

# Robust Point Matching with Distance Profiles

**YoonHaeng Hur**

*Department of Statistics  
University of Chicago  
Chicago, IL 60637, USA*

YOONHAENGHUR@UCHICAGO.EDU

**Yuehaw Khoo**

*Department of Statistics  
University of Chicago  
Chicago, IL 60637, USA*

YKHOO@UCHICAGO.EDU

**Editor:** Jianfeng Lu

## Abstract

Computational difficulty of quadratic matching and the Gromov-Wasserstein distance has led to various approximation and relaxation schemes. One of such methods, relying on the notion of distance profiles, has been widely used in practice, but its theoretical understanding is limited. By delving into the statistical complexity of the previously proposed method based on distance profiles, we show that it suffers from the curse of dimensionality unless we make certain assumptions on the underlying metric measure spaces. Building on this insight, we propose and analyze a modified matching procedure that can be used to robustly match points under a certain probabilistic setting. We demonstrate the performance of the proposed methods using simulations and real data applications to complement the theoretical findings. As a result, we contribute to the literature by providing theoretical underpinnings of the matching procedures based on distance invariants like distance profiles, which have been widely used in practice but rarely analyzed theoretically.

**Keywords:** point matching, robust matching, distance profiles, Gromov-Wasserstein, invariant features

## 1. Introduction

Matching is one of the most fundamental tasks of modern data science, which can be found in numerous applications. For instance, the goal of shape correspondence problems (Van Kaick et al., 2011; Sahillioğlu, 2020) in computer graphics is to match corresponding parts of different geometric shapes, often obtained by 3D scanners or imaging equipment used for Magnetic Resonance Imaging (MRI) or Computed Tomography (CT). Graph matching aims to find corresponding nodes between two or more graphs representing structured data, such as neuronal connectivity in human brains (Sporns et al., 2005) or social networks (Borgatti et al., 2009).

One formal mathematical framework for matching (Mémoli, 2011) is to conceptualize objects as metric measure spaces, which are essentially probability measures defined on suitable metric spaces. In practice, of course, objects are not fully available and thus approximated by suitable point clouds, say,  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_m\}$  on metric spaces  $(\mathcal{X}, d_{\mathcal{X}})$  and  $(\mathcal{Y}, d_{\mathcal{Y}})$ , respectively, and the practical goal is to match these point clouds.

Quadratic matching is one of the most standard approaches to this problem, which formulates and solves a quadratic program based on pairwise distances  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{i,\ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{j,\ell=1}^m$ . When  $n = m$ , quadratic matching finds a one-to-one correspondence that aligns the pairwise distances as closely as possible, which is formulated as a Quadratic Assignment Problem (QAP) (Koopmans and Beckmann, 1957; Loiola et al., 2007). When  $n \neq m$ , the notion of correspondence is relaxed to soft assignments or couplings, leading to the formulation of the Gromov-Wasserstein distance (Mémoli, 2011), which has been increasingly used in applications, such as word alignment (Alvarez-Melis and Jaakkola, 2018), statistical inference (Brécheteau, 2019; Weitkamp et al., 2022), and generative modeling (Bunne et al., 2019; Hur et al., 2024). On the practical side, however, quadratic matching requires solving a quadratic program that is NP-hard in general, and thus is approximated by several methods, such as convex relaxation (Affalo et al., 2015), semidefinite relaxation (Zhao et al., 1998; Kezurer et al., 2015), spectral relaxation (Leordeanu and Hebert, 2005), or entropic regularization (Rangarajan et al., 1999; Solomon et al., 2016; Peyré et al., 2016).

Matching procedures based on pairwise distances  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{i,\ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{j,\ell=1}^m$  are invariant to isometries, which is desirable in practice. For instance, in imaging problems, the images of the objects of interest are often obtained from different angles and thus need to be matched up to rigid transformations. However, such procedures are difficult to analyze theoretically and the literature has been mostly studying theoretical guarantees for other matching procedures based on  $\|X_i - Y_j\|_2$ 's when  $(\mathcal{X}, d_{\mathcal{X}}) = (\mathcal{Y}, d_{\mathcal{Y}})$  is the standard Euclidean space. In this setting, the point clouds lie in the same space and are often viewed as perturbations of one another, where one aims to match the points based on the proximity. Collier and Dalalyan (2016); Kunisky and Niles-Weed (2022); Wang et al. (2022) establish theoretical results on recovering the ground truth matching for the procedures based on  $\|X_i - Y_j\|_2$ 's under some noise, and Galstyan et al. (2022); Minasyan et al. (2023); Galstyan and Minasyan (2023) show matching guarantees in the presence of outliers. Although Theorem 2 of Wang et al. (2022) concerns a procedure using the pairwise distances in the Euclidean setting, it requires solving a nontrivial minimization over permutations, which is computationally challenging, and the result does not consider the presence of outliers.

Motivated by the above challenges, this paper theoretically analyzes matching procedures based on the distances  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{i,\ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{j,\ell=1}^m$ , which are easily implementable in practice. The idea is to associate each point with the set of distances from it to other points, namely,  $X_i$  and  $Y_j$  are associated with the sets of distances to the other points  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{\ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{\ell=1}^m$ , which are the distance profiles of  $X_i$  and  $Y_j$ , respectively. The matching procedure is to find for each  $X_i$ , the point  $Y_j$  having the most similar distance profile to that of  $X_i$ , where the similarity is measured by a suitable quantity between the sets  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{\ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{\ell=1}^m$ . The idea of associating a point with suitable invariants based on the distances to other points from it has a long history in the object matching literature, which includes shape contexts (Belongie et al., 2002; Ruggeri and Saupe, 2008), integral invariants (Manay et al., 2006), invariant histograms (Brinkman and Olver, 2012), the local distribution of distances (Mémoli, 2011), and other related concepts (Gortler et al., 2019; Norelli et al., 2023). However, theoretical results on robustness of such methods are rarely found in the literature. In the context of random graph matching, Ding et al. (2021, 2023) propose and analyze a degree profile matching procedure based on the adjacency matrices  $(A_{i\ell})_{i,\ell=1}^n$  and  $(B_{j\ell})_{j,\ell=1}^m$  of certain correlated random graphs.

However, their analysis focuses on random graph models, such as Erdős-Rényi and correlated Wigner models, where the edges are independently generated, namely,  $(A_{i\ell})_{i,\ell=1}^n$  are mutually independent, which is substantially different from our setting where the distances  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{i,\ell=1}^n$  are not mutually independent.

**Organization** Section 2 presents mathematical background on the Gromov-Wasserstein distance and its computational alternative based on the distance profile, together with the new sample complexity result in Section 2.3. Motivated by the insights built in Section 2, Section 3 proposes and studies the partial matching algorithm based on distance profiles, where we formally show its robustness to noise and outlier components, together with the noise stability analysis for the one-to-one matching procedure. Section 4 demonstrates the performance of the proposed methods using simulations. All technical proofs, lemmas, extensions are deferred to Appendices A, B, C, respectively, to streamline the main text.

