

MOT: Masked Optimal Transport for Partial Domain Adaptation

You-Wei Luo Chuan-Xian Ren*

School of Mathematics, Sun Yat-Sen University, China

luoyw28@mail2.sysu.edu.cn, rchuanx@mail.sysu.edu.cn

Abstract

As an important methodology to measure distribution discrepancy, optimal transport (OT) has been successfully applied to learn generalizable visual models under changing environments. However, there are still limitations, including strict prior assumption and implicit alignment, for current OT modeling in challenging real-world scenarios like partial domain adaptation, where the learned transport plan may be biased and negative transfer is inevitable. Thus, it is necessary to explore a more feasible OT methodology for real-world applications. In this work, we focus on the rigorous OT modeling for conditional distribution matching and label shift correction. A novel masked OT (MOT) methodology on conditional distributions is proposed by defining a mask operation with label information. Further, a relaxed and reweighting formulation is proposed to improve the robustness of OT in extreme scenarios. We prove the theoretical equivalence between conditional OT and MOT, which implies the well-defined MOT serves as a computation-friendly proxy. Extensive experiments validate the effectiveness of theoretical results and proposed model.

1. Introduction

For real-world visual data in unconstrained environments, detecting the shift of data distributions and measuring the statistical discrepancy are essential problems for learning generalizable model [29, 30, 51]. Considerable efforts have been made to build well-defined metrics on distributions, where the optimal transport (OT) based distances have shown significant advantages in sample-wise correlation characterization and geometric interpretability [6, 7, 9, 11, 12, 24, 27]. Under the guarantees of appealing theoretical properties, OT-based modules [8, 10, 11, 16, 20, 24] are extensively explored to learn the transferrable models. These advanced models have also been successfully applied to the computer vision and pattern recognition tasks in

changing environments, e.g., object classification [31, 39], semantic segmentation [52] and medical images [42].

A typical scenario for learning under distribution shift is known as domain adaptation (DA), where the model trained on the labeled source domain P is expected to be transferrable to unlabeled target domain with different data distribution Q . In DA, source domain and target domain share the same label space \mathcal{Y} , then inspired by distribution adaptation theory [1, 48], the OT-based models usually focus on learning invariant properties, e.g., marginal invariant representations [8, 11, 28], adversarial invariant representations [16], conditional invariant representations [19, 24, 38] and joint invariant representations [10, 20].

However, in real-world scenarios, the label space usually changes with the shifting distributions, which induces the partial DA (PDA) problems [2, 3, 16, 21, 25]. Recent theoretical results [38, 49] also imply that the shift on label distributions [47] is non-negligible in application, and the label distribution correction is necessary to achieve a sufficiently small joint risk across domains. Therefore, if label distributions P_Y and Q_Y are different, a correction (e.g., reweighted and relaxation) on source P_Y is usually necessary [22, 47]. Since there are strict constraints on the marginal distributions, classical OT models will learn incorrect sample correlation and further induce the negative transfer problem. Typically, to address these issues, there are two types of variants for OT. 1) Reweighted OT [16, 32, 33] introduces a reweighted operation, which aims to detect the outlier classes (i.e., unseen classes in target domain) and decrease their masses in P_Y . Then the OT assignment, which is constrained by reweighted source distribution, will only transport the information that is shared by both domains. However, the effectiveness of these models directly depends on the reweighted function w , where models could be error-prone when weight w is inaccurate. 2) Unbalanced OT (UOT) [5, 11] introduces relaxation to the strict marginal constraints by penalizing the assignments that don't meet the constraints. Such a relaxation allows the assignments to focus on the correct correlation with low transport cost and ignore the incorrect transportation to outlier classes with higher cost. But, since the relaxation is

*Corresponding Author.

applied to original distributions, the penalty will be unaffordable and the relaxation will fail when cross-domain distributions are significantly different.

Besides, as a sufficient condition to achieve successful DA, the conditional shift correction has received increasing attention recently [14, 19, 23, 38]. Several pioneering works [19, 24, 32] on the OT modeling between conditional distributions have shown great potential in mitigating negative transfer and reducing generalization risk. But, there are still several limitations, i.e., strict assumption on kernel Gaussian prior [24], implicit conditional alignment via reweighted source [32] and implicit proxy with small intra-class discrepancy assumption [19]. Recently, OT with mask is proposed to integrating label information into OT. Zhang et al. [46] first define the mask on transport plan. Gu et al. [15] further establish the mask theory for Gromov-Wasserstein. However, theoretical understanding on the relation between mask mechanism and conditional OT is unexplored. Thus, it is necessary to explore a general and explicit formulation of OT for sufficient conditional distribution matching, and develop computation-friendly algorithms for application.

Generally, to alleviate the difficulties in current OT methodology for label distribution correction and conditional distribution matching, known as generalized label shift (GLS) correction, we are interested in two major problems: 1) rigorous modeling and explicit algorithm for OT between conditional distributions; 2) unbiased and relaxed transport plan learning for extreme generalization scenario, e.g., PDA. In this paper, we propose a novel masked OT (MOT) methodology by introducing a relaxed and reweighted OT formulation and the conditional mask mechanism. Further, we derive an efficient fixed-point method for learning OT assignment, and propose an OT-based invariant risk model to deal with the extreme PDA problem. Our contributions can be summarized as follows.

- The theoretical connection between MOT and conditional OT is proved for commonly used OT formulations. The main result ensures the proposed mask mechanism and MOT model are sufficient to characterize the label-conditioned sample correspondence and measure the conditional discrepancy explicitly.
- A relaxation and reweighting based OT formulation is proposed, which overcomes the sensitivity of current OT modelings to weight estimation and intrinsic label discrepancy. Intuitive analysis is presented to show the advantages of MOT methodology.
- With theoretical guarantees, a computation-friendly empirical estimation with explicit fixed-point algorithm is proposed as proxy for conditional OT. Then an equivalent risk model is proposed to learn minimized risk on both source and transported source domains. The application to PDA is extensively evaluated, which validates the ef-

fectiveness of theoretical results and proposed model.

2. Methodology: Masked OT

Notation. Let X and Y be the covariate and label variable, which take their values from \mathcal{X} and \mathcal{Y} . P and Q denote the distributions of source and target domains, where lowercase letters p and q denote the probability density functions (PDFs) and subscripts represent the corresponding random variables, e.g., P_Y implies the label distribution on source. Given a positive convex function ϕ with $\phi(1) = 0$, the ϕ -divergence is defined as $D_\phi(P||Q) \triangleq \mathbb{E}_Q[\phi(\frac{dP}{dQ})]$. Especially, when $\phi(t) = t \ln t$, $D_\phi(P||Q)$ boils down to the well-known KL divergence. For a learning model $f : \mathcal{X} \rightarrow \mathcal{Y}$, given a loss function $\ell : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$, we denote the risk of f on P_{XY} as $\varepsilon_P(f) = \mathbb{E}_{P_{XY}}[\ell(f(X), Y)]$.

