
Hierarchical Optimal Transport for Multimodal Distribution Alignment

John Lee^{†*}, Max Dabagia[†], Eva L. Dyer^{†‡§}, Christopher J. Rozell^{†§}

[†]School of Electrical and Computer Engineering,

[‡]Coulter Department of Biomedical Engineering

Georgia Institute of Technology, Atlanta, GA, 30332 USA

{john.lee, maxdabagia, evadyer, crozell}@gatech.edu

Abstract

In many machine learning applications, it is necessary to meaningfully aggregate, through alignment, different but related datasets. Optimal transport (OT)-based approaches pose alignment as a divergence minimization problem: the aim is to transform a source dataset to match a target dataset using the Wasserstein distance as a divergence measure under alignment constraints. We introduce a hierarchical formulation of OT which leverages clustered structure in data to improve alignment in noisy, ambiguous, or multimodal settings. To solve this numerically, we propose a distributed ADMM algorithm that exploits the Sinkhorn distance, thus it has an efficient computational complexity that scales quadratically with the size of the largest cluster. When the transformation between two datasets is unitary, we provide performance guarantees that describe *when* and *how well* cluster correspondences can be recovered with our formulation, and then describe the worst-case dataset geometry for such a strategy. We apply this method to synthetic datasets that model data as mixtures of low-rank Gaussians and study the impact that different geometric properties of the data have on alignment. Next, we applied our approach to a neural decoding application where the goal is to predict movement directions and instantaneous velocities from populations of neurons in the macaque primary motor cortex. Our results demonstrate that when clustered structure exists in datasets, and is consistent across trials or time points, a hierarchical alignment strategy that leverages such structure can provide significant improvements in cross-domain alignment.

1 Introduction

In many machine learning applications, it is necessary to meaningfully aggregate, through alignment, different but related datasets (e.g., data across time points or under different conditions or contexts). Alignment is an important problem at the heart of transfer learning [1, 2], point set registration [3, 4, 5], and shape analysis [6, 7, 8], but is generally NP hard. In recent years, distribution alignment methods that use optimal transport (OT) to quantify similarity between two distributions have increased in popularity due to their attractive mathematical properties and impressive performance in a variety of tasks [9, 10]. However, using OT to solve unsupervised distribution alignment problems that must simultaneously match two datasets' distributions (using OT) while also learning a transformation between their latent spaces, is extremely challenging, especially when the data has complicated multi-modal structure. Leveraging additional structure in the problem is thus necessary to regularize OT and constrain the solution space.

Here, we leverage the fact that heterogeneous datasets often admit *clustered* or *multi-subspace* structure to improve OT-based distribution alignment. Our solution to this problem is to simultaneously

*JL is currently with DSO National Laboratories of Singapore.

§Equal contributing senior authors.

estimate the cluster alignment across two datasets using their local geometry, while also solving a global alignment problem to meld these local estimates. While it is advantageous to regularize the OT problem with known cluster pairings [10, 11], we are instead concerned with the substantially harder unsupervised setting where such information is missing. We introduce a hierarchical formulation of OT for clustered and multi-subspace datasets called *Hierarchical Wasserstein Alignment (HiWA)*³.

We empirically show that when data are well approximated with Gaussian mixture models (GMMs) or lie on a union of subspaces, we may leverage existing clustering pipelines (e.g., sparse subspace clustering [12] [13]) to improve alignment. When the transformation between datasets is unitary, we provide analyses that reveal key geometric and sampling insights, as well as perturbation and failure mode analyses. To solve the problem numerically, we propose an efficient distributed ADMM algorithm that also exploits the Sinkhorn distance, thus benefiting from efficient computational complexity that scales quadratically with the size of the largest cluster.

To test and benchmark our approach, we applied it to synthetic data generated from mixtures of low-rank Gaussians and studied the impact of different geometric properties of the data on alignment to confirm the predictions of our theoretical analysis. Next, we applied our approach to a neural decoding application where the goal is to predict movement directions from populations of neurons in the macaque primary motor cortex. Our results demonstrate that when clustered structure exists in neural datasets and is consistent across trials or time points, a hierarchical alignment strategy that leverages such structure can provide significant improvements in unsupervised decoding from ambiguous (symmetric) movement patterns. This suggests OT can be applied to a wider range of neural datasets, and shows that a hierarchical strategy avoids local minima encountered by a global alignment strategy that ignores clustered structure.

2 Background and related work

Transfer learning and distribution alignment. A fundamental goal in transfer learning is to aggregate related datasets by learning a mapping between them. We wish to learn a transformation $T \in \mathcal{T}$, where \mathcal{T} refers to some class of transformations that aligns distributions under a notion of probability divergence $\mathcal{D}(\cdot|\cdot)$ between a target distribution μ and a reference (source) distribution ν :

$$\min_{T \in \mathcal{T}} \mathcal{D}(T(\mu)|\nu). \tag{1}$$

Various probability divergences have been proposed in the literature, such as Euclidean least-squares (when data ordering is known) [14, 15, 16], Kullback-Leibler (KL) [17], maximum mean discrepancy (MMD) [18, 19, 20, 21], and the Wasserstein distance [10], where trade-offs are often statistical (e.g., consistency, sample complexity) versus computational. Alignment problems are ill-posed since the space of \mathcal{T} is large, so *a priori* structure is often necessary to constrain \mathcal{T} based on geometric assumptions. Compact manifolds like the Grassmann or Stiefel [22, 23] are primary choices when little information is present, as they preserve isometry. Non-isometric transformations, though richer, demand much more structure (e.g., manifold or graph structure) [24, 25, 26, 27, 10].

Low-rank and union of subspaces models. Principal components analysis (PCA), one of the most popular methods in data science, assumes a *low-rank* model where the top- k principal components of a dataset provide the optimal rank- k approximation under an Euclidean loss. This has been extended to robust (sparse errors) settings [12], and multi- (union of) subspaces settings where data can be partitioned into disjoint subsets where each subset of data is locally low-rank [28]. Transfer learning methods based on subspace alignment [29, 30, 31] work well with zero-mean unimodal datasets, but struggle on more complicated modalities (e.g., Gaussian mixtures or union of subspaces) due to a mixing of covariances. Related to our work, [32] performs multi-subspace alignment by greedily assigning correspondences between subspaces using chordal distances; this however discards valuable information about a distribution’s shape.

Optimal transport. Optimal transport (OT) [33] is a natural type of divergence for registration problems because it accounts for the underlying geometry of the space. In Euclidean settings, OT gives rise to a metric known as the Wasserstein distance $\mathcal{W}(\mu, \nu)$ which measures the minimum effort required to “displace” points across measures μ and ν (understood here as empirical point clouds). Therefore, OT relieves the need for kernel estimation to create an overlapping support of the measures

³MATLAB code can be found at <https://github.com/siplab-gt/hiwa-matlab>. Neural datasets and Python code are provided at <http://nerdslab.github.io/neuralign>

μ, ν . Despite this attractive property, it has both a poor numerical complexity of $O(n^3 \log n)$ (where n is the sample size) and a dimension-dependent sample complexity of $O(n^{-1/d})$, where the data dimension is d [34, 35]. Recently, an entropically regularized version of OT known as the Sinkhorn distance [36] has emerged as a compelling divergence measure; it not only inherits OT’s geometric properties but also has superior computational and sample complexities of $O(n^2)$ and $O(n^{-1/2})^4$, respectively. It has also become a versatile building block in domain adaptation [10, 38]. Prior art [10] has largely exploited the OT’s push-forward as the alignment map since this map minimizes the OT cost between the source and target distributions while allowing a priori structure to be easily incorporated (e.g., to preserve label/graphical integrity). Such an approach, however, is fundamentally expensive when $d \ll n$ since the primary optimization variable is a large transport coupling (i.e., $\mathbb{R}^{n \times n}$), while in reality the alignment mapping is merely $\mathbb{R}^d \mapsto \mathbb{R}^d$. Moreover, it assumes that the source and target distributions are close in terms of their squared Euclidean distance (i.e., an identity transformation), but this does not generally hold between arbitrary latent spaces.

Hierarchical OT and related work. The idea of learning an affine or unitary transformation to align datasets with an OT-based divergence has previously been studied in [39, 40, 41], a problem known as OT Procrustes. However, these methods don’t use problem-specific or clustered structure in data. Hierarchical OT is a recent generalization of OT [42, 43, 44] that is an effective and efficient way of injecting structure into OT but it has never been used to *jointly* solve alignment problems – our work represents a first attempt at doing so. Thus, a key contribution of this paper is putting both of these two ingredients together to develop a scalable strategy that leverages multimodal structure in data solve the OT Procrustes problem.

3 Hierarchical Wasserstein alignment

Preliminaries and notation. Consider clustered datasets $\{\mathbf{X}_i \in \mathbb{R}^{D \times n_{x,i}}\}_{i=1}^S$ and $\{\mathbf{Y}_j \in \mathbb{R}^{D \times n_{y,j}}\}_{j=1}^S$ whose clusters are denoted with the indices i, j and whose columns are treated as \mathbb{R}^D embedding coordinates. The number of samples in the i -th (j -th) cluster of dataset \mathbf{X} (dataset \mathbf{Y}) is given by $n_{x,i}$ ($n_{y,j}$). We express the empirical measures of clusters \mathbf{X}_i and \mathbf{Y}_j as $\mu_i := \frac{1}{n_{x,i}} \sum_{k=1}^{n_{x,i}} \delta_{\mathbf{X}_i(k)}$ and $\nu_j := \frac{1}{n_{y,j}} \sum_{l=1}^{n_{y,j}} \delta_{\mathbf{Y}_j(l)}$, respectively, where $\delta_{\mathbf{x}}$ refers to a point mass located at coordinate $\mathbf{x} \in \mathbb{R}^D$. The squared 2-Wasserstein distance between μ_i and ν_j is defined as

$$\mathcal{W}_2^2(\mu_i, \nu_j) := \min_{\mathbf{Q} \in \mathcal{U}(n_{x,i}, n_{y,j})} \sum_{k=1}^{n_{x,i}} \sum_{l=1}^{n_{y,j}} \mathbf{Q}(k, l) \|\mathbf{X}_i(k) - \mathbf{Y}_j(l)\|_2^2$$

where \mathbf{Q} is a doubly stochastic matrix that encodes point-wise correspondences (i.e., the (k, l) -th entry describes the flow of mass between $\delta_{\mathbf{X}_i(k)}$ and $\delta_{\mathbf{Y}_j(l)}$), $\mathbf{X}_i(k)$ is the k -th column of matrix \mathbf{X}_i , and the constraint $\mathcal{U}(m, n) := \{\mathbf{Q} \in \mathbb{R}_+^{m \times n} : \mathbf{Q} \mathbf{1}_n = \mathbf{1}_m/m, \mathbf{Q}^\top \mathbf{1}_m = \mathbf{1}_n/n\}$ refers to the *uniform* transport polytope (with $\mathbf{1}_m$ a length m vector containing ones). We will use $\|\cdot\|$ to denote the operator norm, \mathbf{X}^\dagger to denote the pseudo-inverse of \mathbf{X} , and \mathbf{I}_d to denote the $d \times d$ identity matrix.