**Notation** For  $k \in \mathbb{N}$ , let  $[k]$  denote  $\{1, \dots, k\}$ ,  $\Delta_k = \{(p_1, \dots, p_k) \in \mathbb{R}_+^k : \sum_{i=1}^k p_i = 1\}$  denote the probability simplex, and  $\mathcal{S}_k$  denote the set of all permutations on  $[k]$ . For  $a, b \in \mathbb{R}$ , let  $a \vee b = \max\{a, b\}$  and  $a \wedge b = \min\{a, b\}$ . Let  $\mathcal{P}(\mathbb{R}^d)$  denote the set of all Borel probability measures defined on  $\mathbb{R}^d$  and for any  $z \in \mathbb{R}^d$ , let  $\delta_z \in \mathcal{P}(\mathbb{R}^d)$  denote the Dirac measure at  $z$ . We denote by  $W_p$  the Wasserstein- $p$  distance for  $p \geq 1$ . For any measure  $\mu$  on a measurable space  $\mathcal{X}$  and a measurable function  $f: \mathcal{X} \rightarrow \mathbb{R}$ , let  $\mu f = \int_{\mathcal{X}} f d\mu$ .

## 2. Foundations: Distance Profile Distance and Sample Complexity

As mentioned in Section 1, computational difficulty of the Gromov-Wasserstein (GW) distance serves as a major obstacle in practical applications, which has led to the development of various approximation and relaxation schemes. One of such methods, proposed as a tractable lower bound to the GW distance by Mémoli (2011), utilizes the notion of distance profiles. We delineate the mathematical foundations of the distance profile in connection with the GW distance and see how the distance based on it, which we call the distance profile distance (DPD), avoids the computational difficulty of the GW distance. Then, we derive new results on the sample complexity of the DPD and discuss the limitations and potentials of the DPD as a matching method under certain probabilistic settings, which motivates the distance profile matching procedure in the next section.

### 2.1 Background: Gromov-Wasserstein Distance and Computational Difficulty

Following Mémoli (2011), we build on the mathematical framework of object matching based on metric measure spaces. The Gromov-Wasserstein (GW) distance is a generalization of quadratic matching between point clouds to abstract metric measure spaces. The GW distance is defined based on the notion of couplings, which are probability measures on the product space of two metric measure spaces that preserve the marginal distributions of the two spaces. Couplings represent the correspondences between the points in the two metric measure spaces, and the GW distance finds the optimal coupling that best aligns the pairwise distances in the two spaces. Formal definitions are as follows.

**Definition 1** We call  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  a metric measure space if  $\mu$  is a Borel probability measure on a complete and separable metric space  $(\mathcal{X}, d_{\mathcal{X}})$ .

**Definition 2** Let  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  and  $(\mathcal{Y}, d_{\mathcal{Y}}, \nu)$  be two metric measure spaces. We call a Borel probability measure  $\gamma$  on  $\mathcal{X} \times \mathcal{Y}$  a coupling if  $\gamma(A \times \mathcal{Y}) = \mu(A)$  and  $\gamma(\mathcal{X} \times B) = \nu(B)$  for any Borel sets  $A \subset \mathcal{X}$  and  $B \subset \mathcal{Y}$ . We denote the set of all couplings by  $\Pi(\mu, \nu)$ .

**Definition 3** The Gromov-Wasserstein distance of order  $p \geq 1$  between two metric measure spaces  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  and  $(\mathcal{Y}, d_{\mathcal{Y}}, \nu)$  is defined as follows:

$$\text{GW}_p(\mu, \nu) := \left( \inf_{\gamma \in \Pi(\mu, \nu)} \int_{(\mathcal{X} \times \mathcal{Y})^2} |d_{\mathcal{X}}(x, x') - d_{\mathcal{Y}}(y, y')|^p d\gamma(x, y) d\gamma(x', y') \right)^{1/p}. \quad (1)$$

The objective function on the right-hand side of (1) is quadratic in the variable  $\gamma$  and is not convex in general, which makes the minimization computationally challenging. The following discrete case illustrates the computational difficulty of the GW distance in the case of point clouds, which is a common setting in practice.

**Example 1 (Discrete Case)** Consider the point clouds  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_m\}$  on complete and separable metric spaces  $(\mathcal{X}, d_{\mathcal{X}})$  and  $(\mathcal{Y}, d_{\mathcal{Y}})$ , respectively. We can define the GW distance between these point clouds as the GW distance between the metric measure spaces  $(\mathcal{X}, d_{\mathcal{X}}, \hat{\mu})$  and  $(\mathcal{Y}, d_{\mathcal{Y}}, \hat{\nu})$ , where  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$  and  $\hat{\nu} = \frac{1}{m} \sum_{j=1}^m \delta_{Y_j}$  are the corresponding empirical measures. In this case,

$$\text{GW}_p^p(\hat{\mu}, \hat{\nu}) = \min_{\gamma \in \Pi_{n,m}} \sum_{i,k=1}^n \sum_{j,\ell=1}^m |d_{\mathcal{X}}(X_i, X_k) - d_{\mathcal{Y}}(Y_j, Y_{\ell})|^p \gamma_{ij} \gamma_{k\ell}, \quad (2)$$

where  $\Pi_{n,m} = \{\gamma \in \mathbb{R}_+^{n \times m} : \sum_{j=1}^m \gamma_{ij} = \frac{1}{n} \forall i \in [n] \text{ and } \sum_{i=1}^n \gamma_{ij} = \frac{1}{m} \forall j \in [m]\}$  represents the set of all couplings between  $\hat{\mu}$  and  $\hat{\nu}$  as nonnegative matrices—called coupling matrices—of which the sum of each row (resp. column) is  $\frac{1}{n}$  (resp.  $\frac{1}{m}$ ). We again observe that the objective function of (2) is quadratic in the variable  $\gamma$ , which is NP-hard to solve in general.

## 2.2 Distance Profile Distance

Due to the aforementioned computational difficulty of the GW distance, various approximation and relaxation schemes have been proposed in the literature. Mémoli (2011) proposes several relaxations to the GW distance, which provide lower bounds to the GW distance. One of the three lower bounds proposed in Mémoli (2011), which he named the third lower bound based on the ordering, is based on the idea of local distribution of distances, which we generalize as distance profiles below.

**Definition 4** Let  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  be a metric measure space. For  $x \in \mathcal{X}$ , let  $\mu_x$  be the probability measure on  $\mathbb{R}$  that represents the distribution of  $d_{\mathcal{X}}(x, X)$ , where  $X$  is a  $\mathcal{X}$ -valued random variable whose law is  $\mu$ . We call  $\mu_x$  the distance profile of  $\mu$  at point  $x$ .

In a nutshell, for each  $x \in \mathcal{X}$ , the distance profile  $\mu_x$  is the distribution of the distances from  $x$  to the points in  $\mathcal{X}$  with respect to the measure  $\mu$ . This leads to a natural matching procedure as follows: given two metric measure spaces  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  and  $(\mathcal{Y}, d_{\mathcal{Y}}, \nu)$ , we find a coupling that best aligns the distance profiles, namely, associate  $x \in \mathcal{X}$  and  $y \in \mathcal{Y}$  if the

distance profiles  $\mu_x, \nu_y \in \mathcal{P}(\mathbb{R})$  are similar. Here, similarity is measured by any suitable discrepancy on  $\mathcal{P}(\mathbb{R})$ ; we use the Wasserstein distance as a natural choice, which leads to the following definition.

