

Optimal Transport for Domain Adaptation through Gaussian Mixture Models

Anonymous authors

Paper under double-blind review

Abstract

Machine learning systems operate under the assumption that training and test data are sampled from a fixed probability distribution. However, this assumption is rarely verified in practice, as the conditions upon which data was acquired are likely to change. In this context, unsupervised domain adaptation requires minimal access to data from the new conditions for learning models robust to changes in the data distribution. Optimal transport is a theoretically grounded tool for analyzing changes in distribution, especially as it allows the mapping between domains. However, these methods are usually computationally expensive, as their complexity scales cubically with the number of samples. In this work, we explore optimal transport between Gaussian Mixture Models (GMMs), which is conveniently written in terms of the components of source and target GMMs. We experiment with 8 benchmarks, with a total of 84 adaptation tasks, showing that our methods are more efficient than previous methods, and they scale well with both number of samples n and dimensions d .

1 Introduction

Supervised machine learning models are trained with significant amounts of labeled data, constituting a training set. The theory of generalization (Redko et al., 2020), provides a theoretical background that guarantees accurate predictions on unseen samples *from the same distribution*. Nonetheless, these models are often forced to predict on related, but different data samples (Quinero-Candela et al., 2008). This distinction is modeled by a *shift* in the probability distributions generating the data (Sugiyama et al., 2007). This motivates the field of transfer learning (Pan & Yang, 2009), and more specifically the problem of Unsupervised Domain Adaptation (UDA), in which models are adapted from a labeled *source domain*, towards an unlabeled *target domain*, following different distributions.

In this context, Optimal Transport (OT) (Villani et al., 2009; Peyré et al., 2019) is a powerful, theoretically grounded framework for comparing and manipulating probability distributions (Montesuma et al., 2023b). This framework works by computing a transportation strategy, that moves one probability distribution into the other at least effort. Based on this core idea, different methods have been proposed for UDA (Courtney et al., 2017; El Hamri et al., 2022; Chuang et al., 2023; Struckmeier et al., 2023).

However, this methodology faces a few challenges. For instance, OT maps computed between discrete distributions are only defined for samples in the training set. Extrapolating these maps to new samples is the subject of intense research Perrot et al. (2016); Seguy et al. (2017). A possible workaround consists of using Gaussian approximations (Flamary et al., 2019; Struckmeier et al., 2023). While this approach effectively defines a mapping over the whole ambient space, its hypothesis do not reflect the possible sub-populations within the data, which are common in classification problems.

A natural solution to tackle multi-modality in data distributions is using Gaussian Mixture Models (GMMs). A further advantage of this approach is considering the recently proposed GMM OT (GMMOT) by Delon & Desolneux (2020), which establishes an efficient, discrete problem between the components in the GMM. Furthermore, recent works have established the effectiveness of this idea for *multi-source* domain adaptation (Montesuma et al., 2024a), notably through the use of mixture-Wasserstein barycenters.

Nonetheless, some questions on the use of GMMs for UDA remain open. For instance, Delon & Desolneux (2020) propose different mapping strategies between GMMs, but these either fail to map one GMM into the other, or are subject to randomness when sampling transportation maps between GMM components. Furthermore, if we label the components of GMMs, it is natural to *propagate* these labels, in the sense of Redko et al. (2019) towards the target domain. This paper tackles these two questions.

Summary of contributions. In this paper, we propose 2 new strategies for UDA based on GMMs. First, we use basic rules of probability theory for propagating the labels of source domain GMM towards the target domain GMM. We do so, by interpreting the GMMOT plan as the joint probability of source and target GMM components. Second, we map samples from the source domain into the target domain based on the GMMOT plan. For a point in the source domain, our strategy consists of first estimating the component, in the source GMM, most likely to have generated the sample. We then transport this point to components in the target domain, while assigning importance weights based on the GMMOT plan.

Paper organization. The rest of this paper is organized as follows. Section 2 presents a few related works on OT for UDA. Section 3 covers the preliminaries on OT and GMMOT theory. Section 4 covers our methodological contributions. Section 5 details our experiments and discussion on UDA. Finally, section 6 concludes this paper.

2 Related Works

Optimal transport for domain adaptation. Optimal transport has been extensively employed for the design of algorithms (Courty et al., 2017), as well as analyzing the domain adaptation problem (Redko et al., 2017). The key idea of this method is to use the Kantorovich formulation to acquire a matching, known as transport plan, between source and target domain distributions. This matching defines a map between points in the source domain, towards the target domain, called barycentric mapping. Based on this idea, different methods have proposed improvements. For instance, Perrot et al. (2016) proposed learning linear and kernelized extensions of the barycentric map through convex optimization. El Hamri et al. (2022) uses clustering for learning matching with additional structural dependencies. Flamary et al. (2019) uses optimal transport between Gaussian distributions for estimating an affine mapping between source and target domains. More recently, Chuang et al. (2023) proposed a method that leverages kernel density estimation for defining a new optimal transport problem based on information maximization.

Gaussian-mixture based optimal transport. An optimal transport problem involving Gaussian mixtures was initially proposed by Chen et al. (2018), in which a linear program between the components of the two mixtures is solved. This setting was further studied by Delon & Desolneux (2020), who proved an interesting connection to a continuous optimal transport, when the transport plan is constrained to the set of Gaussian mixtures. Based on the framework of Delon & Desolneux (2020), Montesuma et al. (2024a) proposed the extension of multi-source domain adaptation algorithms of Montesuma & Mboula (2021a;b) and Montesuma et al. (2023a). However, these authors focused on performing adaptation through Wasserstein barycenters. Although based on the same framework, our work focuses on *mapping samples and propagating labels* of Gaussian mixtures, especially for single-source domain adaptation.

3 Theoretical Foundations

3.1 Optimal Transport

Founded by Monge (1781), optimal transport is a field of mathematics concerned with transporting mass at least effort. Let \mathcal{X} be a set and $\mathcal{P}(\mathcal{X})$ the set of probability distributions on \mathcal{X} . For $P, Q \in \mathcal{P}(\mathcal{X})$, the Monge formulation of the optimal transport problem is,

$$T^* = \operatorname{arginf}_{T: T_{\#}P=Q} \int_{\mathcal{X}} c(x, T(x)) dP(x), \quad (1)$$

where $T_{\#}P$ denotes the pushforward distribution (Santambrogio, 2015, Problem 1.1) of P by the map T , and $c: \mathcal{X} \times \mathcal{X} \rightarrow \mathbb{R}$, called the ground-cost, denotes the cost of sending x to position $T(x)$.

Although equation 1 is a formal description of the optimal transport problem, the constraint $T_{\#}P = Q$ poses technical difficulties. An alternative description was proposed by Kantorovich (1942), in terms of an optimal transport *plan* γ ,

$$\gamma^* = \operatorname{argmin}_{\gamma \in \Gamma(P, Q)} \int_{\mathcal{X} \times \mathcal{X}} c(x_1, x_2) d\gamma(x_1, x_2), \quad (2)$$

where $\Gamma(P, Q)$ is the set of joint distributions with marginals P and Q . This formulation is simpler to analyze because the constraint $\gamma \in \Gamma(P, Q)$ is linear with respect to the optimization variable γ .

When (\mathcal{X}, d) is a metric space, it is possible to define a distance on $\mathcal{P}(\mathcal{X})$ in terms of d . Let $\alpha \in [1, +\infty)$, and $c(\cdot, \cdot) = d(\cdot, \cdot)^\alpha$. One then has the so-called α -Wasserstein distance:

$$\mathcal{W}_\alpha(P, Q)^\alpha = \operatorname{arginf}_{\gamma \in \Gamma(P, Q)} \int_{\mathcal{X} \times \mathcal{X}} c(x_1, x_2) d\gamma(x_1, x_2). \quad (3)$$

In the following, we do optimal transport on Euclidean spaces, i.e., $\mathcal{X} = \mathbb{R}^d$. In this case, it is natural to use $d(\mathbf{x}_1, \mathbf{x}_2) = \|\mathbf{x}_1 - \mathbf{x}_2\|_2$. Furthermore, we set $\alpha = 2$.

Equations 2 and 3 are linear programs, where the optimization variable is the joint distribution γ . Thus, in its continuous formulation, the optimal transport problem is an infinite-size linear program. In the following, we discuss 3 particular cases where optimal transport either has a closed form, or is approximated by a finite problem, thus tractable by a computer.