Overview. Although unsupervised alignment is challenging due to the presence of local minima, the imposition of additional structure will help to prune them away. Our key insight is that hierarchical structure decomposes a complicated optimization surface into simpler ones that are less prone to local minima. We formulate a hierarchical Wasserstein approach to align datasets with known (or estimated) clusters $\{\mu_i\}_{i=1}^S, \{\nu_j\}_{j=1}^S$ but whose correspondences are unknown. The task therefore is to jointly learn the alignment T and the cluster-correspondences:

$$\min_{\mathbf{P} \in \mathcal{B}_S, T \in \mathcal{T}} \sum_{i=1}^S \sum_{j=1}^S P_{ij} \mathcal{W}_2^2(T(\mu_i), \nu_j), \quad (2)$$

where the matrix \mathbf{P} encodes the strength of correspondences between clusters, with a large P_{ij} value indicating a correspondence between clusters i, j , and a small value indicating a lack thereof. We note that $\mathcal{B}_S := \mathcal{U}(S, S)$ is a special type of transport polytope known as the S -th Birkhoff polytope. Interestingly, this becomes a nested (or block) OT formulation, where correspondences are resolved at two levels: the outer level resolves cluster-correspondences (via \mathbf{P}) while the inner level resolves point-wise correspondences between cluster points (via the Wasserstein distance).

Alignment over the Stiefel manifold. Assuming clusters lie on subspaces and principal angles between subspaces are “well preserved” across \mathbf{X} and \mathbf{Y} (we make this precise in Theorem 4.2), an

⁴Dependent on a regularization parameter [37].

isometric transformation suffices. Hence, we solve (2) with $\mathcal{T} \leftarrow \mathcal{V}_{D,D}$, the Stiefel manifold which is defined as $\mathcal{V}_{k,d} := \{\mathbf{R} \in \mathbb{R}^{k \times d} : \mathbf{R}^\top \mathbf{R} = \mathbf{I}_d\}$. Explicitly, we can re-formulate equation (2) as:

$$\min_{\mathbf{P}, \mathbf{R}, \{\mathbf{Q}_{ij}\}} \sum_{i,j} P_{ij} C_{ij}(\mathbf{R}, \mathbf{Q}_{ij}) \quad \text{s.t.} \quad \mathbf{P} \in \mathcal{B}_S, \quad \mathbf{R} \in \mathcal{V}_{D,D}, \quad \mathbf{Q}_{ij} \in \mathcal{U}(n_{x,i}, n_{y,j}), \quad (3)$$

$$\text{where} \quad C_{ij}(\mathbf{R}, \mathbf{Q}_{ij}) := \frac{1}{D} \sum_{k,l} \mathbf{Q}_{ij}(k,l) \|\mathbf{R}\mathbf{X}_i(k) - \mathbf{Y}_j(l)\|_2^2 \quad (4)$$

measures pairwise cluster divergences using the squared 2-Wasserstein distance under a Stiefel transformation \mathbf{R} acting on the i^{th} cluster.

Finally, we include entropic regularization over transportation couplings \mathbf{P} and all \mathbf{Q}_{ij} 's to modify the Wasserstein distances to Sinkhorn distances, so as to take advantage of its superior computational and sample complexities. Omitting constraints for brevity, our final problem is given as

$$\min_{\mathbf{P}, \mathbf{R}, \{\mathbf{Q}_{ij}\}} \sum_{i,j} \left(P_{ij} C_{ij}(\mathbf{R}, \mathbf{Q}_{ij}) + H_{\gamma_2}(\mathbf{Q}_{ij}) \right) + H_{\gamma_1}(\mathbf{P}), \quad (5)$$

where $\gamma_1, \gamma_2 > 0$ are the entropic regularization parameters and the negative entropy function is defined as $H_\gamma(\mathbf{P}) := \gamma \sum_{i,j} P_{ij} \log P_{ij}$. Parameters γ_1, γ_2 control the correspondence entropy, therefore (5) approximates (3) when $\gamma_1, \gamma_2 > 0$, but reverts to the original problem (3) as $\gamma_1, \gamma_2 \rightarrow 0$.

Distributed ADMM approach. Problem (5) is non-convex due to multilinearity in the objective and its Stiefel manifold domain. Although alternating directions method of multipliers (ADMM) is a convergent convex solver framework [45, 46], it is being applied in increasingly many non-convex settings [47]. Since (5) readily admits a splitting structure that separates the individual C_{ij} blocks, we develop a distributed ADMM approach. We proceed to split (5) as follows:

$$\min_{\mathbf{P}, \tilde{\mathbf{R}}, \{\mathbf{R}_{ij}, \mathbf{Q}_{ij}\}} \sum_{i,j} \left(P_{ij} C_{ij}(\mathbf{R}_{ij}, \mathbf{Q}_{ij}) + H_{\gamma_2}(\mathbf{Q}_{ij}) \right) + H_{\gamma_1}(\mathbf{P}) \quad \text{s.t.} \quad \mathbf{R}_{ij} = \tilde{\mathbf{R}}, \quad \forall i, j,$$

noting that the set constraints are omitted for brevity. The augmented Lagrangian is given by

$$\mathcal{L}_\mu = \sum_{i,j} \left(P_{ij} C_{ij}(\mathbf{R}_{ij}, \mathbf{Q}_{ij}) + \langle \frac{\mu}{D} \Lambda_{ij}, \mathbf{R}_{ij} - \mathbf{R} \rangle + \frac{\mu}{2D} \|\mathbf{R}_{ij} - \tilde{\mathbf{R}}\|_F^2 + H_{\gamma_2}(\mathbf{Q}_{ij}) \right) + H_{\gamma_1}(\mathbf{P}),$$

where $\mu > 0$ is the ADMM parameter and $\{\Lambda_{ij}\}$ are Lagrange multipliers. Full details of the update steps are included in the Supplementary Material. The algorithm may be summarized in two steps (Alg. 1): (i) a distributed step that asks all cluster pairs to individually find their optimal transformations \mathbf{R}_{ij} in parallel, and (ii) a consensus step that aggregates all the locally estimated transformations according to a weighting that is proportional to correspondence strengths P_{ij} .

Parameters. Entropic parameters γ_1, γ_2 relax the one-to-one cluster correspondence assumption, balancing a trade off between alignment precision (small γ) and sample complexity (large γ). Numerically, negative entropy adds strong convexity to the program, reducing sensitivity towards perturbations at the cost of a slower convergence rate. The ADMM parameter μ controls the ‘strength’ of the consensus, or from an algorithmic viewpoint, the gradient step size.

Distributed consensus. Update steps for $\mathbf{Q}_{ij}, \mathbf{R}_{ij}, \mathbf{L}_{ij}$ can be performed in parallel over all cluster pairs (S^2 in total), making it amenable for a distributed implementation. The runtime complexity of this algorithm is presented in the supplementary Materials.

Robustness against initial conditions. We intentionally build robustness against initial conditions by ordering updates for \mathbf{R}_{ij} and \mathbf{Q}_{ij} before \mathbf{P} such that when μ is sufficiently small, the ADMM sequence is influenced more by the data than by initial conditions.

4 Theoretical guarantees for cluster-based alignment

While the previous section explains *how* to align clustered datasets, in this section, we aim to answer the question of *when* and *how well* they can be aligned. We provide necessary conditions for *cluster-based alignability* as well as *alignment perturbation bounds* according to equation (3)’s formulation. To simplify our analysis, we make the following assumptions: (i) each of the clusters contain the same number of datapoints n , (ii) the ground truth cluster correspondences are $\mathbf{P}^* = \mathbf{I}_S/S$ (i.e.,

Algorithm 1 Hierarchical Wasserstein Alignment (HiWA) Algorithm

```

1: procedure HIERARCHICALWASSERSTEINALIGNMENT( $\gamma_1, \gamma_2, \mu, \{\mathbf{X}_i\}_{i=1}^S, \{\mathbf{Y}_j\}_{j=1}^S$ )
2:    $\mathbf{R} \leftarrow \text{random } \mathcal{V}_{D,D}, \mathbf{P} \leftarrow \mathbb{1}_S \mathbb{1}_S^\top / S^2, \mathbf{\Lambda}_{ij} \leftarrow \mathbf{0}, \forall i, j$  ▷ Initialization
3:   while not converged do
4:     for all  $i, j$  in parallel do
5:        $\mathbf{Q}_{ij} \leftarrow \mathbb{1}_{n_{x,i}} \mathbb{1}_{n_{y,j}}^\top / n_{x,i} n_{y,j}$ 
6:       while not converged do
7:          $\mathbf{R}_{ij} \leftarrow \text{STIEFELALIGNMENT}(2P_{ij} \mathbf{Y}_j \mathbf{Q}_{ij}^\top \mathbf{X}_i^\top + \mu(\mathbf{R} - \mathbf{\Lambda}_{ij}))$ 
8:          $\mathbf{Q}_{ij} \leftarrow \text{SINKHORN}(\gamma_2 / P_{ij}, \mathbf{C}(k, l) \leftarrow \frac{1}{D} \|\mathbf{R}_{ij} \mathbf{X}_i(k) - \mathbf{Y}_j(l)\|_2^2)$ 
9:       end while
10:    end for
11:     $\mathbf{P} \leftarrow \text{SINKHORN}(\gamma_1, \mathbf{C}(i, j) \leftarrow C_{ij}(\mathbf{R}_{ij}, \mathbf{Q}_{ij}))$ 
12:     $\mathbf{R} \leftarrow \text{STIEFELALIGNMENT}(\sum_{i,j} \mathbf{R}_{ij} + \mathbf{\Lambda}_{ij})$ 
13:     $\mathbf{\Lambda}_{ij} \leftarrow \mathbf{\Lambda}_{ij} + \mathbf{R}_{ij} - \mathbf{R}, \forall i, j$ 
14:  end while
15: end procedure

```

<pre> 1: procedure SINKHORN($\gamma, \mathbf{C} \in \mathbb{R}^{m \times n}$) 2: $\mathbf{K} \leftarrow \exp(-\mathbf{C}/\gamma), \mathbf{v} \leftarrow \frac{\mathbb{1}_n}{n}$ 3: while not converged do 4: $\mathbf{u} \leftarrow \frac{\mathbb{1}_m}{m} \oslash \mathbf{K} \mathbf{v}$ 5: $\mathbf{v} \leftarrow \frac{\mathbb{1}_n}{n} \oslash \mathbf{K}^\top \mathbf{u}$ 6: end while 7: $\mathbf{P} \leftarrow \text{diag}(\mathbf{u}) \mathbf{K} \text{diag}(\mathbf{v})$ 8: end procedure </pre>	<pre> 1: procedure STIEFELALIGNMENT(\mathbf{A}) 2: $(\mathbf{U}, \mathbf{\Sigma}, \mathbf{V}) \leftarrow \text{SVD}(\mathbf{A})$ 3: $\mathbf{R} \leftarrow \mathbf{U} \mathbf{V}^\top$ 4: end procedure </pre> <p>Notation: \oslash: elementwise division $\exp(\cdot)$: elementwise exponential $\text{diag}(\cdot)$: diagonal matrix of argument</p>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

diagonal containing $1/S$). However, this analysis can be extended to the case where the number of points is unequal without loss of generality. Detailed proofs are given in the Supp. Material.

The following result is a criterion that, if met, ensures the existence of a global minimizer of the cluster-correspondence \mathbf{P}^* . This criterion requires that matched clusters must be closer in Wasserstein distance than mismatched clusters, according to a threshold determined by Wasserstein’s sample complexity (i.e., an asymptotic rate dependent on the clusters’ *sample sizes* and *intrinsic dimensions*). Since these sample complexity results are based on the Wasserstein distance, we expect a less stringent criterion when using the Sinkhorn distance in (5) (due to superior sample complexity [37]).

