adaptation. This is because the classifier-based discrepancy MDD contains more abundant information than domain-discriminator-based discrepancy, owing to the availability of predictive probabilities provided by the classifiers. The rich information in domain discrepancy removes the need for prototype-based distances.

4.4. Ablation studies

4.4.1. IMPACT OF CLASS DIVERSITY AND ALIGNMENT

We analyze the impact of class diversity and alignment by designing experiments along three dimensions: the number of unique labels in each minibatch, whether the classes are aligned, and whether we use pseudo-labels or ground-truth labels when sampling the target domain.




Figure 4. The impact of class diversity and alignment on domain adaptation for Ar→Cl, Office-Home (standard).

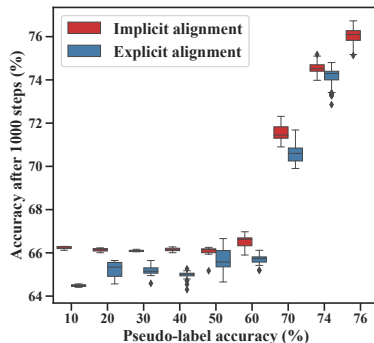


Figure 5. The impact of pseudo-label errors on implicit and explicit alignment, Ar→Cl, Office-Home (standard).

method	acc. (%)
JAN (Long et al., 2017)	61.6
GTA (Sankaranarayanan et al., 2018)	69.5
MCD (Saito et al., 2018)	69.8
CDAN (Long et al., 2018)	70.0
MDD (Zhang et al., 2019b)	74.6
MDD+Explicit Alignment	67.1
MDD+Implicit Alignment	75.8

Table 4. VisDA2017 target accuracy (ResNet-50)

Setup. “Baseline (random)” randomly samples examples of both domains. “Baseline (S-sampled, T-random)” uses N -way sampler for the source domain, and randomly samples the target domain. “Aligned (pseudo-labels)” is the proposed implicit alignment approach. “Aligned (Oracle)” is the oracle form of implicit alignment where the target domain uses ground-truth labels for sampling.

The impact of class diversity. Minibatch-based class diversity determines the sampling distribution of the label space, and a greater diversity corresponds to a more stable measure of this sampling distribution. Figure 4 suggests a positive correlation between the model performance and class diversity: domain adaptation methods do not work well when the class diversity is very low—i.e., only sample 5 classes per batch among the 65 classes—and the alignment-based methods outperform the baseline as we increase class diversity.

The impact of alignment. We confirm the importance of the proposed implicit alignment algorithm from two perspectives. First, “Aligned (oracle)” consistently performs the best, which suggests perfect alignment can provide substantial benefits to unsupervised domain adaptation. Second, the comparison between “Aligned (pseudo-labels)” and “Baseline (S-sampled, T-random)” validates the effectiveness of pseudo-label based implicit alignment, although the pseudo-labels are approximations of the oracle.

Domains	Alignment options		avg. acc.
	masking	sampling	
Rw→Cl	×	×	44.8
	✓	×	44.8
	×	✓	47.4
	✓	✓	50.0
Pr→Rw	×	×	69.3
	✓	×	72.7
	×	✓	72.0
	✓	✓	74.2

Table 5. The impact of different implicit alignment options, i.e., masking in the MDD estimation and sampling class-aligned minibatches, on Office-Home (RS-UT).

4.4.2. ROBUSTNESS TO PSEUDO-LABEL ERRORS

We investigate whether implicit alignment is indeed more robust to pseudo-label errors when compared with explicit alignment. Figure 5 illustrates the relationship between pseudo-label accuracy at training step t and the corresponding subsequent target accuracy at step $t + 1000$, i.e., after 1000 domain adaptation training steps. This process resembles a Markov chain that allows us to analyze the impact of pseudo-label accuracy on the learning dynamics.

It is evident that the drawbacks of explicit alignment are more severe when the pseudo-labels are less accurate, e.g., 10~40%, where implicit alignment has more considerable performance improvements than explicit alignment. This suggests that implicit alignment is more robust to erroneous pseudo-label predictions because it does not require explicit supervision from the pseudo-labels. Implicit and explicit methods eventually converge at 76% and 74%, respectively.

