



**HAL**  
open science

## Mapping conditional distributions for domain adaptation under generalized target shift

Matthieu Kirchmeyer, Alain Rakotomamonjy, Emmanuel de Bezenac, Patrick Gallinari

► **To cite this version:**

Matthieu Kirchmeyer, Alain Rakotomamonjy, Emmanuel de Bezenac, Patrick Gallinari. Mapping conditional distributions for domain adaptation under generalized target shift. 2021. hal-03396183v1

**HAL Id: hal-03396183**

**<https://hal.science/hal-03396183v1>**

Preprint submitted on 25 Oct 2021 (v1), last revised 12 Jan 2023 (v3)

**HAL** is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

---

# MAPPING CONDITIONAL DISTRIBUTIONS FOR DOMAIN ADAPTATION UNDER GENERALIZED TARGET SHIFT

Matthieu Kirchmeyer<sup>a,b</sup>, Alain Rakotomamonjy<sup>b,c</sup>, Emmanuel de Bezenac<sup>a</sup>, Patrick Gallinari<sup>a,b</sup>

<sup>a</sup> Sorbonne Université, CNRS, LIP6, F-75005 Paris, France

<sup>b</sup> Criteo AI Lab, Paris, France

<sup>c</sup> Université de Rouen, LITIS, France

Corresponding author: matthieu.kirchmeyer@gmail.com

## ABSTRACT

We consider the problem of unsupervised domain adaptation (UDA) between a source and a target domain under conditional and label shift a.k.a Generalized Target Shift (GeTarS). Unlike simpler UDA settings, few works have addressed this challenging problem. Recent approaches learn domain-invariant representations, yet they have practical limitations and rely on strong assumptions that may not hold in practice. In this paper, we explore a novel and general approach to align pretrained representations, which circumvents existing drawbacks. Instead of constraining representation invariance, it learns an optimal transport map, implemented as a NN, which maps source representations onto target ones. Our approach is flexible and scalable, it preserves the problem’s structure and it has strong theoretical guarantees under mild assumptions. In particular, our solution is unique, matches conditional distributions across domains, recovers target proportions and explicitly controls the target generalization risk. Through an exhaustive comparison on several datasets, we challenge the state-of-the-art in GeTarS.

## 1 INTRODUCTION

Unsupervised Domain Adaptation (UDA) methods (Pan & Yang, 2010) train a classifier with labelled samples from a source domain  $S$  such that its risk on an unlabelled target domain  $T$  is low. This problem is ill-posed and simplifying assumptions have been considered. Initial contributions focused on two settings which decompose differently the joint distribution over  $X \times Y$ : **covariate shift** (CoS), when input marginals differ across domains while class-posteriors are unchanged i.e.  $p_S(Y|X) = p_T(Y|X)$ ,  $p_S(X) \neq p_T(X)$  and **target shift** (TarS) when label distributions differ while conditionals are unchanged i.e.  $p_S(Y) \neq p_T(Y)$ ,  $p_S(X|Y) = p_T(X|Y)$ . Yet, for real problems, these assumptions are restrictive as it is common that both conditional and label distributions differ i.e.  $p_S(X|Y) \neq p_T(X|Y)$ ,  $p_S(Y) \neq p_T(Y)$ . This problem, named **generalized target shift** (GeTarS), is more complex and has been addressed only recently. A key challenge is to learn how to map the source domain onto the target one to minimize both conditional and label shifts, without using target labels. The current SOTA approach in Gong et al. (2016); Combes et al. (2020); Rakotomamonjy et al. (2020); Shui et al. (2021) learns domain-invariant representations and uses, as importance weights in the training objectives, the estimated class-ratios between domains. However, this approach has several limitations. First, to transfer representations, the domain-invariance constraint breaks the original problem structure and it was shown that this may degrade the discriminativity of target representations (Liu et al., 2019; Chen et al., 2019). Second, generalization guarantees are derived under strong assumptions which may not hold in practice.

In this paper, we address these limitations with a new general approach, named Optimal Sample Transformation and Reweight (OSTAR), which maps pretrained representations using Optimal Transport (OT). OSTAR proposes an alternative to constraining representation invariance and performs jointly three operations: given a pretrained encoder, (i) it learns an OT map, implemented as a neural network (NN), between encoded source and target conditionals, (ii) it estimates target proportions for sample reweighting and (iii) it learns a classifier for the target domain using source labels. OSTAR has several benefits: (i) it is flexible, scalable and preserves target discriminativity and (ii) it provides strong theoretical guarantees under mild assumptions. In summary, our contributions are:

- We propose an approach, `OSTAR`, to align pretrained representations under `GeTarS`. Without constraining representation-invariance, `OSTAR` jointly learns a classifier for inference on the target domain and an OT map, which maps representations of source conditionals to those of target ones under class-reweighting. `OSTAR` preserves target discriminativity and experimentally challenges the state-of-the-art for `GeTarS`.
- `OSTAR` implements its OT map as a NN shared across classes. Our approach is thus flexible and has native regularization biases for stability. Moreover it is scalable and generalizes beyond training samples unlike standard linear programming based OT approaches.
- `OSTAR` has strong theoretical guarantees under mild assumptions: its solution is unique, recovers target proportions and correctly matches source and target conditionals at the optimum. It also explicitly controls the target risk with a new Wasserstein-based bound.

Our paper is organized as follows. In Section 2, we define our problem, approach and assumptions. In Section 3, we derive theoretical results. In Section 4, we describe our implementation. We report in Section 5 experimental results and ablation studies. In Section 6, we present related work.

## 2 PROPOSED APPROACH

In this section, we successively define our problem, present our method, `OSTAR` and its main ideas and introduce our assumptions, used to provide theoretical guarantees for our method.

### 2.1 PROBLEM DEFINITION

Denoting  $\mathcal{X}$  the input space and  $\mathcal{Y} = \{1, \dots, K\}$  the label space, we consider UDA between a source  $S = (\mathcal{X}_S, \mathcal{Y}_S, p_S(X, Y))$  with labelled samples  $\hat{S} = \{(\mathbf{x}_S^{(i)}, y_S^{(i)})\}_{i=1}^n \in (\mathcal{X}_S \times \mathcal{Y}_S)^n$  and a target  $T = (\mathcal{X}_T, \mathcal{Y}_T, p_T(X, Y))$  with unlabelled samples  $\hat{T} = \{\mathbf{x}_T^{(i)}\}_{i=1}^m \in \mathcal{X}_T^m$ . We denote  $\mathcal{Z} \subset \mathbb{R}^d$  a latent space and  $g : \mathcal{X} \rightarrow \mathcal{Z}$  an encoder from  $\mathcal{X}$  to  $\mathcal{Z}$ .  $\mathcal{Z}_S$  and  $\mathcal{Z}_T$  are the encoded source and target input domains,  $\mathcal{Z}_{\hat{S}}$  and  $\mathcal{Z}_{\hat{T}}$  the corresponding training sets and  $Z$  a random variable in this space. The latent marginal probability induced by  $g$  on  $D \in \{S, T\}$  is defined as  $\forall A \subset \mathcal{Z}, p_D^g(A) \triangleq g_{\#}(p_D(A))$ <sup>1</sup>. For convenience,  $\mathbf{p}_D^Y \in \mathbb{R}^K$  denotes the label marginal  $p_D(Y)$  and  $p_D(Z) \triangleq p_D^g(Z)$ . In all generality, conditional distributions in this latent space and label marginals differ across domains; this is the `GeTarS` assumption (Definition 1) illustrated in Figure 1a. This assumption, made in feature space  $\mathcal{Z}$  rather than input space  $\mathcal{X}$ , states that latent representations for both domains from a given class are different with different label proportions. Operating in the latent space has several practical advantages such as improved discriminativity and dimension reduction.

**Definition 1** (`GeTarS`). *GeTarS is characterized by conditional mismatch across domains i.e.  $\exists j \in \{1, \dots, K\}, p_S(Z|Y = j) \neq p_T(Z|Y = j)$  and label shift i.e.  $\mathbf{p}_S^Y \neq \mathbf{p}_T^Y$ .*

Our goal is to learn a classifier in  $\mathcal{Z}$  with low target risk, using source labels. This is challenging as (i) target labels are unknown and (ii) there are two shifts to handle. We will show that this can be achieved with pretrained representations if we recover two key properties: (i) a map which matches source and target conditional distributions and (ii) target proportions to reweight samples by class-ratios and thus account for label shift. Our approach, `OSTAR`, achieves this objective.

### 2.2 MAPPING CONDITIONAL DISTRIBUTIONS UNDER LABEL SHIFT

We now present the components in `OSTAR`, their objective and the various training steps.

**Components** The main components of `OSTAR`, illustrated in Figure 1b, are detailed below. These components are learned and estimated using the algorithm detailed in Section 4. They include:

- a fixed encoder  $g : \mathcal{X} \rightarrow \mathcal{Z}$ , defined in Section 2.1.
- a mapping  $\phi : \mathcal{Z} \rightarrow \mathcal{Z}$ , acting on source samples encoded by  $g$ .
- a label proportion vector  $\mathbf{p}_N^Y$  on the simplex  $\Delta_K$ .
- a classifier  $f_N : \mathcal{Z} \rightarrow \{1, \dots, K\}$  for the target domain in a hypothesis class  $\mathcal{H}$  over  $\mathcal{Z}$ .

<sup>1</sup> $f_{\#}\rho$  is the push-forward measure  $f_{\#}\rho(B) = \rho(f^{-1}(B))$ , for all measurable set  $B$ .

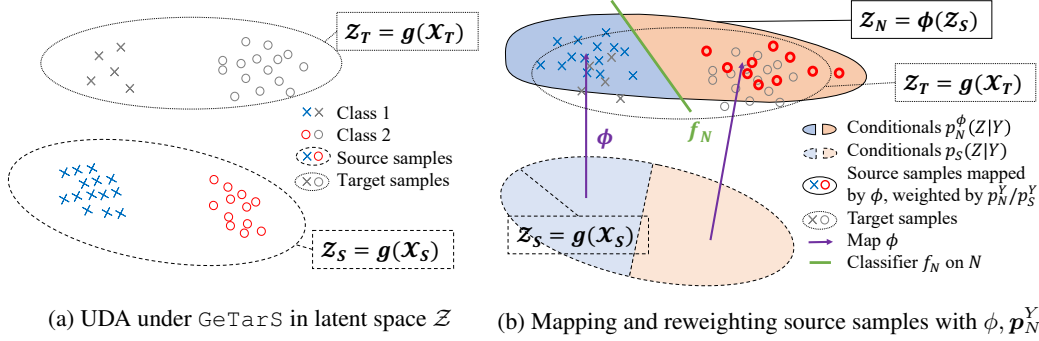


Figure 1: **Illustration of our approach on a 2-class UDA problem** [better viewed in color]. (a) A pretrained encoder  $g$  defines a latent space  $\mathcal{Z}$  with labelled source samples  $\times \circ$  and unlabelled target samples  $\times \circ$  under conditional and label shift (GeTarS). (b) We train a target classifier  $f_N$  on a new domain  $N$ , where labelled samples  $\times \circ$  are obtained by (i) mapping source samples with  $\phi$  acting on conditional distributions and (ii) reweighting these samples by estimated class-ratios  $\mathbf{p}_N^Y/\mathbf{p}_S^Y$ .  $\phi$  should match source and target conditionals and  $\mathbf{p}_N^Y$  should estimate target proportions  $\mathbf{p}_T^Y$ .

**Objective**  $g$  encodes source and target samples in a latent space such that it preserves rich information about the target task and such that the risk on the source domain is small.  $g$  is fixed throughout training to preserve target discriminativity.  $\phi$  should map encoded source conditionals in  $\mathcal{Z}_S$  onto corresponding encoded target ones in  $\mathcal{Z}_T$  to account for conditional shift;  $\mathcal{Z}_N$  denotes the mapped space.  $\mathbf{p}_N^Y$  should estimate the target proportions  $\mathbf{p}_T^Y$  to account for label shift. Components  $(\phi, \mathbf{p}_N^Y)$  define a new labelled domain in latent space  $N = (\mathcal{Z}_N, \mathcal{Y}_N, p_N(Z, Y))$  through a Sample Transformation And Reweight operation of the encoded  $S$  domain, as illustrated in Figure 1b. Indeed, the pushforward by  $\phi$  of encoded source conditionals defines conditionals in domain  $N$ ,  $p_N^\phi(Z|Y)$ :

$$\forall k, p_N^\phi(Z|Y = k) \triangleq \phi_{\#} \left( p_S(Z|Y = k) \right)^1$$

Then,  $\mathbf{p}_N^Y$  weights each conditional in  $N$ . This yields a marginal distribution in  $N$ ,  $p_N^\phi(Z)$ :

$$p_N^\phi(Z) \triangleq \sum_{k=1}^K \mathbf{p}_N^{Y=k} p_N^\phi(Z|Y = k) \quad (1)$$

Finally, classifier  $f_N$  is trained on labelled samples from domain  $N$ . This is possible as each sample in  $N$  is a projection of a labelled sample from  $S$ .  $f_N$  can then be used for inference on  $T$ . We will show that it has low target risk when components  $\phi$  and  $\mathbf{p}_N^Y$  achieve their objectives detailed above.