Theorem 4.1 (Correspondence disambiguity criterion). *Let all clusters be strictly low-rank where the dimension of the i -th cluster in the x -th dataset is $d_{x,i}$. Let $d_{x,i}, d_{y,j} > 4, \forall i, j \in \llbracket S \rrbracket$. Define $\widehat{C}_{ij}^* := \min_{\mathbf{R} \in \mathcal{V}_{D,D}, \mathbf{Q}_{ij} \in \mathcal{B}_n} C_{ij}(\mathbf{R}, \mathbf{Q}_{ij})$. Problem (3) yields the solution $\mathbf{P}^* = \mathbf{I}_S / S$ with probability at least $1 - \delta$ if, $\forall i, j : i \neq j$, the following criterion is satisfied:*

$$\widehat{C}_{ij}^* + \widehat{C}_{ji}^* - \widehat{C}_{ii}^* - \widehat{C}_{jj}^* > B_{x,i}(\delta) + B_{y,i}(\delta) + B_{x,j}(\delta) + B_{y,j}(\delta)$$

$$\text{where } B_{z,k}(\delta) := c_{z,k} n^{-\frac{2}{d_{z,k}}} + \sqrt{\log(1/\delta)/2n}, \quad c_{z,k} = 1458 \left(2 + \frac{1}{3^{d_{z,k}/2-2} - 1} \right).$$

Proof sketch. The proof contains two parts. In the first part, we consider perturbation conditions of the cost matrix \mathbf{C} in a (non-variational) optimal transport program over the Birkhoff polytope. To be unperturbed from $\mathbf{P}^* = \mathbf{I}_S / S$, we require that $C_{ij} + C_{ji} - C_{ii} - C_{jj} > 0, \forall i, j : i \neq j$. In the second part, we extend this condition to the finite-sample regime by utilizing recently developed concentration bounds [35] for the p -Wasserstein distance, which essentially raises the disambiguity lower bound due to finite-sample uncertainty. (Supp. Material, Section 2) \square

Now, even if we know the global correspondence \mathbf{P}^* , we still do not have the full picture about the alignment’s quality. For example, all matching clusters may have very similar covariances, but principal angles between the clusters are “distorted” across the datasets. Our next theorem gives us an upper bound on the alignment error (for unitary transformations), and makes precise the notion of *global structure distortion*.

Theorem 4.2 (Cluster-based alignment perturbation bounds). *Consider data matrices $\{\mathbf{X}_i, \mathbf{Y}_i \in \mathbb{R}^{D \times n}\}_{i=1}^c$ with known point-wise correspondence matrices $\{\mathbf{Q}_{ii} \in \mathcal{B}_n\}_{i=1}^c$. Define matrices*

$$\mathbf{X} := [\mathbf{X}_1 \mathbf{Q}_{11}, \mathbf{X}_2 \mathbf{Q}_{22}, \dots, \mathbf{X}_c \mathbf{Q}_{cc}], \quad \mathbf{Y} := [\mathbf{Y}_1, \mathbf{Y}_2, \dots, \mathbf{Y}_c].$$

Set $\varepsilon^2 := \|\mathbf{Y}^\top \mathbf{Y} - \mathbf{X}^\top \mathbf{X}\|_F$. If the criterion stated in theorem 4.1 is satisfied, \mathbf{X} is full row rank, and $\varepsilon \|\mathbf{X}^\dagger\| \leq \frac{1}{\sqrt{2}} (\|\mathbf{X}\| \|\mathbf{X}^\dagger\|)^{-1/2}$, then

$$\min_{\mathbf{P} \in \mathcal{B}_c, \mathbf{R} \in \mathcal{V}_{D,D}} \sum_{i,j} P_{ij} C_{ij}(\mathbf{R}) \leq (\|\mathbf{X}\| \|\mathbf{X}^\dagger\| + 2)^2 \|\mathbf{X}^\dagger\|^2 \varepsilon^4 + D,$$

where $D = \sum_{i=1}^c \text{tr}(\mathbf{X}_i(\mathbf{I}/n - \mathbf{Q}_{ii} \mathbf{Q}_{ii}^\top) \mathbf{X}_i^\top + (1/n - 1) \mathbf{Y}_i \mathbf{Y}_i^\top)$ is a data-dependent constant.

Proof sketch. We utilize a recent perturbation result on the Procrustes problem (on a Frobenius norm objective) by Arias-Castro et al. [48] and adapt it to our squared 2-Wasserstein objective. (Supp. Material, Section 3) \square

Note that ε plays a major role in the alignment error bound and quantifies the notion of *global structure distortion*, which allows us to understand on how phenomena like covariate shift or misclustering impacts alignment. To shed some light in this regard, we consider a simple analysis on a cluster-pair’s error contribution to ε , denoted as ε_{ij} . Consider the decomposition of the (i, j) -th block of the Gramians related to clusters i and j , where their respective singular value decompositions are $\mathbf{X}_i \mathbf{Q}_{ii} = \mathbf{A}_i \Sigma_{x,i} \mathbf{V}^\top$ and $\mathbf{Y}_j = \mathbf{B}_j \Sigma_{y,j} \mathbf{V}^\top$. Defining the *blockwise* error between clusters i, j as

$$\varepsilon_{ij} := \|\mathbf{Y}_i^\top \mathbf{Y}_j - \mathbf{Q}_{ij}^\top \mathbf{X}_i^\top \mathbf{X}_j \mathbf{Q}_{jj}\|_F = \|\Sigma_{y,i} \mathbf{B}_i^\top \mathbf{B}_j \Sigma_{y,j} - \Sigma_{x,i} \mathbf{A}_i^\top \mathbf{A}_j \Sigma_{x,j}\|_F,$$

two components stand out: (i) *angular shift*, which is characterized by differences in principal angles between $\mathbf{B}_i^\top \mathbf{B}_j$ and $\mathbf{A}_i^\top \mathbf{A}_j$, and (ii) *spectral shift*, which is characterized by differences in spectra.

Finally, we show that the subspace configuration of a dataset’s clusters can also affect alignment. Pretend for a moment that external alignment information were present to aid in the disambiguation between two clusters. The following lemma tells us when such information is useless (Proof in Supp. Material, Section 4).

Lemma 4.3 (Uninformative alignment). *Consider clusters $\mathbf{X}_i, \mathbf{Y}_j \in \mathbb{R}^{D \times n}$ and known point-wise correspondences $\mathbf{Q}_{ij} \in \mathcal{U}(n, n)$. Denote the left and right singular vectors of $\mathbf{Y}_j \mathbf{Q}_{ij}^\top \mathbf{X}_i^\top$ associated with the non-zero singular values as $\tilde{\mathbf{U}}, \tilde{\mathbf{V}} \in \mathbb{R}^{D \times r}$ with $r \leq D$. Define the set of orthogonal transformations that are constrained to agree with known angular directions as*

$$\mathcal{T}(\mathbf{U}', \mathbf{V}') := \{\mathbf{R} \in \mathbb{R}_+^{D \times D} : \mathbf{R}^\top \mathbf{R} = \mathbf{I}, \mathbf{R} \mathbf{V}' = \mathbf{U}'\},$$

where $\mathbf{U}', \mathbf{V}' \in \mathcal{V}_{D,r}$ with $r \leq D$. Given $\mathbf{U}', \mathbf{V}' \in \mathbb{R}^{D \times r'}$ with $r' \leq D$, we have

$$\min_{\mathbf{R} \in \mathcal{T}(\mathbf{U}', \mathbf{V}')} C_{ij}(\mathbf{R}) \geq \min_{\mathbf{R} \in \mathcal{V}_{D,D}} C_{ij}(\mathbf{R}), \quad (6)$$

with equality holding when $\langle \tilde{\mathbf{U}}, \mathbf{U}' \rangle = \langle \tilde{\mathbf{V}}, \mathbf{V}' \rangle$.

Direct consequences of this lemma are the following: When a dataset has equally-spaced subspaces, it has a maximally uninformative geometric configuration since angular information from other clusters (i.e., \mathbf{U}', \mathbf{V}') can never increase the inter-cluster distance C_{ij} (i.e., equality in (6) always holds); it is hence a worst-case scenario for alignment. This also explains why alignment in very high-dimensional space is harder: All subspaces may be orthogonal to each other, and hence offer no “geometric” advantage.

5 Numerical experiments

5.1 Synthetic low-rank Gaussian mixture dataset

In this section, we validate our method as well as demonstrate its limiting characteristics under symmetric-subspace and finite-sample regimes. To generate our synthetic data, we repeat the following procedure for each of the S clusters. We first randomly generate Gaussian distribution parameters $\mu_i \in \mathbb{R}^d, \Sigma_i \in \mathbb{R}^d : \Sigma_i \succeq 0$ (positive semi-definite), then randomly sample n data-points from these parameters, and finally project them into a random subspace $\mathbf{V}_i \in \mathbb{R}^{D \times d}$ in a $D > d$ dimensional embedding. In these experiments, we assume that the clusters are known, but the




Figure 1: *Synthetic experiments*. HiWA was tested in two subspace configurations (a,b): randomly-spaced (average-case, solid) versus equally-spaced (worst-case, dashed) for $S = 5, d = 2, D = 6, n = \{25, 100\}$, where S is the number of clusters, d the dimension of each cluster, D is the embedding dimension, and n is the sample size. As we expect, performance in terms of the (a) alignment and (b) correspondence error is better in the average (vs. worst) case. In (c,d), we report (c) alignment and (d) correspondence errors as d and n varies, and report the error’s 25th/50th/75th percentiles. In (e,f), we show ablation results (50 trials, no random restarts permitted) for semi-supervised HiWA (known clusters), *completely* unsupervised HiWA-SSC (unknown clusters), non-structured Wasserstein alignment (WA), subspace alignment methods (SA [29], CORAL [31]), and iterative closest point (ICP) [49] for $n = 50, d = 2$, and (e) $S = 5, D = 6$, and (f) $S = 2, D = 2$.

cluster-correspondence across datasets is unknown. We measure performance with respect to two metrics: (i) *alignment error*, defined as the relative difference between the recovered versus true rotation acting on the data $\|\widehat{\mathbf{R}}\mathbf{X} - \mathbf{R}^*\mathbf{X}\|_F^2 / \|\mathbf{R}^*\mathbf{X}\|_F^2$, and (ii) *correspondence error*, defined as the sum of absolute differences between the recovered and the true correspondences $\sum_{ij} |\widehat{\mathbf{P}} - \mathbf{P}^*|_{ij}$.

To understand how global geometry impacts alignment, we applied HiWA in two different settings (Figure 1a-b): (i) a *worst-case setting* where subspaces are equally spaced with a subspace similarity of $\|\mathbf{V}_i^\top \mathbf{V}_j\| = 1, \forall i \neq j$, and (ii) the *random setting* where subspaces are randomly selected from the Grassmann manifold. We observe that equally-spaced subspaces have significantly inferior performance when compared to randomly-spaced subspaces, providing some evidence that equally spaced subspaces are indeed the worst-case scenario in alignment, as suggested by Lemma 4.3.

