



# Joint distribution optimal transportation for domain adaptation

Nicolas Courty, Rémi Flamary, Amaury Habrard, Alain Rakotomamonjy

► **To cite this version:**

Nicolas Courty, Rémi Flamary, Amaury Habrard, Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. NIPS 2017, Dec 2017, Los Angeles, United States. <hal-01620589>

**HAL Id: hal-01620589**

**<https://hal.archives-ouvertes.fr/hal-01620589>**

Submitted on 20 Oct 2017

**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.

---

# Joint distribution optimal transportation for domain adaptation

---

**Nicolas Courty\***

Université de Bretagne Sud,  
IRISA, UMR 6074, CNRS,  
courty@univ-ubs.fr

**Rémi Flamary\***

Université Côte d'Azur,  
Lagrange, UMR 7293, CNRS, OCA  
remi.flamary@unice.fr

**Amaury Habrard**

Univ Lyon, UJM-Saint-Etienne, CNRS,  
Lab. Hubert Curien UMR 5516, F-42023  
amaury.habrard@univ-st-etienne.fr

**Alain Rakotomamonjy**

Normandie Université  
Université de Rouen, LITIS EA 4108  
alain.rakoto@insa-rouen.fr

## Abstract

This paper deals with the unsupervised domain adaptation problem, where one wants to estimate a prediction function  $f$  in a given target domain without any labeled sample by exploiting the knowledge available from a source domain where labels are known. Our work makes the following assumption: there exists a non-linear transformation between the joint feature/label space distributions of the two domain  $\mathcal{P}_s$  and  $\mathcal{P}_t$  that can be estimated with optimal transport. We propose a solution of this problem that allows to recover an estimated target  $\mathcal{P}_t^f = (X, f(X))$  by optimizing simultaneously the optimal coupling and  $f$ . We show that our method corresponds to the minimization of a bound on the target error, and provide an efficient algorithmic solution, for which convergence is proved. The versatility of our approach, both in terms of class of hypothesis or loss functions is demonstrated with real world classification and regression problems, for which we reach or surpass state-of-the-art results.

## 1 Introduction

In the context of supervised learning, one generally assumes that the test data is a realization of the same process that generated the learning set. Yet, in many practical applications it is often not the case, since several factors can slightly alter this process. The particular case of visual adaptation [1] in computer vision is a good example: given a new dataset of images without any label, one may want to exploit a different annotated dataset, provided that they share sufficient common information and labels. However, the generating process can be different in several aspects, such as the conditions and devices used for acquisition, different pre-processing, different compressions, etc. Domain adaptation techniques aim at alleviating this issue by transferring knowledge between domains [2]. We propose in this paper a principled and theoretically founded way of tackling this problem.

The domain adaptation (DA) problem is not new and has received a lot of attention during the past ten years. State-of-the-art methods are mainly differing by the assumptions made over the change in data distributions. In the *covariate shift* assumption, the differences between the domains are characterized by a change in the feature distributions  $\mathcal{P}(X)$ , while the conditional distributions  $\mathcal{P}(Y|X)$  remain unchanged ( $X$  and  $Y$  being respectively the instance and label spaces). Importance re-weighting can be used to learn a new classifier (e.g. [3]), provided that the overlapping of the distributions is large

---

\*Both authors contributed equally.

enough. Kernel alignment [4] has also been considered for the same purpose. Other types of method, denoted as *Invariant Components* by Gong and co-authors [5], are looking for a transformation  $\mathcal{T}$  such that the new representations of input data are matching, *i.e.*  $\mathcal{P}_s(\mathcal{T}(X)) = \mathcal{P}_t(\mathcal{T}(X))$ . Methods are then differing by: *i)* The considered class of transformation, that are generally defined as projections (e.g. [6, 7, 8, 9, 5]), affine transform [4] or non-linear transformation as expressed by neural networks [10, 11] *ii)* The types of divergences used to compare  $\mathcal{P}_s(\mathcal{T}(X))$  and  $\mathcal{P}_t(\mathcal{T}(X))$ , such as Kullback Leibler [12] or *Maximum Mean Discrepancy* [9, 5]. Those divergences usually require that the distributions share a common support to be defined. A particular case is found in the use of optimal transport, introduced for domain adaptation by [13, 14].  $\mathcal{T}$  is then defined to be a push-forward operator such that  $\mathcal{P}_s(X) = \mathcal{P}_t(\mathcal{T}(X))$  and that minimizes a global transportation effort or cost between distributions. The associated divergence is the so-called Wasserstein metric, that has a natural Lagrangian formulation and avoids the estimation of continuous distribution by means of kernel. As such, it also alleviates the need for a shared support.

The methods discussed above implicitly assume that the conditional distributions are unchanged by  $\mathcal{T}$ , *i.e.*  $\mathcal{P}_s(Y|\mathcal{T}(X)) \approx \mathcal{P}_t(Y|\mathcal{T}(X))$  but there is no clear reason for this assumption to hold. A more general approach is to adapt both marginal feature and conditional distributions by minimizing a global divergence between them. However, this task is usually hard since no label is available in the target domain and therefore no empirical version  $\mathcal{P}_t(Y|X)$  can be used. This was achieved by restricting to specific class of transformation such as projection [9, 5].

**Contributions and outline.** In this work we propose a novel framework for unsupervised domain adaptation between joint distributions. We propose to find a function  $f$  that predicts an output value given an input  $\mathbf{x} \in \mathcal{X}$ , and that minimizes the optimal transport loss between the joint source distribution  $\mathcal{P}_s$  and an estimated target joint distribution  $\mathcal{P}_t^f = (X, f(X))$  depending on  $f$  (detailed in Section 2). The method is denoted as JDOT for “Joint Distribution Optimal Transport” in the remainder. We show that the resulting optimization problem stands for a minimization of a bound on the target error of  $f$  (Section 3) and propose an efficient algorithm to solve it (Section 4). Our approach is very general and does not require to learn explicitly a transformation, as it directly solves for the best function. We show that it can handle both regression and classification problems with a large class of functions  $f$  including kernel machines and neural networks. We finally provide several numerical experiments on real regression and classification problems that show the performances of JDOT over the state-of-the-art (Section 5).

## 2 Joint distribution Optimal Transport

Let  $\Omega \in \mathbb{R}^d$  be a compact input measurable space of dimension  $d$  and  $\mathcal{C}$  the set of labels.  $\mathcal{P}(\Omega)$  denotes the set of all the probability measures over  $\Omega$ . The standard learning paradigm assumes classically the existence of a set of data  $\mathbf{X}_s = \{\mathbf{x}_i^s\}_{i=1}^{N_s}$  associated with a set of class label information  $\mathbf{Y}_s = \{\mathbf{y}_i^s\}_{i=1}^{N_s}$ ,  $\mathbf{y}_i^s \in \mathcal{C}$  (the learning set), and a data set with unknown labels  $\mathbf{X}_t = \{\mathbf{x}_i^t\}_{i=1}^{N_t}$  (the testing set). In order to determine the set of labels  $\mathbf{Y}_t$  associated with  $\mathbf{X}_t$ , one usually relies on an empirical estimate of the joint probability distribution  $\mathcal{P}(X, Y) \in \mathcal{P}(\Omega \times \mathcal{C})$  from  $(\mathbf{X}_s, \mathbf{Y}_s)$ , and the assumption that  $\mathbf{X}_s$  and  $\mathbf{X}_t$  are drawn from the same distribution  $\mu \in \mathcal{P}(\Omega)$ . In the considered adaptation problem, one assumes the existence of two distinct joint probability distributions  $\mathcal{P}_s(X, Y)$  and  $\mathcal{P}_t(X, Y)$  which correspond respectively to two different *source* and *target* domains. We will write  $\mu_s$  and  $\mu_t$  their respective marginal distributions over  $X$ .

