

Theoretical Analysis of Domain Adaptation with Optimal Transport

Ievgen Redko¹, Amaury Habrard², and Marc Sebban²

¹ Univ.Lyon, INSA-Lyon, Université Claude Bernard Lyon 1, UJM-Saint Etienne
CNRS, Inserm, CREATIS UMR 5220, U1206

F-69266, LYON, France

`ievgen.redko@creatis.insa-lyon.fr`

² Univ. Lyon, UJM-Saint-Etienne

CNRS, Lab. Hubert Curien UMR 5516

F-42023, SAINT-ETIENNE, France

`{amaury.habrard,marc.sebban}@univ.st-etienne.fr`

Abstract. Domain adaptation (DA) is an important and emerging field of machine learning that tackles the problem occurring when the distributions of training (source domain) and test (target domain) data are similar but different. This kind of learning paradigm is of vital importance for future advances as it allows a learner to generalize the knowledge across different tasks. Current theoretical results show that the efficiency of DA algorithms depends on their capacity of minimizing the divergence between source and target probability distributions. In this paper, we provide a theoretical study on the advantages that concepts borrowed from optimal transportation theory [17] can bring to DA. In particular, we show that the Wasserstein metric can be used as a divergence measure between distributions to obtain generalization guarantees for three different learning settings: (i) classic DA with unsupervised target data (ii) DA combining source and target labeled data, (iii) multiple source DA. Based on the obtained results, we motivate the use of the regularized optimal transport and provide some algorithmic insights for multi-source domain adaptation. We also show when this theoretical analysis can lead to tighter inequalities than those of other existing frameworks. We believe that these results open the door to novel ideas and directions for DA.

Keywords: domain adaptation, generalization bounds, optimal transport.

1 Introduction

Many results in statistical learning theory study the problem of estimating the probability that a hypothesis chosen from a given hypothesis class can achieve a small true risk. This probability is often expressed in the form of generalization bounds on the true risk obtained using concentration inequalities with respect to (w.r.t.) some hypothesis class. Classic generalization bounds make the assumption that training and test data follow the same distribution. This assumption, however, can be violated in many real-world applications (e.g., in computer vision, language processing or speech recognition) where training and test data actually follow a related but different probability distribution. One may think of an example, where a spam filter is learned based on the abundant annotated

data collected for one user and is further applied for newly registered user with different preferences. In this case, the performance of the spam filter will deteriorate as it does not take into account the mismatch between the underlying probability distributions. The need for algorithms tackling this problem has led to the emergence of a new field in machine learning called domain adaptation (DA), subfield of transfer learning [18], where the source (training) and target (test) distributions are not assumed to be the same. From a theoretical point of view, existing generalization guarantees for DA are expressed in the form of bounds over the target risk involving the source risk, a divergence between domains and a term λ evaluating the capability of the considered hypothesis class to solve the problem, often expressed as a joint error of the ideal hypothesis between the two domains. In this context, minimizing the divergence between distributions is a key factor for the potential success of DA algorithms. Among the most striking results, existing generalization bounds based on the H-divergence [3] or the discrepancy distance [15] have also an interesting property of being able to link the divergence between the probability distributions of two domains w.r.t. the considered class of hypothesis.

Despite their advantages, the above mentioned divergences do not directly take into account the geometry of the data distribution. Recently, [6,7] has proposed to tackle this drawback by solving the DA problem using ideas from optimal transportation (OT) theory. Their paper proposes an algorithm that aims to reduce the divergence between two domains by minimizing the Wasserstein distance between their distributions. This idea has a very appealing and intuitive interpretation based on the transport of one domain to another. The transportation plan solving OT problem takes into account the geometry of the data by means of an associated cost function which is based on the Euclidean distance between examples. Furthermore, it is naturally defined as an infimum problem over all feasible solutions. An interesting property of this approach is that the resulting solution given by a joint probability distribution allows one to obtain the new projection of the instances of one domain into another directly without being restricted to a particular hypothesis class. This independence from the hypothesis class means that this solution not only ensures successful adaptation but also influences the capability term λ . While showing very promising experimental results, it turns out that this approach, however, has no theoretical guarantees. This paper aims to bridge this gap by presenting contributions covering three DA settings: (i) classic unsupervised DA where the learner has only access to labeled source data and unsupervised target instances, (ii) DA where one has access to labeled data from both source and target domains, (iii) multi-source DA where labeled instances for a set of distinct source domains (more than 2) are available. We provide new theoretical guarantees in the form of generalization bounds for these three settings based on the Wasserstein distance thus justifying its use in DA. According to [26], the Wasserstein distance is rather strong and can be combined with smoothness bounds to obtain convergences in other distances. This important advantage of Wasserstein distance leads to tighter bounds in comparison to other state-of-the-art results and is more computationally attractive.

The rest of this paper is organized as follows: Section 2 is devoted to the presentation of optimal transport and its application in DA. In Section 3, we present the generalization bounds for DA with the Wasserstein distance for both single- and multi-source learning scenarios. Finally, we conclude our paper in Section 4.

2 Definitions and notations

In this section, we first present the formalization of the Monge-Kantorovich [13] optimization problem and show how optimal transportation problem found its application in DA.

2.1 Optimal transport

Optimal transportation theory was first introduced in [17] to study the problem of resource allocation. Assuming that we have a set of factories and a set of mines, the goal of optimal transportation is to move the ore from mines to factories in an optimal way, i.e., by minimizing the overall transport cost. More formally, let $\Omega \subseteq \mathbb{R}^d$ be a measurable space and denote by $\mathcal{P}(\Omega)$ the set of all probability measures over Ω . Given two probability measures $\mu_S, \mu_T \in \mathcal{P}(\Omega)$, the Monge-Kantorovich problem consists in finding a probabilistic coupling γ defined as a joint probability measure over $\Omega \times \Omega$ with marginals μ_S and μ_T for all $x, y \in \Omega$ that minimizes the cost of transport w.r.t. some function $c : \Omega \times \Omega \rightarrow \mathbb{R}_+$:

$$\begin{aligned} & \arg \min_{\gamma} \int_{\Omega_1 \times \Omega_2} c(\mathbf{x}, \mathbf{y})^p d\gamma(\mathbf{x}, \mathbf{y}) \\ & \text{s.t. } \mathbf{P}^{\Omega_1} \# \gamma = \mu_S, \mathbf{P}^{\Omega_2} \# \gamma = \mu_T, \end{aligned}$$

where \mathbf{P}^{Ω_i} is the projection over Ω_i and $\#$ denotes the pushforward measure. This problem admits a unique solution γ_0 which allows us to define the Wasserstein distance of order p between μ_S and μ_T for any $p \in [1; +\infty]$ as follows:

$$W_p^p(\mu_S, \mu_T) = \inf_{\gamma \in \Pi(\mu_S, \mu_T)} \int_{\Omega \times \Omega} c(\mathbf{x}, \mathbf{y})^p d\gamma(\mathbf{x}, \mathbf{y}),$$

where $c : \Omega \times \Omega \rightarrow \mathbb{R}^+$ is a cost function for transporting one unit of mass \mathbf{x} to \mathbf{y} and $\Pi(\mu_S, \mu_T)$ is a collection of all joint probability measures on $\Omega \times \Omega$ with marginals μ_S and μ_T .