Next, we studied the effect of dimensions d and sample size n on the accuracy of alignment (Figure 1 (c-d)). We tested HiWA across various dataset conditions by varying parameters $d = \{2, 3, 4, 5\}$ and $n = \{12, 25, 50, 100, 200\}$ while approximately maintaining the average subspace correlations (i.e., $\mathbb{E}\|\mathbf{V}_i^\top \mathbf{V}_j\|$) by fixing the cluster size $S = 5$ and tuning D to control the subspace spacing. In both cases, sample complexities are better than the theoretical rate of $O(n^{-1/d})$, which is likely due to the Sinkhorn distance’s superior sample complexity. In Figure 1e-f, we conduct an ablation study and evaluate our algorithm against benchmark methods in transfer learning and point set registration in two settings: a simple one in low- d (e) and a harder one in higher- d (f). Specifically, we compare HiWA when clusters are known (but pairwise correspondences are unknown), HiWA with clustering via sparse subspace clustering [12] (HiWA-SSC) to represent *completely* unsupervised alignment, a Wasserstein alignment variant with no cluster-structure (WA) which is akin to OT Procrustes [50, 39, 40, 41], subspace alignment [29], correlation alignment [31], and iterative closest point (ICP) [49]. HiWA exhibits strongest performance, with HiWA-SSC trailing closely behind (since clusters are independently resolved), followed by WA, then other algorithms. Subspace alignment methods have remarkably poor performance in higher dimensions due to their inability to resolve subspace sign ambiguities, while ICP demonstrates its notorious dependence on good initial conditions. These results indicates HiWA’s strong robustness against initial conditions and good scaling properties.

5.2 Neural population decoding example

Decoding intent (e.g., where you want to move your arm) or evoked responses (e.g., what you are looking at or listening to) directly from neural activity is a widely studied problem in neuroscience, and the first step in the design of a brain machine interface (BMI). A critical challenge with BMIs is that neural decoders need to be recalibrated (or re-trained) due to drift in neural responses or electrophysiology measurements/readouts [51]. A recent method for semi-supervised brain decoding finds a transformation between projected neural responses and movements by solving a KL-divergence




Figure 2: *Results on neural decoding dataset:* How distribution alignment is used to translate neural activity into movement – low-dimensional embeddings of neural data are aligned with target movement patterns (a). In (b), we compare the performance (cluster correspondence) of HiWA, WA, and DAD as the number of points in the source dataset decreases. Next, we compared the performance of HiWA with known and estimated clusters (via GMM). Movement patterns in which cluster separability is high and the geometry is preserved across datasets, can be aligned in both cases (green stars). Patterns where separability is low but geometry is useful can be aligned when the cluster arrangements are known are denoted with yellow stars.

minimization problem [52]. Using this approach, one could build robust decoders that work across days and shifts in neural responses through alignment.

We test the utility of hierarchical alignment for neural decoding on datasets collected from the arm region of primary motor cortex of a non-human primate (NHP) during a center out reaching task [52]. After spike sorting and binning the data, we applied factor analysis to reduce the data dimensionality to 3D (source distribution) and applied HiWA to align the neural data to a 3D movement distribution (target distribution) (Figure 2). We compared its performance to (procrustes) Wasserstein alignment (WA) without hierarchical structure, and a baseline brute force search method called distribution alignment decoding (DAD) [52]. We examined the prediction accuracy of the target reach direction for the motor decoding task (i.e., the cluster classification accuracy).

Next, we examined the impact of the sampling density (Figure 2b) on alignment performance. Our results demonstrate that HiWA continues to produce consistent cluster correspondences (> 70% accuracy), even as the number of samples per cluster drops to 8. In comparison, DAD is competitive at larger sample sizes but its performance rapidly drops off as sampling density decreases because it requires estimating a distribution from samples. WA suffers from the presence of many local minima and fails to find the correct cluster correspondences. Our results suggest that HiWA consistently provides stable solutions, outperforming competitor methods for this application.

Finally, to study the impact of local and global geometry on whether an unlabeled source and target can be aligned, we applied HiWA to permutations of eight subsets of reach directions (movement patterns). When just two reach directions are considered (Figure 2c, Columns 1-4), global geometry becomes useless in determining the correct rotation. In this case, we observe that HiWA is only capable of consistent alignment when cluster asymmetries are sufficiently extreme in both the source and target. When three reach directions are considered (Figure 2c, Columns 5-8), the global geometry can be used, yet there still exist symmetrical cases where recovering the correct rotation is difficult without adequate local asymmetries or some supervised (labeled) data to match clusters. These results suggest that hierarchical structure can be critical in resolving ambiguities in alignment of globally symmetric movement distributions.

6 Conclusion

This paper introduces a new method for hierarchical alignment with Wasserstein distances, provided an efficient numerical solution with analytical guarantees. We tested our method and compared its performance against other methods on a synthetic mixture model dataset and on a real neural decoding dataset. Future directions include extensions to non-rigid transformations, and applications to higher dimensional neural datasets that do not rely on external measured behavioral covariates.

Acknowledgments

JL was supported by DSO National Laboratories of Singapore, ED and MD were supported by NSF grant IIS-1755871, and CR was supported by NSF grant CCF-1409422 and CAREER award CCF-1350954.

References

- [1] Sinno Jialin Pan and Qiang Yang. A survey on transfer learning. *IEEE Transactions on Knowledge and Data Engineering*, 22(10):1345–1359, 2009.
- [2] Karl Weiss, Taghi M Khoshgoftaar, and DingDing Wang. A survey of transfer learning. *Journal of Big data*, 3(1):9, 2016.
- [3] Haili Chui and Anand Rangarajan. A new point matching algorithm for non-rigid registration. *Computer Vision and Image Understanding*, 89(2-3):114–141, 2003.
- [4] Andriy Myronenko and Xubo Song. Point set registration: Coherent point drift. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(12):2262–2275, 2010.
- [5] Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. Registration of 3d point clouds and meshes: a survey from rigid to nonrigid. *IEEE Transactions on Visualization and Computer Graphics*, 19(7):1199–1217, 2013.
- [6] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences*, 103(5):1168–1172, 2006.
- [7] Alexander M Bronstein, Michael M Bronstein, Leonidas J Guibas, and Maks Ovsjanikov. Shape google: Geometric words and expressions for invariant shape retrieval. *ACM Transactions on Graphics (TOG)*, 30(1):1, 2011.
- [8] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics*, 31(4):30, 2012.
- [9] Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5-6):355–607, 2019.
- [10] 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.
- [11] Debasmit Das and CS George Lee. Unsupervised domain adaptation using regularized hyper-graph matching. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, pages 3758–3762. IEEE, 2018.
- [12] Ehsan Elhamifar and Rene Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(11):2765–2781, 2013.
- [13] Eva L Dyer, Aswin C Sankaranarayanan, and Richard G Baraniuk. Greedy feature selection for subspace clustering. *The Journal of Machine Learning Research*, 14(1):2487–2517, 2013.
- [14] Xiaoxiao Shi, Qi Liu, Wei Fan, S Yu Philip, and Ruixin Zhu. Transfer learning on heterogenous feature spaces via spectral transformation. In *Data Mining (ICDM), 2010 IEEE 10th International Conference on*, pages 1049–1054. IEEE, 2010.
- [15] Sumit Shekhar, Vishal M Patel, Hien V Nguyen, and Rama Chellappa. Generalized domain-adaptive dictionaries. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 361–368, 2013.
- [16] Yahong Han, Fei Wu, Dacheng Tao, Jian Shao, Yueting Zhuang, and Jianmin Jiang. Sparse unsupervised dimensionality reduction for multiple view data. *IEEE Transactions on Circuits and Systems for Video Technology*, 22(10):1485, 2012.
- [17] Masashi Sugiyama, Shinichi Nakajima, Hisashi Kashima, Paul V Buenau, and Motoaki Kawanabe. Direct importance estimation with model selection and its application to covariate shift adaptation. In *Advances in Neural Information Processing Systems*, pages 1433–1440, 2008.
- [18] Sinno Jialin Pan, Ivor W Tsang, James T Kwok, and Qiang Yang. Domain adaptation via transfer component analysis. *IEEE Transactions on Neural Networks*, 22(2):199–210, 2011.

- [19] Mahsa Baktashmotlagh, Mehrtash T Harandi, Brian C Lovell, and Mathieu Salzmann. Unsupervised domain adaptation by domain invariant projection. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 769–776, 2013.
- [20] Mingsheng Long, Jianmin Wang, Guiguang Ding, Jianguang Sun, and Philip S Yu. Transfer joint matching for unsupervised domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1410–1417, 2014.
- [21] Mingming Gong, Kun Zhang, Tongliang Liu, Dacheng Tao, Clark Glymour, and Bernhard Schölkopf. Domain adaptation with conditional transferable components. In *Proceedings of the 33rd International Conference on Machine Learning*, pages 2839–2848, 2016.
- [22] Raghuraman Gopalan, Ruonan Li, and Rama Chellappa. Domain adaptation for object recognition: An unsupervised approach. In *IEEE International Conference on Computer Vision*, pages 999–1006. IEEE, 2011.
- [23] Boqing Gong, Yuan Shi, Fei Sha, and Kristen Grauman. Geodesic flow kernel for unsupervised domain adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2066–2073. IEEE, 2012.
- [24] Chang Wang and Sridhar Mahadevan. Manifold alignment using procrustes analysis. In *Proceedings of the 25th International Conference on Machine learning*, pages 1120–1127. ACM, 2008.
- [25] Chang Wang and Sridhar Mahadevan. A general framework for manifold alignment. In *2009 AAAI Fall Symposium Series*, 2009.
- [26] Sira Ferradans, Nicolas Papadakis, Gabriel Peyré, and Jean-François Aujol. Regularized discrete optimal transport. *SIAM Journal on Imaging Sciences*, 7(3):1853–1882, 2014.
- [27] Zhen Cui, Hong Chang, Shiguang Shan, and Xilin Chen. Generalized unsupervised manifold alignment. In *Advances in Neural Information Processing Systems*, pages 2429–2437, 2014.
- [28] Yonina C Eldar and Moshe Mishali. Robust recovery of signals from a structured union of subspaces. *IEEE Transactions on Information Theory*, 55(11):5302–5316, 2009.
- [29] Basura Fernando, Amaury Habrard, Marc Sebban, and Tinne Tuytelaars. Unsupervised visual domain adaptation using subspace alignment. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2960–2967, 2013.
- [30] Baochen Sun and Kate Saenko. Subspace distribution alignment for unsupervised domain adaptation. In *BMVC*, volume 4, pages 24–1, 2015.
- [31] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [32] Kowshik Thopalli, Rushil Anirudh, Jayaraman J Thiagarajan, and Pavan Turaga. Multiple subspace alignment improves domain adaptation. *arXiv preprint arXiv:1811.04491*, 2018.
- [33] Leonid Vitalevich Kantorovich. On a problem of monge. *Journal of Mathematical Sciences*, 133(4):1383–1383, 2006.
- [34] Richard M Dudley. The speed of mean glivenko-cantelli convergence. *The Annals of Mathematical Statistics*, 40(1):40–50, 1969.
- [35] Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *arXiv preprint arXiv:1707.00087*, 2017.
- [36] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in Neural Information Processing Systems*, pages 2292–2300, 2013.
- [37] Aude Genevay, Lénaïc Chizat, Francis Bach, Marco Cuturi, and Gabriel Peyré. Sample complexity of sinkhorn divergences. *arXiv preprint arXiv:1810.02733*, 2018.
- [38] Nicolas Courty, Rémi Flamary, Amaury Habrard, and Alain Rakotomamonjy. Joint distribution optimal transportation for domain adaptation. In *Advances in Neural Information Processing Systems*, pages 3730–3739, 2017.
- [39] Meng Zhang, Yang Liu, Huanbo Luan, and Maosong Sun. Earth mover’s distance minimization for unsupervised bilingual lexicon induction. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1934–1945, 2017.

