
Making transport more robust and interpretable by moving data through a small number of anchor points

Chi-Heng Lin¹ Mehdi Azabou^{1 2} Eva L. Dyer^{1 2 3}

Abstract

Optimal transport (OT) is a widely used technique for distribution alignment, with applications throughout the machine learning, graphics, and vision communities. Without any additional structural assumptions on transport, however, OT can be fragile to outliers or noise, especially in high dimensions. Here, we introduce Latent Optimal Transport (LOT), a new approach for OT that simultaneously learns low-dimensional structure in data while leveraging this structure to solve the alignment task. The idea behind our approach is to learn two sets of “anchors” that constrain the flow of transport between a source and target distribution. In both theoretical and empirical studies, we show that LOT regularizes the rank of transport and makes it more robust to outliers and the sampling density. We show that by allowing the source and target to have different anchors, and using LOT to align the latent spaces between anchors, the resulting transport plan has better structural interpretability and highlights connections between both the individual data points and the local geometry of the datasets.

1. Introduction

Optimal transport (OT) (Villani, 2008) is a widely used technique for distribution alignment that learns a *transport plan* which moves mass from one distribution to match another. With recent advances in tools for regularizing and speeding up OT (Cuturi, 2013), this approach has found applications in many diverse areas of machine learning, including domain adaptation (Courty et al.; Courty et al., 2017), genera-

tive modeling (Martin Arjovsky & Bottou, 2017; Tolstikhin et al., 2017), document retrieval (Kusner et al., 2015), computer graphics (Solomon et al., 2014; 2015; Bonneel et al., 2016), and computational neuroscience (Gramfort et al., 2015; Lee et al., 2019).

While the ground metric in OT can be used to impose geometric structure into transport, without any additional assumptions, OT can be fragile to outliers or noise, especially in high dimensions. To overcome this issue, additional structure, either in the data or in the transport plan, can be used to improve alignment or make transport more robust. Examples of methods that incorporate additional structure into OT include approaches that leverage hierarchical structure or cluster consistency (Lee et al., 2019; Yurochkin et al., 2019; Xu et al., 2020), partial class information (Courty et al., 2017; Courty et al.), submodular cost functions (Alvarez-Melis et al., 2018), and low-rank constraints on the transport plan (Forrow et al., 2019; Altschuler et al., 2019). Because of the difficulty of incorporating structure into OT, many of these methods need low-dimensional structure in data to be specified in advance (e.g., estimated clusters or labels).

To simultaneously learn low-dimensional structure and use it to constrain transport, Forrow et al. (2019) recently introduced a statistical approach for OT that builds a factorization of the transport plan to regularize its rank. After factorization, transport from a source to target distribution can be visualized as the flow of mass through a small number of anchors (hubs), which serve as relay stations through which transportation must pass (see Figure 1, a vs. b). Although this idea of moving data through anchors is appealing, in previous work, the anchors used to constrain transport are *shared* by the source and target. As a result, when the source and target contain different structures or experience domain shift (Courty et al.), shared anchors may not provide an adequate representation for both domains simultaneously.

In this work, we propose a new structured transport approach called Latent Optimal Transport (LOT). The main idea behind LOT is to factorize the transport plan into three components, where mass is moved: (i) from individual source points to source anchors, (ii) from the source anchors to target anchors, and (iii) from target anchors to individual target points (Figure 1c-d). The intermediate transport

¹Department of Electrical and Computer Engineering, Georgia Tech, Atlanta, Georgia, USA. ²Machine Learning Program, Georgia Tech, Atlanta, Georgia, USA ³Coulter Department of Biomedical Engineering, Georgia Tech & Emory University, Atlanta, Georgia, USA. Correspondence to: Chi-Heng Lin <cl3385@gatech.edu>, Mehdi Azabou <mazabou@gatech.edu>, Eva L. Dyer <evadyer@gatech.edu>.

plan captures the high-level structural similarity between the source and target, while the outer transport plans cluster data in their respective spaces. In both theoretical and empirical studies, we show that LOT regularizes the rank of transport and has the effect of denoising the transport plan, making it more robust to outliers and sampling. By allowing the source and target to have different anchors and aligning the latent spaces of the anchors, we show that the mapping between datasets can be more easily interpreted.

Specifically, our contributions are as follows. (i) We introduce LOT, a new form of structured transport, and propose an efficient algorithm that solves our proposed objective (Section 3), (ii) Theoretically, we show that LOT can be interpreted as a relaxation to OT, and from a statistical point-of-view, it overcomes the curse of dimensionality in terms of the sampling rate (Section 5), (iii) We study the robustness of the approach to noise, sampling, and various data perturbations when applied to both synthetic data and domain adaptation problems in computer vision (Section 6).

2. Background

Optimal Transport: Optimal transport (OT) (Villani, 2008; Santambrogio, 2015; Peyré et al., 2019) is a distribution alignment technique that learns a transport plan that specifies how to move mass from one distribution to match another. Specifically, consider two sets of data points encoded in matrices, the *source* $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]$ and the *target* $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_m]$, where $\mathbf{x}_i \in \mathcal{X}$, $\mathbf{y}_j \in \mathcal{Y}$, $\forall i, j$. Assume they are endowed with discrete measures $\mu = \sum_{i=1}^N p(\mathbf{x}_i) \delta_{\mathbf{x}_i}$, $\nu = \sum_{j=1}^M p(\mathbf{y}_j) \delta_{\mathbf{y}_j}$, respectively. The cost of transporting \mathbf{x}_i to \mathbf{y}_j is $c(\mathbf{x}_i, \mathbf{y}_j)$, where c denotes some cost function. OT considers the most cost-efficient transport by solving the following problem:¹

$$\text{OT}_{\mathbf{C}}(\mu, \nu) := \min_{\mathbf{P} \geq 0, \mathbf{P}^T \mathbf{1} = \nu} \langle \mathbf{C}, \mathbf{P} \rangle, \quad (1)$$

where $\mathbf{P} := [p(\mathbf{x}_i, \mathbf{y}_j)]_{i,j}$ is the source-to-target transport plan matrix (coupling), and $\mathbf{C} = [c(\mathbf{x}_i, \mathbf{y}_j)]_{i,j}$ is the cost matrix. When $c(\mathbf{x}, \mathbf{y}) = d(\mathbf{x}, \mathbf{y})^p$, where d is a distance function, $\mathcal{W}_p := \text{OT}_{\mathbf{C}}^{1/p}$ defines a distance called the p -Wasserstein distance. The objective in (1) is a linear programming problem, where computation speed can be prohibitive if n is large (Pele & Werman, 2009). A common speedup is to replace the objective by an entropy-regularized proxy,

$$\begin{aligned} \text{OT}_{\mathbf{C}, \varepsilon}(\mu, \nu) &:= \min_{\mathbf{P} \geq 0, \mathbf{P}^T \mathbf{1} = \nu} \langle \mathbf{C}, \mathbf{P} \rangle - \varepsilon \mathbf{H}(\mathbf{P}) \\ &= \min_{\mathbf{P} \geq 0, \mathbf{P}^T \mathbf{1} = \nu} \varepsilon \text{KL}(\mathbf{P} \| \mathbf{K}), \end{aligned} \quad (2)$$

¹The problem can be generalized to setting of continuous measures by $\text{OT}_c(\mu, \nu) = \min_{\gamma \in \mathcal{G}} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\gamma(x, y)$, $\mathcal{G} = \{\gamma : \int_{\mathcal{Y}} d\gamma(x, y) = \mu, \int_{\mathcal{X}} d\gamma(x, y) = \nu\}$.

where \mathbf{K} is the Gibbs kernel induced by the element-wise exponential of the cost matrix $\mathbf{K} := \exp(-\mathbf{C}/\varepsilon)$, $\mathbf{H}(\mathbf{P}) := -\sum_{i,j} \mathbf{P}_{i,j} \log(\mathbf{P}_{i,j})$ is the Shannon entropy, and ε is a user-specified hyperparameter that controls the amount of entropic regularization that is introduced. We can alternatively write the objective function as a minimization of $\varepsilon \text{KL}(\mathbf{P} \| \mathbf{K})$, where KL denotes the Kullback-Leibler divergence. In practice, the entropy-regularized form is often used over the original objective (1) as it admits a fast method called the Sinkhorn algorithm (Cuturi, 2013; Altschuler et al., 2017). Hence, we will use OT to refer to the entropy-regularized form unless specified otherwise in the context.

