

Understanding and Robustifying Sub-domain Alignment for Domain Adaptation

Anonymous authors

Paper under double-blind review

Abstract

In unsupervised domain adaptation (UDA), aligning source and target domains improves the predictive performance of learned models on the target domain. A common methodological improvement in alignment methods is to divide the domains and align sub-domains instead. These sub-domain-based algorithms have demonstrated great empirical success but lack theoretical support. In this work, we establish a rigorous theoretical understanding of the advantages of these methods that have the potential to enhance their overall impact on the field. Our theory uncovers that sub-domain-based methods optimize an error bound that is at least as strong as non-sub-domain-based error bounds and is empirically verified to be much stronger. Furthermore, our analysis indicates that when the marginal weights of sub-domains shift between source and target tasks, the performance of these methods may be compromised. We therefore implement an algorithm to robustify sub-domain alignment for domain adaptation under sub-domain shift, offering a valuable adaptation strategy for future sub-domain-based methods. Empirical experiments across various benchmarks validate our theoretical insights, prove the necessity for the proposed adaptation strategy, and demonstrate the algorithm’s competitiveness in handling label shift.

1 Introduction

Supervised deep learning has achieved unprecedented success in a wide range of real-world applications. However, obtaining labeled data may be costly, labor-intensive, and/or time-consuming in certain applications, particularly in medical and biological domains (Lu et al., 2017; Li et al., 2020). To this end, unsupervised domain adaptation (UDA) transfers knowledge from a labeled source domain to a different but related unlabeled target domain (Farahani et al., 2021). However, efficient UDA is challenging due to the statistical discrepancies between two domains, hereafter referred to as *domain shift* (Wang & Deng, 2018; Sankaranarayanan et al., 2018; Deng et al., 2019). To address this challenge, much of the UDA research has focused on reducing the distributional gap between the source and target domains (Shen et al., 2018; Liu et al., 2016; Isola et al., 2017; Tzeng et al., 2015; 2017; 2020; Ganin & Lempitsky, 2015; Ganin et al., 2016; Peng et al., 2018). Recent methods further partition the data into sub-domains and align the sub-domains instead (Pinheiro, 2018; Long et al., 2018; Deng et al., 2019). In the literature, a *domain* refers to the data distribution of the covariates X (e.g., images from a specific camera), characterized by its feature space $\mathcal{X} \in X$ and a marginal probability distribution. Within a domain, a *class* refers to a specific category (e.g., ‘cat’, ‘dog’). In this work, we mainly use the term *sub-domain* to refer to the overall domain distribution of each class, such that the k -th sub-domain represents the domain conditioned on class label k , i.e., $X|Y = k \sim P(X|Y = k)$. This allows us to perform fine-grained alignment based on different classes, viewed through their conditional distributions.

Notably, while one straightforward definition of the sub-domains is the conditional distributions based on the classification label, other strategies for defining sub-domains include cross-domain adaptive clustering (Li et al., 2021b), classifier-based backprop-induced weighting (Westfechtel et al., 2023), domain consensus clustering (Li et al., 2021a), joint learning of domain-invariant features and classifiers (Shi & Sha, 2012), and the use of deep clustering (Gao et al., 2020). These sub-domain-based algorithms have shown substantial empirical success. However, the benefits of sub-domain alignments have not been rigorously justified.

In this work, we present a theoretical analysis to establish that the sub-domain based methods are in fact optimizing a generalization bound that is at least as strong as (and empirically much stronger than) the full-domain-based objective functions. Our analysis further reveals that when the marginal weights of the sub-domains shift between source and target, the sub-domain based methods can fail. We then present a novel UDA algorithm, *Domain Adaptation via Rebalanced Sub-domain Alignment* (DARSA), that is motivated by our analysis and addresses the case when sub-domain marginal weights shift. DARSA optimizes reweighted classification error and discrepancy between sub-domains of the source and target tasks. The reweighting scheme follows a simple intuition: *important sub-domains in the target domain need more attention*. To illustrate the concept visually, Figure 1 highlights the strengths of sub-domain alignment, providing insight into how our method operates and the benefits it brings.

The contribution of our work is two-fold:

- **Theoretical Contribution:** Our work analyzes and provides a theoretical foundation for sub-domain based methods in domain adaptation, addressing their previous lack of rigorous understanding. Our theoretical framework not only supports our algorithm but can be extended to other methods, contributing to broader impact and value in the field.
- **Algorithmic Contribution:** Our theoretical analysis leads to our algorithm DARSA. DARSA addresses shifted marginal sub-domain weights, which adversely impact existing sub-domain-based methods. We empirically verify its competitive performance under label shifting on various benchmarks, confirming our theoretical insights and validating the proposed adaptation strategy.

2 Related Work

Discrepancy-based Domain Adaptation. UDA commonly tries to reduce the distribution gap between the source and target domains. One approach to achieve this is discrepancy-based methods in the extract feature space (Tzeng et al., 2014; Long et al., 2015; Sun et al., 2016), which often use maximum mean discrepancy (MMD) (Borgwardt et al., 2006). While MMD is a well-known Reproducing Kernel Hilbert Space (RKHS) metric, it is weaker than the Wasserstein-1 distance (Lu & Lu, 2020). Therefore, we use Wasserstein-1 distance in our work. Furthermore, many discrepancy-based methods enforce the sharing of the first few layers of the networks between the source and target domains (HassanPour Zonoozi & Seydi, 2022). In contrast, our method allows a more flexible feature space.

Adversarial-based Domain Adaptation. Adversarial-based domain adaptation methods aim to encourage domain similarity through adversarial learning (Shen et al., 2018; Liu et al., 2016; Isola et al., 2017; Tzeng et al., 2015; 2017; 2020; Ganin & Lempitsky, 2015; Ganin et al., 2016; Peng et al., 2018; Hoffman et al., 2018). These methods are divided into generative methods, which combine discriminative models with a generating process, and non-generative methods, which use a domain confusion loss to learn domain-invariant discriminative features (Wang & Deng, 2018). However, many existing algorithms fail to align multi-modal distributions under label shifting scenarios. Additionally, training adversarial networks can be challenging due to mode collapse and oscillations (Liang et al., 2018).

Sub-domain-based Domain Adaptation. The use of sub-domain adaptation has proven effective in aligning multi-modal distributions, enhancing performance across various tasks (Deng et al., 2019; Long et al., 2018; Pinheiro, 2018; Shi & Sha, 2012; Jiang et al., 2020; Snell et al., 2017). (Deng et al., 2019) introduces the Cluster Alignment with a Teacher (CAT) approach that aligns class-conditional structures across domains. (Long et al., 2018) offers conditional adversarial domain adaptation, enhancing alignment through classifier predictions. (Pinheiro, 2018) proposes an unsupervised domain adaptation approach based on similarity learning, wherein classification is conducted by computing similarities between target domain images and prototype representations of each category. On the other hand, (Shi & Sha, 2012) introduces a method that concurrently learns domain-invariant features and classifiers. (Jiang et al., 2020) elucidates a sampling-based implicit alignment technique, addressing concerns of class imbalance. (Snell et al., 2017) presents prototypical networks designed for few-shot classification, employing distances to class prototype representations for the process. While these methods have demonstrated empirical success, a detailed theoretical perspective on the

