
Algorithm-Dependent Bounds for Representation Learning of Multi-Source Domain Adaptation

Qi Chen
Université Laval

Mario Marchand
Université Laval

Abstract

We use information-theoretic tools to derive a novel analysis of Multi-source Domain Adaptation (MDA) from the representation learning perspective. Concretely, we study joint distribution alignment for supervised MDA with few target labels and unsupervised MDA with pseudo labels, where the latter is relatively hard and less commonly studied. We further provide algorithm-dependent generalization bounds for these two settings, where the generalization is characterized by the mutual information between the parameters and the data. Then we propose a novel deep MDA algorithm, implicitly addressing the target shift through joint alignment. Finally, the mutual information bounds are extended to this algorithm providing a non-vacuous gradient-norm estimation. The proposed algorithm has comparable performance to the state-of-the-art on target-shifted MDA benchmark with improved memory efficiency.

1 INTRODUCTION

The usual machine learning theories assume the test data follows the same distribution as the train data, which is often violated in real-world applications. Such a distribution (domain) shift degrades the algorithm’s performance. So, various methods have been proposed to address this problem through Domain Adaptation (DA) or transfer learning (Huang et al., 2007; Ben-David et al., 2010a; Pan & Yang, 2009). Domain adaptation aims to learn a well-generalized predictor for the target domain \mathcal{T} with data sampled from the source domain \mathcal{S} , where $\mathcal{S} \neq \mathcal{T}$. As deep learning becomes increasingly popular with its superior performance in many complex tasks such as computer vision (Csurka,

2017) and natural language processing (Blitzer, 2008), a series of works on deep domain adaption (Ganin et al., 2016; Tzeng et al., 2017) have gained tremendous success in practice.

The works mentioned above are focused on single-source domain adaptation. However, leveraging knowledge from multiple sources ($\mathcal{S}_{1:N}$) is more attractive in practice. Simply merging all the sources as one single source and then applying the single-source domain adaptation algorithms is obviously suboptimal since some source distributions will generally be more similar to the target than others. In that case, a uniformly merged source can be extremely dissimilar to the target distribution and violate the conditions for successful domain adaptation, e.g., those proposed by Ben-David et al. (2010b).

Hence, to learn from most related sources, many existing works in MDA (Zhao et al., 2018; Wen et al., 2020; Peng et al., 2019; Li et al., 2018) use the divergences between each source marginal and the target marginal $d(\mathcal{S}_i(X), \mathcal{T}(X)), \forall i \in [N]$ for inferring the domain relations. Such marginal alignment approaches have been demonstrated to be problematic even in single-source DA (Zhao et al., 2019) when the label distributions are different $\mathcal{T}(Y) \neq \mathcal{S}(Y)$ — which we denote as the **target shift**¹. Recently, other approaches (Redko et al., 2019a; Shui et al., 2021) have been proposed with an additional algorithmic layer that estimates the label ratio to mitigate the target shift. However, the estimation may be inaccurate in some situations, consequently degrading the performance.

The aforementioned target shift correcting methods are mathematically a special case of the joint distribution alignment. Moreover, when the number N of source domains is very large, the previous pairwise alignment methods will suffer from estimating numerous discriminators w.r.t both memory and sample complexity. Therefore, we propose to learn a joint alignment between a convex combination of source distributions $\mathcal{S}^\alpha = \sum_{i=1}^N \alpha_i \mathcal{S}_i, \forall \alpha \in \Delta_N$ and the target distribution \mathcal{T} , where the vector α of do-

Proceedings of the 26th International Conference on Artificial Intelligence and Statistics (AISTATS) 2023, Valencia, Spain. PMLR: Volume 206. Copyright 2023 by the author(s).

¹We use the terminology of **target shift** in the rest of the paper to avoid confusion with the label shift assumption, where $\mathcal{S}(X|Y) = \mathcal{T}(X|Y), \mathcal{S}(Y) \neq \mathcal{T}(Y)$.

main weights is optimized by the learning algorithm. The alignments are performed on a hidden-layer representation space commonly used for deep learning, while the corresponding analysis is rare. Moreover, the previous theories using Rademacher complexity or VC dimension are not algorithm-dependent. In contrast, we use information-theoretic tools to conduct the algorithm-specific generalization error analysis for representation learning of MDA.

Compared with previous works, the highlighted contributions of this paper are as follows:

Unified Approach We propose a unified approach for supervised and unsupervised MDA that conducts joint alignment w.r.t the representation space, which is appropriate for the target shift problem. The proposed algorithm simultaneously learns the domain weights, and the proposed non-pairwise alignment is more memory efficient than the previous works. Moreover, the proposed algorithm outperforms the previous approaches under a significant target shift in the unsupervised scenario, which is relevant in practice.

Algorithm-Dependent Generalization Bounds We first provide algorithm-dependent generalization bounds for MDA that depend on the mutual information between the model parameters and the input data: the less the model parameters depend on the data, the less the algorithm will overfit. We further apply the bounds on the proposed deep MDA algorithm and obtain a gradient norm estimation, which is used as a regularization coefficient for optimizing the domain weights.

2 RELATED WORK

Since domain adaptation has plenty of works in various settings, we only discuss the most related approaches in this section.

Metrics for Distribution Shift in DA Early theoretical works on single-source DA are based on the \mathcal{H} -divergence (Ben-David et al., 2006, 2010a,b). Then there emerge extensions on exploring different metrics, *e.g.*, discrepancy (Mansour et al., 2009), Wasserstein distance (Courty et al., 2014, 2017), Jensen-Shannon divergence (Shui et al., 2020, 2022b) and more general definitions of divergence like Integral Probability Metrics (IPM) (Zhang et al., 2012) and f -divergence (Acuna et al., 2021). For more comprehensive discussions, please address to Redko et al. (2019b) and Wang et al. (2023).

In this paper, we adopt Wasserstein distance for its numerous nice properties. It’s tighter than the KL divergence under the sub-Gaussian assumption through the transportation cost inequality and can capture the underlying geometry of the data. Courty et al. (2017) first theoretically an-

alyzes joint alignment on *example space* in single-source DA using Wasserstein distance. The algorithm based on solving the optimal transport problem scales quadratically in sample size, which is *intractable on large datasets*. Hence, Damodaran et al. (2018) empirically addresses this difficulty by alternately updating the mini-batch coupling matrix and network parameters to approximate the Wasserstein on the representation space. We theoretically analyze the joint alignment on the representation space w.r.t Wasserstein distance, providing algorithm-dependent bounds.

Information-Theoretic Learning in DA Recently, the information-theoretic analysis introduced by Xu & Raginsky (2017) and Russo & Zou (2019) has been used to provide a rigorous understanding of the generalization capabilities of deep learning models, such as complex meta-learning algorithms (Chen et al., 2021; Jose & Simeone, 2020). In contrast with the conventional VC-dimension and uniform stability bounds, it has the significant advantage of incorporating the dependence on the data distribution, the hypothesis space, and the learning algorithm. Wu et al. (2020, 2022) also use information-theoretic tools to derive bounds for single-source supervised DA. However, their bounds contain a non-optimizable term $D_{\text{KL}}(\mathcal{T}, \mathcal{S})$. On the contrary, we first provide fully algorithm-dependent bounds for both the unsupervised and supervised MDA, without loss of generality for single-source DA.

Multi-Source DA Konstantinov & Lampert (2019) and Shui et al. (2021) consider the *supervised* MDA case when few target labels are available. Konstantinov & Lampert (2019) estimates the domain relations with the pair-wise discrepancy $d(\mathcal{S}_i, \mathcal{T})$. Shui et al. (2021) reweights the losses and Wasserstein distances between domains with label ratio to mitigate the target shift, where the label ratio is estimated with the statistics given the few target labels and the source labels.

For *unsupervised* MDA, most existing approaches can be divided into two categories w.r.t the target shift. Marginal alignment methods make use of different divergences between each source marginal distribution and the target marginal distribution (*i.e.*, $d(\mathcal{S}_i(X), \mathcal{T}(X)), \forall i \in [N]$) to estimate the domain relations. *E.g.*, Zhao et al. (2018) and Li et al. (2018) use the \mathcal{H} -divergence, which is only suitable for binary classification. Wen et al. (2020) and Peng et al. (2019) adopt the discrepancy. The above methods that conduct unsupervised marginal alignments are unstable w.r.t target shift (Zhao et al., 2019).

Redko et al. (2019a) and Shui et al. (2021) consider addressing target shift with special cases of the joint distribution alignment. JCPOT (Redko et al., 2019a) assumes $\mathcal{S}_i(X|Y) = \mathcal{T}(X|Y), \forall i \in [N]$ and uses iterative Bregman projection to solve the constrained Wasserstein barycen-

ter problem, optimizing the class proportion to correct the target shift. However, the above label shift assumption often does not hold, and similar to Courty et al. (2017), the algorithm is *intractable on large datasets*. Shui et al. (2021) assumes the Generalized Label Shift (GLS) condition (Combes et al., 2020) is satisfied and estimates the label ratio with a black box shift estimation (Lipton et al., 2018). In addition, they minimize $d(\mathcal{S}_i(X|Y), \mathcal{T}(X|\hat{Y}))$ on representation space using the Wasserstein distance and the predicted pseudo labels. The theorems in Shui et al. (2021) are proposed for supervised MDA, without formal justification for unsupervised MDA using pseudo labels. MOST (Nguyen et al., 2021) does not optimize the domain relations and simply learns a source domain discriminator to obtain a *weighted ensemble expert* (teacher). A student classifier imitates the teacher by minimizing the Wasserstein distance between the *pseudo* source and *pseudo* target joint distribution using predicted labels.

Our approach differs from all the above-mentioned methods. We do not conduct a pairwise alignment. We define a combined source distribution and directly minimize the Wasserstein distance between the combined source distribution and the target distribution. We do not explicitly estimate the label ratio, and the target shift is tackled implicitly with the joint representation alignment. Moreover, we optimize the task weights α w.r.t the target and source domain shifts with an adaptive information-theoretical regularization coefficient.

3 PROBLEM SETUP

Basic Notations Without specification, we use upper case letters to denote random variables and the corresponding calligraphic letters to denote the corresponding sets on which they are defined, e.g., X, Y on \mathcal{X}, \mathcal{Y} .

Let \mathcal{X} be the *input space*. Let \mathcal{Y} be the *label space* and let $\mathcal{Z} \stackrel{\text{def}}{=} \mathcal{X} \times \mathcal{Y}$ be the *example space*. Then a set of N source distributions $\mathcal{S}_{1:N} \stackrel{\text{def}}{=} \{\mathcal{S}_1, \dots, \mathcal{S}_N\}$ and the target distribution \mathcal{T} are defined on \mathcal{Z} . Now we define a representation learning function $g : \mathcal{U} \times \mathcal{X} \rightarrow \tilde{\mathcal{X}}$ that transforms the inputs to feature representations, where g is parameterized by u defined on \mathcal{U} . Consequently, the distribution on the original example space has an induced distribution on the new space $\tilde{\mathcal{Z}} \stackrel{\text{def}}{=} \tilde{\mathcal{X}} \times \mathcal{Y}$. Given the weight simplex $\Delta_N \stackrel{\text{def}}{=} \{\alpha : \alpha_i \geq 0, \sum_{i=1}^N \alpha_i = 1\}$, we consider a mixture of source distributions $\mathcal{S}^\alpha \stackrel{\text{def}}{=} \sum_{i=1}^N \alpha_i \mathcal{S}_i$ as the *combined source distribution*. The induced distributions on $\tilde{\mathcal{Z}}$ are determined by u and denoted as $\tilde{\mathcal{T}}_u$ and $\tilde{\mathcal{S}}_u^\alpha \stackrel{\text{def}}{=} \sum_{i=1}^N \alpha_i \tilde{\mathcal{S}}_{i,u}$.

Then given the parameterized hypothesis space \mathcal{V} , we define a predictor $h : \mathcal{V} \times \tilde{\mathcal{X}} \rightarrow \mathcal{Y}$ that predicts a label given the feature representation. To measure the performance of the predictor, let $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$ be a positive-valued loss function.

Hence, the **population(true) risk** of $(u, v) \in \mathcal{U} \times \mathcal{V}$ on the target distribution and the combined source distribution are respectively defined as

$$\begin{aligned} R_{\mathcal{T}}(u, v) &\stackrel{\text{def}}{=} \mathbb{E}_{Z \sim \mathcal{T}} \ell(h(v, g(u, X)), Y) \\ &= \mathbb{E}_{\tilde{Z} \sim \tilde{\mathcal{T}}_u} \ell(h(v, \tilde{X}), Y) \\ R_{\mathcal{S}^\alpha}(u, v) &\stackrel{\text{def}}{=} \sum_{i=1}^N \alpha_i \mathbb{E}_{Z \sim \mathcal{S}_i} \ell(h(v, g(u, X)), Y) \\ &= \mathbb{E}_{\tilde{Z} \sim \tilde{\mathcal{S}}_u^\alpha} \ell(h(v, \tilde{X}), Y), \end{aligned}$$

where $Z \stackrel{\text{def}}{=} (X, Y)$, $\tilde{Z} \stackrel{\text{def}}{=} (\tilde{X}, Y)$.

Let the source datasets $S_{1:N}$ and the target dataset T be the examples sampled from the corresponding distributions, where $S_i = \{Z_{i,j}^s\}_{j=1}^{m_i}$, $Z_{i,j}^s = (X_{i,j}^s, Y_{i,j}^s) \sim \mathcal{S}_i, \forall i \in [N]$ and $T = \{Z_j^t\}_{j=1}^{m_t}$, $Z_j^t = (X_j^t, Y_j^t) \sim \mathcal{T}$. We further denote the corresponding dataset for the combined source distribution as S^α . Consequently, the corresponding **empirical risks** are defined as:

$$\begin{aligned} \hat{R}_{\mathcal{T}}(u, v) &\stackrel{\text{def}}{=} \frac{1}{m_t} \sum_{j=1}^{m_t} \ell(h(v, g(u, X_j^t)), Y_j^t) \\ \hat{R}_{\mathcal{S}^\alpha}(u, v) &\stackrel{\text{def}}{=} \sum_{i=1}^N \frac{\alpha_i}{m_i} \sum_{j=1}^{m_i} \ell(h(v, g(u, X_{i,j}^s)), Y_{i,j}^s). \end{aligned}$$

The target labels are not accessible during the training phase for **unsupervised** DA. Thus, several methods (Xie et al., 2018; Courty et al., 2017) apply the predicted labels of target inputs as **pseudo labels**. We formulate the corresponding setting in MDA. Let $f_p : \mathcal{X} \rightarrow \mathcal{Y}$ be a pseudo labeling function that maps the target inputs to the label space. In this paper, we set $f_p(x) = h(v, g(u, x))$ focusing on representation learning. Therefore, $\forall u, v \in \mathcal{U} \times \mathcal{V}$, let $\mathcal{T}_{u,v}$ be the *pseudo target distribution* on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ induced by f_p , and $\tilde{\mathcal{T}}_{u,v}$ be the corresponding distribution on $\tilde{\mathcal{Z}} = \tilde{\mathcal{X}} \times \mathcal{Y}$. So we have a random sample from the pseudo distribution noted as $\hat{Z} \stackrel{\text{def}}{=} (X, \hat{Y}) = (X, f_p(X))$. Finally, let us define the unlabeled target dataset as $T'_X = \{X_j^t\}_{j=1}^{m'_t}$, $X_j^t \sim \mathcal{T}(X)$, where $\mathcal{T}(X)$ denotes the target distribution \mathcal{T} marginalized on \mathcal{X} .

Now we introduce some common assumptions for deriving the multi-source DA bounds in this paper.

Assumption 1. (a) *The representation learning function $g : \mathcal{U} \times \mathcal{X} \rightarrow \tilde{\mathcal{X}}$ is K -Lipschitz² over \mathcal{X} for any $u \in \mathcal{U}$ w.r.t metrics $\rho_{\tilde{x}}$ and ρ_x , i.e., $\rho_{\tilde{x}}(g(u, x), g(u, x')) \leq K\rho_x(x, x')$. (b) *The predictor $h : \mathcal{V} \times \tilde{\mathcal{X}} \rightarrow \mathcal{Y}$ is L -Lipschitz over $\tilde{\mathcal{X}}$ for any $v \in \mathcal{V}$ w.r.t metrics ρ_y and $\rho_{\tilde{x}}$, i.e., $\rho_y(h(v, \tilde{x}), h(v, \tilde{x}')) \leq L\rho_{\tilde{x}}(\tilde{x}, \tilde{x}')$. (c) *The loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}^+$ is assumed to be symmetric, M -Lipschitz w.r.t ρ_y in the first argument and satisfying the triangle inequality.***

²The general definition of Lipschitzness is provided in the Appendix A.2.