### 2.1 Optimal transport in domain adaptation

The Monge problem is seeking for a map  $\mathcal{T}_0 : \Omega \rightarrow \Omega$  that pushes  $\mu_s$  toward  $\mu_t$  defined as:

$$\mathcal{T}_0 = \operatorname{argmin}_{\mathcal{T}} \int_{\Omega} d(\mathbf{x}, \mathcal{T}(\mathbf{x})) d\mu_s(\mathbf{x}), \quad \text{s.t. } \mathcal{T}\#\mu_s = \mu_t,$$

where  $\mathcal{T}\#\mu_s$  the *image measure* of  $\mu_s$  by  $\mathcal{T}$ , verifying:

$$\mathcal{T}\#\mu_s(A) = \mu_t(\mathcal{T}^{-1}(A)), \quad \forall \text{ Borel subset } A \subset \Omega, \quad (1)$$

and  $d : \Omega \times \Omega \rightarrow \mathbb{R}^+$  is a metric. In the remainder, we will always consider without further notification the case where  $d$  is the squared Euclidean metric. When  $\mathcal{T}_0$  exists, it is called an optimal transport map, but it is not always the case (*e.g.* assume that  $\mu_s$  is defined by one Dirac measure and

$\mu_t$  by two). A relaxed version of this problem has been proposed by Kantorovitch [15], who rather seeks for a transport plan (or equivalently a joint probability distribution)  $\gamma \in \mathcal{P}(\Omega \times \Omega)$  such that:

$$\gamma_0 = \operatorname{argmin}_{\gamma \in \Pi(\mu_s, \mu_t)} \int_{\Omega \times \Omega} d(\mathbf{x}_1, \mathbf{x}_2) d\gamma(\mathbf{x}_1, \mathbf{x}_2), \quad (2)$$

where  $\Pi(\mu_s, \mu_t) = \{\gamma \in \mathcal{P}(\Omega \times \Omega) | p^+ \# \gamma = \mu_s, p^- \# \gamma = \mu_t\}$  and  $p^+$  and  $p^-$  denotes the two marginal projections of  $\Omega \times \Omega$  to  $\Omega$ . Minimizers of this problem are called optimal transport plans. Should  $\gamma_0$  be of the form  $(id \times \mathcal{T}) \# \mu_s$ , then the solution to Kantorovich and Monge problems coincide. As such the Kantorovich relaxation can be seen as a generalization of the Monge problem, with less constraints on the existence and uniqueness of solutions [16].

Optimal transport has been used in DA as a principled way to bring the source and target distribution closer [13, 14, 17], by seeking for a transport plan between the empirical distributions of  $\mathbf{X}_s$  and  $\mathbf{X}_t$  and interpolating  $\mathbf{X}_s$  thanks to a barycentric mapping [14], or by estimating a mapping which is not the solution of Monge problem but allows to map unseen samples [17]. Moreover, they show that better constraining the structure of  $\gamma$  through entropic or classwise regularization terms helps in achieving better empirical results.

## 2.2 Joint distribution optimal transport loss

The main idea of this work is to handle a change in both marginal and conditional distributions. As such, we are looking for a transformation  $\mathcal{T}$  that will align directly the joint distributions  $\mathcal{P}_s$  and  $\mathcal{P}_t$ . Following the Kantorovich formulation of (2),  $\mathcal{T}$  will be implicitly expressed through a coupling between both joint distributions as:

$$\gamma_0 = \operatorname{argmin}_{\gamma \in \Pi(\mathcal{P}_s, \mathcal{P}_t)} \int_{(\Omega \times \mathcal{C})^2} \mathcal{D}(\mathbf{x}_1, y_1; \mathbf{x}_2, y_2) d\gamma(\mathbf{x}_1, y_1; \mathbf{x}_2, y_2), \quad (3)$$

where  $\mathcal{D}(\mathbf{x}_1, y_1; \mathbf{x}_2, y_2) = \alpha d(\mathbf{x}_1, \mathbf{x}_2) + \mathcal{L}(y_1, y_2)$  is a joint cost measure combining both the distances between the samples and a loss function  $\mathcal{L}$  measuring the discrepancy between  $y_1$  and  $y_2$ . While this joint cost is specific (separable), we leave for future work the analysis of generic joint cost function. Putting it in words, matching close source and target samples with similar labels costs few.  $\alpha$  is a positive parameter which balances the metric in the feature space and the loss. As such, when  $\alpha \rightarrow +\infty$ , this cost is dominated by the metric in the input feature space, and the solution of the coupling problem is the same as in [14]. It can be shown that a minimizer to (3) always exists and is unique provided that  $\mathcal{D}(\cdot)$  is lower semi-continuous (see [18], Theorem 4.1), which is the case when  $d(\cdot)$  is a norm and for every usual loss functions [19].

In the unsupervised DA problem, one does not have access to labels in the target domain, and as such it is not possible to find the optimal coupling. Since our goal is to find a function on the target domain  $f: \Omega \rightarrow \mathcal{C}$ , we suggest to replace  $y_2$  by a proxy  $f(\mathbf{x}_2)$ . This leads to the definition of the following joint distribution that uses a given function  $f$  as a proxy for  $y$ :

$$\mathcal{P}_t^f = (\mathbf{x}, f(\mathbf{x}))_{\mathbf{x} \sim \mu_t} \quad (4)$$

In practice we consider empirical versions of  $\mathcal{P}_s$  and  $\mathcal{P}_t^f$ , i.e.  $\hat{\mathcal{P}}_s = \frac{1}{N_s} \sum_{i=1}^{N_s} \delta_{\mathbf{x}_i^s, \mathbf{y}_i^s}$  and  $\hat{\mathcal{P}}_t^f = \frac{1}{N_t} \sum_{i=1}^{N_t} \delta_{\mathbf{x}_i^t, f(\mathbf{x}_i^t)}$ .  $\gamma$  is then a matrix which belongs to  $\Delta$ , i.e. the transportation polytope of non-negative matrices between uniform distributions. Since our goal is to estimate a prediction  $f$  on the target domain, we propose to find the one that produces predictions that match optimally source labels to the aligned target instances in the transport plan. For this purpose, we propose to solve the following problem for JDOT:

$$\min_{f, \gamma \in \Delta} \sum_{ij} \mathcal{D}(\mathbf{x}_i^s, \mathbf{y}_i^s; \mathbf{x}_j^t, f(\mathbf{x}_j^t)) \gamma_{ij} \equiv \min_f W_1(\hat{\mathcal{P}}_s, \hat{\mathcal{P}}_t^f) \quad (5)$$

where  $W_1$  is the 1-Wasserstein distance for the loss  $\mathcal{D}(\mathbf{x}_1, y_1; \mathbf{x}_2, y_2) = \alpha d(\mathbf{x}_1, \mathbf{x}_2) + \mathcal{L}(y_1, y_2)$ . We will make clear in the next section that the function  $f$  we retrieve is theoretically sound with respect to the target error. Note that in practice we add a regularization term for function  $f$  in order to avoid overfitting as discussed in Section 4. An illustration of JDOT for a regression problem is given in Figure 1. In this figure, we have very different joint and marginal distributions but we want

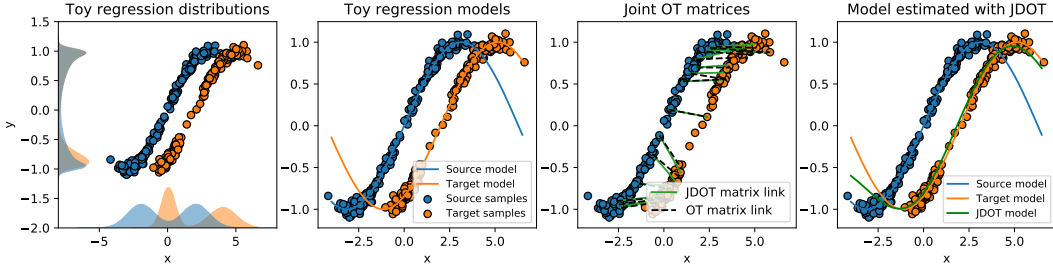


Figure 1: Illustration of JDOT on a 1D regression problem. (left) Source and target empirical distributions and marginals (middle left) Source and target models (middle right) OT matrix on empirical joint distributions and with JDOT proxy joint distribution (right) estimated prediction function  $f$ .

to illustrate that the OT matrix  $\gamma$  obtained using the true empirical distribution  $\mathcal{P}_t$  is very similar to the one obtained with the proxy  $\mathcal{P}_t^f$  which leads to a very good model for JDOT.

**Choice of  $\alpha$ .** This is an important parameter balancing the alignment of feature space and labels. A natural choice of the  $\alpha$  parameter is obtained by normalizing the range of values of  $d(\mathbf{x}_i^s, \mathbf{x}_j^t)$  with  $\alpha = 1/\max_{i,j} d(\mathbf{x}_i^s, \mathbf{x}_j^t)$ . In the numerical experiment section, we show that this setting is very good in two out of three experiments. However, in some cases, better performances are obtained with a cross-validation of this parameter. Also note that  $\alpha$  is strongly linked to the smoothness of the loss  $\mathcal{L}$  and of the optimal labelling functions and can be seen as a Lipschitz constant in the bound of Theorem D.1.