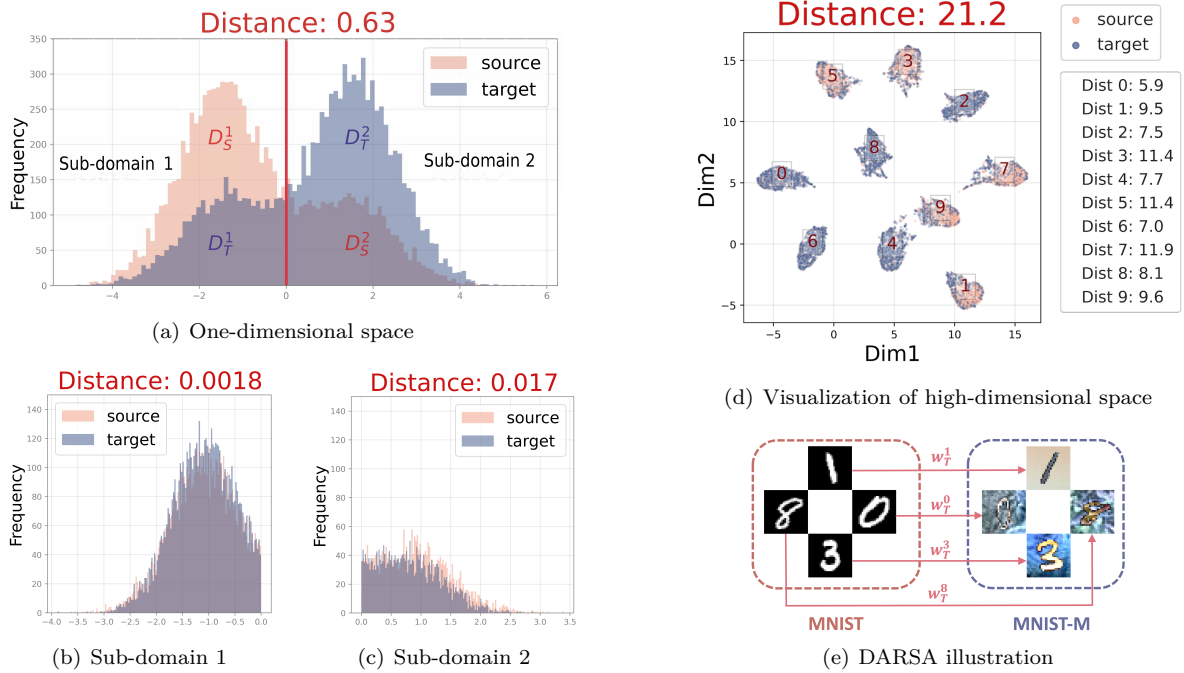


Figure 1: Conceptual overview of our motivation for sub-domain alignment. Listed distances are Wasserstein-1 distances. **(a): Domain Alignment.** The source domain (coral) consists of two Gaussians centered at -1.5 and 1.5 with weights 0.7 and 0.3 , respectively. The target domain (darkblue) is a mixture of two Gaussians centered at -1.4 and 1.6 with inverse weights. For illustrative purposes, we divide the source domain (coral) and target domain (darkblue) into partitions, i.e., sub-domains D_S^1 and D_S^2 for the source domain and sub-domains D_T^1 and D_T^2 for the target domain. Note that while our work formally defines sub-domains as conditional distributions (e.g., based on class labels where applicable), this initial illustration demonstrates the broader concept of aligning finer-grained structures within domains by conceptually partitioning the domains into sub-domains at $x = 0$. **(b-c): Sub-domain Alignments.** The distances between these illustrated sub-domain partitions are trivial compared to the overall domain distance in (a). This serves as conceptual evidence for the benefits of sub-domain alignment. **(d): MNIST to MNIST-M UDA task.** The features are projected to 2-D with UMAP. The legend indicates distances between corresponding sub-domains (with red sub-domain indices labeled in the figure), and the sub-figure title shows the overall domain distance. Here, the sub-domains are defined based on the class (here digits), resulting in the depicted alignment. These sub-domain distances are small compared to the overall distance given at the top. **(e): DARSA illustration** with w_T^k indicating target sub-domain weights, showing DARSA’s applicability under label shifting. This demonstrates our approach using sub-domains defined by digits.

benefits of incorporating sub-domain structures has yet to be fully explored. Our work aims to complement these existing methodologies by providing a comprehensive theoretical understanding of the advantages inherent in these structures.

Theoretical Analysis of Domain Adaptation. Many existing domain adaptation methods are inspired by generalization bounds based on the \mathcal{H} -divergence (Ben-David et al., 2006). The \mathcal{H} -divergence (Ben-David et al., 2006) is a modified version of the total variation distance (L_1) that restricts the hypothesis to a given class. These generalization bounds can be estimated by learning a domain classifier with a finite Vapnik–Chervonenkis (VC) dimension. However, this results in a loose bound for most neural networks (Li et al., 2018). In our method, we use the Wasserstein distance for two reasons. First, the Wasserstein-1 distance is bounded above by the total variation distance (Ben-David et al., 2010). Additionally, the Wasserstein-1 distance is bounded above by the Kullback-Leibler divergence (a special case of the Rényi divergence when α goes to 1 (Fournier & Guillin, 2015)), giving stronger bounds than those presented by Redko et al (Redko et al., 2017) and Mansour et al (Mansour et al., 2012). Additionally, the Wasserstein distance has stable gradients even when the compared distributions are far apart (Gulrajani et al., 2017).

Expanding Theoretical Insights into Domain Adaptation. Our work contributes to the understanding and improvement of sub-domain alignment methods, a type of popular but yet to be rigorously investigated domain adaptation method. In contrast to our work, (Mansour et al., 2009) studies the adaptation performance of various loss functions and models; (Dhouib et al., 2020) focuses on the margin violation rate; (Wang et al., 2022) addresses the problem of learning features that align with human understanding of data; (Zhang et al., 2019) proposes generalization theory for classifiers with scoring function and margin loss; (Germain et al., 2016) studies the generalization theory for the weighted majority vote framework; (Blanchard et al., 2021; Albuquerque et al., 2019; Zhao et al., 2018) focus on the setting with multiple source domain.

3 Preliminaries

Assume a labeled source dataset $\{(x_S^i, y_S^i)\}_{i=1}^{N_S}$ from a source domain X_S with distribution P_S and an unlabeled target dataset $\{x_T^i\}_{i=1}^{N_T}$ from a target domain X_T with distribution P_T . The source dataset has N_S labeled samples, and the target dataset has N_T unlabeled samples. We assume that the samples $x_S^i \in \mathcal{X} \subseteq \mathbb{R}^d$ and $x_T^i \in \mathcal{X} \subseteq \mathbb{R}^d$ are independently drawn from P_S and P_T , respectively. The goal is to learn a classifier $f(x)$ that predicts labels $\{y_T^i\}_{i=1}^{N_T}$ for the target dataset. We further assume that P_S and P_T are probability densities of Borel probability measures in the Wasserstein space $\mathcal{P}_1(\mathbb{R}^d)$, i.e., the space of probability measures with finite first moment.

Sub-domains. We assume that both X_S and X_T are mixtures of K sub-domains. In other words, we have $P_S = \sum_{k=1}^K w_S^k P_S^k$ and $P_T = \sum_{k=1}^K w_T^k P_T^k$ where we use P_S^k and P_T^k to respectively represent the distribution of the k -th sub-domain of the source domain and that of the target domain, and w_S^k/w_T^k correspond to the weights of each sub-domain. Note that $\mathbf{w}_S \doteq [w_S^1, \dots, w_S^K]$ and $\mathbf{w}_T \doteq [w_T^1, \dots, w_T^K]$ belong to Δ_K (the $K-1$ probability simplex). It is straightforward to define sub-domains as conditional distributions, such that the k -th sub-domain is represented as $P_S^k = P(X_S|Y_S = k)$ and $P_T^k = P(X_T|Y_T = k)$, where Y_S and Y_T are the source and target labels, respectively. Note that sub-domain k refers to all the covariates distribution where the class label is k , not just the class itself. However, we note that the framework presented in this work is applicable across various sub-domain methods, that is, any approach assuming that the covariates are mixtures of distribution with $P(X) = \sum_k w^k P^k(X)$ where P^k is the component distribution. While this work focuses on the cases where $P^k(X) = P(X|Y = k)$, all of our theoretical results and algorithms easily extend to other sub-domain methods with different sub-domain definitions.