2.1. Preliminary

GLS Correction and PDA. The label distributions (e.g., class proportions) in real-world scenarios usually induces the increasing discrepancy between marginal distributions P_Y and Q_Y [22, 47]. Such a phenomenon is formally known as label/target shift [22, 47]. Note PDA can be taken as an extreme scenario of label shift, where the proportions of outlier classes will shift to 0 in target distribution Q , i.e., $\text{supp}(q_Y) \subset \text{supp}(p_Y)$. Thus, many PDA methods are developed under the shift correction framework, e.g., reweighted moment matching [34, 45], reweighted OT [16, 33], relaxation OT [11, 28] and reweighted adversarial adaptation [2, 3]. One the other hands, to achieve better generalization, models are usually required to extract invariant discriminant information and intrinsic pattern across domains, which can be theoretically characterized as the matching between conditional distributions $P_{Z|Y}$ and $Q_{Z|Y}$ [14, 18]. These models aim to develop discrepancy-based objectives for learning conditional invariant representation Z , e.g., conditional adversarial loss [23], conditional Bures metric [24], conditional moment matching [40, 50] and local structure transfer [25, 41].

Recently, by merging label shift and conditional shift, GLS is extensively studied to explore the transferability and discriminability of DA models [19, 35, 38, 47]. Theoretically, shifting distributions lead to the biased risk estimation (i.e., $\varepsilon_P(f) \neq \varepsilon_Q(f)$) and cluster misalignment (i.e., $P_{Z|Y} \neq Q_{Z|Y}$), which explicitly degrades generalization performance. Typically, GLS can be mitigated by introducing the reconstructed source P_{ZY}^w parameterized by importance weight w and conditional invariant representation Z . Ideally, the model trained on weighted source P_{ZY}^w is sufficient for small generalization error and successful transfer [38], which implies GLS-based model can be effective in dealing with complex dataset shift scenarios, e.g., PDA.

UOT. Recently, OT-based models have been successfully applied to transfer learning problems [8, 10, 11, 20, 28, 33].

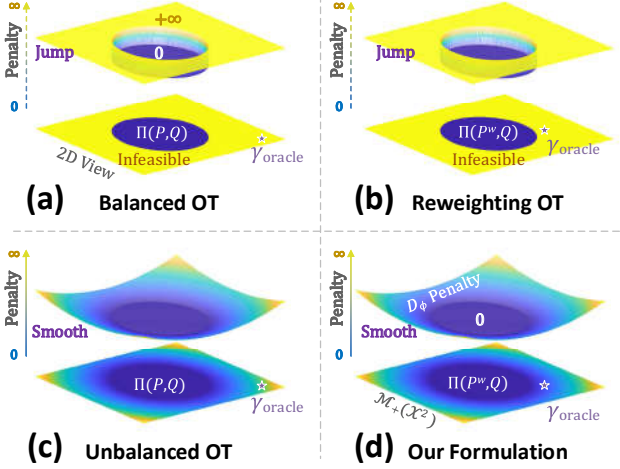


Figure 1. Illustration of different OT formulations. (a)-(b): For OTs with constraints on marginal distributions, though weight w makes ideal solution γ_{oracle} closer to the feasible region Π , γ_{oracle} is still unachievable ($+\infty$ penalty). (c): UOT with relaxation replaces the constraints with smooth penalty, but γ_{oracle} is still far away from the center, which leads to large penalty. (d): Our formulation Eq. (3) with relaxation D_ϕ and reweighted P^w ensures that γ_{oracle} belongs to the feasible region with smaller penalty.

In classical OT, by introducing the entropy regularization $H(\gamma) = \mathbb{E}_\gamma[-\ln(d\gamma)]$, the Sinkhorn divergence [9] between distributions P and Q can be formulated as

$$S^\lambda(P, Q, c) = \min_{\gamma \in \Pi(P, Q)} \int c d\gamma - \lambda H(\gamma), \quad (1)$$

where $\Pi(P, Q)$ is the set of probabilistic couplings over P and Q , λ is parameter of sparsity penalty, c is the cost function for transport, e.g., c is defined on \mathcal{X}^2 for marginal OT [8, 20] and $(\mathcal{X} \times \mathcal{Y})^2$ for joint OT [10, 44].

Further, to deal with more general scenarios, e.g., unequal total mass and label shift, UOT [6, 7, 11] is proposed as a relaxation of classical OT Eq. (1). Specifically, by replacing the strict marginal constraint $\gamma \in \Pi(P, Q)$ with relaxed penalty terms on γ , UOT can be formulated as

$$S^{\lambda, \beta}(P, Q, c) = \min_{\gamma \in \mathcal{M}_+} \int c d\gamma - \lambda H(\gamma) + \beta [D_\phi(\gamma_P \| P) + D_\phi(\gamma_Q \| Q)], \quad (2)$$

where γ_P and γ_Q are the marginals of γ , β is the parameter for marginal penalty, \mathcal{M}_+ is the space of distributions, e.g., $\mathcal{M}_+(\mathcal{X}^2)$ over space \mathcal{X}^2 for marginal UOT [6, 11].

2.2. New Formulation and Analysis

Motivation. Though the reweighted strategy [16, 32, 38, 47] and UOT methodology [11] have been explored for knowledge transfer, there are still potential weaknesses that

we are interested in: 1) though the marginal alignment between P^w and Q mitigates label shift, e.g., MMD [47] and OT [16, 32, 33] on weighted source, it ignores the conditional shift which usually degrades the discriminability; 2) the performance of adaptation with P^w significantly depends on the precision of estimation w as Fig. 1 (b), which may degrade the robustness of shift correction model; 3) the marginal penalty with P in UOT [11] is sensitive to the degree of label shift as Fig. 1 (c), especially for the extreme PDA scenario with large label discrepancy.

To address the challenges above, we first introduce a novel OT with reweighted and relaxation, which reduces the uncertainty of estimation w and alleviates the large penalty on original distribution P simultaneously. Then we propose the *conditional UOT* formulation, which characterizes label-wise correlations and mitigates conditional shift.

Relaxed and Reweighted Formulation. To overcome the bias of risk estimation ε_P with label proportion P_Y on original source domain, the reweighted strategies for modifying distribution are widely applied [19, 35, 38, 47]. To unify the different formulations, we first present a rigorous definition of w -reweighted source as follows.

Definition 1 Given a reweighted function $w : \mathcal{Y} \rightarrow \mathbb{R}_+$ such that $w \cdot p_Y$ is also a PDF on \mathcal{Y} .

- (a) The w -reweighted source P^w is defined as $p_Y^w = w \cdot p_Y$, $p_{X|Y}^w = p_{X|Y}$ and $p_X^w = \int_{\mathcal{Y}} p_{Y=y}^w \cdot p_{X|Y=y} dy$.
(b) The optimal weight w^* satisfies $P_Y^{w^*} = Q_Y$.

An important property of reweighted is that if $P_{X|Y} = Q_{X|Y}$, then dataset shift can be addressed on the w -reweighted source, i.e., $P_{XY}^{w^*} = P_{X|Y} P_Y^{w^*} = Q_{X|Y} Q_Y = Q_{XY}$ and $\varepsilon_{P^{w^*}} = \varepsilon_Q$. Based on the reweighted source, we will next show that the biased sample-wise transports in current OT can be mitigated by the relaxation and reweighted formulation.