4 SUPERVISED MDA

We first consider the supervised MDA regime. In most case, only a *few* amounts of labeled target data is accessible.

4.1 Population Risk Bound

Based on the optimal transport theory (Villani, 2009; Peyré et al., 2017), we obtain the following population risk bound using the Wasserstein distance (Definition A.6).

Theorem 4.1. $\forall (u, v) \in \mathcal{U} \times \mathcal{V}$, and $\forall \alpha \in \Delta_N$, if Assumption 1 is satisfied, then

$$|R_{\mathcal{T}}(u, v) - R_{S^\alpha}(u, v)| \leq \mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) \leq \mathbf{W}_1(\mathcal{T}, S^\alpha),$$

where the first \mathbf{W}_1 distance is defined on the metric space $(\tilde{\mathcal{Z}}, \rho_{\tilde{z}})$, for $\rho_{\tilde{z}}(\tilde{z}, \tilde{z}') = \ell(y, y') + LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}')$, and the second one is defined on metric space (\mathcal{Z}, ρ_z) , for $\rho_z(z, z') = \ell(y, y') + LMK\rho_x(x, x')$.

Theorem 4.1 illustrates the importance of learning the transformation function g , where $\mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha)$ is always smaller than $\mathbf{W}_1(\mathcal{T}, S^\alpha)$. The latter is a fixed term that characterizes the shift between the combined source and the target distribution given α . More specifically, there may not exist an α that gives $\mathbf{W}_1(\mathcal{T}, S^\alpha) = 0$. On the contrary, by optimizing the transformation parameter u , we may obtain $\mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) = 0$. Detailed proof is provided in Appendix C.1.

4.2 Generalization Bound

From Theorem 4.1, the target population risk is bounded by $R_{S^\alpha}(u, v) + \mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha)$. The estimation of the Wasserstein distance requires target data. Considering the extreme case $m_t \rightarrow +\infty$, minimizing the empirical estimation of this upper bound will never be better than directly minimizing the target empirical risk. So we propose to learn by considering the following **combined risk** with parameter $0 \leq \epsilon \leq 1$:

$$R_{S^\alpha, \mathcal{T}}^\epsilon(u, v) \stackrel{\text{def}}{=} (1-\epsilon)R_{\mathcal{T}}(u, v) + \epsilon(R_{S^\alpha}(u, v) + \mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha)).$$

Thus, we have $R_{\mathcal{T}}(u, v) \leq R_{S^\alpha, \mathcal{T}}^\epsilon(u, v)$. Then, let us denote the empirical combined risk as $\hat{R}_{S^\alpha, \mathcal{T}}^\epsilon(u, v) = (1-\epsilon)\hat{R}_{\mathcal{T}}(u, v) + \epsilon\hat{R}_{S^\alpha}(u, v) + \epsilon\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha)$.

We consider a stochastic learning algorithm \mathcal{A} for supervised MDA, which takes the N source datasets and the target dataset as input and then outputs the random parameters $(U, V) = \mathcal{A}(S^\alpha, T) \sim P_{U, V|S^\alpha, T}$, where $P_{U, V|S^\alpha, T}$ is the distribution on $\mathcal{U} \times \mathcal{V}$ induced by \mathcal{A} , given (S^α, T) . The expected generalization gap is then defined as

$$\begin{aligned} \text{gen}(\mathcal{T}, \mathcal{A}) &\stackrel{\text{def}}{=} \mathbb{E}_{U, V, S^\alpha, T}[R_{\mathcal{T}}(U, V) - \hat{R}_{S^\alpha, \mathcal{T}}^\epsilon(U, V)] \\ &\leq \mathbb{E}_{U, V, S^\alpha, T}[R_{S^\alpha, \mathcal{T}}^\epsilon(U, V) - \hat{R}_{S^\alpha, \mathcal{T}}^\epsilon(U, V)]. \end{aligned}$$

To bound the generalization gap for the Wasserstein distance, traditional methods apply the triangle inequality of the Wasserstein distance and then separately bound the gap for the source domain and target domain (Courty et al., 2017; Shui et al., 2021). The proofs are based on previous concentration results (Weed & Bach, 2019) of the Wasserstein distance with the additional assumption of a bounded loss. To obtain a non-vacuous algorithm-dependent bound, we address the problem in a different way. According to the Kantorovich-Rubinstein duality (Villani, 2009), we have

$$\mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) = \sup_{f: \|f\|_{Lip} \leq 1} \mathbb{E}_{\tilde{z} \sim \tilde{\mathcal{T}}_u} f(\tilde{z}) - \mathbb{E}_{\tilde{z}' \sim \tilde{S}_u^\alpha} f(\tilde{z}').$$

The supremum is over all the 1-Lipschitz functions $\mathcal{F} \stackrel{\text{def}}{=} \{f: \tilde{\mathcal{Z}} \rightarrow \mathbb{R}^+, |f(\tilde{z}) - f(\tilde{z}')| \leq \rho_{\tilde{z}}(\tilde{z}, \tilde{z}')\}$. Because of the intractability of learning this class of functions (Arjovsky et al., 2017), let us consider instead a parameterized class of functions $\tilde{\mathcal{F}} \stackrel{\text{def}}{=} \{\tilde{f}: \tilde{\mathcal{Z}} \times \mathcal{V}' \rightarrow \mathbb{R}^+, \|\tilde{f}\|_{Lip} \leq 1\}$ with the parameter set \mathcal{V}' . We assume the above supremum is attained in $\tilde{\mathcal{F}} \subset \mathcal{F}$. Under this assumption, we have

$$\mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) = \sup_{v' \in \mathcal{V}'} \mathbb{E}_{\tilde{z} \sim \tilde{\mathcal{T}}_u} \tilde{f}(v', \tilde{z}) - \mathbb{E}_{\tilde{z}' \sim \tilde{S}_u^\alpha} \tilde{f}(v', \tilde{z}').$$

Now we bring up the sub-Gaussian (Definition A.4) assumptions and provide a mutual information (Definition A.5) bound for supervised MDA.

Assumption 2. $\forall (u, v) \in \mathcal{U} \times \mathcal{V}$, $\ell(h(v, g(u, X)), Y)$ is σ -sub-Gaussian w.r.t $Z \sim \mathcal{T}$ and $Z \sim \mathcal{S}_i, \forall i \in [N]$. And for any proper choice of $\tilde{\mathcal{F}}$ with the supremum attained, $\forall (u, v') \in \mathcal{U} \times \mathcal{V}'$, $\tilde{f}(v', (g(u, X), Y)) \in \tilde{\mathcal{F}}$ is σ' -sub-Gaussian w.r.t $Z \sim \mathcal{T}$ and $Z \sim \mathcal{S}_i, \forall i \in [N]$.

Theorem 4.2. If Assumption 1 and 2 are satisfied, then we can bound the generalization gap of the supervised multi-source domain adaptation algorithm \mathcal{A} for any $\alpha \in \Delta_N$ and $0 \leq \epsilon \leq 1$ with:

$$\begin{aligned} \text{gen}(\mathcal{T}, \mathcal{A}) &\leq \sigma \sqrt{2\left(\frac{(1-\epsilon)^2}{m_t} + \sum_{i=1}^N \frac{\epsilon^2 \alpha_i^2}{m_i}\right) I(U, V; S^\alpha, T)} \\ &\quad + \sigma' \sqrt{2\epsilon^2 \left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right) I(U; S^\alpha, T)}. \end{aligned}$$

The above bound on the **generalization gap** for supervised MDA consists of two terms. The first term is characterized by the mutual information between the input datasets S^α, T and the output hypothesis U, V . The second term, which comes from the estimation of the Wasserstein distance on the representation space, depends on the mutual information between the representation parameter U and the input datasets. Notice that if we consider an arbitrarily random representation learning function (i.e., U does not depend on any data), the mutual information terms are equal to zero. However, the empirical risk will be very large. See detailed proof of Theorem 4.2 in the Appendix C.2.

Exploit Target Data Let us first focus on the case with few target data, *i.e.*, m_t is small. Usually, the number of source data samples is much more significant than m_t . So when $\sum_{i=1}^N \alpha_i^2/m_i \ll 1/m_t$, the generalization gap is dominated by the target data with rate $\mathcal{O}(\sqrt{1/m_t})$. In this case, it's preferable to use the target data for approximating the Wasserstein distance ($\epsilon = 1$) compared to directly train on the target data, because $I(U; S^\alpha, T) \leq I(U, V; S^\alpha, T)$ assuming $\sigma = \sigma'$. On the contrary, when $1/m_t \ll \sum_{i=1}^N \alpha_i^2/m_i$, we prefer to directly train on the target data ($\epsilon = 0$). For other cases, there exists a non-trivial value for ϵ that minimizes the generalization gap w.r.t the algorithm.

5 UNSUPERVISED MDA

For unsupervised domain adaptation, the target labels are not accessible during the training phase. So it's hard to guarantee a successful transfer without any precondition. Ben-David et al. (2010a) identified a necessary condition for successful unsupervised domain adaptation: the ideal joint error must be small enough, otherwise, no classifier can perform well on both the source and target domain. Therefore, a small ideal joint error is commonly assumed in the literature. Based on the same assumption, we provide the theorems for unsupervised MDA with pseudo labels.

Definition 5.1. (*Ideal joint error*) Let the ideal joint hypothesis be $(u^*, v^*) \stackrel{\text{def}}{=} \arg \min_{u, v} R_{S^\alpha}(u, v) + R_{\mathcal{T}}(u, v)$, then the ideal joint error is $R^* = R_{S^\alpha}(u^*, v^*) + R_{\mathcal{T}}(u^*, v^*)$.

5.1 Population Risk Bound

Definition 5.2. (*Joint Approximation Error*) Let $\tilde{x} = g(u, x)$, $\tilde{x}' = g(u, x')$ and the **optimal** coupling on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ given the pseudo labeling function $h(v, g(u, x))$ be

$$\gamma_{u, v}^* \stackrel{\text{def}}{=} \arg \min_{\gamma \in \Pi(\mathcal{T}_{u, v}, S^\alpha)} \int [LM \rho_{\tilde{x}}(\tilde{x}, \tilde{x}') + \ell(\hat{y}, y')] d\gamma(\hat{z}, z')$$

Then the joint approximation error is

$$R_{rep}^*(u, v) \stackrel{\text{def}}{=} LM \int_{\mathcal{Z} \times \mathcal{Z}} [\rho_{\tilde{x}}(g(u^*, x), g(u^*, x')) - \rho_{\tilde{x}}(g(u, x), g(u, x'))] d\gamma_{u, v}^*(\hat{z}, z'),$$

where $R_{rep}^*(u, v)$ characterizes the approximation error for the domain shift on the representation space of u^* (the representation counterpart of the ideal joint hypothesis).

Theorem 5.1. For any given $\alpha \in \Delta_N$ and $\forall u, v \in \mathcal{U} \times \mathcal{V}$, if Assumption 1 is satisfied, then the target population risk can be bounded by:

$$R_{\mathcal{T}}(u, v) \leq \mathbf{W}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha) + R_{rep}^*(u, v) + R^*.$$

See the detailed proof in Appendix D.1. From the above theorem, we notice that if we minimize $\mathbf{W}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha)$, $R_{rep}^*(u, v)$ may increase. Hence, for a successful unsupervised MDA w.r.t joint representation alignment, the proposed theory implies that it is sufficient to have a small $R_{rep}^*(u, v)$ for the (u, v) that minimizes $\mathbf{W}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha)$.

Target Shift An important benefit of the joint alignment methods is that it's favorable compared to the marginal alignment method when the target shift exists. Because with the joint optimal transport, we can automatically correct the target shift by minimizing $\mathbf{W}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha)$ w.r.t v . While the traditional marginal alignment methods that minimize $\mathbf{W}_1(\tilde{S}_u(X), \tilde{\mathcal{T}}_u(X)) + R_{S^\alpha}(u, v)$ cannot correct this shift since neither of the two terms in the objective are related to the divergence between the source and target label distributions.

5.2 Generalization Bound

Similarly to the above supervised MDA method, we use a stochastic algorithm \mathcal{A}^{un} , which takes the N source datasets and the unlabeled target dataset as input and then outputs the random parameters $(U, V) = \mathcal{A}^{un}(S^\alpha, T_X') \sim P_{U, V | S^\alpha, T_X'}$. Combine with Theorem 5.1, the expected generalization gap for unsupervised MDA is defined as:

$$gen(\mathcal{T}, \mathcal{A}^{un}) \stackrel{\text{def}}{=} \mathbb{E}_{U, V, S^\alpha, T_X'} [R_{\mathcal{T}}(U, V) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U, V}, \tilde{S}_U^\alpha)],$$

where $\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U, V}, \tilde{S}_U^\alpha)$ is the empirical estimation of $\mathbf{W}_1(\tilde{\mathcal{T}}_{U, V}, \tilde{S}_U^\alpha)$ that has different forms w.r.t the specific choice of $\tilde{\mathcal{F}}$, as discussed in Sec. 6.1. The definition without specification of $\tilde{\mathcal{F}}$ is provided in the Appendix. Now, consider the following sub-Gaussian assumption.

Assumption 3. For any proper choice of $\tilde{\mathcal{F}}$ with the supremum attained, $\forall (u, v, v') \in \mathcal{U} \times \mathcal{V} \times \mathcal{V}'$, $\tilde{f}(v', (g(u, X), \hat{Y})) \in \tilde{\mathcal{F}}$ is σ' -sub-Gaussian w.r.t $\hat{Z} \sim \mathcal{T}_{u, v}$ and $\forall (u, v') \in \mathcal{U} \times \mathcal{V}'$, $\tilde{f}(v', (g(u, X), Y)) \in \tilde{\mathcal{F}}$ is σ' -sub-Gaussian w.r.t $Z \sim \mathcal{S}_i, \forall i \in [N]$.

Then, we obtain the following algorithm-dependent bound on the generalization gap for joint representation alignment with pseudo labels.

Theorem 5.2. If Assumption 1 and 3 are satisfied, then we have the following bound on the expected generalization gap of the unsupervised multi-source domain adaptation algorithm with pseudo label \mathcal{A}^{un} for any $\alpha \in \Delta_N$:

$$gen(\mathcal{T}, \mathcal{A}^{un}) \leq \sqrt{2\sigma'^2 \left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t} \right) I(U, V; S^\alpha, T_X')} + \mathbb{E}_{U, V} R_{rep}^*(U, V) + R^*$$

Discussion The above bound contains the two terms $\mathbb{E}_{U, V} R_{rep}^*(U, V)$ and R^* , which are assumed to be small

enough for successful unsupervised representation transfer. Then the generalization gap of the joint representation alignment algorithm is upper bounded by the mutual information between (S^α, T'_X) and the algorithm output (U, V) . This is partly because the joint alignment also affects the learning of the predictor parameter V through the pseudo-label predictions. For unsupervised MDA, we often have a large m'_t , where $m'_t \gg m_i, \forall i \in [N]$. So the sample complexity is dominated by $\sum_{i=1}^N \frac{\alpha_i^2}{m_i}$. Consider the special case where all the sources domains have equivalently strong similarity with the target domain such that we have $\alpha_i = 1/N$ and $\sum_{i=1}^N \frac{\alpha_i^2}{m_i} = \frac{\sum_{i=1}^N 1/m_i}{N^2}$. In that case, the generalization gap can be decreased by increasing the number of domains, *i.e.*, $N \rightarrow \infty$. The detailed proof of Theorem 5.2 is provided in Appendix D.2.