Optimal Transport via Factored Couplings: Factored Coupling (FC) is proposed in (Forrow et al., 2019) to reduce the sample complexity of OT in high dimensions. Specifically, it adds an additional constraint to (1) by enforcing the transport plan to be of the following factored form,

$$p(\mathbf{x}_i, \mathbf{y}_j) = \sum_{l=1}^k p(\mathbf{z}_l) p(\mathbf{x}_i | \mathbf{z}_l) p(\mathbf{y}_j | \mathbf{z}_l). \quad (3)$$

This has a nice interpretation: \mathbf{z}_l serves as a common “anchor” that transportation from \mathbf{x}_i to \mathbf{y}_j must pass through. It turns out that FC is closely related to the Wasserstein barycenter problem (Agueh & Carlier, 2011; Cuturi & Doucet, 2014; Cuturi & Peyré, 2016), $\min_{\nu} \sum_{i=1}^N \mathcal{W}_2^2(\mu_i, \nu)$, where ν is the Procrustes mean to distributions μ_i , $i = 1, \dots, N$ with respect to the squared 2-Wasserstein distance. A crucial insight from (Forrow et al., 2019) is that for $N = 2$, the barycenter ν could approximate the optimal anchors to a transport plan of the form (3) that minimizes the objective in (1).

3. Latent Optimal Transport

3.1. Motivation

Most datasets have low-dimensional latent structure, but OT does not naturally use it during transport. This motivates the idea that distribution alignment methods should both *reveal* the latent structure in the data in addition to aligning these latent structures. An illustrative example is provided in Figure 1; here, we show the transport plan for a source (red points) and a target (blue points), both of which exhibit clear cluster structures. Because OT transports points independently, the points can be easily mapped outside of their original cluster (a). In comparison, low-rank OTs (b-d) induce transport plans that are better at preserving clusters. In (b), because factored coupling (FC) transports points via common anchors (black squares), the anchors need to interpolate between both distributions, and it loses the freedom of choosing different structures for the source and target. On the other hand, by specifying different numbers of anchors for the source and target individually (c vs. d), LOT can

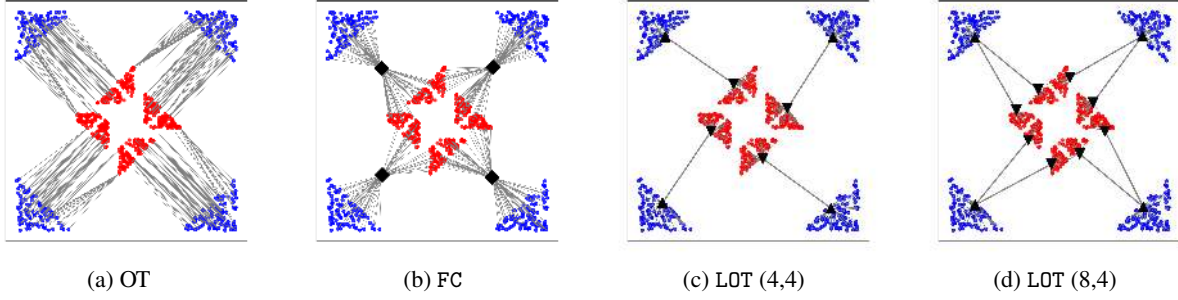


Figure 1: *Comparisons of transport plans obtained for different methods applied to clustered data after domain shift.* Here, we visualize the connection between the source (blue) x and its estimated target (red) $\hat{y} = \arg\max_y p(y|x)$. From left to right, we show the standard OT plan (a) and the factored coupling (FC) approach (b). To the right, we show the result of LOT when we use 4 anchors in the target with the same number in the target (c) and 8 anchors in the target (d).

extract different structures and output different transport plans.

3.2. Problem formulation

Consider data matrices \mathbf{X} and \mathbf{Y} and their measures μ, ν , as detailed in Section 2. We introduce “anchors” through which points must flow, thus constraining the transportation. The anchors are stacked in data matrices $\mathbf{Z}_x := [\mathbf{z}_1^x, \dots, \mathbf{z}_{k_x}^x]$, $\mathbf{Z}_y := [\mathbf{z}_1^y, \dots, \mathbf{z}_{k_y}^y]$. We denote the measures of the source and target anchors as $\mu_z = \sum_{m=1}^{k_x} p(\mathbf{z}_m^x) \delta_{\mathbf{z}_m^x}$ and $\nu_z = \sum_{n=1}^{k_y} p(\mathbf{z}_n^y) \delta_{\mathbf{z}_n^y}$. For any set \mathcal{A} , we further denote $\Delta_{\mathcal{A}}^k := \left\{ \sum_{i=1}^k \omega_i \delta_{\mathbf{a}_i} : \sum_{i=1}^k \omega_i = 1, \omega_i \geq 0, \mathbf{a}_i \in \mathcal{A}, \forall i \right\}$ as the set of probability measures on \mathcal{A} that has discrete support of size up to k . Hence $\mu_z \in \Delta_{\mathcal{Z}_x}^{k_x}, \nu_z \in \Delta_{\mathcal{Z}_y}^{k_y}$, where \mathcal{Z}_x (resp. \mathcal{Z}_y) is the space of source (resp. target) anchors. If we interpret the conditional probability $p(a|b)$ as the strength of transportation from b to a , then, using the chain rule, the concurrence probability $p(\mathbf{x}_i, \mathbf{y}_j)$ of \mathbf{x}_i and \mathbf{y}_j can be written as,

$$\begin{aligned} p(\mathbf{x}_i, \mathbf{y}_j) &= \sum_{m,n} p(\mathbf{x}_i) p(\mathbf{z}_m^x | \mathbf{x}_i) p(\mathbf{z}_n^y | \mathbf{z}_m^x) p(\mathbf{y}_j | \mathbf{z}_n^y) \\ &= \sum_{m,n} p(\mathbf{x}_i, \mathbf{z}_m^x) \frac{p(\mathbf{z}_m^x, \mathbf{z}_n^y)}{p(\mathbf{z}_m^x) p(\mathbf{z}_n^y)} p(\mathbf{z}_n^y, \mathbf{y}_j). \end{aligned} \quad (4)$$

When encoding these probabilities using a transport matrix $\mathbf{P} := [p(\mathbf{x}_i, \mathbf{y}_j)]_{i,j}$, the factorized form (4) can be written as,

$$\mathbf{P} = \mathbf{P}_x \text{diag}(\mathbf{u}_z^{-1}) \mathbf{P}_z \text{diag}(\mathbf{v}_z^{-1}) \mathbf{P}_y, \quad (5)$$

where \mathbf{P}_x encodes transport from source space to source anchor space (i.e., $p(\mathbf{x}_i, \mathbf{z}_m^x)$), \mathbf{P}_z encodes transport from source anchor space to target anchor space, \mathbf{P}_y encodes transport from target anchor space to target space, and $\mathbf{u}_z := [p(\mathbf{z}_1^x), \dots, p(\mathbf{z}_{k_x}^x)]$, $\mathbf{v}_z := [p(\mathbf{z}_1^y), \dots, p(\mathbf{z}_{k_y}^y)]$ encode the latent distributions of anchors. To learn each of

these transport plans, we must first designate the ground metric used to define the cost in each of the three stages. The cost matrices $\mathbf{C}_x, \mathbf{C}_y$ determine how points will be transported to their respective anchors and thus dictate how the data structure will be extracted. We will elaborate on the choice of costs in Section 3.3.

We now formalize our proposed approach to transport in the following definition.