As discussed before, though UOT can deal with the label shift scenarios by relaxing the constraints on marginal distribution, the OT model is still required to learn coupling γ with marginal γ_P close to P . It implies such a model may fail in some extreme scenarios where the ideal marginal are actually away from the original source domain P , e.g., γ should not assign values to the outlier classes in P . Therefore, it is necessary to reformulate the original UOT by modifying the marginal penalty on P as P^w :

$$S^{\lambda, \beta}(P^w, Q, c) = \min_{\gamma \in \mathcal{M}_+(\mathcal{X}^2)} \int c d\gamma - \lambda H(\gamma) + \beta [D_\phi(\gamma_{P^w} \| P^w) + D_\phi(\gamma_Q \| Q)]. \quad (3)$$

An intuitive illustration of Eq. (3) and related OT formulations are shown in Fig. 1, where the ideal γ_{oracle} such that its marginals are Q . For original (i.e., balanced) OT Eq. (1) in Fig 1 (a)-(b), the feasible regions are probabilistic

coupling set Π , which implies ideal solutions are always unachievable unless w is the optimal w^* . Though the marginal constraints are relaxed, and feasible regions are extended to \mathcal{M}_+ in UOT Eq. (2), γ_{oracle} may be still hard to learn since the penalty may be large, i.e., Fig 1 (c). Specifically, since the center with 0 penalty value is $\Pi(P, Q)$, the large discrepancy between P and Q makes γ_{oracle} far away from the low penalty region. For our formulation Eq. (3) in Fig 1 (d), the model inherits the advantages of UOT and reweighted OT. The reweighted source P^w further ensures that γ_{oracle} is feasible, and the penalty value is smaller than UOT. Besides, the relaxation on P^w reduces the risk of estimation w , then the neighbours of Π are also feasible. Thus, the ideal solution is still achievable even if $w \neq w^*$, which overcomes the limitation of reweighted OT. In conclusion, new formulation Eq. (3) ensures the existence of unbiased γ_{oracle} for target domain Q and the feasibility of learning.

Conditional OT. Though the reweighted and relaxation formulation Eq. (3) ensures the possibility of learning unbiased γ for target domain Q , the conditional shift correction for intrinsic structure transfer is still not guaranteed. Therefore, to further achieve the conditional invariant property, we now define conditional transport mechanism to learn label-wise correlations and mitigate the negative transports between inter-class sample pairs.

Definition 2 (Conditional OT) For any non-negative coefficient function $\alpha(\cdot)$ on \mathcal{Y} such that $\text{supp}(\alpha) = \text{supp}(p_Y) \cap \text{supp}(q_Y)$ and $\sum_{y \in \mathcal{Y}} \alpha(y) = 1$, the conditional UOT between P^w and Q is defined as

$$\text{OT}_{\text{cond}}^\alpha(P^w, Q, c) = \sum_{y \in \mathcal{Y}} \alpha(y) \text{OT}(P_{X|y}^w, Q_{X|y}, c). \quad (4)$$

Note that the general formulation Eq. (4) is suitable for different OT-based models, e.g., Kantorovich OT, Sinkhorn OT S^λ and UOT $S^{\lambda, \beta}$. An intuitive explanation for conditional OT is that it can be taken as the sliced OT between conditional distributions $P_{X|y}$ and $Q_{X|y}$ over all label conditions $y \in \mathcal{Y}$. The coefficient α ensures that the conditional transports are only carried out on shared classes, i.e., $\text{supp}(\alpha) = \text{supp}(p_Y) \cap \text{supp}(q_Y)$. In conclusion, such a formulation ensures that the OT problems for all shared classes are considered, i.e., $\alpha(y) > 0$ if y is shared classes.

For empirical application, the conditional OT in Eq. (4) ensures that the conditional transport model is generally feasible for extreme scenarios, e.g., outlier classes in P-DA and unseen classes in open-set scenario. However, since there are multiple (even ‘infinite’ for regression scenario with continuous Y) minimization problems (i.e., $\text{OT}(P_{X|y}^w, Q_{X|y}, c)$ in Eq. (4)), this formulation is not generally applicable and has no closed-form solution. Besides, the coefficient α is a hyperparameter, which is usually hard to estimate. In the next, we will develop an equivalent proxy for conditional OT, i.e., MOT, which provides an interpretable solution for α via the essential of data.

2.3. Theoretical Proxy and Estimation

In this section, we focus on the computation-friendly proxy of conditional OT. Considering the empirical scenario with finite samples, where labeled source data $\mathcal{D}^s = \{\mathbf{x}_i^s, y_i^s\}_{i=1}^n$ and unlabeled target data $\mathcal{D}^t = \{\mathbf{x}_j^t\}_{j=1}^m$ are given, we first define the mask matrix for discrete OT and propose the MOT methodology. Then we prove the theoretical connection between MOT and conditional OT Eq. (4), which ensures that MOT can be a generally applicable proxy with explicit fixed point solution.

Usually, P and Q are denoted as empirical distributions, e.g., $P_X = \frac{1}{n} \sum_i \delta_{\mathbf{x}_i^s}$. Then the discrete formulation of Eq. (3) can be written as

$$S^{\lambda, \beta}(P^w, Q, \mathbf{C}) = \min_{\Gamma \in \mathcal{M}_+(\mathbb{R}^{n \times m})} \langle \Gamma, \mathbf{C} \rangle_F + \lambda \langle \Gamma, \ln \Gamma \rangle_F + \beta [D_\phi(\Gamma_{P^w} \| P^w) + D_\phi(\Gamma_Q \| Q)], \quad (5)$$

where $\Gamma \in \mathbb{R}^{n \times m}$ is known as the transport plan matrix, $\mathbf{C}_{ij} = \|\mathbf{x}_i^s - \mathbf{x}_j^t\|_2^2$ is the pair-wise cost, and $\langle \Gamma, \mathbf{C} \rangle_F = \sum_{ij} \Gamma_{ij} \mathbf{C}_{ij}$ is the Frobenius inner product. For conditional OT, the empirical conditional distributions are defined as $P_{X|Y=l} = \frac{1}{n_l} \sum_i \mathbb{I}_{[y_i^s=l]} \delta_{\mathbf{x}_i^s}$ ($Q_{X|Y}$ is similar), where n_l is the size of the l -th class source data. Then the dimensions of plan and cost matrices in each OT problem of Eq. (4) are $n_l \times m_l$. Therefore, it is clear that the conditional OT induces multiple OT problems between the cross-domain clusters, and the plan in each sub-problem characterizes the class-level dependency between source and target domains.

To overcome slice-wise computation problem of conditional OT, the key idea is incorporating the multiple sub-problems into single OT problem with mask mechanism.

Definition 3 (Masked OT) Given labels $\{y_i^s\}_{i=1}^n$ and $\{y_j^t\}_{j=1}^m$, the mask matrix $\mathbf{M} \in \mathbb{R}^{n \times m}$ is defined as

$$\mathbf{M}_{ij} \triangleq \begin{cases} 1, & \text{if } y_i^s = y_j^t, \\ \infty, & \text{if } y_i^s \neq y_j^t. \end{cases} \quad (6)$$

Then the masked cost is defined as $\tilde{\mathbf{C}} = \mathbf{C} \odot \mathbf{M}$, and the masked OT is formulated as

$$\text{OT}_{\text{mask}}(P^w, Q, \tilde{\mathbf{C}}) = \text{OT}(P^w, Q, \tilde{\mathbf{C}}). \quad (7)$$

With the label-conditioned mask \mathbf{M} , the inter-class distances in the modified cost $\tilde{\mathbf{C}}$ will be enlarged to infinity. Intuitively, such a mask on cost ensures that the transport plan Γ will only assign values to the intra-class sample pair. For example, the masked UOT is formulated as

$$S_{\text{mask}}^{\lambda, \beta}(P^w, Q, \tilde{\mathbf{C}}) = \min_{\Gamma \in \mathcal{M}_+(\mathbb{R}^{n \times m})} \langle \Gamma, \tilde{\mathbf{C}} \rangle_F + \lambda \langle \Gamma, \ln \Gamma \rangle_F + \beta [D_\phi(\Gamma_{P^w} \| P^w) + D_\phi(\Gamma_Q \| Q)]. \quad (8)$$

Now we begin to present the main theoretical results, which connect conditional OT with the computation-friendly OT_{mask} . Proofs are provided in appendix.