Remark 1. In what follows, we consider only the case $p = 1$ but all the obtained results can be easily extended to the case $p > 1$ using Hölder inequality implying for every $p \leq q \Rightarrow W_p \leq W_q$.

In the discrete case, when one deals with empirical measures $\hat{\mu}_S = \frac{1}{N_S} \sum_{i=1}^{N_S} \delta_{x_S^i}$ and $\hat{\mu}_T = \frac{1}{N_T} \sum_{i=1}^{N_T} \delta_{x_T^i}$ represented by the uniformly weighted sums of N_S and N_T Diracs with mass at locations x_S^i and x_T^i respectively, Monge-Kantorovich problem is defined in terms of the inner product between the coupling matrix γ and the cost matrix C :

$$W_1(\hat{\mu}_S, \hat{\mu}_T) = \min_{\gamma \in \Pi(\hat{\mu}_S, \hat{\mu}_T)} \langle C, \gamma \rangle_F$$

where $\langle \cdot, \cdot \rangle_F$ is the Frobenius dot product, $\Pi(\hat{\mu}_S, \hat{\mu}_T) = \{\gamma \in \mathbb{R}_+^{N_S \times N_T} | \gamma \mathbf{1} = \hat{\mu}_S, \gamma^T \mathbf{1} = \hat{\mu}_T\}$ is a set of doubly stochastic matrices and C is a dissimilarity matrix, i.e., $C_{ij} =$




Fig. 1: Blue points are generated to lie inside a square with a side length equal to 1. Red points are generated inside an annulus containing the square. Solution of the regularized optimal transport problem is visualized by plotting dashed and solid lines that correspond to the large and small values given by the optimal coupling matrix γ .

$c(x_S^i, x_T^j)$, defining the energy needed to move a probability mass from x_S^i to x_T^j . Figure 1 shows how the solution of optimal transport between two point-clouds can look like.

It turns out that the Wasserstein distance has been successfully used in various applications, for instance: computer vision [22], texture analysis [21], tomographic reconstruction [12] and clustering [9]. The huge success of algorithms based on this distance is due to [8] who introduced an entropy-regularized version of optimal transport that can be optimized efficiently using matrix scaling algorithm. We are now ready to present the application of OT to DA below.

2.2 Domain adaptation and optimal transport

The problem of DA is formalized as follows: we define a domain as a pair consisting of a distribution μ_D on Ω and a labeling function $f_D : \Omega \rightarrow [0, 1]$. A hypothesis class H is a set of functions so that $\forall h \in H, h : \Omega \rightarrow \{0, 1\}$.

Definition 1. Given a convex loss-function l , the probability according to the distribution μ_D that a hypothesis $h \in H$ disagrees with a labeling function f_D (which can also be a hypothesis) is defined as

$$\epsilon_D(h, f_D) = \mathbb{E}_{x \sim \mu_D} [l(h(x), f_D(x))].$$

When the source and target error functions are defined w.r.t. h and f_S or f_T , we use the shorthand $\epsilon_S(h, f_S) = \epsilon_S(h)$ and $\epsilon_T(h, f_T) = \epsilon_T(h)$. We further denote by $\langle \mu_S, f_S \rangle$ the source domain and $\langle \mu_T, f_T \rangle$ the target domain. The ultimate goal of DA then is to learn a good hypothesis h in $\langle \mu_S, f_S \rangle$ that has a good performance in $\langle \mu_T, f_T \rangle$.

In unsupervised DA problem, one usually has access to a set of source data instances $\mathbf{X}_S = \{\mathbf{x}_S^i \in \mathbb{R}^d\}_{i=1}^{N_S}$ associated with labels $\{y_S^i\}_{i=1}^{N_S}$ and a set of unlabeled target data instances $\mathbf{X}_T = \{\mathbf{x}_T^i \in \mathbb{R}^d\}_{i=1}^{N_T}$. Contrary to the classic learning paradigm, unsupervised DA assumes that the marginal distributions of \mathbf{X}_S and \mathbf{X}_T are different and given by $\mu_S, \mu_T \in \mathcal{P}(\Omega)$.

For the first time, optimal transportation problem was applied to DA in [6,7]. The main underlying idea of their work is to find a coupling matrix that efficiently transports source samples to target ones by solving the following optimization problem :

$$\gamma_o = \arg \min_{\gamma \in \Pi(\hat{\mu}_S, \hat{\mu}_T)} \langle C, \gamma \rangle_F.$$

Once the optimal coupling γ_o is found, source samples \mathbf{X}_S can be transformed into target aligned source samples $\hat{\mathbf{X}}_S$ using the following equation

$$\hat{\mathbf{X}}_S = \text{diag}((\gamma_o \mathbf{1})^{-1}) \gamma_o \mathbf{X}_T.$$

The use of Wasserstein distance here has an important advantage over other distances used in DA (see Section 3.4) as it preserves the topology of the data and admits a rather efficient estimation as mentioned above. Furthermore, as shown in [6,7], it improves current state-of-the-art results on benchmark computer vision data sets and has a very appealing intuition behind.

3 Generalization bounds with Wasserstein distance

In this section, we introduce generalization bounds for the target error when the divergence between tasks' distributions is measured by the Wasserstein distance.

3.1 A bound relating the source and target error

We first consider the case of unsupervised DA where no labelled data are available in the target domain. We start with a lemma that relates the Wasserstein metric with the source and target error functions for an arbitrary pair of hypothesis. Then, we show how the target error can be bounded by the Wasserstein distance for empirical measures. We first present the Lemma that introduces Wasserstein distance to relate the source and target error functions in a Reproducing Kernel Hilbert Space.