**Training** We train OSTAR’s components in two stages. First, we train  $g$  along a source classifier  $f_S$  from scratch by minimizing source classification loss; alternatively,  $g$  can be tailored to specific problems with pretraining if it is possible. Second, we jointly learn  $(f_N, \phi, \mathbf{p}_N^Y)$  to minimize a classification loss in domain  $N$  and to match target conditional and label distributions with those in domain  $N$ . As target conditionals and proportions are unknown, we propose a proxy problem for  $(\phi, \mathbf{p}_N^Y)$  to match instead latent marginals  $p_T(Z)$  and  $p_N^\phi(Z)$  (1). We solve this proxy problem under least action principle measured by a Monge transport cost, denoted  $\mathcal{C}(\phi)$ , as in problem (OT):

$$\begin{aligned} \min_{\phi, \mathbf{p}_N^Y \in \Delta_K} \mathcal{C}(\phi) &\triangleq \sum_{k=1}^K \int_{\mathbf{z} \in \mathcal{Z}} \|\phi(\mathbf{z}) - \mathbf{z}\|_2^2 p_S(\mathbf{z}|Y = k) d\mathbf{z} \\ \text{subject to } p_N^\phi(Z) &= p_T(Z) \end{aligned} \quad (\text{OT})$$

For any function  $\phi$  e.g. parametrized by a NN,  $\mathcal{C}(\phi)$  is the transport cost of encoded source conditionals by  $\phi$ . It uses a cost function,  $c(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^p$ , where without loss of generality  $p = 2$ . The optimal  $\mathcal{C}(\phi)$  is the sum of Wasserstein-2 distances between source conditionals and their mappings, as detailed in Appendix C. Problem (OT) seeks to minimize  $\mathcal{C}(\phi)$  under marginal matching.

Our OT formulation is key to the approach. First, it is at the basis of our theoretical analysis. Under Assumption 2 later defined, the optimal transport cost is the sum of the Wasserstein-2 distance between source and matched target conditionals. We can then provide conditions on these distances in

Assumption 3 to formally define when the solution to problem (OT) correctly matches conditionals. Second, it allows us to learn the OT map with a NN. This NN approach: (i) generalizes beyond training samples and scales up with the number of samples unlike linear programming based OT approaches (Courty et al., 2017b) and (ii) introduces useful stability biases which make learning less prone to numerical instabilities as highlighted in de Bézenac et al. (2021); Karkar et al. (2020).

### 2.3 ASSUMPTIONS

In Section 3 we will introduce the theoretical properties of our method which offers several guarantees. As always, this requires some assumptions which are introduced below. We discuss their relations with assumptions used in related work and detail why they are less restrictive.

**Assumption 1** (Cluster Assumption on S).  $\forall k, p_S^{Y=k} > 0$  and there is a partition of the source domain  $\mathcal{Z}_S$  such that  $\mathcal{Z}_S = \cup_{k=1}^K \mathcal{Z}_S^{(k)}$  and  $\forall k p_S(Z \in \mathcal{Z}_S^{(k)} | Y = k) = 1$ .

Assumption 1, inspired from Chapelle et al. (2010), states that source representations with the same label are within the same cluster. It helps guarantee that only one map is required to match source and target conditionals. Assumption 1 is for instance satisfied when the classification loss on the source domain is zero which corresponds to the training criterion of our encoder  $g$ . Interestingly other approaches e.g. Combes et al. (2020); Rakotomamonjy et al. (2020) assume that it also holds for target representations; this is harder to induce as target labels are unknown.

**Assumption 2** (Conditional matching). A mapping  $\phi$  solution to our matching problem in problem (OT) maps a source conditional to a target one i.e.  $\forall k \exists j \phi_{\#}(p_S(Z|Y = k)) = p_T(Z|Y = j)$ .

Assumption 2 establishes that  $\phi$  solution to problem (OT) performs optimal assignment between conditionals. It is less restrictive than alternatives: the *ACons* assumption in Zhang et al. (2013); Gong et al. (2016) states the existence of a map matching the right conditional pairs i.e.  $j = k$  in Assumption 2, while the *GLS* assumption of Combes et al. (2020) imposes that  $p_S(Z|Y) = p_T(Z|Y)$ . *GLS* is thus included in Assumption 2 when  $j = k$  and  $\phi = \text{Id}$  and is more restrictive.

**Assumption 3** (Cyclical monotonicity between S and T). For all  $K$  elements permutation  $\sigma$ , conditional probabilities in the source and target domains satisfy  $\sum_{k=1}^K \mathcal{W}_2(p_S(Z|Y = k), p_T(Z|Y = \sigma(k))) \leq \sum_{k=1}^K \mathcal{W}_2(p_S(Z|Y = k), p_T(Z|Y = \sigma(k)))$  with  $\mathcal{W}_2$ , the Wasserstein-2 distance.

Assumption 3, introduced in Rakotomamonjy et al. (2020), formally defines settings where conditionals are guaranteed to be correctly matched with an optimal assignment in the latent space under Assumption 2. One sufficient condition for yielding this assumption is when  $\forall k, j, \mathcal{W}_2(p_S(Z|Y = k), p_T(Z|Y = k)) \leq \mathcal{W}_2(p_S(Z|Y = k), p_T(Z|Y = j))$ . This last condition is typically achieved when conditionals between source and target of the same class are “sufficiently near” to each other.

**Assumption 4** (Conditional linear independence on T).  $\{p_T(Z|Y = k)\}_{i=k}^K$  are linearly independent implying  $\forall k, \nexists \alpha \in \left\{ \Delta_K | p_T(Z|Y = k) = \sum_{j=1, j \neq k}^K \alpha_j p_T(Z|Y = j) \right\}$ .

Assumption 4 is standard and seen in TarS to guarantee correctly estimating target proportions (Redko et al., 2019; Garg et al., 2020). It discards pathological cases like when  $\exists(i, j, k) \exists(a, b)$  s.t.  $a + b = 1, p_T(Z|Y = i) = a \times p_T(Z|Y = j) + b \times p_T(Z|Y = k)$ . It is milder than its alternative *A2Cons* in Zhang et al. (2013); Gong et al. (2016) which states linear independence of linear combinations of source and target conditionals, in  $\mathcal{X}$  respectively  $\mathcal{Z}$ .

## 3 THEORETICAL RESULTS

We present our theoretical results, with proofs in Appendix D, for *OSTAR* under our mild assumptions in Section 2.3. We first analyze in Section 3.1 the properties of the solution to our problem in (OT). Then, we show in Section 3.2 that given the learned components  $g, \phi$  and  $\mathbf{p}_N^Y$ , the target generalization error of a classifier in the hypothesis space  $\mathcal{H}$  can be upper-bounded by different terms including its risk on domain  $N$  and the Wasserstein-1 distance between marginals in  $N$  and  $T$ .

### 3.1 PROPERTIES OF THE SOLUTION TO THE OT ALIGNMENT PROBLEM

**Proposition 1** (Unicity and match). For any encoder  $g$  which defines  $\mathcal{Z}$  satisfying Assumption 1, 2, 3, 4, there is a unique solution  $(\phi, \mathbf{p}_N^Y)$  to (OT) and  $\phi_{\#}(p_S(Z|Y)) = p_T(Z|Y)$  and  $\mathbf{p}_N^Y = \mathbf{p}_T^Y$ .

Given an encoder  $g$  satisfying our assumptions, Proposition 1 shows two strong results. First, the solution to problem (OT) exists and is unique. Second, we prove that this solution defines a domain  $N$ , via the sample transformation and reweight operation defined in (1), where encoded conditionals and label proportions are equal to target ones. For comparison, Combes et al. (2020); Shui et al. (2021) recover target proportions only under *GLS* i.e. when conditionals are already matched, while we match both conditional and label proportions under the more general *GeTarS*.

### 3.2 CONTROLLED TARGET GENERALIZATION RISK

We now characterize the generalization properties of a classifier  $f_N$  trained on this domain  $N$  with a new general upper-bound. First, we introduce some notations; given an encoder  $g$  onto a latent space  $\mathcal{Z}$ , we define the risk of a classifier  $f$  as  $\epsilon_D^g(f) \triangleq \mathbb{E}_{\mathbf{z} \sim p_D(Z,Y)}[f(\mathbf{z}) \neq y]$  with  $D \in \{S, T, N\}$ .

**Theorem 1** (Target risk upper-bound). *Given a fixed encoder  $g$  defining a latent space  $\mathcal{Z}$ , two domains  $N$  and  $T$  satisfying cyclical monotonicity in  $\mathcal{Z}$ , assuming that we have  $\forall k, \mathbf{p}_N^{Y=k} > 0$ , then  $\forall f_N \in \mathcal{H}$  where  $\mathcal{H}$  is a set of  $M$ -Lipschitz continuous functions over  $\mathcal{Z}$ , we have*

$$\begin{aligned} \epsilon_T^g(f_N) \leq & \underbrace{\epsilon_N^g(f_N)}_{\text{Classification (C)}} + \underbrace{\frac{2M}{\min_{k=1}^K \mathbf{p}_N^{Y=k}} \mathcal{W}_1(p_N(Z), p_T(Z))}_{\text{Alignment (A)}} + \\ & 2M \underbrace{\left(1 + \frac{1}{\min_{k=1}^K \mathbf{p}_N^{Y=k}}\right) \mathcal{W}_1\left(\sum_{k=1}^K \mathbf{p}_T^{Y=k} p_T(Z|Y=k), \sum_{k=1}^K \mathbf{p}_N^{Y=k} p_T(Z|Y=k)\right)}_{\text{Label (L)}} \end{aligned} \quad (2)$$

We first analyze our upper-bound in (2). The target risk of  $f_N$  is controlled by three main terms: the first (C) is the risk of  $f_N$  on domain  $N$ , the second (A) is the Wasserstein-1 distance between latent marginals of domain  $N$  and  $T$ , the third term (L) measures a divergence between label distributions using, as a proxy, two ad-hoc marginal distributions. There are two other terms in (2): first, a Lipschitz-constant  $M$  that can be made small by implementing  $f_N$  as a NN with piece-wise linear activations and regularized weights; second the minimum of  $\mathbf{p}_N^Y$ , which says that generalization is harder when a target class is less represented. We learn *OSTAR*'s components to minimize the r.h.s of (2). Terms (C) and (A) can be explicitly minimized, respectively by training a classifier  $f_N$  on domain  $N$  and by learning  $(\phi, \mathbf{p}_N^Y)$  to match marginals of domains  $N$  and  $T$ . Term (L) is harder to control, yet its minimization is naturally handled by *OSTAR*. Indeed, term (L) is minimal when  $\mathbf{p}_N^Y = \mathbf{p}_T^Y$  under Assumption 4 per Redko et al. (2019). With *OSTAR*, this sufficient condition is guaranteed in Proposition 1 by the solution to problem (OT) under our mild assumptions in Section 2.3. This solution defines a domain  $N$  for which  $\mathbf{p}_N^Y = \mathbf{p}_T^Y$  and term (A) equals zero.

We now detail the originality of this result over existing Wasserstein-based generalization bounds. First, our upper-bound is general and can be explicitly minimized even when target labels are unknown unlike Shui et al. (2021) or Combes et al. (2020) which require knowledge of  $\mathbf{p}_T^Y$  or its perfect estimation. Combes et al. (2020) claims that correct estimation of  $\mathbf{p}_T^Y$  i.e. (L) close to zero, is achieved under *GLS* i.e.  $p_S(Z|Y) = p_T(Z|Y)$ . However, *GLS* is hardly guaranteed when the latent space is learned: a sufficient condition in Combes et al. (2020) requires knowing  $\mathbf{p}_T^Y$ , which is unrealistic in UDA. Second, our bound is simpler than the one in Rakotomamonjy et al. (2020): in particular, it removes several redundant terms which are unmeasurable due to unknown target labels.

## 4 IMPLEMENTATION

Our solution, detailed below, minimizes the generalization bound in (2) with *OSTAR*. It implements components  $f_N, g$  with NNs and the map  $\phi$  with a residual NN. Our solution jointly solves with pretrained representations (i) a classification problem to account for term (C) in (2) and (ii) an alignment problem (OT) to account for terms (A) and (L) in (2). Our pseudo-code is in Appendix E.

**Encoder initialization** Prior to learning  $\phi, \mathbf{p}_N^Y, f_N$ , we first learn the encoder  $g$  jointly with a source classifier  $f_S$  to yield a zero source classification loss via (3). With  $\mathcal{L}_{ce}$  the cross-entropy loss,

$$\min_{f_S, g} \mathcal{L}_c^g(f_S, S) \triangleq \min_{f_S, g} \frac{1}{n} \sum_{i=1}^n \mathcal{L}_{ce}(f_S \circ g(\mathbf{x}_S^{(i)}), y_S^{(i)}) \quad (3)$$

$g$  is then fixed to preserve the original problem structure. This initialization step helps enforce Assumption 1 and could alternatively be replaced by directly using a pretrained encoder if available.