**Relation to other optimal transport based DA methods.** Previous DA methods based on optimal transport [14, 17] do not only differ by the nature of the considered distributions, but also in the way the optimal plan is used to find  $f$ . They learn a complex mapping between the source and target distributions when the objective is only to estimate a prediction function  $f$  on target. To do so, they rely on a barycentric mapping that minimizes only approximately the Wasserstein distance between the distributions. As discussed in Section 4, JDOT uses the optimal plan to propagate and fuse the labels from the source to target. Not only are the performances enhanced, but we also show how this approach is more theoretically well grounded in next section 3.

**Relation to Transport  $L^p$  distances.** Recently, Thorpe and co-authors introduced the Transportation  $L^p$  distance [20]. Their objective is to compute a meaningful distance between multi-dimensional signals. Interestingly their distance can be seen as optimal transport between two distributions of the form (4) where the functions are known and the label loss  $\mathcal{L}$  is chosen as a  $L^p$  distance. While their approach is inspirational, JDOT is different both in its formulation, where we introduce a more general class of loss  $\mathcal{L}$ , and in its objective, as our goal is to estimate the target function  $f$  which is not known *a priori*. Finally we show theoretically and empirically that our formulation addresses successfully the problem of domain adaptation.

### 3 A Bound on the Target Error

Let  $f$  be an hypothesis function from a given class of hypothesis  $\mathcal{H}$ . We define the expected loss in the target domain  $err_T(f)$  as  $err_T(f) \stackrel{\text{def}}{=} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{P}_t} \mathcal{L}(y, f(\mathbf{x}))$ . We define similarly  $err_S(f)$  for the source domain. We assume the loss function  $\mathcal{L}$  to be bounded, symmetric,  $k$ -lipschitz and satisfying the triangle inequality.

To provide some guarantees on our method, we consider an adaptation of the notion probabilistic Lipschitzness introduced in [21, 22] which assumes that two close instances must have the same labels with high probability. It corresponds to a relaxation of the classic Lipschitzness allowing one to model the marginal-label relatedness such as in Nearest-Neighbor classification, linear classification or cluster assumption. We propose an extension of this notion in a domain adaptation context by assuming that a labeling function must comply with two close instances of each domain w.r.t. a coupling  $\Pi$ .

**Definition (Probabilistic Transfer Lipschitzness)** Let  $\mu_s$  and  $\mu_t$  be respectively the source and target distributions. Let  $\phi : \mathbb{R} \rightarrow [0, 1]$ . A labeling function  $f : \Omega \rightarrow \mathbb{R}$  and a joint distribution  $\Pi(\mu_s, \mu_t)$  over  $\mu_s$  and  $\mu_t$  are  $\phi$ -Lipschitz transferable if for all  $\lambda > 0$ :

$$Pr_{(\mathbf{x}_1, \mathbf{x}_2) \sim \Pi(\mu_s, \mu_t)} [|f(\mathbf{x}_1) - f(\mathbf{x}_2)| > \lambda d(\mathbf{x}_1, \mathbf{x}_2)] \leq \phi(\lambda).$$

Intuitively, given a deterministic labeling functions  $f$  and a coupling  $\Pi$ , it bounds the probability of finding pairs of source-target instances labelled differently in a  $(1/\lambda)$ -ball with respect to  $\Pi$ .

We can now give our main result (simplified version):

**Theorem 3.1** *Let  $f$  be any labeling function of  $\in \mathcal{H}$ . Let  $\Pi^* = \operatorname{argmin}_{\Pi \in \Pi(\mathcal{P}_s, \mathcal{P}_t^f)} \int_{(\Omega \times \mathcal{C})^2} \alpha d(\mathbf{x}_s, \mathbf{x}_t) + \mathcal{L}(y_s, y_t) d\Pi(\mathbf{x}_s, y_s; \mathbf{x}_t, y_t)$  and  $W_1(\hat{\mathcal{P}}_s, \hat{\mathcal{P}}_t^f)$  the associated 1-Wasserstein distance. Let  $f^* \in \mathcal{H}$  be a Lipschitz labeling function that verifies the  $\phi$ -probabilistic transfer Lipschitzness (PTL) assumption w.r.t.  $\Pi^*$  and that minimizes the joint error  $\operatorname{err}_S(f^*) + \operatorname{err}_T(f^*)$  w.r.t all PTL functions compatible with  $\Pi^*$ . We assume the input instances are bounded s.t.  $|f^*(\mathbf{x}_1) - f^*(\mathbf{x}_2)| \leq M$  for all  $\mathbf{x}_1, \mathbf{x}_2$ . Let  $\mathcal{L}$  be any symmetric loss function,  $k$ -Lipschitz and satisfying the triangle inequality. Consider a sample of  $N_s$  labeled source instances drawn from  $\mathcal{P}_s$  and  $N_t$  unlabeled instances drawn from  $\mu_t$ , and then for all  $\lambda > 0$ , with  $\alpha = k\lambda$ , we have with probability at least  $1 - \delta$  that:*

$$\operatorname{err}_T(f) \leq W_1(\hat{\mathcal{P}}_s, \hat{\mathcal{P}}_t^f) + \sqrt{\frac{2}{c'} \log\left(\frac{2}{\delta}\right)} \left( \frac{1}{\sqrt{N_s}} + \frac{1}{\sqrt{N_t}} \right) + \operatorname{err}_S(f^*) + \operatorname{err}_T(f^*) + kM\phi(\lambda).$$

The detailed proof of Theorem D.1 is given in the supplementary material. The previous bound on the target error above is interesting to interpret. The first two terms correspond to the objective function (5) we propose to minimize accompanied with a sampling bound. The last term  $\phi(\lambda)$  assesses the probability under which the probabilistic Lipschitzness does not hold. The remaining two terms involving  $f^*$  correspond to the joint error minimizer illustrating that domain adaptation can work only if we can predict well in both domains, similarly to existing results in the literature [23, 24]. If the last terms are small enough, adaptation is possible if we are able to align well  $\mathcal{P}_s$  and  $\mathcal{P}_t^f$ , provided that  $f^*$  and  $\Pi^*$  verify the PTL. Finally, note that  $\alpha = k\lambda$  and tuning this parameter is thus actually related to finding the Lipschitz constants of the problem.

## 4 Learning with Joint Distribution OT

In this section, we provide some details about the JDOT's optimization problem given in Equation (5) and discuss algorithms for its resolution. We will assume that the function space  $\mathcal{H}$  to which  $f$  belongs is either a RKHS or a function space parametrized by some parameters  $\mathbf{w} \in \mathbb{R}^p$ . This framework encompasses linear models, neural networks, and kernel methods. Accordingly, we are going to define a regularization term  $\Omega(f)$  on  $f$ . Depending on how  $\mathcal{H}$  is defined,  $\Omega(f)$  is either a non-decreasing function of the squared-norm induced by the RKHS (so that the representer theorem is applicable) or a squared-norm on the vector parameter. We will further assume that  $\Omega(f)$  is continuously differentiable. As discussed above,  $f$  is to be learned according to the following optimization problem

$$\min_{f \in \mathcal{H}, \gamma \in \Delta} \sum_{i,j} \gamma_{i,j} (\alpha d(\mathbf{x}_i^s, \mathbf{x}_j^t) + \mathcal{L}(y_i^s, f(\mathbf{x}_j^t))) + \lambda \Omega(f) \quad (6)$$

where the loss function  $\mathcal{L}$  is continuous and differentiable with respects to its second variable. Note that while the above problem does not involve any regularization term on the coupling matrix  $\gamma$ , it is essentially for the sake of simplicity and readability. Regularizers like entropic regularization [25], which is relevant when the number of samples is very large, can still be used without significant change to the algorithmic framework.

**Optimization procedure.** According to the above hypotheses on  $f$  and  $\mathcal{L}$ , Problem (6) is smooth and the constraints are separable according to  $f$  and  $\gamma$ . Hence, a natural way to solve the problem (6) is to rely on alternate optimization w.r.t. both parameters  $\gamma$  and  $f$ . This algorithm well-known as Block Coordinate Descent (BCD) or Gauss-Seidel method (the pseudo code of the algorithm is given in appendix). Block optimization steps are discussed with further details in the following.