Probabilistic Classifier Discrepancy. For a distribution \mathcal{D} , we define the discrepancy between two functions f and g as:

$$\gamma_{\mathcal{D}}(f, g) = \mathbb{E}_{x \sim \mathcal{D}} [|f(x) - g(x)|].$$

We use g_T and g_S to represent the true labeling functions of the target and source domains, respectively. We use $\gamma_S(f) \doteq \gamma_{P_S}(f, g_S)$ and $\gamma_T(f) \doteq \gamma_{P_T}(f, g_T)$ to respectively denote the discrepancies of a hypothesis f to the true labeling function for the source and target domains.

Wasserstein Distance. The Kantorovich-Rubenstein dual representation of the Wasserstein-1 distance (Villani, 2009) between two distributions P_S and P_T is defined as

$$W_1(P_S, P_T) = \sup_{\|f\|_L \leq 1} \mathbb{E}_{x \sim P_S} [f(x)] - \mathbb{E}_{x \sim P_T} [f(x)],$$

where the supremum is over the set of 1-Lipschitz functions (all Lipschitz functions f with Lipschitz constant $L \leq 1$). For notational simplicity, we use $D(X_1, X_2)$ to denote a distance between the distributions of any pair of random variables X_1 and X_2 . For instance, $W_1(\Phi(X_S), \Phi(X_T))$ denotes the Wasserstein-1 distance between the distributions of the random variables $\Phi(X_S)$ and $\Phi(X_T)$ for any transformation Φ .

4 Understanding Sub-domain-based Methods

We now present our theoretical analysis of sub-domain-based methods. We first present a generalization bound for domain adaptation that is closely related to existing work, and then establish a novel generalization bound for sub-domain-based methods, aligning with the objectives used by these existing methods. Furthermore, we demonstrate that the sub-domain-based generalization bound is at least as strong as the

non-sub-domain-based generalization bound, which establishes a rigorous theoretical understanding of the advantages of these methods. Our analysis also uncovers that when the marginal weights of sub-domains shift between the source and the target task, sub-domain methods can potentially fail.

4.1 Generalization Bounds for Domain Adaptation

Before presenting our novel theoretical results about sub-domain-based domain adaptation, we first present an upper bound closely related to Ben-David et al. (2010) and Li et al. (2018) Theorem A.8. It is worth noting that we use the Wasserstein-1 distance in our analysis, as it provides a stronger bound than the total variation distance Redko et al. (2017) employed by Ben-David et al. (2010). The proof of Theorem 4.1 is deferred to the Appendix A.4.

Theorem 4.1 (Full Domain Generalization Bound). *For a hypothesis $f : \mathcal{X} \rightarrow [0, 1]$,*

$$\gamma_T(f) \leq \gamma_S(f) + (\lambda + \lambda_H)W_1(P_S, P_T) + \gamma^*, \quad (1)$$

where $\gamma^* = \min_{f \in \mathbb{H}} \gamma_S(f) + \gamma_T(f)$, \mathbb{H} is a hypothesis class included in the set of λ_H -Lipschitz functions, and the true functions g_T and g_S are both λ -Lipschitz functions (as defined in Appendix A.1).

Remark 4.2. The upper bound in Theorem 4.1 consists of three components: (i) $\gamma_S(f)$ is the performance of the hypothesis on the source domain, (ii) $W_1(P_S, P_T)$ is the distance between the source and the target domains, and (iii) γ^* is a constant related to the difference between the source and the target problems that cannot be addressed by domain adaptation. For succinctness and clarity of the following analysis, we assume without loss of generality that $\lambda + \lambda_H \leq 1$, simplifying the bound to

$$\gamma_T(f) \leq \gamma_S(f) + W_1(P_S, P_T) + \gamma^*. \quad (2)$$

Numerous works attempt to solve the domain adaptation problem by designing algorithms that minimize similar generalization bounds to the one in equation 2, e.g., Theorem 1 in Ben-David et al. (2010). These approaches consist of two components: (i) a mapping $\Phi : \mathcal{X} \rightarrow \mathcal{H}$ that transforms the original problem by embedding X_S and X_T into a shared hidden space \mathcal{H} , and (ii) a hypothesis $h : \mathcal{H} \rightarrow [0, 1]$ for prediction. Since $\gamma_T(h \circ \Phi) = \gamma_{\Phi(X_T)}(h)$, with Theorem 4.1, we have a generalization bound of the function $h \circ \Phi : \mathcal{X} \rightarrow [0, 1]$ on the original target problem:

$$\gamma_T(h \circ \Phi) = \gamma_{\Phi(X_T)}(h) \leq \gamma_{\Phi(X_S)}(h) + W_1(\Phi(X_S), \Phi(X_T)) + \gamma_{\Phi}^*. \quad (3)$$

If the distance between $\Phi(X_S)$ and $\Phi(X_T)$, i.e., $W_1(\Phi(X_S), \Phi(X_T))$, is close and the classification error of h on the transformed source problem, i.e., $\gamma_{\Phi(X_S)}(h)$, remains low, then the performance of the hypothesis $h \circ \Phi$ on the original target problem can be guaranteed. This motivation has led to a variety of domain adaptation frameworks with objectives of the following format:

$$\min_{\substack{\Phi: \mathcal{X} \rightarrow \mathcal{H} \\ h: \mathcal{H} \rightarrow [0, 1]}} \gamma_{\Phi(X_S)}(h) + \alpha D(\Phi(X_S), \Phi(X_T)), \quad (4)$$

where $\gamma_{\Phi(X_S)}(h)$ is the classification error of h on the transformed source problem, D is a distance between distributions and α is the balancing weight. In this work, we use Wasserstein-1 distance.

4.2 Analysis of Sub-domain-based Methods

We first present several results that will be used to build the main theorem. These results themselves may be of interest.

First of all, Theorem 4.1 directly leads to the following proposition:

Proposition 4.3 (Individual Sub-domain Generalization Bound). *For $k \in \{1, \dots, K\}$, where K represents the total number of distinct sub-domains, for sub-domain X_S^k with distribution P_S^k and X_T^k with distribution P_T^k , it holds any $f \in \mathbb{H}$ that*

$$\gamma_T^k(f) \leq \gamma_S^k(f) + W_1(P_S^k, P_T^k) + (\gamma^k)^*, \quad (5)$$

where $\gamma_S^k(f)/\gamma_T^k(f)$ is the performance of the hypothesis on the sub-domain X_S^k/X_T^k , $(\gamma^k)^* = \min_{f \in \mathbb{H}} \gamma_S^k(f) + \gamma_T^k(f)$, \mathbb{H} is a hypothesis class included in the set of λ_H -Lipschitz functions, the true functions g_T and g_S are both λ -Lipschitz functions, and $\lambda + \lambda_H \leq 1$.

The second result below shows that the classification error of any hypothesis f on a domain can be decomposed into a weighted sum of the classification errors of f on its sub-domains (proofs deferred to the Appendix [A.5](#)).

Lemma 4.4 (Decomposition of the Classification Error). *For any hypothesis $f \in \mathbb{H}$,*

$$\gamma_S(f) = \sum_{k=1}^K w_S^k \gamma_S^k(f), \gamma_T(f) = \sum_{k=1}^K w_T^k \gamma_T^k(f). \quad (6)$$

With above results, we present a generalization bound with sub-domain information (proofs deferred to the Appendix [A.6](#)).