**Joint alignment and classification** We then solve alternatively (i) a classification problem on domain  $N$  w.r.t.  $f_N$  to account for term (C) in (2) and (ii) the problem (OT) w.r.t.  $(\phi, \mathbf{p}_N^Y)$  to account for term (A). Term (L) in (2) is handled by matching  $\mathbf{p}_N^Y$  and  $\mathbf{p}_T^Y$  through the minimization of problem (OT).  $\mathbf{p}_N^Y$  is estimated with the confusion-based approach in Lipton et al. (2018), while  $\phi$  minimizes the Lagrangian relaxation, with hyperparameter  $\lambda_{OT}$ , of the constrained optimization problem (OT). The equality constraint in (OT) is measured through a Wasserstein distance as our bound in (2) is tailored to this distance, however, we are not restricted by this choice as other discrepancy measures are also possible. The objective based on (2) corresponds to the following optimization problem:

$$\begin{aligned} \min_{\phi, f_N} \mathcal{L}_c^g(f_N, N) + \lambda_{OT} \mathcal{L}_{OT}^g(\phi) + \mathcal{L}_{wd}^g(\phi, \mathbf{p}_N^Y) \\ \text{subject to } \mathbf{p}_N^Y = \arg \min_{\mathbf{p} \geq 0, \mathbf{p} \in \Delta_K} \frac{1}{2} \|\hat{\mathbf{p}}_T^Y - \hat{\mathbf{C}} \frac{\mathbf{p}}{\mathbf{p}_S^Y}\|_2^2 \end{aligned} \quad \text{[Label proportion estimation]} \quad (\text{CAL})$$

$$\text{where } \mathcal{L}_c^g(f_N, N) \triangleq \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{p}_N^{y_S^{(i)}}}{\mathbf{p}_S^{y_S^{(i)}}} \mathcal{L}_{ce}(f_N \circ \phi \circ g(\mathbf{x}_S^{(i)}), y_S^{(i)}) \quad \text{[Classification loss in } \mathcal{Z}_{\hat{N}}] \quad (4)$$

$$\text{and } \mathcal{L}_{OT}^g(\phi) \triangleq \sum_{k=1}^K \frac{1}{\#\{y_S^{(i)} = k\}_{i \in \{1 \dots n\}}} \sum_{\substack{y_S^{(i)} = k \\ i \in \{1 \dots n\}}} \|\phi(\mathbf{z}_S^{(i)}) - \mathbf{z}_S^{(i)}\|_2^2 \quad \text{[Objective function of (OT)]} \quad (5)$$

$$\text{and } \mathcal{L}_{wd}^g(\phi, \mathbf{p}_N^Y) \triangleq \sup_{\|v\|_L \leq 1} \frac{1}{n} \sum_{i=1}^n \frac{\mathbf{p}_N^{y_S^{(i)}}}{\mathbf{p}_S^{y_S^{(i)}}} v \circ \phi(\mathbf{z}_S^{(i)}) - \frac{1}{m} \sum_{j=1}^m v(\mathbf{z}_T^{(j)}) \quad \text{[Relaxed equality constraint in (OT)]} \quad (6)$$

$\hat{\mathbf{C}}$  is the confusion matrix of  $f_N$  on domain  $N$ ,  $\hat{\mathbf{p}}_T^Y \in \Delta_K$  is the target label proportions estimated with  $f_N$ .  $\mathcal{L}_c^g(f_N, N)$  (4) is the classification loss of  $f_N$  on  $\mathcal{Z}_{\hat{N}}$ , derived in Appendix D, which minimizes term (C) in (2). Note that samples in domain  $N$  are obtained by mapping source samples with  $\phi$ ; they are reweighted to account for label shift.  $\mathcal{L}_{OT}^g$  (5) defines the transport cost of  $\phi$  on  $\mathcal{Z}_{\hat{S}}$ . Implementing  $\phi$  with a ResNet performed better than standard MLPs, thus we minimize in practise the dynamical transport cost, better tailored to residual maps and used in de Bézenac et al. (2021); Karkar et al. (2020).  $\mathcal{L}_{wd}^g$  (6) is the empirical form of the dual Wasserstein-1 distance between  $p_N^\phi(Z)$  and  $p_T(Z)$  and seeks at enforcing the equality constraint in problem (OT) i.e. minimizing term (A). OSTAR’s assumptions also guarantee that term (L) is small at the optimum.

**Improve target discriminativity** OSTAR solves a transfer problem with pretrained representations, but does not directly operate on the discriminativity of target representations. This may harm performance when the encoder is not well suited for the target. Domain-invariant methods are less prone to this issue as they handle jointly these two problems. In our approach, target discriminativity is assessed by the minimum of term (C) in (2) on domain  $N$ , built from the solution to problem (OT), for which conditionals and proportions are matched with target ones. This value depends on the encoder and may be high when representations are fixed. Nevertheless, it can be reduced by better enforcing class separability for target representations. To achieve this goal, we propose an extension of (CAL) using Information Maximization (IM), not considered in existing domain-invariant GeTarS methods. IM, used for UDA in Liang et al. (2020), refines the decision boundaries of  $f_N$  with two terms on target samples.  $\mathcal{L}_{ent}^g(f_N, T)$  (7) is the conditional entropy of  $f_N$  which favors low-density separation between classes. Denoting  $\delta_k(\cdot)$  the  $k$ -th component of the softmax function,

$$\mathcal{L}_{ent}^g(f_N, T) = \sum_{i=1}^m \sum_{k=1}^K \delta_k(f_N \circ g(\mathbf{x}_T^{(i)})) \log(\delta_k(f_N \circ g(\mathbf{x}_T^{(i)}))) \quad (7)$$

$\mathcal{L}_{div}^g(f_N, T)$  (8) promotes diversity by regularizing the average output of  $f_N \circ g$  to be uniform on  $T$ . It avoids predictions from collapsing to the same class thus softens the effect of conditional entropy.

$$\mathcal{L}_{div}^g(f_N, T) = \sum_{k=1}^K \hat{p}_k \log \hat{p}_k = D_{KL}(\hat{p}, \frac{1}{K} \mathbf{1}_K) - \log K; \hat{p} = \mathbb{E}_{\mathbf{x}_T \in \mathcal{X}_T} [\delta(f_N \circ g(\mathbf{x}_T))] \quad (8)$$

Our variant introduces two additional steps in the learning process. First, the latent space is fixed and we optimize  $f_N$  with IM in (SS). Then, we optimize the representations in (SSg) while avoiding modifying source representations by including  $\mathcal{L}_c^g(f_S, S)$  (3) with a fixed source classifier  $f_S$ .

$$\min_{f_N} \mathcal{L}_c^g(f_N, N) + \mathcal{L}_{ent}^g(f_N, T) + \mathcal{L}_{div}^g(f_N, T) \quad (\text{SS})$$

$$\min_{f_N, g} \mathcal{L}_c^g(f_N, N) + \mathcal{L}_{ent}^g(f_N, T) + \mathcal{L}_{div}^g(f_N, T) + \mathcal{L}_c^g(f_S, S) \quad (\text{SSg})$$

## 5 EXPERIMENTAL RESULTS

We now present our experimental results on several UDA problems under GeTarS and show that OSTAR outperforms recent SOTA baselines. The GeTarS assumption is particularly relevant on our datasets as encoded conditional distributions do not initially match as seen in Appendix Figure 3.

**Experimental setup** We consider: (i) an academic benchmark `Digits` with two adaptation problems between `USPS` (Hull, 1994) and `MNIST` (LeCun et al., 1998), (ii) a synthetic to real images adaptation benchmark `VisDA12` (Peng et al., 2017) and (iii) two object categorizations problems `Office31` (Saenko et al., 2010), `OfficeHome` (Venkateswara et al., 2017) with respectively six and twelve adaptation problems. The original datasets have fairly balanced classes, thus source and target label distributions are similar. This is why we subsample our datasets to make label proportions dissimilar across domains as detailed in Appendix Table 6; prefix "s" corresponds to subsampled datasets. For `Digits`, we subsample the target domain and investigate three settings - balanced, mild and high as Rakotomamonjy et al. (2020). For other datasets, we modify the source domain by considering 30% of the samples coming from the first half of their classes as Combes et al. (2020). We compare OSTAR against a `Source` baseline trained using only source samples without adaptation and to various UDA methods: two `CoS` methods and three recent SOTA GeTarS models. Chosen UDA methods learn invariant representations with reweighting in GeTarS models or without reweighting in other baselines. The two `CoS` baselines are `DANN` (Ganin et al., 2016) which approaches  $\mathcal{H}$ -divergence and  $\text{WD}_{\beta=0}$  (Shen et al., 2018) which computes Wasserstein distance. The GeTarS baselines are (Wu et al., 2019; Rakotomamonjy et al., 2020; Combes et al., 2020). We use Wasserstein distance to learn invariant representations such that differences lie mostly on the strategy to account for label shift. Wu et al. (2019), denoted  $\text{WD}_{\beta}$ , performs asymmetric alignment with parameter  $\beta$ , for which we test different values ( $\beta \in \{1, 2\}$ ). The remaining models, `MARSc`, `MARSG` (Rakotomamonjy et al., 2020) and `IW-WD` (Combes et al., 2020), estimate target proportions respectively with optimal assignment or with the estimator in Lipton et al. (2018) also used in OSTAR. We report `DI-Oracle`, an oracle which learns invariant representations with Wasserstein distance and makes use of true class-ratios. All baselines are reimplemented for a fair comparison with the same NN architectures detailed in Appendix F and without runtime difference for OSTAR.

**Results** We report in Table 1 mean and standard deviations for balanced accuracy over 10 runs. It corresponds to the average recall obtained on each class and is suited for imbalanced problems (Broderson et al., 2010). Additionally, we report in Figure 4 in the Appendix due to space limits, the  $\ell_1$  error between true and estimated target proportions for GeTarS baselines. First, we note that low estimation error of  $p_T^Y$  is correlated to high accuracy for all models and that `DI-Oracle`, which uses true class-ratios, upper-bounds domain-invariant approaches. This shows the importance of correctly estimating  $p_T^Y$ . Second, we note that OSTAR clearly outperforms or equals the baselines on both balanced accuracy and proportion estimation. It even improves `DI-Oracle` for balanced accuracy despite not using true class-ratios. This (i) shows the benefits of not constraining domain-invariance which may degrade target discriminativity especially under label estimation errors (ii) validates our theoretical results which show that OSTAR controls the target risk and recovers target proportions. OSTAR also matches conditionals, as visualized in Appendix Figure 3.

**Ablation studies** We perform two ablation studies. First, we evaluate the effect of IM in Table 2 and Appendix Table 4. We successively add (SS) (column 3) then (SSg) (column 4) to (CAL) (column 2). Without IM (i.e. with (CAL)), OSTAR improves the `Source` baseline (column 2 in Table 1), as it minimizes both conditional and label shifts. Adding (SS) then (SSg) improves OSTAR and challenges all of our baselines. We also evaluated the effect of adding IM to our baselines in Appendix Table 5, even if this is not part of their original work. IM generally improves performance,



Table 1: Balanced accuracy ( $\uparrow$ ) over 10 runs. The best performing model is indicated in **bold**.