6 UNIFIED APPROACH FOR DEEP MDA

6.1 Information-Theoretic MDA (IMDA)

Combine Supervised and Unsupervised MDA The supervised MDA algorithm does not exploit the information conveyed in the unlabeled target data. At the same time, the unsupervised MDA algorithms require a small ideal joint error, which may not be satisfied in practice. Hence, we consider combining the two methods that both conduct joint representation alignment. Then we propose a general approach by balancing the confidence of the two schemes through a weighting parameter $0 \leq \tau \leq 1$. When τ is neither 0 nor 1, the proposed algorithm becomes a semi-supervised transfer learning algorithm. So we define the following empirical risk:

$$\begin{aligned} \hat{R}_{\epsilon, \tau}^\alpha(u, v) &\stackrel{\text{def}}{=} \tau(1 - \epsilon)\hat{R}_T(u, v) + \tau\epsilon\hat{R}_{S^\alpha}(u, v) \\ &+ \tau\epsilon\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) + (1 - \tau)\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha) \end{aligned}$$

Choice of $\tilde{\mathcal{F}}$ Since $\ell(h(v, \cdot), \cdot), \forall v \in \mathcal{V}$ satisfies the Lipschitz condition with $|\ell(h(v, \tilde{x}), y) - \ell(h(v, \tilde{x}'), y')| \leq \rho_{\tilde{z}}(\tilde{z}, \tilde{z}')$ for the $\rho_{\tilde{z}}$ used in the two Wasserstein distances, we can choose $\tilde{\mathcal{F}} = \{\tilde{f} = \ell(h(v', \cdot), \cdot) : v' \in \mathcal{V}' = \mathcal{V}\}$. As discussed in Section 4.2, We assume that the supremum over the space of all the 1-Lipschitz functions $\tilde{\mathcal{F}}$ is attained with some $\tilde{f} \in \tilde{\mathcal{F}}$. Thus, the model contains two predictors $h(v, \cdot)$ and $h(v', \cdot)$ using the same network structure ($\mathcal{V}' = \mathcal{V}$) that make predictions from the feature representation $g(u, x)$. The former minimize the empirical risk $\tau(1 - \epsilon)\hat{R}_T(u, v) + \tau\epsilon\hat{R}_{S^\alpha}(u, v)$ w.r.t v . The other auxiliary predictor maximizes the empirical estimation of the two Wasserstein distances w.r.t v' . Then we minimize the sum of these empirical estimates w.r.t u, v . The above approach is different from the marginal representation alignment method, where the \mathbf{W}_1 distance is defined on $\tilde{\mathcal{X}}$ and the corresponding \tilde{f} is the domain discriminator adopted by DANN-like methods (Wen et al., 2020; Shui et al., 2021).

Optimize the Wasserstein Distances With the choice of \tilde{f} described above, the empirical risk of the pseudo target distribution is:

$$\hat{R}_{\mathcal{T}_{u, v}}(u, v, v') = \frac{1}{m'_t} \sum_{j=1}^{m'_t} \ell(h(v', g(u, X_j^t)), h(v, g(u, X_j^t)))$$

Therefore, we have the following approximation of the two Wasserstein distances:

$$\begin{aligned} \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) &= \max_{v'} [\hat{R}_T(u, v') - \hat{R}_{S^\alpha}(u, v')] \\ \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha) &= \max_{v'} [\hat{R}_{\mathcal{T}_{u, v}}(u, v, v') - \hat{R}_{S^\alpha}(u, v')] \end{aligned}$$

Then we get the two mini-max optimization objectives $\min_u \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha)$ and $\min_{u, v} \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{u, v}, \tilde{S}_u^\alpha)$, which can be implemented with a gradient reversal layer (Ganin et al., 2016).

Joint Optimization with Stochastic Gradient Algorithm

Pensia et al. (2018) provides methods bounding the mutual information with gradient updates for noisy iterative algorithms, which is especially suitable for deep learning, where the Stochastic Gradient Descent (SGD) and its variants are often applied. We apply this technique to analyze the aforementioned multi-source transfer scheme using the Stochastic Gradient Langevin Dynamics (SGLD) algorithm (Welling & Teh, 2011) — a stochastic variant of SGD with independent noise injection at each step.

At each iteration k , a batch of source data $S_{B_{1:N}}^k$ is sampled from the source datasets $S_{1:N}$. Simultaneously, a batch of labeled target data T_B^k and a batch of unlabeled target data $T_{X_B}^k$ are sampled from the corresponding target datasets. Subsequently, we update the representation parameter U and predictor parameter V with the following updates:

$$U_k = U_{k-1} - \eta_u^k G_u^k + \xi_u^k; \quad V_k = V_{k-1} - \eta_v^k G_v^k + \xi_v^k,$$

where $\xi_u^k \sim N(0, \sigma_k^2 I_{d_u})$ and $\xi_v^k \sim N(0, \sigma_k^2 I_{d_v})$ are injected isotropic Gaussian noise of variance σ_k^2 . Parameters η_u^k and η_v^k are, respectively, the learning rate for U and V at each step. G_u^k and G_v^k are, respectively, the gradient estimation of $\nabla_U \hat{R}_{\epsilon, \tau}^\alpha(U, V)$ and $\nabla_V \hat{R}_{\epsilon, \tau}^\alpha(U, V)$ using the batch datasets sampled at each iteration. Note that the sampling strategy needs to be agnostic to the previous iterates of the parameters.

6.2 Theoretical Guarantee via Gradient Norm Bound

We now present the gradient norm bound for the proposed IMDA algorithm.

Theorem 6.1. *Following Theorem 4.1, 5.1 4.2, 5.2, define a weight parameter τ that balances the unsupervised and supervised MDA. Adopt the aforementioned choice of $\tilde{\mathcal{F}}$ with $\sigma' = \sigma$ and the SGLD updates described above. Then we can obtain the gradient norm bound for the expected*

target risk, $\forall \alpha \in \Delta_N, \tau, \epsilon \in [0, 1]$, we have:

$$\begin{aligned} \mathbb{E}_{U,V,S^\alpha,T,T'_X} R_{\mathcal{T}}(U,V) &\leq \mathbb{E}_{U,V,S^\alpha,T,T'_X} \hat{R}_{\epsilon,\tau}^{\alpha}(U,V) \\ &+ \tau\sigma \sqrt{2\left(\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i}\right)(\delta_u + \delta_v)} \\ &+ \tau\epsilon\sigma \sqrt{2\left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right)\delta_u + (1-\tau)\mathbb{E}_{U,V} R_{rep}^*(U,V)} \\ &+ (1-\tau)\left(\sigma \sqrt{2\left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right)(\delta_u + \delta_v)} + R^*\right), \end{aligned}$$

$\delta_u \stackrel{def}{=} \sum_{k=1}^K \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2}{2\sigma_k^2}$ and $\delta_v \stackrel{def}{=} \sum_{k=1}^K \frac{(\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2}{2\sigma_k^2}$ are the accumulated gradient norm for U and V , respectively.

Theorem 6.1 is quite general. When $\tau = 1$, it covers the supervised MDA. When $\tau = 0$, it covers the unsupervised setting. In other cases, it can be applied for semi-supervised transfer learning. We can see that for semi-supervised transfer, the assumption for a small $R^* + \mathbb{E}_{U,V} R_{rep}^*(U,V)$ is weakened. When we have many labeled target data (m_t large), we can choose a large τ . Moreover, if we have only one source, α becomes a scalar and equals 1. Thus the bound can also be applied to single-source domain adaptation. See the proof in Appendix E.1.

6.3 Implementation and Discussion

Optimize the Domain Weights α From Theorem 6.1, for given u, v, v' and ϵ, τ , we can optimize α w.r.t the related term in the upper bound. So we get the following optimization problem with constants $C_0, C_1 > 0$:

$$\begin{aligned} \min_{\alpha} &\left((\epsilon\tau + C_0(1-\tau))\hat{R}_{S^\alpha}(u,v) \right. \\ &\quad \left. - (\epsilon\tau + 1 - \tau)\hat{R}_{S^\alpha}(u,v') \right. \\ &\quad \left. + C_1((1-\tau + \tau\epsilon)\sqrt{\delta_u + \delta_v} + \tau\epsilon\sqrt{\delta_u})R(\alpha) \right), \\ R(\alpha) &= \sqrt{\sum_{i=1}^N \frac{\alpha_i^2}{m_i}}, \text{ s.t. } \forall i \in [N], \alpha_i \geq 0, \sum_{i=1}^N \alpha_i = 1, \end{aligned}$$

which can be solved with a standard convex optimization toolbox. We fix α at each training epoch to optimize the network parameters u, v, v' by minimizing the empirical risk $\hat{R}_{\epsilon,\tau}^{\alpha}(u,v)$. Then, at the end of each epoch, we optimize α given the updated u, v, v' with the above convex optimization objective. We add $C_0\hat{R}_{S^\alpha}(u,v)$ for unsupervised MDA ($\tau = 0$) as approximation of $R_{S^\alpha}(u^*, v^*)$ in R^* when optimizing α , which does not exist in the optimization objective for u, v . Other higher order terms of α in δ_u, δ_v are ignored.

We emphasize that the above objective is different from Shui et al. (2021), which also optimizes the domain weights

α . The coefficient before the regularization term $R(\alpha)$ adapts with the gradient steps, which has an interesting interpretation. When the gradient step grows, both the model parameters and the domain weights α may over-fit to the existing datasets. Thus, the adaptive regularization coefficient of $R(\alpha)$ can gradually control the possibility of over-fitting to specific source distributions.

Complexity Comparison For memory complexity, the proposed method only requires **one** duplicate predictor. In contrast, WADN (Shui et al., 2021) requires N domain discriminator and $N|\mathcal{Y}|$ class feature centroids. MDAN and DARN (Wen et al., 2020) need N domain discriminators. M³SDA (Peng et al., 2019) and MDMN (Li et al., 2018) need N^2 domain discriminators. For time complexity, the proposed method needs to compute $\mathcal{O}(N)$ mini-batch risks and the corresponding backward gradients at each batch for optimizing the model parameters, which is the same as the other methods. Moreover, after each epoch, we need to optimize the domain weights with $\mathcal{O}(N)$ time complexity, which is an improvement compared to the $\mathcal{O}(N|\mathcal{Y}|)$ time complexity of WADN, where the $|\mathcal{Y}|$ comes from estimating the label ratio.

7 EXPERIMENTS

In this section, we test IMDA (Sec. 6) on the target shift MDA benchmark proposed by Shui et al. (2021). Experimental details and additional results can be found in Appendix F and G.

7.1 Datasets and Baselines

Amazon Review Dataset Blitzer et al. (2007) consists of positive and negative reviews from four domains: "Books", "DVD", "Electronics" and "Kitchen". We use each domain in turns as the target domain and the other three as the source domains. The data are pre-processed with the same approach as Chen et al. (2012), where the top-5000 frequent unigrams/bigrams are used as the bag-of-words features. The original Amazon Review data is label balanced. To create the *target shift* data, we follow the strategy of Shui et al. (2021) by randomly dropping 50% negative reviews in the source domains while keeping the target domain unchanged.

Digits Dataset contains four different tasks for digit recognition. MNIST (LeCun, 1998) and USPS (Hull, 1994) are standard handwriting digits datasets. Street View House Number (SVHN) (Netzer et al., 2011) is a digits dataset consisting of Google Street View images of house numbers. SYNTH (Ganin & Lempitsky, 2015) is a synthetic dataset applying various transformations on SVHN. As Shui et al. (2021), we also drop 50% of the data on digits 5-9 of all the sources with the target label distribution

Table 1: Accuracy(%) of Unsupervised MDA on Target Shift Data (Drop Rate 50%): Amazon Review(Left), Digits(Right)

Method	Target Domain				Average	Method	Target Domain				Average
	Books	DVD	Electronics	Kitchen			MNIST	SVHN	SYNTH	USPS	
Source	68.15 \pm 1.37	69.51 \pm 0.74	82.09 \pm 0.88	75.30 \pm 1.29	73.81	Source	84.93 \pm 1.50	67.14 \pm 1.40	78.11 \pm 1.31	86.02 \pm 1.12	79.05
DANN	65.59 \pm 1.35	67.23 \pm 0.71	80.49 \pm 1.11	74.71 \pm 1.53	72.00	DANN	86.99 \pm 1.53	69.56 \pm 2.26	78.73 \pm 1.30	86.81 \pm 1.74	80.52
MDAN	68.77 \pm 2.31	67.81 \pm 2.46	80.96 \pm 0.77	75.67 \pm 1.96	73.30	MDAN	87.86 \pm 2.24	69.13 \pm 1.56	79.77 \pm 1.69	86.50 \pm 1.59	80.81
MDMN	70.56 \pm 1.05	69.64 \pm 0.73	82.71 \pm 0.71	77.05 \pm 0.78	74.99	MDMN	87.31 \pm 1.88	69.84 \pm 1.59	80.27 \pm 0.88	86.61 \pm 1.41	81.00
M ³ SDA	69.09 \pm 1.26	68.67 \pm 1.37	81.34 \pm 0.66	76.10 \pm 1.47	73.79	M ³ SDA	87.22 \pm 1.70	68.89 \pm 1.93	80.01 \pm 1.77	86.39 \pm 1.68	80.87
DARN	69.88 \pm 1.91	69.63 \pm 1.09	80.83 \pm 1.13	77.47 \pm 1.05	74.45	DARN	86.96 \pm 1.27	68.99 \pm 1.76	80.65 \pm 1.12	86.89 \pm 1.64	80.87
WADN	75.74 \pm 1.60	78.46 \pm 1.72	82.10 \pm 2.09	81.05 \pm 2.25	79.34	WADN	89.04 \pm 0.86	71.68 \pm 1.27	82.06 \pm 0.97	90.06 \pm 1.12	83.21
IMDA(Ours)	75.21 \pm 2.3	78.15 \pm 1.79	83.93 \pm 0.24	81.44 \pm 1.44	79.68	IMDA(Ours)	89.26 \pm 0.89	76.6 \pm 0.86	81.76 \pm 0.69	84.45 \pm 1.2	83.02

Table 2: Accuracy(%) of Supervised MDA on Target Shift Data (Drop Rate 50%): Amazon Review(Left), Digits(Right)

Method	Target Domain				Average	Method	Target Domain				Average
	Books	DVD	Electronics	Kitchen			MNIST	SVHN	SYNTH	USPS	
Source + Tar	72.59 \pm 1.89	73.02 \pm 1.84	81.59 \pm 1.58	77.03 \pm 1.73	76.06	Source + Tar	79.63 \pm 1.74	56.48 \pm 1.90	69.64 \pm 1.38	86.29 \pm 1.56	73.01
DANN	67.35 \pm 2.28	66.33 \pm 2.42	78.03 \pm 1.72	74.31 \pm 1.71	71.50	DANN	86.77 \pm 1.30	69.13 \pm 1.09	78.82 \pm 1.35	86.54 \pm 1.03	80.32
MDAN	68.70 \pm 2.99	69.30 \pm 2.21	78.78 \pm 2.21	74.07 \pm 1.89	72.71	MDAN	86.93 \pm 1.05	68.25 \pm 1.53	79.80 \pm 1.17	86.23 \pm 1.41	80.30
MDMN	69.19 \pm 2.09	68.71 \pm 2.39	81.88 \pm 1.46	78.51 \pm 1.91	74.57	MDMN	77.59 \pm 1.36	69.62 \pm 1.26	78.93 \pm 1.64	87.26 \pm 1.13	78.35
M ³ SDA	69.28 \pm 1.78	67.40 \pm 0.46	76.28 \pm 0.81	76.50 \pm 1.19	72.36	M ³ SDA	85.88 \pm 2.06	68.84 \pm 1.05	76.29 \pm 0.95	87.15 \pm 1.10	79.54
RLUS	71.83 \pm 1.71	69.64 \pm 2.39	81.98 \pm 1.04	78.69 \pm 1.15	75.54	RLUS	87.61 \pm 1.08	70.50 \pm 0.94	79.52 \pm 1.30	86.70 \pm 1.13	81.08
MME	69.66 \pm 0.58	71.36 \pm 0.96	78.88 \pm 1.51	76.64 \pm 1.73	74.14	MME	87.24 \pm 0.95	65.20 \pm 1.35	80.31 \pm 0.60	87.88 \pm 0.76	80.16
DARN	69.48 \pm 2.28	69.59 \pm 2.90	80.66 \pm 1.38	77.25 \pm 0.94	74.25	DARN	86.56 \pm 1.48	68.76 \pm 1.36	80.4 \pm 1.27	86.82 \pm 1.09	80.64
WADN	74.52 \pm 1.45	77.39 \pm 1.04	81.40 \pm 1.61	81.52 \pm 1.64	78.71	WADN	88.02 \pm 1.27	70.54 \pm 1.06	81.58 \pm 0.96	90.48 \pm 1.11	82.66
IMDA(Ours)	75.17 \pm 0.94	77.92 \pm 1.73	83.04 \pm 1.34	80.11 \pm 1.17	79.06	IMDA(Ours)	84.82 \pm 1.15	75.46 \pm 1.2	79.87 \pm 1.04	91.6 \pm 0.96	82.94




Figure 1: Average Test Accuracy over All Domains w.r.t Different Levels of Target Shift on Amazon Dataset for (a) Supervised MDA with Few Target Labels and (b) Unsupervised MDA with Pseudo Labels. (c) Visualization of Domain Weights and Regularization Coefficients in Different Training Epochs on the Amazon Dataset for Supervised MDA with a Drop Rate of 0.4.

unchanged to create the *target shift* data.