Empirical Case. If we have samples $\{\mathbf{x}_i^{(P)}\}_{i=1}^n$ and $\{\mathbf{x}_j^{(Q)}\}_{j=1}^m$ with probabilities p_i and q_j respectively, we can make empirical approximations for P and Q ,

$$\hat{P}(\mathbf{x}) = \sum_{i=1}^n p_i \delta(\mathbf{x} - \mathbf{x}_i^{(P)}). \quad (4)$$

The approximation in equation 4 is at the core of discrete optimal transport Peyré et al. (2019). If we plug equation 4 into equation 2, the optimal transport problem becomes computable, i.e., it turns into a linear programming problem with $n \times m$ variables,

$$\gamma^* = \arg \min_{\gamma \in \Gamma(\hat{P}, \hat{Q})} \sum_{i=1}^n \sum_{j=1}^m \gamma_{ij} C_{ij}, \quad (5)$$

where $C_{ij} = c(\mathbf{x}_i^{(P)}, \mathbf{x}_j^{(Q)})$. As a linear program, one should keep in mind that solving equation 5 has a complexity of $\mathcal{O}(n^3 \log n)$. This complexity can be reduced by regularizing the problem in terms of the OT plan entropy,

$$\gamma^* = \arg \min_{\gamma \in \Gamma(\hat{P}, \hat{Q})} \sum_{i=1}^n \sum_{j=1}^m \gamma_{ij} C_{ij} + \epsilon \sum_{i=1}^n \sum_{j=1}^m \gamma_{ij} (\log \gamma_{ij} - 1), \quad (6)$$

which can be solved through the celebrated Sinkhorn algorithm (Cuturi, 2013). A solution can be found with complexity $\mathcal{O}(Ln^2)$, where L is the number of iterations of the algorithm. While the Kantorovich formulation is useful, the notion of map between P and Q is lost. We can recover a correspondence through the barycentric map,

$$T_\gamma(\mathbf{x}_i^{(P)}) = \min_{\mathbf{x} \in \mathcal{X}} \sum_{j=1}^m \gamma_{ij} c(\mathbf{x}, \mathbf{x}_j^{(Q)}). \quad (7)$$

Gaussian Case. When $P = \mathcal{N}(\mu^{(P)}, \Sigma^{(P)})$ (resp., Q), equation 1 has a closed and affine form Takatsu (2011), $T^*(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, where,

$$\mathbf{A} = (\Sigma^{(P)})^{-\frac{1}{2}} ((\Sigma^{(P)})^{\frac{1}{2}} \Sigma^{(Q)} (\Sigma^{(P)})^{\frac{1}{2}})^{\frac{1}{2}} (\Sigma^{(P)})^{-\frac{1}{2}}, \text{ and } \mathbf{b} = \mu^{(Q)} - \mathbf{A}\mu^{(P)}, \quad (8)$$

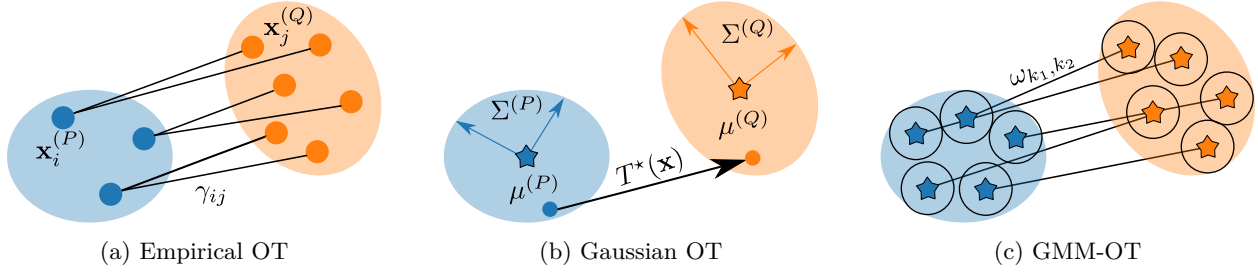


Figure 1: **Comparison of different ways of solving OT.** In empirical OT, P and Q are approximated non-parametrically through their samples. In Gaussian OT, P and Q are Gaussian distributions, and the mapping between these distributions is affine. In GMM-OT, P and Q are assumed to be GMMs, and an OT plan between components, ω , defines an OT plan between samples, γ .

and the Wasserstein distance takes the form,

$$\mathcal{W}_2(P, Q)^2 = \|\mu^{(P)} - \mu^{(Q)}\|_2^2 + \text{Tr}(\Sigma^{(P)} + \Sigma^{(Q)} + ((\Sigma^{(P)})^{1/2} \Sigma^{(Q)} (\Sigma^{(P)})^{1/2})^{1/2}). \quad (9)$$

In contrast with empirical optimal transport, under the Gaussian hypothesis, computing the Wasserstein distance and a mapping between P and Q has sample-free complexity. Indeed, the complexity of equations 8 and 9 is dominated by computing the square-root of the covariance matrix with complexity $\mathcal{O}(d^3)$.

Furthermore, in high dimensions and when only a few data points are available, estimating full covariance matrices is challenging. In these cases, it is useful to assume axis-aligned Gaussians, i.e., Σ is a diagonal matrix with diagonal elements $\Sigma_{ii} = \sigma_i^2$. In this case, equations 8 and 9 can be further simplified,

$$\mathbf{A} = \text{diag}(\sigma^{(Q)}/\sigma^{(P)}), \mathbf{b} = \mu^{(Q)} - \mathbf{A}\mu^{(P)}, \text{ and } \mathcal{W}_2(P, Q)^2 = \|\mu^{(P)} - \mu^{(Q)}\|_2^2 + \|\sigma^{(P)} - \sigma^{(Q)}\|_2^2. \quad (10)$$

Gaussian Mixture Case. A Gaussian mixture corresponds to $P = \sum_{k=1}^K \pi_k^{(P)} P_k$, where $P_k = \mathcal{N}(\mu_k^{(P)}, \Sigma_k^{(P)})$. As in Delon & Desolneux (2020), we denote by $\text{GMM}_d(K)$ the set of distributions $P \in \mathcal{P}(\mathbb{R}^d)$ written as a mixture of at most K components. In this framework Delon & Desolneux (2020) explores the optimal transport problem 2 under the constraint that γ is a GMM as well, i.e., $\gamma \in \Gamma(P, Q) \cap \text{GMM}_{2d}(\infty)$. This formulation is interesting because it is equivalent to a discrete and hierarchical problem (Delon & Desolneux, 2020, Proposition 4) in terms of the GMMs' components

$$\omega^* = \min_{\omega \in \Gamma(\pi^{(P)}, \pi^{(Q)})} \sum_{k_1=1}^{K_P} \sum_{k_2=1}^{K_Q} \omega_{k_1, k_2} \mathcal{W}_2(P_{k_1}, Q_{k_2})^2. \quad (11)$$

We call this problem GMMOT. This latter equation is an hierarchical optimal transport problem, i.e., a problem that involves itself an inner optimal transport. As in the previous cases, we have a notion of distance related to equation 11, $\mathcal{MW}_2(P, Q)^2 = \sum_{k_1=1}^{K_1} \sum_{k_2=1}^{K_2} \omega_{k_1, k_2} \mathcal{W}_2(P_{k_1}, Q_{k_2})^2$. We show an overview of different strategies for solving OT, under different assumptions, in Figure 1.

Remark 3.1. (*Computational Complexity*) The overall complexity of the GMMOT in equation 11 is $\mathcal{O}(K^3 \log K)$, or $\mathcal{O}(LK^2)$, for the linear programming, and Sinkhorn algorithms respectively. However, one should keep in mind that the ground-cost matrix must be computed beforehand, as given by equation 9. This involves a complexity that scales with the dimension of the ambient space, due the matrix inversions and square-roots. The complexity of these operations is $\mathcal{O}(K^2 d^3)$ in general. However, assuming diagonal covariance matrices, the computational complexity is drastically reduced to $\mathcal{O}(K^2 d)$, i.e., the complexity of computing K^2 Euclidean distances between d -dimensional vectors.

The closely related method of El Hamri et al. (2022) also solves a reduced OT problem between clusters of data. However, the inner Wasserstein distances are still empirical. This means that, for computing the ground-cost matrix, one has $\mathcal{O}(K^2 m^3 \log m)$ complexity, where m is the maximum number of samples on each class over the domains. Overall, our work fits into the ongoing discussion on whether use non-parametric

models, and have a dimension-free complexity, or using a parametric approach with Gaussian or GMMs, and get a sample-free complexity, modulo the complexity of fitting the model.

Remark 3.2. (Parameter Estimation) Besides the computational advantage, using diagonal covariance matrices yields a simpler estimation problem for GMMs. This is pivotal in Domain Adaptation (DA), as the target domain likely does not have enough samples for the accurate estimation of complete covariance matrices. Furthermore, data is oftentimes high-dimensional, as feature spaces commonly involve thousands of features. This is a sharp contrast with previous works in GMMOT, such as Delon & Desolneux (2020) and Chen et al. (2018), which involved a few dimensions.

3.2 Learning Theory and Domain Adaptation

In this paper, we deal with DA for classification. This latter problem can be formalized mathematically, as the learning of a function $h : \mathcal{X} \rightarrow \mathcal{Y}$, from a feature space \mathcal{X} (e.g., \mathbb{R}^d) to a label space, $\mathcal{Y} = \{1, \dots, n_c\}$ through samples of a probability distribution. As reviewed by Redko et al. (2020), from the point of view of probability, there multiple ways of formalizing this problem. Here, we use the Empirical Risk Minimization (ERM) framework of Vapnik (2013).