Lemma 1. *Let $\mu_S, \mu_T \in \mathcal{P}(\Omega)$ be two probability measures on \mathbb{R}^d . Assume that the cost function $c(\mathbf{x}, \mathbf{y}) = \|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_{\mathcal{H}_{k_l}}$, where \mathcal{H}_{k_l} is a Reproducing Kernel Hilbert Space (RKHS) equipped with kernel $k_l : \Omega \times \Omega \rightarrow \mathbb{R}$ induced by $\phi : \Omega \rightarrow \mathcal{H}_{k_l}$ and $k_l(\mathbf{x}, \mathbf{y}) = \langle \phi(\mathbf{x}), \phi(\mathbf{y}) \rangle_{\mathcal{H}_{k_l}}$. Assume further that the loss function $l_{h,f} : x \rightarrow l(h(x), f(x))$ is convex, symmetric, bounded, obeys the triangular equality and has the parametric form $|h(x) - f(x)|^q$ for some $q > 0$. Assume also that kernel k_l in the RKHS \mathcal{H}_{k_l} is square-root integrable w.r.t. both μ_S, μ_T for all $\mu_S, \mu_T \in \mathcal{P}(\Omega)$ where Ω is separable and $0 \leq k_l(\mathbf{x}, \mathbf{y}) \leq K, \forall \mathbf{x}, \mathbf{y} \in \Omega$. Then the following holds*

$$\epsilon_T(h, h') \leq \epsilon_S(h, h') + W_1(\mu_S, \mu_T)$$

for every hypothesis h', h .

Proof. As this Lemma plays a key role in the following sections, we give its proof here. We assume that $l_{h,f} : x \rightarrow l(h(x), f(x))$ in the definition of $\epsilon(h)$ is a convex

loss-function defined $\forall h, f \in \mathcal{F}$ where \mathcal{F} is a unit ball in the RKHS \mathcal{H}_k . Considering that $h, f \in \mathcal{F}$, the loss function l is a non-linear mapping of the RKHS \mathcal{H}_k for the family of losses $l(h(x), f(x)) = |h(x) - f(x)|^q$ ³. Using results from [23], one may show that $l_{h,f}$ also belongs to the RKHS \mathcal{H}_{k_l} admitting the reproducing kernel k_l and that its norm obeys the following inequality:

$$\|l_{h,f}\|_{\mathcal{H}_{k_l}}^2 \leq \|h - f\|_{\mathcal{H}_k}^{2q}.$$

This result gives us two important properties of $l_{f,h}$ that we use further:

- $l_{h,f}$ belongs to the RKHS that allows us to use the reproducing property;
- the norm $\|l_{h,f}\|_{\mathcal{H}_{k_l}}$ is bounded.

For simplicity, we can assume that $\|l_{h,f}\|_{\mathcal{H}_{k_l}}$ is bounded by 1. This assumption can be verified by imposing the appropriate bounds on the norms of h and f and is easily extendable to the case when $\|l_{h,f}\|_{\mathcal{H}_{k_l}} \leq M$ by scaling as explained in [15, Proposition 2]. We also note that q does not necessarily have to appear in the final result as we seek to bound the norm of l and not to give an explicit expression for it in terms of $\|h\|_{\mathcal{H}_k}$, $\|f\|_{\mathcal{H}_k}$ and q . Now the error function defined above can be also expressed in terms of the inner product in the corresponding Hilbert space, i.e.⁴:

$$\epsilon_S(h, f_S) = \mathbb{E}_{x \sim \mu_S} [l(h(x), f_S(x))] = \mathbb{E}_{x \sim \mu_S} [\langle \phi(x), l \rangle_{\mathcal{H}}].$$

We define the target error in the same manner:

$$\epsilon_T(h, f_T) = \mathbb{E}_{y \sim \mu_T} [l(h(y), f_T(y))] = \mathbb{E}_{y \sim \mu_T} [\langle \phi(y), l \rangle_{\mathcal{H}}].$$

With the definitions introduced above, the following holds:

$$\begin{aligned} \epsilon_T(h, h') &= \epsilon_T(h, h') + \epsilon_S(h, h') - \epsilon_S(h, h') \\ &= \epsilon_S(h, h') + \mathbb{E}_{y \sim \mu_T} [\langle \phi(y), l \rangle_{\mathcal{H}}] - \mathbb{E}_{x \sim \mu_S} [\langle \phi(x), l \rangle_{\mathcal{H}}] \\ &= \epsilon_S(h, h') + \langle \mathbb{E}_{y \sim \mu_T} [\phi(y)] - \mathbb{E}_{x \sim \mu_S} [\phi(x)], l \rangle_{\mathcal{H}} \\ &\leq \epsilon_S(h, h') + \|l\|_{\mathcal{H}} \|\mathbb{E}_{y \sim \mu_T} [\phi(y)] - \mathbb{E}_{x \sim \mu_S} [\phi(x)]\|_{\mathcal{H}} \\ &\leq \epsilon_S(h, h') + \left\| \int_{\Omega} \phi d(\mu_S - \mu_T) \right\|_{\mathcal{H}}. \end{aligned}$$

Second line is obtained by using the reproducing property applied to l , third line follows from the properties of the expected value. Fourth line here is due to the properties of the inner-product while fifth line is due to $\|l_{h,f}\|_{\mathcal{H}} \leq 1$. Now using the definition of the joint distribution we have the following:

$$\left\| \int_{\Omega} \phi d(\mu_S - \mu_T) \right\|_{\mathcal{H}} = \left\| \int_{\Omega \times \Omega} (\phi(\mathbf{x}) - \phi(\mathbf{y})) d\gamma(\mathbf{x}, \mathbf{y}) \right\|_{\mathcal{H}}$$

³ If $h, f \in \mathcal{H}$ then $h - f \in \mathcal{H}$ implying that $l(h(x), f(x)) = |h(x) - f(x)|^q$ is a nonlinear transform for $h - f \in \mathcal{H}$.