Setting	Source	DANN	$WD_{\beta=0}$	$WD_{\beta=1}$	$WD_{\beta=2}$	MARSG	MARSc	IW-WD	OSTAR	DI-Oracle
sDigits - MNIST→USPS										
balanced	76.89 ± 5.8	93.97 ± 0.9	94.42 ± 0.6	82.05 ± 6.3	75.14 ± 6.7	94.19 ± 1.8	<b>96.44 ± 0.3</b>	<b>96.10 ± 0.3</b>	<b>96.91 ± 0.3</b>	96.95 ± 0.2
mild	81.97 ± 3.7	93.23 ± 1.2	91.52 ± 0.6	84.07 ± 5.0	78.29 ± 7.2	94.78 ± 1.0	95.18 ± 0.9	94.72 ± 0.4	<b>96.18 ± 1.0</b>	97.22 ± 0.2
high	84.34 ± 2.3	93.03 ± 1.0	91.14 ± 0.7	85.15 ± 2.9	78.90 ± 3.8	94.50 ± 1.3	95.07 ± 0.6	94.60 ± 0.8	<b>96.06 ± 0.6</b>	96.87 ± 0.4
sDigits - USPS→MNIST										
balanced	63.01 ± 6.5	87.64 ± 1.7	90.84 ± 1.3	83.54 ± 3.0	77.00 ± 7.6	90.16 ± 2.5	93.37 ± 2.5	95.68 ± 0.6	<b>98.11 ± 0.2</b>	96.84 ± 0.2
mild	62.81 ± 2.3	87.00 ± 2.1	88.07 ± 1.2	77.83 ± 5.4	79.41 ± 7.4	89.93 ± 2.4	93.20 ± 2.8	92.73 ± 1.5	<b>97.44 ± 0.5</b>	95.75 ± 0.3
high	64.04 ± 5.4	86.37 ± 1.4	87.04 ± 1.4	79.17 ± 6.0	74.47 ± 7.4	88.24 ± 3.3	91.54 ± 0.9	90.81 ± 1.5	<b>97.08 ± 0.6</b>	95.87 ± 0.3
VisDA12										
VisDA	48.63 ± 1.0	53.72 ± 0.9	57.40 ± 1.1	47.56 ± 0.8	36.21 ± 1.8	55.62 ± 1.6	55.33 ± 0.8	51.88 ± 1.6	<b>59.24 ± 0.5</b>	57.61 ± 0.3
sVisDA	42.46 ± 1.4	47.57 ± 0.9	47.32 ± 1.4	41.48 ± 1.6	31.83 ± 3.0	55.00 ± 1.9	51.86 ± 2.0	50.65 ± 1.5	<b>58.84 ± 1.0</b>	55.77 ± 1.1
soffice31										
sA-D	80.71 ± 0.5	82.39 ± 0.4	81.76 ± 0.4	75.98 ± 1.2	68.64 ± 2.4	<b>84.54 ± 1.0</b>	<b>84.10 ± 0.8</b>	81.83 ± 0.5	<b>84.17 ± 0.7</b>	87.74 ± 0.6
sD-W	89.08 ± 0.4	88.70 ± 0.2	88.98 ± 0.2	88.53 ± 0.2	88.97 ± 0.1	91.03 ± 0.4	90.76 ± 0.4	88.17 ± 0.3	<b>94.13 ± 0.2</b>	91.31 ± 0.2
sW-A	58.91 ± 0.2	58.87 ± 0.1	59.18 ± 0.2	60.70 ± 0.3	60.95 ± 0.2	63.94 ± 0.1	63.80 ± 0.3	60.25 ± 0.2	<b>69.99 ± 0.1</b>	63.92 ± 0.2
sW-D	95.64 ± 0.2	97.26 ± 0.3	97.13 ± 0.3	95.99 ± 0.3	95.57 ± 0.5	97.96 ± 0.1	<b>98.16 ± 0.2</b>	97.53 ± 0.2	<b>98.47 ± 0.2</b>	98.35 ± 0.0
sD-A	53.41 ± 0.9	57.45 ± 0.2	57.81 ± 0.2	58.24 ± 0.2	58.61 ± 0.3	62.12 ± 0.2	62.13 ± 0.4	60.03 ± 0.2	<b>65.00 ± 0.5</b>	62.57 ± 0.3
sA-W	69.23 ± 0.5	72.09 ± 0.5	72.60 ± 0.3	65.94 ± 0.9	61.64 ± 7.2	81.60 ± 0.5	81.05 ± 0.7	75.84 ± 0.7	<b>83.91 ± 0.5</b>	82.51 ± 0.5
sofficeHome										
sA-C	44.44 ± 0.3	46.08 ± 0.3	41.74 ± 1.7	40.90 ± 0.8	39.22 ± 1.1	47.19 ± 0.3	46.94 ± 0.2	45.29 ± 0.1	<b>48.43 ± 0.2</b>	48.09 ± 0.2
sA-P	58.96 ± 0.3	59.96 ± 0.2	54.67 ± 1.8	52.18 ± 2.3	46.29 ± 1.4	62.17 ± 0.2	61.97 ± 0.2	59.46 ± 0.3	<b>69.52 ± 0.4</b>	63.59 ± 0.2
sA-R	67.10 ± 0.2	67.42 ± 0.2	65.40 ± 0.6	62.52 ± 1.7	60.51 ± 1.9	68.66 ± 0.3	68.62 ± 0.3	67.76 ± 0.2	<b>73.29 ± 0.3</b>	69.85 ± 0.1
sC-A	35.54 ± 2.3	35.47 ± 1.7	37.34 ± 2.0	36.81 ± 1.5	33.15 ± 2.3	<b>46.03 ± 0.2</b>	<b>46.10 ± 0.2</b>	44.18 ± 0.1	<b>46.47 ± 0.3</b>	46.94 ± 0.2
sC-P	52.48 ± 2.1	50.56 ± 0.9	53.53 ± 0.1	49.96 ± 1.4	44.67 ± 1.5	59.82 ± 0.1	59.82 ± 0.1	58.67 ± 0.1	<b>63.37 ± 0.1</b>	60.14 ± 0.1
sC-R	54.99 ± 1.7	54.22 ± 0.8	54.69 ± 0.5	51.34 ± 2.5	45.16 ± 3.5	62.69 ± 0.1	62.41 ± 0.1	60.74 ± 0.2	<b>63.12 ± 0.2</b>	62.80 ± 0.2
sP-A	38.10 ± 4.5	36.36 ± 3.3	48.24 ± 0.2	47.24 ± 0.3	48.20 ± 0.3	47.78 ± 0.3	46.10 ± 0.9	45.68 ± 0.3	<b>50.84 ± 0.3</b>	50.11 ± 0.4
sP-C	34.16 ± 4.6	33.00 ± 1.9	39.10 ± 0.2	40.36 ± 0.4	37.41 ± 0.3	42.41 ± 0.3	41.92 ± 0.3	38.22 ± 0.2	<b>44.15 ± 0.3</b>	43.10 ± 0.2
sP-R	66.28 ± 3.4	59.19 ± 0.8	70.01 ± 0.1	68.78 ± 0.2	66.61 ± 0.3	70.00 ± 0.4	69.37 ± 0.9	69.43 ± 0.4	<b>73.95 ± 0.3</b>	71.26 ± 0.4
sR-P	66.67 ± 5.4	70.97 ± 0.6	73.47 ± 0.3	72.66 ± 0.7	71.76 ± 0.7	72.62 ± 0.9	72.72 ± 1.1	72.90 ± 0.7	<b>75.58 ± 0.6</b>	74.17 ± 0.7
sR-A	48.59 ± 6.7	51.90 ± 1.1	<b>56.97 ± 0.3</b>	<b>57.02 ± 0.5</b>	55.38 ± 1.1	54.02 ± 0.7	53.37 ± 1.3	53.44 ± 0.6	<b>56.28 ± 0.5</b>	57.68 ± 0.6
sR-C	39.36 ± 2.5	45.33 ± 0.8	46.47 ± 0.4	47.11 ± 0.5	45.38 ± 1.3	45.81 ± 1.2	45.30 ± 1.2	42.66 ± 1.0	<b>49.07 ± 0.9</b>	47.86 ± 0.3

marginally on Digits, more significantly on VisDA and Office. However, the performance remains below the ones of OSTAR. Second, we evaluate in Table 3 the effect of our transport cost in problem (CAL). Proposition 1 shows that the OT formulation guarantees recovering target proportions and matching conditionals. We consider MNIST→USPS under various shifts (Table 3a) and initialization gains i.e. the standard deviation of the weights of the NN (Table 3b). In Table 3a, we note that  $\lambda_{OT} \neq 0$  in problem (CAL) improves balanced accuracy (left) and  $\ell_1$  estimation error (middle) over  $\lambda_{OT} = 0$  over all shifts, especially high ones. This improvement is correlated with lower mean and standard deviation of the normalized transport cost per sample (right). We observe the same trends when changing initialization gains in Table 3b. This comforts our theoretical results and shows the advantages of OT regularization biases for performance and stability.

Table 2: Effect of semi-supervised learning on balanced accuracy ( $\uparrow$ ). The best model is in **bold**.

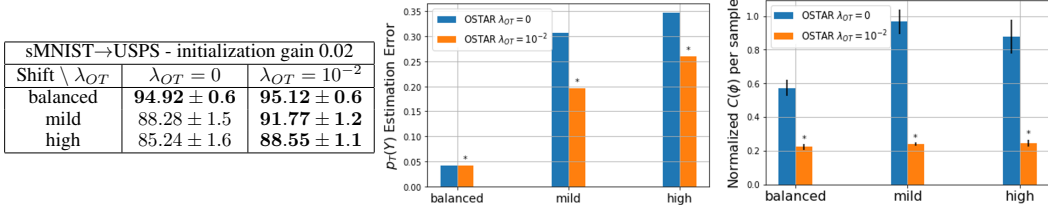
Setting \ Objective	(CAL)	+ (SS)	+ (SSg)
sMNIST→USPS			
balanced	95.12 ± 0.6	96.68 ± 0.1	<b>96.91 ± 0.3</b>
mild	91.77 ± 1.2	95.39 ± 1.4	<b>96.18 ± 1.0</b>
high	88.55 ± 1.1	95.70 ± 0.8	<b>96.06 ± 0.6</b>
VisDA12			
VisDA	50.37 ± 0.6	52.54 ± 0.3	<b>59.24 ± 0.5</b>
sVisDA	49.05 ± 0.9	53.37 ± 0.6	<b>58.84 ± 1.0</b>

## 6 RELATED WORK

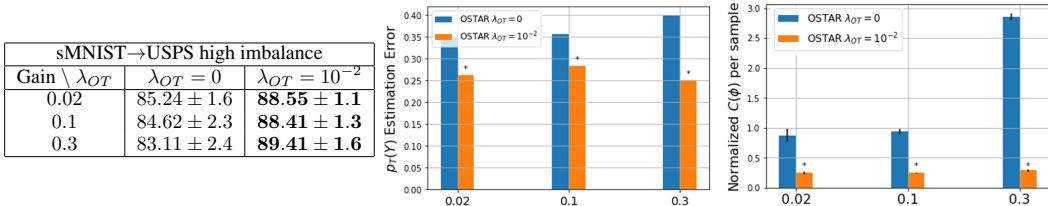
**UDA** Existing approaches train a classifier using source labels while handling distribution shifts. We review two main approaches: the first learns invariant representations and the second learns a map between fixed samples from source and target domains. While domain-invariant approaches are SOTA, mapping-based approaches avoid structure issues posed by the invariance constraint which may degrade target discriminativity (Liu et al., 2019; Chen et al., 2019). For CoS, domain-invariant methods directly match marginal distributions in latent space (Ganin et al., 2016; Shen et al., 2018; Long et al., 2015), while mapping-based approaches learn a mapping in input space (Courty et al., 2017b; Hoffman et al., 2018). For TarS and GeTarS, label shift requires estimating  $p_T^Y$  to reweight source samples by estimated class-ratios (Zhao et al., 2019). An alternative is to use a fixed weight (Wu et al., 2019). When conditionals are unchanged i.e. TarS,  $p_T^Y$  can be recovered without needs for alignment (Lipton et al., 2018; Redko et al., 2019). Under GeTarS, there is the additional

Table 3: Effect of transport cost on balanced accuracy  $\uparrow$  (left),  $\ell_1$  estimation error  $\downarrow$  (middle) and normalized transport cost  $\mathcal{L}_{OT}^g/n \downarrow$  (right). The best model for accuracy is in **bold** and with a “\*\*”.

(a) Under varying shifts of sMNIST $\rightarrow$ USPS at fixed initialization gain



(b) Under varying initialization gains (i.e. stdev of the NN’s weights) on highly imbalanced sMNIST $\rightarrow$ USPS



difficulty of matching conditionals. The SOTA methods for GeTarS are domain-invariant with sample reweighting (Rakotomamonjy et al., 2020; Combes et al., 2020; Gong et al., 2016). An earlier mapping-based approach was proposed in Zhang et al. (2013) to align conditionals, also under reweighting. Yet, it is not flexible, operates on high dimensional input spaces and does not scale up to large label spaces as it considers a linear map for each pair of conditionals. Estimators used in GeTarS are confusion-based in Combes et al. (2020); Shui et al. (2021); derived from optimal assignment in Rakotomamonjy et al. (2020) or from the minimization of a reweighted MMD between marginals in Zhang et al. (2013); Gong et al. (2016). OSTAR is a new mapping-based approach for GeTarS with improvements over these approaches as detailed in this paper. The improvements are both practical (flexibility, scalability, stability) and theoretical. CTC (Gong et al., 2016) also operates on representations, yet OSTAR simplifies learning by clearly separating alignment and encoding operations, intertwined in CTC as both encoder and maps are trained to align.

**OT approaches for UDA** They usually apply linear programming (LP) to compute a transport plan between empirical distributions. LP is used in CoS to align source and target joint distributions (Courty et al., 2017a; Damodaran et al., 2018) or to compute a barycentric mapping in input space (Courty et al., 2017b); and in TarS to compute a reweighted Wasserstein distance between source and target marginals whose minimum yields target proportions estimates (Redko et al., 2019). An alternative to LP is to minimize a dual Wasserstein-1 distance with adversarial training; this is done for CoS in Shen et al. (2018) and, more recently, for GeTarS in Rakotomamonjy et al. (2020); Shui et al. (2021). Rakotomamonjy et al. (2020) formulates two separate OT problems for class-ratio estimation and conditional alignment and Shui et al. (2021) is the OT extension of Combes et al. (2020) to multi-source UDA. OSTAR is the first mapping-based OT approach for GeTarS and has several benefits: (i) OSTAR keeps representation separate which avoids deteriorating target discriminativity, (ii) it solves a single OT problem unlike Rakotomamonjy et al. (2020), (iii) it implements the OT map with a NN which generalizes beyond training samples and scales up with the number of samples unlike barycentric OT maps; (iv) it adds encoding which improves efficiency of matching by improving discriminativity and reducing dimensionality, unlike Courty et al. (2017b).

## 7 CONCLUSION

We introduce OSTAR, a new general approach to align pretrained representations under GeTarS, which does not constrain representation invariance. OSTAR learns a flexible and scalable map between conditional distributions in the latent space. This map, implemented as a ResNet, solves an OT matching problem with native regularization biases. Our approach provides strong generalization guarantees under mild assumptions as it explicitly minimizes a new upper-bound to the target risk. Experimentally, it challenges recent invariant GeTarS methods on several UDA benchmarks.