Let P be a probability distribution. Given a loss function $\mathcal{L} : \mathcal{Y} \times \mathcal{Y} \rightarrow \mathbb{R}_+$ and a family of classifiers $\mathcal{H} \subset \mathcal{Y}^{\mathcal{X}}$, one may define a notion of disagreement between $h, h' \in \mathcal{H}$,

$$\mathcal{R}_P(h, h') = \mathbb{E}_{\mathbf{x} \sim P} [\mathcal{L}(h(\mathbf{x}), h'(\mathbf{x}))]. \quad (12)$$

Equation 12 defines the so-called risk of h with respect h' . Given a ground-truth labeling function $h_0 : \mathcal{X} \rightarrow \mathcal{Y}$, classification can be phrased in terms of minimizing the risk of h with respect the ground-truth h_0 , i.e.,

$$h^* = \arg \min_{h \in \mathcal{H}} \mathcal{R}_P(h, h_0).$$

Henceforth, we adopt $\mathcal{R}_P(h) = \mathcal{R}_P(h, h_0)$ in short. This formalization equates the problem of learning a classifier with an optimization problem. Nonetheless, in practice one does not have access to *a priori* knowledge from P nor h_0 . In a more realistic scenario, one has samples $\{\mathbf{x}_i^{(P)}, y_i^{(P)}\}_{i=1}^n$, where $\mathbf{x}_i^{(P)} \stackrel{iid}{\sim} P$, and $y_i^{(P)} = h_0(\mathbf{x}_i^{(P)})$. Based on these samples, one may estimate the risk empirically, by resorting to the approximation in equation 4,

$$\hat{\mathcal{R}}_P(h) = \frac{1}{n} \sum_{i=1}^n \mathcal{L}(h(\mathbf{x}_i^{(P)}), y_i^{(P)}), \text{ and } \hat{h} = \arg \min_{h \in \mathcal{H}} \hat{\mathcal{R}}_P(h). \quad (13)$$

In machine learning literature, the true risk $\mathcal{R}_P(h)$ is called *generalization error*, i.e., the error that a classifier h makes on samples from the distribution P . In contrast, algorithms usually minimize the *training error*, $\hat{\mathcal{R}}_P(h)$, i.e., the error of h on the particular examples available during training.

A key limitation of the presented theory is its assumption that data originates from a single probability distribution P . As discussed by Quinonero-Candela et al. (2008), this is seldom happens in practice. For instance, in fault diagnosis, process conditions influence the statistical properties of measured signals (Montesuma et al., 2022). As a result, generalization must be carried to a new, related probability distribution. This problem is known in the literature as DA, a sub-field within transfer learning.

As discussed by Pan & Yang (2009), in transfer learning, a domain is a pair $\mathcal{D} = (\mathcal{X}, P)$ of a feature space and a probability distribution over \mathcal{X} . Likewise, a task is a pair $\mathcal{T} = (\mathcal{Y}, h_0)$ of a label space and a ground-truth labeling function. Transfer learning is characterized by different source and target domains and tasks, i.e., $(\mathcal{D}_S, \mathcal{T}_S, \mathcal{D}_T, \mathcal{T}_T)$, where at least one element from the source is different from the target. Within the taxonomy of transfer learning (Pan & Yang, 2009, Figure 2), DA assumes different domains, $\mathcal{D}_S \neq \mathcal{D}_T$ and a single task $\mathcal{T} = \mathcal{T}_S = \mathcal{T}_T$.

In this paper, we deal primarily with *distributional shift*. In this case, we assume $\mathcal{X}_S = \mathcal{X}_T = \mathbb{R}^d$, so that source and target domains are characterized by different probability distributions $P \neq Q$. Furthermore, we place ourselves in the *unsupervised DA* setting, that is, we assume labeled source domain data,

$\{\mathbf{x}_i^{(P)}, y_i^{(P)}\}_{i=1}^n$, and unlabeled target domain data $\{\mathbf{x}_j^{(Q)}\}_{j=1}^m$. Our goal is to use these samples to learn a classifier h that works well on the target domain, i.e., that achieves small target risk \mathcal{R}_Q .

Optimal Transport for Domain Adaptation was proposed by Courty et al. (2017), and primarily tries to match samples from P to those of Q based on the empirical OT problem in equation 5. After acquiring γ^* , the authors propose mapping samples from P towards those of Q via the barycentric map. This strategy effectively constitutes a new dataset $\{T_{\gamma^*}(\mathbf{x}_i^{(P)}), y_i^{(P)}\}_{i=1}^n$, where T_{γ^*} is defined by equation 7. Note, here, that under T_{γ^*} , the source domain points *carry their labels*, to the target domain. This operation is valid, as long as the conditionals $P(Y|X) = Q(Y|T(X))$, which is restrictive, but reasonable under the covariate shift hypothesis. A few problems plague this strategy. First, T_{γ^*} is only defined on the support of P . The mapping of new points has been extensively studied in the literature (Perrot et al., 2016; Seguy et al., 2017; Chuang et al., 2023). Second, it is desirable to have mappings with additional structure with respect the classes in DA. This problem is partially solved by considering special regularization schemes, as covered in Courty et al. (2017). Third, this method is not scalable with respect the number of samples n , due its prohibitive complexity $\mathcal{O}(n^3 \log n)$. In this paper, we offer a solution for the aforementioned problems through the GMM-OT framework of Delon & Desolneux (2020).

4 Domain Adaptation via Optimal Transport between Gaussian Mixtures

In this section, we develop new tools for DA through OT between GMMs. As we discussed in our preliminaries section, we are particularly interested in DA for classification. In this sense, we need to define labels for the components of GMMs in the transportation problem. We propose doing so through an heuristic. Instead of modeling $P(X)$ through a GMM, we propose modeling each $P_y = P(X|Y = y)$, for $y = 1, \dots, n_c$. As a result, we fit a GMM to the data $\mathbf{X}^{(P_y)} = \{\mathbf{x}_i^{(P)}\}_{i: y_i^{(P)}=y}$, using $cpc = K_P/n_c$ components. Here, we conveniently choose K_P as a multiple of n_c , for ensuring that cpc is an integer. After fitting a GMM to each P_y , we create P by concatenating the components of $\{P_y\}_{y=1}^{n_c}$,

$$\mu^{(P)} = \{\{\mu_k^{(P_y)}\}_{k=1}^{cpc}\}_{y=1}^{n_c}, \Sigma^{(P)} = \{\{\Sigma_k^{(P_y)}\}_{k=1}^{cpc}\}_{y=1}^{n_c}, \text{ and } \pi^{(P)} = \left\{ \left\{ \frac{\pi_k^{(P_y)}}{\sum_{k=1}^{cpc} \sum_{y=1}^{n_c} \pi_k^{(P_y)}} \right\}_{k=1}^{cpc} \right\}_{y=1}^{n_c}, \quad (14)$$

where $\pi_k^{(P_y)}$ is the weight of k -th component of the y -th GMM. Furthermore, we create a one-hot encoded vector of component labels, denoted $\nu^{(P)} \in (\Delta_{n_c})^K$, where $\nu_{k,y}^{(P)} = 1$ if the k -th component comes from the y -th class, and 0 otherwise. We interpret the vector $\nu_k^{(P)}$ as the conditional probability $P(Y|K = k)$.

Given this strategy, we recall that, in DA, we have a labeled source dataset $\{\mathbf{x}_i^{(P)}, y_i^{(P)}\}_{i=1}^n$, and an unlabeled target dataset, $\{\mathbf{x}_j^{(Q)}\}_{j=1}^m$. The starting point of our strategy is defining GMMs for the source and the target. Since labeled data is available on the source domain, we define a labeled source GMM, denoted by $P = \{\pi_k^{(P)}, \mu_k^{(P)}, \Sigma_k^{(P)}, \nu_k^{(P)}\}_{k=1}^{K_P}$. For the target domain, we do not have labels. As a result, we fit a GMM to the features using a standard Expectation Maximization (EM) (Dempster et al., 1977). We thus have $Q = \{\pi_k^{(Q)}, \mu_k^{(Q)}, \Sigma_k^{(Q)}\}_{k=1}^{K_Q}$.

In the following, we discuss two strategies for domain adaptation. The first, based on label propagation, leverages the optimal transport plan *between components* to define pseudo-labels for the components of the target domain GMM. The second, based on mapping estimation, leverages the hierarchical nature of the GMMOT problem for defining a map between source and target domain.

4.1 Label Propagation and Maximum a Posterior Estimation

Recall that, in equation 11, the result of the GMMOT problem is a transportation plan ω , between components of the GMMs P and Q . As a result, ω has as marginals the probability vectors $\pi^{(P)}$ and $\pi^{(Q)}$. Furthermore, in the original probabilistic view of GMMs, $\pi_{k_1}^{(P)} = P(K = k_1)$, whereas $\pi_{k_2}^{(Q)} = Q(K = k_2)$. Given this interpretation, it is natural to interpret the transportation plan ω as a *joint distribution* $P(K_1, K_2)$.