- [40] David Alvarez-Melis, Stefanie Jegelka, and Tommi S Jaakkola. Towards optimal transport with global invariances. *arXiv preprint arXiv:1806.09277*, 2018.
- [41] Edouard Grave, Armand Joulin, and Quentin Berthet. Unsupervised alignment of embeddings with wasserstein procrustes. *arXiv preprint arXiv:1805.11222*, 2018.
- [42] Mikhail Yurochkin, Sebastian Claiçi, Edward Chien, Farzaneh Mirzazadeh, and Justin Solomon. Hierarchical optimal transport for document representation. *arXiv preprint arXiv:1906.10827*, 2019.
- [43] Bernhard Schmitzer and Christoph Schnörr. A hierarchical approach to optimal transport. In *International Conference on Scale Space and Variational Methods in Computer Vision*, pages 452–464. Springer, 2013.
- [44] David Alvarez-Melis, Tommi S Jaakkola, and Stefanie Jegelka. Structured optimal transport. *arXiv preprint arXiv:1712.06199*, 2017.
- [45] Jonathan Eckstein and Dimitri P Bertsekas. On the douglas—rachford splitting method and the proximal point algorithm for maximal monotone operators. *Mathematical Programming*, 55(1-3):293–318, 1992.
- [46] Stephen Boyd, Neal Parikh, Eric Chu, Borja Peleato, Jonathan Eckstein, et al. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.
- [47] Yu Wang, Wotao Yin, and Jinshan Zeng. Global convergence of admm in nonconvex nonsmooth optimization. *Journal of Scientific Computing*, 78(1):29–63, 2019.
- [48] Ery Arias-Castro, Adel Javanmard, and Bruno Pelletier. Perturbation bounds for procrustes, classical scaling, and trilateration, with applications to manifold learning. *arXiv preprint arXiv:1810.09569*, 2018.
- [49] Paul J Besl and Neil D McKay. Method for registration of 3-d shapes. In *Sensor Fusion IV: Control Paradigms and Data Structures*, volume 1611, pages 586–607. International Society for Optics and Photonics, 1992.
- [50] Anand Rangarajan, Haili Chui, and Fred L Bookstein. The softassign procrustes matching algorithm. In *Biennial International Conference on Information Processing in Medical Imaging*, pages 29–42. Springer, 1997.
- [51] Chethan Pandarinath, K Cora Ames, Abigail A Russo, Ali Farshchian, Lee E Miller, Eva L Dyer, and Jonathan C Kao. Latent factors and dynamics in motor cortex and their application to brain–machine interfaces. *Journal of Neuroscience*, 38(44):9390–9401, 2018.
- [52] Eva L Dyer, Mohammad Gheshlaghi Azar, Matthew G Perich, Hugo L Fernandes, Stephanie Naufel, Lee E Miller, and Konrad P Körding. A cryptography-based approach for movement decoding. *Nature Biomedical Engineering*, 1(12):967, 2017.

Hierarchical Optimal Transport for Multimodal Distribution Alignment: Supplementary Material

1 ADMM algorithmic details

1.1 Derivation of ADMM update steps

ADMM admits the following sequence of updates:

$$(\mathbf{R}_{ij}^{(t+1)}, \mathbf{Q}_{ij}^{(t+1)}) \leftarrow \arg \min_{\substack{\mathbf{R}_{ij} \in S(D, D) \\ \mathbf{Q}_{ij} \in \mathcal{U}(n_{x,i}, n_{y,j})}} P_{ij}^{(t)} C_{ij}(\mathbf{R}_{ij}, \mathbf{Q}_{ij}) + \frac{\mu}{2D} \|\mathbf{R}_{ij} - \mathbf{R}^{(t)} + \mathbf{\Lambda}_{ij}^{(t)}\|_F^2 + H_{\gamma_2}(\mathbf{Q}_{ij}), \quad (1)$$

$$\mathbf{P}^{(t+1)} \leftarrow \arg \min_{\mathbf{P} \in \mathcal{B}_c} \sum_{i,j} P_{ij} C_{ij}(\mathbf{R}_{ij}^{(t+1)}, \mathbf{Q}_{ij}^{(t+1)}) + H_{\gamma_1}(\mathbf{P}), \quad (2)$$

$$\mathbf{R}^{(t+1)} \leftarrow \arg \min_{\mathbf{R} \in S(D, D)} \sum_{i,j} \|\mathbf{R}_{ij}^{(t+1)} - \mathbf{R} + \mathbf{\Lambda}_{ij}^{(t)}\|_F^2, \quad (3)$$

$$\mathbf{\Lambda}_{ij}^{(t+1)} \leftarrow \mathbf{\Lambda}_{ij}^{(t)} + \mathbf{R}_{ij}^{(t+1)} - \mathbf{R}. \quad (4)$$

Update (1) involves an alternating minimization over \mathbf{Q}_{ij} and \mathbf{R}_{ij} whereby the first variable is fixed while the second is minimized, followed by the second fixed and the first minimized, and the procedure is repeated until convergence is achieved. When solving for \mathbf{R}_{ij} we have the following Stiefel manifold optimization:

$$\begin{aligned} \mathbf{R}_{ij}^{(t+1)} &\leftarrow \arg \min_{\mathbf{R}_{ij} \in S(D, D)} P_{ij}^{(t)} C_{ij}(\mathbf{R}_{ij}, \mathbf{Q}_{ij}) + \frac{\mu}{2D} \|\mathbf{R}_{ij} - \mathbf{R}^{(t)} + \mathbf{\Lambda}_{ij}^{(t)}\|_F^2 \\ &= \arg \max_{\mathbf{R}_{ij} \in S(D, D)} \text{tr} \left((2P_{ij}^{(t)} \mathbf{Y}_j \mathbf{Q}_{ij}^\top \mathbf{X}_i^\top + \mu(\mathbf{R}^{(t)} - \mathbf{\Lambda}_{ij}^{(t)})) \mathbf{R}_{ij}^\top \right) = \mathbf{U} \mathbf{V}^\top, \end{aligned}$$

where $2P_{ij}^{(t)} \mathbf{Y}_j \mathbf{Q}_{ij}^\top \mathbf{X}_i^\top + \mu(\mathbf{R}^{(t)} - \mathbf{\Lambda}_{ij}^{(t)}) = \mathbf{U} \mathbf{D} \mathbf{V}^\top$ is its SVD. We employ the Sinkhorn algorithm (Algorithm 1 of [1]) to solve for \mathbf{Q}_{ij} , using an entropic parameter of $\gamma_2/P_{ij}^{(t)}$ and uniform marginals. $\mathbf{R}_{ij}^{(t+1)}$ and $\mathbf{Q}_{ij}^{(t+1)}$ are retrieved once the alternating minimization converges.

Update (2) also employs the Sinkhorn algorithm over the cost matrix generated by $C_{ij}(\mathbf{R}_{ij}^{(t+1)}, \mathbf{Q}_{ij}^{(t+1)})$ using variables found in update (1), along with an entropic parameter of γ_1 and uniform marginals.

Update (3) is a consensus update over a Stiefel manifold optimization:

$$\mathbf{R}^{(t+1)} \leftarrow \arg \min_{\mathbf{R} \in S(D, D)} \sum_{i,j} \|\mathbf{R}_{ij}^{(t+1)} - \mathbf{R} + \mathbf{\Lambda}_{ij}^{(t)}\|_F^2 = \mathbf{U} \mathbf{V}^\top,$$

where $\sum_{i,j} \mathbf{R}_{ij}^{(t+1)} + \mathbf{\Lambda}_{ij}^{(t)} = \mathbf{U} \mathbf{D} \mathbf{V}^\top$ is its SVD.

1.2 Computational complexity of distributed ADMM algorithm

The main computational complexity of the ADMM algorithm comes from line 8, which demands the solving of S^2 Sinkhorn problems per ADMM iteration. Fortunately, these Sinkhorn problems (i.e., update steps for $\mathbf{Q}_{ij}, \mathbf{R}_{ij}, \mathbf{L}_{ij}$) may be conducted in parallel, making it amenable for a distributed implementation. When fully parallelized, the algorithm has a per-iteration computational complexity of $O(n_i n_j)$, where n_i, n_j refers to the number of points in the largest clusters of \mathbf{X}, \mathbf{Y} respectively (compared to vanilla Sinkhorn’s $O(n_x n_y)$ complexity where n_x, n_y refers to the *total* number of points in respective datasets, assuming $D \ll \max(n_i, n_j)$). In Figure 1, we record run times of two versions of the same algorithm, with and without parallelism, operating on a similar test dataset, at varying cluster sizes S (50 data points per cluster). 10 random trials were conducted and the means and standard deviations were reported (error lines). As expected, the parallel implementation scales better. These results were computed using an i7 intel processor with 6 cores clocked at 2.6 GHz.

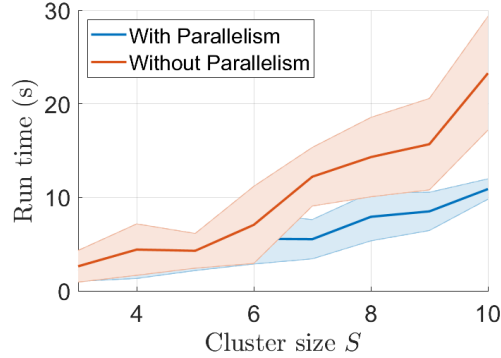


Figure 1: Run time comparison between a parallelized and a non-parallelized implementation: We varied the number of clusters – the main bottleneck to run time – to compare the runtimes of the parallelized and non-parallelized implementations.

1.3 Stopping criteria.

In lines 3 and 6 of Algorithm 1, possible stopping criteria are (i) $\|\mathbf{R}^{(t+1)} - \mathbf{R}^{(t)}\|_F \leq \tau$ where the difference is between the current and previous iteration’s transformation and τ is the tolerance, and (ii) $t \leq T$ where T is the maximum number of iterations.

2 Proof of Theorem 4.1

2.1 Part 1: Geometric perturbations conditions over the Birkhoff polytope

First, we provide the following lemma that illuminates some basic geometrical insights of the general OT cost matrix, whose coupling is in the Birkhoff polytope $\mathcal{B}_c := \{\mathbf{P} \in \mathbb{R}^{c \times c} : \mathbf{P}\mathbb{1}_c = \mathbf{P}^\top \mathbb{1}_c = \mathbb{1}_c\}$. We define the OT program with respect to a cost matrix \mathbf{C} as

$$L(\mathbf{C}) := \arg \min_{\mathbf{P} \in \mathcal{B}_c} \langle \mathbf{P}, \mathbf{C} \rangle. \quad (5)$$

The following lemma describes the conditions on \mathbf{C} for \mathbf{P} to remain unperturbed at \mathbf{P}^* .