---

**Acknowledgement** This work is partially funded by RAIMO ANR-20-CHIA-0021-01 and OAT-MIL ANR-17-CE23-0012 Projects of the French National Research Agency (ANR).

## REFERENCES

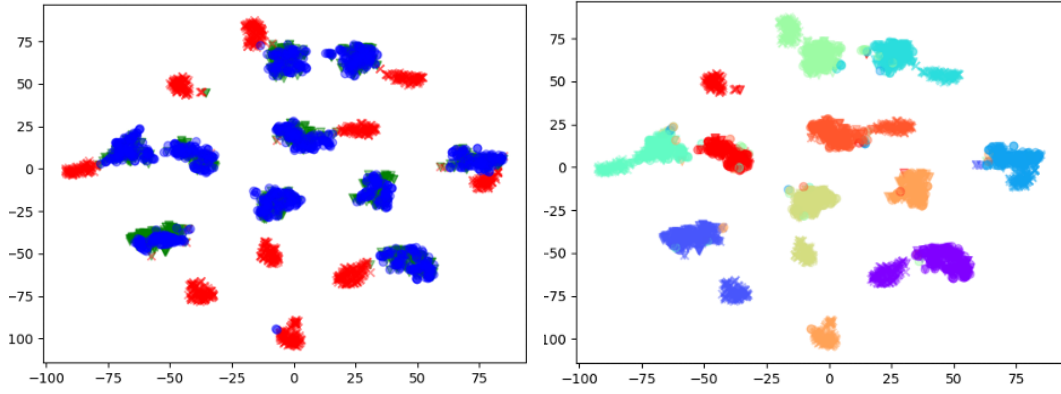
- Y. Brenier. Polar factorization and monotone rearrangement of vector-valued functions. 1991.
- K. H. Brodersen, C. S. Ong, K. E. Stephan, and J. M. Buhmann. The balanced accuracy and its posterior distribution. In *20th ICPR*, pp. 3121–3124, 2010.
- Olivier Chapelle, Bernhard Schölkopf, and Alexander Zien. *Semi-Supervised Learning*. The MIT Press, 1st edition, 2010. ISBN 0262514125.
- Xinyang Chen, Sinan Wang, Mingsheng Long, and Jianmin Wang. Transferability vs. discriminability: Batch spectral penalization for adversarial domain adaptation. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 1081–1090. PMLR, 09–15 Jun 2019. URL <https://proceedings.mlr.press/v97/chen19i.html>.
- Remi Tachet des Combes, Han Zhao, Yu-Xiang Wang, and Geoff Gordon. Domain adaptation with conditional distribution matching and generalized label shift. In *Advances in Neural Information Processing Systems*, 2020.
- Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems 30*, pp. 3730–3739. Curran Associates, Inc., 2017a.
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, 2017b. doi: 10.1109/TPAMI.2016.2615921.
- Bharath Bhushan Damodaran, Benjamin Kellenberger, Remi Flamary, Devis Tuia, and Nicolas Courty. DeepJDOT: Deep Joint Distribution Optimal Transport for Unsupervised Domain Adaptation. In *ECCV*, volume 11208 of *LNCS*, pp. 467–483, Munich, Germany, September 2018. Springer.
- Emmanuel de Bézenac, Ibrahim Ayed, and Patrick Gallinari. Cyclegan through the lens of (dynamical) optimal transport. In *ECML-PKDD*, 2021.
- Yaroslav Ganin, Evgeniya Ustinova, Hana Ajakan, Pascal Germain, Hugo Larochelle, François Laviolette, Mario March, and Victor Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.
- Saurabh Garg, Yifan Wu, Sivaraman Balakrishnan, and Zachary Lipton. A unified view of label shift estimation. In *NeurIPS*, volume 33, pp. 3290–3300. Curran Associates, Inc., 2020.
- Mingming Gong, Kun Zhang, Tongliang Liu, Dacheng Tao, Clark Glymour, and Bernhard Schölkopf. Domain adaptation with conditional transferable components. In *Proceedings of The 33rd ICML*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2839–2848, New York, New York, USA, 20–22 Jun 2016. PMLR.
- Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *NeurIPS*, pp. 5767–5777. Curran Associates, 2017.
- Judy Hoffman, Eric Tzeng, Taesung Park, Jun-Yan Zhu, Phillip Isola, Kate Saenko, Alexei Efros, and Trevor Darrell. CyCADA: Cycle-consistent adversarial domain adaptation. In *Proceedings of the 35th ICML*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1989–1998, Stockholm, Sweden, 10–15 Jul 2018. PMLR.
- J. J. Hull. A database for handwritten text recognition research. *IEEE TPAMI*, 16(5):550–554, May 1994. ISSN 0162-8828. doi: 10.1109/34.291440.
- Skander Karkar, Ibrahim Ayed, Emmanuel de Bezenac, and Patrick Gallinari. A Principle of Least Action for the Training of Neural Networks. In *ECML PKDD*, Ghent, Belgium, September 2020.

- 
- Yann LeCun, Leon Bottou, Yoshua Bengio, and P Haffner. Gradient-based learning applied to document recognition. 1998.
- Jian Liang, Dapeng Hu, and Jiashi Feng. Do we really need to access the source data? Source hypothesis transfer for unsupervised domain adaptation. In Hal Daumé III and Aarti Singh (eds.), *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6028–6039. PMLR, 13–18 Jul 2020. URL <http://proceedings.mlr.press/v119/liang20a.html>.
- Zachary C. Lipton, Yu-Xiang Wang, and Alexander J. Smola. Detecting and correcting for label shift with black box predictors. In Jennifer G. Dy and Andreas Krause (eds.), *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pp. 3128–3136. PMLR, 2018. URL <http://proceedings.mlr.press/v80/lipton18a.html>.
- Hong Liu, Mingsheng Long, Jianmin Wang, and Michael Jordan. Transferable adversarial training: A general approach to adapting deep classifiers. In *Proceedings of 36th ICML*, volume 97 of *Proceedings of ML Research*, pp. 4013–4022. PMLR, 09-15 Jun 2019.
- Mingsheng Long, Yue Cao, Jianmin Wang, and Michael I. Jordan. Learning transferable features with deep adaptation networks. In *Proceedings of the 32nd International Conference on Machine Learning, ICML'15*, pp. 97–105. JMLR.org, 2015.
- S. J. Pan and Q. Yang. A survey on transfer learning. *TKDE*, Oct 2010. doi: 10.1109/TKDE.2009.191.
- 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.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11 (5-6):355–602, 2019.
- N. Pustelnik, C. Chaux, and J. Pesquet. Parallel proximal algorithm for image restoration using hybrid regularization. *IEEE TIP*, 20(9):2450–2462, 2011.
- Alain Rakotomamonjy, Rémi Flamary, Gilles Gasso, Mokhtar Z. Alaya, Maxime Berar, and Nicolas Courty. Match and reweight strategy for generalized target shift. *Arxiv 2006.08161*, 2020.
- Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. Optimal transport for multi-source domain adaptation under target shift. volume 89 of *Proceedings of Machine Learning Research*, pp. 849–858. PMLR, 16–18 Apr 2019.
- Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In Kostas Daniilidis, Petros Maragos, and Nikos Paragios (eds.), *Computer Vision – ECCV 2010*, pp. 213–226, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- Jian Shen, Yanru Qu, Weinan Zhang, and Yong Yu. Wasserstein distance guided representation learning for domain adaptation. In *AAAI-18, 30th IAAI-18, 8th AAAI Symposium on EAAI-18, New Orleans, USA, February 2-7*, pp. 4058–4065. AAAI Press, 2018.
- Changjian Shui, Zijian Li, Jiaqi Li, Christian Gagné, Charles X Ling, and Boyu Wang. Aggregating from multiple target-shifted sources. In Marina Meila and Tong Zhang (eds.), *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9638–9648. PMLR, 18–24 Jul 2021. URL <http://proceedings.mlr.press/v139/shui21a.html>.
- Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *(IEEE) CVPR*, 2017.
- C Villani. *Optimal transport – Old and new*, volume 338, pp. xxii+973. 01 2008. doi: 10.1007/978-3-540-71050-9.

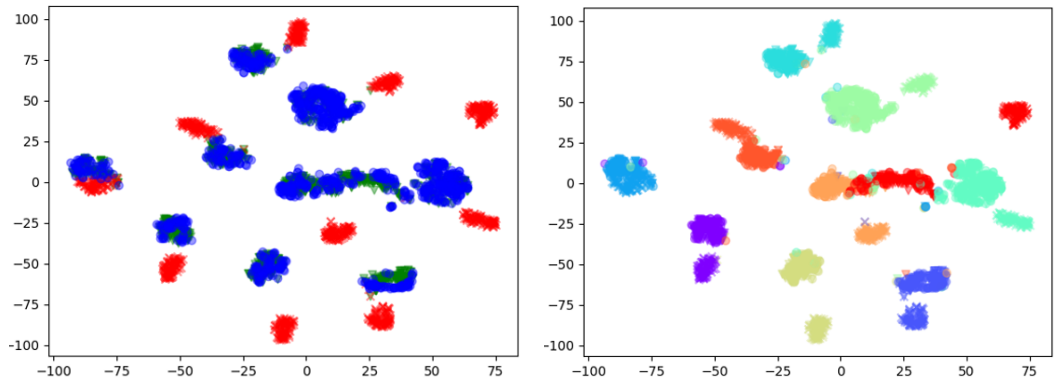
- 
- Yifan Wu, Ezra Winston, Divyansh Kaushik, and Zachary Lipton. Domain adaptation with asymmetrically-relaxed distribution alignment. In *ICML*, volume 97 of *Proceedings of ML Research*, pp. 6872–6881, Long Beach, CA, USA, 09–15 Jun 2019. PMLR.
- Kun Zhang, Bernhard Schölkopf, Krikamol Muandet, and Zhikun Wang. Domain adaptation under target and conditional shift. In *Proceedings of the 30th ICML*, volume 28 of *Proceedings of Machine Learning Research*, pp. 819–827, Atlanta, Georgia, USA, 17–19 Jun 2013.
- Han Zhao, Remi Tachet Des Combes, Kun Zhang, and Geoffrey Gordon. On learning invariant representations for domain adaptation. In *Proceedings of the 36th ICML*, volume 97 of *Proceedings of Machine Learning Research*, pp. 7523–7532. PMLR, 09–15 Jun 2019.

## A ADDITIONAL VISUALISATION

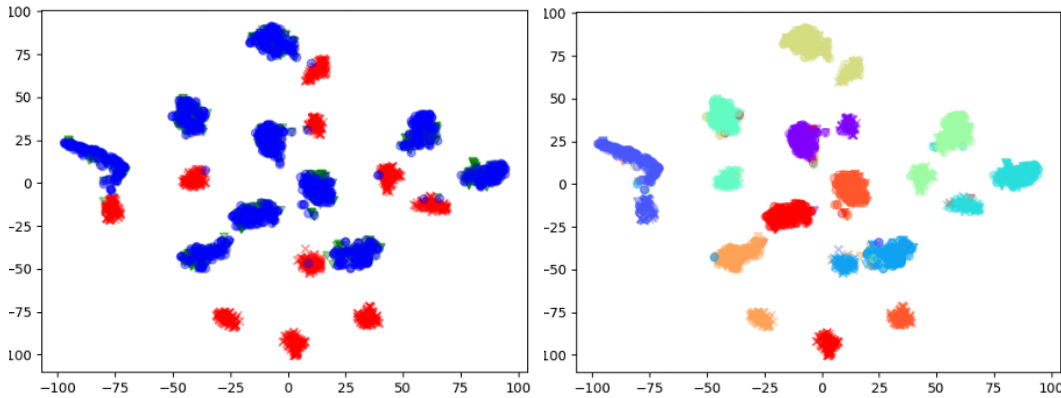
We visualize how OSTAR maps source representations onto target ones under GeTarS. We note that OSTAR (i) maps source conditionals to target ones (blue and green points are matched c.f. left), (ii) matches conditionals of the same class ("v" and "o" of the same colour are matched c.f. right).