Although many recent techniques attempt to address pseudo-label bias in explicit alignment, they depend on the assumption that probabilities of model predictions are well-calibrated during training. They fail to address ill-calibrated probabilities (Guo et al., 2017), where the model tends to make confident mistakes on the target domain. Moreover, given that models do not initially perform well when training begins, for a random classifier, implicit alignment selects random samples that is equivalent to training without sampling. In contrast, explicit alignment optimizes model parameters from these random labels explicitly.

4.4.3. ABLATION STUDY ON MDD

Table 5 presents the ablation study on Office-Home (RS-UT) that aims to assess the impact of different implicit alignment options: alignment in the domain divergence estimations in Section 3.2.1 (i.e., *masking* in MDD) and alignment in the input space in Section 3.2.1 (i.e., *sampling* class-conditioned examples). We observe that both alignment techniques are essential for domain adaptation because alignment should




Figure 6. Interactions between *within-domain class imbalance* and *between-domain class distribution shift*.

Table 6. Per-class average accuracy (%) with *mismatched prior* where the source domain is balanced while the target domain is imbalanced.

method	SVHN→MNIST		MNIST→SVHN	
	mild	extreme	mild	extreme
source only	67.4±7.3	66.3±3.3	32.5±2.9	28.2±2.3
DANN	78.2±2.8	59.1±0.8	20.9±6.0	20.5±3.1
DANN+implicit	88.6±0.7	82.2±2.1	32.4±2.1	28.9±3.3

be enforced consistently across all aspects of adaptation. We report similar findings, in the supplementary material, on Office-Home (standard).

4.4.4. GENERALIZATION:

IMPLICIT ALIGNMENT ALSO IMPROVES DANN

We design additional experiments to further demonstrate the effectiveness of the proposed approach on a *different domain adaptation algorithm*—DANN—on two synthetic domains with *different degrees of class imbalance*: “mild” (light-tailed class imbalance from a triangular-like distribution) and “extreme” (heavy-tailed class imbalance from a Pareto distribution). We synthetically manipulate the class distributions of SVHN and MNIST to simulate various interactions between *within-domain class imbalance* and *between-domain class distribution shift*. As illustrated in Fig 6, we simulate three types of distribution shift when $p_S(y) \neq p_T(y)$ (i) **source-balanced, target-imbalanced**; (ii) **source-imbalanced, target-balanced**; (iii) **both-imbalanced**.

Table 6, 7 and 8 present the results for the abovementioned scenarios and all experiments are repeated five times. The proposed implicit alignment approach significantly improves the performance of DANN regardless of the degree of imbalance or the type of distribution shift. Besides, implicit alignment offers greater improvements over DANN when the degree of imbalance is more severe, i.e., comparing “mild” with “extreme”. Implicit alignment overcomes this limitation of DANN and greatly improves the performance of the challenging task between SVHN and MNIST. We conclude that the proposed approach is independent of the choice of domain adaptation algorithms and helps both MDD and DANN.

Note that the aim of this subsection is to show that implicit

Table 7. Per-class average accuracy (%) with *mismatched prior* where the source domain is imbalanced while the target domain is balanced.

method	SVHN→MNIST		MNIST→SVHN	
	mild	extreme	mild	extreme
source only	65.2±2.1	53.3±1.3	31.6±3.3	32.8±0.9
DANN	82.0±0.7	52.3±2.3	23.4±3.6	25.9±0.5
DANN+implicit	91.0±1.9	87.1±2.6	34.9±0.5	31.1±2.9

Table 8. Per-class average accuracy (%) with *mismatched prior* where both domains are imbalanced.

method	SVHN→MNIST		MNIST→SVHN	
	mild	extreme	mild	extreme
source only	60.9±5.2	51.2±5.9	30.6±1.3	27.1±1.7
DANN	67.6±0.8	40.5±5.5	23.4±1.6	18.8±2.9
DANN+implicit	88.6±0.6	70.5±3.6	36.3±2.5	27.9±2.4

alignment could help improve DANN on the digits dataset. We are currently working on integrating implicit alignment into the current state-of-the-art methods (Kumar et al., 2018; Shu et al., 2018; Wu et al., 2019) on this dataset.