Baselines We first introduce the baselines for unsupervised MDA. **Source** use all the labeled source datasets to train the predictor without adaptation. **DANN** (Ganin et al., 2016) is a single-source DA method, thus conducting marginal distribution matching between the merged source data and the target data. **MDAN** (Zhao et al., 2018), **MDMN** (Li et al., 2018), **M³SDA** (Peng et al., 2019) and **DARN** (Wen et al., 2020) are marginal distribution matching methods. **WADN** (Shui et al., 2021) first estimates the label ratio for correcting the target shift, then align the conditional distributions. Thus it also performs a (distinct) joint distribution matching method. For supervised MDA, **Source + Tar** is trained on merging the source data and the few labeled target data. Two more baselines, **RLUS** (Kon-

stantinov & Lampert, 2019) and **MME** (Saito et al., 2019), are also considered.

7.2 Results and Analysis

Results on Benchmark We compare our empirical results on the datasets mentioned above with the baselines provided by Shui et al. (2021). For supervised MDA, we consider the case where only a few labeled target samples are accessible. So we randomly sample 10% labeled target data as the training set, and the rest 90% are used as the test set. For unsupervised MDA, the above target data without labels are provided for training, and the results are tested on another unseen target dataset. We have *reproduced* the results of WADN and DARN using the code provided in these two papers and directly used the results presented in

Shui et al. (2021) for the other baselines. The comparison is shown in Tab. 1 for unsupervised MDA and Tab. 2 for supervised MDA. The reproduced results and those obtained by the proposed approach (IMDA) are averaged over ten runs. Each approach’s means and standard deviations are reported with the best value shown in bold.

From the tables, we can see that IMDA has improved on "Electronics" and "Kitchen" compared to the state-of-the-art (WADN) while having a slight drop on "Books" and "DVD" for unsupervised MDA. For supervised MDA, it improves on "Electronics" and "Books." The average Accuracy has a 0.3 – 0.4% *improvement* w.r.t the WADN. For the Digits dataset, the proposed approach has significant improvement on "SVHN" while a drop is shown on either "MNIST" or "USPS" for supervised and unsupervised MDA. W.r.t the average accuracy, we have *comparable results* to WADN with a slight improvement on supervised MDA and a slight decrease on unsupervised MDA. In addition, the worse performance of **Source + Tar** compared to all the divergence-based methods validates our statement for exploiting the target data in Theorem 4.2.

Ablation Study for Target Shift We conduct the ablation study on the Amazon dataset to show the effectiveness of IMDA for mitigating the target shift. We adjust the drop rate for negative reviews from 10% – 90%. The results are shown in Fig. 1 (a) and (b) for supervised MDA with few target labels and unsupervised MDA with pseudo labels, respectively.

Fig 1 (b) illustrates that IMDA can tackle significant target shift for unsupervised MDA and outperforms WADN, which needs, additionally, an estimate of the label ratio with the pseudo labels. The performance of WADN drops because the predictor is often inaccurate under a significant target shift due to an incorrect label ratio estimate, which in turn degrades the predictor’s performance.

For the supervised setting, only a *few labeled target samples* are provided. IMDA performs as well as WADN with a drop rate smaller than 50%, while its performance degrades when the drop rate goes up. The result is consistent with the proposed theory. When the source and target joint distributions differ significantly w.r.t the label distribution, correcting the target shift with simple statistics for the label ratio given the true target labels is more efficient than minimizing the Wasserstein distance, which requires adequate data.

Visualization of Domain Relations We show, in the heat-map of Fig. 1(c), the evolution of domain weights and the regularization coefficient w.r.t the training epoch, where a darker color means a larger weight. Moreover, we can see an adaptive effect of avoiding over-fitting to specific domains, *e.g.*, the weight of "DVD" was tuned down in some epochs when estimating "Electronics."

8 DISCUSSION AND CONCLUSION

Limitation Kumar et al. (2020) and Wang et al. (2022) conduct unsupervised DA with self-training when the target domain continuously shifts (gradual DA /continuous test-time DA). In this scenario, there often exists the phenomenon of error accumulation. IMDA also suffers from this issue due to the utilization of pseudo labels. Wang et al. (2022) proposes two methods to mitigate the error accumulation – the weight-averaged and augmentation-average pseudo labels. A similar weight-averaged approach can be used in IMDA to improve the algorithm’s stability.

Future Application in High-level Computer Vision Tasks Complex computer vision tasks like semantic segmentation often require structural adaptation. So extending the traditional feature alignment methods needs a specific feature map design. *E.g.*, ADVENT (Vu et al., 2019) uses the weighted self-information space. However, IMDA can be directly extended to the segmentation task using the mini-max objective w.r.t the empirical risks, where the structural information is contained in the definition of \mathcal{Y} .

Future Application in Fairness Recently, the ethical problems related to machine learning algorithms have gained growing concern in the community, especially regarding how to avoid implicit discrimination. Since unfair predictions are often associated with distribution shifts, domain adaptation methods are considered for improving fairness in Zhao & Gordon (2019); Schumann et al. (2019). The proposed approach can also be applied or incorporated with other fair learning algorithms (Hardt et al., 2016; Barocas et al., 2017; Liu et al., 2019; Shui et al., 2022a,c).

To conclude, this paper conducts an information-theoretic analysis of representation learning for both supervised and unsupervised MDA. We first provide fully algorithm-dependent bounds for MDA that apply to deep learning-based algorithms. Then, we propose a novel algorithm (IMDA) that mitigates the target shift and empirically achieves comparable performance to the previous state-of-the-art with improved memory efficiency. Notably, the proposed algorithm outperforms the previous works w.r.t a significant target shift in the more practically realistic unsupervised learning scenario. Finally, the proposed approach is fundamental and can be easily extended to other applications, *e.g.*, complex cv tasks and fairness.

Acknowledgements

We appreciate constructive feedback from anonymous reviewers and meta-reviewers. This work is supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) Discovery Grant, the Collaborative Research and Development Grant from SSQ Assurances and NSERC, and the China Scholarship Council.

References

- David Acuna, Guojun Zhang, Marc T Law, and Sanja Fidler. f-domain adversarial learning: Theory and algorithms. In *International Conference on Machine Learning*, pp. 66–75. PMLR, 2021.
- Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- Solon Barocas, Moritz Hardt, and Arvind Narayanan. Fairness in machine learning. *Nips tutorial*, 1:2017, 2017.
- 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, 2010a.
- Shai Ben-David, Tyler Lu, Teresa Luu, and Dávid Pál. Impossibility theorems for domain adaptation. In *International Conference on Artificial Intelligence and Statistics*, pp. 129–136, 2010b.
- John Blitzer. *Domain adaptation of natural language processing systems*. PhD thesis, University of Pennsylvania, 2008.
- John Blitzer, Mark Dredze, and Fernando Pereira. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th annual meeting of the association of computational linguistics*, pp. 440–447, 2007.
- Stéphane Boucheron, Gábor Lugosi, and Pascal Massart. *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press, 2013.
- Minmin Chen, Zhixiang Xu, Kilian Weinberger, and Fei Sha. Marginalized denoising autoencoders for domain adaptation. *arXiv preprint arXiv:1206.4683*, 2012.
- Qi Chen, Changjian Shui, and Mario Marchand. Generalization bounds for meta-learning: An information-theoretic analysis. *Advances in Neural Information Processing Systems*, 34:25878–25890, 2021.
- Remi Tachet des Combes, Han Zhao, Yu-Xiang Wang, and Geoff Gordon. Domain adaptation with conditional distribution matching and generalized label shift. *arXiv preprint arXiv:2003.04475*, 2020.
- Nicolas Courty, Rémi Flamary, and Devis Tuia. Domain adaptation with regularized optimal transport. In *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2014, Nancy, France, September 15-19, 2014. Proceedings, Part I 14*, pp. 274–289. Springer, 2014.
- Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. *Advances in Neural Information Processing Systems*, 30, 2017.
- Gabriela Csurka. A comprehensive survey on domain adaptation for visual applications. *Domain adaptation in computer vision applications*, pp. 1–35, 2017.
- Bharath Bhushan Damodaran, Benjamin Kellenberger, Rémi Flamary, Devis Tuia, and Nicolas Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 447–463, 2018.
- John Duchi. Lecture notes for statistics 311/electrical engineering 377. URL: [https://stanford.edu/class/stats311/Lectures/full notes.pdf](https://stanford.edu/class/stats311/Lectures/full%20notes.pdf). Last visited on, 2:23, 2016.
- 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.
- 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.
- Moritz Hardt, Eric Price, and Nati Srebro. Equality of opportunity in supervised learning. *Advances in neural information processing systems*, 29, 2016.
- Jiayuan Huang, Arthur Gretton, Karsten Borgwardt, Bernhard Schölkopf, and Alex J Smola. Correcting sample selection bias by unlabeled data. In *Advances in neural information processing systems*, pp. 601–608, 2007.
- Jonathan J. Hull. A database for handwritten text recognition research. *IEEE Transactions on pattern analysis and machine intelligence*, 16(5):550–554, 1994.
- Sharu Theresa Jose and Osvaldo Simeone. Information-theoretic generalization bounds for meta-learning and applications. *arXiv preprint arXiv:2005.04372*, 2020.
- Nikola Konstantinov and Christoph Lampert. Robust learning from untrusted sources. In *International conference on machine learning*, pp. 3488–3498. PMLR, 2019.
- Ananya Kumar, Tengyu Ma, and Percy Liang. Understanding self-training for gradual domain adaptation. In *International Conference on Machine Learning*, pp. 5468–5479. PMLR, 2020.
- Yann LeCun. The mnist database of handwritten digits. <http://yann.lecun.com/exdb/mnist/>, 1998.

- Yitong Li, David E Carlson, et al. Extracting relationships by multi-domain matching. *Advances in Neural Information Processing Systems*, 31, 2018.
- Zachary Lipton, Yu-Xiang Wang, and Alexander Smola. Detecting and correcting for label shift with black box predictors. In *International conference on machine learning*, pp. 3122–3130. PMLR, 2018.
- Lydia T Liu, Max Simchowitz, and Moritz Hardt. The implicit fairness criterion of unconstrained learning. In *International Conference on Machine Learning*, pp. 4051–4060. PMLR, 2019.
- Yishay Mansour, Mehryar Mohri, and Afshin Ros-tamizadeh. Domain adaptation: Learning bounds and algorithms. *arXiv preprint arXiv:0902.3430*, 2009.
- Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bis-sacco, Bo Wu, and Andrew Y Ng. Reading digits in nat-ural images with unsupervised feature learning. 2011.
- Tuan Nguyen, Trung Le, He Zhao, Quan Hung Tran, Truyen Nguyen, and Dinh Phung. Most: Multi-source domain adaptation via optimal transport for student-teacher learning. In *Uncertainty in Artificial Intelli-gence*, pp. 225–235. PMLR, 2021.
- Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data en-gineering*, 22(10):1345–1359, 2009.
- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.
- Ankit Pensia, Varun Jog, and Po-Ling Loh. Generaliza-tion error bounds for noisy, iterative algorithms. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pp. 546–550. IEEE, 2018.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport. *Center for Research in Economics and Statis-tics Working Papers*, (2017-86), 2017.
- Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. Optimal transport for multi-source domain adap-tation under target shift. In *The 22nd International Con-ference on Artificial Intelligence and Statistics*, pp. 849–858. PMLR, 2019a.
- Ievgen Redko, Emilie Morvant, Amaury Habrard, Marc Sebban, and Younes Bennani. *Advances in domain adaptation theory*. Elsevier, 2019b.
- Daniel Russo and James Zou. How much does your data exploration overfit? controlling bias via information us-age. *IEEE Transactions on Information Theory*, 66(1): 302–323, 2019.
- Kuniaki Saito, Donghyun Kim, Stan Sclaroff, Trevor Dar-rell, and Kate Saenko. Semi-supervised domain adap-tation via minimax entropy. In *Proceedings of the IEEE/CVF International Conference on Computer Vi-sion*, pp. 8050–8058, 2019.
- Candice Schumann, Xuezhi Wang, Alex Beutel, Jilin Chen, Hai Qian, and Ed H Chi. Transfer of ma-chine learning fairness across domains. *arXiv preprint arXiv:1906.09688*, 2019.
- Changjian Shui, Qi Chen, Jun Wen, Fan Zhou, Christian Gagné, and Boyu Wang. Beyond h-divergence: Do-main adaptation theory with jensen-shannon divergence. *arXiv preprint arXiv:2007.15567*, 6, 2020.
- Changjian Shui, Zijian Li, Jiaqi Li, Christian Gagné, Charles X Ling, and Boyu Wang. Aggregating from mul-tiple target-shifted sources. In *International Conference on Machine Learning*, pp. 9638–9648. PMLR, 2021.
- Changjian Shui, Qi Chen, Jiaqi Li, Boyu Wang, and Chris-tian Gagné. Fair representation learning through implicit path alignment. In *International Conference on Machine Learning*, pp. 20156–20175. PMLR, 2022a.
- Changjian Shui, Qi Chen, Jun Wen, Fan Zhou, Christian Gagné, and Boyu Wang. A novel domain adaptation the-ory with jensen–shannon divergence. *Knowledge-Based Systems*, 257:109808, 2022b.
- Changjian Shui, Gezheng Xu, Qi Chen, Jiaqi Li, Charles Ling, Tal Arbel, Boyu Wang, and Christian Gagné. On learning fairness and accuracy on multiple subgroups. *arXiv preprint arXiv:2210.10837*, 2022c.
- Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Dar-rell. Adversarial discriminative domain adaptation. In *Proceedings of the IEEE Conference on Computer Vi-sion and Pattern Recognition*, pp. 7167–7176, 2017.
- Cédric Villani. *Optimal transport: old and new*, volume 338. Springer, 2009.
- Tuan-Hung Vu, Himalaya Jain, Maxime Bucher, Matthieu Cord, and Patrick Pérez. Advent: Adversarial entropy minimization for domain adaptation in semantic seg-mentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2517–2526, 2019.
- Boyu Wang, Jorge A Mendez, Changjian Shui, Fan Zhou, Gezheng Xu, Christian Gagné, and Eric Eaton. Gap min-imization for knowledge sharing and transfer. *Journal of Machine Learning Research*, 24(33):1–57, 2023.
- Qin Wang, Olga Fink, Luc Van Gool, and Dengxin Dai. Continual test-time domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7201–7211, 2022.
- Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- Max Welling and Yee W Teh. Bayesian learning via stochastic gradient langevin dynamics. In *Proceedings*

of the 28th international conference on machine learning (ICML-11), pp. 681–688, 2011.

Junfeng Wen, Russell Greiner, and Dale Schuurmans. Domain aggregation networks for multi-source domain adaptation. In *International Conference on Machine Learning*, pp. 10214–10224. PMLR, 2020.

Xuetong Wu, Jonathan H Manton, Uwe Aickelin, and Jingge Zhu. Information-theoretic analysis for transfer learning. In *2020 IEEE International Symposium on Information Theory (ISIT)*, pp. 2819–2824. IEEE, 2020.

Xuetong Wu, Jonathan H Manton, Uwe Aickelin, and Jingge Zhu. An information-theoretic analysis for transfer learning: Error bounds and applications. *arXiv preprint arXiv:2207.05377*, 2022.

Shaoan Xie, Zibin Zheng, Liang Chen, and Chuan Chen. Learning semantic representations for unsupervised domain adaptation. In *International Conference on Machine Learning*, pp. 5423–5432, 2018.

Aolin Xu and Maxim Raginsky. Information-theoretic analysis of generalization capability of learning algorithms. In *Advances in Neural Information Processing Systems*, pp. 2524–2533, 2017.

Chao Zhang, Lei Zhang, and Jieping Ye. Generalization bounds for domain adaptation. *Advances in neural information processing systems*, 25, 2012.