Definition 1. Let $\mathbf{C}_x, \mathbf{C}_y$ denote the cost matrices between the source/target and their representative anchors, and let \mathbf{C}_z denote the cost matrix between anchors. We define the latent optimal transport (LOT) problem as,

$$\begin{aligned} \text{OT}^L(\mu, \nu) := & \inf_{\mu_z \in \Delta_{\mathcal{Z}_x}^{k_x}, \nu_z \in \Delta_{\mathcal{Z}_y}^{k_y}} \{ \text{OT}_{\mathbf{C}_x}(\mu, \mu_z) \\ & + \text{OT}_{\mathbf{C}_z}(\mu_z, \nu_z) + \text{OT}_{\mathbf{C}_y}(\nu_z, \nu) \}, \end{aligned}$$

where \mathcal{Z}_x and \mathcal{Z}_y are the latent spaces of the source and target anchors, respectively.²

The intuition behind Def. 1 is that we use $\text{OT}_{\mathbf{C}_x}(\mu, \mu_z)$ and $\text{OT}_{\mathbf{C}_y}(\nu_z, \nu)$ to capture group structure in each space, and then $\text{OT}_{\mathbf{C}_z}(\mu_z, \nu_z)$ to align the source and target by determining the transportation across anchors. Hence, LOT can be interpreted as an optimization of joint clustering and alignment. The flexibility of cost matrices allows LOT to capture different structures and induce different transport plans. In Section 5, we further show that LOT can be regarded as a relaxation of an OT problem.

Remark 1. In Forrow et al. (2019), the authors introduce the notion of the transport rank for a transport plan \mathbf{P} as the minimum number of product probability measures that its corresponding coupling can be composed from, i.e., $p(\mathbf{x}, \mathbf{y}) = \sum_{i=1}^r \lambda_i (p_i(\mathbf{x}) \otimes p_i(\mathbf{y}))$, $\lambda_i \geq 0, \forall i$. In general, given a transportation plan \mathbf{P} , the transport rank $\text{rank}_+(\mathbf{P})$ is lower bounded by its usual matrix rank

²This definition extends naturally to continuous measures by replacing cost matrix \mathbf{C} with cost function c .

$\text{rank}(\mathbf{P})$. In the case of LOT, the transport plan induced by Def. 1 satisfies $\text{rank}(\mathbf{P}) \leq \text{rank}_+(\mathbf{P}) \leq \min(k_x, k_y)$. Thus, by selecting a small number of anchors we naturally induce a low-rank solution for transport.

Next, we show some properties of LOT that highlight its similarity to a metric.

Proposition 1. Suppose the latent spaces $\mathcal{Z}_x = \mathcal{Z}_y$ are the same as the original data spaces $\mathcal{X} = \mathcal{Y}$, and the cost matrices are defined by $\mathbf{C}_x[a, b] = \mathbf{C}_z[a, b] = \mathbf{C}_y[a, b] = d(a, b)^p$, where $p \geq 1$ and d is some distance function. If we define the latent Wasserstein discrepancy as $\mathcal{W}_p^L := (\text{OT}^L)^{1/p}$, then there exist $\kappa > 0$ such that, for any μ, ν and ζ having latent distributions of support sizes up to k , the discrepancy satisfies,

- $\mathcal{W}_p^L(\mu, \nu) \geq 0$
- $\mathcal{W}_p^L(\mu, \nu) = \mathcal{W}_p^L(\nu, \mu)$
- $\mathcal{W}_p^L(\mu, \nu) \leq \kappa (\mathcal{W}_p^L(\mu, \zeta) + \mathcal{W}_p^L(\zeta, \nu))$

The low-rank nature of LOT has a biasing effect that results in $\mathcal{W}_p^L(\mu, \mu) > 0$ for a general μ . We can debias it by defining its variant $\tilde{\mathcal{W}}_p^L(\mu, \nu) := \left((\mathcal{W}_p^L(\mu, \nu))^p - \min_{z_k \in \Phi_x} \mathcal{W}_p^p(\mu, z_k) - \min_{z'_k \in \Phi_y} \mathcal{W}_p^p(\nu, z'_k) \right)^{1/p}$, where $\Phi_x = \Delta_{\mathcal{Z}_x}^{k_x}$, $\Phi_y = \Delta_{\mathcal{Z}_y}^{k_y}$. The following property connects $\tilde{\mathcal{W}}_p^L(\mu, \nu)$ to k-means clustering.

Corollary 1. Under the assumptions of Proposition 1, if $p = 2$ and $k_x = k_y = k$, then $\forall \mu, \nu$, we have $\tilde{\mathcal{W}}_2^L(\mu, \nu) \geq 0$. Furthermore, $\tilde{\mathcal{W}}_2^L(\mu, \nu) > 0$ if their k-means centroids or sizes of their k-means clusters differ.

3.3. Establishing a ground metric

In what follows, we will focus on the Euclidean space $\mathcal{X} = \mathcal{Y} = \mathbb{R}^d$. Instead of considering every source-to-target distance to build our transportation cost, we can use anchors as proxies for each point. A well-established way of encoding the distance that each point needs to travel to get to its nearest anchor, is to define the cost as:

$$\mathbf{C}_x = d_{\mathbf{M}_x}, \mathbf{C}_z = d_{\mathbf{M}_z}, \mathbf{C}_y = d_{\mathbf{M}_y}, \quad (6)$$

where $d_{\mathbf{M}}$ denotes the Mahalanobis distance: $d_{\mathbf{M}}^2(\mathbf{x}, \mathbf{y}) := (\mathbf{x} - \mathbf{y})^T \mathbf{M} (\mathbf{x} - \mathbf{y})$ and \mathbf{M} is some positive semidefinite matrix. The Mahalanobis distance generalizes the squared Euclidean distance and allows us to consider different costs based on correlations between features. The framework of Mahalanobis distance benefits from efficient metric learning techniques (Cuturi & Avis, 2014); recent research also establishes connections between it and robust OT (Paty & Cuturi, 2019; Dhoubib et al., 2020). When a simple L2-distance

is used ($\mathbf{M} = \mathbf{I}$), we will denote this specific variant as LOT-L2.

When LOT moves source points through anchors, the anchors impose a type of bottleneck, and this results in a loss of information that makes it difficult to estimate the corresponding point in the target space. In cases where accurate point-to-point alignment is desired, we propose an alternative strategy for defining the cost matrix \mathbf{C}_z . The idea is to represent an anchor as the distribution of points assigned to it. Specifically, we represent $\mathbf{z}^x, \mathbf{z}^y$ as measures in \mathbb{R}^d : $\tilde{\mathbf{z}}^x = \sum_{i=1}^N \mathbf{P}_x(\mathbf{x}_i | \mathbf{z}^x) \delta_{\mathbf{x}_i}$, $\tilde{\mathbf{z}}^y = \sum_{j=1}^M \mathbf{P}_y(\mathbf{y}_j | \mathbf{z}^y) \delta_{\mathbf{y}_j}$. Then we measure the cost between anchors as the squared Wasserstein distance between their respective distributions,

$$\mathbf{C}_z := [\mathcal{W}_2^2(\mathbf{P}_x(\cdot | \mathbf{z}_m^x), \mathbf{P}_y(\cdot | \mathbf{z}_n^y))]_{m,n}. \quad (7)$$

Besides the quantity itself, the transport plan returned by calculating \mathbf{C}_z is also very important as it provides accurate point-to-point maps. Since the cost matrix is now a function of \mathbf{P}_x and \mathbf{P}_y , we use an additional alternating scheme to solve the problem: we alternate between updating \mathbf{C}_z while keeping \mathbf{P}_x and \mathbf{P}_y fixed, and then updating $\mathbf{P}_x, \mathbf{P}_y, \mathbf{P}_z$ while keeping \mathbf{C}_z fixed. An efficient algorithm is presented in Appendix B.3 to reduce the computation complexity. This variant, LOT-WA, can yield better performance in downstream tasks that require precise alignment at the cost of additional computation.

3.4. Algorithm

In the rest of this section, we will develop our main approach for solving the problem in Def. 1. We provide an outline of the algorithm in Algorithm 1 and an implementation of the algorithm in Python at: <http://nerdslab.github.io/latentOT>.