Lemma 2.1. Define the set of vertices on the c -th Birkhoff polytope \mathcal{B}_c that are within a δ -ball from \mathbf{P}^* as

$$\mathcal{P}_\delta(\mathbf{P}^*) = \{\mathbf{P} \in \mathcal{B}_c \setminus \{\mathbf{P}^*\} : \|\mathbf{P} - \mathbf{P}^*\|_F \leq \delta\}. \quad (6)$$

Define the set of matrices that denote directions from \mathbf{P}^* to each neighboring vertex \mathbf{P}_i as

$$\mathcal{V}_\delta(\mathbf{P}^*) = \{\mathbf{V} \in \mathbb{R}^{c \times c} : \mathbf{V} = \mathbf{P} - \mathbf{P}^*, \mathbf{P} \in \mathcal{P}_\delta(\mathbf{P}^*)\}. \quad (7)$$

For the linear program’s solution $\mathbf{P}^* = L(\mathbf{C})$ to remain unchanged, $\forall \mathbf{V} \in \mathcal{V}_\delta(\mathbf{P}^*)$ for $\delta = 2$, the cost matrix \mathbf{C} should satisfy

$$\langle \mathbf{C}, \mathbf{V} \rangle > 0. \quad (8)$$

Proof. Birkhoff-von Neuman's theorem [2] states that the optimal transport solution must lie on the convex hull of the c -th Birkhoff polytope \mathcal{B}_c , and that its vertices are in fact permutation matrices. We therefore say that an LP solution \mathbf{P}^* is a vertex on \mathcal{B}_c . The outline of this proof is straightforward: so that \mathbf{P}^* remains unchanged, \mathbf{C} should not cause \mathbf{P}^* to move to an adjacent edge of the Birkhoff polytope, nor should it cause it to extend beyond its adjacent edge because then, the neighboring vertex would assume the new solution. For the rest of this proof, we shall let $\mathbf{P}^* = \mathbf{I}$ without any loss of generality. We define the set of nearest neighbors to \mathbf{P}^* , which are simply permutation matrices that can be described as taking \mathbf{P}^* and exchanging any two columns. For a \mathbf{P} with any two columns of \mathbf{P}^* exchanged, notice that the difference matrix $\mathbf{V} = \mathbf{P} - \mathbf{P}^*$ is a symmetric matrix of mostly zeros except for two off-diagonal $+1$ entries and two diagonal -1 entries, hence $\|\mathbf{P} - \mathbf{P}^*\|_F = 2$. Formally, we describe the set of nearest neighbors with $\mathcal{P}_\delta(\mathbf{P}^*)$ with $\delta = 2$, defined by (6). Next, we define edges that are adjacent to \mathbf{P}^* using the set $\mathcal{V}_\delta(\mathbf{P}^*)$ also with $\delta = 2$, defined by (7). Note that there are $K := \binom{c}{2} - 1 = \frac{c(c-1)}{2} - 1$ neighboring vertices and adjacent edges.

We will now show how perturbing \mathbf{C} in just one direction \mathbf{V} changes \mathbf{P}^* . First, consider a cost matrix that produce \mathbf{P}^* , which is defined as $\mathbf{C}^* := \sum_{\mathbf{V}_i \in \mathcal{V}_\delta(\mathbf{P}^*)} \mathbf{V}_i$, meaning that it is equi-angle from all $\mathbf{V} \in \mathcal{V}_\delta(\mathbf{P}^*)$, for $\delta = 2$. By enumerating over all \mathbf{V} , we may derive $\mathbf{C}^* = \mathbb{1}\mathbb{1}^\top - c\mathbf{I}$, along with the fact that $\frac{\langle \mathbf{C}^*, \mathbf{V} \rangle}{\|\mathbf{C}^*\|_F \|\mathbf{V}\|_F} = \frac{1}{\sqrt{c-1}}$. Define $\mathbf{P}_1 = L(\mathbf{C}_1)$ with \mathbf{C}_1 such that $\langle \mathbf{V}_1, \mathbf{C}_1 \rangle \leq 0$ and $\frac{\langle \mathbf{V}_1, \mathbf{C}_1 \rangle}{\|\mathbf{V}_1\|_F \|\mathbf{C}_1\|_F} = \frac{1}{\sqrt{c-1}}$, $\forall i \in \llbracket K \rrbracket$. Define $\mathbf{P}_2 = L(\mathbf{C}_2)$ with \mathbf{C}_2 such that $\frac{\langle \mathbf{V}_1, \mathbf{C}_2 \rangle}{\|\mathbf{V}_1\|_F \|\mathbf{C}_2\|_F} = \frac{1}{\sqrt{c-1}}$, $\forall i \in \llbracket K \rrbracket$. We make the claim that $\mathbf{P}_1 = \mathbf{P}_2$ and proceed with a proof by contradiction. As mentioned before, $\forall \mathbf{V} \in \mathcal{V}_\delta(\mathbf{P}^*)$ for $\mathbf{P}^* = \mathbf{I}$, $\delta = 2$ has exactly four non-zero entries, i.e., $V_{ii} = V_{jj} = -1$ and $V_{ij} = V_{ji} = +1$, where $i, j \in \llbracket c \rrbracket, i \neq j$. Writing out $\langle \mathbf{V}, \mathbf{C} \rangle$ explicitly, we have $\langle \mathbf{V}, \mathbf{C} \rangle = -(C_{ii} + C_{jj}) + (C_{ij} + C_{ji})$. To ensure that \mathbf{C}_1 and \mathbf{C}_2 does not interact with other edges of the polytope, we fix $\langle \mathbf{V}_i, \mathbf{C} \rangle = \langle \mathbf{V}_i, \mathbf{C}^* \rangle = c$ for $i = 2, \dots, K$. Since we constructed \mathbf{C}_1 and \mathbf{C}_2 to differ only by the condition $\langle \mathbf{V}_1, \mathbf{C}_1 \rangle \leq 0$ or $\langle \mathbf{V}_1, \mathbf{C}_2 \rangle > 0$, and any \mathbf{V}_i affects only four entries of the cost matrix, we may greatly simplify our analysis of \mathbf{C} and \mathbf{P} to only these four entries. As such, we extract these four entries of \mathbf{P} , represent it using $\hat{\mathbf{P}} \in \mathcal{B}_2$, and parameterize using $t \in [0, 1]$ it as

$$\hat{\mathbf{P}}(t) = t \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} + (1-t) \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}. \quad (9)$$

With this parameterized form, we may reexpress the optimization for $L(\mathbf{C})$ as

$$\min_{t \in [0,1]} t(C_{ii} + C_{jj}) + (1-t)(C_{ij} + C_{ji}) = \min_{t \in [0,1]} \left(\frac{C_{ii} + C_{jj}}{C_{ij} + C_{ji}} - 1 \right) t$$

The above minimization has three cases. If $C_{ii} + C_{jj} = C_{ij} + C_{ji}$ then there exists no unique solution t . If $C_{ii} + C_{jj} > C_{ij} + C_{ji}$ then $t = 0$. If $C_{ii} + C_{jj} < C_{ij} + C_{ji}$ then $t = 1$. The first two cases directly corresponds to $\langle \mathbf{V}_1, \mathbf{C}_1 \rangle \leq 0$, while the third case corresponds to $\langle \mathbf{V}_1, \mathbf{C}_2 \rangle > 0$. The fact that t is not consistent between all cases demonstrates a contradiction. Moreover, $t = 1$ produces the solution $\hat{\mathbf{P}} = \mathbf{I}$, and if this holds for all $\mathbf{V} \in \mathcal{V}_\delta(\mathbf{P}^*)$, then all off-diagonal entries must be zero and therefore $\mathbf{P}^* = \mathbf{I}$ must be the minimizer for $L(\mathbf{C})$. \square

A direct consequence of lemma 2.1 is the following.

Corollary 2.2. *The solution to the linear program defined by (5) is $\mathbf{P}^* = \mathbf{I}$ if the linear cost matrix \mathbf{C} satisfies the following property*

$$C_{ij} + C_{ji} - C_{ii} - C_{jj} > 0, \quad \forall i, j \in \llbracket c \rrbracket, i \neq j. \quad (10)$$

Proof. Analyzing (8), we observe that any $\mathbf{V} \in \mathcal{V}_\delta(\mathbf{P}^*)$ for $\delta = 2$ has only four symmetric non-negative entries, which we condense \mathbf{V} and \mathbf{C} into $\mathbb{R}^{2 \times 2}$ matrices at these four support locations respectively as $\hat{\mathbf{V}} = \begin{bmatrix} -1 & +1 \\ +1 & -1 \end{bmatrix}$ and $\hat{\mathbf{C}} = \begin{bmatrix} C_{ii} & C_{ji} \\ C_{ij} & C_{jj} \end{bmatrix}$. It thus follows that an explicit computation produces:

$$\langle \mathbf{P}, \mathbf{C} \rangle \geq 0 \Rightarrow \langle \hat{\mathbf{P}}, \hat{\mathbf{C}} \rangle = (C_{ij} + C_{ji}) - (C_{ii} + C_{jj}) \geq 0.$$

Since the set $\mathcal{V}_\delta(\mathbf{P}^*)$ spans all permutations between $i, j \in \llbracket c \rrbracket, i \neq j$, we conclude with (10). \square

In the variational setting, $C_{ij}(\mathbf{R}, \mathbf{Q}_{ij})$'s are themselves linearly coupled with \mathbf{R} and \mathbf{Q}_{ij} . The following proposition introduces a trivial criterion on $C_{ij}(\mathbf{R}, \mathbf{Q}_{ij})$'s to guarantee that \mathbf{P} remains unperturbed from \mathbf{P}^* . For pedagogical reasons, we shall assume that $n \rightarrow \infty$ for this proposition but subsequently relax this.

Proposition 2.3 (Rotationally invariant disambiguity criterion). *Problem (3) yields the solution \mathbf{P}^* if, $\forall i, j : i \neq j$, the following criterion is satisfied:*

$$\min_{\mathbf{R}, \mathbf{Q}_{ij}} C_{ij}(\mathbf{R}, \mathbf{Q}_{ij}) + \min_{\mathbf{R}, \mathbf{Q}_{ji}} C_{ji}(\mathbf{R}, \mathbf{Q}_{ji}) - \min_{\mathbf{R}, \mathbf{Q}_{ii}} C_{ii}(\mathbf{R}, \mathbf{Q}_{ii}) - \min_{\mathbf{R}, \mathbf{Q}_{jj}} C_{jj}(\mathbf{R}, \mathbf{Q}_{jj}) > 0. \quad (11)$$

Proof. Consider a set of c clusters where clusters i, j satisfy

$$C_{ij}(\mathbf{R}'_{ij}, \mathbf{Q}'_{ij}) + C_{ji}(\mathbf{R}'_{ji}, \mathbf{Q}'_{ji}) - C_{ii}(\mathbf{R}', \mathbf{Q}'_{ii}) - C_{jj}(\mathbf{R}', \mathbf{Q}'_{jj}) \leq 0, \quad (12)$$

where

$$\begin{aligned} (\mathbf{R}'_{ij}, \mathbf{Q}'_{ij}) &:= \arg \min_{\mathbf{R} \in S(d, d), \mathbf{Q}_{ij} \in \mathcal{B}_n} C_{ij}(\mathbf{R}), \\ (\mathbf{R}'_{ji}, \mathbf{Q}'_{ji}) &:= \arg \min_{\mathbf{R} \in S(d, d), \mathbf{Q}_{ji} \in \mathcal{B}_n} C_{ji}(\mathbf{R}), \\ (\mathbf{R}', \mathbf{Q}'_{ii}, \mathbf{Q}'_{jj}) &:= \arg \min_{\mathbf{R} \in S(d, d), \mathbf{Q}_{ii}, \mathbf{Q}_{jj} \in \mathcal{B}_n} C_{ii}(\mathbf{R}, \mathbf{Q}_{ii}) + C_{jj}(\mathbf{R}, \mathbf{Q}_{jj}). \end{aligned}$$

Since the 2-Wasserstein is a valid metric, its distance between any two clusters must satisfy $C_{ij}(\mathbf{R}) \geq 0$ with equality holding if and only if the clusters are exactly similar. If $\mathbf{P}^* = \mathbf{I}$, similar clusters are denoted with matching indices, and it must follow that $C_{ii}(\mathbf{R}') + C_{jj}(\mathbf{R}') = 0$. This implies that (12) must be false since $C_{ij}(\mathbf{R}'_{ij}), C_{ji}(\mathbf{R}'_{ji}) > 0$ for mismatched clusters. Due to this contradiction, the disambiguity criterion (11) must hold for all cluster pairs $i, j : i \neq j$. \square

This proposition provides a disambiguity criterion, requiring that matched clusters (i.e., C_{ii}, C_{jj}) should be more similar than mismatched clusters (i.e., C_{ij}, C_{ji}) up to some disambiguity *threshold* (in the case of $n \rightarrow \infty$, the threshold is 0). To extend this proposition to the finite-sample regime, we require a higher disambiguity-threshold to offset uncertainty due to finite samples.