Han Zhao and Geoff Gordon. Inherent tradeoffs in learning fair representations. *Advances in neural information processing systems*, 32, 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.

Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *International Conference on Machine Learning*, pp. 7523–7532. PMLR, 2019.

A DEFINITIONS

Definition A.1. (*Coupling*). Let (\mathcal{X}, μ) and (\mathcal{Y}, ν) be two probability spaces. Coupling μ, ν means constructing two random variables X and Y on some probability space (\mathcal{Z}, π) , such that $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, $(\text{proj}_{\mathcal{X}})_{\#}\pi = \mu$ and $(\text{proj}_{\mathcal{Y}})_{\#}\pi = \nu$, which means that π is the joint measure on $\mathcal{X} \times \mathcal{Y}$ with marginals μ, ν on \mathcal{X} and \mathcal{Y} respectively. The couple (X, Y) is called a coupling of (μ, ν) .

Definition A.2. (*Lipschitzness*). Given two metric spaces $(\mathcal{X}, \rho_x), (\mathcal{Y}, \rho_y)$, a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ is L -Lipschitz w.r.t \mathcal{X} if $\forall x, x' \in \mathcal{X}, \rho_y(f(x), f(x')) \leq L\rho_x(x, x')$. Specifically, if $\mathcal{X} \subset \mathbb{R}^{d_1}, \mathcal{Y} \subset \mathbb{R}^{d_2}$, we often choose $\rho_x(x, x') = \|x - x'\|$ and $\rho_y(\cdot, \cdot) = \|y - y'\|$ w.r.t norm $\|\cdot\|$.

Definition A.3. (*Transportation Cost Inequality*). A probability measure μ on (\mathcal{X}, ρ) satisfies an L^p transportation inequality with constant $c > 0$, if for every probability $\nu \ll \mu$, we have:

$$\mathbf{W}_p(\mu, \nu) \leq \sqrt{2cD_{\text{KL}}(\nu\|\mu)}$$

Definition A.4. (*Sub-Gaussian*). Define the cumulant generating function (CGF) of random variable X as $\psi_X(\lambda) \stackrel{\text{def}}{=} \log \mathbb{E}[e^{\lambda(X - \mathbb{E}[X])}]$. X is said to be σ -sub-Gaussian if

$$\psi_X(\lambda) \leq \frac{\lambda^2 \sigma^2}{2}, \forall \lambda \in \mathbb{R}.$$

Definition A.5. (*Mutual Information*). Let X and Y be arbitrary random variables and D_{KL} denote the KL divergence. The mutual information between X and Y is defined as:

$$I(X; Y) \stackrel{\text{def}}{=} D_{\text{KL}}(P(X, Y) \| P(X)P(Y))$$

Definition A.6. (*Wasserstein- p Distance*). Let the two distributions defined on the same Polish metric space (\mathcal{X}, ρ) , where $\rho(\cdot, \cdot)$ is a metric and $p \in [1, +\infty)$, $\Pi(\mu, \nu)$ is the set of all the couplings (see Definition A.1) of μ, ν . The Wasserstein distance with order p between μ and ν is defined as:

$$\mathbf{W}_p(\mu, \nu) \stackrel{\text{def}}{=} \inf_{\pi \in \Pi(\mu, \nu)} \left[\int_{\mathcal{X} \times \mathcal{X}} \rho(x, x')^p d\pi(x, x') \right]^{1/p}.$$

If $\rho(x, x') = \|x - x'\|$ and $p = 1$, the above definition is the Earth-Mover(EM) or Wasserstein-1 distance.

B TECHNICAL LEMMAS

Lemma B.1. (Donsker-Varadhan representation) Let P and Q be two probability measures defined on a set \mathcal{X} . Let $g : \mathcal{X} \rightarrow \mathbb{R}$ be a measurable function, and let $\mathbb{E}_{x \sim Q}[\exp g(x)] \leq \infty$. Then we have

$$D_{KL}(P||Q) = \sup_g \{\mathbb{E}_{x \sim P}[g(x)] - \log \mathbb{E}_{x \sim Q}[\exp g(x)]\}.$$

The proof of the above Lemma can be found in Corollary 4.15 (Boucheron et al., 2013) or Theorem 4.1 (Duchi, 2016).

Lemma B.2. Let X_1, X_2, \dots, X_n be independent σ_i -sub-Gaussian random variables. Then

$$\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right) \right] \leq \exp \left(\frac{\lambda^2 \sum_{i=1}^n \sigma_i^2}{2} \right), \forall \lambda \in \mathbb{R}$$

that is, $\sum_{i=1}^n X_i$ is $\sqrt{\sum_{i=1}^n \sigma_i^2}$ -sub-Gaussian.

Proof. $X_i, \forall i = \{1, 2, \dots, n\}$ is σ_i -sub-Gaussian, from Definition A.4, we have $\mathbb{E}[\exp(\lambda(X_i - \mathbb{E}X_i))] \leq \exp(\frac{\lambda^2 \sigma_i^2}{2})$.

Since the variables are independent, we have

$$\mathbb{E} \left[\exp \left(\lambda \sum_{i=1}^n (X_i - \mathbb{E}[X_i]) \right) \right] = \prod_{i=1}^n \mathbb{E}[\exp(\lambda(X_i - \mathbb{E}X_i))] \leq \prod_{i=1}^n \exp\left(\frac{\lambda^2 \sigma_i^2}{2}\right) = \exp\left(\frac{\lambda^2 \sum_{i=1}^n \sigma_i^2}{2}\right)$$

□

Lemma B.3. For independent variables X and Y , we have the following entropy inequality:

$$\max\{H(X), H(Y)\} \leq H(X + Y) \leq H(X, Y) = H(X) + H(Y)$$

Proof. Let $Z = X + Y$, from the definition of mutual information, we can obtain the well-known result that "conditioning reduces entropy". $I(Z; X) = H(Z) - H(Z|X) \geq 0$, then $H(Z) \geq H(Z|X)$. Similarly, we have $H(Z) \geq H(Z|Y)$. X, Y are independent, so we can obtain $H(Z|X) = H(Y|X) = H(Y)$ and $H(Z|Y) = H(X|Y) = H(X)$. Consequently, we have proved that $H(X + Y) \geq \max\{H(X), H(Y)\}$.

To prove the right-hand side, we need first prove that the entropy of a function of random variables is smaller than the joint entropy of the random variables. Denote $Z = f(X, Y) = X + Y$, we have

$$I((X, Y); Z) = H(Z) - H(Z|(X, Y)) = H(Z) = H(X, Y) - H((X, Y)|Z),$$

where the above equation is obtained with the fact that $H(Z|(X, Y)) = 0$.

From the definition of the entropy, we have $H((X, Y)|Z) \geq 0$. Combine the independence of X, Y , so we can get $H(Z) \leq H(X, Y) = H(X) + H(Y)$. □

C MISSING PROOFS IN SECTION 4

C.1 Proof of Theorem 4.1

Theorem. $\forall (u, v) \in \mathcal{U} \times \mathcal{V}$, and $\forall \alpha \in \Delta_N$, if Assumption 1 is satisfied, then

$$|R_{\mathcal{T}}(u, v) - R_{\mathcal{S}^\alpha}(u, v)| \leq \mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{\mathcal{S}}_u^\alpha) \leq \mathbf{W}_1(\mathcal{T}, \mathcal{S}^\alpha),$$

where the first \mathbf{W}_1 distance is defined on the metric space $(\tilde{\mathcal{Z}}, \rho_{\tilde{z}})$, for $\rho_{\tilde{z}}(\tilde{z}, \tilde{z}') = \ell(y, y') + LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}')$, and the second one is defined on metric space (\mathcal{Z}, ρ_z) , for $\rho_z(z, z') = \ell(y, y') + LMK\rho_x(x, x')$.

Proof. For any $\pi \in \Pi(\tilde{\mathcal{T}}_u, \tilde{\mathcal{S}}_u^\alpha)$, we have

$$\begin{aligned} |R_{\mathcal{T}}(u, v) - R_{\mathcal{S}^\alpha}(u, v)| &= |\mathbb{E}_{\tilde{Z} \sim \tilde{\mathcal{T}}_u} \ell(h(v, \tilde{X}), Y) - \mathbb{E}_{\tilde{Z}' \sim \tilde{\mathcal{S}}_u^\alpha} \ell(h(v, \tilde{X}'), Y')| \\ &= \left| \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} \ell(h(v, \tilde{x}), y) - \ell(h(v, \tilde{x}'), y') d\pi(\tilde{z}, \tilde{z}') \right| \\ &\leq \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} |\ell(h(v, \tilde{x}), y) - \ell(h(v, \tilde{x}'), y')| d\pi(\tilde{z}, \tilde{z}') \\ &= \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} |\ell(h(v, \tilde{x}), y) - \ell(h(v, \tilde{x}), y') + \ell(h(v, \tilde{x}), y') - \ell(h(v, \tilde{x}'), y')| d\pi(\tilde{z}, \tilde{z}') \\ &\leq \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} |\ell(h(v, \tilde{x}), y) - \ell(h(v, \tilde{x}), y')| + |\ell(h(v, \tilde{x}), y') - \ell(h(v, \tilde{x}'), y')| d\pi(\tilde{z}, \tilde{z}') \\ &\leq \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} (\ell(y, y') + M\rho_y(h(v, \tilde{x}), h(v, \tilde{x}'))) d\pi(\tilde{z}, \tilde{z}') \\ &\leq \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} (\ell(y, y') + LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}')) d\pi(\tilde{z}, \tilde{z}') \end{aligned}$$

The first two equalities hold with the definition of the population risk and Definition A.1 for coupling. The first inequality is obtained with Jensen's inequality for absolute function. And the second to last inequality is derived using the triangle inequality of the loss function and the Lipschitzness in Assumption 1.

Let the metric on $\tilde{\mathcal{Z}}$ be $\rho_{\tilde{z}}(\tilde{z}, \tilde{z}') = \ell(y, y') + LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}')$.

$$\begin{aligned} \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} \rho_{\tilde{z}}(\tilde{z}, \tilde{z}') d\pi(\tilde{z}, \tilde{z}') &= \int_{\mathcal{Z} \times \mathcal{Z}} (\ell(y, y') + LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}')) d\pi(\tilde{z}, \tilde{z}') \\ &= \int_{\mathcal{Z} \times \mathcal{Z}} (\ell(y, y') + LM\rho_{\tilde{x}}(g(u, x), g(u, x'))) d\gamma(z, z') \\ &\leq \int_{\mathcal{Z} \times \mathcal{Z}} (\ell(y, y') + LMK\rho_x(x, x')) d\gamma(z, z') \end{aligned}$$

Let the metric on \mathcal{Z} be $\rho_z(z, z') = \ell(y, y') + LMK\rho_x(x, x')$.

Let the optimal coupling on the representation space be $\pi_u^* = \arg \min_{\pi \in \Pi(\tilde{\mathcal{T}}_u, \tilde{\mathcal{S}}_u^\alpha)} \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} \rho_{\tilde{z}}(\tilde{z}, \tilde{z}') d\pi(\tilde{z}, \tilde{z}')$ and the corresponding coupling on the original example space be γ_u^* .

Let the optimal coupling on the original example space be $\gamma^* = \arg \min_{\gamma \in \Pi(\mathcal{T}, \mathcal{S}^\alpha)} \int_{\mathcal{Z} \times \mathcal{Z}} \rho_z(z, z') d\gamma(z, z')$ and the corresponding coupling on the representation space be π_u . So we have:

$$\begin{aligned} |R_{\mathcal{T}}(u, v) - R_{\mathcal{S}^\alpha}(u, v)| &\leq \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} \rho_{\tilde{z}}(\tilde{z}, \tilde{z}') d\pi_u^*(\tilde{z}, \tilde{z}') = \mathbf{W}_1(\tilde{\mathcal{T}}_u, \tilde{\mathcal{S}}_u^\alpha) \\ &\leq \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} \rho_{\tilde{z}}(\tilde{z}, \tilde{z}') d\pi_u(\tilde{z}, \tilde{z}') \leq \int_{\mathcal{Z} \times \mathcal{Z}} \rho_z(z, z') d\gamma^*(z, z') = \mathbf{W}_1(\mathcal{T}, \mathcal{S}^\alpha) \end{aligned}$$

The last inequality is obtained since $\pi_u \neq \pi_u^*$ may not be the optimal coupling on the representation space. \square

C.2 Proof of Theorem 4.2

Theorem. *If Assumption 1 and 2 are satisfied, then we can bound the generalization gap of the supervised multi-source domain adaptation algorithm \mathcal{A} for any $\alpha \in \Delta_N$ and $0 \leq \epsilon \leq 1$ with:*

$$\text{gen}(\mathcal{T}, \mathcal{A}) \leq \sigma \sqrt{2\left(\frac{(1-\epsilon)^2}{m_t} + \sum_{i=1}^N \frac{\epsilon^2 \alpha_i^2}{m_i}\right) I(U, V; S^\alpha, T)} + \sigma' \sqrt{2\epsilon^2 \left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right) I(U; S^\alpha, T)}.$$

Proof.

$$\begin{aligned} \text{gen}(\mathcal{T}, \mathcal{A}) &\leq \mathbb{E}_{U, V, S^\alpha, T} [R_{S^\alpha, \mathcal{T}}^\epsilon(U, V) - \hat{R}_{S^\alpha, \mathcal{T}}^\epsilon(U, V)] \\ &= \mathbb{E}_{U, V, S^\alpha, T} [(1-\epsilon)(R_{\mathcal{T}}(U, V) - \hat{R}_{\mathcal{T}}(U, V)) + \epsilon(R_{S^\alpha}(U, V) - \hat{R}_{S^\alpha}(U, V))] \\ &\quad + \epsilon \mathbb{E}_{U, V, S^\alpha, T} [\mathbf{W}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha)] \end{aligned}$$

Step 1 Bounding the Empirical Risks

$$B_r = \mathbb{E}_{U, V, S^\alpha, T} \left[(1-\epsilon)[R_{\mathcal{T}}(U, V) - \hat{R}_{\mathcal{T}}(U, V)] + \epsilon[R_{S^\alpha}(U, V) - \hat{R}_{S^\alpha}(U, V)] \right]$$

$(U, V) \in \mathcal{U} \times \mathcal{V}$ are random variables outputs by the supervised MDA algorithm, which depend on the input datasets $S_{1:N}$, T and a fixed α (not random). Let (\tilde{U}, \tilde{V}) be an independent copy of (U, V) such that $(\tilde{U}, \tilde{V}) \perp\!\!\!\perp (S_{1:N}, T)$. Then let us define

$$\begin{aligned} f(U, V, S^\alpha, T) &\stackrel{\text{def}}{=} (1-\epsilon)\hat{R}_{\mathcal{T}}(U, V) + \epsilon\hat{R}_{S^\alpha}(U, V) \\ &= (1-\epsilon) \frac{1}{m_t} \sum_{j=1}^{m_t} \ell(h(V, g(U, X_j^t)), Y_j^t) + \epsilon \sum_{i=1}^N \frac{\alpha_i}{m_i} \sum_{j=1}^{m_i} \ell(h(V, g(U, X_{i,j}^s)), Y_{i,j}^s) \end{aligned}$$

Since the loss function is σ -sub-Gaussian (Assumption 2), applying Lemma B.2, we can obtain that $f(\tilde{U}, \tilde{V}, S^\alpha, T)$ is $\sqrt{\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i}}$ σ -sub-Gaussian.

$$\begin{aligned} I(U, V; S^\alpha, T) &= D_{\text{KL}}(P_{U, V, S^\alpha, T} \| P_{U, V} P_{S^\alpha, T}) \\ &= \sup_g \{ \mathbb{E}_{U, V, S^\alpha, T} \mathbf{g}(U, V, S^\alpha, T) - \log \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T} [\exp^{\mathbf{g}(\tilde{U}, \tilde{V}, S^\alpha, T)}] \} \\ &\geq \lambda \mathbb{E}_{U, V, S^\alpha, T} [f(U, V, S^\alpha, T)] - \lambda \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T} [f(\tilde{U}, \tilde{V}, S^\alpha, T)] - \psi_{\tilde{U}, \tilde{V}, S^\alpha, T}(\lambda), \forall \lambda \in \mathbb{R} \\ &\geq \lambda \mathbb{E}_{U, V, S^\alpha, T} [f(U, V, S^\alpha, T)] - \lambda \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T} [f(\tilde{U}, \tilde{V}, S^\alpha, T)] - \frac{\lambda^2 \sigma^2 \left(\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right)}{2} \end{aligned}$$

We have:

$$\begin{aligned} \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T} [f(\tilde{U}, \tilde{V}, S^\alpha, T)] &= \epsilon \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha} \hat{R}_{S^\alpha}(\tilde{U}, \tilde{V}) + (1-\epsilon) \mathbb{E}_{\tilde{U}, \tilde{V}, T} \hat{R}_{\mathcal{T}}(\tilde{U}, \tilde{V}) \\ &= \epsilon \mathbb{E}_{\tilde{U}, \tilde{V}} R_{S^\alpha}(\tilde{U}, \tilde{V}) + (1-\epsilon) \mathbb{E}_{\tilde{U}, \tilde{V}} R_{\mathcal{T}}(\tilde{U}, \tilde{V}) \\ &= \epsilon \mathbb{E}_{U, V, S^\alpha, T} R_{S^\alpha}(U, V) + (1-\epsilon) \mathbb{E}_{U, V, S^\alpha, T} R_{\mathcal{T}}(U, V) \end{aligned}$$

Place the above term into the inequality, we get $-\lambda B_r - \frac{\lambda^2 \sigma^2 \left(\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right)}{2} \leq I(U, V; S^\alpha, T)$

Consequently, we have $|B_r| \leq \sqrt{2\left(\sigma^2 \left(\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right) I(U, V; S^\alpha, T)}\right)}$.