Theorem 1 (Proxy) Assume $\text{supp}(q_Y) \subseteq \text{supp}(p_Y)$, the following identities hold.

(a) **Kantorovich OT:**

$$\text{OT}_{\text{mask}}(P^{w^*}, Q, \tilde{\mathbf{C}}) = \text{OT}_{\text{cond}}^{q_Y}(P^{w^*}, Q, \mathbf{C}).$$

(b) **Sinkhorn OT:**

$$S_{\text{mask}}^\lambda(P^{w^*}, Q, \tilde{\mathbf{C}}) + \lambda H(Q_Y) = S_{\text{cond}}^{\lambda, q_Y}(P^{w^*}, Q, \mathbf{C}).$$

(c) **UOT:** there exists non-negative $\alpha(\cdot)$ on \mathcal{Y} such that $\text{supp}(\alpha) = \text{supp}(q_Y)$, $\sum_{y \in \mathcal{Y}} \alpha(y) = 1$ and

$$S_{\text{mask}}^{\lambda, \beta}(P^{w^*}, Q, \tilde{\mathbf{C}}) + C_0(\alpha, Q_Y) = S_{\text{cond}}^{\lambda, \beta, \alpha}(P^{w^*}, Q, \mathbf{C}),$$

where C_0 is a constant depending only on α and Q_Y .

Thm. 1 shows that MOT serves as a well-defined approximation for conditional OT. Specifically, for Kantorovich OT, MOT OT_{mask} exactly equals to conditional OT $\text{OT}_{\text{cond}}^\alpha$; for Sinkhorn OT and UOT, MOT is equivalent (up to a constant) to conditional OT $\text{OT}_{\text{cond}}^\alpha$, and minimizing the computation-friendly MOT is sufficient to ensure a small $\text{OT}_{\text{cond}}^\alpha$. We also present an intuitive illustration for the connection between $\text{OT}_{\text{cond}}^\alpha$ and OT_{mask} in Fig. 2. In conclusion, the main results imply that MOT can also achieve conditional transport for class-level knowledge transfer, while overcoming the weaknesses of slice-wise computation and hyperparameter α in $\text{OT}_{\text{cond}}^\alpha$.

3. Algorithm and Application

Based on the theoretical guarantee of MOT, we now present the numerical algorithm for solving Eq. (8) explicitly. Then a conditional invariant model with risk minimization on the reweighted source domain and transported source domain is proposed for PDA problem.

3.1. Algorithm: A Fixed Point Method

In this section, we focus on the numerical solution for the masked UOT problem in Eq. (8). Note that such a solution is also applicable for the Sinkhorn OT, since $S_{\text{mask}}^{\lambda, \beta}$ will boil down to S_{mask}^λ when $\beta = \infty$ (i.e., strict constraints),

As studied by Chizat et al. [6, Thm. 3.11], if $D_\phi = \text{KL}$, the original UOT problem Eq. (5) admits a fixed point solution with explicit computation. Specifically, defining the kernel $\mathcal{K} = \exp(-\mathbf{C}/\lambda) \in \mathbb{R}^{n \times m}$, the fixed point iterations for non-negative scaling vectors is formulated as

$$\mathbf{s}_1^{(i)} = \left(\frac{p_X^w}{\mathcal{K} \mathbf{s}_2^{(i-1)}} \right)^{\frac{\beta}{\beta + \lambda}}, \quad \mathbf{s}_2^{(i)} = \left(\frac{q_X}{\mathcal{K}^\top \mathbf{s}_1^{(i)}} \right)^{\frac{\beta}{\beta + \lambda}}. \quad (9)$$

Then the sequence $\text{diag}(\mathbf{s}_1^{(i)}) \mathcal{K} \text{diag}(\mathbf{s}_2^{(i)})$ will converge to the solution of Eq. (5) (i.e., optimal transport plan Γ^*).

Similarly, the masked UOT Eq. (8) can also be solved by the fixed point method. The major difference is that the

Algorithm 1 Fixed Point Method for Eq. (8)

Input: observations $\{(\mathbf{x}_i^s, y_i^s)\}_{i=1}^n$ of P and $\{(\mathbf{x}_i^t, y_i^t)\}_{i=1}^m$ of Q , maximum iteration I_{max} , importance weight w , penalty parameter λ, β ;

Output: transport plan Γ ;

- 1: Initialize $\mathbf{s}_2^{(0)} = [1, 1, \dots, 1]^\top / m \in \mathbb{R}^m$;
 - 2: Compute p_X^w as Def. 1 and $\tilde{\mathcal{K}}$ as Eq. (10);
 - 3: **for** $i = 1, 2, \dots, I_{\text{max}}$ **do**
 - 4: $\mathbf{s}_1^{(i)} \leftarrow \left(\frac{p_X^w}{\tilde{\mathcal{K}} \mathbf{s}_2^{(i-1)}} \right)^{\frac{\beta}{\beta + \lambda}}, \quad \mathbf{s}_2^{(i)} \leftarrow \left(\frac{q_X}{\tilde{\mathcal{K}}^\top \mathbf{s}_1^{(i)}} \right)^{\frac{\beta}{\beta + \lambda}}$;
 - 5: **end for**
 - 6: $\tilde{\Gamma} \leftarrow \text{diag}(\mathbf{s}_1^{(I_{\text{max}})}) \tilde{\mathcal{K}} \text{diag}(\mathbf{s}_2^{(I_{\text{max}})})$.
-

kernel $\tilde{\mathcal{K}}$ induced by label-conditioned masked cost $\tilde{\mathbf{C}}$ is derived as

$$\tilde{\mathcal{K}}_{ij} = \begin{cases} 0, & \text{if } \tilde{\mathbf{C}}_{ij} = \infty, \\ \exp(-\mathbf{C}_{ij}/\lambda), & \text{else.} \end{cases} \quad (10)$$

Then the transport plan for masked UOT can be deduced by replacing the kernel \mathcal{K} in scaling iteration Eq. (9) with the masked kernel $\tilde{\mathcal{K}}$. Note that since $\Gamma^* \approx \text{diag}(\mathbf{s}_1^{(i)}) \tilde{\mathcal{K}} \text{diag}(\mathbf{s}_2^{(i)})$, then the transport mass Γ_{ij}^* will be 0 if $\tilde{\mathcal{K}}_{ij} = 0$. Intuitively, the mask mechanism ensures that the assignments between the inter-class samples will be significantly reduced. We summarize the fixed point method for masked UOT Eq. (8) in Alg. 1, where the masked Sinkhorn OT ($\beta = \infty$) can be solved by setting $\beta/(\beta + \lambda) = 1$ in step 4. Note that since $\tilde{\mathcal{K}}$ is block diagonal (up to permutations on samples), the iterations in step 4 can be effectively implemented by iterating the block matrices parallelly.

3.2. Application: PDA

In this section, we focus on the OT-based modeling for extreme PDA scenario, and propose MOT-based model to learn the cross-domain invariant knowledge.

Motivated by the discussion of successful DA in Sec. 2.1, we are interested in two problems: 1) learning conditional invariant representation Z via OT; 2) learning transferrable model with unbiased risk estimation. With these goals in mind, we decompose the learning model f as $f_c \circ f_r$, where $f_r : X \mapsto Z$ is representation learner and $f_c : Z \mapsto Y$ is task learner. These compositions are parameterized by neural networks, and the overall framework can be summarized as following two alternative processes (shown as Fig. 2).