**Definition 5** *The distance profile distance (DPD) of order  $p \geq 1$  between metric measure spaces  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  and  $(\mathcal{Y}, d_{\mathcal{Y}}, \nu)$  is defined as follows:*

$$\mathfrak{T}_p(\mu, \nu) := \left( \inf_{\gamma \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} W_p^p(\mu_x, \nu_y) d\gamma(x, y) \right)^{1/p}. \quad (3)$$

**Remark 6** *In Definition 5, while we use  $W_p^p$  to measure the discrepancy between the distance profiles  $\mu_x, \nu_y \in \mathcal{P}(\mathbb{R})$ , any suitable discrepancy measure on  $\mathcal{P}(\mathbb{R})$  can be used in place of  $W_p^p$ . The choice  $W_p^p$  leads to a nice metric property, namely,  $\mathfrak{T}_p$  is a pseudometric on the set of metric measure spaces, and also naturally connects to the Gromov-Wasserstein distance. Indeed, one can show that  $\mathfrak{T}_p \leq \text{GW}_p$ . See Mémoli (2011) and Mémoli and Needham (2022) for technical details.*

Unlike the GW distance, the objective function on the right-hand side of (3) is linear in the variable  $\gamma$ , which makes the minimization computationally tractable. The following discrete case illustrates the computational advantage of the distance profile distance in the case of point clouds.

**Example 2** *Consider the setting of Example 1. Again, we can define the distance profile distance between the point clouds  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_m\}$  as the distance profile distance  $\mathfrak{T}_p(\hat{\mu}, \hat{\nu})$ . In this case, the distance profile of  $\hat{\mu}$  at  $x \in \mathcal{X}$  is  $\hat{\mu}_x = \frac{1}{n} \sum_{\ell=1}^n \delta_{d_{\mathcal{X}}(x, X_{\ell})}$ . Distance profiles of  $\hat{\nu}$  are defined analogously. Then, one can see that the DPD is given as*

$$\mathfrak{T}_p^p(\hat{\mu}, \hat{\nu}) = \min_{\gamma \in \Pi_{n,m}} \sum_{i=1}^n \sum_{j=1}^m W_p^p(\hat{\mu}_{X_i}, \hat{\nu}_{Y_j}) \gamma_{ij}. \quad (4)$$

where  $\Pi_{n,m}$  is as defined in Example 1. Computational advantage of  $\mathfrak{T}_p$  over  $\text{GW}_p$  is clear:  $\mathfrak{T}_p$  can be computed by solving a linear program, while  $\text{GW}_p$  requires solving a non-convex quadratic program. Unlike (2), the minimization (4) is a linear program that is solvable by polynomial-time algorithms.

### 2.3 Sample Complexity: Challenges and Opportunities

In practice, we have access to point clouds  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_m\}$  that approximate some underlying objects represented as metric measure spaces  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  and  $(\mathcal{Y}, d_{\mathcal{Y}}, \nu)$ , respectively. Hence, the quantity  $\mathfrak{T}_p(\hat{\mu}, \hat{\nu})$  in Example 2 is an approximation of  $\mathfrak{T}_p(\mu, \nu)$ . An important statistical question is the rate of convergence of  $\mathfrak{T}_p(\hat{\mu}, \hat{\nu})$  to  $\mathfrak{T}_p(\mu, \nu)$  in terms of the sample sizes  $n$  and  $m$  when the point clouds are independent samples from  $\mu$  and  $\nu$ . Faster convergence rates are desirable as they indicate that the distance profile distance is a good approximation of the population version  $\mathfrak{T}_p(\mu, \nu)$  with a small sample size, which can further expedite the computation of the distance profile distance in practice.

It turns out, however, that the convergence rate of  $\mathfrak{T}_p(\hat{\mu}, \hat{\nu})$  to  $\mathfrak{T}_p(\mu, \nu)$  can be slow in the high-dimensional setting, just as the standard Wasserstein distance  $W_p$  suffers from

the curse of dimensionality when estimated from samples (Dudley, 1969; Dereich et al., 2013; Boissard and Le Gouic, 2014; Fournier and Guillin, 2015; Weed and Bach, 2019). We derive this result for  $p = 1$  as the leading case, focusing on  $\mathbb{E}|\mathfrak{T}_1(\mu, \nu) - \mathfrak{T}_1(\hat{\mu}, \hat{\nu})|$ . While the result is pessimistic in its full generality, we find that a better rate of convergence can be achieved when the underlying metric spaces have certain structures, controlled by the effective dimension of the metric measure space that can be much smaller than the ambient dimension.

**Setting** Throughout this section, we mainly consider the setting of Example 2 where the point clouds are viewed as independent samples from the underlying metric measure spaces. It suffices to analyze the rate of convergence of  $\mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu})$  due to the triangle inequality:  $\mathbb{E}|\mathfrak{T}_1(\mu, \nu) - \mathfrak{T}_1(\hat{\mu}, \hat{\nu})| \leq \mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu}) + \mathbb{E}\mathfrak{T}_1(\nu, \hat{\nu})$ . Hence, we focus on  $\mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu})$ .

### 2.3.1 GENERAL CASE: CURSE OF DIMENSIONALITY

We first observe that

$$\begin{aligned} \mathfrak{T}_1(\mu, \hat{\mu}) &= \inf_{\gamma \in \Pi(\mu, \hat{\mu})} \int_{\mathcal{X} \times \mathcal{X}} W_1(\mu_x, \hat{\mu}_{x'}) \, d\gamma(x, x') \\ &\leq \inf_{\gamma \in \Pi(\mu, \hat{\mu})} \int_{\mathcal{X} \times \mathcal{X}} (W_1(\mu_x, \mu_{x'}) + W_1(\mu_{x'}, \hat{\mu}_{x'})) \, d\gamma(x, x') \\ &\leq \underbrace{\inf_{\gamma \in \Pi(\mu, \hat{\mu})} \int_{\mathcal{X} \times \mathcal{X}} W_1(\mu_x, \mu_{x'}) \, d\gamma(x, x')}_{(*)} + \underbrace{\int_{\mathcal{X}} W_1(\mu_{x'}, \hat{\mu}_{x'}) \, d\hat{\mu}(x')}_{(**)}. \end{aligned} \tag{5}$$

It turns out that the expectation of  $(**)$  enjoys a fast convergence rate. The idea is that

$$\mathbb{E}(**) = \mathbb{E} \left( \frac{1}{n} \sum_{i=1}^n W_1(\mu_{X_i}, \hat{\mu}_{X_i}) \right) \leq \mathbb{E} \left( \sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) \right), \tag{6}$$

where the last term can be shown to converge at the rate of  $O(n^{-1/2})$  as stated below.

**Theorem 7** *Let  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  be a metric measure space and  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$  be the empirical measure based on  $X_1, \dots, X_n$  that are i.i.d. from  $\mu$ . Suppose  $\mathcal{X}$  is a closed subset of  $\mathbb{R}^d$  and  $d_{\mathcal{X}}$  is the standard Euclidean distance. Then,*

$$\mathbb{E} \left[ \sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) \right] \leq C \sqrt{\frac{d+1}{n}}, \tag{7}$$

where  $C$  is an absolute constant that is independent of  $n$ ,  $d$ , and  $\mu$ .