Solving with fixed  $f$  boils down to a classical OT problem with a loss matrix  $\mathbf{C}$  such that  $C_{i,j} = \alpha d(\mathbf{x}_i^s, \mathbf{x}_j^t) + \mathcal{L}(y_i^s, f(\mathbf{x}_j^t))$ . We can use classical OT solvers such as the network simplex algorithm, but other strategies can be considered, such as regularized OT [25] or stochastic versions [26].

The optimization problem with fixed  $\gamma$  leads to a new learning problem expressed as

$$\min_{f \in \mathcal{H}} \sum_{i,j} \gamma_{i,j} \mathcal{L}(y_i^s, f(\mathbf{x}_j^t)) + \lambda \Omega(f) \quad (7)$$

Note how the data fitting term elegantly and naturally encodes the transfer of source labels  $y_i^s$  through estimated labels of test samples with a weighting depending on the optimal transport matrix. However, this comes at the price of having a quadratic number  $N_s N_t$  of terms, which can be considered as computationally expensive. We will see in the sequel that we can benefit from the structure of the chosen loss to greatly reduce its complexity. In addition, we emphasize that when  $\mathcal{H}$  is a RKHS, owing to kernel trick and the representer theorem, problem (7) can be re-expressed as an optimization problem with  $N_t$  number of parameters all belonging to  $\mathbb{R}$ .

Let us now discuss briefly the convergence of the proposed algorithm. Owing to the 2-block coordinate descent structure, to the differentiability of the objective function in Problem (6) and constraints on  $f$  (or its kernel trick parameters) and  $\gamma$  are closed, non-empty and convex, convergence result of Grippo et al. [27] on 2-block Gauss-Seidel methods directly applies. It states that if the sequence  $\{\gamma^k, f^k\}$  produced by the algorithm has limit points then every limit point of the sequence is a critical point of Problem (6).

**Estimating  $f$  for least square regression problems.** We detail the use of JDOT for transfer least-square regression problem i.e when  $\mathcal{L}$  is the squared-loss. In this context, when the optimal transport matrix  $\gamma$  is fixed the learning problem boils down to

$$\min_{f \in \mathcal{H}} \sum_j \frac{1}{n_t} \|\hat{y}_j - f(\mathbf{x}_j^t)\|^2 + \lambda \|f\|^2 \quad (8)$$

where the  $\hat{y}_j = n_t \sum_j \gamma_{i,j} y_i^s$  is a weighted average of the source target values. Note that this simplification results from the properties of the quadratic loss and that it may not occur for more complex regression loss.

**Estimating  $f$  for hinge loss classification problems.** We now aim at estimating a multiclass classifier with a one-against-all strategy. We suppose that the data fitting is the binary squared hinge loss of the form  $\mathcal{L}(y, f(\mathbf{x})) = \max(0, 1 - yf(\mathbf{x}))^2$ . In a One-Against-All strategy we often use the binary matrices  $\mathbf{P}$  such that  $P_{i,k}^s = 1$  if sample  $i$  is of class  $k$  else  $P_{i,k}^s = 0$ . Denote as  $f_k \in \mathcal{H}$  the decision function related to the  $k$ -vs-all problem. The learning problem (7) can now be expressed as

$$\min_{f_k \in \mathcal{H}} \sum_{j,k} \hat{P}_{j,k} \mathcal{L}(1, f_k(\mathbf{x}_j^t)) + (1 - \hat{P}_{j,k}) \mathcal{L}(-1, f_k(\mathbf{x}_j^t)) + \lambda \sum_k \|f_k\|^2 \quad (9)$$

where  $\hat{\mathbf{P}}$  is the transported class proportion matrix  $\hat{\mathbf{P}} = \frac{1}{N_t} \gamma^\top \mathbf{P}^s$ . Interestingly this formulation illustrates that for each target sample, the data fitting term is a convex sum of hinge loss for a negative and positive label with weights in  $\gamma$ .

## 5 Numerical experiments

In this section we evaluate the performance of our method (**JDOT**) on two different transfer tasks of classification and regression on real datasets <sup>2</sup>.

**Caltech-Office classification dataset.** This dataset [28] is dedicated to visual adaptation. It contains images from four different domains: *Amazon*, the *Caltech-256* image collection, *Webcam* and *DSLR*. Several features, such as presence/absence of background, lightning conditions, image quality, etc.) induce a distribution shift between the domains, and it is therefore relevant to consider a domain adaptation task to perform the classification. Following [14], we choose deep learning features to represent the images, extracted as the weights of the fully connected 6th layer of the DECAF convolutional neural network [29], pre-trained on ImageNet. The final feature vector is a sparse 4096 dimensional vector.

<sup>2</sup>Open Source Python implementation of JDOT: <https://github.com/rflamary/JDOT>

Table 1: Accuracy on the Caltech-Office Dataset. Best value in bold.

Domains	Base	SurK	SA	ARTL	OT-IT	OT-MM	JDOT
caltech→amazon	92.07	91.65	90.50	92.17	89.98	<b>92.59</b>	91.54
caltech→webcam	76.27	77.97	81.02	80.00	80.34	78.98	<b>88.81</b>
caltech→dslr	84.08	82.80	85.99	88.54	78.34	76.43	<b>89.81</b>
amazon→caltech	84.77	84.95	85.13	85.04	85.93	<b>87.36</b>	85.22
amazon→webcam	79.32	81.36	<b>85.42</b>	79.32	74.24	85.08	84.75
amazon→dslr	86.62	87.26	<b>89.17</b>	85.99	77.71	79.62	87.90
webcam→caltech	71.77	71.86	75.78	72.75	<b>84.06</b>	82.99	82.64
webcam→amazon	79.44	78.18	81.42	79.85	89.56	90.50	<b>90.71</b>
webcam→dslr	96.18	95.54	94.90	<b>100.00</b>	99.36	99.36	98.09
dslr→caltech	77.03	76.94	81.75	78.45	<b>85.57</b>	83.35	84.33
dslr→amazon	83.19	82.15	83.19	83.82	<b>90.50</b>	<b>90.50</b>	88.10
dslr→webcam	96.27	92.88	88.47	<b>98.98</b>	96.61	96.61	96.61
<b>Mean</b>	83.92	83.63	85.23	85.41	86.02	86.95	<b>89.04</b>
<b>Mean rank</b>	5.33	5.58	4.00	3.75	3.50	2.83	2.50
<b>p-value</b>	< 0.01	< 0.01	0.01	0.04	0.25	0.86	–

Table 2: Accuracy on the Amazon review experiment. Maximum value in bold font.

Domains	NN	DANN	JDOT (mse)	JDOT (Hinge)
books→dvd	0.805	<b>0.806</b>	0.794	0.795
books→kitchen	0.768	0.767	0.791	<b>0.794</b>
books→electronics	0.746	0.747	0.778	<b>0.781</b>
dvd→books	0.725	0.747	0.761	<b>0.763</b>
dvd→kitchen	0.760	0.765	0.811	<b>0.821</b>
dvd→electronics	0.732	0.738	0.778	<b>0.788</b>
kitchen→books	0.704	0.718	<b>0.732</b>	0.728
kitchen→dvd	0.723	0.730	0.764	<b>0.765</b>
kitchen→electronics	<b>0.847</b>	0.846	0.844	0.845
electronics→books	0.713	0.718	0.740	<b>0.749</b>
electronics→dvd	0.726	0.726	<b>0.738</b>	0.737
electronics→kitchen	0.855	0.850	0.868	<b>0.872</b>
<b>Mean</b>	0.759	0.763	0.783	<b>0.787</b>
<b>p-value</b>	0.004	0.006	0.025	–

We compare our method with four other methods: the surrogate kernel approach ([4], denoted **SurK**), subspace adaptation for its simplicity and good performances on visual adaptation ([8], **SA**), Adaptation Regularization based Transfer Learning ([30], **ARTL**), and the two variants of regularized optimal transport [14]: entropy-regularized **OT-IT** and classwise regularization implemented with the Majoration-Minimization algorithm **OT-MM**, that showed to give better results in practice than its group-lasso counterpart. The classification is conducted with a SVM together with a linear kernel for every method. Its results when learned on the source domain and tested on the target domain are also reported to serve as baseline (**Base**). All the methods have hyper-parameters, that are selected using the reverse cross-validation of Zhong and colleagues [31]. The dimension  $d$  for **SA** is chosen from  $\{1, 4, 7, \dots, 31\}$ . The entropy regularization for **OT-IT** and **OT-MM** is taken from  $\{10^2, \dots, 10^5\}$ ,  $10^2$  being the minimum value for the Sinkhorn algorithm to prevent numerical errors. Finally the  $\eta$  parameter of **OT-MM** is selected from  $\{1, \dots, 10^5\}$  and the  $\alpha$  in **JDOT** from  $\{10^{-5}, 10^{-4}, \dots, 1\}$ .