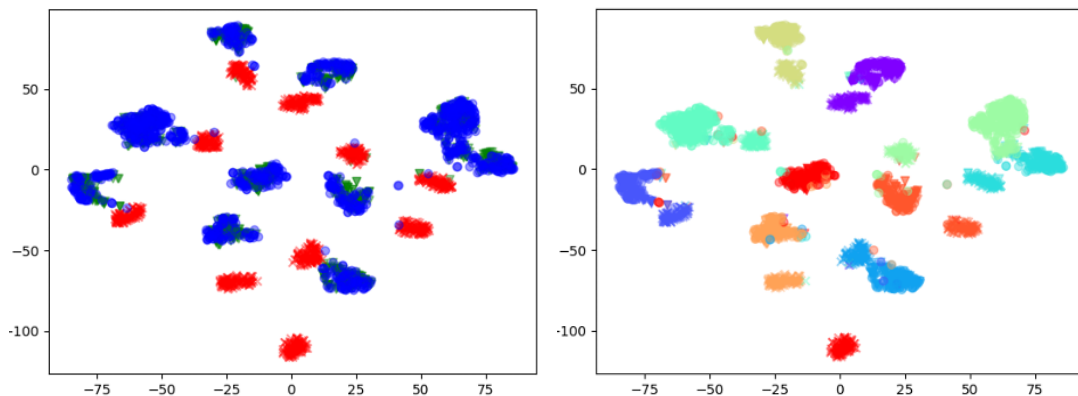
(a) sMNIST→USPS balanced



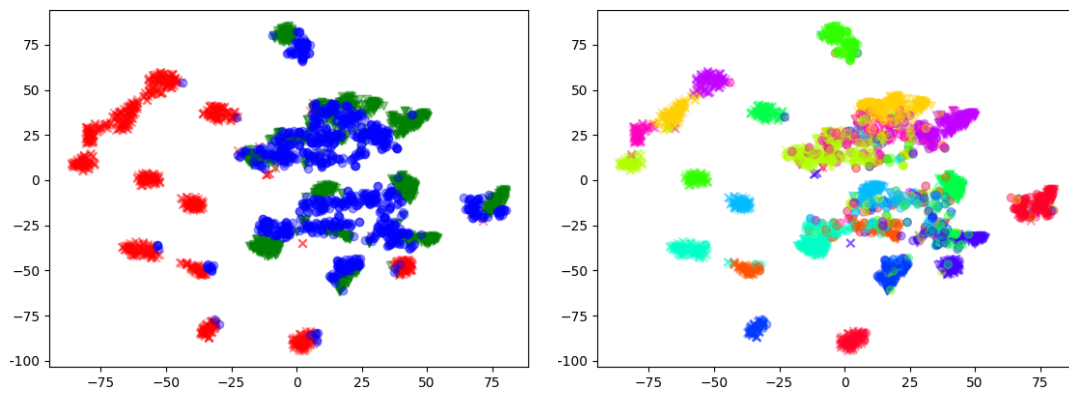
(b) sMNIST→USPS high



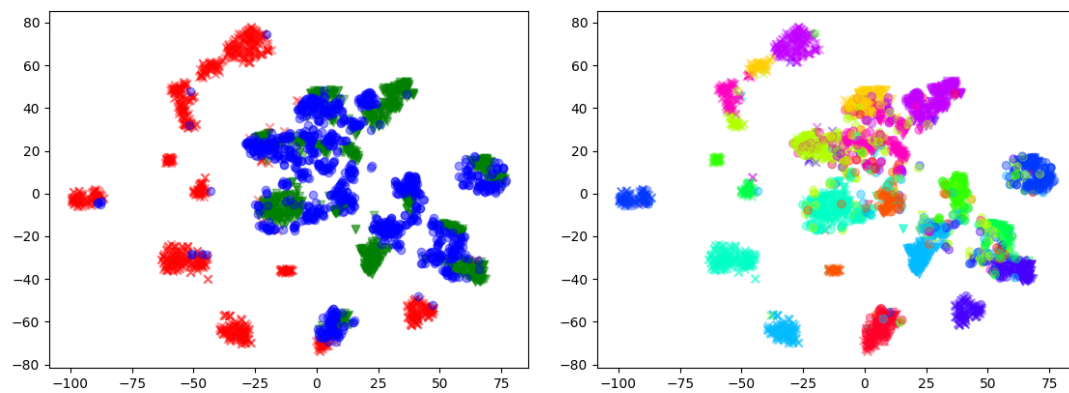
(c) sUSPS→MNIST balanced



(a) sUSPS→MNIST high



(b) VisDA



(c) sVisDA

Figure 3: t-SNE feature visualizations for OSTAR on various datasets and label imbalance. Crosses "x" denote source samples, circles "o" target samples and triangles "v" transported source samples. On the left, source samples are red, target samples blue and transported source samples green. On the right, samples from the same class have the same colour regardless of domain.

## B ADDITIONAL RESULTS

**Target label proportion estimation** We report the  $\ell_1$  estimation errors for our GeTarS methods.

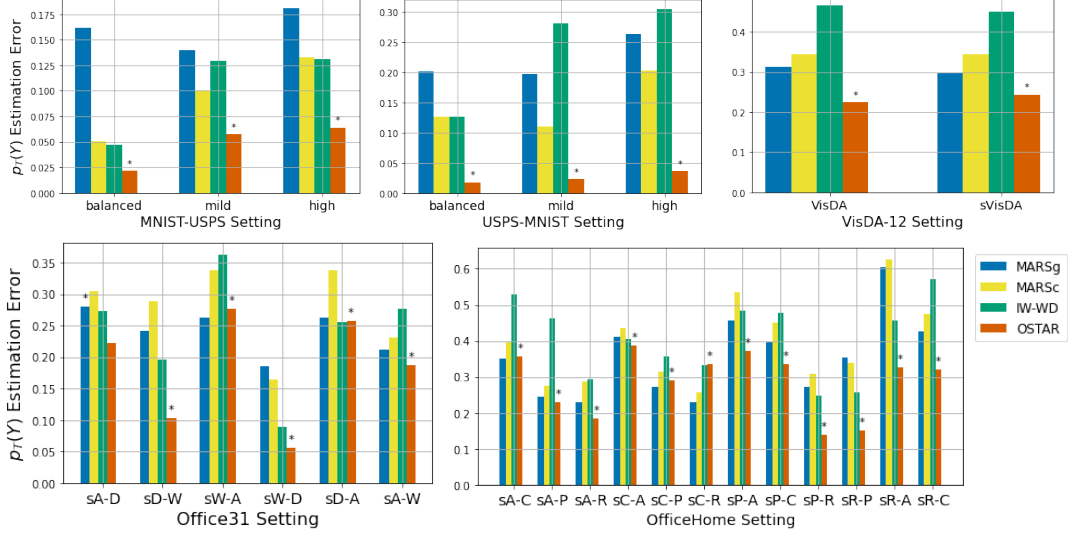


Figure 4:  $\ell_1$  estimation error of  $p_T^Y$  ( $\downarrow$ ). The best model for balanced accuracy is indicated with ”\*”.

**Ablation study for semi-supervised learning** We report additional results for this ablation study.

Table 4: Semi-supervised learning for OSTAR and balanced accuracy ( $\uparrow$ ). Best results are in **bold**.

Setting \ Objective	(CAL)	+ (SS)	+ (SSg)
sUSPS→MNIST			
balanced	88.19 ± 1.1	97.16 ± 0.3	<b>98.11 ± 0.2</b>
mild	88.34 ± 1.3	96.34 ± 0.2	<b>97.44 ± 0.5</b>
high	84.87 ± 2.3	95.61 ± 0.4	<b>97.08 ± 0.6</b>
sOffice31			
sA-D	81.52 ± 0.7	83.18 ± 0.2	<b>84.17 ± 0.7</b>
sD-W	89.94 ± 0.8	89.50 ± 0.8	<b>94.13 ± 0.2</b>
sW-A	59.62 ± 0.6	60.06 ± 0.4	<b>69.99 ± 0.1</b>
sW-D	96.39 ± 0.6	97.44 ± 0.2	<b>98.47 ± 0.2</b>
sD-A	54.38 ± 1.1	56.58 ± 0.6	<b>65.00 ± 0.5</b>
sA-W	75.30 ± 1.0	81.32 ± 0.8	<b>83.91 ± 0.5</b>

Table 5: Semi-supervised learning for MARSc and IW-WD GeTarS baselines and balanced accuracy ( $\uparrow$ ) on several representative datasets. Best results are in **bold**.

Setting \ Model	MARSc	MARSc + IM	IW-WD	IW-WD + IM	OSTAR + IM
sMNIST→USPS					
balanced	96.44 ± 0.3	<b>97.92 ± 0.2</b>	96.10 ± 0.3	<b>97.91 ± 0.1</b>	96.91 ± 0.3
mild	95.18 ± 0.9	95.47 ± 1.1	94.72 ± 0.4	95.74 ± 0.6	<b>96.18 ± 1.0</b>
high	95.07 ± 0.6	93.76 ± 0.5	94.60 ± 0.8	91.73 ± 0.6	<b>96.06 ± 0.6</b>
sUSPS→MNIST					
balanced	93.37 ± 2.5	93.03 ± 1.9	95.68 ± 0.6	96.17 ± 0.5	<b>98.11 ± 0.2</b>
mild	93.20 ± 2.8	94.60 ± 1.7	92.73 ± 1.5	92.65 ± 1.0	<b>97.44 ± 0.5</b>
high	91.54 ± 0.9	90.16 ± 2.0	90.81 ± 1.5	91.26 ± 1.1	<b>97.08 ± 0.6</b>
VisDA12					
VisDA	55.33 ± 0.8	57.57 ± 0.8	51.88 ± 1.6	57.63 ± 0.1	<b>59.24 ± 0.5</b>
sVisDA	51.86 ± 2.0	57.06 ± 0.8	50.65 ± 1.5	57.62 ± 0.7	<b>58.84 ± 1.0</b>
sOffice31					
sW-A	63.80 ± 0.3	68.12 ± 0.5	60.25 ± 0.2	67.42 ± 0.8	<b>69.99 ± 0.1</b>
sA-W	81.05 ± 0.7	81.83 ± 1.9	75.84 ± 0.7	82.34 ± 1.6	<b>83.91 ± 0.5</b>
sOfficeHome					
sR-P	72.72 ± 1.1	<b>75.17 ± 0.6</b>	72.90 ± 0.7	74.94 ± 0.6	<b>75.58 ± 0.6</b>
sR-A	53.37 ± 1.3	54.20 ± 1.3	53.44 ± 0.6	54.50 ± 1.1	<b>56.28 ± 0.5</b>
sR-C	45.30 ± 1.2	48.17 ± 1.2	42.66 ± 1.0	47.93 ± 1.7	<b>49.07 ± 0.9</b>



## C BACKGROUND ON OPTIMAL TRANSPORT

OT was introduced to find a transportation map minimizing the cost of displacing mass from one configuration to another (Villani, 2008). For a comprehensive introduction, we refer to Peyré & Cuturi (2019). Formally, let  $\alpha$  and  $\beta$  be absolutely continuous distributions compactly supported in  $\mathbb{R}^d$  and  $c : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$  a cost function. Consider a map  $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$  that satisfies  $\phi_{\#}\alpha = \beta$ , i.e. that pushes  $\alpha$  to  $\beta$ . We remind that for a function  $f$ ,  $f_{\#}\rho$  is the push-forward measure  $f_{\#}\rho(B) = \rho(f^{-1}(B))$ , for all measurable set  $B$ . The total transportation cost depends on the contributions of costs for transporting each point  $\mathbf{x}$  to  $\phi(\mathbf{x})$  and the Monge OT problem is:

$$\begin{aligned} \min_{\phi} \mathcal{C}_{\text{monge}}(\phi) &= \int_{\mathbb{R}^d} c(\mathbf{x}, \phi(\mathbf{x})) d\alpha(\mathbf{x}) \\ \text{s.t. } \phi_{\#}\alpha &= \beta \end{aligned} \quad (9)$$

$c(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_2^p$  induces the  $p$ -Wasserstein distance,  $\mathcal{W}_p(\alpha, \beta) = \min_{\phi_{\#}\alpha = \beta} \mathcal{C}_{\text{monge}}(\phi)^{1/p}$ . When  $p = 1$ ,  $\mathcal{W}_1$  can be expressed in the dual form  $\mathcal{W}_1(\alpha, \beta) = \sup_{\|v\|_L \leq 1} \mathbb{E}_{\mathbf{x} \sim \alpha} v(\mathbf{x}) - \mathbb{E}_{\mathbf{y} \sim \beta} v(\mathbf{y})$  where  $\|v\|_L$  is the Lipschitz constant of function  $v$ .

## D PROOFS

**Proposition (1).** *For any encoder  $g$  which defines  $\mathcal{Z}$  satisfying Assumption 1, 2, 3, 4, there is an unique solution  $(\phi, \mathbf{p}_N^Y)$  to (OT) and  $\phi_{\#}(p_S(Z|Y)) = p_T(Z|Y)$  and  $\mathbf{p}_N^Y = \mathbf{p}_T^Y$ .*

*Proof.* Fixing  $\mathcal{Z}$  satisfying Assumption 1, 3 and 4, we first show that there exists a solution  $(\phi, \mathbf{p}_N^Y)$  to (OT). Following Brenier (1991) as  $\mathcal{Z} \subset \mathbb{R}^d$ , we can find  $K$  unique Monge maps  $\{\hat{\phi}^{(k)}\}_{k=1}^K$  s.t.  $\forall k \hat{\phi}_{\#}^{(k)}(p_S(Z|Y = k)) = p_T(Z|Y = k)$  with respective transport costs  $\mathcal{W}_1(p_S(Z|Y = k), p_T(Z|Y = k))$ . Let's define  $\hat{\phi}$  as  $\forall k \hat{\phi}|_{\mathcal{Z}_S^{(k)}} = \hat{\phi}^{(k)}$  where  $\cup_{k=1}^K \mathcal{Z}_S^{(k)}$  is the partition of  $\mathcal{Z}_S$  in Assumption 1.  $(\hat{\phi}, \mathbf{p}_T^Y)$  satisfies the equality constraint to (OT), thus we easily deduce existence.