5. Related Work

We review related work on unsupervised domain adaptation and discuss their relations with our proposed method.

Instance-based importance-weighting (Chawla et al., 2002; Kouw & Loog, 2019; Azzizadenesheli et al., 2019) aims to minimize the target error directly from the source domain data, weighted at the example level or class level. It first estimates of the probability of a source example, or class, belonging to the target domain (Iyer et al., 2014; Lipton et al., 2018), then uses that probability as important-weighting to train the classifier from the source domain data. Unlike our approach, importance-weighting only uses the source data to train the classifier without learning domain invariant representations.

Feature-based distribution adaptation is the prevailing approach to domain adaptation that aims to minimize the distribution discrepancy between the source and target domains. The domain difference can be measured in various ways, such as Maximum Mean Discrepancy (MMD) (Borgwardt et al., 2006), which is further minimized to achieve domain invariance. The minimization of such discrepancy can be carried out by directly minimizing the distance (Tzeng et al., 2014) or with the help of adversarial learning (Ganin et al., 2016).

Classifier-based distribution adaptation is a strong competitor to feature-based adaptation. It aims to minimize the discrepancy between two classifiers so that the learned representations respect the decision boundary of the

classification task (Saito et al., 2018; Zhang et al., 2019b). We show that the proposed approach is beneficial to both classifier-based discrepancy MDD (Zhang et al., 2019b) and feature-based discrepancy DANN (Ganin et al., 2016).

Feature-classifier joint distribution adaptation aims to align the joint distribution between features and their corresponding predictions (Long et al., 2013; Tsai et al., 2018). The joint distribution can be represented in a multilinear map between features and classifier predictions (Long et al., 2018), or the Cartesian product between the domain space and class space (Cicek & Soatto, 2019). In our work, we implicitly align the joint distribution with the factorization $p(x, y) = p(x|y)p(y)$ from a sampling perspective where $p(y)$ is the pre-specified alignment distribution in the label space, and $p(x|y)$ represents class-conditioned sampling.

Explicit class-conditioned domain alignment, or class prototype alignment, introduces a loss function that minimizes the distances of class-level prototypes between the source and target domains (Snell et al., 2017; Pinheiro, 2018; Pan et al., 2019; Deng et al., 2019). It is prone to error accumulation due to its reliance on explicit optimization of model parameters from the pseudo-labels. A variety of recent methods have been proposed to mitigate these limitations by estimating batch-level statistics (Xie et al., 2018) and introducing an easy-to-hard curriculum that favors confident predictions (Chen et al., 2019a). Nevertheless, these algorithms suffer from ill-calibrated probabilities in the form of confident mistakes, and more work is needed to improve model calibration so as to better utilize explicit alignment.

Self-training (Nigam & Ghani, 2000) is a special form of co-training (Blum & Mitchell, 1998) where the model iteratively uses its predictions, i.e., pseudo-labels, as explicit supervision to re-train itself. The use of pseudo-labels has become an emerging trend in domain adaptation, because they provide estimations of the target domain label distribution that can be exploited by training algorithms. Apart from class prototype based methods (Chen et al., 2011; Saito et al., 2017; Zhang et al., 2018; Deng et al., 2019) for explicit alignment, (Wen et al., 2019) proposed the use of uncertainty estimates of the target domain predictions as second-order statistics to promote feature-label joint adaptation. For semantic segmentation tasks, (Zou et al., 2018) proposed to iteratively generate pseudo-labels in the target domain and re-train the model on these labels; (Zhang et al., 2019a) proposed to use pseudo-labels to encourage examples to cluster together if they belong to the same class; (Chen et al., 2019b) applied entropy minimization (Grandvalet & Bengio, 2005) on the pseudo-labels to encourage class overlap between domains. A main bottleneck for this approach is the bias in pseudo-label predictions. Directly optimizing these labels is prone to “entropy over-minimization” (Zou et al., 2019) and negative transfer (Lifshitz & Wolf, 2020) where the

model overfits to mistakes in the pseudo-labels. Moreover, the pseudo-labels are likely to suffer from ill-calibrated probabilities (Guo et al., 2017), especially for deep learning methods. The resulting misleadingly confident mistakes exacerbate the critical problem of error accumulation in pseudo-label bias. In contrast, our proposed method removes the need for direct supervision from pseudo-labels, and as a result is more robust to bias in how these labels are produced.