⁴ For the sake of simplicity, we will further write \mathcal{H} meaning \mathcal{H}_{k_l} and l meaning $l_{f,h}$.

$$\leq \int_{\Omega \times \Omega} \|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_{\mathcal{H}} d\gamma(\mathbf{x}, \mathbf{y}).$$

As the last inequality holds for any γ , we obtain the final result by taking the infimum over γ from the right-hand side, i.e.:

$$\int_{\Omega} \phi d(\mu_S - \mu_T) \|_{\mathcal{H}} \leq \inf_{\gamma \in \Pi(\mu_S, \mu_T)} \int_{\Omega \times \Omega} \|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_{\mathcal{H}} d\gamma(\mathbf{x}, \mathbf{y}).$$

which gives

$$\epsilon_T(h, h') \leq \epsilon_S(h, h') + W_1(\mu_S, \mu_T).$$

□

Remark 2. We note that the functional form of the loss-function $l(h(x), f(x)) = |h(x) - f(x)|^q$ is just an example that was used as the basis for the proof. According to [23, Appendix 2], we may also consider more general nonlinear transformations of h and f that satisfy the assumption imposed on $l_{h,f}$ above. These transformations may include a product of hypothesis and labeling functions and thus the proposed results is valid for hinge-loss too.

This lemma makes use of the Wasserstein distance to relate the source and target errors. The assumption made here is to specify that the cost function $c(\mathbf{x}, \mathbf{y}) = \|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_{\mathcal{H}}$. While it may seem too restrictive, this assumption is, in fact, not that strong. Using the properties of the inner-product, one has:

$$\begin{aligned} \|\phi(\mathbf{x}) - \phi(\mathbf{y})\|_{\mathcal{H}} &= \sqrt{\langle \phi(\mathbf{x}) - \phi(\mathbf{y}), \phi(\mathbf{x}) - \phi(\mathbf{y}) \rangle_{\mathcal{H}}} \\ &= \sqrt{k(\mathbf{x}, \mathbf{x}) - 2k(\mathbf{x}, \mathbf{y}) + k(\mathbf{y}, \mathbf{y})}. \end{aligned}$$

Now it can be shown that for any given positive-definite kernel k there is a distance c (used as a cost function in our case) that generates it and vice versa (see Lemma 12 from [24]).

In order to prove our next theorem, we present first an important result showing the convergence of the empirical measure $\hat{\mu}$ to its true associated measure w.r.t. the Wasserstein metric. This concentration guarantee allows us to propose generalization bounds based on the Wasserstein distance for finite samples rather than true population measures. Following [4], it can be specialized for the case of W_1 as follows⁵

Theorem 1 ([4], Theorem 1.1). *Let μ be a probability measure in \mathbb{R}^d so that for some $\alpha > 0$, we have that $\int_{\mathbb{R}^d} e^{\alpha \|x\|^2} d\mu < \infty$ and $\hat{\mu} = \frac{1}{N} \sum_{i=1}^N \delta_{x_i}$ be its associated empirical measure defined on a sample of independent variables $\{x_i\}_{i=1}^N$ drawn from μ . Then for any $d' > d$ and $\zeta' < \sqrt{2}$ there exists some constant N_0 depending on d' and some square exponential moment of μ such that for any $\varepsilon > 0$ and $N \geq N_0 \max(\varepsilon^{-(d'+2)}, 1)$*

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

where d', ζ' can be calculated explicitly.

⁵ We present the original version of this Theorem in the Supplementary material.

The convergence guarantee of this theorem can be further strengthened as shown in [11] but we prefer this version for the ease of reading. We can now use it in combination with the previous Lemma to prove the following theorem.

Theorem 2. *Under the assumptions of Lemma 1, let \mathbf{X}_S and \mathbf{X}_T be two samples of size N_S and N_T drawn i.i.d. from μ_S and μ_T respectively. Let $\hat{\mu}_S = \frac{1}{N_S} \sum_{i=1}^{N_S} \delta_{x_S^i}$ and $\hat{\mu}_T = \frac{1}{N_T} \sum_{i=1}^{N_T} \delta_{x_T^i}$ be the associated empirical measures. Then for any $d' > d$ and $\varsigma' < \sqrt{2}$ there exists some constant N_0 depending on d' such that for any $\delta > 0$ and $\min(N_S, N_T) \geq N_0 \max(\delta^{-(d'+2)}, 1)$ with probability at least $1 - \delta$ for all h the following holds:*

$$\epsilon_T(h) \leq \epsilon_S(h) + W_1(\hat{\mu}_S, \hat{\mu}_T) + \sqrt{2 \log \left(\frac{1}{\delta} \right) / \varsigma'} \left(\sqrt{\frac{1}{N_S}} + \sqrt{\frac{1}{N_T}} \right) + \lambda,$$

where λ is the combined error of the ideal hypothesis h^* that minimizes the combined error of $\epsilon_S(h) + \epsilon_T(h)$.

Proof.

$$\begin{aligned} \epsilon_T(h) &\leq \epsilon_T(h^*) + \epsilon_T(h^*, h) = \epsilon_T(h^*) + \epsilon_S(h, h^*) + \epsilon_T(h^*, h) - \epsilon_S(h, h^*) \\ &\leq \epsilon_T(h^*) + \epsilon_S(h, h^*) + W_1(\mu_S, \mu_T) \\ &\leq \epsilon_T(h^*) + \epsilon_S(h) + \epsilon_S(h^*) + W_1(\mu_S, \mu_T) \\ &= \epsilon_S(h) + W_1(\mu_S, \mu_T) + \lambda \\ &\leq \epsilon_S(h) + W_1(\mu_S, \hat{\mu}_S) + W_1(\hat{\mu}_S, \mu_T) + \lambda \\ &\leq \epsilon_S(h) + \sqrt{2 \log \left(\frac{1}{\delta} \right) / N_S \varsigma'} + W_1(\hat{\mu}_S, \hat{\mu}_T) + W_1(\hat{\mu}_T, \mu_T) + \lambda \\ &\leq \epsilon_S(h) + W_1(\hat{\mu}_S, \hat{\mu}_T) + \lambda + \sqrt{2 \log \left(\frac{1}{\delta} \right) / \varsigma'} \left(\sqrt{\frac{1}{N_S}} + \sqrt{\frac{1}{N_T}} \right). \end{aligned}$$

Second and fourth lines are obtained using the triangular inequality applied to the error function. Third inequality is a consequence of Lemma 1. Fifth line follows from the definition of λ , sixth, seventh and eighth lines use the fact that Wasserstein metric is a proper distance and Theorem 1. \square