Theorem 4.5 (Sub-domain-based Generalization Bound).

$$\gamma_T(f) \leq \sum_{k=1}^K w_T^k \gamma_S^k(f) + \sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) + \sum_{k=1}^K w_T^k (\gamma^k)^\star. \quad (7)$$

In particular, in a balanced domain adaptation setting where for all k , $w_S^k = w_T^k$, we have that

$$\gamma_T(f) \leq \gamma_S(f) + \sum_{k=1}^K w_S^k W_1(P_S^k, P_T^k) + \sum_{k=1}^K w_S^k (\gamma^k)^\star. \quad (8)$$

Remark 4.6. Note that the format of the RHS of equation [8](#) is reminiscent of the objectives used by the majority of the sub-domain-based methods.

We next show that, under reasonable assumptions, the weighted sum of distances between corresponding sub-domains of the source and target domains is at most as large as the distance between the marginal distribution of the source domain and that of the target domain.

Theorem 4.7 (Benefits of Sub-domain Alignment). *Under the following assumptions:*

A1. For all k , P_S^k / P_T^k are Gaussian distributions with mean m_S^k / m_T^k and covariance Σ_S^k / Σ_T^k .

A2. Distance between the paired source-target sub-domain is less or equal to distance between the non-paired source-target sub-domain, i.e., $W_1(P_S^k, P_T^k) \leq W_1(P_S^k, P_T^{k'})$ for $k \neq k'$.

A3. There exists a small constant $\epsilon > 0$, such that $\max_{1 \leq k \leq K} (\text{tr}(\Sigma_S^k)) \leq \epsilon$ and $\max_{1 \leq k \leq K} (\text{tr}(\Sigma_T^{k'})) \leq \epsilon$. Then the following inequality holds:

$$\sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) \leq W_1(P_S, P_T) + \delta_c, \quad (9)$$

where δ_c is $4\sqrt{\epsilon}$. In particular, when $w_S^k = w_T^k$ for all k ,

$$\sum_{k=1}^K w_S^k W_1(P_S^k, P_T^k) \leq W_1(P_S, P_T) + \delta_c. \quad (10)$$

Proof. Note that $\mathbf{w}_S \doteq [w_S^1, \dots, w_S^K]$ and $\mathbf{w}_T \doteq [w_T^1, \dots, w_T^K]$ belong to Δ^K (the $K-1$ probability simplex). $\Pi(w_S, w_T)$ represents the simplex $\Delta^{K \times K}$ with marginals \mathbf{w}_S and \mathbf{w}_T . With $w \in \Pi(\mathbf{w}_S, \mathbf{w}_T)$, we can write out w_T^k as $\sum_{k'=1}^K w_{k,k'}$, then based on assumption A.2, we have:

$$\begin{aligned} \sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) &= \sum_{k=1}^K \sum_{k'=1}^K w_{k,k'} W_1(P_S^k, P_T^k) \\ &\leq \sum_{k=1}^K \sum_{k'=1}^K w_{k,k'} W_1(P_S^k, P_T^{k'}). \end{aligned}$$

Thus we have $(MW_1(P_S, P_T))$ defined in Appendix [A.7](#),

$$\begin{aligned} \sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) &\leq \min_{w \in \Pi(\mathbf{w}_S, \mathbf{w}_T)} \sum_{k=1}^K \sum_{k'=1}^K w_{k,k'} W_1(P_S^k, P_T^{k'}) \\ &= MW_1(P_S, P_T). \end{aligned} \quad (11)$$

Also we prove in Theorem [A.10](#) that:

$$MW_1(P_S, P_T) \leq W_1(P_S, P_T) + 4\sqrt{\epsilon}.$$

Then we conclude our proof and show that:

$$\sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) \leq MW_1(P_S, P_T) \leq W_1(P_S, P_T) + 4\sqrt{\epsilon} = W_1(P_S, P_T) + \delta_c. \quad (12)$$

□

Remark 4.8. In Appendix B, we provide empirical evidence to verify that these assumptions are satisfied on real-world datasets. We note that the assumption of a Gaussian distribution for X^k is not unreasonable since it is often the result of a complex transformation, Φ , and the Central Limit Theorem indicates that the outcome of such a transformation is approximately normally distributed under regularity assumptions (please see Appendix B.1 for empirical evidence).

Remark 4.9. δ_c is a constant dependent only on the variance of the features but not the covariance between features in different dimensions. Moreover, the inequality holds empirically without δ_c as demonstrated in Figure 3, as well as Figure 7 and Figure 8 in Appendix F.2.

4.3 Challenges of Imbalanced UDA

Theorem 4.7 shows that the objective function of sub-domain methods is at least as strong as the objective function of domain alignment methods, explaining its improved performance. However, if the marginal weights of the sub-domain shifts, i.e., $w_S^k \neq w_T^k$, the inequality in equation 10 is not likely to hold and the framework can collapse. One such example is the scenario of shifted label distributions where w_T^k and w_S^k (class weights for target and source domains) can be vastly different.

To overcome this, we propose to minimize an objective with the simple intuition that *important sub-domains in the target domain need more attention*. With this motivation, we propose the following objective function for UDA with shifted label distribution:

$$\mathcal{L}(f) = \sum_{k=1}^K w_T^k \gamma_S^k(f). \quad (13)$$

In particular, \mathcal{L} reweighs the losses of sub-domains so that the sub-domain with more weight in the target domain can be emphasized more. We next prove that through the proposed approach (proofs deferred to the Appendix A.12), we can again obtain a sub-domain-based generalization bound that is at least as strong as the full domain generalization bound without the sub-domain information.

Theorem 4.10. Let $\mathbb{H} \doteq \{f|f : \mathcal{X} \rightarrow [0, 1]\}$ denote a hypothesis space. Under the assumptions in Theorem 4.7, for any $f \in \mathbb{H}$ such that:

$$\sum_{k=1}^K w_T^k \gamma_S^k(f) \leq \sum_{k=1}^K w_S^k \gamma_S^k(f), \quad (14)$$

then we have $\sum_{k=1}^K w_T^k (\gamma^k)^* \leq \gamma^*$. Further, let

$$\epsilon_c(f) \doteq \sum_{k=1}^K w_T^k \gamma_S^k(f) + \sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) + \sum_{k=1}^K w_T^k (\gamma^k)^*$$

denote the sub-domain-based generalization bound and let

$$\epsilon_g(f) \doteq \gamma_S(f) + W_1(P_S, P_T) + \gamma^*$$

denote the generalization bound without any sub-domain information, we have,

$$\epsilon_c(f) \leq \epsilon_g(f) + \delta_c.$$

Remark 4.11. In Section 6.1 and Appendix F.2, we provide extensive empirical evidence to establish that equation 14 can easily hold, as the left hand side is the optimization objective. Moreover, in these sections, we offer empirical evidence to further verify the value of this theoretical result by showing that our proposed bound is empirically much stronger than the existing one.

Inspired by our analysis, we propose a framework, *Domain Adaptation with Rebalanced Sub-domain Alignment* (DARSA), for imbalanced UDA, a special case of the sub-domain weight shifting scenario where the class weights of the target domain shifts from that of the source domain.