Now, let  $(\phi, \mathbf{p}_N^Y)$  be a solution to (OT), let's show unicity. We first show that  $\phi = \hat{\phi}$ . Under Assumption 2, (OT) is the Monge formulation of the optimal assignment problem between  $\{p_S(Z|Y = k)\}_{k=1}^K$  and  $\{p_T(Z|Y = k)\}_{k=1}^K$  with  $\mathbf{C}$  the cost matrix defined by  $\mathbf{C}_{ij} = \mathcal{W}_2(p_S(Z|Y = i), p_T(Z|Y = j))$ . At the optimum, the transport cost is related to the Wasserstein distance between source conditionals and their corresponding target conditionals i.e.  $\mathcal{C}(\phi) = \sum_{k=1}^K \mathcal{W}_2(p_S(Z|Y = k), \phi_{\#}(p_S(Z|Y = k)))$ . Suppose  $\exists(i, j), j \neq i$  s.t.  $\phi_{\#}(p_S(Z|Y = i)) = p_T(Z|Y = j)$  and  $\phi_{\#}(p_S(Z|Y = j)) = p_T(Z|Y = i)$  and  $\forall k \neq i, j, \phi_{\#}(p_S(Z|Y = k)) = p_T(Z|Y = k)$ . Assumption 3 implies  $\sum_{k=1}^K \mathcal{W}_2(p_S(Z|Y = k), p_T(Z|Y = k)) \leq \sum_{k \neq i, j} \mathcal{W}_2(p_S(Z|Y = k), p_T(Z|Y = k)) + \mathcal{W}_2(p_S(Z|Y = i), p_T(Z|Y = j)) + \mathcal{W}_2(p_S(Z|Y = j), p_T(Z|Y = i))$ . Thus  $\mathcal{C}(\hat{\phi}, \mathcal{Z}) \leq \mathcal{C}(\phi, \mathcal{Z})$  whereas  $\phi$ , solution to (OT), has minimal transport cost. Thus  $\phi = \hat{\phi}$ .

Now let's show  $\mathbf{p}_N^Y = \mathbf{p}_T^Y$  under Assumption 4. We inject  $\phi_{\#}(p_S(Z|Y)) = p_T(Z|Y)$  into (OT),

$$\sum_{k=1}^K \mathbf{p}_N^{Y=k} p_T(Z|k) = \sum_{k=1}^K \mathbf{p}_T^{Y=k} p_T(Z|k) \Leftrightarrow \sum_{k=1}^K (\mathbf{p}_N^{Y=k} - \mathbf{p}_T^{Y=k}) p_T(Z|k) = 0 \Leftrightarrow \mathbf{p}_N^Y = \mathbf{p}_T^Y$$

□

**Theorem (1).** *Given a fixed encoder  $g$  defining a latent space  $\mathcal{Z}$ , two domains  $N$  and  $T$  satisfying cyclical monotonicity in  $\mathcal{Z}$ , assuming that we have  $\forall k, \mathbf{p}_N^{Y=k} > 0$ , then  $\forall f_N \in \mathcal{H}$  where  $\mathcal{H}$  is a set of  $M$ -Lipschitz continuous functions over  $\mathcal{Z}$ , we have*

$$\begin{aligned} \epsilon_T^g(f_N) &\leq \underbrace{\epsilon_N^g(f_N)}_{\text{Classification (C)}} + \underbrace{\frac{2M}{\min_{k=1}^K \mathbf{p}_N^{Y=k}} \mathcal{W}_1(p_N(Z), p_T(Z))}_{\text{Alignment (A)}} \\ &+ \underbrace{2M \left(1 + \frac{1}{\min_{k=1}^K \mathbf{p}_N^{Y=k}}\right) \mathcal{W}_1\left(\sum_{k=1}^K \mathbf{p}_N^{Y=k} p_T(Z|Y = k), \sum_{k=1}^K \mathbf{p}_T^{Y=k} p_T(Z|Y = k)\right)}_{\text{Label (L)}} \end{aligned} \quad (2)$$

*Proof.* We first recall that  $\epsilon_D^g(f_N) = \mathbb{E}_{(\mathbf{z}, y) \in p_D(Z, Y)} \mathcal{L}(f_N(\mathbf{z}), y)$  where  $\mathcal{L}$  is the 0/1 loss. For conciseness,  $\forall \mathbf{z} \in \mathcal{Z}$ ,  $\mathbf{p}_T^{\mathbf{z}|Y}, \mathbf{p}_N^{\mathbf{z}|Y}, \mathcal{L}^{\mathbf{z}, k}$  will refer to the vector of  $[p_T(\mathbf{z}|k)]_{k=1}^K, [p_N(\mathbf{z}|k)]_{k=1}^K, [\mathcal{L}(f_N(\mathbf{z}), k)]_{k=1}^K$  respectively. In the following,  $\odot$  denotes the element-wise product operator between two vectors of the same size.

$$\begin{aligned}
\forall f_N, \epsilon_T^g(f_N) &= \epsilon_N^g(f_N) + \epsilon_T^g(f_N) - \epsilon_N^g(f_N) \\
&\leq \epsilon_N^g(f_N) + \int_{\mathcal{Z}} \sum_{k=1}^K \left( p_T(\mathbf{z}, k) - p_N(\mathbf{z}, k) \right) \times \mathcal{L}(f_N(\mathbf{z}), k) d\mathbf{z} dy \\
&\leq \epsilon_N^g(f_N) + \int_{\mathcal{Z}} \sum_{k=1}^K \left[ \mathbf{p}_T^{Y=k} p_T(\mathbf{z}|k) - \mathbf{p}_N^{Y=k} p_N(\mathbf{z}|k) \right] \times \mathcal{L}(f_N(\mathbf{z}), k) d\mathbf{z} \\
&\leq \epsilon_N^g(f_N) + \int_{\mathcal{Z}} \mathbf{p}_T^{Y\top} \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) - \mathbf{p}_N^{Y\top} \left( \mathbf{p}_N^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z} \\
&\leq \epsilon_N^g(f_N) + \int_{\mathcal{Z}} \mathbf{p}_T^{Y\top} \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) - \mathbf{p}_N^{Y\top} \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) + \mathbf{p}_N^{Y\top} \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) - \mathbf{p}_N^{Y\top} \left( \mathbf{p}_N^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z} \\
&\leq \epsilon_N^g(f_N) + \int_{\mathcal{Z}} \left( \mathbf{p}_T^{Y\top} - \mathbf{p}_N^{Y\top} \right) \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) + \mathbf{p}_N^{Y\top} \left( \left( \mathbf{p}_T^{\mathbf{z}|Y} - \mathbf{p}_N^{\mathbf{z}|Y} \right) \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z} \\
&\leq \epsilon_N^g(f_N) + \int_{\mathcal{Z}} \left( \mathbf{p}_T^{Y\top} - \mathbf{p}_N^{Y\top} \right) \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z} + \int_{\mathcal{Z}} \mathbf{p}_N^{Y\top} \left( \left( \mathbf{p}_T^{\mathbf{z}|Y} - \mathbf{p}_N^{\mathbf{z}|Y} \right) \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z}
\end{aligned}$$

We now introduce a preliminary result from Shen et al. (2018).  $\forall f_N \in \mathcal{H}$   $M$ -Lipschitz continuous,

$$\epsilon_N^g(f_N) - \epsilon_T^g(f_N) \leq 2M \cdot \mathcal{W}_1(p_N^g(Z), p_T^g(Z))$$

Assuming that  $h$  is  $M$ -Lipschitz continuous we apply this result in the following

$$\begin{aligned}
\int_{\mathcal{Z}} \left( \mathbf{p}_T^{Y\top} - \mathbf{p}_N^{Y\top} \right) \left( \mathbf{p}_T^{\mathbf{z}|Y} \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z} &= \int_{\mathcal{Z}} \sum_{k=1}^K \left( \mathbf{p}_T^{Y=k} - \mathbf{p}_N^{Y=k} \right) p_T(\mathbf{z}|k) \times \mathcal{L}(f_N(\mathbf{z}), k) d\mathbf{z} \\
&= \epsilon_T^g(f_N) - \epsilon_N^g(f_N) \quad \text{where } p_{\bar{T}}(Z) = \sum_{k=1}^K \mathbf{p}_N^{Y=k} p_T(Z|k) \\
&\leq 2M \cdot \mathcal{W}_1(p_{\bar{T}}(Z), p_T(Z))
\end{aligned}$$

$$\begin{aligned}
\int_{\mathcal{Z}} \mathbf{p}_N^{Y\top} \left( \left( \mathbf{p}_T^{\mathbf{z}|Y} - \mathbf{p}_N^{\mathbf{z}|Y} \right) \odot \mathcal{L}^{\mathbf{z}, k} \right) d\mathbf{z} &= \int_{\mathcal{Z}} \sum_{k=1}^K \mathbf{p}_N^{Y=k} \left( p_T(\mathbf{z}|k) - p_N(\mathbf{z}|k) \right) \times \mathcal{L}(f_N(\mathbf{z}), k) d\mathbf{z} \\
&\leq \sum_{k=1}^K \int_{\mathcal{Z}} \left( p_T(\mathbf{z}|k) - p_N(\mathbf{z}|k) \right) \times \mathcal{L}(f_N(\mathbf{z}), k) d\mathbf{z} \quad \forall k \mathbf{p}_N^{Y=k} \leq 1 \\
&\leq 2M \sum_{k=1}^K \mathcal{W}_1(p_T(Z|k), p_N(Z|k))
\end{aligned}$$

Thus,  $\forall f_N$   $M$ -Lipschitz continuous

$$\epsilon_T^g(f_N) \leq \epsilon_N^g(f_N) + 2M \times \sum_{k=1}^K \mathcal{W}_1(p_T(Z|k), p_N(Z|k)) + 2M \times \mathcal{W}_1(p_{\bar{T}}(Z), p_T(Z))$$

We rewrite the second term to involve directly latent marginals. Proposition 2 in Rakotomamonjy et al. (2020) shows that under cyclical monotonicity, if  $\forall k, \mathbf{p}_N^{Y=k} > 0$ ,

$$\mathcal{W}_1 \left( \sum_{k=1}^K \mathbf{p}_N^{Y=k} p_T(Z|k), p_N(Z) \right) = \sum_{k=1}^K \mathbf{p}_N^{Y=k} \mathcal{W}_1(p_N(Z|k), p_T(Z|k))$$

This allows to write

$$\begin{aligned} \min_{k=1}^K \mathbf{p}_N^{Y=k} \sum_{k=1}^K \mathcal{W}_1(p_N(Z|k), p_T(Z|k)) &\leq \sum_{k=1}^K \mathbf{p}_N^{Y=k} \mathcal{W}_1(p_N(Z|k), p_T(Z|k)) \\ &= \mathcal{W}_1\left(\sum_{k=1}^K \mathbf{p}_N^{Y=k} p_T(Z|k), p_N(Z)\right) = \mathcal{W}_1(p_{\tilde{T}}(Z), p_N(Z)) \end{aligned}$$

We then use the triangle inequality for the Wasserstein distance  $\mathcal{W}_1$

$$\begin{aligned} \epsilon_T^g(f_N) &\leq \epsilon_N^g(f_N) + \frac{2M}{\min_{k=1}^K \mathbf{p}_N^{Y=k}} \mathcal{W}_1(p_{\tilde{T}}(Z), p_N(Z)) + 2M \times \mathcal{W}_1(p_{\tilde{T}}(Z), p_T(Z)) \\ &\leq \epsilon_N^g(f_N) + \frac{2M}{\min_{k=1}^K \mathbf{p}_N^{Y=k}} \mathcal{W}_1(p_N(Z), p_T(Z)) + 2M\left(1 + \frac{1}{\min_{k=1}^K \mathbf{p}_N^{Y=k}}\right) \mathcal{W}_1(p_{\tilde{T}}(Z), p_T(Z)) \end{aligned}$$

□

**Derivation of the reweighted classification loss (C)**  $\epsilon_N^g(f_N)$  Let  $\mathcal{L}_{ce}$  be the cross-entropy loss. Given a classifier  $h$ , feature extractor  $g$  and domain  $N$ , the mapping of domain  $S$  by  $(\phi, \mathbf{p}_N^Y)$ ,

$$\begin{aligned} \epsilon_N^g(f_N) &= \int_{\mathcal{Z}, \mathcal{Y}} p_N^\phi(\mathbf{z}, y) \mathcal{L}_{ce}(f_N(\mathbf{z}), y) d\mathbf{z}dy \\ &= \int_{\mathcal{Z}, \mathcal{Y}} \mathbf{p}_N^{Y=y} p_N^\phi(\mathbf{z}|y) \mathcal{L}_{ce}(f_N(\mathbf{z}), y) d\mathbf{z}dy \\ &= \int_{\mathcal{Z}, \mathcal{Y}} \frac{\mathbf{p}_N^{Y=y}}{\mathbf{p}_S^{Y=y}} \mathbf{p}_S^{Y=y} p_N^\phi(\mathbf{z}|y) \mathcal{L}_{ce}(f_N(\mathbf{z}), y) d\mathbf{z}dy \quad \forall y \mathbf{p}_S^{Y=y} > 0 \\ &= \int_{\mathcal{Z}, \mathcal{Y}} \frac{\mathbf{p}_N^{Y=y}}{\mathbf{p}_S^{Y=y}} \mathbf{p}_S^{Y=y} \phi_\#(p_S(\mathbf{z}|y)) \mathcal{L}_{ce}(f_N(\mathbf{z}), y) d\mathbf{z}dy \end{aligned}$$

## E PSEUDO-CODE

We detail in Algorithm 1 our pseudo-code and in Algorithm 2 how we minimize (CAL) with respect to  $(\phi, f_N)$  using the dual form of Wasserstein-1 distance (6). Our method is based on a standard backpropagation strategy with gradient descent and uses gradient penalty (Gulrajani et al., 2017).