**Remark 8** *For each fixed  $x \in \mathcal{X}$ , as  $\mu_x$  is a probability measure on  $\mathbb{R}$ , one can deduce from the standard results (Dudley, 1969; Fournier and Guillin, 2015) that  $W_1(\mu_x, \hat{\mu}_x)$  enjoys a fast rate of convergence in terms of the sample size  $n$ , which is  $O(n^{-1/2})$ . Theorem 7 shows that taking the supremum over  $x \in \mathcal{X}$  does not slow down the rate of convergence. We present the extension of Theorem 7 to general metric measure spaces in Appendix C.1.*

While Theorem 7 shows fast convergence of (\*\*), the term (\*) in (5) is the main bottleneck for the rate of convergence of  $\mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu})$ . Unfortunately, the rate of convergence of the expectation of (\*) can be as slow as that of  $\mathbb{E}W_1(\mu, \hat{\mu})$ , which is known to be  $O(n^{-1/d})$  for  $\mathcal{X} \subset \mathbb{R}^d$ . To see this, deduce (\*)  $\leq W_1(\mu, \hat{\mu})$  from Proposition 9.

**Proposition 9** *Let  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  be a metric measure space. Then,  $W_1(\mu_x, \mu_{x'}) \leq d_{\mathcal{X}}(x, x')$  for  $x, x' \in \mathcal{X}$ .*

In summary, despite the fast convergence rate of the term (\*\*) in (5), the term (\*) in (5) can lead to a slow rate of convergence of  $\mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu})$  as  $O(n^{-1/d})$ , thereby suffering from the curse of dimensionality.

### 2.3.2 FASTER RATE BY INVARIANCE AND EFFECTIVE DIMENSION

It turns out that when  $\mu$  has a certain structure, we can derive a faster rate of convergence for (\*). Namely, if  $\mu$  is a mixture of rotationally invariant distributions like a Gaussian mixture, then the effective dimension of  $\mu$  is the number of the components, which can be much smaller than the ambient dimension  $d$ . The key is that  $W_1(\mu_x, \mu_{x'})$ , the distance between distance profiles, captures this structure and thus is determined by the effective dimension, rather than the ambient dimension.

**Proposition 10** *Suppose  $\mu$  is a probability measure on  $\mathbb{R}^d$  given as a mixture of rotationally invariant distributions, namely,  $\mu = \sum_{k=1}^t p_k \mu_k$ , where  $(p_1, \dots, p_t) \in \Delta_t$ , and  $\mu_1, \dots, \mu_t \in \mathcal{P}(\mathbb{R}^d)$  are rotationally invariant with respect to the centers  $\theta_1, \dots, \theta_t \in \mathbb{R}^d$ , respectively. Then, for any  $x, x' \in \mathbb{R}^d$ ,*

$$W_1(\mu_x, \mu_{x'}) \leq \sum_{k=1}^t p_k \left| \|x - \theta_k\|_2 - \|x' - \theta_k\|_2 \right|. \quad (8)$$

Combining Proposition 10 with Theorem 7, we can derive the following convergence rate of  $\mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu})$  depending on the effective dimension  $t$  of  $\mu$ .

**Theorem 11** *Suppose  $\mu$  is a probability measure on  $\mathbb{R}^d$  given as a mixture of rotationally invariant distributions, namely,  $\mu = \sum_{k=1}^t p_k \mu_k$ , where  $(p_1, \dots, p_t) \in \Delta_t$ , and  $\mu_1, \dots, \mu_t \in \mathcal{P}(\mathbb{R}^d)$  are rotationally invariant with respect to the centers  $\theta_1, \dots, \theta_t \in \mathbb{R}^d$ , respectively. Let  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$  be the empirical measure based on  $X_1, \dots, X_n$  that are i.i.d. from  $\mu$ . Then,*

$$\mathbb{E}\mathfrak{T}_1(\mu, \hat{\mu}) \leq O(n^{-1/t}) + C \sqrt{\frac{d+1}{n}},$$

where  $C$  is an absolute constant that is independent of  $n$ ,  $d$ , and  $\mu$ .

**Summary and Insights** The above results provide insights into how to use distance profiles for matching. Due to the curse of dimensionality, it is impractical to estimate the distance profile distance  $\mathfrak{T}_1$  between high-dimensional metric measure spaces by random samples. Hence, it is more suitable to use and study  $\mathfrak{T}_1$  under a different setting than the sampling scenario; the next section conducts such a study based on the model with

fixed high-dimensional locations corrupted by noise. Meanwhile, we have seen that a better rate can be achieved under a mixture model with rotationally invariant components. The high-level idea is that when such a structure is present, points from the same component are treated as the same point so that the effective dimension is essentially the number of components. This is pertinent to partial matching, where we want to allow similar points from one object to be matched to the same point in another object, not necessarily requiring a one-to-one correspondence or a coupling. In the next section, we study the distance profile matching procedure that allows for such flexibility, with a robustness property.

### 3. Distance Profile Matching and Its Theoretical Guarantees

This section introduces and analyzes the distance profile matching algorithms, inspired by the insights on the distance profile distance  $\mathfrak{T}_1$  in the previous section. We have learned that using  $\mathfrak{T}_1$  directly for matching in the random sampling setting can be impractical. Hence, we propose a flexible matching procedure—Algorithm 1—modified from  $\mathfrak{T}_1$  that allows for partial matching and derive its robustness property. Meanwhile, Algorithm 2 is a special case of  $\mathfrak{T}_1$  to find a one-to-one correspondence, which we analyze under a fixed-location model with noise, deviating from the random sampling setting where  $\mathfrak{T}_1$  is impractical.

#### 3.1 Distance Profile Matching Algorithms

##### 3.1.1 DISTANCE PROFILE MATCHING ALGORITHM FOR PARTIAL MATCHING

Algorithm 1 follows the idea of  $\mathfrak{T}_1$  in that it compares the distance profiles of two point clouds, but it does not require finding an optimal coupling which was the main bottleneck of the sample complexity of  $\mathfrak{T}_1$  in the random sampling setting. Instead, for each point  $X_i$ , we find  $Y_j$  whose distance profile is closest to the distance profile of  $X_i$  under  $W_1$  as shown in Step 2 of Algorithm 1. Here, note that  $W(i, j) = W_1(\hat{\mu}_{X_i}, \hat{\nu}_{Y_j})$ , where  $\hat{\mu}_{X_i}$  and  $\hat{\nu}_{Y_j}$  are the distance profiles based on the empirical measures introduced in Example 2. Therefore, the key difference from  $\mathfrak{T}_1$  is that instead of finding a coupling  $\gamma$  as in (4), we find  $\pi$  that matches each  $X_i$  to a point  $Y_{\pi(i)}$  closes to  $X_i$  in terms of the distance profiles.

---

#### Algorithm 1 Distance Profile Matching

---

**Require:** Pairwise distances  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{i,\ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{j,\ell=1}^m$  from two point clouds  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_m\}$  on metric spaces  $(\mathcal{X}, d_{\mathcal{X}})$  and  $(\mathcal{Y}, d_{\mathcal{Y}})$ , respectively.