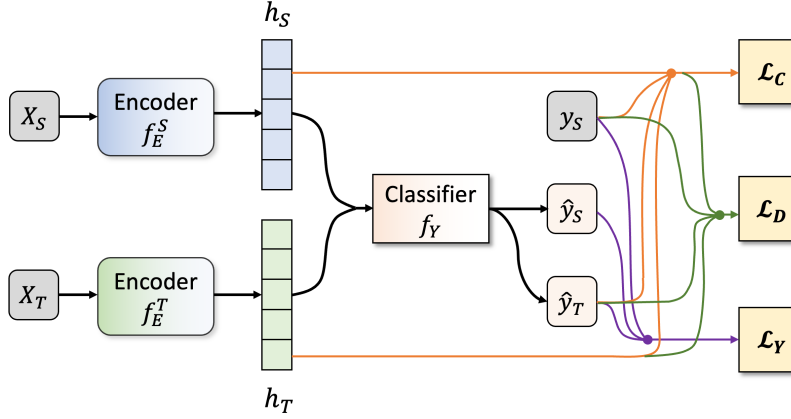


Figure 2: The DARSA framework. Orange lines representing the clustering loss \mathcal{L}_C , green lines indicating domain discrepancy \mathcal{L}_D , and purple lines indicating source classification loss \mathcal{L}_Y .

5 Methods

In DARSA, we divide the source domains into sub-domains based on class labels, and divide target domains into sub-domains using predicted class labels (serving as pseudo labels, which have shown success in previous research (Deng et al., 2019; Lee et al., 2013)) for unlabeled target domains. Motivated by Theorem 4.10, the framework of DARSA, shown in Figure 2, is composed of a source encoder f_E^S parameterized by θ_E^S , a target encoder f_E^T parameterized by θ_E^T , and a classifier f_Y parameterized by θ_Y . The pseudo-code for DARSA can be found in Appendix D.

The objective function of DARSA is defined as follows:

$$\min_{\theta_Y, \theta_E^S, \theta_E^T} \lambda_Y \mathcal{L}_Y + \lambda_D \mathcal{L}_D + \mathcal{L}_C, \quad (15)$$

where \mathcal{L}_Y , \mathcal{L}_D , \mathcal{L}_C are losses described below with relative weights given by λ_Y and λ_D .

Weighted source domain classification error \mathcal{L}_Y . The weighted source domain classification error in Theorem 4.10 can be further expressed as:

$$\begin{aligned} \sum_{k=1}^K w_T^k \gamma_S^k(f) &= \sum_{k=1}^K w_T^k \int P_S(x|c=k) |f(x) - g_S(x)| dx \\ &= \sum_{k=1}^K w_T^k \int \frac{P_S(c=k|x) P_S(x)}{P_S(c=k)} |f(x) - g_S(x)| dx = \sum_{k=1}^K \frac{w_T^k}{w_S^k} \mathbb{E}_{x \sim D_S} w_S^k(x) |f(x) - g_S(x)|, \end{aligned} \quad (16)$$

where variable c represents class, $w_T^k = P_T(c=k)$, $w_S^k = P_S(c=k)$, $w_S^k(x) = P_S(c=k|x)$. We set $P_S(c=k|x) = 1$ only when data point x is in class k , otherwise $P_S(c=k|x) = 0$. w_S^k can be set to the marginal source label distribution, and w_T^k can be estimated from the target predictions. From equation 16, $\mathcal{L}_Y(\theta_Y, \theta_E^S)$ is defined as:

$$\mathcal{L}_Y(\theta_Y, \theta_E^S) = \frac{1}{N_S} \sum_{x^i \in \mathcal{X}_S} \mathbb{1}_{y^i=k} \frac{w_T^k}{w_S^k} \ell(\hat{y}^i, y^i),$$

where $\hat{y}^i = f_Y(f_E^S(x^i))$ is the predicted label and ℓ can be any non-negative loss function (e.g., cross-entropy loss for classification tasks).

Weighted source-target subdomain discrepancy \mathcal{L}_D . The weighted source-target domain discrepancy in Theorem 4.10 can be further expressed as:

$$\mathcal{L}_D(\theta_E^S, \theta_E^T, \theta_Y) = \sum_{k=1}^K w_T^k W_1(P_S^k, P_T^k) = \sum_{k=1}^K w_T^k W_1(f_E^S(x_S^k), f_E^T(x_T^k)), \quad (17)$$

where x_S^k are source samples with labels $y_S = k$, and x_T^k are target samples with predicted labels $\hat{y}_T = k$. We leverage the Sinkhorn algorithm (Cuturi, 2013) to approximate the Wasserstein metric.

Clustering loss \mathcal{L}_C . The clustering loss $\mathcal{L}_C = \lambda_c \mathcal{L}_{intra} + \lambda_a \mathcal{L}_{inter}$ is comprised of two components: the intra-clustering loss, \mathcal{L}_{intra} , and the inter-clustering loss, \mathcal{L}_{inter} . The role of \mathcal{L}_{intra} is to satisfy the assumption A.3 in Theorem 4.7. It encourages embeddings of the same label to cluster tightly together, while also pushing embeddings of different labels to separate by at least a user-specified distance, m (Luo et al., 2018). The inter-clustering loss \mathcal{L}_{inter} further enhances sub-domain alignment by aligning the centroids of source sub-domains with those of their corresponding target sub-domains in the representation space. We define \mathcal{L}_{intra} and \mathcal{L}_{inter} as follows:

$$\mathcal{L}_{intra}(\theta_E^S, \theta_E^T, \theta_Y) = \mathcal{L}_{intra}(f_E^S(\mathcal{X}_S)) + \mathcal{L}_{intra}(f_E^T(\mathcal{X}_T)), \quad (18)$$

$$\mathcal{L}_{intra}(f_E^S(\mathcal{X})) = \frac{1}{N^2} \sum_{i,j=1}^N \left[\delta_{ij} D_{ij} + (1 - \delta_{ij}) \max(0, m - D_{ij}) \right];$$

$$\mathcal{L}_{inter}(\theta_E^S, \theta_E^T, \theta_Y) = \frac{1}{K} \sum_{k=1}^K \|\mathcal{C}(f_E^S(x_T^k)) - \mathcal{C}(f_E^T(x_T^k))\|^2, \quad (19)$$

where N represents the number of samples in the domain \mathcal{X} and $\mathcal{C}(\cdot)$ calculates the centroids of the sub-domains, $\delta_{ij} = 1$ only if x_i and x_j have the same label; otherwise, $\delta_{ij} = 0$. We use the ground truth label or the predicted label if x is in source domain or target domain, respectively. m is a pre-defined distance controlling how separated each sub-domain should be. $D_{ij} = \|f_E(x_i) - f_E(x_j)\|^2$ represents distance between x_i and x_j .

6 Experiments

In this section, we verify our theoretical results and assess DARSA’s efficacy through real-world experiments. We begin by empirically confirming the superiority of the sub-domain-based generalization bound (Theorem 4.10) in Section 6.1. Then, we verify that the assumptions for Theorem 4.10 are empirically satisfied on real-world datasets (details in Appendix B). Next, we demonstrate the vital role of subdomain weight re-balancing in Section 6.2 and show DARSA’s robustness to minor weight estimation discrepancies. Lastly, given that our theoretical analysis guarantees that DARSA should have competitive performance in scenarios where the number of classes is not overwhelming, we evaluate DARSA on real-world datasets with this property. Comparing with other state-of-the-art UDA baselines, we verify the correctness of our analysis as well as an advantage of DARSA that its strong performance can be guaranteed on particular real-world applications such as those in medical and operations research. We base the following confirmatory experiments on two sets of datasets.

Experiments on the Digits Datasets. In our Digits datasets experiments, we evaluate our performance across four datasets: MNIST (M) (LeCun et al., 1998), MNIST-M (MM) (Ganin et al., 2016), USPS (U), and SVHN (S), all modified to induce label distribution shifts. Here, the parameter α denotes the class imbalance rate, representing a ratio such as $1:\alpha$ and $\alpha:1$ for the odd:even distribution in the source and target datasets, respectively. Weak and strong imbalance correspond to $\alpha = 3$ and $\alpha = 8$. For comprehensive details, refer to Appendix F.