A first immediate consequence of this theorem is that it justifies the use of the optimal transportation in DA context. However, we would like to clarify the fact that the bound does not suggest that minimization of the Wasserstein distance can be done independently from the minimization of the source error nor it says that the joint error given by the lambda term becomes small. First, it is clear that the result of W_1 minimization provides a transport of the source to the target such as W_1 becomes small when computing the distance between newly transported sources and target instances. Under the hypothesis that class labeling is preserved by transport, i.e. $P_{\text{source}}(y|x_s) = P_{\text{target}}(y|\text{Transport}(x_s))$, the adaptation can be possible by minimizing W_1 only. However, this is not a reasonable assumption in practice. Indeed, by minimizing the W_1 distance only, it is possible that

the obtained transformation transports one positive and one negative source instance to the same target point and then the empirical source error cannot be properly minimized. Additionally, the joint error will be affected since no classifier will be able to separate these source points. We can also think of an extreme case where the positive source examples are transported to negative target instances, in that case the joint error λ will be dramatically affected. A solution is then to regularize the transport to help the minimization of the source error which can be seen as a kind of joint optimization. This idea was partially implemented as a class-labeled regularization term added to the original optimal transport formulation in [6,7] and showed good empirical results in practice. The proposed regularized optimization problem reads

$$\min_{\gamma \in \Pi(\hat{\mu}_S, \hat{\mu}_T)} \langle C, \gamma \rangle_F - \frac{1}{\lambda} E(\gamma) + \eta \sum_j \sum_{\mathcal{L}} \|\gamma(I_{\mathcal{L}}, j)\|_q^p.$$

Here, the second term $E(\gamma) = -\sum_{i,j}^{N_S, N_T} \gamma_{i,j} \log(\gamma_{i,j})$ is the regularization term that allows one to solve optimal transportation problem efficiently using Sinkhorn-Knopp matrix scaling algorithm [25]. Second regularization term $\eta \sum_j \sum_c \|\gamma(I_c, j)\|_q^p$ is used to restrict source examples of different classes to be transported to the same target examples by promoting group sparsity in the matrix γ thanks to $\|\cdot\|_q^p$ with $q = 1$ and $p = \frac{1}{2}$. In some way, this regularization term influences the capability term by ensuring the existence of a good hypothesis that will be able to be discriminant on both source and target domains data. Another recent paper of [28] also suggests that transport regularization is important for the use of OT in domain adaptation tasks. Thus, we conclude that the regularized transport formulations such as the one of [6,7] can be seen as algorithmic solutions for controlling the trade-off between the terms of the bound.

Assuming that $\epsilon_S(h)$ is properly minimized, only λ and the Wasserstein distance between empirical measures defined on the source and target samples have an impact on the potential success of adaptation. Furthermore, the fact that the Wasserstein distance is defined in terms of the optimal coupling used to solve the DA problem and is not restricted to any particular hypothesis class directly influences λ as discussed above. We now proceed to give similar bounds for the case where one has access to some labeled instances in the target domain.

3.2 A learning bound for the combined error

In semi-supervised DA, when we have access to an additional small set of labeled instances in the target domain, the goal is often to find a trade-off between minimizing the source and the target errors depending on the number of instances available in each domain and their mutual correlation. Let us now assume that we possess βn instances drawn independently from μ_T and $(1 - \beta)n$ instances drawn independently from μ_S and labeled by f_S and f_T , respectively. In this case, the empirical combined error [2] is defined as a convex combination of errors on the source and target training data:

$$\hat{\epsilon}_\alpha(h) = \alpha \hat{\epsilon}_T(h) + (1 - \alpha) \hat{\epsilon}_S(h),$$

where $\alpha \in [0, 1]$.

The use of the combined error is motivated by the fact that if the number of instances in the target sample is small compared to the number of instances in the source domain (which is usually the case in DA), minimizing only the target error may not be appropriate. Instead, one may want to find a suitable value of α that ensures the minimum of $\hat{\epsilon}_\alpha(h)$ w.r.t. a given hypothesis h . We now prove a theorem for the combined error similar to the one presented in [2].

Theorem 3. *Under the assumptions of Theorem 2 and Lemma 1, let D be a labeled sample of size n with βn points drawn from μ_T and $(1 - \beta)n$ from μ_S with $\beta \in (0, 1)$, and labeled according to f_S and f_T . If \hat{h} is the empirical minimizer of $\hat{\epsilon}_\alpha(h)$ and $h_T^* = \min_h \epsilon_T(h)$ then for any $\delta \in (0, 1)$ with probability at least $1 - \delta$ (over the choice of samples),*

$$\epsilon_T(\hat{h}) \leq \epsilon_T(h_T^*) + c_1 + 2(1 - \alpha)(W_1(\hat{\mu}_S, \hat{\mu}_T) + \lambda + c_2),$$

where

$$c_1 = 2\sqrt{\frac{2K \left(\frac{(1-\alpha)^2}{1-\beta} + \frac{\alpha^2}{\beta} \right) \log(2/\delta)}{n}} + 4\sqrt{K/n} \left(\frac{\alpha}{n\beta\sqrt{\beta}} + \frac{(1-\alpha)}{n(1-\beta)\sqrt{1-\beta}} \right),$$

$$c_2 = \sqrt{2 \log\left(\frac{1}{\delta}\right) / \varsigma'} \left(\sqrt{\frac{1}{N_S}} + \sqrt{\frac{1}{N_T}} \right).$$

Proof.

$$\begin{aligned} \epsilon_T(\hat{h}) &\leq \epsilon_\alpha(\hat{h}) + (1 - \alpha)(W_1(\mu_S, \mu_T) + \lambda) \\ &\leq \hat{\epsilon}_\alpha(\hat{h}) + \sqrt{\frac{2K \left(\frac{(1-\alpha)^2}{1-\beta} + \frac{\alpha^2}{\beta} \right) \log(2/\delta)}{n}} + (1 - \alpha)(W_1(\mu_S, \mu_T) + \lambda) \\ &\quad + 2\sqrt{K/n} \left(\frac{\alpha}{n\beta\sqrt{\beta}} + \frac{(1-\alpha)}{n(1-\beta)\sqrt{1-\beta}} \right) \\ &\leq \hat{\epsilon}_\alpha(h_T^*) + \sqrt{\frac{2K \left(\frac{(1-\alpha)^2}{1-\beta} + \frac{\alpha^2}{\beta} \right) \log(2/\delta)}{n}} + (1 - \alpha)(W_1(\mu_S, \mu_T) + \lambda) \\ &\quad + 2\sqrt{K/n} \left(\frac{\alpha}{n\beta\sqrt{\beta}} + \frac{(1-\alpha)}{n(1-\beta)\sqrt{1-\beta}} \right) \\ &\leq \epsilon_\alpha(h_T^*) + 2\sqrt{\frac{2K \left(\frac{(1-\alpha)^2}{1-\beta} + \frac{\alpha^2}{\beta} \right) \log(2/\delta)}{n}} + (1 - \alpha)(W_1(\mu_S, \mu_T) + \lambda) \\ &\quad + 4\sqrt{K/n} \left(\frac{\alpha}{n\beta\sqrt{\beta}} + \frac{(1-\alpha)}{n(1-\beta)\sqrt{1-\beta}} \right) \end{aligned}$$