Figure 2: **Illustration of the label propagation strategy.** (a) Shows the source and target GMMs. GMMs are represented through their means (stars) and covariance matrices (ellipses). The labels of components are represented through colors. Since the target domain is unlabeled, their means are gray colored. (b) Shows the OT plan between components, ω . ω allows us to propagate the labels of source domain GMM towards the target. (c) Shows the obtained target GMM. Finally, (d) shows the MAP classifier.

Note that we can estimate the probability $Q(Y|K_2)$ using the law of total probabilities,

$$Q(Y|K_2 = k_2) = \sum_{k_1=1}^{K_P} P(Y|K_1 = k_1, K_2 = k_2)P(K_1 = k_1|K_2 = k_2).$$

Here, assuming that Y and K_2 are conditionally independent given K_1 , we have,

$$Q(Y|K_2 = k_2) = \sum_{k_1=1}^{K_P} P(Y|K_1 = k_1)P(K_1 = k_1|K_2 = k_2),$$

This assumption plays the same role as covariate shift hypothesis (Sugiyama et al., 2007) in conventional DA works. On an intuitive level, our assumption explicit the fact that K_2 is redundant with respect to K_1 . The conditional $P(K_1|K_2) = P(K_1, K_2)/P(K_2)$ can be computed through the optimal transport plan, i.e.,

$$\hat{\nu}_{k_2}^{(Q)} = \frac{1}{\pi_{k_2}^{(Q)}} \sum_{k_1=1}^{K_P} \omega_{k_1, k_2} \nu_{k_1}^{(P)}. \quad (15)$$

Equation 15 is known in the OT literature as *label propagation*, and, as we discussed in the related works section, has been used extensively in the context of empirical OT. Given the estimated labels $\hat{\nu}_{k_2}^{(Q)}$, we effectively defined a labeled GMM for the target domain. Based on this GMM, we can perform Maximum A Posteriori (MAP) estimation to define a classifier in the target domain,

$$\begin{aligned} \hat{h}_{MAP}(\mathbf{x}) &= \arg \max_{y=1, \dots, n_c} Q(Y = y|X = \mathbf{x}) = \arg \max_{y=1, \dots, n_c} \sum_{k=1}^{K_Q} Q(Y = y|X = \mathbf{x}, K = k)Q(K = k|X = \mathbf{x}), \\ &= \arg \max_{y=1, \dots, n_c} \sum_{k=1}^{K_Q} \left(\frac{\pi_k^{(Q)} Q_k(\mathbf{x})}{\sum_{k'=1}^{K_Q} \pi_{k'}^{(Q)} Q_{k'}(\mathbf{x})} \right) \nu_k^{(Q)}, \end{aligned} \quad (16)$$

where, from the second to the third equality we assumed that Y and K are conditionally independent given X . This hypothesis is intuitive, as classes and components are representing the same structures within the data points. We show an illustration of these ideas in Figure 2.

4.2 Mapping Estimation

In this section, we propose a new mapping estimation technique between two GMMs, P and Q . As discussed in (Delon & Desolneux, 2020, Section 6.3), this problem is not straightforward. Indeed, since γ is a GMM, it

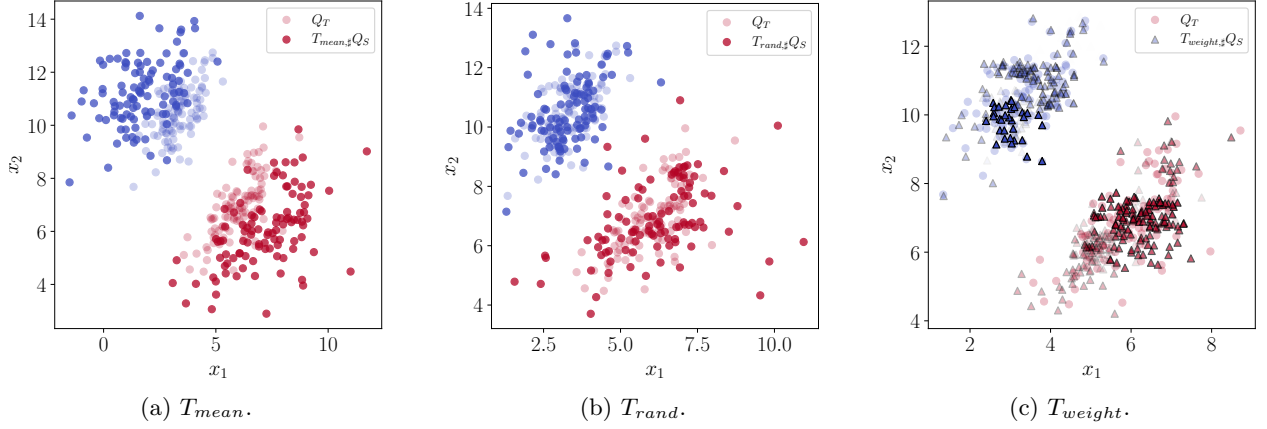


Figure 3: **Mapping estimation using GMMOT.** In (a) and (b), we show the T_{mean} and T_{rand} strategies of Delon & Desolneux (2020), respectively. In (c), we show our strategy T_{weight} . Our mapping reduces randomness by first estimating the component k_1 most likely to have generated $\mathbf{x}^{(P)}$. Then, we weight the importance of transported samples by ω_{k_1, k_2} .

cannot be, in general, written as $(Id, T)_{\#}P$. In Delon & Desolneux (2020), the authors consider two strategies: a mean map and a random map. In the first case, samples are mapped $T_{mean}(\mathbf{x}^{(P)}) = \mathbb{E}_{\mathbf{x}^{(Q)} \sim \gamma(\cdot | \mathbf{x}^{(P)})}[\mathbf{x}^{(Q)}]$, but it may end up not actually matching P with Q (see, e.g., (Delon & Desolneux, 2020, Section 6.3)). In the second case, Delon & Desolneux (2020) defines

$$T_{rand}(\mathbf{x}^{(P)}) = T_{k_1, k_2}(\mathbf{x}^{(P)}) \text{ with probability } p_{k_1, k_2}(\mathbf{x}^{(P)}) = \omega_{k_1, k_2}^* \frac{\mathcal{N}(\mathbf{x}^{(P)} | \mu_{k_1}^{(P)}, \Sigma_{k_1}^{(P)})}{\sum_k p_k \mathcal{N}(\mathbf{x}^{(P)} | \mu_k^{(P)}, \Sigma_k^{(P)})},$$

which has the advantage of matching P with Q . However, as noted by Delon & Desolneux (2020), T_{rand} usually leads to irregular mappings due the sampling procedure of indices (k_1, k_2) as shown in (Delon & Desolneux, 2020, Figure 7). To tackle this issue, we propose a new notion of mapping called T_{weight} .

We put forth a third strategy for mapping P into Q . Our intuition is twofold. First, we can increase the regularity of T_{rand} , by estimating the component k_1 that most likely originated $\mathbf{x}^{(P)}$, that is,

$$k_1 := \arg \max_{k=1, \dots, K_P} P(K = k | X = \mathbf{x}^{(P)}) = \frac{\pi_k^{(P)} P_k(\mathbf{x}^{(P)})}{\sum_{k'=1}^{K_P} \pi_{k'}^{(P)} P_{k'}(\mathbf{x})}.$$

Second, we map $\mathbf{x}^{(P)}$ into the components of Q . Note that, even if $K_P = K_Q$, the optimal transport plan ω may split the mass of P_{k_1} into several $Q_{k'_2}$. This is due the different marginals \mathbf{p} and \mathbf{q} . As a result, we produce $\{T_{k_1, k_2}(\mathbf{x}^{(P)})\}_{k_2: \omega_{k_1, k_2} \geq \tau}$, i.e., we map $\mathbf{x}^{(P)}$ to all components Q_{k_2} such that $\omega_{k_1, k_2} \geq \tau \geq 0$. In principle, one may choose $\tau = 0$ and filter only the components that are not matched with P_{k_1} . Third, we further weight the importance of generated samples, by using ω_{k_1, k_2} as sample weights. At the end, we generate a weighted dataset $\{(\omega_{k_1, k_2}, T_{k_1, k_2}(\mathbf{x}_i^{(P)}), y_i^{(P)})\}_{i=1}^m$, where m is the total amount of samples generated. We call our overall mapping T_{weight} .

The mapping we just defined has a few interesting properties. First, it is a piece-wise affine map, as each T_{k_1, k_2} is affine. This property contrast with the Gaussian hypothesis, which defines an affine map between P and Q . Second, with respect the transportation of samples, our mapping strategy is naturally *group-sparse*, in the sense of Courty et al. (2017). This claim comes from the fact that samples in P are transported based on the Gaussian component they belong to. Third, our mapping is defined on the whole support of the GMM P . In contrast, empirical OT is only defined on the samples $\{\mathbf{x}_i^{(P)}\}_{i=1}^n$ of P .