Transport Assignment Learning. In this process, we aim to learn an ideal transport plan Γ for characterizing the cross-domain sample correspondence. Firstly, we estimate weight w with BBSE [22] algorithm, whose convergence property is theoretically ensured. Then the label-




Figure 2. Illustration of the MOT-based model for PDA. **Stage Red**: transport assignment learning for obtaining ideal transport plan $\tilde{\Gamma}$ with well-defined proxy OT_{mask} ; **Stage Black**: conditional alignment and risk minimization for learning conditional invariant representation Z and transferrable model $f_c \circ f_r$.

conditioned masked UOT Eq. (8) is employed, i.e.,

$$\mathcal{L}_{\text{cond}}(f_r, \Gamma) = \langle \Gamma, \tilde{C} + \lambda \ln \Gamma \rangle_F + \beta [D_\phi(\Gamma_{P_Z^w} \| P_Z^w) + D_\phi(\Gamma_{Q_Z} \| Q_Z)], \quad (11)$$

where the cost matrix $C_{ij} = \|\mathbf{z}_i^s - \mathbf{z}_j^t\|_2^2$. Note that this transport assignment learning is implemented in representation space Z with learner f_r , which ensures the model can dynamically measure and bridge the conditional discrepancy during training procedure. Since there are no labels for target domain, we use the pseudo labels $\hat{y} = f(\mathbf{x})$ for estimating the masked cost \tilde{C} . In this stage, the transport plan will be learned with fixed representation learner f_r , i.e., compute $\tilde{\Gamma} = \arg \min_{\Gamma} \mathcal{L}_{\text{cond}}$ according to Alg. 1.

Conditional Alignment and Risk Minimization. In this process, we aim to estimate unbiased risk and cross-domain conditional discrepancy based on the learned ideal assignment $\tilde{\Gamma}$. Then the model f will be optimized according to the estimated objectives.

For conditional alignment, the masked UOT loss $\mathcal{L}_{\text{cond}}$ Eq. (11) with fixed $\tilde{\Gamma}$ is sufficient to achieve class-level transport. Theoretically, it can also mitigate the conditional distribution shift as shown in Thm. 1. Besides, since $\mathcal{L}_{\text{cond}}$ is built on reweighted source domain, the risk of negative transfer can be significantly reduced.

For risk minimization, apart from the risk estimation on reweighted source P^w , we propose to learn risk on transported source \tilde{P}^w based on the learned ideal assignment $\tilde{\Gamma}$. Such transport establishes the explicit connection between empirical risk and target data. Specifically, given $\tilde{\Gamma}$, the source data can be represented by the target samples via the following *barycenter map problem*:

$$\tilde{\mathbf{z}}_i^s = \arg \min_{\mathbf{z}} \sum_{j=1}^m \tilde{\Gamma}_{ij} \|\mathbf{z} - \mathbf{z}_j^t\|_2^2. \quad (12)$$

As shown by Courty et al. [8], an analytic solution for Eq. (12), called *barycenter mapping* ψ , can be written as

$$\tilde{\mathbf{z}}_i^s = \psi_{\tilde{\Gamma}_{i,:}}(\mathbf{Z}^t) = (\langle \tilde{\Gamma}_{i,:}, \mathbf{1}_m \rangle)^{-1} \sum_{j=1}^m \tilde{\Gamma}_{ij} \mathbf{z}_j^t. \quad (13)$$

Then we denote observations of transported source \tilde{P}^w as $\{(\tilde{\mathbf{z}}_i^s, y_i^s)\}_{i=1}^n$, and propose to learn empirical risk on \tilde{P}^w as

$$\mathbb{E}_{\tilde{P}^w}[\ell(f_c(\tilde{\mathbf{z}}^s), y^s)] = \mathbb{E}_Q[\ell(f_c \circ \psi \circ f_r(\mathbf{x}^t), y^s)]. \quad (14)$$

An intuitive explanation for barycenter mapping ψ is that it considers the minimized cost for reconstructing source sample \mathbf{z}^s in the representation space of target domain, i.e., the range space $\text{Im}(\mathbf{Z}^t)$. Such a reconstruction ensures that the target representation space can be supervisedly optimized via its transport correspondence with source data, i.e., $\psi \circ f_r(\mathbf{x}^t)$ in transported source risk Eq. (14). With P^w and \tilde{P}^w , we can formulate the unbiased empirical risk estimations for task learning as

$$\mathcal{L}_{\text{risk}}(f_r, f_c) = \mathbb{E}_{P^w}[\ell(f(\mathbf{x}^s), y^s)] + \mathbb{E}_{\tilde{P}^w}[\ell(f_c(\tilde{\mathbf{z}}^s), y^s)].$$

Finally, the optimization objective in conditional alignment and risk minimization stage can be summarized as

$$\min_{f_r, f_c} \mathcal{L}(f_r, f_c) = \mathcal{L}_{\text{risk}}(f_r, f_c) + \eta \mathcal{L}_{\text{cond}}(f_r, \tilde{\Gamma}). \quad (15)$$

4. Experiments

In this section, we validate MOT methodology and evaluate the proposed PDA model by conducting experiments on standard PDA datasets, i.e., Office-Home [39], VisDA-2017 [31], Office-31 [36] and ImageCLEF [4]. Details of datasets and implementations are provided in appendix.

Proxy Analysis. To validate the proxy property of masked OT in Thm. 1, we compare values of $\text{OT}_{\text{cond}}^\alpha$ and OT_{mask} on Office-Home. Results in Fig. 3a-3d imply that the absolute difference between $\text{OT}_{\text{cond}}^\alpha$ and OT_{mask} (with constant in Thm. 1) is almost zero and is negligible compared to the scales of OT values. These results demonstrate that the theoretical results are valid, and $\text{OT}_{\text{cond}}^\alpha$ and proxy OT_{mask} are indeed equivalent in application.

OT Formulations. To compare the existing OT formulations with proposed modelings, we visualize the plan values (darker color is higher in value) and report the masses assigned to outlier, (shared) inter-class and intra-class on Office-Home. Results in Fig. 3e-3h demonstrate that reweighted OT actually reduces assignments to outlier classes, but estimated w may be less accurate for shared classes and inter-class mass is large. Though UOT learns better assignments for shared classes, there is large mass for outlier classes since marginal penalty may be unaffordable for extreme scenario. In Fig. 3g, our reweighted and relaxation formulation Eq. (3) is superior and the intra-class mass is larger. Further, MOT ensures a significant block diagonal structure for transport plan in Fig. 3h, where the inter-class mass and outlier mass are significantly reduced by penalty on cost matrix. These results validate that MOT methodology can learn better plan for task knowledge transfer.