Reinforced sample selection (Dong & Xing, 2018) is proposed for one-shot domain adaptation where a model actively selects labeled examples to train the domain adaptation model. In comparison, the advantage of our approach is in its simplicity that no reinforcement learning is required to obtain the sampling strategy.

6. Conclusion and Future Work

We introduce an approach for unsupervised domain adaptation—with a strong focus on practical considerations of within-domain class imbalance and between-domain class distribution shift—from a class-conditioned domain alignment perspective. We show theoretically that the proposed implicit alignment provides a more reliable measure of empirical domain divergence which facilitates adversarial domain-invariant representation learning, that would otherwise be hampered by the class-misaligned domain divergence. We show that our proposed approach leads to superior UDA performance under extreme within-domain class imbalance and between-domain class distribution shift, as well as competitive results on standard UDA tasks. We emphasize that the proposed method is robust to pseudo-label bias, simple to implement, has a unified training objective, and does not require additional parameter tuning. We also show that the proposed approach is orthogonal to the choice of domain adaptation algorithms and offers consistent improvements to feature-based and classifier-based domain adaptation algorithms.

Future work includes extensions to cost-sensitive learning for domain adaptation to account for the misclassification cost of different classes under both covariate and class distribution shift. It is also important to analyze the probability calibration of different domain adaptation models and develop well-calibrated methods for more effective use of pseudo-labels.

Acknowledgements

We thank the anonymous reviewers for providing thoughtful feedback. The authors also thank Lisa Di Jorio, Tanya Nair, Francis Dutil, Cecil Low-Kam, Nicolas Chapados, and the Imagia team for their support. Xiang Jiang acknowledges the support of NVIDIA Corporation with the donation of the Titan X GPU used for this research.

References

- Azizzadenesheli, K., Liu, A., Yang, F., and Anandkumar, A. Regularized learning for domain adaptation under label shifts. In *International Conference on Learning Representations*, 2019. URL <https://openreview.net/forum?id=rJl0r3R9KX>.
- Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine learning*, 79(1-2):151–175, 2010.
- Blum, A. and Mitchell, T. Combining labeled and unlabeled data with co-training. In *Proceedings of the eleventh annual conference on Computational learning theory*, pp. 92–100. ACM, 1998.
- Borgwardt, K. M., Gretton, A., Rasch, M. J., Kriegel, H.-P., Schölkopf, B., and Smola, A. J. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006.
- Cao, Z., Ma, L., Long, M., and Wang, J. Partial adversarial domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 135–150, 2018.
- Chawla, N. V. Data mining for imbalanced datasets: An overview. In *Data mining and knowledge discovery handbook*, pp. 875–886. Springer, 2009.
- Chawla, N. V., Bowyer, K. W., Hall, L. O., and Kegelmeyer, W. P. Smote: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 16:321–357, 2002.
- Chen, C., Xie, W., Huang, W., Rong, Y., Ding, X., Huang, Y., Xu, T., and Huang, J. Progressive feature alignment for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 627–636, 2019a.
- Chen, M., Weinberger, K. Q., and Blitzer, J. Co-training for domain adaptation. In *Advances in neural information processing systems*, pp. 2456–2464, 2011.
- Chen, M., Xue, H., and Cai, D. Domain adaptation for semantic segmentation with maximum squares loss. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2090–2099, 2019b.
- Chen, X., Wang, S., Long, M., and Wang, J. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In *International Conference on Machine Learning*, pp. 1081–1090, 2019c.
- Cicek, S. and Soatto, S. Unsupervised domain adaptation via regularized conditional alignment. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.
- Deng, Z., Luo, Y., and Zhu, J. Cluster alignment with a teacher for unsupervised domain adaptation. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 9944–9953, 2019.
- Dong, N. and Xing, E. P. Domain adaption in one-shot learning. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 573–588. Springer, 2018.
- Ganin, Y., Ustinova, E., Ajakan, H., Germain, P., Larochelle, H., Laviolette, F., Marchand, M., and Lempitsky, V. Domain-adversarial training of neural networks. *The Journal of Machine Learning Research*, 17(1):2096–2030, 2016.
- Grandvalet, Y. and Bengio, Y. Semi-supervised learning by entropy minimization. In *Advances in neural information processing systems*, pp. 529–536, 2005.
- Guo, C., Pleiss, G., Sun, Y., and Weinberger, K. Q. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 1321–1330. JMLR. org, 2017.
- He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Heckman, J. J. Sample selection bias as a specification error. *Econometrica: Journal of the econometric society*, pp. 153–161, 1979.
- Iyer, A., Nath, S., and Sarawagi, S. Maximum mean discrepancy for class ratio estimation: Convergence bounds and kernel selection. In *International Conference on Machine Learning*, pp. 530–538, 2014.
- Japkowicz, N. and Stephen, S. The class imbalance problem: A systematic study. *Intelligent data analysis*, 6(5): 429–449, 2002.
- Kouw, W. M. and Loog, M. A review of domain adaptation without target labels. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- Kumar, A., Sattigeri, P., Wadhawan, K., Karlinsky, L., Feris, R., Freeman, B., and Wornell, G. Co-regularized alignment for unsupervised domain adaptation. In *Advances in Neural Information Processing Systems*, pp. 9345–9356, 2018.
- Liang, J., He, R., Sun, Z., and Tan, T. Distant supervised centroid shift: A simple and efficient approach to visual domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2975–2984, 2019a.