Step 2 Bounding the Wasserstein Distance

$$B_w = \mathbb{E}_{U, V, S^\alpha, T} \left[\epsilon \mathbf{W}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) - \epsilon \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) \right] = \mathbb{E}_{U, S^\alpha, T} \left[\epsilon \mathbf{W}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) - \epsilon \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) \right]$$

Let

$$f'(U, S^\alpha, T) = \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) = \sup_{v'} \left[\frac{1}{m_t} \sum_{j=1}^{m_t} \tilde{f}(v', (g(U, X_j^t), Y_j^t)) - \sum_{i=1}^N \frac{\alpha_i}{m_i} \sum_{j=1}^{m_i} \tilde{f}(v', (g(U, X_{i,j}^s), Y_{i,j}^s)) \right],$$

$$\begin{aligned}
I(U; S^\alpha, T) &= D_{\text{KL}}(P_{U, S^\alpha, T} \| P_U P_{S^\alpha, T}) \\
&= \sup_{\mathfrak{g}} \{ \mathbb{E}_{U, S^\alpha, T} \mathfrak{g}(U, S^\alpha, T) - \log \mathbb{E}_{\tilde{U}, S^\alpha, T} [\exp^{\mathfrak{g}(\tilde{U}, S^\alpha, T)}] \} \\
&\geq \lambda \mathbb{E}_{U, S^\alpha, T} [f'(U, S^\alpha, T)] - \lambda \mathbb{E}_{\tilde{U}, S^\alpha, T} [f'(\tilde{U}, S^\alpha, T)] - \psi_{\tilde{U}, S^\alpha, T}(\lambda), \forall \lambda \in \mathbb{R}
\end{aligned}$$

\mathcal{V}' is assumed to be the subset that does not affect the supremum. So we have:

$$\begin{aligned}
\mathbb{E}_{\tilde{U}, S^\alpha, T} [f'(\tilde{U}, S^\alpha, T)] &= \mathbb{E}_{\tilde{U}} \sup_{v'} \mathbb{E}_{S^\alpha, T} \left[\frac{1}{m_t} \sum_{j=1}^{m_t} \tilde{f}(v', (g(\tilde{U}, X_j^t), Y_j^t)) - \sum_{i=1}^N \frac{\alpha_i}{m_i} \sum_{j=1}^{m_i} \tilde{f}(v', (g(\tilde{U}, X_{i,j}^s), Y_{i,j}^s)) \right] \\
&= \mathbb{E}_{\tilde{U}} \mathbf{W}_1(\tilde{\mathcal{T}}_{\tilde{U}}, \tilde{\mathcal{S}}_{\tilde{U}}) = \mathbb{E}_{U, S^\alpha, T} \mathbf{W}_1(\tilde{\mathcal{T}}_U, \tilde{\mathcal{S}}_U)
\end{aligned}$$

From Assumption 2, \tilde{f} is σ' -sub-Gaussian for any u and v' , so the supremum does not affect the sub-Gaussianity. Apply Lemma B.2, we know that $\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_u, \tilde{\mathcal{S}}_u)$ is $\sqrt{\frac{1}{m_t} + \sum_{i=1}^N \frac{\alpha_i^2}{m_i}} \sigma'$ -sub-Gaussian.

With the same proof process as bounding the empirical risks, we obtain that:

$$|B_w| \leq \sqrt{2\sigma'^2 \epsilon^2 \left(\frac{1}{m_t} + \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right) I(U; S^\alpha, T)}$$

Finally, we have:

$$\begin{aligned}
gen(\mathcal{T}, \mathcal{A}) &\leq B_r + B_w \\
&\leq \sqrt{2(\sigma^2 \frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i}) I(U, V; S^\alpha, T)} + \sqrt{2\sigma'^2 \epsilon^2 \left(\frac{1}{m_t} + \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right) I(U; S^\alpha, T)}
\end{aligned}$$

Conclude the proof. □

D MISSING PROOFS IN SECTION 5

D.1 Proof of Theorem 5.1

Definition D.1. (Joint Approximation Error) Let $\tilde{x} = g(u, x)$, $\tilde{x}' = g(u, x')$, then let the **optimal** coupling on $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$ given the pseudo labeling function $h(v, g(u, x))$ be

$$\gamma_{u,v}^* \stackrel{\text{def}}{=} \arg \min_{\gamma \in \Pi(\mathcal{T}_{u,v}, \mathcal{S}^\alpha)} \int_{\mathcal{Z} \times \mathcal{Z}} [LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}') + \ell(\hat{y}, y')] d\gamma(\hat{z}, z')$$

Then the joint approximation error is

$$R_{rep}^*(u, v) \stackrel{\text{def}}{=} LM \int_{\mathcal{Z} \times \mathcal{Z}} [\rho_{\tilde{x}}(g(u^*, x), g(u^*, x')) - \rho_{\tilde{x}}(g(u, x), g(u, x'))] d\gamma_{u,v}^*(\hat{z}, z'),$$

where $R_{rep}^*(u, v)$ characterizes the approximation error for the domain shift on the representation space of u^* (the representation counterpart of the ideal joint hypothesis).

Theorem. For any given $\alpha \in \Delta_N$ and $\forall u, v \in \mathcal{U} \times \mathcal{V}$, if Assumption 1 is satisfied, then the target population risk can be bounded by:

$$R_{\mathcal{T}}(u, v) \leq \mathbf{W}_1(\tilde{\mathcal{T}}_{u,v}, \tilde{\mathcal{S}}_u^\alpha) + R_{rep}^*(u, v) + R^*.$$

Proof. Let $u^*, v^* = \arg \min_{u,v} R_{\mathcal{S}^\alpha}(u, v) + R_{\mathcal{T}}(u, v)$ be the ideal joint hypothesis, we have:

$$\begin{aligned} R_{\mathcal{T}}(u, v) &= \mathbb{E}_{Z \sim \mathcal{T}} \ell(h(v, g(u, X)), Y) \\ &\leq \mathbb{E}_{Z \sim \mathcal{T}} \ell(h(v, g(u, X)), h(v^*, g(u^*, X))) + \mathbb{E}_{Z \sim \mathcal{T}} \ell(h(v^*, g(u^*, X)), Y) \\ &= \mathbb{E}_{\hat{Z} \sim \mathcal{T}_{u,v}} \ell(h(v^*, g(u^*, X)), \hat{Y}) + R_{\mathcal{T}}(u^*, v^*) \\ &= \mathbb{E}_{Z \sim \mathcal{T}_{u,v}} \ell(h(v^*, g(u^*, X)), Y) - \mathbb{E}_{Z' \sim \mathcal{S}^\alpha} \ell(h(v^*, g(u^*, X')), Y') \\ &\quad + \mathbb{E}_{Z' \sim \mathcal{S}^\alpha} \ell(h(v^*, g(u^*, X')), Y') + R_{\mathcal{T}}(u^*, v^*) \\ &\leq |\mathbb{E}_{\hat{Z} \sim \mathcal{T}_{u,v}} \ell(h(v^*, g(u^*, X)), \hat{Y}) - \mathbb{E}_{Z' \sim \mathcal{S}^\alpha} \ell(h(v^*, g(u^*, X')), Y')| \\ &\quad + R_{\mathcal{S}^\alpha}(u^*, v^*) + R_{\mathcal{T}}(u^*, v^*). \end{aligned}$$

For alignment on representation space, we have $\forall \gamma \in \Pi(\mathcal{T}_{u,v}, \mathcal{S}^\alpha)$ that:

$$\begin{aligned} &|\mathbb{E}_{Z \sim \mathcal{T}_{u,v}} \ell(h(v^*, g(u^*, X)), \hat{Y}) - \mathbb{E}_{Z' \sim \mathcal{S}^\alpha} \ell(h(v^*, g(u^*, X')), Y')| \\ &\leq \int_{\mathcal{Z} \times \mathcal{Z}} |\ell(h(v^*, g(u^*, x)), \hat{y}) - \ell(h(v^*, g(u^*, x')), y')| d\gamma(\hat{z}, z') \\ &\leq \int_{\mathcal{Z} \times \mathcal{Z}} |\ell(h(v^*, g(u^*, x)), y') - \ell(h(v^*, g(u^*, x')), y')| \\ &\quad + |\ell(h(v^*, g(u^*, x)), \hat{y}) - \ell(h(v^*, g(u^*, x)), y')| d\gamma(\hat{z}, z') \\ &\leq \int_{\mathcal{Z} \times \mathcal{Z}} LM\rho_{\tilde{x}}(g(u^*, x), g(u^*, x')) + \ell(\hat{y}, y') d\gamma(\hat{z}, z') \\ &= \int_{\mathcal{Z} \times \mathcal{Z}} LM\rho_{\tilde{x}}(g(u^*, x), g(u^*, x')) - LM\rho_{\tilde{x}}(g(u, x), g(u, x')) d\gamma(\hat{z}, z') \\ &\quad + \int_{\mathcal{Z} \times \mathcal{Z}} LM\rho_{\tilde{x}}(g(u, x), g(u, x')) + \ell(\hat{y}, y') d\gamma(\hat{z}, z') \end{aligned}$$

Let the metric on $\tilde{\mathcal{Z}}$ be $\rho_{\tilde{x}}(\hat{z}, \hat{z}') = LM\rho_{\tilde{x}}(\tilde{x}, \tilde{x}') + \ell(\hat{y}, y')$. Let the optimal coupling on the transformed joint space $\tilde{\mathcal{Z}}$ be $\pi_{u,v}^* = \arg \min_{\pi \in \Pi(\tilde{\mathcal{T}}_{u,v}, \tilde{\mathcal{S}}_u^\alpha)} \int_{\tilde{\mathcal{Z}} \times \tilde{\mathcal{Z}}} \rho_{\tilde{x}}(\hat{z}, \hat{z}') d\pi(\hat{z}, \hat{z}')$, so the corresponding coupling on the original example space is:

$$\gamma_{u,v}^* = \arg \min_{\gamma \in \Pi(\mathcal{T}_{u,v}, \mathcal{S}^\alpha)} \int_{\mathcal{Z} \times \mathcal{Z}} LM\rho_{\tilde{x}}(g(u, x), g(u, x')) + \ell(h(v, g(u, x')), y') d\gamma(\hat{z}, z').$$

Hence,

$$R_{\mathcal{T}}(u, v) \leq \mathbf{W}_1(\tilde{\mathcal{T}}_{u,v}, \tilde{\mathcal{S}}_u^\alpha) + R_{rep}^*(u, v) + R^* .$$

where $R_{rep}^*(u, v) = \int_{\mathcal{Z} \times \mathcal{Z}} LM \rho_{\tilde{x}}(g(u^*, x), g(u^*, x')) - LM \rho_{\tilde{x}}(g(u, x), g(u, x')) d\gamma_{u,v}^*(\hat{z}, z')$ is the joint approximation error (Definition D.1), which is controlled by the representation learning function $g(u, x)$ and the optimal coupling $\gamma_{u,v}^*(\hat{z}, z')$. \square

D.2 Proof of Theorem 5.2

Theorem. *If Assumption 1 and 3 are satisfied, then we have the following bound on the expected generalization gap of the unsupervised multi-source domain adaptation algorithm with pseudo label \mathcal{A}^{un} for any $\alpha \in \Delta_N$:*

$$gen(\mathcal{T}, \mathcal{A}^{un}) \leq \sqrt{2\sigma'^2 \left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t'} \right) I(U, V; S^\alpha, T_X') + \mathbb{E}_{U,V} R_{rep}^*(U, V) + R^*}$$

Proof. By definition:

$$gen(\mathcal{T}, \mathcal{A}^{un}) \stackrel{\text{def}}{=} \mathbb{E}_{U,V,S^\alpha,T_X'} [R_{\mathcal{T}}(U, V) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha)]$$

Combine Theorem 5.1, we have

$$gen(\mathcal{T}, \mathcal{A}^{un}) \leq \mathbb{E}_{U,V,S^\alpha,T_X'} [\mathbf{W}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha)] + \mathbb{E}_{U,V} R_{rep}^*(U, V) + R^*$$

Let

$$\begin{aligned} f'(U, V, S^\alpha, T_X') &= \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) \\ &= \sup_{v'} \left[\frac{1}{m_t'} \sum_{j=1}^{m_t'} \tilde{f}(v', (g(U, X_j^t), h(V, g(U, X_j^t)))) - \sum_{i=1}^N \frac{\alpha_i}{m_i} \sum_{j=1}^{m_i} \tilde{f}(v', (g(U, X_{i,j}^s), Y_{i,j}^s)) \right] \end{aligned}$$

$$\begin{aligned} I(U, V; S^\alpha, T_X') &= D_{\text{KL}}(P_{U,V,S^\alpha,T} \| P_{U,V} P_{S^\alpha,T_X'}) \\ &= \sup_{\mathfrak{g}} \{ \mathbb{E}_{U,V,S^\alpha,T_X'} \mathfrak{g}(U, V, S^\alpha, T_X') - \log \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T_X'} [\exp^{\mathfrak{g}(\tilde{U}, \tilde{V}, S^\alpha, T_X')}] \} \\ &\geq \lambda \mathbb{E}_{U,V,S^\alpha,T_X'} [f(U, V, S^\alpha, T_X')] - \lambda \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T_X'} [f(\tilde{U}, \tilde{V}, S^\alpha, T_X')] - \psi_{\tilde{U}, \tilde{V}, S^\alpha, T_X'}(\lambda), \forall \lambda \in \mathbb{R} \end{aligned}$$

\mathcal{V}' is assumed to be the set that does not affect the supremum, so we have:

$$\begin{aligned} \mathbb{E}_{\tilde{U}, \tilde{V}, S^\alpha, T_X'} [f'(\tilde{U}, \tilde{V}, S^\alpha, T_X')] &= \mathbb{E}_{\tilde{U}, \tilde{V}} \sup_{v'} \mathbb{E}_{\hat{Z} \sim \mathcal{T}_{\tilde{U}, \tilde{V}}} \tilde{f}(v', (g(\tilde{U}, X), \hat{Y})) - \mathbb{E}_{Z' \sim S^\alpha} \tilde{f}(v', (g(\tilde{U}, X'), Y')) \\ &= \mathbb{E}_{\tilde{U}, \tilde{V}} \mathbf{W}_1(\tilde{\mathcal{T}}_{\tilde{U}, \tilde{V}}, \tilde{\mathcal{S}}_U^\alpha) = \mathbb{E}_{U,V,S^\alpha,T_X'} \mathbf{W}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) \end{aligned}$$

Then from Assumption 3, \tilde{f} is σ' -sub-Gaussian for any u, v, v' under $\hat{Z} \sim \mathcal{T}_{u,v}$, and σ' -sub-Gaussian for any u, v' under $Z \sim \mathcal{S}_i, \forall i \in [N]$. Apply Lemma B.2, we have $\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha)$ is $\sqrt{\frac{1}{m_t'} + \sum_{i=1}^N \frac{\alpha_i^2}{m_i}}$ σ' -sub-Gaussian. So we can obtain $-\lambda \mathbb{E}_{U,V,S^\alpha,T_X'} \left(\mathbf{W}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) \right) - \lambda^2 \sigma'^2 \left(\frac{1}{m_t'} + \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right) / 2 \leq I(U, V; S^\alpha, T_X')$.