The classification accuracy for all the methods is reported in Table 1. We can see that **JDOT** is consistently outperforming the baseline (5 points in average), indicating that the adaptation is successful in every cases. Its mean accuracy is the best as well as its average ranking. We conducted a Wilcoxon signed-rank test to test if **JDOT** was statistically better than the other methods, and report the p-value in the tables. This test shows that **JDOT** is statistically better than the considered methods, except for OT based ones that where state of the art on this dataset [14].

**Amazon review classification dataset** We now consider the *Amazon review dataset* [32] which contains online reviews of different products collected on the Amazon website. Reviews are encoded with bag-of-word unigram and bigram features as input. The problem is to predict positive (higher than 3 stars) or negative (3 stars or less) notation of reviews (binary classification). Since different



Table 3: Comparison of different methods on the Wifi localization dataset. Maximum value in bold.

Domains	<b>KRR</b>	<b>SurK</b>	<b>DIP</b>	<b>DIP-CC</b>	<b>GeTarS</b>	<b>CTC</b>	<b>CTC-TIP</b>	<b>JDOT</b>
t1 → t2	80.84±1.14	90.36±1.22	87.98±2.33	91.30±3.24	86.76 ± 1.91	89.36±1.78	89.22±1.66	<b>93.03 ± 1.24</b>
t1 → t3	76.44±2.66	<b>94.97±1.29</b>	84.20±4.29	84.32±4.57	90.62±2.25	94.80±0.87	92.60 ± 4.50	90.06 ± 2.01
t2 → t3	67.12±1.28	85.83 ± 1.31	80.58 ± 2.10	81.22 ± 4.31	82.68 ± 3.71	87.92 ± 1.87	<b>89.52 ± 1.14</b>	86.76 ± 1.72
hallway1	60.02 ± 2.60	76.36 ± 2.44	77.48 ± 2.68	76.24± 5.14	84.38 ± 1.98	86.98 ± 2.02	86.78 ± 2.31	<b>98.83±0.58</b>
hallway2	49.38 ± 2.30	64.69±0.77	78.54 ± 1.66	77.8± 2.70	77.38 ± 2.09	87.74 ± 1.89	87.94 ± 2.07	<b>98.45±0.67</b>
hallway3	48.42 ± 1.32	65.73 ± 1.57	75.10± 3.39	73.40± 4.06	80.64 ± 1.76	82.02± 2.34	81.72 ± 2.25	<b>99.27±0.41</b>

words are employed to qualify the different categories of products, a domain adaptation task can be formulated if one wants to predict positive reviews of a product from labelled reviews of a different product. Following [33, 11], we consider only a subset of four different types of product: books, DVDs, electronics and kitchens. This yields 12 possible adaptation tasks. Each domain contains 2000 labelled samples and approximately 4000 unlabelled ones. We therefore use these unlabelled samples to perform the transfer, and test on the 2000 labelled data.

The goal of this experiment is to compare to the state-of-the-art method on this subset, namely Domain adversarial neural network ([11], denoted **DANN**), and to show the versatility of our method that can adapt to any type of classifier. The neural network used for all methods in this experiment is a simple 2-layer model with sigmoid activation function in the hidden layer to promote non-linearity. 50 neurons are used in this hidden layer. For **DANN**, hyper-parameters are set through the reverse cross-validation proposed in [11], and following the recommendation of authors the learning rate is set to  $10^{-3}$ . In the case of **JDOT**, we used the heuristic setting of  $\alpha = 1 / \max_{i,j} d(\mathbf{x}_i^s, \mathbf{x}_j^t)$ , and as such we do not need any cross-validation. The squared Euclidean norm is used for both metric in feature space and we test as loss functions both mean squared errors (mse) and Hinge losses. 10 iterations of the block coordinate descent are realized. For each method, we stop the learning process of the network after 5 epochs. Classification accuracies are presented in table 2. The neural network (**NN**), trained on source and tested on target, is also presented as a baseline. **JDOT** surpasses **DANN** in 11 out of 12 tasks (except on books→dvd). The Hinge loss is better in than mse in 10 out of 12 cases, which is expected given the superiority of the Hinge loss on classification tasks [19].

**Wifi localization regression dataset** For the regression task, we use the cross-domain indoor Wifi localization dataset that was proposed by Zhang and co-authors [4], and recently studied in [5]. From a multi-dimensional signal (collection of signal strength perceived from several access points), the goal is to locate the device in a hallway, discretized into a grid of 119 squares, by learning a mapping from the signal to the grid element. This translates as a regression problem. As the signals were acquired at different time periods by different devices, a shift can be encountered and calls for an adaptation. In the remaining, we follow the exact same experimental protocol as in [4, 5] for ease of comparison. Two cases of adaptation are considered: **transfer across periods**, for which three time periods t1, t2 and t3 are considered, and **transfer across devices**, where three different devices are used to collect the signals in the same straight-line hallways (hallway1-3), leading to three different adaptation tasks in both cases.

We compare the result of our method with several state-of-the-art methods: kernel ridge regression with RBF kernel (**KRR**), surrogate kernel ([4], denoted **SurK**), domain-invariant projection and its cluster regularized version ([7], denoted respectively **DIP** and **DIP-CC**), generalized target shift ([34], denoted **GeTarS**), and conditional transferable components, with its target information preservation regularization ([5], denoted respectively **CTC** and **CTC-TIP**). As in [4, 5], the hyper-parameters of the competing methods are cross-validated on a small subset of the target domain. In the case of **JDOT**, we simply set the  $\alpha$  to the heuristic value of  $\alpha = 1 / \max_{i,j} d(\mathbf{x}_i^s, \mathbf{x}_j^t)$  as discussed previously, and  $f$  is estimated with kernel ridge regression.

Following [4], the accuracy is measured in the following way: the prediction is said to be correct if it falls within a range of three meters in the transfer across periods, and six meters in the transfer across devices. For each experiment, we randomly sample sixty percent of the source and target domain, and report the mean and standard deviation of ten repetitions accuracies in Table 3. For transfer across periods, **JDOT** performs best in one out of three tasks. For transfer across devices, the superiority of **JDOT** is clearly assessed, for it reaches an average score  $> 98\%$ , which is at least ten points ahead of the best competing method for every task. Those extremely good results could be explained by the fact that using optimal transport allows to consider large shifts of distribution, for which divergences (such as maximum mean discrepancy used in **CTC**) or reweighting strategies can not cope with.

## 6 Discussion and conclusion

We have presented in this paper the Joint Distribution Optimal Transport for domain adaptation, which is a principled way of performing domain adaptation with optimal transport. JDOT assumes the existence of a transfer map that transforms a source domain joint distribution  $\mathcal{P}_s(X, Y)$  into a target domain equivalent version  $\mathcal{P}_t(X, Y)$ . Through this transformation, the alignment of both feature space and conditional distributions is operated, allowing to devise an efficient algorithm that simultaneously optimizes for a coupling between  $\mathcal{P}_s$  and  $\mathcal{P}_t$  and a prediction function that solves the transfer problem. We also proved that learning with **JDOT** is equivalent to minimizing a bound on the target distribution. We have demonstrated through experiments on classical real-world benchmark datasets the superiority of our approach w.r.t. several state-of-the-art methods, including previous work on optimal transport based domain adaptation, domain adversarial neural networks or transfer components, on a variety of task including classification and regression. We have also showed the versatility of our method, that can accommodate with several types of loss functions (mse, hinge) or class of hypothesis (including kernel machines or neural networks). Potential follow-ups of this work include a semi-supervised extension (using unlabelled examples in source domain) and investigating stochastic techniques for solving efficiently the adaptation. From a theoretical standpoint, future works include a deeper study of probabilistic transfer lipschitzness and the development of guarantees able to take into the complexity of the hypothesis class and the space of possible transport plans.