2.2 Part 2: Disambiguity criterion in the finite-sample regime

We utilize a recent p -Wasserstein concentration bound by Weed and Bach [3] that describes finite sample behavior on the Wasserstein distance for data embedded in high-dimensional space, but whose clusters are themselves approximately low-dimensional. We will proceed our analysis with the language of probability measures μ to make our analysis consistent with [3]. We thus define the equivalent measure analogs as follows.

Definition 2.4. Let clusters $\mathbf{X}_i \in \mathbb{R}^{D \times n_{x,i}}$ and $\mathbf{Y}_j \in \mathbb{R}^{D \times n_{y,j}}$ be respectively denoted by empirical measures as

$$\hat{\mu}_{x,i} := \frac{1}{n_{x,i}} \sum_{k=1}^{n_{x,i}} \delta_{\mathbf{X}_i(k)}, \quad \hat{\mu}_{y,i} := \frac{1}{n_{y,j}} \sum_{k=1}^{n_{y,j}} \delta_{\mathbf{Y}_j(k)},$$

where $\delta_{\mathbf{X}_i(k)}$ refers to a discrete point located at $\mathbf{X}_i(k)$. At the limit, we denote the measures as

$$\mu_{x,i} := \lim_{n_{x,i} \rightarrow \infty} \frac{1}{n_{x,i}} \sum_{k=1}^{n_{x,i}} \delta_{\mathbf{X}_i(k)}, \quad \mu_{y,j} := \lim_{n_{y,j} \rightarrow \infty} \frac{1}{n_{y,j}} \sum_{k=1}^{n_{y,j}} \delta_{\mathbf{Y}_j(k)}.$$

Definition 2.5. Denote a linear transformation \mathbf{R} applied on the measure as

$$\mathbf{R} \circ \hat{\mu}_{x,i} := \frac{1}{n_{x,i}} \sum_{k=1}^{n_{x,i}} \delta_{\mathbf{R}\mathbf{X}_i(k)}, \quad \mathbf{R} \circ \mu_{x,i} := \lim_{n_{x,i} \rightarrow \infty} \frac{1}{n_{x,i}} \sum_{k=1}^{n_{x,i}} \delta_{\mathbf{R}\mathbf{X}_i(k)}.$$

The transformed inter-cluster distance between clusters may thus be denoted as

$$\hat{C}'_{ij}(\mathbf{R}) := \mathcal{W}_2^2(\mathbf{R} \circ \hat{\mu}_{x,i}, \hat{\mu}_{y,i}), \quad \tilde{C}'_{ij}(\mathbf{R}) := \mathcal{W}_2^2(\mathbf{R} \circ \mu_{x,i}, \mu_{y,i}).$$

Now, we may proceed to state results from [3]. The following result pertains to the sample complexity of measures μ in \mathbb{R}^D that are supported on an approximately low-dimensional set in \mathbb{R}^d , where $d \ll D$. First we require some definitions.

Definition 2.6. Given a set $S \subseteq \mathbb{R}^d$, let $\mathcal{N}_\varepsilon(S)$ denote the ε -covering number of set S , which is defined as the minimum number m of closed balls B_1, \dots, B_m of diameter ε such that $S \subseteq \bigcup_{1 \leq i \leq m} B_i$.

Definition 2.7. For any set $S \subseteq \mathbb{R}^d$, the ε -fattening of S is $S_\varepsilon := \{y : D(y, S) \leq \varepsilon\}$.

Proposition 2.8 (Weed and Bach [3], Proposition 16). Let S be a set that satisfies $\mathcal{N}_{\varepsilon'}(S) \leq (3\varepsilon')^{-d}$ for all $\varepsilon' \leq 1/27$ and for some $d > 2p$. Suppose there exists a positive constant σ such that μ satisfies $\mu(S_\varepsilon) \geq 1 - e^{-\varepsilon^2/2\sigma^2}$ for all $\varepsilon > 0$. If $p \log \frac{1}{\sigma} \geq 1/18$, then for all $n \leq (18p\sigma^2 \log \frac{1}{\sigma})^{-d/2}$,

$$\mathbb{E}[\mathcal{W}_p^p(\mu, \hat{\mu}_n)] \leq cn^{-p/d},$$

where $c = 27^p(2 + \frac{1}{3^{d/2-p-1}})$.

This proposition states that the degree that μ is concentrated (as parameterized by σ) around set S (approximately supported in low-dimensions) affects how ‘‘long’’ (in terms of n) we can enjoy the fast convergence rate of $n^{-p/d}$. We will leverage this result to obtain the following theorem on cluster correspondence disambiguity with respect to sample complexity.

Corollary 2.9. Let sets $S_{x,i}$ and $S_{y,j}$ satisfy the conditions for S in proposition 2.8 for some $\sigma_{x,i}, \sigma_{y,j} > 0$ and $d_{x,i}, d_{y,j} > 4$. If $\log \frac{1}{\sigma_{x,i}} \geq \frac{1}{36}$ and $\log \frac{1}{\sigma_{y,j}} \geq \frac{1}{36}$, then for all $n_{x,i} \leq (36\sigma_{x,i}^2 \log \frac{1}{\sigma_{x,i}})^{-d_{x,i}/2}$ and $n_{y,j} \leq (36\sigma_{y,j}^2 \log \frac{1}{\sigma_{y,j}})^{-d_{y,j}/2}$,

$$\mathbb{E}[|\tilde{C}_{ij}(\mathbf{R}) - \hat{C}_{ij}(\mathbf{R})|] \leq c_{x,i}n_{x,i}^{-2/d_{x,i}} + c_{y,j}n_{y,j}^{-2/d_{y,j}},$$

where

$$c_{z,k} = 729(2 + \frac{1}{3^{d_{z,k}/2-2} - 1}).$$

Proof. Denote μ_i, μ_j as measures and $\hat{\mu}_i, \hat{\mu}_j$ as their empirical estimates. By the triangle inequality,

$$\begin{aligned} \mathbb{E}[\mathcal{W}_2^2(\hat{\mu}_i, \hat{\mu}_j)] &\leq \mathbb{E}[\mathcal{W}_2^2(\hat{\mu}_i, \mu_i) + \mathcal{W}_2^2(\mu_j, \hat{\mu}_j)] \\ &\leq \mathbb{E}[\mathcal{W}_2^2(\hat{\mu}_i, \mu_i) + \mathcal{W}_2^2(\mu_i, \mu_j) + \mathcal{W}_2^2(\hat{\mu}_j, \mu_j)] \\ \Rightarrow \mathbb{E}[|\mathcal{W}_2^2(\mu_i, \mu_j) - \mathcal{W}_2^2(\hat{\mu}_i, \hat{\mu}_j)|] &\leq \mathbb{E}[\mathcal{W}_2^2(\hat{\mu}_i, \mu_i)] + \mathbb{E}[\mathcal{W}_2^2(\hat{\mu}_j, \mu_j)] \\ \Rightarrow \mathbb{E}[|\tilde{C}_{ij}(\mathbf{R}) - \hat{C}_{ij}(\mathbf{R})|] &\leq c_{x,i}n_{x,i}^{-2/d_{x,i}} + c_{y,j}n_{y,j}^{-2/d_{y,j}}, \end{aligned}$$

where the last line is a direct application of definition 2.5 and proposition 2.8. \square

Lemma 2.10. Let $S_{x,i}, S_{y,i}, S_{x,j}, S_{y,j}$ be sets that satisfy the conditions for S in proposition 2.8 for some $\sigma_{x,i}, \sigma_{y,i}, \sigma_{x,j}, \sigma_{y,j} > 0$ and $d_{x,i}, d_{y,i}, d_{x,j}, d_{y,j} > 4$. If $\log \frac{1}{\sigma_{x,i}}, \log \frac{1}{\sigma_{y,i}}, \log \frac{1}{\sigma_{x,j}}, \log \frac{1}{\sigma_{y,j}} \geq \frac{1}{36}$, then for all $n_{x,i} \leq (36\sigma_{x,i}^2 \log \frac{1}{\sigma_{x,i}})^{-d_{x,i}/2}$, $n_{y,i} \leq (36\sigma_{y,i}^2 \log \frac{1}{\sigma_{y,i}})^{-d_{y,i}/2}$, $n_{x,j} \leq (36\sigma_{x,j}^2 \log \frac{1}{\sigma_{x,j}})^{-d_{x,j}/2}$, and $n_{y,j} \leq (36\sigma_{y,j}^2 \log \frac{1}{\sigma_{y,j}})^{-d_{y,j}/2}$, cluster correspondences in problem (3) may be disambiguated to achieve $\mathbf{P}^* = \mathbf{I}$ when the following criterion is fulfilled for all $i, j \in [c] : i \neq j$:

$$\mathbb{E}[\min_{\mathbf{R}} \hat{C}_{ij}(\mathbf{R}) + \min_{\mathbf{R}} \hat{C}_{ji}(\mathbf{R}) - \min_{\mathbf{R}} \hat{C}_{ii}(\mathbf{R}) - \min_{\mathbf{R}} \hat{C}_{jj}(\mathbf{R})] > 2(B_{x,i} + B_{y,i} + B_{x,j} + B_{y,j}).$$

where the constants are defined as

$$B_{z,k} := c_{z,k}n_{z,k}^{-2/d_{z,k}}, \quad c_{z,k} := 729\left(2 + \frac{1}{3^{d_{z,k}/2-2} - 1}\right).$$

Proof. This follows a direct application of criterion (11) and corollary 2.9. \square

Proposition 2.11 (Weed and Bach [3], Proposition 20). For all $n \geq 0$ and $0 \leq p < \infty$,

$$\mathbb{P}[\mathcal{W}_p^p(\mu, \hat{\mu}_n) \geq \mathbb{E}\mathcal{W}_p^p(\mu, \hat{\mu}_n) + t] \leq \exp(-2nt^2).$$

Theorem 2.12. Define $\hat{C}_{ij}^* := \min_{\mathbf{R} \in \mathcal{S}(D,D)} \hat{C}_{ij}(\mathbf{R})$. If the conditions in Lemma 2.10 are satisfied then problem (3) yields the solution $\mathbf{P}^* = \mathbf{I}$ with probability at least $1 - \delta$ if, $\forall i, j : i \neq j$, the following criterion is satisfied:

$$\hat{C}_{ij}^* + \hat{C}_{ji}^* - \hat{C}_{ii}^* - \hat{C}_{jj}^* > 2(B_{x,i}(\delta) + B_{y,i}(\delta) + B_{x,j}(\delta) + B_{y,j}(\delta))$$

where

$$B_{z,k}(\delta) := c_{z,k} n_{z,k}^{-\frac{2}{d_{z,k}}} + \sqrt{\log(1/\delta)/2n_{z,k}}, \quad c_{z,k} := 729 \left(2 + \frac{1}{3^{d_{z,k}/2-2} - 1} \right), \quad (13)$$

where $d_{z,k}$ refers to the intrinsic dimension of the k -th cluster from the z -th dataset.