5 Experiments

Overview of Benchmarks. Our experimental validation is divided into two parts. First, we evaluate our methods in commonly used visual domain adaptation benchmarks. These are composed by office-like benchmarks, such as Caltech-Office (Gong et al., 2012), ImageCLEF (Caputo et al., 2014), Office31 (Saenko et al., 2010) and Office-Home (Venkateswara et al., 2017). We also explore a large scale adaptation experiment with the *MNIST*, *USPS*, *SVHN* experiment, first considered by Seguy et al. (2017). The goal is to stress-test OT-based methods in terms of their scalability with respect number of samples. In this case, replicate the experimental setting of Seguy et al. (2017). Furthermore, we consider 3 Cross-Domain Fault Diagnosis (CDFD) benchmarks, namely, Case Western Reserve University (CWRU), the Continuous Stirred Tank Reactor (CSTR) (Pilario & Cao, 2017; Montesuma et al., 2022) and the Tennessee Eastman (TE) processes Reinartz et al. (2021); Montesuma et al. (2024b). More details about these are given in the appendix. In each benchmark, we consider an adaptation task as a pair (S, T) of different domains.

Baselines and Backbones. For each visual domain adaptation task, with the exception of the *MNIST*, *USPS*, *SVHN* benchmark, we proceed as follows. First, we pre-train a ResNet He et al. (2016) using source domain data. The pre-trained network is then used to extract 2048-dimensional feature vectors for samples in source, and target domains. These feature serve as the basis for each domain adaptation algorithm. With the exception of GMM-OTDA_{MAP}, performance on the target domain is based on the generalization of a 1-layer neural network trained with transformed data. For GMM-OTDA_{MAP} we use the MAP strategy described in equation (16).

For the *MNIST*, *USPS*, *SVHN* benchmark, we follow Seguy et al. (2017) to obtain comparable results. For $M \rightarrow U$ and $U \rightarrow M$ we downsize MNIST to the resolution of USPS, i.e., we downscale images to (16, 16). For $M \rightarrow S$, we upscale MNIST images to match the resolution of SVHN, i.e., (32, 32), then we use features extracted from the last layer of a LeNet5. We refer readers to Seguy et al. (2017) and Struckmeier et al. (2023) for further details on these benchmarks.

For CDFD benchmarks, we follow the experimental settings of Montesuma et al. (2023a), Montesuma et al. (2022) and Montesuma et al. (2024b), which roughly follows a similar idea to visual adaptation tasks. This means that we pre-train a neural network with source domain data, then use its encoder for feature extraction. We refer readers to the original papers, and our appendix, for further information. To summarize our experimentation, there are in total 8 benchmarks, and 84 domain adaptation tasks.

Tested methods. Our main point of comparison is with other OT-based DA methods. We compare our GMM-OT for DA (OTDA) strategy with other OT methods for DA, namely, we consider the OTDA strategy of Courty et al. (2017) (Exact and Sinkhorn), the linear mapping estimation of Flamary et al. (2019), the mapping estimation technique of Perrot et al. (2016) (Linear and Kernel), and the InfoOT strategy of Chuang et al. (2023) (barycentric and conditional mappings).

For the scalability experiment using the *MNIST*, *USPS*, *SVHN* benchmark, we consider the large scale OT methods of Seguy et al. (2017), denoted as Alg. 1 and 2. Furthermore, for completeness, we consider the Linearly Alignable Optimal Transport (LaOT) strategy of Struckmeier et al. (2023).

5.1 Scalability with respect d

In this section, our goal is to evaluate how our method scales with the data dimensionality d . To do so, we evaluate methods based on their performance on visual adaptation benchmarks. The goal is to classify images into categories, based on 2048-dimensional vectors from ResNets (He et al., 2016) fine-tuned on the source domain of each adaptation task. We summarize our results in Figure 4. The detailed results may be found in the appendix, i.e., Table 5.

Over Caltech-Office and ImageCLEF, our methods outperform other state-of-the-art methods. For Office 31 and Office-Home, Info-OT_c and OTDA_{affine} proved to be more effect than our methods, but ours still ranks second. For Office 31, the density estimation strategy of Info-OT_c proves effective in finding a better map between the domains. For Office-Home, the baseline is already one of the best performing methods. An affine transformation is therefore sufficient for an effective adaptation.

Nonetheless, our method surpasses empirical OT over all tested benchmarks, especially preventing negative transfer in the Office-Home benchmark. Likewise, approximating class conditional distributions $P(X|Y)$ through GMMs proves effective over empirical OT. Indeed, our method improves over HOT-DA of El Hamri et al. (2022), which is based on empirical distributions.

These experiments prove that our method can effectively perform UDA between high-dimensional distributions. Note that, for these benchmarks, we use 2048-dimensional vectors, which is by far the largest dimensionality values considered in this study.




Figure 4: Average adaptation performance over 4 visual domain adaptation benchmarks.

5.2 Scalability with respect n

In this section, we use the *MNIST*, *USPS*, *SVHN* benchmark. Our goal is to evaluate how our method scales with the number of samples n . As we discussed in Remark 3.1, a major advantage of the GMM formulation is reducing OT complexity from $\mathcal{O}(n^3 \log n)$ to $\mathcal{O}(K^3 \log K)$, i.e., we replace number of samples by number of components, which are orders of magnitude inferior. To give a practical comparison, while *MNIST* has $n = 6 \times 10^5$ samples, we represent its probability distribution through a GMM with $K = 10^2$ components. This modeling choice improves scalability, especially when the components of GMMs have diagonal covariance matrices.

With respect performance, table 1 shows that GMM-OTDA has a similar performance to LaOT, i.e., performance degrades on $M \rightarrow U$, but increases on $U \rightarrow M$ and $M \rightarrow S$. Curiously, this corresponds to the case where a simpler dataset (e.g., *MNIST*) is transferred to a more complex dataset (e.g., *SVHN*). The performance similarity is not surprising, since LaOT and GMM-OTDA work under similar principles. However, the GMM modeling, again, proves superior to the Gaussian hypothesis, as we are able to improve performance on $U \rightarrow M$ and $M \rightarrow S$ tasks.

Algorithm	$M \rightarrow U$	$U \rightarrow M$	$M \rightarrow S$
Baseline	73.47	36.97	54.33
OTDA _{EMD}	57.75	52.46	-
OTDA _{Sink}	68.75	57.35	-
Alg. 1 with H	68.84	57.55	58.87
Alg. 1 with ℓ_2 .	67.80	57.47	60.56
Alg. 1 + 2 with H	77.92	60.02	61.11
Alg. 1 + 2 with ℓ_2 .	72.61	60.50	62.88
LaOT	72.57	62.28	60.36
GMM-OTDA (ours)	71.83	63.11	87.19

Table 1: Large scale OT experiment. We consider the adaptation between 3 digit recognition benchmarks, namely, *USPS*, *MNIST* and *SVHN*. Overall, GMM-OTDA largely outperforms other methods on harder adaptation tasks, i.e., $U \rightarrow M$ and $M \rightarrow S$.

5.3 Cross-Domain Fault Diagnosis

Case Western Reserve University Benchmark. With respect other benchmarks, the CWRU has the most number of samples per domain, i.e., 8000. In this case, InfoOT was intractable due its computational complexity. Furthermore, this benchmark illustrate the advantage of employing a grouping technique for enforcing the class structure in OT.

Task	Baseline	OTDA _{EMD}	OTDA _{Sink}	OTDA _{affine}	InfoOT _b	InfoOT _c	HOT-DA	GMM-OTDA _{MAP}	GMM-OTDA _T
$A \rightarrow B$	51.12	72.00	75.19	78.12	-	-	69.88	79.75	79.75
$A \rightarrow C$	62.88	94.12	100.00	95.62	-	-	100.00	99.94	100.00
$B \rightarrow A$	42.50	76.12	78.50	75.88	-	-	79.75	80.00	80.00
$B \rightarrow C$	37.44	77.62	78.88	75.38	-	-	79.81	79.56	79.94
$C \rightarrow A$	52.81	98.38	99.25	94.12	-	-	98.75	99.12	99.88
$C \rightarrow B$	55.62	70.25	74.50	75.50	-	-	83.12	79.75	80.00
Avg.	50.40	81.42	84.39	82.44	-	-	85.22	86.35	86.59

Table 2: Experimental results on the CWRU benchmark. For each task (i.e., each row), we highlight the best performing method in bold. Overall, InfoOT did not have a reasonable running time due the large number of samples on each domain.

Continuous Stirred Tank Reactor. As covered in Montesuma et al. (2022), the CSTR process carries an exothermic reaction $A \rightarrow B$. The reactor is jacketed, and an inflow of coolant is controlled by a Proportional, Integral, Derivative (PID) controller as described in Pilario & Cao (2017). From this reactor, a set of 7 variables are measured throughout simulation, corresponding to different temperatures, concentrations and flow-rates. We refer readers to Montesuma et al. (2022) for further details. Associated with this process, there are a set of 12 different faults, ranging from process and sensors faults, and input disturbances. On top of these 12 faults, there is the no-fault scenario, characterizing a classification problem with 13 classes.