Experiments on the TST Dataset. We use the Tail Suspension Test (TST) dataset (Gallagher et al., 2017) of local field potentials (LFPs) from 26 mice with two genetic backgrounds: Clock- $\Delta 19$ (a bipolar disorder model) and wildtype. This dataset is publicly available (Carlson et al., 2023). Our study involves two domain adaptation tasks, predicting the current condition - home cage (HC), open field (OF), or tail-suspension (TS) - from one genotype to the other. We subsample datasets to induce label distribution shifts with imbalance rate = 2. For comprehensive details, refer to Appendix G.

6.1 Empirical Analysis of our Proposed Generalization Bound

We first verify the pivotal result in Theorem 4.10 that the sub-domain based generalization bound is at least as tight as the non-sub-domain bound. We empirically evaluate the proposed bound on the Digits datasets under weak imbalance. As shown in Figure 3, our empirical results demonstrate that the sub-domain-based generalization bound in Theorem 4.5 is empirically much stronger than the non-sub-domain-based bound in Theorem 4.1, corroborating our insights for the effectiveness of sub-domain based methods. Additional

¹The code to replicate all experiments is available at: <https://anonymous.4open.science/r/DARSA/>

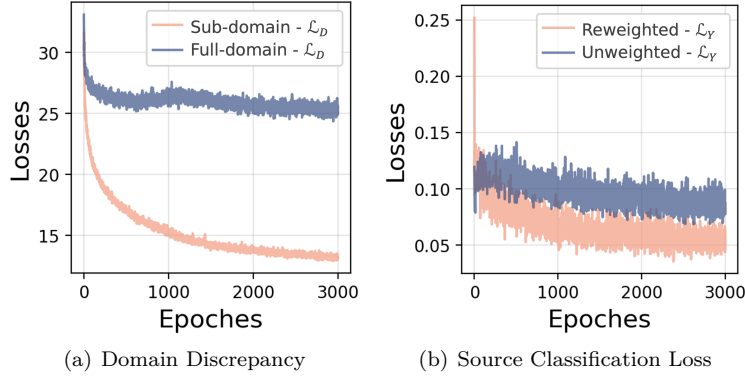


Figure 3: For MNIST to MNIST-M task under weak imbalance. (a) Compare the domain discrepancy term (\mathcal{L}_D) in our proposed bound to that in Theorem 4.1. (b) Compare the source classification term (\mathcal{L}_Y) in our proposed bound to that in Theorem 4.1

experiments on the other UDA tasks in the Digits datasets under weak and strong imbalance also support this claim, and full results are in Appendix F.2.

6.2 Importance of Re-weighting

Here, we experiment on the Digits datasets under weak imbalance to demonstrate the importance of (i) weights re-weighting and (ii) the accuracy of target sub-domain weights estimation. We compare DARSA with one variation of DARSA which employs uniform weights for all sub-domains and another variation which swaps sub-domain weights estimation of source with target. We also include two other baselines where the weights of the target domain are chosen to be deviating from the truth. Specifically, we compare DARSA with the following configurations:

- DARSA: Full algorithm where weights are inferred.
- DARSA Oracle: Utilizing true values of w_T^k .
- DARSA Small Divergence: Setting w_T^k to be 20% divergent from true values.
- DARSA Large Divergence: Setting w_T^k to be 50% divergent from true values.
- DARSA Flip: Swapping w_T^k with w_S^k , effectively flipping importance weighting.
- DARSA Uniform: Assigning uniform weights for all sub-domains.

The results of these experiments are in Table 1. We verify the importance of subdomain weights re-balancing by showing that the performance of DARSA degrades significantly without the weights re-balancing or wrong sub-domain weights, further corroborating the value of our insights. Additionally, while the oracle case provides the best performance, inferring the weights in the DARSA algorithm provides nearly the same quality of predictions. In addition, we found our method, DARSA is robust to minor divergence in weights estimation and varying imbalance rates.

Table 1: Evaluation of the importance of re-weighting on Digits datasets under weak imbalance. Performance is measured by prediction accuracy (%) on the target domain.

	M → MM	MM → M	U → M	S → M
DARSA Oracle	96.2	98.4	92.7	92.6
DARSA Uniform	67.9	96.6	75.9	71.7
DARSA Small Divergence	95.6	98.3	91.4	92.4
DARSA Large Divergence	85.0	98.2	86.1	85.2
DARSA Flip	55.7	65.7	57.4	65.7
DARSA	96.0	98.8	92.6	90.1

Table 2: Summary of UDA results on the Digits datasets with shifted label distribution, measured in terms of prediction accuracy (%) on the target domain.

	M \rightarrow MM $\alpha = 3$	MM \rightarrow M $\alpha = 3$	U \rightarrow M $\alpha = 3$	S \rightarrow M $\alpha = 3$	M \rightarrow MM $\alpha = 8$	MM \rightarrow M $\alpha = 8$	U \rightarrow M $\alpha = 8$	S \rightarrow M $\alpha = 8$
DANN (Ganin et al., 2016)	63.1	93.0	59.8	64.9	61.1	90.2	49.1	57.3
DSN (Bousmalis et al., 2016)	62.3	98.4	59.9	15.2	57.5	95.3	30.3	17.8
ADDA (Tzeng et al., 2017)	88.2	90.7	44.8	42.4	47.9	89.4	45.7	45.3
pixelDA (Bousmalis et al., 2017)	95.0	96.0	72.0	68.0	81.0	95.6	29.2	60.4
CDAN (Long et al., 2018)	58.7	96.0	42.0	38.3	37.1	90.6	34.8	32.5
WDGRL (Shen et al., 2018)	60.4	93.6	63.9	64.3	22.3	91.4	46.7	52.2
MCD (Saito et al., 2018)	58.1	98.2	74.6	75.5	37.4	97.5	76.1	66.7
CAT (Deng et al., 2019)	54.1	95.4	81.0	65.8	48.9	93.8	61.3	62.2
MDD (Zhang et al., 2019)	48.7	97.7	82.3	62.4	47.6	93.6	83.2	64.5
DRANet (Lee et al., 2021)	95.2	97.8	86.5	40.2	63.3	96.1	54.2	31.3
Source Only	47.9	91.5	40.8	53.7	39.6	88.4	27.8	47.2
DARSA	96.0	98.8	92.6	90.1	78.8	97.3	87.9	83.5

Table 3: Summary of UDA results on the TST datasets with shifted label distribution, measured in terms of prediction accuracy (%) on the target domain.

	DANN	WDGRL	DSN	ADDA	CAT	CDAN	Source only	DARSA
Clock- Δ 19 to Wildtype	79.9	79.6	79.4	75.1	77.3	75.0	73.8	86.6
Wildtype to Clock- Δ 19	81.5	79.5	80.9	72.6	78.6	73.6	70.4	84.8

6.3 DARSA on Real-world Datasets

We now compare DARSA with many competing algorithms on these two datasets. Full details on the experiments, the rationale for competing algorithms choices, and their settings are in Appendix F and Appendix G for the Digits and TST datasets, respectively.

Digits. Results shown in Table 2 demonstrates DARSA’s competitiveness in handling label shifting. Additionally, DARSA performs well with varying imbalance rates (Appendix Table 6) and competes favorably in scenarios without label distribution shifts (Appendix Table 7).

TST. As demonstrated in Table 3, DARSA achieves competitive performance on this biologically relevant task. For comprehensive experimental details, refer to Appendix G.