Proof. For some measure μ and its empirical finite-sample estimate $\hat{\mu}_n$, proposition 2.11 may be equivalently stated with the choice of $t = \sqrt{\log \frac{1}{\delta}/2n}$ as:

$$|\mathcal{W}_p^p(\mu, \hat{\mu}_n) - \mathbb{E}\mathcal{W}_p^p(\mu, \hat{\mu}_n)| \leq \sqrt{\log(1/\delta)/2n},$$

holds with at least probability $1 - \delta$. Under the conditions stated in proposition (2.8), and combining its result with the above relation, we have

$$\mathcal{W}_p^p(\mu, \hat{\mu}_n) \leq cn^{-p/d} + \sqrt{\log(1/\delta)/2n},$$

where $c = 27^p(2 + \frac{1}{3^{d/2-p}-1})$. Combining this for all terms in the left-hand side of (13) yields the stated result. \square

2.3 Putting everything together

The final proof of Theorem 4.1 is a simplified version of Theorem 2.12's since we assert a stronger (but cleaner) exact low-rank assumption to streamline communication. When the data is *exactly* supported in low-dimensions (as opposed to approximately), the ε -fattening disappears (i.e., $\varepsilon \rightarrow 0$) thus any positive $\sigma < \varepsilon$ will send $n \rightarrow \infty$, implying that the rapid convergence in dimensions $d \ll D$ holds for $n \rightarrow \infty$. Hence an identical result holds, with the sole condition that $d > 4$. \square

3 Proof of Theorem 4.2

We apply a very recent perturbation bound for the Procrustes problem developed by Arias-Castro et al. [4] to subsequently state a cluster-based alignment bound. First, we outline the perturbation bound for the classical Procrustes problem below.

Theorem 3.1 (Procrustes perturbation bounds, Theorem 1 [4]). *Consider short matrices $\mathbf{X}, \mathbf{Y} \in \mathbb{R}^{d \times n}$ with $d < n$ and \mathbf{X} having full rank. Set $\varepsilon^2 = \|\mathbf{Y}^\top \mathbf{Y} - \mathbf{X}^\top \mathbf{X}\|_p$, where $\|\cdot\|_p$ denotes the Schatten p -norm. Denote the singular value decomposition of $\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$, where $\mathbf{\Sigma}$ contains diagonal elements $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d > 0 = \dots = 0$, and let \mathbf{X}^\dagger be the pseudo-inverse of \mathbf{X} , i.e., $\mathbf{X}^\dagger = \mathbf{U}\mathbf{\Sigma}^\dagger\mathbf{V}^\top$, where $\mathbf{\Sigma}^\dagger = \text{diag}(\sigma_1^{-1}, \sigma_2^{-1}, \dots, \sigma_d^{-1}, 0, \dots, 0)$. If $\|\mathbf{X}^\dagger\| \varepsilon \leq \frac{1}{\sqrt{2}}(\|\mathbf{X}\| \|\mathbf{X}^\dagger\|)^{-1/2}$ then*

$$\min_{\mathbf{R} \in \mathcal{S}(d,d)} \|\mathbf{R}\mathbf{X} - \mathbf{Y}\|_p \leq (\|\mathbf{X}\| \|\mathbf{X}^\dagger\| + 2) \|\mathbf{X}^\dagger\| \varepsilon^2.$$

Directly applying Theorem 3.1 using the Schatten 2-norm (i.e., the Frobenius norm) yields

$$\begin{aligned} \min_{\mathbf{R} \in \mathcal{S}(d,d)} \|\mathbf{R}\mathbf{X} - \mathbf{Y}\|_F^2 &= \min_{\mathbf{R} \in \mathcal{S}(d,d)} \sum_{i=1}^c \text{tr}(\mathbf{Q}_{ii}^\top \mathbf{X}_i^\top \mathbf{X}_i \mathbf{Q}_{ii} + \mathbf{Y}_i^\top \mathbf{Y}_i - 2\mathbf{Y}_i \mathbf{Q}_{ii}^\top \mathbf{X}_i^\top \mathbf{R}^\top) \\ &\leq ((\|\mathbf{X}\| \|\mathbf{X}^\dagger\| + 2) \|\mathbf{X}^\dagger\| \varepsilon^2)^2 := B^2, \\ \Rightarrow \max_{\mathbf{R} \in \mathcal{S}(d,d)} \sum_{i=1}^c \text{tr}(2\mathbf{Y}_i \mathbf{Q}_{ii}^\top \mathbf{X}_i^\top \mathbf{R}^\top) &\geq \text{tr}(\mathbf{X}\mathbf{X}^\top + \mathbf{Y}\mathbf{Y}^\top) - B^2, \end{aligned}$$

where $\varepsilon^2 := \|\mathbf{Y}^\top \mathbf{Y} - \mathbf{X}^\top \mathbf{X}\|_F$ and under the conditions that \mathbf{X} is full rank and $\|\mathbf{X}^\dagger\| \varepsilon \leq \frac{1}{\sqrt{2}}(\|\mathbf{X}\| \|\mathbf{X}^\dagger\|)^{-1/2}$. When the criterion given by corollary 2.3 is satisfied, we are guaranteed

cluster-correspondences $\mathbf{P}^* = \mathbf{I}$. Therefore $\sum_{ij} P_{ij} C_{ij}(\mathbf{R}) = \sum_i C_{ii}(\mathbf{R})$. We utilize this to lower bound $\sum_i C_{ii}(\mathbf{R})$ as follows:

$$\begin{aligned}
& \min_{\mathbf{R} \in S(d,d)} \sum_{i=1}^c C_{ii}(\mathbf{R}) \\
&= \min_{\mathbf{R} \in S(d,d)} \sum_{i=1}^c \text{tr}(\mathbf{X}_i \text{diag}(\mathbf{Q}_{ii} \mathbb{1}) \mathbf{X}_i^\top + \mathbf{Y}_i \text{diag}(\mathbf{Q}_{ii}^\top \mathbb{1}) \mathbf{Y}_i^\top - 2\mathbf{Y}_i \mathbf{Q}_{ii}^\top \mathbf{X}_i^\top \mathbf{R}^\top) \\
&= \sum_{i=1}^c \text{tr}(\mathbf{X}_i \text{diag}(\mathbf{Q}_{ii} \mathbb{1}) \mathbf{X}_i^\top + \mathbf{Y}_i \text{diag}(\mathbf{Q}_{ii}^\top \mathbb{1}) \mathbf{Y}_i^\top) - \max_{\mathbf{R} \in S(d,d)} \sum_{i=1}^c \text{tr}(2\mathbf{Y}_i \mathbf{Q}_{ii}^\top \mathbf{X}_i^\top \mathbf{R}^\top) \\
&\leq B^2 + \sum_{i=1}^c \text{tr}(\mathbf{X}_i (\text{diag}(\mathbf{Q}_{ii} \mathbb{1}) - \mathbf{Q}_{ii} \mathbf{Q}_{ii}^\top) \mathbf{X}_i^\top + \mathbf{Y}_i (\text{diag}(\mathbf{Q}_{ii} \mathbb{1}) - \mathbf{I}) \mathbf{Y}_i^\top) \\
&= B^2 + \sum_{i=1}^c \text{tr}(\mathbf{X}_i (1/n - \mathbf{Q}_{ii} \mathbf{Q}_{ii}^\top) \mathbf{X}_i^\top + \mathbf{Y}_i (1/n - \mathbf{I}) \mathbf{Y}_i^\top).
\end{aligned}$$

□

4 Proof of Lemma 4.3

To ease notation, let $\mathbf{A} = \mathbf{Y}_j \mathbf{Q}_{ij}^\top \mathbf{X}_i^\top$, and $\mathbf{C} = \frac{1}{n} \mathbf{X}_i \mathbf{X}_i^\top + \frac{1}{n} \mathbf{Y}_j \mathbf{Y}_j^\top$. Let \mathbf{A} be decomposed by the singular-value decomposition as $\mathbf{A} = \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top$. Lastly, let $\mathbf{U} = [\mathbf{U}', \mathbf{U}']$ and $\mathbf{V} = [\mathbf{V}', \mathbf{V}']$, where $\langle \mathbf{U}'', \mathbf{U}' \rangle = \langle \mathbf{V}'', \mathbf{V}' \rangle = 0$. Then it follows that

$$\begin{aligned}
& \text{tr}(\mathbf{U}^\top \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top \mathbf{V}) \leq \text{tr}(\tilde{\Sigma}) = \text{tr}(\tilde{\mathbf{U}}^\top \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top \tilde{\mathbf{V}}) \quad (14) \\
& \text{tr}(\mathbf{U}^\top \mathbf{A} \mathbf{V}) \leq \text{tr}(\tilde{\mathbf{U}}^\top \mathbf{A} \tilde{\mathbf{V}}) \\
& \text{tr}(\mathbf{C} - 2\mathbf{U}'^\top \mathbf{A} \mathbf{V}' - 2\mathbf{U}''^\top \mathbf{A} \mathbf{V}'') \geq \text{tr}(\mathbf{C} - 2\tilde{\mathbf{U}}^\top \mathbf{A} \tilde{\mathbf{V}}) \\
& \min_{\substack{\mathbf{U}'', \mathbf{V}'' \in S(d,d-r): \\ \langle \mathbf{U}', \mathbf{U}'' \rangle = \langle \mathbf{V}', \mathbf{V}'' \rangle = 0}} \text{tr}(\mathbf{C} - 2\mathbf{A} \mathbf{V}' \mathbf{U}'^\top - 2\mathbf{A} \mathbf{V}'' \mathbf{U}''^\top) \geq \min_{\tilde{\mathbf{U}}, \tilde{\mathbf{V}} \in S(d,d)} \text{tr}(\mathbf{C} - 2\tilde{\mathbf{A}} \tilde{\mathbf{U}} \tilde{\mathbf{V}}^\top) \\
& \min_{\mathbf{R} \in \mathcal{T}(\mathbf{U}', \mathbf{V}')} C_{ij}(\mathbf{R}) \geq \min_{\mathbf{R} \in S(d,d)} C_{ij}(\mathbf{R}).
\end{aligned}$$

What remains is for us to show the condition for equality. From (14), we have that

$$\text{tr}(\mathbf{U}^\top \tilde{\mathbf{U}} \tilde{\Sigma} \tilde{\mathbf{V}}^\top \mathbf{V}) = \text{tr}(\tilde{\mathbf{V}}^\top \mathbf{V} \mathbf{U}^\top \tilde{\mathbf{U}} \tilde{\Sigma}) \leq \text{tr}(\tilde{\Sigma}),$$

with equality holding if $\tilde{\mathbf{V}}^\top \mathbf{V} \mathbf{U}^\top \tilde{\mathbf{U}} = \mathbf{I}$, implying that $\tilde{\mathbf{U}}^\top \mathbf{U} = \tilde{\mathbf{V}}^\top \mathbf{V}$, which imply that

$$\langle \tilde{\mathbf{U}}, \mathbf{U}' \rangle = \langle \tilde{\mathbf{V}}, \mathbf{V}' \rangle, \quad \langle \tilde{\mathbf{U}}, \mathbf{U}'' \rangle = \langle \tilde{\mathbf{V}}, \mathbf{V}'' \rangle,$$

which are obtained via the substitutions $\mathbf{U} = [\mathbf{U}', \mathbf{U}']$ and $\mathbf{V} = [\mathbf{V}', \mathbf{V}']$. □

References

- [1] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. In *Advances in neural information processing systems*, pages 2292–2300, 2013.
- [2] John Von Neumann. A certain zero-sum two-person game equivalent to the optimal assignment problem. *Contributions to the Theory of Games*, 2(5-12):7, 1953.
- [3] Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance. *arXiv preprint arXiv:1707.00087*, 2017.
- [4] Ery Arias-Castro, Adel Javanmard, and Bruno Pelletier. Perturbation bounds for procrustes, classical scaling, and trilateration, with applications to manifold learning. *arXiv preprint arXiv:1810.09569*, 2018.