## Acknowledgements

This work benefited from the support of the project OATMIL ANR-17-CE23-0012 of the French National Research Agency (ANR), the Normandie Projet GRR-DAISI, European funding FEDER DAISI and CNRS funding from the Défi Imag'In. The authors also wish to thank Kai Zhang and Qiaojun Wang for providing the Wifi localization dataset.

## References

- [1] V. M. Patel, R. Gopalan, R. Li, and R. Chellappa. Visual domain adaptation: an overview of recent advances. *IEEE Signal Processing Magazine*, 32(3), 2015.
- [2] S. J. Pan and Q. Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.
- [3] M. Sugiyama, S. Nakajima, H. Kashima, P.V. Buenau, and M. Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *NIPS*, 2008.
- [4] K. Zhang, V. W. Zheng, Q. Wang, J. T. Kwok, Q. Yang, and I. Marsic. Covariate shift in Hilbert space: A solution via surrogate kernels. In *ICML*, 2013.
- [5] M. Gong, K. Zhang, T. Liu, D. Tao, C. Glymour, and B. Schölkopf. Domain adaptation with conditional transferable components. In *ICML*, volume 48, pages 2839–2848, 2016.
- [6] B. Gong, Y. Shi, F. Sha, and K. Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *CVPR*, 2012.
- [7] M. Baktashmotlagh, M. Harandi, B. Lovell, and M. Salzmann. Unsupervised domain adaptation by domain invariant projection. In *ICCV*, pages 769–776, 2013.
- [8] B. Fernando, A. Habrard, M. Sebban, and T. Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *ICCV*, 2013.
- [9] M. Long, J. Wang, G. Ding, J. Sun, and P. Yu. Transfer joint matching for unsupervised domain adaptation. In *CVPR*, pages 1410–1417, 2014.
- [10] Y. Ganin and V. Lempitsky. Unsupervised domain adaptation by backpropagation. In *ICML*, pages 1180–1189, 2015.
- [11] Y. Ganin, E. Ustinova, H. Ajakan, P. Germain, H. Larochelle, F. Laviolette, M. Marchand, and V. Lempitsky. Domain-adversarial training of neural networks. *Journal of Machine Learning Research*, 17(59):1–35, 2016.

- [12] S. Si, D. Tao, and B. Geng. Bregman divergence-based regularization for transfer subspace learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(7):929–942, July 2010.
- [13] N. Courty, R. Flamary, and D. Tuia. Domain adaptation with regularized optimal transport. In *ECML/PKDD*, 2014.
- [14] N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.
- [15] L. Kantorovich. On the translocation of masses. *C.R. (Doklady) Acad. Sci. URSS (N.S.)*, 37:199–201, 1942.
- [16] F. Santambrogio. Optimal transport for applied mathematicians. *Birkäuser, NY*, 2015.
- [17] M. Perrot, N. Courty, R. Flamary, and A. Habrard. Mapping estimation for discrete optimal transport. In *NIPS*, pages 4197–4205, 2016.
- [18] C. Villani. *Optimal transport: old and new*. Grund. der mathematischen Wissenschaften. Springer, 2009.
- [19] Lorenzo Rosasco, Ernesto De Vito, Andrea Caponnetto, Michele Piana, and Alessandro Verri. Are loss functions all the same? *Neural Computation*, 16(5):1063–1076, 2004.
- [20] M. Thorpe, S. Park, S. Kolouri, G. Rohde, and D. Slepcev. A transportation  $l^p$  distance for signal analysis. *CoRR*, abs/1609.08669, 2016.
- [21] R. Urner, S. Shalev-Shwartz, and S. Ben-David. Access to unlabeled data can speed up prediction time. In *Proceedings of ICML*, pages 641–648, 2011.
- [22] S. Ben-David, S. Shalev-Shwartz, and R. Urner. Domain adaptation—can quantity compensate for quality? In *Proc of ISAIM*, 2012.
- [23] Y. Mansour, M. Mohri, and A. Rostamizadeh. Domain adaptation: Learning bounds and algorithms. In *Proc. of COLT*, 2009.
- [24] S. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. Wortman Vaughan. A theory of learning from different domains. *Machine Learning*, 79(1-2):151–175, 2010.
- [25] M. Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NIPS*, 2013.
- [26] A. Genevay, M. Cuturi, G. Peyré, and F. Bach. Stochastic optimization for large-scale optimal transport. In *NIPS*, pages 3432–3440, 2016.
- [27] Luigi Grippo and Marco Sciandrone. On the convergence of the block nonlinear gauss–seidel method under convex constraints. *Operations research letters*, 26(3):127–136, 2000.
- [28] K. Saenko, B. Kulis, M. Fritz, and T. Darrell. Adapting visual category models to new domains. In *ECCV*, LNCS, pages 213–226, 2010.
- [29] J. Donahue, Y. Jia, O. Vinyals, J. Hoffman, N. Zhang, E. Tzeng, and T. Darrell. Decaf: A deep convolutional activation feature for generic visual recognition. In *ICML*, 2014.
- [30] M. Long, J. Wang, G. Ding, S. Jialin Pan, and P.S. Yu. Adaptation regularization: A general framework for transfer learning. *IEEE TKDE*, 26(7):1076–1089, 2014.
- [31] E. Zhong, W. Fan, Q. Yang, O. Verscheure, and J. Ren. Cross validation framework to choose amongst models and datasets for transfer learning. In *ECML/PKDD*, 2010.
- [32] J. Blitzer, R. McDonald, and F. Pereira. Domain adaptation with structural correspondence learning. In *Proc. of the 2006 conference on empirical methods in natural language processing*, pages 120–128, 2006.
- [33] M. Chen, Z. Xu, K. Weinberger, and F. Sha. Marginalized denoising autoencoders for domain adaptation. In *ICML*, 2012.
- [34] K. Zhang, M. Gong, and B. Schölkopf. Multi-source domain adaptation: A causal view. In *AAAI Conference on Artificial Intelligence*, pages 3150–3157, 2015.
- [35] F. Bolley, A. Guillin, and C Villani. Quantitative concentration inequalities for empirical measures on non-compact spaces. *Probability Theory and Related Fields*, pages 541–593, 2007.

## A Illustration on a simple example

We illustrate the behavior of our method on a 3-class toy example (Figure 2). We consider a classification problem using the hinge loss and  $\mathcal{H}$  is a Reproducing Kernel Hilbert Space. Source domain samples are drawn from three different 2D Gaussian distributions with different centers and standard deviations. The target domain is obtained rotating the source distribution by  $\pi/4$  radian. Two types of kernel are considered: linear and RBF. In Figure 2.a, one can observe on the first column of images that using directly a classifier learned on the source domain leads to bad performances because of the rotation. We then show the iterations of the block coordinate descent which allows one to recover the true labels of the target domain. It is also interesting to examine the impact of the  $\alpha$  parameter on the success of the method. In Figure 2.b, we show the evolution of classification accuracy for six different  $\alpha$  in the case of RBF kernel. Relying mostly on the label cost ( $\alpha = \{0.1\}$ ) leads to a deterioration of the final accuracy. Using only the input space distance ( $\alpha = \{50, 100\}$ ), which is equivalent to [14], allows a performance gain. But it is clear that using both losses with  $\alpha = \{0.5, 1, 10\}$  leads to the best performance. Also note the small number of iterations required ( $< 10$ ) for achieving a steady state.




Figure 2: **Illustration on a toy example.** (a): Decision boundaries for linear and RBF kernels on selected iterations. The source domain is depicted with crosses, while the target domain samples are class-colored circles. (b): Evolution of the accuracy along 15 iterations of the method for different values of the  $\alpha$  parameter;

## B Block coordinate descent algorithm for solving JDOT

We give in algorithm 1 an overview of the block coordinate descent algorithm used for solving JDOT.

---

### Algorithm 1 Optimization with Block Coordinate Descent

---

```

Initialize function  $f^0$  and set  $k = 1$ 
Set  $\alpha$  and  $\lambda$ 
while not converged do
     $\gamma^k \leftarrow$  Solve OT problem (3 in paper) with fixed  $f^{k-1}$ 
     $f^k \leftarrow$  Solve learning problem (7 in paper) with fixed  $\gamma^k$ 
     $k \leftarrow k + 1$ 
end while

```

---

## C BCD iterations on real data

We report in Table 4, for a fixed set of parameter (no CV), the evolution of the empirical error along the iterations of the 15 first iterations of the BCD on a real dataset. We can see that generally the result stabilizes at around 10 iterations. We can also observe that the increase in performance is not monotonic, contrary to the toy example.