Target Domain	1	2	3	4	5	6	Score
Reaction Order (N)	1.0	1.0	1.0	0.5	1.5	2.0	
Parameter Noise (η)	10%	15%	20%	15%	15%	15%	
Baseline	69.23	67.30	73.07	53.84	63.46	57.69	64.10
OTDA _{EMD}	71.15	71.15	71.15	61.53	57.69	50.00	63.78
OTDA _{Sink}	67.31	67.31	69.23	55.76	51.92	53.84	60.89
OTDA _{Affine}	65.38	71.15	71.15	61.54	65.38	53.84	64.74
InfoOT _b	67.31	67.31	67.31	50.00	51.92	40.07	58.65
InfoOT _c	71.15	67.31	71.15	53.84	51.92	48.07	60.57
HOT-DA	55.77	40.38	44.23	55.77	40.38	42.31	46.47
GMM-OTDA _{MAP}	78.84	73.07	73.07	61.54	57.69	55.77	66.67
GMM-OTDA _T	76.92	73.07	73.07	63.46	53.84	50.00	65.04

Table 3: Average classification accuracy with confidence intervals over a 5-fold cross-validation.

The different domains in this benchmark correspond to changes in the process conditions. These are of 2 kinds. First, one introduces a noise, η , in the process parameters (e.g., reactor or jacket volume), reflecting the possible uncertainty in the mathematical modeling of the reactor. Second, one changes the reaction order, N , of the reaction $A \rightarrow B$, which drastically changes the dynamics of the system. As Montesuma et al. (2022), we consider $\eta \in \{0.1, 0.15, 0.2\}$, and $N \in \{1, 0.5, 1.5, 2\}$. The source domain corresponds to $N = 1$, $\eta = 0.0$, whereas the 6 different targets correspond to combinations of η and N . These are shown in Table 3.

The CSTR benchmark poses a few challenges. First, it has a small number of samples on domains. While the source is composed of 1300 samples, each target only has 260. Second, each sample lies in a 1400-dimensional space. Third, target domains are noisy, due the parameter noise η . As a result, most methods have difficulty

in adapting, and performance usually degrades for more intense shifts (e.g., $N = 2.0$ and $\eta = 0.15$). However, GMM-OTDA_{MAP} and GMM-OTDA_T outperform other methods.

Tennessee Eastman Process. Our last experiment consists of the TE process, a benchmark widely used by the chemical engineering community (Reinartz et al., 2021). This benchmark has the largest number of domains, i.e., 6. Each of these domains is characterized by a different mode of production for the products of a chemical reaction, which affects the measured signals from the chemical plant. We summarize our results in Figure 5, which comprises the 30 adaptation tasks.

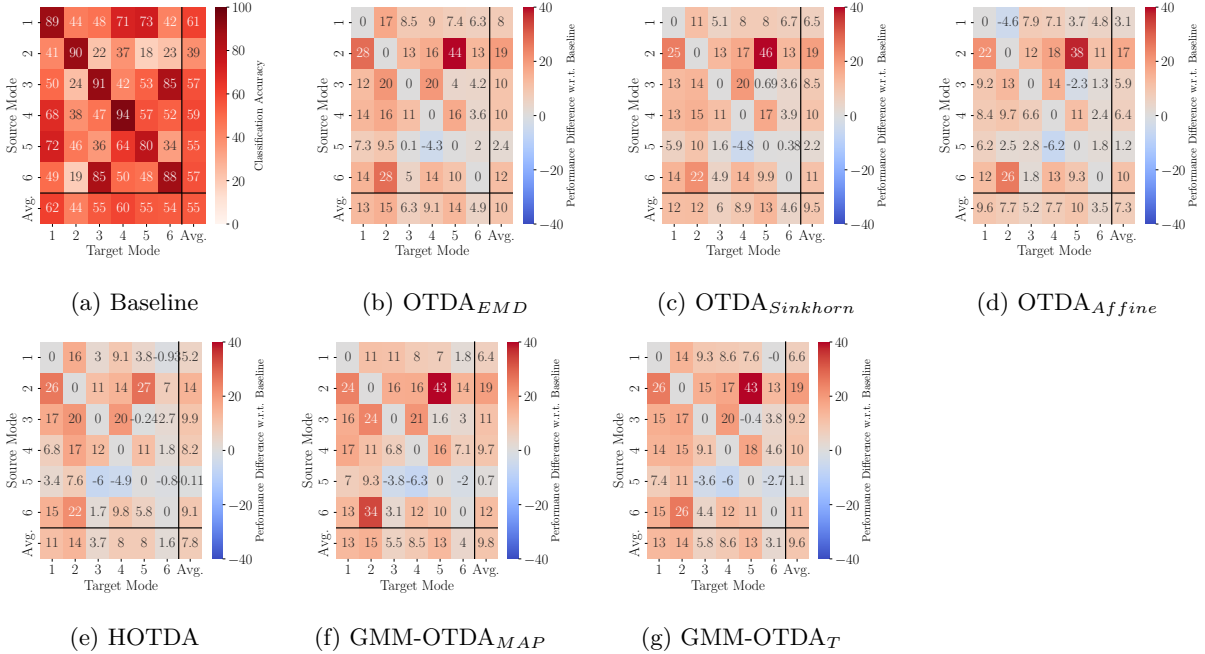


Figure 5: Domain adaptation results on the Tennessee Eastman Process. In (a), we show the baseline adaptation tasks, where each row represents a source domain, and each column represents a target domain. From (b) to (f), we show the performance offset with respect (a) of adaptation algorithms.

In comparison with other benchmarks, the TE feature vectors have fewer dimensions (i.e., 128). As a result, OTDA_{EMD} is the best performing method. However, GMM-OTDA manages to improve over OTDA_{affine} and HOT-DA. Overall, our methods are especially better on harder adaptation tasks, such as $6 \rightarrow 2$ and $3 \rightarrow 2$. We refer readers to the exploratory data analysis of Montesuma et al. (2024b) for further insights on why these adaptation tasks are harder.

6 Conclusion

In this paper, we consider the GMMOT framework of Delon & Desolneux (2020) as a candidate for UDA. Based on probability and OT theory, we devise 2 new effective strategies for UDA, namely, label propagation, and mapping estimation. The label propagation interprets the optimal transport plan between GMM components as a joint probability distribution over source-target component pairs. This modeling choice allows us to predict the label of target GMM components through a label propagation equation similar to Redko et al. (2019). Furthermore, we propose a mapping strategy that transports samples from the same component together through an affine map. This has a double advantage. First, it enforces group sparsity (Courty et al., 2017). Second, it has an analytical form in terms of GMM parameters. We show through a series of 84 UDA tasks that our methods outperform, or are competitive with the state-of-the-art, while being scalable with both number of samples n , and number of dimensions d . Our work further confirms previous studies on the intersection of GMMs and UDA, such as Montesuma et al. (2024a), showing that the GMMOT framework of Delon & Desolneux (2020) is a powerful candidate for domain adaptation.

References

- Barbara Caputo, Henning Müller, Jesus Martinez-Gomez, Mauricio Villegas, Burak Acar, Novi Patricia, Neda Marvasti, Suzan Üsküdarlı, Roberto Paredes, Miguel Cazorla, et al. Imageclef 2014: Overview and analysis of the results. In *Information Access Evaluation. Multilinguality, Multimodality, and Interaction: 5th International Conference of the CLEF Initiative, CLEF 2014, Sheffield, UK, September 15-18, 2014. Proceedings 5*, pp. 192–211. Springer, 2014.
- Yongxin Chen, Tryphon T Georgiou, and Allen Tannenbaum. Optimal transport for gaussian mixture models. *IEEE Access*, 7:6269–6278, 2018.
- Ching-Yao Chuang, Stefanie Jegelka, and David Alvarez-Melis. Infofoot: Information maximizing optimal transport. In *International Conference on Machine Learning*, pp. 6228–6242. PMLR, 2023.
- Nicolas Courty, Rémi Flamary, Devis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, 2017. doi: 10.1109/TPAMI.2016.2615921.
- Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- Julie Delon and Agnes Desolneux. A wasserstein-type distance in the space of gaussian mixture models. *SIAM Journal on Imaging Sciences*, 13(2):936–970, 2020.
- Arthur P Dempster, Nan M Laird, and Donald B Rubin. Maximum likelihood from incomplete data via the em algorithm. *Journal of the royal statistical society: series B (methodological)*, 39(1):1–22, 1977.
- Mourad El Hamri, Younes Bennani, and Issam Falih. Hierarchical optimal transport for unsupervised domain adaptation. *Machine Learning*, 111(11):4159–4182, 2022.
- Rémi Flamary, Karim Lounici, and André Ferrari. Concentration bounds for linear monge mapping estimation and optimal transport domain adaptation. *arXiv preprint arXiv:1905.10155*, 2019.
- Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *2012 IEEE conference on computer vision and pattern recognition*, pp. 2066–2073. IEEE, 2012.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- L Kantorovich. On the transfer of masses (in russian). In *Doklady Akademii Nauk*, volume 37, pp. 227–229, 1942.
- Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Histoire de l’Académie Royale des Sciences de Paris*, 1781.
- Eduardo Fernandes Montesuma and Fred Maurice Ngole Mboula. Wasserstein barycenter transport for acoustic adaptation. In *ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3405–3409, May 2021a. doi: 10.1109/ICASSP39728.2021.9414199.
- Eduardo Fernandes Montesuma and Fred Ngolè Mboula. Wasserstein barycenter for multi-source domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 16785–16793, June 2021b.
- Eduardo Fernandes Montesuma, Michela Mulas, Francesco Corona, and Fred-Maurice Ngole Mboula. Cross-domain fault diagnosis through optimal transport for a cstr process. *IFAC-PapersOnLine*, 55(7):946–951, 2022.
- Eduardo Fernandes Montesuma, Fred Ngolè Mboula, and Antoine Souloumiac. Multi-source domain adaptation through dataset dictionary learning in wasserstein space. In *26th European Conference on Artificial Intelligence (ECAI)*, pp. 1739–1746. IOS Press, 2023a.