**Require:** (Optional) A threshold  $\rho > 0$ . Set  $\rho = \infty$  if not provided.

1: For each pair  $(i, j) \in [n] \times [m]$ , define

$$W(i, j) := W_1 \left( \frac{1}{n} \sum_{\ell=1}^n \delta_{d_{\mathcal{X}}(X_i, X_\ell)}, \frac{1}{m} \sum_{\ell=1}^m \delta_{d_{\mathcal{Y}}(Y_j, Y_\ell)} \right). \quad (9)$$

2: Define a map  $\pi: [n] \rightarrow [m]$  by  $\pi(i) = \arg \min_{j \in [m]} W(i, j)$ .

3: Define  $I \subset [n]$  by  $I = \{i \in [n] : W(i, \pi(i)) < \rho\}$ .

4: **return**  $\pi$  and  $I$ .

---

Note that  $\pi$  may not be injective or surjective, meaning that some points may not be matched or some points may be matched to multiple points. The motivation behind this is simple: in partial matching, (1) we want some points to be ignored or matched arbitrarily, and (2) we may allow similar points to be matched to the same point. The underlying objects that the points  $X_i$ 's and  $Y_j$ 's approximate, say, probability measures  $\mu$  and  $\nu$  defined on  $\mathbb{R}^d$ , respectively, are not identical in practice. Usually, they coincide up to several outlier components, which we can formally state as  $\mu = \lambda\rho + (1 - \lambda)\mu_o$  and  $\nu = \lambda\rho + (1 - \lambda)\nu_o$  for some  $\lambda \in (0, 1)$  and  $\rho, \mu_o, \nu_o \in \mathcal{P}(\mathbb{R}^d)$ . Here,  $\lambda\rho$  represents the common part of  $\mu$  and  $\nu$  that we want to match, while  $(1 - \lambda)\mu_o$  and  $(1 - \lambda)\nu_o$  are the outlier components of  $\mu$  and  $\nu$ , respectively. We are mainly interested in matching  $X_i$ 's and  $Y_j$ 's from the common part, while ignoring or arbitrarily matching the points from the outlier components.

The threshold  $\rho$  is an optional parameter that can be used to obtain the subset  $I \subset [n]$  of indices  $i$ 's such that the discrepancy  $W(i, \pi(i))$  is below  $\rho$ . Small  $\rho$  leads to a more conservative matching that only focuses on the points that are closely matched, which can be used for other downstream tasks. For instance, one can solve the orthogonal Procrustes problem (Wahba, 1965; Yang and Carlone, 2019) using the matched pairs  $\{(X_i, Y_{\pi(i)})\}_{i \in I}$  to find the optimal rotation and translation aligning them as closely as possible, which theoretically works with just  $|I| = d$  pairs in  $\mathbb{R}^d$ , or  $|I| = O(d)$  due to potential noise or degeneracy in practice. The obtained transformation can then be used as an initialization of the registration procedures, such as the Iterative Closest Point (ICP) algorithm (Besl and McKay, 1992), to find the optimal transformation that aligns the whole point clouds. We will utilize this idea in the experiments in Section 4.2.2.

### 3.1.2 DISTANCE PROFILE MATCHING BY ASSIGNMENT

Algorithm 2 is a special case of  $\mathfrak{T}_1$  that finds a one-to-one correspondence, namely, assignment, when the two point clouds have the same number of points. In Step 2 of Algorithm 2, we find a permutation  $\pi \in \mathcal{S}_n$  that minimizes the total discrepancy between the distance profiles of  $X_i$  and  $Y_{\pi(i)}$ 's, which is formulated as the standard linear assignment problem. By the standard linear programming theory, one can deduce that the minimum of this linear assignment problem is equivalent to the minimum of the linear program over the Birkhoff polytope, which is  $\mathfrak{T}_1$  defined in (4) when  $n = m$ . Hence, Algorithm 2 is essentially  $\mathfrak{T}_1$ , which, as discussed, can be impractical in the random sampling setting. Instead, it can be useful in the fixed-location model with noise, which we will analyze in Section 3.3.

---

#### Algorithm 2 Distance Profile Matching by Assignment

---

**Require:** Pairwise distances  $\{d_{\mathcal{X}}(X_i, X_\ell)\}_{i, \ell=1}^n$  and  $\{d_{\mathcal{Y}}(Y_j, Y_\ell)\}_{j, \ell=1}^n$  from two point clouds  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_n\}$  on metric spaces  $(\mathcal{X}, d_{\mathcal{X}})$  and  $(\mathcal{Y}, d_{\mathcal{Y}})$ , respectively.

- 1: For each pair  $(i, j) \in [n] \times [n]$ , define  $W(i, j)$  as in (9).
  - 2: Solve  $\hat{\pi} = \arg \min_{\pi \in \mathcal{S}_n} \sum_{i=1}^n W(i, \pi(i))$ .
  - 3: **return**  $\hat{\pi}$ .
-

### 3.1.3 COMPUTATIONAL COMPLEXITY

The Wasserstein-1 distance between two sets of points in  $\mathbb{R}$  admits a closed-form expression based on the corresponding cumulative distribution functions, which can be computed efficiently (Peyré and Cuturi, 2019). In this case, the computational complexity of Algorithms 1 and 2 is computed as follows. For Step 1, assuming  $n \geq m$ , computing the Wasserstein-1 distances between two sets of points in  $\mathbb{R}$  for  $nm$  pairs takes  $O(n^3 \log(n))$  in total; when  $n = m$  as in Algorithm 2, the exact computation is possible solely based on sorting. Meanwhile, Step 2 takes  $O(nm \log(m))$  as it requires sorting  $m$  numbers  $n$  times; for Algorithm 2, it takes  $O(n^3)$  to solve the linear assignment problem. Hence, the overall complexity of Algorithm 1 or 2 is  $O(n^3 \log(n))$  assuming  $n \geq m$ , which is nearly cubic in the number of points  $n$ ; however, as the input dimension of the algorithm is  $n^2$ , we can say that the complexity scales nearly as the input dimension  $n^2$  raised to the power of  $\frac{3}{2}$ .

**Remark 12** *While we mainly use the Wasserstein-1 distance  $W_1$  to compare the distance profiles, one may use any well-defined discrepancy measure in the place of  $W_1$  in (9). The choice  $W_1$  is natural in connection with the previous section (Remark 6), which is theoretically appealing and practically useful. Lastly, we comment that Algorithms 1 and 2 are implementable without knowing the points  $X_i$ 's and  $Y_j$ 's as long as the pairwise distances  $\{d_X(X_i, X_\ell)\}_{i,\ell=1}^n$  and  $\{d_Y(Y_j, Y_\ell)\}_{j,\ell=1}^m$  are provided. This is because the distance profiles are solely determined by the pairwise distances.*

### 3.2 Partial Matching Analysis: Robustness to Outlier Components

This section analyzes the robustness of Algorithm 1 in the presence of outlier components. Algorithm 1 is designed to match the points corresponding to the same region of the objects, while ignoring or arbitrarily matching the points from the outlier components. In order to theoretically analyze robustness to such components, we consider appropriate model assumptions on the underlying objects  $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$  that capture the nature of the outlier components that we want to ignore in the partial matching setting. Motivated by several practical considerations, we model  $\mu, \nu$  as follows.