$$\begin{aligned}
&\leq \epsilon_T(h_T^*) + 2\sqrt{\frac{2K\left(\frac{(1-\alpha)^2}{1-\beta} + \frac{\alpha^2}{\beta}\right)\log(2/\delta)}{n}} + 2(1-\alpha)(W_1(\mu_S, \mu_T) + \lambda) \\
&\quad + 4\sqrt{K/n}\left(\frac{\alpha}{n\beta\sqrt{\beta}} + \frac{(1-\alpha)}{n(1-\beta)\sqrt{1-\beta}}\right) \\
&\leq \epsilon_T(h_T^*) + c_1 + 2(1-\alpha)(W_1(\hat{\mu}_S, \hat{\mu}_T) + \lambda + c_2).
\end{aligned}$$

The proof follows the standard theory of uniform convergence for empirical risk minimizers where lines 1 and 5 are obtained by observing that $|\epsilon_\alpha(h) - \epsilon_T(h)| = |\alpha\epsilon_T(h) + (1-\alpha)\epsilon_S(h) - \epsilon_T(h)| = |(1-\alpha)(\epsilon_S(h) - \epsilon_T(h))| \leq (1-\alpha)(W_1(\mu_T, \mu_S) + \lambda)$ where the last inequality comes from line 4 of the proof of Theorem 2, line 3 follows from the definition of \hat{h} and h_T^* and line 6 is a consequence of Theorem 1. Finally, lines 2 and 4 are obtained based on the concentration inequality obtained for $\epsilon_\alpha(h)$. Due to the lack of space, we put this result in the Supplementary material. \square

This theorem shows that the best hypothesis that takes into account both source and target labeled data (i.e., $0 \leq \alpha < 1$) performs at least as good as the best hypothesis learned on target data instances alone ($\alpha = 1$). This result agrees well with the intuition that semi-supervised DA approaches should be at least as good as unsupervised ones.

4 Multi-source domain adaptation

We now consider the case where not one but many source domains are available during the adaptation. More formally, we define N different source domains (where T can either be or not a part of this set). For each source j , we have a labelled sample S_j of size $n_j = \beta_j n$ ($\sum_{j=1}^N \beta_j = 1$, $\sum_{j=1}^N n_j = n$) drawn from the associated unknown distribution μ_{S_j} and labelled by f_j . We now consider the empirical weighted multi-source error of a hypothesis h defined for some vector $\alpha = \{\alpha_1, \dots, \alpha_N\}$ as follows:

$$\hat{\epsilon}_\alpha(h) = \sum_{j=1}^N \alpha_j \hat{\epsilon}_{S_j}(h),$$

where $\sum_{j=1}^N \alpha_j = 1$ and each α_j represents the weight of the source domain S_j .

In what follows, we show that generalization bounds obtained for the weighted error give some interesting insights into the application of the Wasserstein distance to multi-source DA problems.

Theorem 4. *With the assumptions from Theorem 2 and Lemma 1, let S be a sample of size n , where for each $j \in \{1, \dots, N\}$, $\beta_j n$ points are drawn from μ_{S_j} and labelled according to f_j . If \hat{h}_α is the empirical minimizer of $\hat{\epsilon}_\alpha(h)$ and $h_T^* = \min_h \epsilon_T(h)$ then for any fixed α and $\delta \in (0, 1)$ with probability at least $1 - \delta$ (over the choice of samples),*

$$\epsilon_T(\hat{h}_\alpha) \leq \epsilon_T(h_T^*) + c_1 + 2 \sum_{j=1}^N \alpha_j (W_1(\hat{\mu}_j, \hat{\mu}_T) + \lambda_j + c_2),$$

where

$$c_1 = 2\sqrt{\frac{2K \sum_{j=1}^N \frac{\alpha_j^2}{\beta_j} \log(2/\delta)}{n}} + 2\sqrt{\sum_{j=1}^N \frac{K\alpha_j}{\beta_j n}},$$

$$c_2 = \sqrt{2 \log\left(\frac{1}{\delta}\right)} / \zeta' \left(\sqrt{\frac{1}{N_{S_j}}} + \sqrt{\frac{1}{N_T}} \right),$$

where $\lambda_j = \min_h (\epsilon_{S_j}(h) + \epsilon_T(h))$ represents the joint error for each source domain j .

Proof. The proof of this Theorem is very similar to the proof of Theorem 4. The final result is obtained by applying the concentration inequality for $\epsilon_\alpha(h)$ (instead of those used for $\epsilon_\alpha(\hat{h})$ in the proof of Theorem 4) and by using the following inequality that can be obtained easily by following the principle of the proof of [2, Theorem 4]:

$$|\epsilon_\alpha(h) - \epsilon_T(h)| \leq \sum_{j=1}^N \alpha_j (W_1(\mu_j, \mu_T) + \lambda_j),$$

where $\lambda_j = \min_h (\epsilon_{S_j}(h) + \epsilon_T(h))$. For the sake of completeness, we present the concentration inequality for $\epsilon_\alpha(h)$ in the Supplementary material. \square

While the results for multi-source DA may look like a trivial extension of the theoretical guarantees for the case of two domains, they can provide a very fruitful implication on their own. As in the previous case, we consider that the potential term that should be minimized in this bound by a given multi-source DA algorithm is the term $\sum_{j=1}^N \alpha_j W_1(\hat{\mu}_j, \hat{\mu}_T)$.