Iter	caltech $\rightarrow$ amazon	dslr $\rightarrow$ amazon	webcam $\rightarrow$ caltech
0	89.14	80.9	75.6
1	91.75	86.22	80.23
2	91.44	86.95	81.75
3	91.54	87.68	82.64
4	91.44	87.68	83.26
5	91.65	88.0	83.35
6	91.86	88.1	83.17
7	92.17	87.89	83.08
8	92.28	87.58	83.26
9	92.28	87.58	83.26
10	92.28	87.68	83.35
11	92.28	87.79	83.44
12	92.28	87.79	83.44
13	92.28	87.79	83.44
14	92.28	87.79	83.44

Table 4: Accuracy of the estimated model along BCD iterations on Caltech-office dataset

## D Proof of Theorem 3.1

We first recall some hypothesis used for this theorem.

$\mathcal{H} \subset \mathcal{C}^\Omega$  is the hypothesis class.  $\mathcal{L} : \mathcal{C} \times \mathcal{C} \rightarrow \mathbb{R}^+$  is the loss function measuring the discrepancy between two labels. This loss is assumed to be symmetric, bounded and  $k$ -lipschitz in its second argument, *i.e.* there exists  $k$  such that for any  $y_1, y_2, y_3 \in \mathcal{C}$ :

$$|\mathcal{L}(y_1, y_2) - \mathcal{L}(y_1, y_3)| \leq k|y_2 - y_3|.$$

$\mathcal{P}_t$  and  $\mathcal{P}_s$  are respectively the target and source distributions over  $\Omega \times \mathcal{C}$ , with  $\mu_t$  and  $\mu_s$  the respective marginals over  $\Omega$ . The expected loss in the target domain  $err_T(f)$  is defined for any  $f \in \mathcal{H}$  as

$$err_T(f) \stackrel{\text{def}}{=} \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{P}_t} \mathcal{L}(y, f(\mathbf{x})).$$

We can similarly define  $err_S(f)$  in the source domain and the expected inter function loss  $err_T(f, g) = \mathbb{E}_{(\mathbf{x}, y) \sim \mathcal{P}_t} \mathcal{L}(g(\mathbf{x}), f(\mathbf{x}))$ .

The proxy  $\mathcal{P}_t^f$  over  $\Omega \times \mathcal{C}$  of  $\mathcal{P}_t$  w.r.t. to  $\mu_t$  and  $f$  is defined as:  $\mathcal{P}_t^f = (\mathbf{x}, f(\mathbf{x}))_{\mathbf{x} \sim \mu_t}$ .

We consider the following transport loss function:

$$W_1(\mathcal{P}_s, \mathcal{P}_t^f) = \inf_{\Pi \in \Pi(\mathcal{P}_s, \mathcal{P}_t^f)} \int_{(\Omega \times \mathcal{C})^2} \alpha d(\mathbf{x}_s, \mathbf{x}_t) + \mathcal{L}(y_s, y_t) d\Pi((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t)).$$

We now recall the definition of the theorem with all the assumptions.

**Theorem D.1** *Let  $\mathcal{H} \subset \mathcal{C}^\Omega$  be the hypothesis class where  $\Omega$  is a compact measurable space of finite dimension accompanied with a metric  $d$ , and  $\mathcal{C}$  is the output space. Let  $f$  be any labeling function of  $\mathcal{H}$ . Let  $\mathcal{P}_s, \mathcal{P}_t, \mathcal{P}_t^f$  be three probability distributions over  $\Omega \times \mathcal{C}$  with bounded support, with  $\mathcal{P}_t^f$  defined w.r.t. the marginal  $\mu_t$  of  $\mathcal{P}_t$  and  $f$ , accompanied with a sample of  $N_s$  labeled source instances drawn from  $\mathcal{P}_s$  and  $N_t$  unlabeled instances drawn from  $\mu_t$  and labeled by  $f$ , such that  $\mathcal{P}_s$  and  $\mathcal{P}_t^f$  and the associated samples follow the assumptions of Theorem E.1. Let  $\Pi^* = \operatorname{argmin}_{\Pi \in \Pi(\mathcal{P}_s, \mathcal{P}_t^f)} \int_{(\Omega \times \mathcal{C})^2} \alpha d(\mathbf{x}_s, \mathbf{x}_t) + \mathcal{L}(y_s, y_t) d\Pi((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t))$ . Let  $f^*$  be a Lipschitz labeling function of  $\mathcal{H}$ , that verifies the  $\phi$ -probabilistic transfer Lipschitzness (PTL) assumption with respect to  $\Pi^*$  and that minimizes the joint error  $err_S(f^*) + err_T(f^*)$  w.r.t all compatible PTL functions with  $\Pi^*$ . We assume the instance space  $\mathcal{X} \subseteq \Omega$  is bounded<sup>3</sup> such that*

<sup>3</sup>Since the input space is bounded by say a constant  $K$ :  $\|\mathbf{x}\| \leq K$ , since  $f^*$  is supposed  $l$ -Lipschitz, then we have for any  $\mathbf{x}_1, \mathbf{x}_2$ :  $|f(\mathbf{x}_1) - f(\mathbf{x}_2)| \leq l\|\mathbf{x}_1 - \mathbf{x}_2\| \leq 2lK = M$ .

$|f^*(\mathbf{x}_1) - f^*(\mathbf{x}_2)| \leq M$  for all  $\mathbf{x}_1, \mathbf{x}_2 \in \mathcal{X}^2$ . Let  $\mathcal{L}$  be any loss function symmetric,  $k$ -lipschitz and that satisfies the triangle inequality. Then, there exists,  $c'$  and  $N$ , such that for  $N_s > N$  and  $N_t > N$ , for all  $\lambda > 0$ , with  $\alpha = k\lambda$ , we have with probability at least  $1 - \delta$ :

$$\begin{aligned} \text{err}_T(f) &\leq W_1(\hat{\mathcal{P}}_s, \hat{\mathcal{P}}_t^f) + \sqrt{\frac{2}{c'} \log\left(\frac{2}{\delta}\right)} \left( \frac{1}{\sqrt{N_s}} + \frac{1}{\sqrt{N_t}} \right) \\ &\quad + \text{err}_S(f^*) + \text{err}_T(f^*) + kM\phi(\lambda). \end{aligned}$$

### Proof

$$\begin{aligned} \text{err}_T(f) &= E_{(\mathbf{x}, y) \sim P_t} \mathcal{L}(y, f(\mathbf{x})) \\ &\leq E_{(\mathbf{x}, y) \sim P_t} \mathcal{L}(y, f^*(\mathbf{x})) + \mathcal{L}(f^*(\mathbf{x}), f(\mathbf{x})) \\ &= E_{(\mathbf{x}, y) \sim P_t} \mathcal{L}(f(\mathbf{x}), f^*(\mathbf{x})) + \text{err}_T(f^*) \end{aligned} \tag{10}$$

$$= E_{(\mathbf{x}, y) \sim \mathcal{P}_t^f} \mathcal{L}(f(\mathbf{x}), f^*(\mathbf{x})) + \text{err}_T(f^*) \tag{11}$$

$$\begin{aligned} &= \text{err}_{T^f}(f^*) - \text{err}_S(f^*) + \text{err}_S(f^*) + \text{err}_T(f^*) \\ &\leq |\text{err}_{T^f}(f^*) - \text{err}_S(f^*)| + \text{err}_S(f^*) + \text{err}_T(f^*) \end{aligned} \tag{12}$$

Line (10) is due to the symmetry of the loss. Line (11) comes from the fact that:

$$E_{(\mathbf{x}, y) \sim \mathcal{P}_t} \mathcal{L}(f(\mathbf{x}), f^*(\mathbf{x})) = E_{(\mathbf{x}, f(\mathbf{x})) \sim \mathcal{P}_t^f} \mathcal{L}(f(\mathbf{x}), f^*(\mathbf{x})) \stackrel{\text{def}}{=} \text{err}_{T^f}(f^*(\mathbf{x})).$$