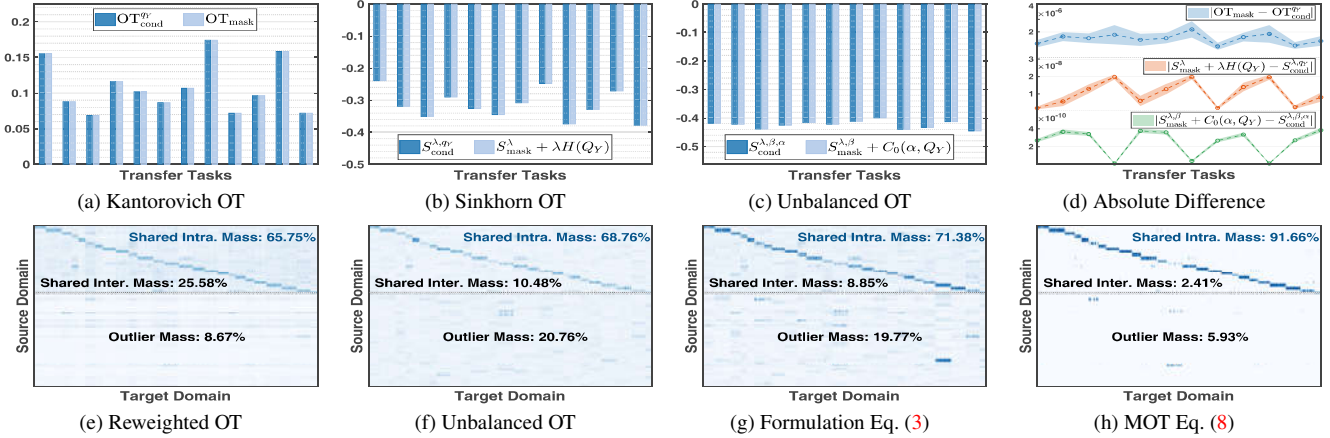


Figure 3. (a)-(d): Numerical validation of the proxy theorem on all transfer tasks of Office-Home, where $OT_{\text{cond}}^{\alpha}$ and OT_{mask} are actually equivalent in empirical scenarios. (e)-(h): Comparison of different OT formulations, where proposed methods Eq. (3) and Eq. (8) are better.

Modules			Office Home	VisDA 2017	Office Image 31	Image CLEF	Mean
OT	$\mathbb{E}_{\tilde{P}_w}[\ell]$	$\mathcal{L}_{\text{cond}}$					
ROT	✓	✓	73.9	73.1	93.7	88.2	82.2
UOT	✓	✓	69.1	67.6	92.5	87.1	79.1
Eq. (3)	✓	✓	74.2	76.8	94.3	89.2	83.6
MOT	✓		76.6	84.9	95.4	90.4	86.8
MOT		✓	74.6	89.3	97.5	91.9	88.3
MOT	✓	✓	80.6	92.4	98.4	93.6	91.3

Table 1. Ablation study (ResNet-50).

Ablation Study. To evaluate the effectiveness of different modules, we present results of ablation experiments in Tab. 1. We first consider different OT formulations for model proposed in Sec. 3.2. The 1st-3rd rows imply that proposed formulation Eq. (3) achieve higher accuracies than reweighted OT (ROT) and UOT. Further, the 6th row implies mask mechanism can improve accuracies significantly with label-conditioned knowledge transfer. These results validate the superiority of new OT formulations. Besides, we consider the transported source risk $\mathbb{E}_{\tilde{P}_w}[\ell]$ in Eq. (14) and conditional alignment $\mathcal{L}_{\text{cond}}$ in Eq. (15). Comparing the results in 4th-6th, we observe that both $\mathbb{E}_{\tilde{P}_w}[\ell]$ and $\mathcal{L}_{\text{cond}}$ can improve model performance, and the full MOT model is significantly better. These results demonstrate that proposed invariant learning model with transported risk minimization is effective in dealing with PDA problem.

Feature Visualization. To analyze the quality of learned representations, we use t-SNE [26] to visualize the 2-D features of different OT formulations on Office-Home. For reweighted OT and UOT, though the negative impacts of outlier classes are reduced as Fig. 4a-4b, the discriminability for shared classes are not sufficiently learned, and some samples are transferred to incorrect classes with marginal adaptation. For formulation Eq. (3), the representations are more compact, and fewer samples are overlapping with outlier classes. Further, the conditional alignment in MOT ensures the class-wise alignment for cross domain samples,

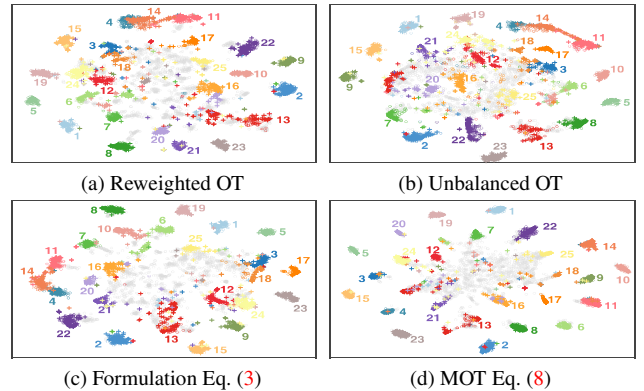


Figure 4. t-SNE [26] visualization of learned representations, where shared classes have different color and outlier classes are gray. 'o': source samples; '+': target samples; 'number': labels.

then the intra-class compactness and inter-class separability are significantly improved as Fig. 4d. These results demonstrate that the proposed OT formulations and invariant risk learning model indeed learn correct transport relations and discriminative representations.

Different PDA Scenarios. We evaluate OT-based models under different PDA scenarios by varying the number of shared classes on Office-Home. The results in Fig 5a-5b imply that all models achieve high accuracy when there are less shared classes, since there are relatively more samples/knowledge on the source domain. With the increase of shared classes, the label discrepancy is smaller, but the target classification task is more challenging since there are more target samples. Thus, we can observe that the accuracies increase at interval [20, 30] and then decrease with the increasing target samples. Generally, MOT is superior to other OT modelings, which demonstrate the effectiveness of proposed methodology in different learning scenarios.

Comparison with SOTA. To evaluate the model performance, we compare MOT with several SOTA PDA methods, and present results in Tab. 2-3. 1) For adversarial adapta-

Methods	Office-Home												VisDA-2017	
	A→C	A→P	A→R	C→A	C→P	C→R	P→A	P→C	P→R	R→A	R→C	R→P	Mean	S→R
Source [17]	46.3	67.5	75.9	59.1	59.9	62.7	58.2	41.8	74.9	67.4	48.2	74.2	61.4	45.3
DANN [13]	43.8	67.9	77.5	63.7	59.0	67.6	56.8	37.1	76.4	69.2	44.3	77.5	61.7	51.0
PADA [2]	52.0	67.0	78.7	52.2	53.8	59.0	52.6	43.2	78.8	73.7	56.6	77.1	62.1	53.5
ETN [3]	59.2	77.0	79.5	62.9	65.7	75.0	68.3	55.4	84.4	75.7	57.7	84.5	70.5	-
SAFN [43]	58.9	76.3	81.4	70.4	73.0	77.8	72.4	55.3	80.4	75.8	60.4	79.9	71.8	67.7
DRCN [21]	51.6	75.8	82.0	62.9	65.1	72.9	67.4	50.0	81.0	76.4	57.7	79.3	68.5	58.2
DMP [25]	59.0	81.2	86.3	68.1	72.8	78.8	71.2	57.6	84.9	77.3	61.5	82.9	73.5	72.7
JUMBOT [11]	62.7	77.5	84.4	76.0	73.3	80.5	74.7	60.8	85.1	80.2	66.5	83.9	75.5	84.0 [†]
AR [16]	67.4	85.3	90.0	77.3	70.6	85.2	79.0	64.8	89.5	80.4	66.2	86.4	78.3	88.7
m-POT [28]	64.6	80.6	87.2	76.4	77.6	83.6	77.1	63.7	87.6	81.4	68.5	87.4	78.0	87.0 [†]
MOT	63.1	86.1	92.3	78.7	85.4	89.6	79.8	62.3	89.7	83.8	67.0	89.6	80.6	92.4

Table 2. Classification accuracies (%) on Office-Home and VisDA-2017 datasets (ResNet-50). [†] Results reported by Salvador et al. [37].