(1) Optimizing $\mathbf{P}_x, \mathbf{P}_y$ and \mathbf{P}_z : To begin, we assume that the anchors and cost matrices $\mathbf{C}_x, \mathbf{C}_z, \mathbf{C}_y$ are already specified. Let $\mathbf{K}_x, \mathbf{K}_z, \mathbf{K}_y$ be the Gibbs kernels induced from the cost matrices $\mathbf{C}_x, \mathbf{C}_z, \mathbf{C}_y$ as in (2). The optimization problem can be written as,

$$\begin{aligned} \min_{\mathbf{u}_z, \mathbf{v}_z, \mathbf{P}_x, \mathbf{P}_z, \mathbf{P}_y} \quad & \sum_{i \in \{x, y, z\}} \varepsilon_i \text{KL}(\mathbf{P}_i \| \mathbf{K}_i), \\ \text{subject to: } \quad & \mathbf{P}_x \mathbf{1} = \mu, \mathbf{P}_x^T \mathbf{1} = \mathbf{u}_z, \mathbf{P}_z \mathbf{1} = \mathbf{u}_z, \\ & \mathbf{P}_z^T \mathbf{1} = \mathbf{v}_z, \mathbf{P}_y \mathbf{1} = \mathbf{v}_z, \mathbf{P}_y^T \mathbf{1} = \nu. \end{aligned} \quad (8)$$

This is a Bregman projection problem with affine constraints. An iterative projection procedure can thus be applied to solve the problem (Benamou et al., 2015). We present the procedure as **UPDATEPLAN** in Algorithm 1, where $\mathbf{P}_x, \mathbf{P}_z, \mathbf{P}_y$ are successively projected onto the constrained sets of fixed marginal distributions. We defer the detailed derivation to Appendix B.1.

Algorithm 1 Latent Optimal Transport - LOT

Input: Data matrices \mathbf{X}, \mathbf{Y} ; metric costs $\mathbf{M}_x, \mathbf{M}_y, \mathbf{M}_z$; entropy regularization parameters $\varepsilon_x, \varepsilon_y, \varepsilon_z$; initial anchors $\mathbf{Z}_x, \mathbf{Z}_y$.

Output: Transport plans $\mathbf{P}_x, \mathbf{P}_y, \mathbf{P}_z$; source and target anchors $\mathbf{Z}_x, \mathbf{Z}_y$.

Initialize: $\mathbf{P}_x, \mathbf{P}_y, \mathbf{P}_z, \mathbf{Z}_x, \mathbf{Z}_y$

- 1: **while** not converging **do**
- 2: $(\text{vec}(\mathbf{Z}_x), \text{vec}(\mathbf{Z}_y)) \leftarrow \text{Eqn. (9)}$
- 3: $\mathbf{K}_x = \{\exp(-\|\mathbf{X}[i] - \mathbf{Z}_x[j]\|_{\mathbf{M}_x}^2 / \varepsilon_x)\}_{i,j}$
- 4: $\mathbf{K}_y = \{\exp(-\|\mathbf{Y}[i] - \mathbf{Z}_y[j]\|_{\mathbf{M}_y}^2 / \varepsilon_y)\}_{j,q}$
- 5: $\mathbf{K}_z = \{\exp(-\|\mathbf{Z}_x[i] - \mathbf{Z}_y[j]\|_{\mathbf{M}_z}^2 / \varepsilon_z)\}_{p,q}$
- 6: $\mathbf{P}_x, \mathbf{P}_y, \mathbf{P}_z \leftarrow \text{UPDATEPLAN}(\mathbf{K}_x, \mathbf{K}_y, \mathbf{K}_z)$
- 7: **end while**

Return: $\mathbf{P}_x, \mathbf{P}_y, \mathbf{P}_z, \mathbf{Z}_x, \mathbf{Z}_y$

UPDATEPLAN ($\mathbf{K}_x, \mathbf{K}_y, \mathbf{K}_z$)

Initialize: $\alpha_x \leftarrow \mathbf{1}_N; \beta_x \leftarrow \mathbf{1}_{k_1}; \alpha_y \leftarrow \mathbf{1}_{k_2};$
 $\beta_y \leftarrow \mathbf{1}_M; \alpha_z \leftarrow \mathbf{1}_{k_1}; \beta_z \leftarrow \mathbf{1}_{k_2}$

- 1: **while** not converging **do**
- 2: $\alpha_x \leftarrow \mu \odot \mathbf{K}_x \beta_x; \beta_y \leftarrow \nu \odot \mathbf{K}_y^T \alpha_y$
- 3: $\mathbf{u}_z \leftarrow ((\alpha_z \odot \mathbf{K}_z \beta_z) \odot (\beta_x \odot \mathbf{K}_x^T \alpha_x))^{\frac{1}{2}}$
- 4: $\beta_x \leftarrow \mathbf{u}_z \odot \mathbf{K}_x^T \alpha_x; \alpha_z \leftarrow \mathbf{u}_z \odot \mathbf{K}_z \beta_z$
- 5: $\mathbf{v}_z \leftarrow ((\alpha_y \odot \mathbf{K}_y \beta_y) \odot (\beta_z \odot \mathbf{K}_z^T \alpha_z))^{\frac{1}{2}}$
- 6: $\beta_z \leftarrow \mathbf{v}_z \odot \mathbf{K}_z^T \alpha_z; \alpha_y \leftarrow \mathbf{v}_z \odot \mathbf{K}_y \beta_y$
- 7: **end while**

Return: $\mathbf{P}_i = \text{diag}(\alpha_i) \mathbf{K}_i \text{diag}(\beta_i), i \in \{x, y, z\}$

(2) Optimizing the anchor locations: Now we consider the case where we are free to select the anchor locations in \mathbb{R}^d . We consider the class of Mahalanobis costs described in Section 3.3. Let $\mathbf{M}_x, \mathbf{M}_z, \mathbf{M}_y$ be the Mahalanobis matrices correspond to $\mathbf{C}_x, \mathbf{C}_z$, and \mathbf{C}_y , respectively.

Given the transport plans generated after solving (8), we can derive the the first-order stationary condition of OT^L with respect to \mathbf{Z}_x and \mathbf{Z}_y . Let

$$\mathbf{A} = \begin{bmatrix} D(\mathbf{u}_z) \otimes (\mathbf{M}_x + \mathbf{M}_z) & \mathbf{P}_z \otimes \mathbf{M}_z \\ -\mathbf{P}_z^T \otimes \mathbf{M}_z & D(\mathbf{v}_z) \otimes (\mathbf{M}_y + \mathbf{M}_z) \end{bmatrix}$$

The update formula is given by

$$\begin{bmatrix} \text{vec}(\mathbf{Z}_x^*) \\ \text{vec}(\mathbf{Z}_y^*) \end{bmatrix} = \mathbf{A}^{-1} \times \begin{bmatrix} (\mathbf{P}_x^T \otimes \mathbf{M}_x) \text{vec}(\mathbf{X}) \\ (\mathbf{P}_y \otimes \mathbf{M}_y) \text{vec}(\mathbf{Y}) \end{bmatrix}, \quad (9)$$

where $\text{vec}(\cdot)$ denotes the operator converting a matrix to a column vector, and $D(\cdot)$ denotes the operator converting a vector to a diagonal matrix. We defer the detailed derivation to Appendix B.2. Pseudo-code for the combined scheme can be found in Algorithm 1.

(3) Robust estimation of data transport: LOT provides robust transport in the target domain by aligning the

data through anchors, which can facilitate regression, and classification in downstream applications. We denote the centroids of the source and target by $\mathbf{Q}_x = \text{diag}(\mathbf{u}_z^{-1}) \mathbf{P}_x^T \mathbf{X}^T$, $\mathbf{Q}_y = \text{diag}(\mathbf{v}_z^{-1}) \mathbf{P}_y \mathbf{Y}^T$. We propose the estimator $\hat{\mathbf{X}} := \sum_{m,n} p(\mathbf{z}_m^x, \mathbf{z}_n^y | \mathbf{x}) (\mathbf{Q}_m^y - \mathbf{Q}_n^x) = \text{diag}(\mu^{-1}) \mathbf{P}_x \text{diag}((\mathbf{P}_z \mathbf{1})^{-1}) \mathbf{P}_z (\mathbf{Q}_y - \mathbf{Q}_x)$. In contrast to factored coupling (Forrow et al., 2019), where $\mathbf{Z}_x = \mathbf{Z}_y$, LOT is robust even when the source and target have different structures (see Table 1 MNIST-DU, Figure 4).