Now, we have

$$\begin{aligned} &|\text{err}_{T^f}(f^*) - \text{err}_S(f^*)| \\ &= \left| \int_{\Omega \times \mathcal{C}} \mathcal{L}(y, f^*(\mathbf{x})) (\mathcal{P}_t^f(\mathbf{X} = \mathbf{x}, Y = y) - \mathcal{P}_s(\mathbf{X} = \mathbf{x}, Y = y)) d\mathbf{x} dy \right| \\ &= \left| \int_{\Omega \times \mathcal{C}} \mathcal{L}(y, f^*(\mathbf{x})) d(\mathcal{P}_t^f - \mathcal{P}_s) \right| \\ &\leq \int_{(\Omega \times \mathcal{C})^2} \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_t)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) \right| d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) \end{aligned} \tag{13}$$

$$\begin{aligned} &= \int_{(\Omega \times \mathcal{C})^2} \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_t)) - \mathcal{L}(y_t^f, f^*(\mathbf{x}_s)) + \right. \\ &\quad \left. \mathcal{L}(y_t^f, f^*(\mathbf{x}_s)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) \right| d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) \\ &\leq \int_{(\Omega \times \mathcal{C})^2} \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_t)) - \mathcal{L}(y_t^f, f^*(\mathbf{x}_s)) \right| \\ &\quad + \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_s)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) \right| d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) \\ &\leq \int_{(\Omega \times \mathcal{C})^2} k |f^*(\mathbf{x}_t) - f^*(\mathbf{x}_s)| + \\ &\quad \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_s)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) \right| d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) \end{aligned} \tag{14}$$

$$\begin{aligned} &\leq k * M * \phi(\lambda) + \int_{(\Omega \times \mathcal{C})^2} k\lambda d(\mathbf{x}_t, \mathbf{x}_s) + \\ &\quad \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_s)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) \right| d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) \end{aligned} \tag{15}$$

$$\leq \int_{(\Omega \times \mathcal{C})^2} \alpha d(\mathbf{x}_s, \mathbf{x}_t) + \mathcal{L}(y_t^f, y_s) d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) + k * M * \phi(\lambda) \tag{16}$$

$$\leq \int_{(\Omega \times \mathcal{C})^2} \alpha d(\mathbf{x}_s, \mathbf{x}_t) + \mathcal{L}(y_s, y_t^f) d\Pi^*((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) + k * M * \phi(\lambda) \tag{17}$$

$$= W_1(\mathcal{P}_s, \mathcal{P}_t^f) + k * M * \phi(\lambda). \tag{18}$$

Line (13) is a consequence of the duality form of the Kantorovitch-Rubinstein theorem saying that for any coupling  $\Pi \in \Pi(P_s, P_t^f)$ , we have:

$$\begin{aligned} & \left| \int_{\Omega \times \mathcal{C}} \mathcal{L}(y, f^*(\mathbf{x})) d(\mathcal{P}_t^f - \mathcal{P}_s) \right| \\ &= \left| \int_{(\Omega \times \mathcal{C})^2} \mathcal{L}(y_t^f, f^*(\mathbf{x}_t)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) d\Pi((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)) \right| \\ &\leq \int_{(\Omega \times \mathcal{C})^2} \left| \mathcal{L}(y_t^f, f^*(\mathbf{x}_t)) - \mathcal{L}(y_s, f^*(\mathbf{x}_s)) \right| d\Pi((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t^f)). \end{aligned}$$

Since the inequality is true for any coupling, it is then also true for  $\Pi^*$ . Inequality (14) is due to the  $k$ -lipschitzness of the loss  $\mathcal{L}$  in its second argument. Inequality (15) uses the fact that  $f^*$  and  $\Pi^*$  verify the probabilistic transfer Lipschitzness property with probability  $1 - \phi(\lambda)$ , additionally, taking into account that the deviation between 2 instances with respect to  $f^*$  is bounded by  $M$  we have the additional term  $kM\phi(\lambda)$  that covers the regions where the PTL does not hold. (16) is obtained by the symmetry of  $d$ , the use of triangle inequality on  $\mathcal{L}$  and by replacing  $k\lambda$  by  $\alpha$ . Other inequalities above are due the use of triangle inequality or properties of the absolute value. The last line (18) is due to the definition of  $\Pi^*$ .

Now, note that by the use of triangle inequality:

$$W_1(\mathcal{P}_s, \mathcal{P}_t^f) \leq W_1(\mathcal{P}_s, \hat{\mathcal{P}}_s) + W_1(\hat{\mathcal{P}}_s, \hat{\mathcal{P}}_t^f) + W_1(\hat{\mathcal{P}}_t^f, \mathcal{P}_t^f) \quad (19)$$

$$\leq W_1(\hat{\mathcal{P}}_s, \hat{\mathcal{P}}_t^f) + \sqrt{\frac{2}{c'} \log\left(\frac{2}{\delta}\right)} \left( \frac{1}{\sqrt{N_s}} + \frac{1}{\sqrt{N_t}} \right). \quad (20)$$

Indeed, the cost function  $\mathcal{D}((\mathbf{x}_s, y_s), (\mathbf{x}_t, y_t)) = \alpha d(\mathbf{x}_1, \mathbf{x}_2) + \mathcal{L}(y_1, y_2)$  defines a distance over  $(\Omega \times \mathcal{L})^2$ , assuming that  $\mathcal{P}_s$  and  $\mathcal{P}_t^f$  have bounded support and the fact that our loss function is bounded, we can apply Theorem E.1 (presented below) on  $W_1(\mathcal{P}_s, \hat{\mathcal{P}}_s)$  and  $W_1(\hat{\mathcal{P}}_t^f, \mathcal{P}_t^f)$  above (with probability  $\delta/2$  each). The two settings may have different constants  $N$  and  $c'$  and we consider the maximum  $N$  and the minimum  $c'$  that comply with both cases.

Combining inequalities (12), (18), inequality (20) and the use of the union bound, the theorem holds with probability at least  $1 - \delta$  for any  $f \in \mathcal{H}$ .  $\square$

Note that, additionally to the analysis in the paper, a link can be made with classic generalization bounds when the two distributions are equal, *i.e.*  $\mathcal{P}_s = \mathcal{P}_t$ . Indeed, if we can choose  $f^*$  as the true labeling function on source/target domains such that  $f^*$  is strongly  $\phi$ -lipschitz w.r.t.  $\Pi^*$  (*i.e.*  $\phi(\lambda)$  is almost 0), then the bound is similar to a classic generalization bound: terms involving  $f^*$  are null and using the same sample for source and target  $d(\mathbf{x}_1, \mathbf{x}_2) = 0$  w.r.t the best alignment. Thus, it remains only the label loss which corresponds to a classic supervised learning loss.

## E Empirical concentration result for Wasserstein distance

We give now the result from Bolley and co-authors used in the previous section.

**Theorem E.1 (from [35], Theorem 1.1.)** *Let  $\mu$  be a probability measure in  $Z$  so that for some  $\alpha > 0$  we have for any  $\mathbf{z}' \int_{\mathbb{R}^d} e^{\alpha \text{dist}(\mathbf{z}, \mathbf{z}')^2} d\mu < \infty$  and  $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N \delta_{z_i}$  be the associated empirical measure defined on a sample of independent variables  $\{\mathbf{z}_i\}_{i=1}^N$  drawn from  $\mu$ . Then, for any  $d' > \dim(Z)$  and  $c' < c$ , there exists some constant  $N_0$  depending on  $d'$  and some square exponential moments of  $\mu$  such that for any  $\epsilon > 0$  and  $N \geq N_0 \max(\epsilon^{-(d'+2)}, 1)$ ,*

$$P[W_1(\mu, \hat{\mu}) > \epsilon] \leq \exp\left(-\frac{c'}{2} N \epsilon^2\right)$$

where  $c'$  can be calculated explicitly.

Note that  $c$  is such that  $\mu$  verifies for any measure  $\nu$  the Talagrand (transport) inequality  $T_1(c)$  :  $W_1(\mu, \nu) \leq \sqrt{\frac{2}{c} H(\nu|\mu)}$  with  $H$  is the relative entropy.  $T_1(c)$  holds when for some  $\alpha > 0$  and for any  $\mathbf{z}'$ :  $\int_{\mathbb{R}^d} e^{\alpha \text{dist}(\mathbf{z}, \mathbf{z}')^2} d\mu(\mathbf{z}) < \infty$ , and  $c$  can be found explicitly [35].