Assume that $\hat{\mu}$ is an arbitrary unknown empirical probability measure on \mathbb{R}^d . Using the triangle inequality and bearing in mind that $\alpha_j \leq 1$ for all j , we can bound this term as follows:

$$\sum_{j=1}^N \alpha_j W_1(\hat{\mu}_j, \hat{\mu}_T) \leq \left(\sum_{j=1}^N \alpha_j W_1(\hat{\mu}_j, \hat{\mu}) \right) + N W_1(\hat{\mu}, \hat{\mu}_T).$$

Now, let us consider the following optimization problem

$$\inf_{\hat{\mu} \in \mathcal{P}(\Omega)} \frac{1}{N} \sum_{j=1}^N \alpha_j W_1(\hat{\mu}_j, \hat{\mu}) + W_1(\hat{\mu}, \hat{\mu}_T). \quad (1)$$

In this formulation, the first term $\frac{1}{N} \sum_{j=1}^N \alpha_j W_1(\hat{\mu}_j, \hat{\mu})$ corresponds exactly to the problem known in the literature as the Wasserstein barycenters problem [1] that can be defined for W_1 as follows.

Definition 2. For N probability measures $\mu_1, \mu_2, \dots, \mu_N \in \mathcal{P}(\Omega)$, an empirical Wasserstein barycenter is a minimizer $\mu_N^* \in \mathcal{P}(\Omega)$ of $J_N(\mu) = \min_{\mu} \frac{1}{N} \sum_{i=1}^N a_i W_1(\mu, \mu_i)$, where for all i , $a_i > 0$ and $\sum_{i=1}^N a_i = 1$.

The second term $W_1(\hat{\mu}, \hat{\mu}_T)$ of Equation 1 finds the probability coupling that transports the barycenter to the target distribution. Altogether, this bound suggests that in order to adapt in the multi-source learning scenario, one can proceed by finding a barycenter of the source probability distributions and transport it to the target probability distribution.

On the other hand, the optimization problem related to the Wasserstein barycenters is closely related to the Multimarginal optimal transportation problem [19] where the goal is to find a probabilistic coupling that aligns N distinct probability measures. Indeed, as shown in [1], for a quadratic Euclidean cost function the solution μ_N^* of the barycenter problem in the Wasserstein space is given by the following equation:

$$\mu_N^* = \sum_{k \in \{k_1, \dots, k_N\}} \gamma_k \delta_{A_k(x)},$$

where $A_k(x) = \sum_{j=1}^N \gamma_j x_{k_j}$ and $\gamma \in \mathbb{R}^{\prod_{j=1}^N n_j}$ is an optimal coupling solving for all $k \in \{1, \dots, N\}$ the multimarginal optimal transportation problem with the following cost:

$$c_k = \sum \frac{a_j}{2} \|x_{k_j} - A_k(x)\|^2.$$

We note that this reformulation is particularly useful when the source distributions are assumed to be Gaussians. In this case, there exists a closed form solution for the multimarginal optimal transportation problem [14] and thus for Wasserstein barycenters problem too. Finally, it is also worth noticing that the optimization problem Equation 1 has already been introduced to solve the multiview learning task [12]. In their formulation, the second term is referred to as an a priori knowledge about the barycenter which, in our case, is explicitly given by the target probability measure simultaneously.

5 Comparison to other existing bounds

As mentioned in the introduction, there are numerous papers that proposed DA generalization bounds. The main difference between them lies in the distance used to measure the divergence between source and target probability distributions. The seminal work of [3] considered a modification of the total variation distance called H-divergence given by the following equation:

$$d_H(p, q) = 2 \sup_{h \in H} |p(h(x) = 1) - q(h(x) = 1)|.$$

On the other hand, [15] and [5] proposed to replace it with the discrepancy distance:

$$\text{disc}(p, q) = \max_{h, h' \in H} |\epsilon_p(h, h') - \epsilon_q(h, h')|.$$

The latter one was shown to be tighter in some plausible scenarios. A more recent work on generalization bounds using integral probability metric

$$D_{\mathcal{F}}(p, q) = \sup_{f \in \mathcal{F}} \left| \int f dp - \int f dq \right|$$

and Rényi divergence

$$D_\alpha(p\|q) = \frac{1}{\alpha - 1} \log \left(\sum_{i=1}^n \frac{p_i^\alpha}{q_i^{\alpha-1}} \right)$$

were presented in [27] and [16], respectively. [27] provides a comparative analysis of discrepancy and integral metric based bounds and shows that the former are less tight. [16] derives the domain adaptation bounds in multisource scenario by assuming that the good hypothesis can be learned as a weighted convex combination of hypothesis from all the sources available. Considering a reasonable amount of previous work on the subject, a natural question about the tightness of the DA bounds based on the Wasserstein metric introduced above arises in spite of the Theorem 3.

The answer to this question is partially given by the Csiszàr-Kullback-Pinsker inequality [20] defined for any two probability measures $p, q \in \mathcal{P}(\Omega)$ as follows:

$$W_1(p, q) \leq \text{diam}(\Omega) \|p - q\|_{\text{TV}} \leq \sqrt{2 \text{diam}(\Omega) \text{KL}(p\|q)},$$

where $\text{diam}(\Omega) = \sup_{x, y \in \Omega} \{d(x, y)\}$ and $\text{KL}(p\|q)$ is the Kullback-Leibler divergence.

A first consequence of this inequality shows that the Wasserstein distance not only appears naturally and offers algorithmic advantages in DA but also gives tighter bounds than total variation distance (L1) used in [2, Theorem 1]. On the other hand, it is also tighter than bounds presented in [16] as the Wasserstein metric can be bounded by the Kullback-Leibler divergence which is a special case of Rényi divergence when $\alpha \rightarrow 1$ as shown in [10]. Regarding the discrepancy distance and omitting the hypothesis class restriction, one has $d_{\min} \text{disc}(p, q) \leq W_1(p, q)$, where $d_{\min} = \min_{x \neq y \in \Omega} \{d(x, y)\}$. This inequality, however, is not very informative as minimum distance between two distinct points can be dramatically small thus making it impossible to compare the considered distances directly.

Regarding computational guarantees, we note that the H-divergence used in [3] is defined as the error of the best hypothesis distinguishing between the source and target domain samples pseudo-labeled with 0's and 1's and thus presents an intractable problem in practice. For the discrepancy distance, authors provided a linear time algorithm for its calculation in 1D case and showed that in other cases it scales as $O(N_S^2 d^{2.5} + N_T d^2)$ when the squared loss is used [15]. In its turn, the Wasserstein distance with entropic regularization can be calculated based on the linear time Sinkhorn-Knopp algorithm regardless the choice of the cost function c that presents a clear advantage over the other distances considered above.

Finally, none of the distances previously introduced in the generalization bounds for DA take into account the geometry of the space meaning that the Wasserstein distance is a powerful and precise tool to measure the divergence between domains.