(4) Implementation details: LOT has two primary hyperparameters that must be specified: (i) the number of the source and target anchors k_x, k_y and (ii) the regularization parameter ε . For details on the tuning of these parameters, please refer to Appendix F. In practice, we use centroids from k-means clustering (Arthur & Vassilvitskii, 2006) to initialize the anchors, and for all the experiments we have conducted, LOT typically converges within 20 iterations.

4. Related Work

Interpolation between factored coupling and k-means clustering: Assume we select the Mahalanobis matrices of the costs defined in Section 3.3 to be $\mathbf{M}_x = \mathbf{M}_y = \mathbf{I}$, and $\mathbf{M}_z = \lambda \mathbf{I}$. If we let $\lambda \rightarrow \infty$ when estimating the transport between source and target anchors, the anchors merge, and our approach reduces to the case of factored coupling (Forrow et al., 2019). At the other end, if we let $\lambda \rightarrow 0$, then LOT becomes separable, and the middle term vanishes. In this case, each remaining term exactly corresponds to a pure clustering task, and LOT reduces to k-means clustering (Arthur & Vassilvitskii, 2006).

Relationship to OT-based clustering methods: Many methods that combine OT and clustering (Li & Wang, 2008; Ye et al., 2017; Ho et al., 2017; Dessein et al., 2017; Genevay et al., 2019; Alvarez-Melis & Fusi, 2020) focus on using the Wasserstein distance to identify barycenters that serve as the centroids of clusters. When finding barycenters for the source and target separately, this could be seen as LOT with $\mathbf{C}_z = 0$ and $\mathbf{C}_x, \mathbf{C}_y$ defined using a squared L2 distance. In other related work (Laclau et al., 2017), co-clustering is applied to a transport plan as a post-processing operation, and no additional regularization on the transportation cost in OT is imposed. In contrast, our approach induces explicit regularization by separately defining cost matrices for the transport between the source/target points and their anchors. This yields a transport plan guided by a cluster-level matching.

Relationship to hierarchical OT : Hierarchical OT (Chen et al., 2018; Lee et al., 2019; Yurochkin et al., 2019; Xu et al., 2020) transports points by moving them within some predetermined subgroup simultaneously based on either their class label or pre-specified structures, and then forms a

matching of these subgroups using the Wasserstein distance. The resulting problem solves a multi-layer OT problem which gives rise to its name. With a Wasserstein distance used to build the \mathbf{C}_z cost matrix, LOT effectively reduces to hierarchical OT for fixed and hard-class assignment \mathbf{P}_x and \mathbf{P}_y . However, a crucial difference between LOT and hierarchical OT lies in that the latter imposes the *known* structure information. In contrast, LOT discovers this structure by simultaneously learning \mathbf{P}_x and \mathbf{P}_y .

Transportation with anchors: The notion of moving data points with anchors to match points in heterogeneous spaces has appeared in other work (Sato et al., 2020; Manay et al., 2006). These approaches map each point from one domain into a distribution of the costs, which effectively builds up a common representation for the points from both spaces. In contrast to this work, we use the anchors to encourage clustering of data and to impose rank constraints on the transport plan.

5. Theoretical Analysis

LOT as a relaxation of OT: We now ask how the optimal value of our original rank-constrained objective in (8) is related to the transportation cost defined in entropy-regularized OT. It turns out their objectives are connected by an inequality described below (see Appendix A for a proof).

Proposition 2. *Let \mathbf{P} be a transport plan of the form in (5). Assume that \mathbf{K} is some Gibbs kernel that satisfies,*

$$\mathbf{K}_x \mathbf{K}_z \mathbf{K}_y \leq \mathbf{K}, \quad (10)$$

where the inequality is over each entry. Then we have,

$$\begin{aligned} \varepsilon KL(\mathbf{P} \parallel \mathbf{K}) &\leq \varepsilon (KL(\mathbf{P}_x \parallel \mathbf{K}_x) + KL(\mathbf{P}_z \parallel \mathbf{K}_z) \\ &\quad + KL(\mathbf{P}_y \parallel \mathbf{K}_y)) + \varepsilon (\mathbf{H}(\mathbf{u}_z) + \mathbf{H}(\mathbf{v}_z)), \end{aligned} \quad (11)$$

where $\mathbf{H}(\mathbf{a}) := -\sum_i \mathbf{a}_i \log \mathbf{a}_i$ denotes the entropy.

The proposition shows that an OT objective, corresponding to a kernel \mathbf{K} (resp. \mathbf{C}), can be upper bounded by three sub-OT problems defined by subsequent kernels $\mathbf{K}_x, \mathbf{K}_z, \mathbf{K}_y$ (resp. $\mathbf{C}_x, \mathbf{C}_z, \mathbf{C}_y$) that satisfies (10) (resp. $\exp(-\mathbf{C}_x/\varepsilon) \exp(-\mathbf{C}_z/\varepsilon) \exp(-\mathbf{C}_y/\varepsilon) \leq \exp(-\mathbf{C}/\varepsilon)$).

Let us compare the upper bound given by Proposition 2 with Def. 1 and ignore the entropy terms; we recognize that it is precisely the entropy-regularized objective of LOT. In other words, with suitable cost matrices satisfying (10), LOT could be interpreted as a relaxation of an OT problem in a decomposed form. We then ask what $\mathbf{C}_x, \mathbf{C}_z, \mathbf{C}_y$ should be to satisfy (10). In cases where cost \mathbf{C} is defined by the L^p -norm to the power p , the following corollary shows that the same form suffices.

Corollary 2. *Let $d(\mathbf{x}, \mathbf{y}) := \|\mathbf{x} - \mathbf{y}\|_p^p$. Consider an optimal transport problem $\text{OT}_{\mathbf{C}, \varepsilon}$ with cost $\mathbf{C}[i, j] = d(\mathbf{x}_i, \mathbf{y}_j)$,*

where $p \geq 1$. Then for a sufficiently small ε , the latent optimal transport OT^L with cost matrices, $\mathbf{C}_x[i, m] = 3^{p-1}d(\mathbf{x}_i - \mathbf{z}_m^x)$, $\mathbf{C}_z[m, n] = 3^{p-1}d(\mathbf{z}_m^x - \mathbf{z}_n^y)$, $\mathbf{C}_y[n, j] = 3^{p-1}d(\mathbf{z}_n^y - \mathbf{y}_j)$ minimizes an upper bound of the entropy-regularized OT objective in (2).

Corollary (2) provides natural costs for LOT to be posed as a relaxation to a OT problem with L^p norm. More generally, finding the optimal cost functions that obey (10) and minimize the gap in the inequality in Proposition 2 is outside the scope of this work but would be an interesting topic for future investigation.

Sampling complexity: Below we analyze LOT from a statistical point of view. Specifically, we bound the sampling rate of OT^L in Def. 1 when the true distributions μ and ν are estimated by their empirical distributions.

Proposition 3. *Suppose X and Y have distributions μ and ν supported on a compact region Ω in \mathbb{R}^d , the cost functions $c_x(\cdot, \cdot)$ and $c_y(\cdot, \cdot)$ are defined as the squared Euclidean distance, and $\hat{\mu}, \hat{\nu}$ are empirical distributions of n and m i.i.d. samples from μ and ν , respectively. If the spaces for latent distributions are equal to $\mathcal{Z}_x = \mathcal{Z}_y = \mathbb{R}^d$, and there are k_x and k_y anchors in the source and target, respectively, then with probability at least $1 - \delta$,*

$$\text{Err} \leq C \sqrt{\frac{k_{\max}^3 d \log k_{\max} + \log(2/\delta)}{N}}, \quad (12)$$

where $\text{Err} = |\text{OT}^L(\mu, \nu) - \text{OT}^L(\hat{\mu}, \hat{\nu})|$, $k_{\max} = \max\{k_x, k_y\}$, $N = \min\{n, m\}$ and $C \geq 0$ is some constant not depending on N .