Consequently, we have:

$$|\mathbb{E}_{U,V,S^\alpha,T_X'} \left(\mathbf{W}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) \right)| \leq \sqrt{2\sigma'^2 \left(\frac{1}{m_t'} + \sum_{i=1}^N \frac{\alpha_i^2}{m_i} \right) I(U, V; S^\alpha, T_X')}$$

Thus,

$$\begin{aligned} gen(\mathcal{T}, \mathcal{A}^{un}) &\leq \mathbb{E}_{U,V,S^\alpha,T_X'} [\mathbf{W}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha) - \hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{U,V}, \tilde{\mathcal{S}}_U^\alpha)] + \mathbb{E}_{U,V} R_{rep}^*(U, V) + R^* \\ &\leq \sqrt{2\sigma'^2 \left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t'} \right) I(U, V; S^\alpha, T_X') + \mathbb{E}_{U,V} R_{rep}^*(U, V) + R^*} \end{aligned}$$

Conclude the proof. \square

E MISSING PROOFS IN SECTION 6

E.1 Proof of Theorem 6.1

Theorem. *Following Theorem 4.1, 5.1 4.2, 5.2, define a weight parameter τ that balances the unsupervised and supervised MDA. Adopt the aforementioned choice of $\tilde{\mathcal{F}}$ with $\sigma' = \sigma$ and the SGLD updates described above. Then we can obtain the gradient norm bound for the expected target risk, $\forall \alpha \in \Delta_N, \tau, \epsilon \in [0, 1]$, we have:*

$$\begin{aligned} \mathbb{E}_{U,V,S^\alpha,T,T'_X} R_{\mathcal{T}}(U,V) &\leq \mathbb{E}_{U,V,S^\alpha,T,T'_X} \hat{R}_{\epsilon,\tau}^\alpha(U,V) + \tau\sigma \sqrt{2\left(\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i}\right)(\delta_u + \delta_v)} \\ &\quad + \tau\epsilon\sigma \sqrt{2\left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right)\delta_u + (1-\tau)\mathbb{E}_{U,V} R_{rep}^*(U,V)} \\ &\quad + (1-\tau)\left(\sigma \sqrt{2\left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right)(\delta_u + \delta_v) + R^*}\right), \end{aligned}$$

$\delta_u \stackrel{\text{def}}{=} \sum_{k=1}^K \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2}{2\sigma_k^2}$ and $\delta_v \stackrel{\text{def}}{=} \sum_{k=1}^K \frac{(\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2}{2\sigma_k^2}$ are the accumulated gradient norm for U and V , respectively.

Proof. Combine Theorem 4.1, 4.2, 5.1 and 5.2 we have:

$$\begin{aligned} \mathbb{E}_{U,V,S^\alpha,T,T'_X} R_{\mathcal{T}}(U,V) &\leq \mathbb{E}_{U,V,S^\alpha,T,T'_X} \hat{R}_{\epsilon,\tau}^\alpha(U,V) + \tau \sqrt{2\epsilon^2\sigma^2\left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right)I(U;S^\alpha,T)} \\ &\quad + \tau \sqrt{2\sigma^2\left(\frac{(1-\epsilon)^2}{m_t} + \epsilon^2 \sum_{i=1}^N \frac{\alpha_i^2}{m_i}\right)I(U,V;S^\alpha,T) + (1-\tau)(R^* + \mathbb{E}_{U,V} R_{rep}^*(U,V))} \\ &\quad + (1-\tau) \sqrt{2\sigma^2\left(\sum_{i=1}^N \frac{\alpha_i^2}{m_i} + \frac{1}{m_t}\right)I(U,V;S^\alpha,T'_X)}, \end{aligned}$$

We use SGLD to optimize the empirical risk:

$$\begin{aligned} \hat{R}_{\epsilon,\tau}^\alpha(u,v) &\stackrel{\text{def}}{=} \tau(1-\epsilon)\hat{R}_{\mathcal{T}}(u,v) + \tau\epsilon\hat{R}_{S^\alpha}(u,v) \\ &\quad + \tau\epsilon\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_u, \tilde{S}_u^\alpha) + (1-\tau)\hat{\mathbf{W}}_1(\tilde{\mathcal{T}}_{u,v}, \tilde{S}_u^\alpha) \end{aligned}$$

Since we jointly minimize the objective function w.r.t u, v and maximize w.r.t v' by definition of the Wasserstein distances, we analyze U, V given a fixed v' updated from previous steps. Given the batch size $|B|$, let the batch estimation of the empirical risks be:

$$\begin{aligned} \hat{R}_{T_B^k}(U,V) &= \frac{1}{|B|} \sum_{j=1}^{|B|} \ell(h(V, g(U, X_j^t)), Y_j^t) \\ \hat{R}_{S_{B_{1:N}}^\alpha}(U,V) &= \sum_{i=1}^N \frac{\alpha_i}{|B|} \sum_{j=1}^{|B|} \ell(h(V, (g(U, X_{i,j}^s))), Y_{i,j}^s) \\ \hat{R}_{T_{X_B}^k}(U,V,v') &= \frac{1}{|B|} \sum_{j=1}^{|B|} \ell(h(v', g(U, X_j^t)), h(V, g(U, X_j^t))) \\ \hat{\mathbf{W}}_1^k(\tilde{\mathcal{T}}_U, \tilde{S}_U^\alpha) &= \hat{R}_{T_B^k}(U,v') - \hat{R}_{S_{B_{1:N}}^\alpha}(U,v') \\ \hat{\mathbf{W}}_1^k(\tilde{\mathcal{T}}_{U,V}, \tilde{S}_U^\alpha) &= \hat{R}_{T_{X_B}^k}(U,V,v') - \hat{R}_{S_{B_{1:N}}^\alpha}(U,v') \end{aligned}$$

For updating U , note the gradients w.r.t three data batches as:

$$G_u(U_{k-1}, V_{k-1}, T_B^k) = \tau(1-\epsilon)\nabla_U \hat{R}_{T_B^k}(U_{k-1}, V_{k-1}) + \tau\epsilon\nabla_U \hat{R}_{T_B^k}(U, v')$$

$$G_u(U_{k-1}, V_{k-1}, T_{X_B}^k) = (1 - \tau) \nabla_U \hat{R}_{T_{X_B}^k}(U_{k-1}, V_{k-1}, v')$$

$$G_u(U_{k-1}, V_{k-1}, S_{B_{1:N}}^k) = \tau \epsilon \nabla_U \hat{R}_{S_{B_{1:N}}^k}^\alpha(U_{k-1}, V_{k-1}) - (1 - \tau + \epsilon \tau) \nabla_U \hat{R}_{S_{B_{1:N}}^k}^\alpha(U_{k-1}, v')$$

Similarly, for updating V , we have:

$$G_v(U_{k-1}, V_{k-1}, T_B^k) = \tau(1 - \epsilon) \nabla_V \hat{R}_{T_B^k}(U_{k-1}, V_{k-1})$$

$$G_v(U_{k-1}, V_{k-1}, T_{X_B}^k) = (1 - \tau) \nabla_V \hat{R}_{T_{X_B}^k}(U_{k-1}, V_{k-1}, v')$$

$$G_v(U_{k-1}, V_{k-1}, S_{B_{1:N}}^k) = \tau \epsilon \nabla_V \hat{R}_{S_{B_{1:N}}^k}^\alpha(U_{k-1}, V_{k-1})$$

Denote the overall gradient updates for U and V respectively as:

$$G_u^k = G_u(U_{k-1}, V_{k-1}, T_B^k) + G_u(U_{k-1}, V_{k-1}, T_{X_B}^k) + G_u(U_{k-1}, V_{k-1}, S_{B_{1:N}}^k)$$

$$G_v^k = G_v(U_{k-1}, V_{k-1}, T_B^k) + G_v(U_{k-1}, V_{k-1}, T_{X_B}^k) + G_v(U_{k-1}, V_{k-1}, S_{B_{1:N}}^k)$$

Then the updates with noise injected:

$$\begin{aligned} U_k &= U_{k-1} - \eta_k^u G_u^k + \xi_k^u \\ V_k &= V_{k-1} - \eta_k^v G_v^k + \xi_k^v, \end{aligned}$$

where $\xi_k^u \sim N(0, \sigma_k^2 I_{d_u})$ and $\xi_k^v \sim N(0, \sigma_k^2 I_{d_v})$.

Define the sequence of representation parameter and the predictor parameter for K iterations as $(U, V)^{[K]} = (U^{[K]}, V^{[K]}) = ((U_1, V_1), (U_2, V_2), \dots, (U_K, V_K))$. The sequence of samplings for each dataset are defined as $T_B^{[K]} = (T_B^1, \dots, T_B^K)$, $S_{B_{1:N}}^{[K]} = (S_{B_{1:N}}^1, \dots, S_{B_{1:N}}^K)$ and $T_{X_B}^{[K]} = (T_{X_B}^1, \dots, T_{X_B}^K)$. The final output of the joint optimization algorithm $(U, V) = f((U, V)^{[K]})$.

$$\begin{array}{c} T \rightarrow T_B^{[K]} \\ \downarrow \\ S_{1:N} \rightarrow S_{B_{1:N}}^{[K]} \rightarrow (U, V)^{[K]} \rightarrow (U, V) \\ \uparrow \\ T'_X \rightarrow T_{X_B}^{[K]} \end{array}$$

Given specific α , let us apply the data processing inequality and the chain rule, then we have

$$I(U, V; S^\alpha, T) = I(U, V; S_{1:N}, T) \leq I((U^{[K]}, V^{[K]}); S_{B_{1:N}}^{[K]}, T_B^{[K]}) = \sum_{k=1}^K I(U_k, V_k; S_{B_{1:N}}^{[K]}, T_B^{[K]} | U^{[k-1]}, V^{[k-1]})$$

Moreover, since the sampling strategy is agnostic to the previous iterates of the parameters and previous samplings, so we can obtain:

$$\begin{aligned} I(U_k, V_k; S_{B_{1:N}}^{[K]}, T_B^{[K]} | U^{[k-1]}, V^{[k-1]}) &= I(U_k, V_k; S_{B_{1:N}}^k, T_B^k | U^{k-1}, V^{k-1}) \\ &= H(U_k, V_k | U_{k-1}, V_{k-1}) - H(U_k, V_k | U_{k-1}, V_{k-1}, S_{B_{1:N}}^k, T_B^k) \end{aligned}$$

For any $(U_{k-1}, V_{k-1}) = (u_{k-1}, v_{k-1})$, since the joint entropy of a set of variables is less than or equal to the sum of the individual entropies of the variables in the set, we have

$$H(U_k, V_k | U_{k-1} = u_{k-1}, V_{k-1} = v_{k-1}) \leq H(-\eta_k^u G_u^k + \xi_k^u) + H(-\eta_k^v G_v^k + \xi_k^v)$$

Apply Lemma B.3, we can obtain

$$\begin{aligned}
 H(U_k, V_k | U_{k-1} = u_{k-1}, V_{k-1} = v_{k-1}, S_{B_{1:N}}^k, T_B^k) & \\
 &= H(-\eta_k^u G_u(u_{k-1}, v_{k-1}, T_{X_B}^k) + \xi_u^k, -\eta_k^v G_v(u_{k-1}, v_{k-1}, T_{X_B}^k) + \xi_v^k) \\
 &\geq H(-\eta_k^u G_u(u_{k-1}, v_{k-1}, T_{X_B}^k) - \eta_k^v G_v(u_{k-1}, v_{k-1}, T_{X_B}^k) + \xi_u^k + \xi_v^k) \\
 &\geq H(\xi_u^k) + H(\xi_v^k) = \frac{d_u}{2} \log(2\pi e \sigma_k^2) + \frac{d_v}{2} \log(2\pi e \sigma_k^2).
 \end{aligned}$$

The first inequality is obtained because the entropy of a function of random variables is smaller than the entropy of the random variables (See proof in Lemma B.3). The second inequality comes from the independence between the gradients and the injected noise.

Moreover, we can obtain $\mathbb{E}(\|-\eta_k^u G_u^k + \xi_u^k\|_2^2) = \mathbb{E}(\|\eta_k^u G_u^k\|_2^2 + \|\xi_u^k\|_2^2) = (\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2 + d_u \sigma_k^2$ and $\mathbb{E}(\|-\eta_k^v G_v^k + \xi_v^k\|_2^2) = \mathbb{E}(\|\eta_k^v G_v^k\|_2^2 + \|\xi_v^k\|_2^2) = (\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2 + d_v \sigma_k^2$ with the independence. Since Gaussian distribution has the largest entropy among the variables with the same second order moment, so we can further get:

$$\begin{aligned}
 H(U_k, V_k | U_{k-1} = u_{k-1}, V_{k-1} = v_{k-1}) - H(U_k, V_k | U_{k-1} = u_{k-1}, V_{k-1} = v_{k-1}, S_{B_{1:N}}^K, T_B^k) & \\
 \leq H(-\eta_k^u G_u^k + \xi_u^k) + H(-\eta_k^v G_v^k + \xi_v^k) - (H(\xi_u^k) + H(\xi_v^k)) & \\
 \leq \frac{d_u}{2} \log(2\pi e \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2 + d_u \sigma_k^2}{d_u}) + \frac{d_v}{2} \log(2\pi e \frac{(\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2 + d_v \sigma_k^2}{d_v}) & \\
 \quad - \frac{d_u}{2} \log(2\pi e \sigma_k^2) - \frac{d_v}{2} \log(2\pi e \sigma_k^2) & \\
 \leq \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2 + (\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2}{2\sigma_k^2} &
 \end{aligned}$$

The above bound holds for all u_{k-1}, v_{k-1} , thus we can integrate the bound to conclude that

$$H(U_k, V_k | U_{k-1}, V_{k-1}) - H(U_k, V_k | U_{k-1}, V_{k-1}, S_{B_{1:N}}^K, T_B^k) \leq \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2 + (\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2}{2\sigma_k^2},$$

So we get $I(U, V; S^\alpha, T) \leq \sum_{k=1}^K \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2 + (\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2}{2\sigma_k^2} = \delta_u + \delta_v$

Similarly, we can obtain the following results with the same proof process:

$$\begin{aligned}
 I(U; S^\alpha, T) &\leq \sum_{k=1}^K \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2}{2\sigma_k^2} = \delta_u \\
 I(U, V; S^\alpha, T'_X) &\leq \sum_{k=1}^K \frac{(\eta_u^k)^2 \mathbb{E}\|G_u^k\|_2^2 + (\eta_v^k)^2 \mathbb{E}\|G_v^k\|_2^2}{2\sigma_k^2} = \delta_u + \delta_v
 \end{aligned}$$

Place the above terms into the inequality presented at the beginning of this section, we can conclude the proof. \square

F EXPERIMENTAL DETAILS

We provide the experimental details and some additional results in this section. The pseudo code ³ for the proposed IMDA algorithm is presented in Sec. H.

F.1 Amazon Review

As presented in the main paper, the original dataset is pre-processed to 5000-dimension bag-of words features following Chen et al. (2012). And the target shift data is created by randomly dropping 50% negative reviews.

Model Structure *Representation learner*: [5000, 1000, 500, 100] MLP net using 0.7 dropout rate with Relu activation added after each hidden layer and finally output a 100-dimension feature representation.

Predictor and duplicate predictor: [100, 2] linear transformation followed by a log softmax layer transforming the 100-dimension feature to 2-class log probabilities.

Loss function: we choose the "negative log-likelihood loss" as the loss function.

Computing Resources The experiments were run on a server with 6 CPUs and 1 GPU of 32GB memory.

Experimental Setting The Amazon review dataset contains 6465 samples for "books", 5586 samples for "dvd", 7681 samples for "electronics" and 7945 samples for "kitchen". We randomly sample 2000 examples for each domain. In unsupervised MDA, the 2000 target samples without labels are applied in the training phase, and the rest samples are used as the test set. For supervised MDA with few target labels, we randomly sample 10% examples from the 2000 target samples using as the train data, then the rest 90% are used as the test set.

We provide the main hyper-parameters for supervised and unsupervised settings, which are chosen based on cross-validation. Other hyper-parameters are set as the default value provided in the code.