Office-31	A→WD	W→W	DA→DA	DD→DD	AW→AW	A→A	Mean
Source [17]	75.6	96.3	98.1	83.4	83.9	85.0	87.1
DANN [13]	73.6	96.3	98.7	81.5	82.8	86.1	86.5
PADA [2]	86.5	99.3	100	82.2	92.7	95.4	92.7
SAFN [43]	87.5	96.6	99.4	89.8	92.6	92.7	93.1
DRCN [21]	90.8	100	100	94.3	95.2	94.8	95.9
DMP [25]	96.6	100	100	96.4	95.1	95.4	97.2
AR [16]	93.5	100	99.7	96.8	95.5	96.0	96.9
MOT	99.3	100	100	98.7	96.1	96.4	98.4

ImageCLEF	I→P	P→I	I→C	C→I	C→P	P→C	Mean
Source [17]	78.3	86.9	91.0	84.3	72.5	91.5	84.1
DANN [13]	78.1	86.3	91.3	84.0	72.1	90.3	83.7
PADA [2]	81.7	92.1	94.6	89.8	77.7	94.1	88.3
SAFN [43]	79.5	90.7	93.0	90.3	77.8	94.0	87.5
DMP [25]	82.4	94.5	96.7	94.3	78.7	96.4	90.5
MOT	87.7	95.0	98.0	95.0	87.0	98.7	93.6

Table 3. Classification accuracies (%) on Office-31 and ImageCLEF datasets (ResNet-50).




Figure 5. Accuracy curves and 95% confidence intervals of OT-based models by varying the number of shared classes.

tion, PADA improves DANN by introducing the reweighted source P^w to adversarial training. The better performance validates the effectiveness of label distribution correction for PADA. Compared with these adversarial models with marginal alignment, MOT learns conditional invariant representations with masked mechanism. 2) For metric-based adaptation models SAFN and DMP, though label-conditioned information is incorporated, the explicit conditional distribution alignment is not ensured. 3) Compared with other OT models (JUMBOT, AR and m-POT), MOT overcomes the limitations of existing OT models as dis-

cussed before. Specifically, AR is reweighted-based model, which applies P^w to the dual OT problem, but the performance may be sensitive to the precision of w . JUMBOT and m-POT are respectively unbalanced transport model and partial transport model, which can be both regarded as relaxation-based OT without strict marginal constraints, but the relaxation may be less effective when label shift is severe, e.g., PADA. Fortunately, our formulation Eq. (3) integrating the advantages of relaxation and reweighted, and mask mechanism further ensures the sufficiency of conditional shift correction. Therefore, MOT achieve significant improvements (about 1%~4%) on four PADA datasets, which validates the effectiveness of proposed methodology.

5. Conclusion

In this paper, we systematically study the limitations in current OT modelings for real-world learning scenarios, and then develop novel OT methodology with theoretical guarantees. For theoretical aspect, a novel formulation is proposed with reweighted and relaxation operations, and a new OT variant called MOT is proposed by exploring the mask mechanism; the equivalent relation between MOT and conditional OT is proved, which implies the computation-friendly MOT can also characterize the conditional information in transportation. For methodology, we propose an invariant learning model with MOT, whose effectiveness is validated by extensively numerical experiments and analysis. An interesting future direction is studying theory and algorithm for masked partial OT and MOT with soft labels.

Acknowledgement

This work is supported in part by the National Key R&D Program of China under Grant 2021YFA1003001; in part by the National Natural Science Foundation of China under Grant 61976229; in part by the Guangdong Basic and Applied Basic Research Foundation under Grant 2023B1515020004; and in part by the Open Research Projects of Zhejiang Lab under Grant 2021KH0AB08.

References

- [1] Shai Ben-David, John Blitzer, Koby Crammer, Alex Kulesza, Fernando Pereira, and Jennifer Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010. [1](#)
- [2] Zhangjie Cao, Lijia Ma, Mingsheng Long, and Jianmin Wang. Partial adversarial domain adaptation. In *ECCV*, pages 135–150, 2018. [1](#), [2](#), [8](#)
- [3] Zhangjie Cao, Kaichao You, Mingsheng Long, Jianmin Wang, and Qiang Yang. Learning to transfer examples for partial domain adaptation. In *CVPR*, pages 2985–2994, 2019. [1](#), [2](#), [8](#)
- [4] Barbara Caputo, Henning Müller, Jesus Martinez-Gomez, Mauricio Villegas, Burak Acar, Novi Patricia, Neda Mastviti, Suzan Üsküdarlı, Roberto Paredes, Miguel Cazorla, et al. Imageclef 2014: Overview and analysis of the results. In *International Conference of the Cross-Language Evaluation Forum for European Languages*, pages 192–211, 2014. [6](#)
- [5] Laetitia Chapel, Rémi Flamary, Haoran Wu, Cédric Févotte, and Gilles Gasso. Unbalanced optimal transport through non-negative penalized linear regression. In *NeurIPS*, volume 34, pages 23270–23282, 2021. [1](#)
- [6] Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Scaling algorithms for unbalanced optimal transport problems. *Mathematics of Computation*, 87(314):2563–2609, 2018. [1](#), [3](#), [5](#)
- [7] Lenaïc Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. Unbalanced optimal transport: Dynamic and Kantorovich formulations. *Journal of Functional Analysis*, 274(11):3090–3123, 2018. [1](#), [3](#)
- [8] Nicolas Courty, Rmi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE TPAMI*, 39(9):1853–1865, 2017. [1](#), [2](#), [3](#), [6](#)
- [9] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NeurIPS*, volume 26, 2013. [1](#), [3](#)
- [10] Bharath Bhushan Damodaran, Benjamin Kellenberger, Rémi Flamary, Devis Tuia, and Nicolas Courty. Deepjdot: Deep joint distribution optimal transport for unsupervised domain adaptation. In *ECCV*, pages 447–463, 2018. [1](#), [2](#), [3](#)
- [11] Kilian Fatras, Thibault Séjourné, Rémi Flamary, and Nicolas Courty. Unbalanced minibatch optimal transport; applications to domain adaptation. In *ICML*, pages 3186–3197. PMLR, 2021. [1](#), [2](#), [3](#), [8](#)
- [12] Kilian Fatras, Younes Zine, Rémi Flamary, Remi Gribonval, and Nicolas Courty. Learning with minibatch wasserstein: asymptotic and gradient properties. In *AISTAT*, pages 2131–2141, 2020. [1](#)
- [13] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015. [8](#)
- [14] Mingming Gong, Kun Zhang, Tongliang Liu, Dacheng Tao, Clark Glymour, and Bernhard Schölkopf. Domain adaptation with conditional transferable components. In *ICML*, pages 2839–2848, 2016. [2](#)
- [15] Xiang Gu, Yucheng Yang, Wei Zeng, Jian Sun, and Zongben Xu. Keypoint-guided optimal transport with applications in heterogeneous domain adaptation. In *NeurIPS*, 2022. [2](#)
- [16] Xiang Gu, Xi Yu, Jian Sun, Zongben Xu, et al. Adversarial reweighting for partial domain adaptation. In *NeurIPS*, volume 34, pages 14860–14872, 2021. [1](#), [2](#), [3](#), [8](#)
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, pages 770–778, 2016. [8](#)
- [18] Xiang Jiang, Qicheng Lao, Stan Matwin, and Mohammad Havaei. Implicit class-conditioned domain alignment for unsupervised domain adaptation. In *ICML*, pages 4816–4827, 2020. [2](#)
- [19] Matthieu Kirchmeyer, Alain Rakotomamonjy, Emmanuel de Bezenac, and Patrick Gallinari. Mapping conditional distributions for domain adaptation under generalized target shift. In *ICLR*, 2022. [1](#), [2](#), [3](#)
- [20] Mengxue Li, Yi-Ming Zhai, You-Wei Luo, Peng-Fei Ge, and Chuan-Xian Ren. Enhanced transport distance for unsupervised domain adaptation. In *CVPR*, pages 13936–13944, 2020. [1](#), [2](#), [3](#)
- [21] Shuang Li, Chi Harold Liu, Qiuxia Lin, Qi Wen, Limin Su, Gao Huang, and Zhengming Ding. Deep residual correction network for partial domain adaptation. *IEEE TPAMI*, 43(7):2329–2344, 2021. [1](#), [8](#)
- [22] Zachary Lipton, Yu-Xiang Wang, and Alexander Smola. Detecting and correcting for label shift with black box predictors. In *ICML*, pages 3122–3130. PMLR, 2018. [1](#), [2](#), [5](#)
- [23] Mingsheng Long, Zhangjie Cao, Jianmin Wang, and Michael I Jordan. Conditional adversarial domain adaptation. In *NeurIPS*, volume 31, 2018. [2](#)
- [24] You-Wei Luo and Chuan-Xian Ren. Conditional Bures metric for domain adaptation. In *CVPR*, pages 13989–13998, 2021. [1](#), [2](#)
- [25] You-Wei Luo, Chuan-Xian Ren, Dao-Qing Dai, and Hong Yan. Unsupervised domain adaptation via discriminative manifold propagation. *IEEE TPAMI*, 44(3):1653–1669, 2022. [1](#), [2](#), [8](#)
- [26] Laurens van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *JMLR*, 9(11):2579–2605, 2008. [7](#)
- [27] Khai Nguyen, Dang Nguyen, Quoc Nguyen, Tung Pham, Hung Bui, Dinh Phung, Trung Le, and Nhat Ho. On transportation of mini-batches: A hierarchical approach. In *ICML*, page 34, 2022. [1](#)
- [28] Khai Nguyen, Dang Nguyen, Tung Pham, Nhat Ho, et al. Improving mini-batch optimal transport via partial transportation. In *ICML*, pages 16656–16690, 2022. [1](#), [2](#), [8](#)
- [29] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE TKDE*, 22(10):1345–1359, 2009. [1](#)
- [30] Vishal M Patel, Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Visual domain adaptation: A survey of recent advances. *IEEE Signal Processing Magazine*, 32(3):53–69, 2015. [1](#)
- [31] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017. [1](#), [6](#)