As shown in (Weed et al., 2019), the general sampling rate of a plug-in OT scales with $N^{\frac{1}{d}}$, suffering from the ‘‘curse of dimensionality’’. On the other hand, as evidence from (Forrow et al., 2019), structural optimal transport paves ways to overcome the issue. In particular, LOT achieves $N^{-\frac{1}{2}}$ scaling by regularizing the transport rank.

Time complexity: We can bound the time complexity as $O(T_i + T_{\text{bcd}}(T_k + T_{\text{au}} + T_{\text{pu}}))$, where T_i is the initialization complexity, e.g., if we use k-means, then it equals to $O(nk_x dT_x + mk_y dT_y)$ where T_x and T_y are the iteration numbers of the Floyd algorithm applied to the source and target, respectively, T_{bcd} is the total number of iterations of block-coordinate descent, $T_k = O(nk_x + mk_y)$ is the computation time for updating the kernels, $T_{\text{au}} = O((k_x + k_y)^3 + d(nk_x + mk_y))$ is the complexity of updating anchors, and T_{pu} is the complexity for updating plans. Because our updates are based on iterative Bregman projections similar to the Sinkhorn algorithm, it has complexity comparable to OT. Therefore, the overall complexity of LOT is approximately T_{bcd} times of OT, assuming $n, m \geq d(k_x + k_y)$. Empirically, T_{bcd} depends on the structure of data, but we observed that it is usually under

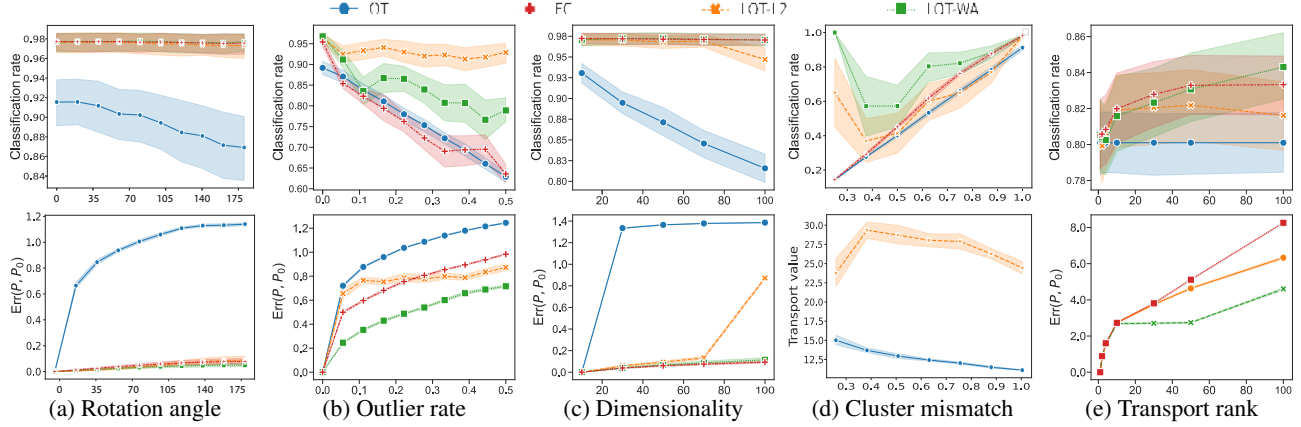


Figure 2: **Results on Gaussian mixture models.** In (a), we apply a rotation between the source and target, in (b) we add outliers, in (c) we vary the ambient dimension, in (d) the target is set to have 8 components, and we vary the number of components in the source to simulate source-target mismatch, in (e) we fix the rank to 10 and vary the number of factors (anchors) used in the approximation. Throughout, we simulate data according to a GMM and evaluate performance by measuring the classification accuracy (top) and computing the deviation between the transport plans before and after the perturbations with respect to the Fröbenius norm (bottom).

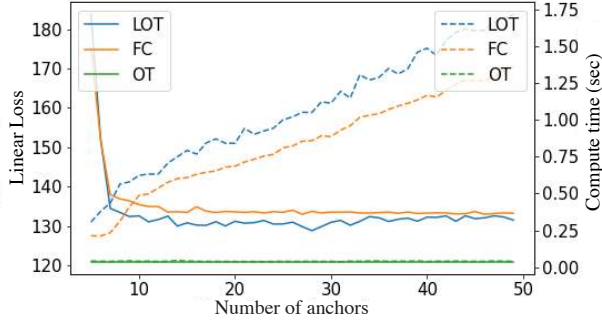


Figure 3: **Comparisons of the time complexity and loss.** The figure compares the time complexity (dashed) and linear loss (solid) of LOT, FC, and OT in the setting of the 7-component GMM model.

20. Note that the same applies to FC with $k_x = k_y = k$. In Figure 3, we complement our analysis by simulating a comparison of the time complexity for LOT and FC vs. OT in the setting of a 7-component Gaussian mixture model. We can see the compute time of LOT scales similarly to FC.

Transport cost: We also compare the transport loss returned by LOT (blue), FC (orange), and OT (green) as a function of the number of anchors in Figure 3. For a fair comparison, we considered a balanced scenario where 7-component GMMs generate the source and target. The anchors of the source and target are chosen to be equal for LOT. The result shows that the losses are indeed higher for LOT and FC compared to OT but are fairly insensitive with to the chosen number of anchors. Moreover, we find that LOT has a slightly lower loss compared to FC even when we choose the number of source and target anchors to be equal.

6. Experiments

In this section, we conduct empirical investigations. Details of hyperparameter tuning can be found in Appendix F.

E1) Testing robustness to various data perturbations: To better understand how different types of domain shift impact the transport plans generated by our approach, we considered different transformations between the source and target. To create synthetic data for this task, we generated multiple clusters/components using a k -dimensional Gaussian with random mean and covariance sampled from a Wishart distribution, randomly projected to a 5-dimensional subspace. The source and target are generated independently: we randomly sample a fixed number of points according to the true distribution for each cluster. We compared the performance of the LOT variants proposed in Section 3.3: LOT-L2 (orange curves) and LOT-WA (green curves) with baselines OT (blue curves) and rank regularized factored coupling (FC) (Graf & Luschgy, 2007) (red curves) in terms of their (i) classification rates and (ii) deviation from the original transport plan without perturbations, which we compute as $\text{Err}(\mathbf{P} - \mathbf{P}_0) = \|\mathbf{P} - \mathbf{P}_0\|_F / \|\mathbf{P}_0\|_F$, where \mathbf{P}_0 is the transport plan obtained before perturbations. The results are averaged over 20 runs, and a 75% confidence interval is used. See Appendix E for further details.

When compared with OT, both our method and FC provide more stable class recovery, even with significant amounts of perturbations (Figure 2). When we examine the error term in the transport plan, we observe that, in most cases, the OT plan deviates rapidly, even for small amounts of perturbations. Both FC and LOT appear to have similar performances across rotations while OT’s performance decreases quickly. In experiment (b), we found that both LOT variants provide

substantial improvements on classification subject to outliers, implying the applicability of LOT for noisy data. In experiment (c), we study LOT in the high-dimensional setting; we find that LOT-WA behaves similarly to FC with some degradation in performance after the dimension increases beyond 70. Next, in experiment (d), we fix the number of components in the target to be 10, while varying the number in the source from 4 to 10. In contrast to the outlier experiment in (b), LOT-WA shows more resilience to mismatches between the source and target. At the bottom of plot (d), we show the 2-Wasserstein distance (blue) and latent Wasserstein discrepancy (orange) defined in Proposition 1. This shows that the latent Wasserstein discrepancy does indeed provide an upper bound on the 2-Wasserstein distance. Finally, we look at the effect of transport rank on LOT and FC in (e). The plot shows that the slope for LOT is flatter than FC while maintaining similar performances.