- Liang, J., He, R., Sun, Z., and Tan, T. Exploring uncertainty in pseudo-label guided unsupervised domain adaptation. *Pattern Recognition*, 96:106996, 2019b.
- Lifshitz, O. and Wolf, L. A sample selection approach for universal domain adaptation. *arXiv preprint arXiv:2001.05071*, 2020.
- Lipton, Z. C., Wang, Y.-X., and Smola, A. Detecting and correcting for label shift with black box predictors. *arXiv preprint arXiv:1802.03916*, 2018.
- Long, M., Wang, J., Ding, G., Sun, J., and Yu, P. S. Transfer feature learning with joint distribution adaptation. In *Proceedings of the IEEE international conference on computer vision*, pp. 2200–2207, 2013.
- Long, M., Cao, Y., Wang, J., and Jordan, M. I. Learning transferable features with deep adaptation networks. In *Proceedings of the 32nd International Conference on International Conference on Machine Learning-Volume 37*, pp. 97–105. JMLR. org, 2015.
- Long, M., Zhu, H., Wang, J., and Jordan, M. I. Deep transfer learning with joint adaptation networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2208–2217. JMLR. org, 2017.
- Long, M., Cao, Z., Wang, J., and Jordan, M. I. Conditional adversarial domain adaptation. In *Advances in Neural Information Processing Systems*, pp. 1640–1650, 2018.
- Luo, Z., Zou, Y., Hoffman, J., and Fei-Fei, L. F. Label efficient learning of transferable representations across domains and tasks. In *Advances in Neural Information Processing Systems*, pp. 165–177, 2017.
- Nigam, K. and Ghani, R. Analyzing the effectiveness and applicability of co-training. In *Cikm*, volume 5, pp. 3, 2000.
- Pan, S. J. and Yang, Q. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10): 1345–1359, 2009.
- Pan, Y., Yao, T., Li, Y., Wang, Y., Ngo, C.-W., and Mei, T. Transferrable prototypical networks for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 2239–2247, 2019.
- Pei, Z., Cao, Z., Long, M., and Wang, J. Multi-adversarial domain adaptation. In *Thirty-Second AAAI Conference on Artificial Intelligence*, 2018.
- Peng, X., Usman, B., Kaushik, N., Hoffman, J., Wang, D., and Saenko, K. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- Pinheiro, P. O. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8004–8013, 2018.
- Quionero-Candela, J., Sugiyama, M., Schwaighofer, A., and Lawrence, N. D. *Dataset shift in machine learning*. The MIT Press, 2009.
- Russakovsky, O., Deng, J., Su, H., Krause, J., Satheesh, S., Ma, S., Huang, Z., Karpathy, A., Khosla, A., Bernstein, M., et al. Imagenet large scale visual recognition challenge. *International journal of computer vision*, 115 (3):211–252, 2015.
- Saenko, K., Kulis, B., Fritz, M., and Darrell, T. Adapting visual category models to new domains. In *European conference on computer vision*, pp. 213–226. Springer, 2010.
- Saito, K., Ushiku, Y., and Harada, T. Asymmetric tri-training for unsupervised domain adaptation. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pp. 2988–2997. JMLR. org, 2017.
- Saito, K., Watanabe, K., Ushiku, Y., and Harada, T. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3723–3732, 2018.
- Sankaranarayanan, S., Balaji, Y., Castillo, C. D., and Chellappa, R. Generate to adapt: Aligning domains using generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 8503–8512, 2018.
- Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2): 227–244, 2000.
- Shu, R., Bui, H. H., Narui, H., and Ermon, S. A dirt-t approach to unsupervised domain adaptation. *arXiv preprint arXiv:1802.08735*, 2018.
- Snell, J., Swersky, K., and Zemel, R. Prototypical networks for few-shot learning. In *Advances in Neural Information Processing Systems*, pp. 4077–4087, 2017.
- Tan, S., Peng, X., and Saenko, K. Generalized domain adaptation with covariate and label shift co-alignment. *arXiv preprint arXiv:1910.10320*, 2019.
- Torralba, A., Efros, A. A., et al. Unbiased look at dataset bias. In *CVPR*, volume 1, pp. 7. Citeseer, 2011.
- Tsai, Y.-H., Hung, W.-C., Schuster, S., Sohn, K., Yang, M.-H., and Chandraker, M. Learning to adapt structured