**Assumption 1** *Suppose that  $\mu$  and  $\nu$  are mixture distributions consisting of  $t$  and  $s$  components, respectively, defined as follows: for some  $(p_1, \dots, p_t) \in \Delta_t$  and  $(q_1, \dots, q_s) \in \Delta_s$ ,*

$$\mu = \sum_{k=1}^t p_k \mu_k \quad \text{and} \quad \nu = \sum_{k=1}^s q_k \nu_k,$$

where  $\mu_k, \nu_k \in \mathcal{P}(\mathbb{R}^d)$  are sub-Gaussian<sup>1</sup> with variance proxies  $\sigma_k^2, \tau_k^2$  and means  $\theta_k, \eta_k \in \mathbb{R}^d$ , respectively. For some  $K \leq \min\{t, s\}$ , we assume that the first  $K$  components of  $\mu$  and  $\nu$  are identical, that is,  $p_k = q_k$  and  $\mu_k = \nu_k$  for  $k \in [K]$ .

Assumption 1 is equivalent to letting  $\mu = \lambda\rho + (1 - \lambda)\mu_o$  and  $\nu = \lambda\rho + (1 - \lambda)\nu_o$ , where the common part  $\lambda\rho$  is the mixture consisting of the first  $K$  components and the other parts, namely, the outlier parts  $(1 - \lambda)\mu_o, (1 - \lambda)\nu_o$  are the mixtures of the remaining

---

1. See Lemma 20 for the precise definition.

components. The motivation behind Assumption 1 is that the objects  $\mu, \nu$  are composed of several regions, where each region is represented as a sub-Gaussian distribution so that the mean corresponds to the center of the region and the variance proxy roughly represents the size of the region. The points corresponding to the same region are to be matched, while we ignore matching the points from the outlier components.

**Remark 13 (Scope of Assumption 1)** *Gaussian mixture models are the most common instance of Assumption 1, where  $\mu_k$ 's and  $\nu_k$ 's are Gaussian. By allowing these components to be sub-Gaussian, Assumption 1 encompasses a wider class of distributions. For instance, stochastic ball models (Nellore and Ward, 2015; Iguchi et al., 2017)—widely used to evaluate clustering algorithms—are special cases of Assumption 1 where  $\mu_k$ 's and  $\nu_k$ 's are supported on balls. Also, Assumption 1 allows the components to be low-dimensional, which can model 3D objects given as smooth densities on the surface of the objects.*

Now, we apply Algorithm 1 to the point clouds  $\{X_1, \dots, X_n\}$  and  $\{Y_1, \dots, Y_m\}$  that approximate the objects  $\mu$  and  $\nu$ , respectively. We take a viewpoint that the points  $X_1, \dots, X_n$  and  $Y_1, \dots, Y_m$  are independent samples from  $\mu$  and  $\nu$ , respectively. For  $i \in [n]$  and  $j \in [m]$ , let  $t(i) \in [t]$  and  $s(j) \in [s]$  denote the components where  $X_i$  and  $Y_j$  are from, respectively. Our goal is to match  $X_i$  and  $Y_j$  if  $t(i) = s(j) \in [K]$ , namely, they are from the same region. Under this probabilistic setting, we show that Algorithm 1 outputs the matching  $\pi$  that has the following guarantee: for  $X_i$  such that  $t(i) \in [K]$ , we have  $s(\pi(i)) = t(i)$ , namely, its match  $Y_{\pi(i)}$  is from the corresponding component. More precisely, we show that the event

$$\bigcap_{i \in [n]: t(i) \in [K]} (s(\pi(i)) = t(i)), \quad (10)$$

which represents the successful matching of the points from the common components, holds with high probability.

**Theorem 14** *Suppose  $\mu$  and  $\nu$  satisfy Assumption 1. Let  $X_1, \dots, X_n$  and  $Y_1, \dots, Y_m$  be independent samples from  $\mu$  and  $\nu$ , respectively. For  $i \in [n]$  and  $j \in [m]$ , let  $t(i) \in [t]$  and  $s(j) \in [s]$  denote the components where  $X_i$  and  $Y_j$  are from, respectively. Let  $\pi: [n] \rightarrow [m]$  be the output of Algorithm 1. Then, for any  $\delta \in (0, 1)$ , we have*

$$\mathbb{P} \left( \bigcap_{i \in [n]: t(i) \in [K]} (s(\pi(i)) = t(i)) \right) \geq 1 - \delta \quad (11)$$

if

$$\begin{aligned} & \min_{\substack{\alpha \in [K] \\ \beta \in [s] \setminus \{\alpha\}}} \bar{W}(\alpha, \beta) - R \cdot \left( 1 - \sum_{k \in [K]} p_k \right) \\ & \geq \max \left( 8R \sqrt{\frac{\log K + \log(t + s - 2) + \log \frac{4}{\delta}}{2 \min(n, m)}}, 32\Gamma \sqrt{2d + \log \frac{4(\max(n, m))^2}{\delta}} \right), \end{aligned} \quad (12)$$

where  $R = \max\{\|x - y\|_2 : x, y \in \{\theta_1, \dots, \theta_t, \eta_1, \dots, \eta_s\}\}$ ,  $\Gamma = \max\{\sigma_1, \dots, \sigma_t, \tau_1, \dots, \tau_s\}$ , and

$$\bar{W}(\alpha, \beta) := W_1 \left( \sum_{k=1}^t p_k \delta_{\|\theta_\alpha - \theta_k\|_2}, \sum_{k=1}^s q_k \delta_{\|\eta_\beta - \eta_k\|_2} \right) \quad \forall (\alpha, \beta) \in [t] \times [s]. \quad (13)$$

The main takeaway of Theorem 14 is that we have the desired guarantee (11) if *separation among the components compared to the proportion of outlier components  $\geq$  some threshold depending on the noise and sampling error,*

which is exactly what (12) states. Here, the left-hand side of (12) captures the separation in terms of (13) in comparison to the proportion  $1 - \sum_{k \in [K]} p_k$  of the outlier components, while the right-hand side is a threshold determined by the noise and sampling error.

It is worth noting that the right-hand side of (12) contains the term  $\Gamma\sqrt{d}$ , which is sensible as a bound on the typical size of the noise displacement from the centers of the components. This is a minimal condition given that both sides of (12) scale with typical Euclidean distances in  $\mathbb{R}^d$ . See Figure 1. Each sub-Gaussian component corresponds—in the worst case sense—to a ball of radius  $\Gamma\sqrt{d}$ , which has to be well-separated from the other balls to ensure that it is a distinct component.




Figure 1: Illustration of the sub-Gaussian mixtures. For simplicity, each component is represented as a ball centered at the corresponding center, where the radius of the ball is roughly  $\Gamma\sqrt{d}$ . All of them are contained in a larger ball of diameter  $R$ .

### 3.3 One-to-One Matching Analysis: Noise Stability

We now analyze the noise stability of Algorithm 2. Again, remember that Algorithm 2 is a special case of  $\mathfrak{T}_1$  that finds a one-to-one correspondence, namely, assignment, which is not recommended in the random sampling setting of the previous section. Instead, this algorithm can be useful in the fixed-location model with noise; as we analyze below, we find that the matching procedure based on distance profiles is robust to the additive noise.