- [32] Alain Rakotomamonjy, Rémi Flamary, Gilles Gasso, M El Alaya, Maxime Berar, and Nicolas Courty. Optimal transport for conditional domain matching and label shift. *Machine Learning*, 111(5):1651–1670, 2022. 1, 2, 3
- [33] Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. Optimal transport for multi-source domain adaptation under target shift. In *AISTAT*, pages 849–858. PMLR, 2019. 1, 2, 3
- [34] Chuan-Xian Ren, Pengfei Ge, Peiyi Yang, and Shuicheng Yan. Learning target-domain-specific classifier for partial domain adaptation. *IEEE TNNLS*, 32(5):1989–2001, 2021. 2
- [35] Chuan-Xian Ren, Xiao-Lin Xu, and Hong Yan. Generalized conditional domain adaptation: A causal perspective with low-rank translators. *IEEE TCYB*, 50(2):821–834, 2018. 2, 3
- [36] Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *ECCV*, pages 213–226, 2010. 6
- [37] Tiago Salvador, Kilian Fatras, Ioannis Mitliagkas, and Adam Oberman. A reproducible and realistic evaluation of partial domain adaptation methods. *arXiv preprint arXiv:2210.01210*, 2022. 8
- [38] Remi Tachet des Combes, Han Zhao, Yu-Xiang Wang, and Geoffrey J Gordon. Domain adaptation with conditional distribution matching and generalized label shift. In *NeurIPS*, volume 33, pages 19276–19289, 2020. 1, 2, 3
- [39] Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *CVPR*, pages 5018–5027, 2017. 1, 6
- [40] Pengfei Wei, Yiping Ke, Xinghua Qu, and Tze-Yun Leong. Subdomain adaptation with manifolds discrepancy alignment. *IEEE TCYB*, 52(11):11698–11708, 2022. 2
- [41] Haifeng Xia, Taotao Jing, and Zhengming Ding. Maximum structural generation discrepancy for unsupervised domain adaptation. *IEEE TPAMI*, 2022. 2
- [42] Geng-Xin Xu, Chen Liu, Jun Liu, Zhongxiang Ding, Feng Shi, Man Guo, Wei Zhao, Xiaoming Li, Ying Wei, Yaozong Gao, Chuan Xian Ren, and Dinggang Shen. Cross-site severity assessment of covid-19 from ct images via domain adaptation. *IEEE Transactions on Medical Imaging*, 41(1):88–102, 2021. 1
- [43] Ruijia Xu, Guanbin Li, Jihan Yang, and Liang Lin. Larger norm more transferable: An adaptive feature norm approach for unsupervised domain adaptation. In *ICCV*, pages 1426–1435, 2019. 8
- [44] Renjun Xu, Pelen Liu, Liyan Wang, Chao Chen, and Jindong Wang. Reliable weighted optimal transport for unsupervised domain adaptation. In *CVPR*, pages 4394–4403, 2020. 3
- [45] Hongliang Yan, Yukang Ding, Peihua Li, Qilong Wang, Yong Xu, and Wangmeng Zuo. Mind the class weight bias: Weighted maximum mean discrepancy for unsupervised domain adaptation. In *CVPR*, pages 2272–2281, 2017. 2
- [46] Jiying Zhang, Xi Xiao, Long-Kai Huang, Yu Rong, and Yatao Bian. Fine-tuning graph neural networks via graph topology induced optimal transport. In *IJCAI*, 2022. 2
- [47] Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *ICML*, pages 819–827. PMLR, 2013. 1, 2, 3
- [48] Yuchen Zhang, Tianle Liu, Mingsheng Long, and Michael Jordan. Bridging theory and algorithm for domain adaptation. In *ICML*, pages 7404–7413. PMLR, 2019. 1
- [49] Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *ICML*, pages 7523–7532, 2019. 1
- [50] Yongchun Zhu, Fuzhen Zhuang, Jindong Wang, Guolin Ke, Jingwu Chen, Jiang Bian, Hui Xiong, and Qing He. Deep subdomain adaptation network for image classification. *IEEE TNNLS*, 32(4):1713–1722, 2021. 2
- [51] Fuzhen Zhuang, Zhiyuan Qi, Keyu Duan, Dongbo Xi, Yongchun Zhu, Hengshu Zhu, Hui Xiong, and Qing He. A comprehensive survey on transfer learning. *Proceedings of the IEEE*, 109(1):43–76, 2020. 1
- [52] Yang Zou, Zhiding Yu, BVK Kumar, and Jinsong Wang. Unsupervised domain adaptation for semantic segmentation via class-balanced self-training. In *ECCV*, pages 289–305, 2018. 1