E2) Domain adaptation application: In our next experiment, we used LOT to correct for domain shift in a neural network that is trained on one dataset but underperforms on a new but similar dataset (Table 1, Figure 4). MNIST and USPS are two handwritten digits datasets that are semantically similar but that have different pixel-level distributions and thus introduce domain shift (Figure 4a). We train a multi-layer perceptron (MLP) on the training set of the MNIST dataset, freeze the network, and use it for the remaining experiments. The classifier achieves 100% training accuracy and a 98% validation accuracy on MNIST but only achieves 79.3% accuracy on the USPS validation set. We project MNIST’s training samples in the classifier’s output space (logits) and consider the 10D projection to be the target distribution. Similarly, we project images from the USPS dataset in the network’s output space to get our source distribution. We study the performance of LOT in correcting the classifier’s outputs and compare with FC, k-means OT (KOT) (Forrow et al., 2019), and subspace alignment (SA) (Fernando et al., 2013).

In Table 1, we summarize the results of our comparisons on the domain adaptation task (MNIST-USPS). Our results suggest that both FC and LOT perform pretty well on this task, with LOT beating FC by 2% in terms of their final classification accuracy. We also show that LOT does better than naive KOT. In Figure 4a, we use Isomap to project the distribution of USPS images as well as the alignment results for LOT, FC, and OT. For both LOT and FC, we also display the anchors; note that for LOT, we have two different sets of anchors (source, red; target, blue). This example highlights the alignment of the anchors in our approach and contrasts it with that of FC.

Taking inspiration from studies in self-supervised learning (Doersch et al., 2015; He et al., 2020) that use different transformations of an input image (e.g., masking parts of the

Table 1: *Results for concept drift and domain adaptation for handwritten digits.* The classification accuracy and L2-error are computed after transport for MNIST to USPS (left) and coarse dropout (right). Our method is compared with the accuracy before alignment (Original), entropy-regularized OT, k-means plus OT (KOT), and subspace alignment (SA).

	MNIST-USPS	MNIST-DU	
	Accuracy	Accuracy	L2 error
Original	79.3	72.6	0.72
OT	76.9	61.5	0.71
KOT	79.4	60.9	0.73
SA	81.3	72.3	-
FC	84.1	67.2	0.59
LOT-WA	86.2	77.7	0.56

image) to build invariances into a network’s representations, here we ask how augmentations of the images introduce domain shift and whether our approach can correct/identify it. To test this, we apply coarse dropout on test samples in MNIST and feed them to the classifier to get a new source distribution. We do this in a balanced (all digits in source and target) and an unbalanced setting (2, 4, 8 removed from source, all digits in target). The results of the unbalanced dropout are summarized in Table 1 (MNIST-DU), and the other results are provided in Table S1 in the Appendix. In this case, we have the features of the testing samples pre-transformation, and thus, we can compare the transported features to the ground truth features in terms of their point-to-point error (L2 distance). In the unbalanced case, we observe even more significant gaps between FC and LOT, as the source and target datasets have different structures. To quantify these different class-level errors, we compare the confusion matrices for the classifier’s output after alignment (Figure 4b). By examining the columns corresponding to the removed digits, we see that FC is more likely to misclassify these images. Our results suggest that LOT has comparable performance with FC in a balanced setting and outperforms FC in an unbalanced case.

The decomposition in both LOT and FC allows us to visualize transport between the source, anchors, and the target (Figure 4c-d, S2). This visualization highlights the interpretability of the transport plans learned via our approach, with the middle transport plan P_z providing a concise map of interactions between class manifolds in the unbalanced setting. With LOT (Figure 4c), we find that each source anchor is mapped to the correct target anchor, with some minor interactions with the target anchors corresponding to the removed digits. In comparison, FC (Figure 4d, S2) has more spurious interactions between source, anchors, and target.

E3) Robustness to sampling: We examined the robustness to sampling for the MNIST to USPS example (Figure 5).

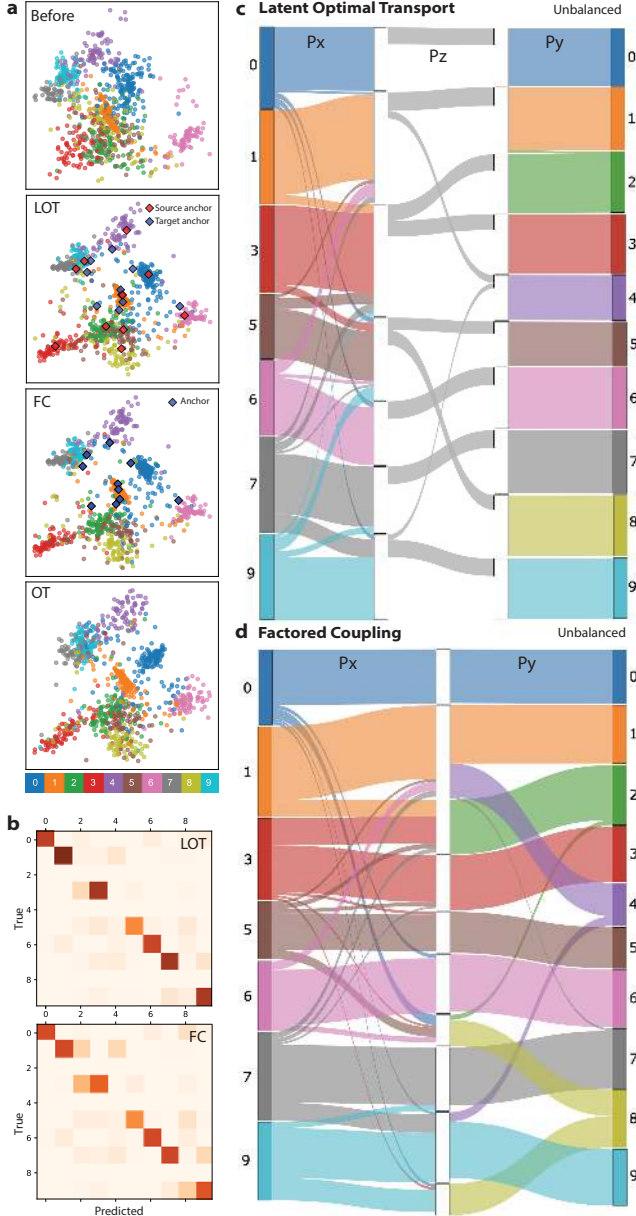


Figure 4: *Visualization of results on handwritten digits and examples of domain shift.* (a) 2D projections of representations formed in deep neural network before (top) and after different alignment methods (LOT, FC, OT). (b) confusion matrices for LOT (top) and FC (bottom) after alignment. The transport plans are visualized for LOT (c) and FC (d) in the unbalanced case.

In this case, we find that LOT has a stable alignment as we subsample the source dataset, with very little degradation in classification accuracy, even as we reduce the source to only 20 samples. We also observe a significant gap between LOT and other approaches in this experiment, with more than a 10% gap between FC and LOT when very few source samples are provided. Our results demonstrate that LOT is




Figure 5: *LOT provides robust alignment, even when given very few samples.* We compare our method with OT and FC on the MNIST-USPS domain adaptation task when different numbers of USPS samples are available. Reported classification rates are averaged over 50 random sets.

robust to subsampling, providing empirical evidence for Proposition 3.

7. Discussion

In this paper, we introduced LOT, a new form of structured transport leading to an approach for jointly clustering and aligning data. We provided an efficient optimization method to solve for the transport, and studied its statistical rate via theoretical analysis and robustness to data perturbations with empirical experiments. In the future, we would like to explore the application of LOT to non-Euclidean spaces, and incorporate metric learning into our framework.

Acknowledgements

This project is supported by the NIH award 1R01EB029852-01, NSF awards IIS-1755871 and CCF-1740776, as well as generous gifts from the Alfred Sloan Foundation and the McKnight Foundation.