Ablation. To assess the impact of each component within our objective function (Section 5), we conduct an ablation study under weak imbalance. As demonstrated in Table 4, the ablation analysis confirms that each component in our objective function contributes to the overall performance. Therefore, we recommend the use of all components for optimal results. In addition, we have included feature space visualizations in Appendix C and Appendix Figure 9 which demonstrate that the learned representation of DARSA has improved separation when using all the components, supporting the effectiveness of the proposed objective function.

7 Limitations

However, despite demonstrating empirical improvements over existing methods, DARSA may be less effective in applications with a large number of class labels (e.g., 1,000). In such cases, aligning the corresponding sub-domains (classes) becomes significantly more challenging because the alignment relies on pseudo-labels, which tend to be less reliable when the label space is large. Moreover, effective and stable alignment of sub-domains requires an increased batch size for SGD. This can cause difficulty when the computational resource is limited. On the other hand, DARSA may not achieve the best performance when the domain shift is extreme (e.g., when the probability of observing certain sub-domains is extremely low or even zero). In these scenarios, accurately estimating the sub-domain Wasserstein distance and marginal weights is difficult due to high variance from the scarcity of samples. However, we note that these limitations are not unique to DARSA: most UDA methods relying on pseudo-labels face similar challenges under large label spaces or extreme domain shifts. In terms of the theoretical results, our theorems assume that sub-domain embeddings

Table 4: Ablation study results. Each row represents a configuration with different λ values. The last column reports the prediction accuracy (%) for each configuration.

Experiment	λ_Y	λ_D	λ_a	λ_c	Accuracy
M \rightarrow MM	0.4	0.35	0.9	1	96.0
M \rightarrow MM	0	0.35	0.9	1	61.3
M \rightarrow MM	0.4	0	0.9	1	72.5
M \rightarrow MM	0.4	0.35	0	1	61.9
M \rightarrow MM	0.4	0.35	0.9	0	33.5
MM \rightarrow M	1	0.5	1	1	98.8
MM \rightarrow M	0	0.5	1	1	96.7
MM \rightarrow M	1	0	1	1	98.4
MM \rightarrow M	1	0.5	1	0	15
MM \rightarrow M	1	0.5	0	1	98.2
U \rightarrow M	1	0.5	1	1	92.6
U \rightarrow M	0	0.5	1	1	65.9
U \rightarrow M	1	0	1	1	85.8
U \rightarrow M	1	0.5	0	1	76.2
U \rightarrow M	1	0.5	1	0	58.4
S \rightarrow M	0.95	0.11	0.11	0.3	90.1
S \rightarrow M	0	0.11	0.11	0.3	77.9
S \rightarrow M	0.95	0	0.11	0.3	86.1
S \rightarrow M	0.95	0.11	0.11	0	64.3
S \rightarrow M	0.95	0.11	0	0.3	84.9

are well-separated and that each sub-domain in the source corresponds to the closest sub-domain in the target. This requirement is equivalent to requiring that both the source and target domains themselves must be well-separated (i.e., classified accurately with high probability). Hence, when either domain is inherently difficult to classify, our theoretical guarantees may not hold, and the benefit of DARSA is less significant.

8 Conclusion

Sub-domain-based algorithms have demonstrated considerable empirical success across various applications in domain adaptation. However, a comprehensive theoretical understanding of their advantages had been elusive. This work addresses this gap and presents a substantial contribution by providing a rigorous theoretical perspective on the benefits of sub-domain-based methods, thereby potentially enhancing their overall impact in the field. Moreover, our analysis leads to an algorithm DARSA with improved robustness to the shift of sub-domain weights and label distributions. Additionally, our framework can be extended to data integration and causal effect estimation by reframing these tasks as distribution alignment problems.

9 Reproducibility Statement

Rigorous definitions and complete proofs of our theoretical analysis are included in the Appendix [A](#), with empirical evidence to verify assumptions in Appendix [B](#). The code to replicate all experiments is available at: <https://anonymous.4open.science/r/DARSA/>. Full details on the experiments, competing algorithms, and their settings are in Appendix [F](#) and Appendix [G](#) for the Digits and TST dataset, respectively. The MNIST, BSDS500, USPS, and SVHN datasets are publicly available with an open-access license. The Tail Suspension Test (TST) dataset ([Gallagher et al., 2017](#)) is available to download at <https://research.repository.duke.edu/concern/datasets/zc77sr31x?locale=en> for free under a Creative Commons BY-NC Attribution-NonCommercial 4.0 International license. The experiments are conducted on a computer cluster equipped with a NVIDIA GeForce RTX 2080 Ti that has a memory capacity of 11019MiB.

References

- Ali Abedi, QM Wu, Ning Zhang, and Farhad Pourpanah. Euda: An efficient unsupervised domain adaptation via self-supervised vision transformer. *arXiv preprint arXiv:2407.21311*, 2024.
- Isabela Albuquerque, João Monteiro, Tiago H Falk, and Ioannis Mitliagkas. Adversarial target-invariant representation learning for domain generalization. *arXiv preprint arXiv:1911.00804*, 8, 2019.
- Jason Altschuler, Jonathan Niles-Weed, and Philippe Rigollet. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. *Advances in neural information processing systems*, 30, 2017.
- Pablo Arbelaez, Michael Maire, Charles Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*, 33(5):898–916, 2010.
- Eytan Bakshy, Max Balandat, and Kostya Kashin. Open-sourcing ax and botorch: New ai tools for adaptive experimentation. URL <https://ai.facebook.com/blog/open-sourcing-ax-and-botorch-new-ai-tools-for-adaptive-experimentation>.
- Shai Ben-David, John Blitzer, Koby Crammer, and Fernando Pereira. Analysis of representations for domain adaptation. *Advances in neural information processing systems*, 19, 2006.
- Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.
- Gilles Blanchard, Aniket Anand Deshmukh, Ürün Dogan, Gyemin Lee, and Clayton Scott. Domain generalization by marginal transfer learning. *The Journal of Machine Learning Research*, 22(1):46–100, 2021.
- Karsten M Borgwardt, Arthur Gretton, Malte J Rasch, Hans-Peter Kriegel, Bernhard Schölkopf, and Alex J Smola. Integrating structured biological data by kernel maximum mean discrepancy. *Bioinformatics*, 22(14):e49–e57, 2006.
- Konstantinos Bousmalis, George Trigeorgis, Nathan Silberman, Dilip Krishnan, and Dumitru Erhan. Domain separation networks. *Advances in neural information processing systems*, 29, 2016.
- Konstantinos Bousmalis, Nathan Silberman, David Dohan, Dumitru Erhan, and Dilip Krishnan. Unsupervised pixel-level domain adaptation with generative adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3722–3731, 2017.
- David Carlson, Sunil Kumar, and Kafui Dzirasa. Multi-region local field potential recordings during a tail-suspension test. *Duke Research Data Repository*, 2023.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- Julie Delon and Agnes Desolneux. A wasserstein-type distance in the space of gaussian mixture models. *SIAM Journal on Imaging Sciences*, 13(2):936–970, 2020.
- Zhijie Deng, Yucen Luo, and Jun Zhu. Cluster alignment with a teacher for unsupervised domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 9944–9953, 2019.
- Sofien Dhoubib, Ievgen Redko, and Carole Lartzien. Margin-aware adversarial domain adaptation with optimal transport. In *International conference on machine learning*, pp. 2514–2524. PMLR, 2020.
- Abolfazl Farahani, Sahar Voghoei, Khaled Rasheed, and Hamid R Arabnia. A brief review of domain adaptation. *Advances in data science and information engineering*, pp. 877–894, 2021.
- Nicolas Fournier and Arnaud Guillin. On the rate of convergence in wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3):707–738, 2015.
- Neil Gallagher, Kyle R Ulrich, Austin Talbot, Kafui Dzirasa, Lawrence Carin, and David E Carlson. Cross-spectral factor analysis. *Advances in neural information processing systems*, 30, 2017.

- Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *International conference on machine learning*, pp. 1180–1189. PMLR, 2015.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario Marchand, and Victor Lempitsky. Domain-adversarial training of neural networks. *The journal of machine learning research*, 17(1):2096–2030, 2016.
- Boyan Gao, Yongxin Yang, Henry Gouk, and Timothy M Hospedales. Deep clustering for domain adaptation. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4247–4251. IEEE, 2020.
- Pascal Germain, Amaury Habrard, François Laviolette, and Emilie Morvant. A new pac-bayesian perspective on domain adaptation. In *International conference on machine learning*, pp. 859–868. PMLR, 2016.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. *Advances in neural information processing systems*, 30, 2017.
- Mahta HassanPour Zonoozi and Vahid Seydi. A survey on adversarial domain adaptation. *Neural Processing Letters*, pp. 1–41, 2022.
- Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. Cycada: Cycle-consistent adversarial domain adaptation. In *International conference on machine learning*, pp. 1989–1998. Pmlr, 2018.
- Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1125–1134, 2017.
- Xiang Jiang, Qicheng Lao, Stan Matwin, and Mohammad Havaei. Implicit class-conditioned domain alignment for unsupervised domain adaptation. In *International Conference on Machine Learning*, pp. 4816–4827. PMLR, 2020.
- Ying Jin, Ximei Wang, Mingsheng Long, and Jianmin Wang. Minimum class confusion for versatile domain adaptation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, pp. 464–480. Springer, 2020.
- Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3, pp. 896, 2013.
- Seunghun Lee, Sunghyun Cho, and Sunghoon Im. Dranet: Disentangling representation and adaptation networks for unsupervised cross-domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 15252–15261, 2021.
- Benjamin Letham, Brian Karrer, Guilherme Ottoni, and Eytan Bakshy. Constrained bayesian optimization with noisy experiments. 2019.
- Guangrui Li, Guoliang Kang, Yi Zhu, Yunchao Wei, and Yi Yang. Domain consensus clustering for universal domain adaptation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9757–9766, 2021a.
- Jichang Li, Guanbin Li, Yemin Shi, and Yizhou Yu. Cross-domain adaptive clustering for semi-supervised domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2505–2514, 2021b.
- Xiaomeng Li, Lequan Yu, Hao Chen, Chi-Wing Fu, Lei Xing, and Pheng-Ann Heng. Transformation-consistent self-ensembling model for semisupervised medical image segmentation. *IEEE Transactions on Neural Networks and Learning Systems*, 32(2):523–534, 2020.

- Yitong Li, David E Carlson, et al. Extracting relationships by multi-domain matching. *Advances in Neural Information Processing Systems*, 31, 2018.
- Kevin J Liang, Chunyuan Li, Guoyin Wang, and Lawrence Carin. Generative adversarial network training is a continual learning problem. *arXiv preprint arXiv:1811.11083*, 2018.
- Ming-Yu Liu, Thomas Breuel, and Jan Kautz. Coupled generative adversarial networks. *arXiv preprint arXiv:1606.07536*, 2016.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael Jordan. Learning transferable features with deep adaptation networks. In *International conference on machine learning*, pp. 97–105. PMLR, 2015.
- Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. *Advances in neural information processing systems*, 31, 2018.
- Le Lu, Yefeng Zheng, Gustavo Carneiro, and Lin Yang. Deep learning and convolutional neural networks for medical image computing. *Advances in computer vision and pattern recognition*, 10:978–3, 2017.
- Yulong Lu and Jianfeng Lu. A universal approximation theorem of deep neural networks for expressing probability distributions. *Advances in neural information processing systems*, 33:3094–3105, 2020.
- Yucen Luo, Jun Zhu, Mengxi Li, Yong Ren, and Bo Zhang. Smooth neighbors on teacher graphs for semi-supervised learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8896–8905, 2018.
- Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Multiple source adaptation and the rényi divergence. *arXiv preprint arXiv:1205.2628*, 2012.
- Khai Nguyen, Tongzheng Ren, Huy Nguyen, Litu Rout, Tan Nguyen, and Nhat Ho. Hierarchical sliced wasserstein distance. *arXiv preprint arXiv:2209.13570*, 2022.
- Papers with Code. Domain Adaptation. <https://paperswithcode.com/task/domain-adaptation>, 2023. Online; accessed 24-Sept-2023.
- Ofir Pele and Michael Werman. Fast and robust earth mover’s distances. In *2009 IEEE 12th international conference on computer vision*, pp. 460–467. IEEE, 2009.
- Kuan-Chuan Peng, Ziyang Wu, and Jan Ernst. Zero-shot deep domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 764–781, 2018.
- Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- Pedro O Pinheiro. Unsupervised domain adaptation with similarity learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8004–8013, 2018.
- Ievgen Redko, Amaury Habrard, and Marc Sebban. Theoretical analysis of domain adaptation with optimal transport. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 737–753. Springer, 2017.
- Tim Sainburg, Leland McInnes, and Timothy Q Gentner. Parametric umap embeddings for representation and semisupervised learning. *Neural Computation*, 33(11):2881–2907, 2021.
- Kuniaki Saito, Kohei Watanabe, Yoshitaka Ushiku, and Tatsuya Harada. Maximum classifier discrepancy for unsupervised domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3723–3732, 2018.

- Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3752–3761, 2018.
- Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.
- Yuan Shi and Fei Sha. Information-theoretical learning of discriminative clusters for unsupervised domain adaptation. *arXiv preprint arXiv:1206.6438*, 2012.
- Jake Snell, Kevin Swersky, and Richard Zemel. Prototypical networks for few-shot learning. *Advances in neural information processing systems*, 30, 2017.
- Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 30, 2016.
- Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014.
- Eric Tzeng, Judy Hoffman, Trevor Darrell, and Kate Saenko. Simultaneous deep transfer across domains and tasks. In *Proceedings of the IEEE international conference on computer vision*, pp. 4068–4076, 2015.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7167–7176, 2017.
- Eric Tzeng, Coline Devin, Judy Hoffman, Chelsea Finn, Pieter Abbeel, Sergey Levine, Kate Saenko, and Trevor Darrell. Adapting deep visuomotor representations with weak pairwise constraints. In *Algorithmic Foundations of Robotics XII*, pp. 688–703. Springer, 2020.
- Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Haohan Wang, Zeyi Huang, Hanlin Zhang, Yong Jae Lee, and Eric P Xing. Toward learning human-aligned cross-domain robust models by countering misaligned features. In *Uncertainty in Artificial Intelligence*, pp. 2075–2084. PMLR, 2022.
- Mei Wang and Weihong Deng. Deep visual domain adaptation: A survey. *Neurocomputing*, 312:135–153, 2018.
- Thomas Westfechtel, Hao-Wei Yeh, Qier Meng, Yusuke Mukuta, and Tatsuya Harada. Backprop induced feature weighting for adversarial domain adaptation with iterative label distribution alignment. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 392–401, 2023.
- YiFan Zhang, Xue Wang, Jian Liang, Zhang Zhang, Liang Wang, Rong Jin, and Tieniu Tan. Free lunch for domain adversarial training: Environment label smoothing. *arXiv preprint arXiv:2302.00194*, 2023.
- Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *International conference on machine learning*, pp. 7404–7413. PMLR, 2019.
- Han Zhao, Shanghang Zhang, Guanhang Wu, José MF Moura, Joao P Costeira, and Geoffrey J Gordon. Adversarial multiple source domain adaptation. *Advances in neural information processing systems*, 31, 2018.