- output space for semantic segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7472–7481, 2018.
- Tzeng, E., Hoffman, J., Zhang, N., Saenko, K., and Darrell, T. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- Tzeng, E., Hoffman, J., Saenko, K., and Darrell, T. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7167–7176, 2017.
- Venkateswara, H., Eusebio, J., Chakraborty, S., and Panchanathan, S. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5018–5027, 2017.
- Webb, G. I. and Ting, K. M. On the application of roc analysis to predict classification performance under varying class distributions. *Machine learning*, 58(1):25–32, 2005.
- Wen, J., Zheng, N., Yuan, J., Gong, Z., and Chen, C. Bayesian uncertainty matching for unsupervised domain adaptation. *arXiv preprint arXiv:1906.09693*, 2019.
- Wu, Y., Winston, E., Kaushik, D., and Lipton, Z. Domain adaptation with asymmetrically-relaxed distribution alignment. *arXiv preprint arXiv:1903.01689*, 2019.
- Xie, S., Zheng, Z., Chen, L., and Chen, C. Learning semantic representations for unsupervised domain adaptation. In *International Conference on Machine Learning*, pp. 5419–5428, 2018.
- Zhang, Q., Zhang, J., Liu, W., and Tao, D. Category anchor-guided unsupervised domain adaptation for semantic segmentation. In *Advances in Neural Information Processing Systems*, pp. 433–443, 2019a.
- Zhang, W., Ouyang, W., Li, W., and Xu, D. Collaborative and adversarial network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3801–3809, 2018.
- Zhang, Y., Liu, T., Long, M., and Jordan, M. Bridging theory and algorithm for domain adaptation. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7404–7413, Long Beach, California, USA, 09–15 Jun 2019b. PMLR.
- Zou, Y., Yu, Z., Vijaya Kumar, B., and Wang, J. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 289–305, 2018.
- Zou, Y., Yu, Z., Liu, X., Kumar, B. V., and Wang, J. Confidence regularized self-training. In *The IEEE International Conference on Computer Vision (ICCV)*, October 2019.