**Fixed-Location Model with Noise** We consider two sets of locations  $\{\theta_1, \dots, \theta_n\}$  and  $\{\eta_1, \dots, \eta_n\}$  in  $\mathbb{R}^d$  and assume  $\eta_{\pi^*(i)} = T(\theta_i)$  for  $i \in [n]$ , where  $T: \mathbb{R}^d \rightarrow \mathbb{R}^d$  is an unknown rigid transformation and  $\pi^*$  is an unknown permutation on  $[n]$  representing the correspondence, say,  $\theta_i$  and  $\eta_{\pi^*(i)}$  are to be matched. Unlike the previous model in Section 3.2, we assume there is a ground truth one-to-one correspondence  $\pi^*$  between two sets of locations,

which we want to recover from the following perturbed locations  $X_i = \theta_i + \xi_i$  and  $Y_i = \eta_i + \zeta_i$ , where  $\xi_i$  and  $\zeta_i$  are independent mean-zero random vectors following sub-Gaussian distributions with variance proxies  $\sigma_i^2$  and  $\tau_i^2$ , respectively.<sup>2</sup> For instance, for  $d = 3$ , we can think of  $X_i$ 's and  $Y_i$ 's as the coordinates obtained by scanning some objects of interest but from different angles, say, frontal and lateral views, where the noise is added to the ground truth coordinates  $\theta_i$ 's and  $\eta_i$ 's during the scanning process.

The goal is to recover the true permutation  $\pi^*$  from the perturbed locations  $X_i$ 's and  $Y_i$ 's. We show that the output  $\hat{\pi}$  of Algorithm 2 recovers  $\pi^*$  with high probability. As discussed in Section 3.2, several conditions determine the success of the matching. Again, the configuration of the locations  $\theta_1, \dots, \theta_n, \eta_1, \dots, \eta_n$  is crucial for the matching to be well-defined. To this end, we suppose the following configuration condition:

$$\Phi := \min_{i \neq j} W_1 \left( \frac{1}{n} \sum_{k=1}^n \delta_{\|\theta_i - \theta_k\|_2}, \frac{1}{n} \sum_{k=1}^n \delta_{\|\theta_j - \theta_k\|_2} \right) > 0. \quad (14)$$

The condition (14) ensures that  $\pi^*$  is the unique optimal matching in the following sense:

$$\pi^* = \arg \min_{\pi \in \mathcal{S}_n} \sum_{i=1}^n W_1 \left( \frac{1}{n} \sum_{k=1}^n \delta_{\|\theta_i - \theta_k\|_2}, \frac{1}{n} \sum_{k=1}^n \delta_{\|\eta_{\pi(i)} - \eta_k\|_2} \right) =: \arg \min_{\pi \in \mathcal{S}_n} \mathcal{Q}(\pi).$$

To see this, note that  $\mathcal{Q}(\pi^*) = 0$ , while  $\mathcal{Q}(\pi) > 0$  for  $\pi \neq \pi^*$  as  $\Phi > 0$ . The following theorem shows that the output  $\hat{\pi}$  of Algorithm 2 recovers  $\pi^*$  with high probability given this condition (14) and the variances  $\sigma_i$ 's and  $\tau_i$ 's are sufficiently small.

**Theorem 15** *Suppose that the locations  $\theta_1, \dots, \theta_n, \eta_1, \dots, \eta_n \in \mathbb{R}^d$  satisfy*

$$\eta_{\pi^*(i)} = T(\theta_i) \quad \forall i \in [n], \quad (15)$$

*for some unknown rigid transformation  $T: \mathbb{R}^d \rightarrow \mathbb{R}^d$  and some unknown permutation  $\pi^*$  on  $[n]$ . Now, suppose we have the following noisy locations:*

$$X_i = \theta_i + \xi_i \quad \text{and} \quad Y_i = \eta_i + \zeta_i, \quad (16)$$

*where  $\xi_i$  and  $\zeta_i$  are independent mean-zero random vectors following sub-Gaussian distributions with variance proxies  $\sigma_i^2$  and  $\tau_i^2$ , respectively. Let  $\hat{\pi}$  be the output of Algorithm 2. Then, for any  $\delta \in (0, 1)$ , we have*

$$\mathbb{P}(\hat{\pi} = \pi^*) \geq 1 - \delta$$

*if for  $\Phi$  defined in (14), the following condition holds:*

$$\max\{\sigma_1, \dots, \sigma_n, \tau_1, \dots, \tau_n\} \leq \frac{\Phi}{16\sqrt{2d + \log(2n^2/\delta)}}. \quad (17)$$

---

2. Technically, one may think of this model as the previous model in Section 3.2 with  $t = s = K$ , namely, no outlier components. However, the key difference is that we no longer sample randomly from  $\mu$  and  $\nu$ , which is not the setting we want to use Algorithm 2 for as we previously discussed.

For the case where there is no rigid transformation, namely,  $T$  is the identity map, the noise stability guarantee in the above model is well-studied in the literature, known as the feature matching problem. In this case, the simplest approach is to solve the following discrete optimal transport problem (Peyré and Cuturi, 2019):

$$\hat{\pi}^{\text{OT}} = \arg \min_{\pi \in \mathcal{S}_n} \sum_{i=1}^n \|X_i - Y_{\pi(i)}\|_2^2. \quad (18)$$

It is well-studied (Collier and Dalalyan, 2016; Galstyan et al., 2022) that such a procedure is indeed guaranteed to recover  $\pi^*$  with high probability if the separation among  $\theta_i$ 's is large enough compared to the noise level. However, in the presence of an unknown rigid transformation  $T$ , such a procedure is not guaranteed to succeed, while  $\hat{\pi}$  in Theorem 15 is invariant to rigid transformations. We demonstrate this point in Section 4.

## 4. Simulations

### 4.1 Synthetic Data

#### 4.1.1 GAUSSIAN MIXTURE MATCHING

We first provide a simulation study to demonstrate the performance of the matching procedure discussed in Section 3.2. Figure 2(a) shows the probability of the perfect matching, namely, the left-hand side of (11). Here,  $d = 3$  and  $K = 10$ , and we consider the clean case  $t = s = K$  and the case with one outlier component  $t = s = K + 1$  for the sample size  $n = m \in \{500, 1000\}$ ; also, we let  $p_k = q_k = \frac{1}{t}$  and consider Gaussian mixtures, namely,  $\mu_k = N(\theta_k, \sigma^2 I_d)$  and  $\nu_k = N(\eta_k, \sigma^2 I_d)$ , where we vary  $\sigma$ . We can observe that the probability of the perfect matching decreases as  $\sigma$  grows for both cases. Though the presence of an outlier component deteriorates the matching performance as expected, we can verify that Algorithm 1 is robust to the outlier components in that the recovery probability is sufficiently high when  $\sigma$  is small. Comparing the plots for  $n = m = 500$  and  $n = m = 1000$ , we can see that the matching performance is better for the larger sample size as it better approximates the population-level distance profiles.