- Eduardo Fernandes Montesuma, Fred Ngole Mboula, and Antoine Souloumiac. Recent advances in optimal transport for machine learning. *arXiv preprint arXiv:2306.16156*, 2023b.
- Eduardo Fernandes Montesuma, Fred Ngolè Mboula, and Antoine Souloumiac. Lighter, better, faster multi-source domain adaptation with gaussian mixture models and optimal transport. In Albert Bifet, Jesse Davis, Tomas Krilavičius, Meelis Kull, Eirini Ntoutsi, and Indrė Žliobaitė (eds.), *Machine Learning and Knowledge Discovery in Databases. Research Track*, pp. 21–38, Cham, 2024a. Springer Nature Switzerland. ISBN 978-3-031-70365-2.
- Eduardo Fernandes Montesuma, Michela Mulas, Fred Ngolè Mboula, Francesco Corona, and Antoine Souloumiac. Benchmarking domain adaptation for chemical processes on the tennessee eastman process. In *Machine Learning for Chemistry and Chemical Engineering (ML4CCE) Workshop*, 2024b.
- Simno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359, 2009.
- Xingchao Peng, Qinxun Bai, Xide Xia, Zijun Huang, Kate Saenko, and Bo Wang. Moment matching for multi-source domain adaptation. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 1406–1415, 2019.
- Michaël Perrot, Nicolas Courty, Rémi Flamary, and Amaury Habrard. Mapping estimation for discrete optimal transport. *Advances in Neural Information Processing Systems*, 29, 2016.
- Gabriel Peyré, Marco Cuturi, et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Karl Ezra Pilario and Yi Cao. Process incipient fault detection using canonical variate analysis. In *2017 23rd International Conference on Automation and Computing (ICAC)*, pp. 1–6. IEEE, 2017.
- Joaquin Quinonero-Candela, Masashi Sugiyama, Anton Schwaighofer, and Neil D Lawrence. *Dataset shift in machine learning*. Mit Press, 2008.
- Ievgen Redko, Amaury Habrard, and Marc Sebban. Theoretical analysis of domain adaptation with optimal transport. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pp. 737–753. Springer, 2017.
- Ievgen Redko, Nicolas Courty, Rémi Flamary, and Devis Tuia. Optimal transport for multi-source domain adaptation under target shift. In *The 22nd International Conference on artificial intelligence and statistics*, pp. 849–858. PMLR, 2019.
- Ievgen Redko, Emilie Morvant, Amaury Habrard, Marc Sebban, and Younès Bennani. A survey on domain adaptation theory: learning bounds and theoretical guarantees. *arXiv preprint arXiv:2004.11829*, 2020.
- Christopher Reinartz, Murat Kulahci, and Ole Ravn. An extended tennessee eastman simulation dataset for fault-detection and decision support systems. *Computers & Chemical Engineering*, 149:107281, 2021.
- Kate Saenko, Brian Kulis, Mario Fritz, and Trevor Darrell. Adapting visual category models to new domains. In *European conference on computer vision*, pp. 213–226. Springer, 2010.
- Filippo Santambrogio. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015.
- Vivien Seguy, Bharath Bhushan Damodaran, Rémi Flamary, Nicolas Courty, Antoine Rolet, and Mathieu Blondel. Large-scale optimal transport and mapping estimation. *arXiv preprint arXiv:1711.02283*, 2017.
- Oliver Struckmeier, Ievgen Redko, Anton Mallasto, Karol Arndt, Markus Heinonen, and Ville Kyrki. Learning representations that are closed-form monge mapping optimal with application to domain adaptation. *Transactions on Machine Learning Research*, 2023. ISSN 2835-8856. URL <https://openreview.net/forum?id=nOIGfQnFZm>.

- Masashi Sugiyama, Matthias Krauledat, and Klaus-Robert Müller. Covariate shift adaptation by importance weighted cross validation. *Journal of Machine Learning Research*, 8(5), 2007.
- Asuka Takatsu. Wasserstein geometry of Gaussian measures. *Osaka Journal of Mathematics*, 48(4):1005 – 1026, 2011.
- Laurens Van der Maaten and Geoffrey Hinton. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Vladimir Vapnik. *The nature of statistical learning theory*. Springer science & business media, 2013.
- Hemanth Venkateswara, Jose Eusebio, Shayok Chakraborty, and Sethuraman Panchanathan. Deep hashing network for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5018–5027, 2017.
- Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

A Additional Details about Benchmarks

Benchmark	Domains	Backbone	# Samples	# Classes
ImageCLEF	Caltech (C)	ResNet 50	600	12
	Bing (B)		600	
	ImageNet (I)		600	
	Pascal (P)		600	
	Total		2400	
Caltech-Office 10	Amazon (A)	ResNet 101	958	10
	dSLR (D)		157	
	Webcam (W)		295	
	Caltech (C)		1123	
	Total		2533	
Office 31	Amazon (A)	ResNet 50	2817	31
	dSLR (D)		498	
	Webcam (W)		795	
	Total		4110	
Office-Home	Art (Ar)	ResNet 101	2427	65
	Clipart (Cl)		4365	
	Product (Pr)		4439	
	Real World (Rw)		4357	
	Total		15588	

(a) Visual Domain Adaptation Benchmarks

Benchmark	Domains	Backbone	# Samples	# Classes
CWRU	1772rpm (A)	MLP	8000	10
	1750rpm (B)		8000	
	1730rpm (C)		8000	
	Total		24000	
TEP	Mode 1	Fully Convolutional	2900	29
	Mode 2		2845	
	Mode 3		2899	
	Mode 4		2865	
	Mode 5		2883	
	Mode 6		2897	
	Total		17289	
CSTR	$N = 1.0, \epsilon = 0.00$	-	1300	13
	$N = 1.0, \epsilon = 0.10$		260	
	$N = 1.0, \epsilon = 0.15$		260	
	$N = 0.5, \epsilon = 0.15$		260	
	$N = 1.5, \epsilon = 0.15$		260	
	$N = 2.0, \epsilon = 0.15$		260	
	Total		2860	

(b) Cross-Domain Fault Diagnosis Benchmarks

Table 4: Overview of Visual Domain Adaptation and Cross-Domain Fault Diagnosis benchmarks

In table 4, we show an overview of the used benchmarks. We run our experiments on 4 visual DA datasets, and 3 CDFD datasets. For vision, we use Residual Networks (ResNets) pre-trained on ImageNET as the backbone. For each adaptation task (e.g., $C \rightarrow B$ in ImageCLEF) we fine tune the network using labeled source domain data. For all methods, we extract the features of source, and target domain, using the fine-tuned checkpoint. The size of the ResNet is used to agree with previous research, such as Peng et al. (2019) and Montesuma et al. (2023a).

For CDFD, we considered the same setting as previous works using these benchmarks, such as Montesuma et al. (2023a), Montesuma et al. (2024b) and Montesuma et al. (2022), for CWRU, TE process and CSTR respectively. For CWRU, we extract windows out of raw signals of size 2048, then get the frequency representation for these windows using a fast Fourier transform. These are treated as 2048 feature vectors, that are then fed to a neural network. In the TE process, we consider the same setting of Montesuma et al. (2024b), i.e., we use a fully convolutional neural net. For the CSTR, we directly use the signals, concatenated into a 1400-dimensional vector as the features. The data for these benchmarks is publicly available here¹, and here².

Note that, for each dataset, there are $n_{domains}(n_{domains} - 1)$ adaptation tasks, except for CSTR, which is a multi-target benchmark (i.e., a single source, and 6 targets). As a result, we have $12 \times 3 + 6 \times 2 + 30 + 6 = 84$ adaptation tasks.