References

- Agueh, M. and Carlier, G. Barycenters in the wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2): 904–924, 2011.
- Altschuler, J., Niles-Weed, J., and Rigollet, P. Near-linear time approximation algorithms for optimal transport via sinkhorn iteration. In *Advances in neural information processing systems*, pp. 1964–1974, 2017.
- Altschuler, J., Bach, F., Rudi, A., and Niles-Weed, J. Massively scalable sinkhorn distances via the nyström method. In *Advances in Neural Information Processing Systems*, pp. 4427–4437, 2019.
- Alvarez-Melis, D. and Fusi, N. Geometric dataset distances

- via optimal transport. *arXiv preprint arXiv:2002.02923*, 2020.
- Alvarez-Melis, D., Jaakkola, T., and Jegelka, S. Structured optimal transport. In *Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics*, pp. 1771–1780, 2018.
- Alvarez-Melis, D., Jegelka, S., and Jaakkola, T. S. Towards optimal transport with global invariances. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, pp. 1870–1879. PMLR, 2019.
- Arthur, D. and Vassilvitskii, S. k-means++: The advantages of careful seeding. Technical report, Stanford, 2006.
- Benamou, J.-D., Carlier, G., Cuturi, M., Nenna, L., and Peyré, G. Iterative bregman projections for regularized transportation problems. *SIAM Journal on Scientific Computing*, 37(2):A1111–A1138, 2015.
- Bonneel, N., Peyré, G., and Cuturi, M. Wasserstein barycentric coordinates: histogram regression using optimal transport. *ACM Trans. Graph.*, 35(4):71–1, 2016.
- Caliński, T. and Harabasz, J. A dendrite method for cluster analysis. *Communications in Statistics*, 3(1):1–27, 1974. doi: 10.1080/03610927408827101. URL <https://www.tandfonline.com/doi/abs/10.1080/03610927408827101>.
- Chen, Y., Georgiou, T. T., and Tannenbaum, A. Optimal transport for gaussian mixture models. *IEEE Access*, 7: 6269–6278, 2018.
- Courty, N., Flamary, R., and Tuia, D. Domain adaptation with regularized optimal transport.
- Courty, N., Flamary, R., Tuia, D., and Rakotomamonjy, A. Optimal transport for domain adaptation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 39(9):1853–1865, 2017.
- Cuturi, M. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pp. 2292–2300, 2013.
- Cuturi, M. and Avis, D. Ground metric learning. *J. Mach. Learn. Res.*, 15(1):533–564, 2014.
- Cuturi, M. and Doucet, A. Fast computation of wasserstein barycenters. pp. 685–693, 2014.
- Cuturi, M. and Peyré, G. A smoothed dual approach for variational wasserstein problems. *SIAM Journal on Imaging Sciences*, 9(1):320–343, 2016.
- Dessein, A., Papadakis, N., and Deledalle, C.-A. Parameter estimation in finite mixture models by regularized optimal transport: A unified framework for hard and soft clustering. *arXiv preprint arXiv:1711.04366*, 2017.
- Dhouib, S., Redko, I., Kerdoncuff, T., Emonet, R., and Sebban, M. A swiss army knife for minimax optimal transport. In *Proceedings of the 37th International Conference on Machine Learning*, 2020.
- Doersch, C., Gupta, A., and Efros, A. A. Unsupervised visual representation learning by context prediction. In *IEEE International Conference on Computer Vision*, pp. 1422–1430, 2015.
- Dykstra, R. L. An algorithm for restricted least squares regression. *Journal of the American Statistical Association*, 78(384):837–842, 1983.
- Fernando, B., Habrard, A., Sebban, M., and Tuytelaars, T. Unsupervised visual domain adaptation using subspace alignment. In *IEEE International Conference on Computer Vision*, pp. 2960–2967, 2013.
- Flamary, R. and Courty, N. POT Python optimal transport library, 2017. URL <https://pythonot.github.io/>.
- Forrow, A., Hütter, J.-C., Nitzan, M., Rigollet, P., Schiebinger, G., and Weed, J. Statistical optimal transport via factored couplings. In *Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics*, pp. 2454–2465, 2019.
- Genevay, A., Dulac-Arnold, G., and Vert, J.-P. Differentiable deep clustering with cluster size constraints. *arXiv preprint arXiv:1910.09036*, 2019.
- Graf, S. and Luschgy, H. *Foundations of quantization for probability distributions*. Springer, 2007.
- Gramfort, A., Peyré, G., and Cuturi, M. Fast optimal transport averaging of neuroimaging data. In *Information Processing in Medical Imaging*, pp. 261–272. Springer, 2015.
- He, K., Fan, H., Wu, Y., Xie, S., and Girshick, R. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9729–9738, 2020.
- Ho, N., Nguyen, X., Yurochkin, M., Bui, H. H., Huynh, V., and Phung, D. Multilevel clustering via wasserstein means. *arXiv preprint arXiv:1706.03883*, 2017.
- Kusner, M., Sun, Y., Kolkin, N., and Weinberger, K. From word embeddings to document distances. In *Proceedings of the Thirty-Forth International Conference on Artificial Intelligence and Statistics*, pp. 957–966, 2015.

- Laclau, C., Redko, I., Matei, B., Bennani, Y., and Brault, V. Co-clustering through optimal transport. *arXiv preprint arXiv:1705.06189*, 2017.
- Lee, J., Dabagia, M., Dyer, E., and Rozell, C. Hierarchical optimal transport for multimodal distribution alignment. In *Advances in Neural Information Processing Systems*, pp. 13474–13484, 2019.
- Li, J. and Wang, J. Z. Real-time computerized annotation of pictures. *IEEE transactions on pattern analysis and machine intelligence*, 30(6):985–1002, 2008.
- Liero, M., Mielke, A., and Savaré, G. Optimal entropy-transport problems and a new hellinger–kantorovich distance between positive measures. *Inventiones mathematicae*, 211(3):969–1117, 2018.
- Manay, S., Cremers, D., Hong, B.-W., Yezzi, A. J., and Soatto, S. Integral invariants for shape matching. *IEEE Transactions on pattern analysis and machine intelligence*, 28(10):1602–1618, 2006.
- Martin Arjovsky, S. and Bottou, L. Wasserstein generative adversarial networks. In *Proceedings of the 34th International Conference on Machine Learning*, 2017.
- Paty, F.-P. and Cuturi, M. Subspace robust wasserstein distances. 2019.
- Pele, O. and Werman, M. Fast and robust earth mover’s distances. In *2009 IEEE 12th International Conference on Computer Vision*, pp. 460–467. IEEE, 2009.
- Peyré, G., Cuturi, M., et al. Computational optimal transport: With applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Rousseeuw, P. J. Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. *Journal of Computational and Applied Mathematics*, 20:53 – 65, 1987. ISSN 0377-0427. doi: [https://doi.org/10.1016/0377-0427\(87\)90125-7](https://doi.org/10.1016/0377-0427(87)90125-7). URL <http://www.sciencedirect.com/science/article/pii/0377042787901257>.
- Santambrogio, F. Optimal transport for applied mathematicians. *Birkhäuser, NY*, 55(58-63):94, 2015.
- Sato, R., Cuturi, M., Yamada, M., and Kashima, H. Fast and robust comparison of probability measures in heterogeneous spaces. *arXiv preprint arXiv:2002.01615*, 2020.
- Solomon, J., Rustamov, R., Guibas, L., and Butscher, A. Earth mover’s distances on discrete surfaces. *ACM Trans. Graph.*, 33(4):1–12, 2014.
- Solomon, J., De Goes, F., Peyré, G., Cuturi, M., Butscher, A., Nguyen, A., Du, T., and Guibas, L. Convolutional wasserstein distances: Efficient optimal transportation on geometric domains. *ACM Trans. Graph.*, 34(4):1–11, 2015.
- Soltanolkotabi, M., Candes, E. J., et al. A geometric analysis of subspace clustering with outliers. *The Annals of Statistics*, 40(4):2195–2238, 2012.
- Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. Wasserstein auto-encoders. *arXiv preprint arXiv:1711.01558*, 2017.
- Villani, C. *Optimal transport: old and new*, volume 338. Springer Science & Business Media, 2008.
- Weed, J., Bach, F., et al. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- Xu, H., Luo, D., Henao, R., Shah, S., and Carin, L. Learning autoencoders with relational regularization. 2020.
- Ye, J., Wu, P., Wang, J. Z., and Li, J. Fast discrete distribution clustering using wasserstein barycenter with sparse support. *IEEE Transactions on Signal Processing*, 65(9): 2317–2332, 2017.
- Yurochkin, M., Claici, S., Chien, E., Mirzazadeh, F., and Solomon, J. M. Hierarchical optimal transport for document representation. In *Advances in Neural Information Processing Systems*, pp. 1601–1611, 2019.