In the supervised MDA, we set l2_scale= 0.5 (the C_1 constant in the convex optimization objective for α), and the learning rate of a Adadelta optimizer is $\eta = 0.5$. The penalty for the \mathbf{W}_1 distance is set to W1_sup_coef= 0.01. Finally, the network is trained for 40 epochs with a mini-batch size of 20.

In the unsupervised MDA, as presented in the main paper, the empirical risk of the pseudo target distribution is: $\hat{R}_{\mathcal{T}_{u,v}}(u, v, v') = \frac{1}{m_t} \sum_{j=1}^{m_t} \ell(h(v', g(u, X_j^t)), h(v, g(u, X_j^t)))$. There does not exist a direct realization of the above loss. So we implement it with $\ell(h(v', g(U, X_j^t)), \hat{Y}_j^t) + \ell(h(V, g(U, X_j^t)), \hat{Y}_j^{t'})$, where the \hat{Y}_j^t and $\hat{Y}_j^{t'}$ are the predicted label of corresponding predictor. Then we apply two different penalty weights W1_discri_coef1 and W1_discri_coef2 for the two losses respectively. We set l2_scale= 1, learning rate $\eta = 0.8$, W1_discri_coef1= 0.06, $C_0=W1_discri_coef2= 1.2$. Finally, the network is trained for 50 epochs and the mini-batch size is 20.

For reproducing the results of WADN (Shui et al., 2021) and DARN (Wen et al., 2020), we used the code repos <https://github.com/cjshui/WADN> and <https://github.com/junfengwen/DARN>. The hyper-parameters are set as provided in the respective paper.

F.2 Digits

Model Structure *Representation learner* is consisted of 3 stacked modules : 'conv':[3, 3, 64], 'relu', 'maxpooling':[2,2,0]; 'conv':[3, 3, 128], 'relu', 'maxpooling':[2,2,0]; 'conv':[3, 3, 256], 'relu', 'maxpooling':[2,2,0].

Predictor and duplicate predictor: [2304, 512, 100, 10] MLP net using "Relu" activation after each hidden layer is followed by a log softmax layer. The predictors transform the 2304-dimension feature to 10-class log probabilities. There is no dropout.

Loss function: we choose the "negative log-likelihood loss" as the loss function.

Computing Resources The experiments were run on a server with 6 CPUs and 1 GPU of 32GB memory.

³The code is released at <https://github.com/livreQ/IMDA>

Experimental Setting Digits dataset has four domains: "MNIST," "USPS," "SVHN," and "SYNTH." "MNIST" and "USPS" are handwriting datasets that are very close. "SYNTH" is transformed from "SVHN," so they are also more similar to each other. Each domain has different train-test split when downloaded. The respective training sample sizes are 60000, 7219, 73257, 479400, and the respective test sample sizes are 10000, 2017, 26032, 9553. Following the same procedure of Shui et al. (2021), we randomly select 7000 samples for each domain, then create a target shifted distribution for each source, making each source contain 5300 samples, and the target domain contains 7000 examples.

In unsupervised MDA, the 7000 target samples without labels are applied in the training phase, and the test set is an unseen portion that contains 2000 samples. For supervised MDA with few target labels, we randomly sample 10% examples from the 7000 target samples using as the train data, then the rest 90% are used as the test set.

We provide the main hyper-parameters for supervised and unsupervised settings, which are chosen based on cross-validation. Other hyper-parameters are set as the default value provided in the code.

In the supervised MDA, we set $l2_scale=1.5, 1$, and the learning rate of a Adadelta optimizer is $\eta = 0.5, 0.2$. The penalty for the \mathbf{W}_1 distance is set to $W1_sup_coef=0.01, 0.03$. Finally, the network is trained for 70 epochs with a mini-batch size of 128. In the unsupervised MDA, we set $l2_scale=1.5$, learning rate $\eta = 0.4$, $W1_discri_coef1=0.002$, $C_0=W1_discri_coef2=0.9$. Finally, the network is trained for 15 epochs and the mini-batch size is 128.

F.3 Gradient Penalty

We implemented two types of gradient penalties in the code. The first is to ensure the Lipschitzness of \tilde{f} for the Wasserstein distances. We adopt the method proposed by Gulrajani et al. (2017). For example, with $(\tilde{X}, Y) \sim \mathcal{T}_u$ and $(\tilde{X}', Y') \sim \mathcal{S}_u^\alpha$, we generate the interpolated feature $\tilde{X}_{int} = \lambda\tilde{X} + (1 - \lambda)\tilde{X}'$, $\lambda \sim \text{Unif}[0, 1]$. The gradient penalty is implemented by adding the regularization term $\|\nabla_{\tilde{X}_{int}} h(v, \tilde{X}_{int})\|_2^2$. The second gradient penalty comes from the bound in Theorem 6.1. We consider the gradient norm in each batch as a regularization term. Note this gradient is w.r.t the model parameters, not the feature or input.

G ADDITIONAL RESULTS

G.1 Amazon Review

Now we present some additional results. At first, we illustrate the ablation study w.r.t different drop rates on Amazon dataset for each domain in unsupervised MDA. In Fig. 2, we can see the proposed IMDA algorithm consistently outperforms the previous state-of-art on every single domain.

We further provide additional t_SNE visualization of amazon data in Fig. 3, where an illustration of target shift with drop rate of 50% is also included. From the figure, we observe that the geometry information of the data is retained using the Wasserstein distance, and the distribution shift is decreased after adaptation.

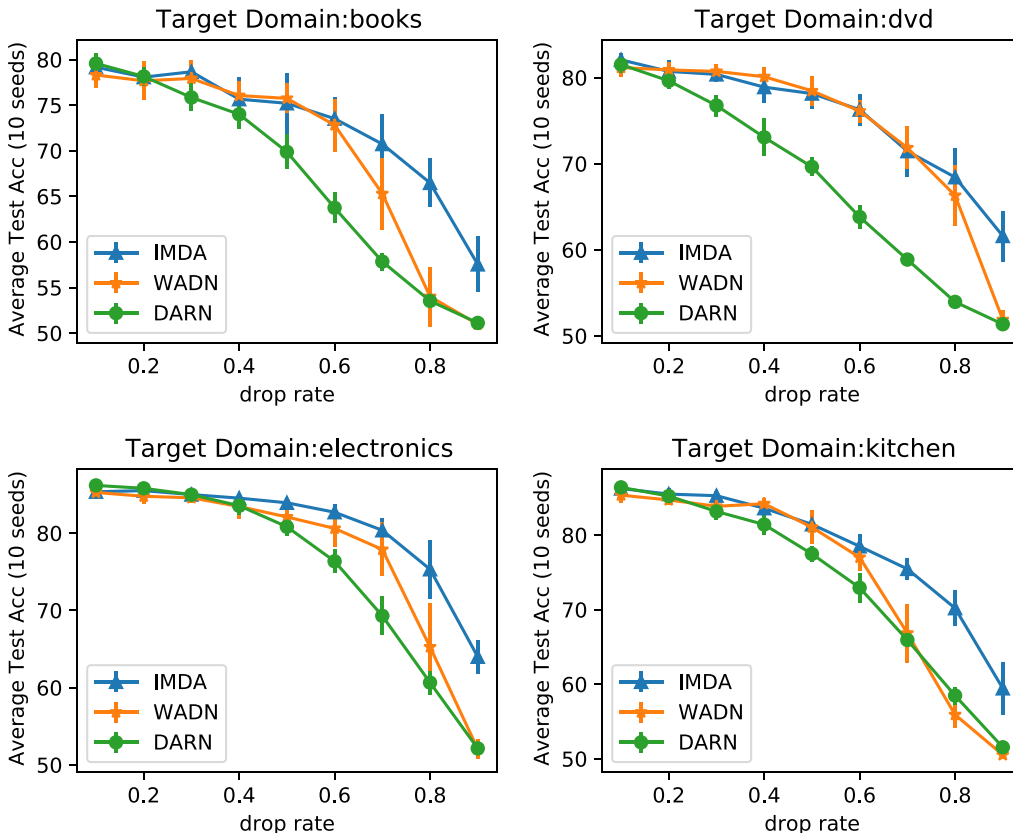


Figure 2: Test Accuracy for Each Domain w.r.t Different Levels of Target Shift on Amazon Review Dataset for Unsupervised MDA

G.2 Digits

The Amazon review dataset’s four domains are homogeneous tasks with similar difficulty. Otherwise, the four digits domains are heterogeneous, where the two handwriting digit recognition tasks are much easier than the two home number recognition tasks. To learn more reliable relations, we set uniform weight at the beginning of the training process (5 epochs) to make the complex tasks sufficiently learned, avoiding an unfair low weight due to the poor results.

From the dataset description, we know MNIST and USPS are handwriting digits and are more similar to each other than other datasets. Moreover, SVHN and SYNTH are more correlated since SYNTH is obtained by adding transformation on SVHN. We observe the apparent similarity in Fig. 4 (a). The algorithm gives high weight to USPS when learning MNIST and high weight to SVHN when learning SYNTH, and vice versa. Fig. 4 illustrates the minimax optimization process for the Wasserstein distance. In unsupervised MDA, we do not directly optimize the combined empirical risk $\hat{R}_{S^\alpha}(u, v)$, the evolution of its estimation continuously decreases, which shows the effectiveness of minimizing $\mathbf{W}_1(\tilde{\mathcal{T}}_{u,v}, \tilde{\mathcal{S}}_u)$.


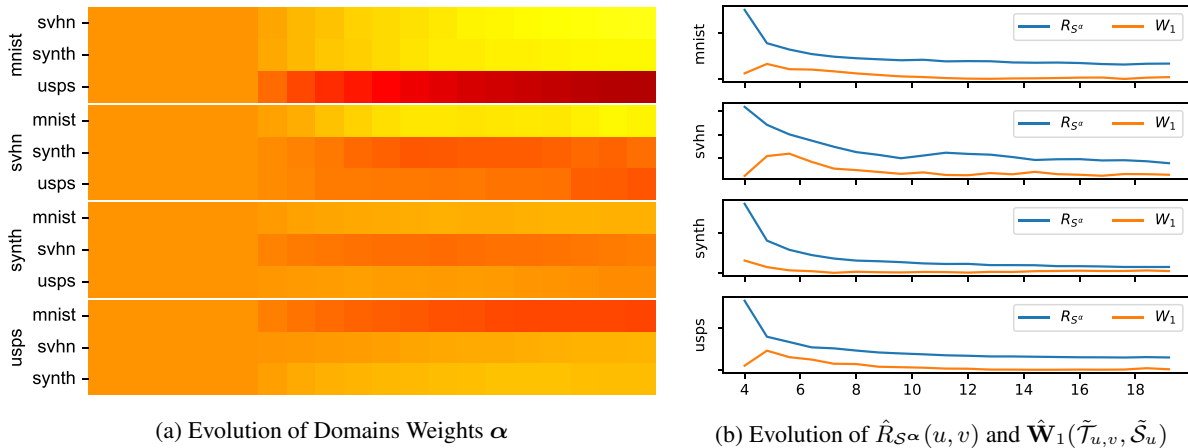


Figure 3: T_SNE & Target Shift on Amazon Review Dataset

We also test the single-source DA method DeepJDOT (Damodaran et al., 2018) (merging the sources as one) and MOST (Nguyen et al., 2021) on the target shifted digits data. The result for the target domain MNIST is illustrated in Tab. 3. We do not consider adding a more detailed comparison of these two methods for the following reasons. First, we have already compared other similar single-source DA approaches that using the merged source data. Second, the original MOST algorithm consist of a mixture of several optimization objectives, but the theoretical result is only related to one objective. Therefore, adjusting hyper-parameters for these combinations is meaningless for this paper’s theoretical nature.

Table 3: Accuracy(%) on Digits Dataset with a Drop Rate of 50%

Method	target: MNIST
MOST	88.23
DeepJDOT	87.5
IMDA	89.26



(a) Evolution of Domains Weights α

(b) Evolution of $\hat{R}_S^\alpha(u, v)$ and $\hat{W}_1(\tilde{\mathcal{T}}_{u,v}, \tilde{\mathcal{S}}_u)$

Figure 4: Visualization of the Evolution of Different Terms w.r.t Training Epochs on Digits Dataset for Unsupervised MDA with a Drop Rate of 0.5

H ALGORITHM

Algorithm 1: Information-Theoretic Multi-Source Domain Adaptation (IMDA)

Input: Labeled source samples $S_{1:N}$, few labeled target samples T , unlabeled target samples T'_X , weights ϵ, τ , learning rate η , variance σ , batch size B , constants C_0, C_1 , moving average weight $0 < C < 1$;

Output: Representation parameter U , predictor parameter V , duplicate predictor parameter v' , domain weights α ;

for $t \leftarrow 1$ **to** T **do**

if $t == 1$ **then**

 Initialize domain weights $\alpha = \{1/N, \dots, 1/N\}$, randomly initialize U_0, V_0, v'_0 ;
 $\delta_u = 0, \delta_v = 0$;

else

 Initialize U_0, V_0, v'_0 with last epoch outputs;

end

for $k \leftarrow 1$ **to** K **do**

 Sample batch data $T_B^k, S_{B_{1:N}}^k, T_{X_B}^k$ from the corresponding datasets;

 Evaluate gradients:

$$\begin{aligned} G_u^k &= \tau(1 - \epsilon)\nabla_U \hat{R}_{T_B^k}(U_{k-1}, V_{k-1}) \\ &\quad + \tau\epsilon\nabla_U \hat{R}_{T_B^k}(U, v'_{k-1}) + (1 - \tau)\nabla_U \hat{R}_{T_{X_B}^k}(U_{k-1}, V_{k-1}, v'_{k-1}) \\ &\quad + \tau\epsilon\nabla_U \hat{R}_{S_{B_{1:N}}^k}(U_{k-1}, V_{k-1}) - (1 - \tau + \epsilon\tau)\nabla_U \hat{R}_{S_{B_{1:N}}^k}(U_{k-1}, v'_{k-1}) \\ G_v^k &= \tau(1 - \epsilon)\nabla_V \hat{R}_{T_B^k}(U_{k-1}, V_{k-1}) \\ &\quad + (1 - \tau)\nabla_V \hat{R}_{T_{X_B}^k}(U_{k-1}, V_{k-1}, v'_{k-1}) + \tau\epsilon\nabla_V \hat{R}_{S_{B_{1:N}}^k}(U_{k-1}, V_{k-1}) \\ G_{v'}^k &= \tau\epsilon[\nabla_{v'} \hat{R}_{T_B^k}(U, v'_{k-1}) - \nabla_{v'} \hat{R}_{S_{B_{1:N}}^k}(U, v'_{k-1})] \\ &\quad + (1 - \tau)[\nabla_{v'} \hat{R}_{T_{X_B}^k}(U, V, v'_{k-1}) - \nabla_{v'} \hat{R}_{S_{B_{1:N}}^k}(U, v'_{k-1})] \end{aligned}$$

 Update:

$$U_k = U_{k-1} - \eta G_u^k + \xi_k^u, V_k = V_{k-1} - \eta G_v^k + \xi_k^v, v'_k = v'_{k-1} + \eta G_{v'}^k;$$

$$\delta_u = \delta_u + \frac{\eta^2 \mathbb{E} \|G_u^k\|_2^2}{2\sigma^2}, \delta_v = \delta_v + \frac{\eta^2 \mathbb{E} \|G_v^k\|_2^2}{2\sigma^2};$$

end

$U = U_K, V = V_K, v' = v'_K$;

 Solve

$$\begin{aligned} \alpha' &= \min_{\alpha} \left((\epsilon\tau + C_0(1 - \tau)) \hat{R}_{S_{B_{1:N}}^k}(U, V) - (\epsilon\tau + 1 - \tau) \hat{R}_{S_{B_{1:N}}^k}(U, v') \right. \\ &\quad \left. + C_1((1 - \tau + \tau\epsilon)\sqrt{\delta_u + \delta_v} + \tau\epsilon\sqrt{\delta_u})R(\alpha) \right), \end{aligned}$$

$$R(\alpha) = \sqrt{\sum_{i=1}^N \frac{\alpha_i^2}{m_i}}, \text{ s.t. } \forall i \in [N], \alpha_i \geq 0, \sum_{i=1}^N \alpha_i = 1,$$

 Update $\alpha = C\alpha + (1 - C)\alpha'$;

end