¹<https://www.kaggle.com/datasets/eddardd/tennessee-eastman-process-domain-adaptation>

²<https://www.kaggle.com/datasets/eddardd/continuous-stirred-tank-reactor-domain-adaptation>

B Additional Experiments

B.1 Detailed Results

Benchmark	Task	Baseline	OTDA _{EMD}	OTDA _{Sink}	OTDA _{affine}	InfoOT _b	InfoOT _c	HOT-DA	GMM-OTDA _{MAP}	GMM-OTDA _T
Caltech-Office	$A \rightarrow D$	87.10	77.42	87.10	93.55	93.55	83.87	96.77	93.55	80.65
	$A \rightarrow W$	91.53	93.22	96.61	96.61	96.61	94.92	96.61	93.22	91.53
	$A \rightarrow C$	88.44	91.56	73.33	91.11	87.11	90.67	74.67	88.44	88.44
	$D \rightarrow A$	88.54	90.10	93.23	92.19	91.15	92.71	96.88	96.35	95.83
	$D \rightarrow W$	98.31	93.22	93.22	94.92	100.00	91.53	94.92	100.00	98.31
	$D \rightarrow C$	75.56	67.11	17.33	71.56	68.44	75.11	51.11	86.67	71.56
	$W \rightarrow A$	85.94	82.81	15.62	81.25	60.94	86.46	69.79	87.50	84.90
	$W \rightarrow D$	100.00	93.55	93.55	96.77	90.32	93.55	90.32	96.77	93.55
	$W \rightarrow C$	84.89	87.56	87.56	88.44	87.56	87.56	84.89	87.56	88.00
	$C \rightarrow A$	98.44	96.88	98.44	98.44	98.44	96.88	98.44	98.44	98.44
	$C \rightarrow D$	93.55	87.10	90.32	93.55	87.10	87.10	100.0	93.55	80.65
	$C \rightarrow W$	91.53	94.92	94.92	94.92	93.22	93.22	98.31	96.61	94.92
	Avg.	90.32	87.95	78.44	91.11	87.87	89.46	87.72	93.22	88.90
ImageCLEF	$B \rightarrow C$	92.50	95.00	95.00	94.17	97.50	96.67	96.67	95.00	94.17
	$B \rightarrow I$	90.00	89.17	89.17	91.67	94.17	91.67	95.00	95.00	93.33
	$B \rightarrow P$	68.33	69.17	70.00	71.67	73.33	75.83	74.17	72.50	74.17
	$C \rightarrow B$	65.00	65.83	65.83	65.00	51.67	62.50	62.50	65.00	66.67
	$C \rightarrow I$	89.17	96.67	96.67	95.00	92.50	97.50	95.83	94.17	96.67
	$C \rightarrow P$	71.67	74.17	73.33	71.67	75.00	75.83	72.50	70.83	75.83
	$I \rightarrow B$	68.33	70.00	68.33	70.00	65.00	66.67	61.67	67.50	70.00
	$I \rightarrow C$	93.33	95.83	95.83	95.83	95.83	96.67	95.83	95.00	95.83
	$I \rightarrow P$	71.67	74.17	75.00	73.33	73.33	71.67	72.50	73.33	75.00
	$P \rightarrow B$	67.50	69.17	66.67	68.33	57.50	65.83	62.50	62.50	64.17
	$P \rightarrow C$	95.00	95.83	95.83	95.00	96.67	96.67	96.67	95.00	95.00
	$P \rightarrow I$	90.83	92.50	92.50	90.83	95.83	95.83	95.00	94.17	95.00
	Avg.	80.28	82.29	82.01	81.88	80.69	82.78	81.74	81.67	82.99
Office 31	$A \rightarrow D$	66.07	68.75	69.64	69.64	75.89	76.79	72.32	69.64	72.32
	$A \rightarrow W$	76.02	74.27	80.12	80.12	79.53	79.53	73.68	76.61	80.70
	$D \rightarrow A$	65.68	65.85	67.77	66.90	67.60	66.20	61.15	68.29	73.52
	$D \rightarrow W$	94.15	95.32	98.25	98.25	95.91	97.08	84.80	98.83	95.32
	$W \rightarrow A$	63.41	66.90	67.42	65.51	67.60	67.60	61.67	66.38	65.68
	$W \rightarrow D$	96.43	90.18	92.86	95.54	87.50	91.96	81.25	91.96	91.07
	Avg.	76.96	76.88	79.34	79.32	79.00	79.86	72.48	78.62	79.77
Office-Home	$Ar \rightarrow Cl$	55.10	54.98	54.87	56.24	17.41	53.95	47.88	53.95	57.96
	$Ar \rightarrow Pr$	70.95	68.69	71.96	71.96	30.97	70.27	67.23	74.89	74.10
	$Ar \rightarrow Rw$	79.68	79.68	80.83	80.71	40.53	80.25	76.00	77.96	82.43
	$Cl \rightarrow Ar$	63.51	60.62	63.09	62.68	31.34	62.68	53.81	59.79	64.33
	$Cl \rightarrow Pr$	69.26	66.89	68.81	70.72	41.78	68.92	63.51	70.05	71.73
	$Cl \rightarrow Rw$	72.68	69.92	71.18	72.33	38.81	71.18	67.97	68.66	74.63
	$Pr \rightarrow Ar$	66.80	62.47	64.12	66.39	32.16	64.33	55.88	57.73	62.89
	$Pr \rightarrow Cl$	36.88	38.83	25.32	38.83	8.59	30.93	23.71	30.70	31.62
	$Pr \rightarrow Rw$	78.76	77.84	79.22	79.22	47.99	78.30	71.64	73.59	80.94
	$Rw \rightarrow Ar$	72.99	71.96	72.37	73.81	51.13	70.72	62.47	66.39	69.69
	$Rw \rightarrow Cl$	53.15	57.85	57.39	56.93	37.00	56.59	47.65	50.63	56.36
	$Rw \rightarrow Pr$	82.32	80.86	81.87	82.21	64.86	81.31	72.30	80.41	82.09
	Avg.	66.84	65.88	65.92	67.67	36.88	65.79	59.17	63.73	67.40

Table 5: Single-source domain adaptation results. We compare 8 methods over 5 benchmarks, with a total of 42 adaptation tasks.

B.2 Ablations and Visualization

Ablating the number of components and entropic regularization. In this experiment, we ablate the two parameters of our methods, namely, the number of components K , and the entropic regularization ϵ . Recall that we normalize the ground-cost by the maximum value, i.e., $\hat{C}_{ij} = C_{ij}/(\max_{ij} C_{ij})$, which improves the numerical stability of the Sinkhorn algorithm. We summarize our results in Figure 6.

For the number of components K , the relationship with performance is mostly clear. Indeed, except for $\epsilon = 10^{-1}$ on the mapping strategy, *using more components enhances performance*. Note that, even though this implies a more complex GMMs, we still have far less components than samples ($n = 600$ per domain, i.e., 5 times more samples than components).




Figure 6: Ablation on number of components K , and entropic penalty ϵ , for the MAP estimation strategy based on labeled propagation (left), and the mapping estimation strategy (right).

For the entropic penalty ϵ , we have two drastically different scenarios. For the MAP estimation, using higher entropic regularization coefficients improve performance, whereas the mapping strategy works better for smaller regularization coefficients (or exact OT). While this may seem contradictory, we note that, in the mapping strategy, we actually filter out irrelevant matchings between components based on a parameter τ . However, the entropic regularization is known to generate smoother couplings, which means that more entries of ω are non-zero, or possibly greater than a fixed τ . As a consequence, the mapping strategy ends up behaving like T_{rand} , which causes a bad reconstruction for the target domain. Naturally, this effect gets amplified with more components in both GMMs, as there are more possible matchings.




Figure 7: Source and target domain samples alongside the centroids (denoted by stars) found through EM. colors reflect the different classes.

Visualizing components and mapped samples. In this experiment, we use the CWRU benchmark adaptation task $A \rightarrow C$. We start by embedding the source and target domain data with the t-Stochastic Neighbor Embeddings (SNE) technique of Van der Maaten & Hinton (2008). We do this in 2 separate plots, where we concatenate the source domain features with the centroids obtained by running the EM algorithm (resp. target). This visualization is shown in Figure 7. Next, we map samples from the source to the target domain, with the various strategies described in this section, with the exception of InfoOT, which did not had a reasonable running time. These are shown in Figure 8.

Overall, while the exact OT solution provides a measure that better reflects the feature positions, it mixes the classes, as evidenced by the orange and blue classes being mapped to the same place, as well as the green and violet classes. This phenomenon does not happen for other methods, at the cost of having mapped




Figure 8: t-SNE visualisation of mapped samples to the target domain, on the CWRU task $A \rightarrow C$.

points distributed in a different way. However, for $\text{OTDA}_{\text{sink}}$ and GMM-OTDA_T , the mapped points better respect the class boundaries. For $\text{OTDA}_{\text{affine}}$, note that the mapped distribution does not actually match the target. Overall, we achieve a better mapping through the GMM modeling.