6 Conclusion

In this paper, we studied the problem of DA in the optimal transportation context. Motivated by the existing algorithmic advances in domain adaptation, we presented the

generalization bounds for both single and multi-source learning scenarios where the distance between source and target probability distributions is measured by the Wasserstein metric. Apart from the distance term that taken alone justifies the use of optimal transport in domain adaptation, the obtained bounds also included the capability term depicting the existence of a good hypothesis for both source and target domains. A direct consequence of its appearance in the bounds is the need to regularize optimal transportation plan in a way that allows to ensure efficient learning in the source domain once the interpolation was done. This regularization, achieved in [6,7] by the means of the class-based regularization, thus can be also viewed as an implication of the obtained results. Furthermore, it explains the superior performance of both class-based and Laplacian regularized optimal transport in domain adaptation compared to its simple entropy regularized form. On the other hand, we also showed that the use of the Wasserstein distance leads to tighter bounds compared to the bounds based on the total variation distance and Rényi divergence and is more computationally attractive than some other existing results. From the analysis of the bounds obtained for the multi-source DA, we derived a new algorithmic idea that suggests the minimization of two terms: first term corresponds to the Wasserstein barycenter problem calculated on the empirical source measures while the second one solves the optimal transport problem between this barycenter and the empirical target measure.

Future perspectives of this work are many and concern both the derivation of new algorithms for domain adaptation and the demonstration of new theoretical results. First of all, we would like to study the extent to which the cost function used in the derivation of the bounds can be used on actual real-world DA problems. This distance, defined as a norm of difference between two feature maps, can offer a flexibility in the calculation of the optimal transport metric due to its kernel representation. Secondly, we aim to produce new concentration inequalities for the λ term that will allow to bound the true best joint hypothesis by its empirical counter-part. These concentration inequalities will allow to access the adaptability of two domains from the given labelled samples while the speed of convergence may show how many data instances from the source domains is needed to obtain a reliable estimate of λ . Finally, the introduction of the Wasserstein distance to the bounds means that new DA algorithms can be designed based on the other optimal coupling techniques. These include, for instance, Knothe-Rosenblatt coupling and Moser coupling.

Acknowledgments. This work was supported in part by the French ANR project LIVES ANR-15-CE23-0026-03.

References

1. M. Agueh and G. Carlier. Barycenters in the Wasserstein space. *SIAM J. Math. Analysis*, 43(2):904–924, 2011.
2. Sh. Ben-David, J. Blitzer, K. Crammer, A. Kulesza, F. Pereira, and J. Vaughan. A theory of learning from different domains. *Machine Learning*, 79:151–175, 2010.
3. Sh. Ben-David, J. Blitzer, K. Crammer, and O Pereira. Analysis of representations for domain adaptation. In *NIPS*, 2007.

4. Fr. Bolley, Ar. Guillin, and C. Villani. Quantitative concentration inequalities for empirical measures on non-compact spaces. *Prob. Theory and Related Fields*, 137(3–4):541–593, 2007.
5. C. Cortes and M. Mohri. Domain adaptation and sample bias correction theory and algorithm for regression. *Theoretical Computer Science*, 519:103–126, 2014.
6. N. Courty, R. Flamary, and D. Tuia. Domain adaptation with regularized optimal transport. In *Proceedings of ECML/PKDD 2014*, pages 1–16, 2014.
7. N. Courty, R. Flamary, D. Tuia, and A. Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2016.
8. Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *NIPS*, pages 2292–2300, 2013.
9. Marco Cuturi and Arnaud Doucet. Fast computation of wasserstein barycenters. In *ICML*, pages 685–693, 2014.
10. Ying Ding. Wasserstein-divergence transportation inequalities and polynomial concentration inequalities. *Statistics and Probability Letters*, 94(C):77–85, 2014.
11. Nicolas Fournier and Arnaud Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3–4):707, 2015.
12. M. Bergounioux I. Abraham, R. Abraham and G. Carlier. Tomographic reconstruction from a few views: a multi-marginal optimal transport approach. *Applied Mathematics and Optimization*, page 1–19, 2016.
13. Leonid Kantorovich. On the translocation of masses. In *C.R. (Doklady) Acad. Sci. URSS(N.S.)*, volume 37(10), page 199–201, 1942.
14. M. Knott and C. S. Smith. On a generalization of cyclic-monotonicity and distances among random vectors. *Linear Algebra and its Appl.*, 199:363–371, 1994.
15. Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Domain adaptation: Learning bounds and algorithms. In *COLT*, 2009.
16. Yishay Mansour, Mehryar Mohri, and Afshin Rostamizadeh. Multiple source adaptation and the rényi divergence. In *UAI*, pages 367–374, 2009.
17. Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l'Académie Royale des Sciences*, pages 666–704, 1781.
18. Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2010.
19. Brendan Pass. Uniqueness and monge solutions in the multimarginal optimal transportation problem. *SIAM Journal on Mathematical Analysis*, 43(6):2758–2775, 2011.
20. M. S. Pinsker and Amiel Feinstein. Information and information stability of random variables and processes, 1964.
21. Julien Rabin, Gabriel Peyré, Julie Delon, and Marc Bernot. Wasserstein barycenter and its application to texture mixing. In *SSVM*, volume 6667, pages 435–446, 2011.
22. Yossi Rubner, Carlo Tomasi, and Leonidas J. Guibas. The earth mover’s distance as a metric for image retrieval. *Int. J. Comput. Vision*, 40(2):99–121, 2000.
23. Saburo Saitoh. *Integral Transforms, Reproducing Kernels and their Applications*. Pitman Research Notes in Mathematics Series, 1997.
24. D. Sejdinovic, B. K. Sriperumbudur, A. Gretton, and K. Fukumizu. Equivalence of distance-based and RKHS-based statistics in hypothesis testing. *Ann. Stat.*, 41(5):2263–2291, 2013.
25. Richard Sinkhorn and Paul. Knopp. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific J. Math*, 21:343–348, 1967.
26. Cédric Villani. *Optimal transport : old and new*. Grundlehren der mathematischen Wissenschaften. Springer, Berlin, 2009.
27. Chao Zhang, Lei Zhang, and Jieping Ye. Generalization bounds for domain adaptation. In *NIPS*, 2012.
28. Michaël Perrot, Nicolas Courty, Rémi Flamary, and Amaury Habrard. Mapping Estimation for Discrete Optimal Transport. In *NIPS*, 4197–4205, 2016.