---

**Algorithm 1** Training and inference procedure for OSTAR

---

**Training:**

$$\hat{S} = \{\mathbf{x}_S^{(i)}, y_S^{(i)}\}_{i=1}^n, \hat{T} = \{\mathbf{x}_T^{(i)}\}_{i=1}^m, \mathcal{Z}_{\hat{N}} = \{\phi \circ g(\mathbf{x}_S^{(i)}), y_S^{(i)}\}_{i=1}^n, \mathcal{Z}_{\hat{T}} = \{g(\mathbf{x}_T^{(i)})\}_{i=1}^m$$

$f_S, f_N \in \mathcal{H}$  classifiers;  $g$  feature extractor;  $\phi$  latent domain-mapping,  $v$  critic.

$N_e$ : number of epochs,  $N_u$ : epoch to update  $\mathbf{p}_N^Y$ ,  $N_g$ : epoch to update  $g$

- 1: Train  $f_S, g$  on  $\hat{S}$  to minimize source classification loss ▷ (3)
- 2: Initialize  $\mathbf{p}_N^Y = \frac{1}{K} \mathbf{1}_K$
- 3: **for**  $n_{epoch} \leq N_e$  **do**
- 4:   **if**  $n_{epoch} \bmod N_u = 0$  **then**
- 5:     Compute  $\mathbf{p}_N^Y$  with estimator in Lipton et al. (2018) on  $(\mathcal{Z}_{\hat{N}}, \mathcal{Z}_{\hat{T}})$  ▷ (CAL) w.r.t.  $\mathbf{p}_N^Y$
- 6:     Average  $\mathbf{p}_N^Y$  with cumulative moving average
- 7:   **if**  $n_{epoch} \leq N_g$  **then** Train  $\phi, v, f_N$  with  $(\hat{S}, \hat{T})$  ▷ (CAL) + (SS) w.r.t.  $\phi, f_N$
- 8:   **else** Train  $\phi, v, f_N, g$  with  $(\hat{S}, \hat{T})$  ▷ (CAL) + (SSg) w.r.t.  $\phi, f_N$

**Inference:** Score  $\mathbf{x}_T$  with  $f_N \circ g(\mathbf{x}_T)$

---

---

**Algorithm 2** Minimize (CAL) w.r.t.  $(\phi, f_N)$ 

---

$\widehat{S} = \{\mathbf{x}_S^{(i)}, y_S^{(i)}\}_{i=1}^n$ ,  $\widehat{T} = \{\mathbf{x}_T^{(i)}\}_{i=1}^m$ ,  $g$  feature extractor,  $\phi$  domain-mapping,  $v$  critic,  $f_N$  classifier. Parameters of  $\phi, v, f_N$ :  $\theta_\phi, \theta_v, \theta_{f_N}$  and learning rates  $\alpha_\phi, \alpha_v, \alpha_{f_N}$ .

$N_{iter}$ : batches per epoch,  $N_b$ : batch size,  $N_v$ : critic iterations

- 1: **for**  $n_{iter} < N_{iter}$  **do**
  - 2:   Sample minibatches  $\mathbf{x}_S^B, y_S^B = \{\mathbf{x}_S^{(i)}, y_S^{(i)}\}_{i=1}^{N_b}$ ,  $\mathbf{x}_T^B = \{\mathbf{x}_T^{(i)}\}_{i=1}^{N_b}$  from  $\widehat{S}, \widehat{T}$
  - 3:   Compute  $\mathbf{z}_S^B = g(\mathbf{x}_S^B)$ ,  $\mathbf{z}_N^B = \phi \circ g(\mathbf{x}_S^B)$  and  $\mathbf{z}_T^B = g(\mathbf{x}_T^B)$
  - 4:   Compute class ratios:  $\mathbf{w}_Y = \mathbf{p}_N^Y / \mathbf{p}_S^Y$
  - 5:   **for**  $n_v < N_v$  **do**
  - 6:     Sample random points  $\mathbf{z}^{B'}$  from the lines between  $(\mathbf{z}_N^B, \mathbf{z}_T^B)$  pairs
  - 7:     Compute gradient penalty  $\mathcal{L}_{grad}$  with  $\mathbf{z}_N^B, \mathbf{z}_T^B, \mathbf{z}^{B'}$  (Gulrajani et al., 2017)
  - 8:     Compute  $\mathcal{L}_{wd}^g = \sum_{i=1}^{N_b} \mathbf{w}_{y_S^{(i)}} v(\mathbf{z}_N^{(i)}) - \frac{1}{N_b} \sum_{i=1}^{N_b} v(\mathbf{z}_T^{(i)})$  (6)
  - 9:      $\theta_v \leftarrow \theta_v - \alpha_v \nabla_{\theta_v} [\mathcal{L}_{wd}^g - \mathcal{L}_{grad}]$
  - 10:    Compute  $\mathcal{L}_{OT}^g = \sum_{k=1}^K \frac{1}{\#\{y_S^{(i)} = k\}_{i \in \{1 \dots N_b\}}} \sum_{y_S^{(i)} = k, i \in \{1 \dots N_b\}} \|\phi(\mathbf{z}_S^{(i)}) - \mathbf{z}_S^{(i)}\|_2^2$
  - 11:     $\theta_\phi \leftarrow \theta_\phi - \alpha_\phi \nabla_{\theta_\phi} [\mathcal{L}_{wd}^g + \mathcal{L}_{OT}^g]$
  - 12:    Compute  $\mathcal{L}_c^g(f_N, N) = \frac{1}{N_b} \sum_{i=1}^{N_b} \mathbf{w}_{y_S^{(i)}} \mathcal{L}_{ce}(f_N \circ \phi \circ g(\mathbf{x}_S^{(i)}), y_S^{(i)})$
  - 13:     $\theta_{f_N} \leftarrow \theta_{f_N} - \alpha_{f_N} \nabla_{\theta_{f_N}} \mathcal{L}_c^g(f_N, N)$
- 

## F EXPERIMENTAL SETUP

**Datasets** We consider the following UDA problems:

- **Digits** is a synthetic binary adaptation problem. We consider adaptation between MNIST and USPS datasets. We consider a subsampled version of the original datasets with the following number of samples per domain: 10000-2609 for MNIST→USPS, 5700-20000 for USPS→MNIST. The feature extractor is learned from scratch.
- **VisDA12** is a 12-class adaptation problem between simulated and real images. We consider a subsampled version of the original problem using 9600 samples per domain and use pre-trained ImageNet ResNet-50 features <http://csr.bu.edu/ftp/visda17/clf/>.
- **Office31** is an object categorization problem with 31 classes. We do not sample the original dataset. There are 3 domains: Amazon (A), DSLR (D) and WebCam (W) and we consider all pairwise source-target domains. We use pre-trained ImageNet ResNet-50 features <https://github.com/jindongwang/transferlearning/blob/master/data/dataset.md>.
- **OfficeHome** is another object categorization problem with 65 classes. We do not sample the original dataset. There are 4 domains: Art (A), Product (P), Clipart (C), Realworld (R) and we consider all pairwise source-target domains. We use pre-trained ImageNet ResNet-50 features <https://github.com/jindongwang/transferlearning/blob/master/data/dataset.md>.

**Imbalance settings** We consider different class-ratios between domains to simulate label-shift and denote with a "s" prefix, the subsampled datasets. For **Digits**, we explicitly provide the class-ratios as Rakotomamonjy et al. (2020) (e.g. for high imbalance, class 2 accounts for the 7% of target samples while class 4 accounts for 22% of target samples). For **Visda12**, **Office31** and **OfficeHome**, subsampled datasets only consider a small percentage of source samples for the first half classes as Combes et al. (2020) (e.g. **sOffice31** considers 30% of source samples in classes below 15 and uses all source samples from other classes and all target samples).

Table 6: Label imbalance settings

Configuration	$\mathbf{p}_S^Y$	$\mathbf{p}_T^Y$
sDigits balanced	$\{\frac{1}{10}, \dots, \frac{1}{10}\}$	$\{\frac{1}{10}, \dots, \frac{1}{10}\}$
sDigits mild	$\{\frac{1}{10}, \dots, \frac{1}{10}\}$	$\{0, 1, 2, 3, 6\} = 0.06, \{4, 5\} = 0.2, \{7, 8, 9\} = 0.1$
sDigits high	$\{\frac{1}{10}, \dots, \frac{1}{10}\}$	$\{0, 1, 2, 3, 6, 7, 8, 9\} = 0.07, \{4, 5\} = 0.22$
VisDA12	$\{0 - 11\} : 100\%$	$\{0 - 11\} : 100\%$
sVisDA12	$\{0 - 5\} : 30\% \{5 - 11\} : 100\%$	$\{0 - 11\} : 100\%$
sOffice31	$\{0 - 15\} : 30\% \{15 - 31\} : 100\%$	$\{0 - 31\} : 100\%$
sOfficeHome	$\{0 - 32\} : 30\% \{33 - 64\} : 100\%$	$\{0 - 64\} : 100\%$

**Hyperparameters** Domain-invariant methods weight alignment over classification; we tuned the corresponding hyperparameter for  $WD_{\beta=0}$  in the range  $[10^{-4}, 10^{-2}]$  and used the one that achieves the best performance on other models. We also tuned  $\lambda_{OT}$  in problem (CAL) and fixed it to  $10^{-2}$  on Digits and  $10^{-5}$  on VisDA12, Office31 and OfficeHome. Batch size is  $N_b = 200$  and all models are trained using Adam with learning rate tuned in the range  $[10^{-4}, 10^{-3}]$ . We initialize NN for classifiers and feature extractors with a normal prior with zero mean and gain 0.02 and  $\phi$  with orthogonal initialization with gain 0.02.

**Training procedure** We fix  $N_e$  the number of epochs to 50 on Digits, 150 on VisDA12 and 100 on Office31, OfficeHome; OSTAR requires smaller  $N_e$  to converge. Critic iterations are fixed to  $N_v = 5$  which worked best for all baseline models; for OSTAR higher values performed better. For all models, we initialize  $f_S, g$  for 10 epochs with (3). Then, we perform alignment either through domain-invariance or with a domain-mapping until we reach the total number of epochs. GeTarS models IW-WD, MARSc, MARSG, OSTAR perform reweighting with estimates refreshed every  $N_u = 2$  epochs in the first 10 alignment epochs, every  $N_u = 5$  epochs after. OSTAR minimizes (CAL) + (SS) for 10 epochs on Digits, Office31, OfficeHome and 5 epochs on VisDA12, then minimizes (CAL) + (SSg) for remaining epochs.

**Architectures** For Digits, our feature extractor  $g$  is composed of three convolutional layers with respectively 64, 64, 128 filters of size  $5 \times 5$  interleaved with batch norm, max-pooling and ReLU. Our classifiers ( $f_S, f_N$ ) are three-layered fully-connected networks with 100 units interleaved with batch norm, ReLU. Our discriminators are three-layered NN with 100 units and ReLU activation. For VisDA12 and Office31, OfficeHome, we consider pre-trained 2048 features obtained from a ResNet-50 followed by 2 fully-connected networks with ReLU and 100 units for VisDA12, 256 units for Office31, OfficeHome. Discriminators are 2-layer fully-connected networks with respectively 100/1 units on VisDA12, 256/1 units on Office31, OfficeHome interleaved with ReLU. Classifiers are 2-layer fully-connected networks with 100/ $K$  units on VisDA12, single layer fully-connected network with  $K$  units on Office31, OfficeHome.  $\phi$  is a ResNet with 10 blocks of two fully-connected layers with ReLU and batch-norm.

**Implementation of target proportion estimators** OSTAR and IW-WD use the confusion based estimator in Lipton et al. (2018) and solve a convex optimization problem ((4) in Combes et al. (2020) and CAL w.r.t  $\mathbf{p}_N^Y$  for OSTAR) which has a unique solution if the soft confusion matrix  $\mathbf{C}$  is of full rank. We implement the same optimization problem using the parallel proximal method from Pustelnik et al. (2011) instead of `cvxopt`<sup>2</sup> used in Combes et al. (2020). MARSc and MARSG (Rakotomamonjy et al., 2020) use linear programming with `POT`<sup>3</sup> to estimate proportions with optimal assignment between conditional distributions. Target conditionals are obtained with hierarchical clustering or with a Gaussian Mixture Model (GMM) using `sklearn`<sup>4</sup>. MARSG has some computational overhead due to the GMM.

<sup>2</sup><http://cvxopt.org/>

<sup>3</sup><https://pythonot.github.io/>

<sup>4</sup><https://scikit-learn.org/>