Although the analysis in Section 3.2 was based on the probability of the perfect matching, it is reasonable to expect matching errors in practice. Figure 2(b) shows the matching accuracy for the same setting as Figure 2(a) with  $n = m = 1000$ , where the matching accuracy is the proportion of correctly matched  $X_i$ 's from the first  $K$  components:

$$\frac{|\{i \in [n] : t(i) \in [K] \text{ and } s(\pi(i)) = t(i)\}|}{|\{i \in [n] : t(i) \in [K]\}|}, \quad (19)$$

where  $t, s$  are defined as in Theorem 14. Again, we can see that the proposed method performs reasonably well. Figure 2(c) compares the case with one outlier component ( $K = 10$  and  $t = s = 11$ ) to the case with more outlier components but with the same proportion ( $K = 20$  and  $t = s = 22$ ) and to the case with more outlier components with larger proportions ( $K = 20$  and  $t = s = 24$ ). We can observe that the matching accuracy decreases as the proportion of outlier components increases, which is consistent with the theoretical results in Section 3.2. Even when the proportion  $(t - K)/K$  is the same, the

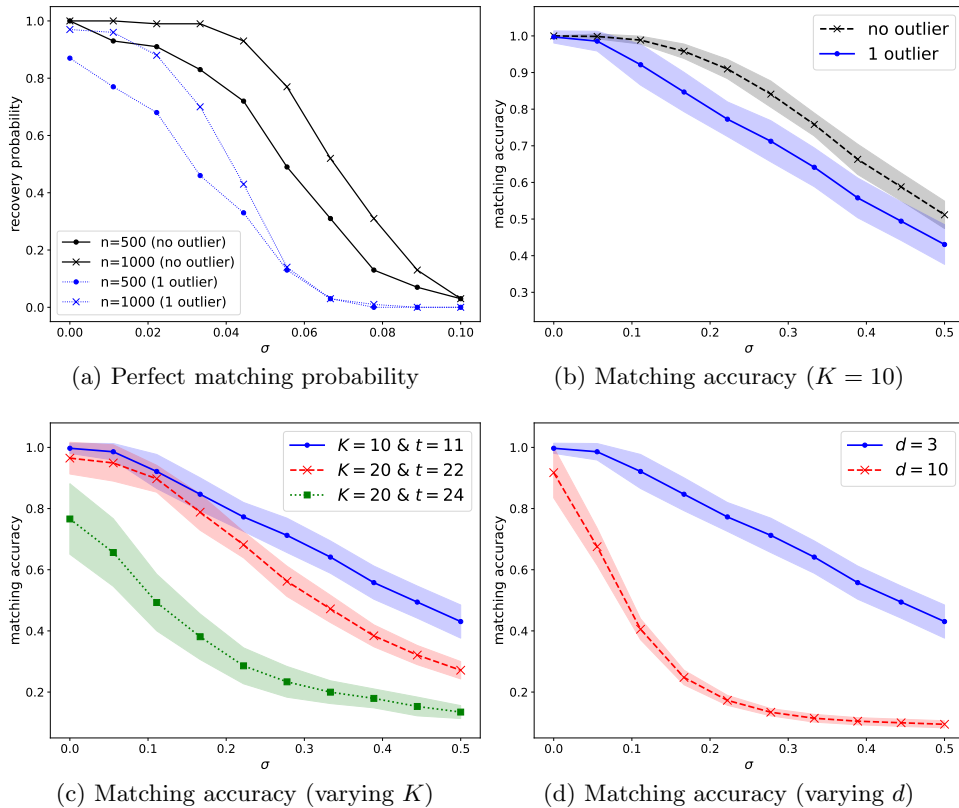


Figure 2: Gaussian mixture matching results with  $p_k = q_k = \frac{1}{t}$ ,  $\mu_k = N(\theta_k, \sigma^2 I_d)$ ,  $\nu_k = N(\eta_k, \sigma^2 I_d)$ , and varying  $\sigma$ . (a) shows the probability—the left-hand side of (11)—with  $d = 3$ ,  $K = 10$ , and the sample size  $n = m \in \{500, 1000\}$ . (b) shows the matching accuracy (19) for the same setting as (a), while (c), (d) show the matching accuracy for different settings, varying  $K, d$ , respectively, with  $n = m = 1000$  for (b)-(d). For all plots, the sampling from  $\mu, \nu$  is repeated 100 times. For (a), the markers show the probability estimated from the 100 realizations; for (b)-(d), they show the average of the matching accuracy over 100 realizations, with the standard deviations shown as the shaded regions. The blue solid lines of (b)-(d) are the same, showing the case of  $K = 10$  and  $t = s = 11$ .

matching accuracy is lower for larger  $t$  as the larger effective dimensions of the underlying distributions  $\mu, \nu$  lead to larger estimation errors given the same sample size. Figure 2(d) compares the case with  $d = 3$  to the case with  $d = 10$  for  $K = 10$  and  $t = s = 11$ . We can see that the matching accuracy is lower and deteriorates much faster for larger  $d$ .

#### 4.1.2 NOISE STABILITY

Next, we provide simulation results demonstrating the noise stability discussed in Section 3.3. We first obtain  $n$  locations  $\theta_1, \dots, \theta_n \in \mathbb{R}^d$  and define the other locations  $\eta_1, \dots, \eta_n$  as

in (15) using some rigid transformation  $T$  and a permutation  $\pi^*$ . Then, we generate the points according to (16) and obtain  $\hat{\pi}$  by applying Algorithm 2; we estimate  $\mathbb{P}(\hat{\pi} = \pi^*)$  by repeating this for 100 times. Figure 3 shows the results, where the noise variables  $\xi_i$ 's and  $\zeta_i$ 's follow  $N(0, \sigma^2 I_d)$ , and we vary  $\sigma$  and compare with the probability  $\mathbb{P}(\hat{\pi}^{\text{OT}} = \pi^*)$  for  $\hat{\pi}^{\text{OT}}$  obtained by the procedure (18). As shown in Figure 3(a), when the rigid transformation is based on a rotation matrix that only affects the first two coordinates, both methods can recover the true permutation to some extent for small enough  $\sigma$ ; in this case, as the rotation matrix is relatively close to the identity matrix, the standard optimal transport method (18) shows good performance. For a rotation matrix that changes all the coordinates, however, the standard optimal transport method fails to recover the true permutation as shown in Figure 3(b), while the distance profile matching method performs almost as in Figure 3(a).

Lastly, we briefly comment that the plots of Figure 2(a) and Figure 3 suggest the existence of a phase transition in the noise level, where the performance of the matching methods might change dramatically under certain settings. We leave this for future work.

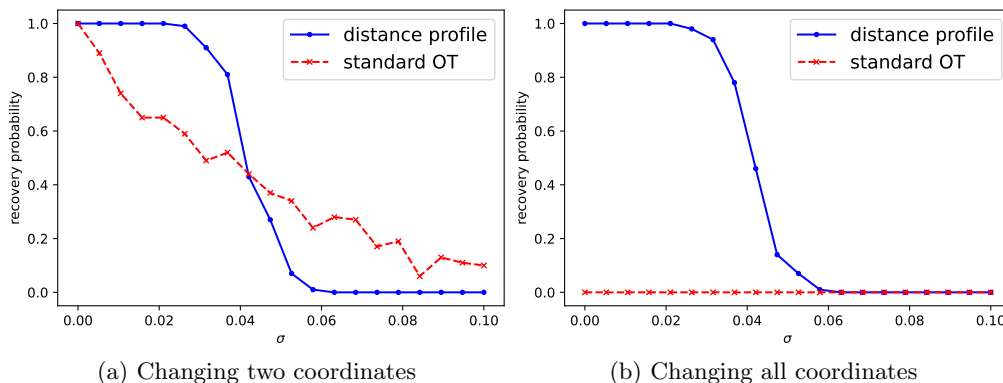


Figure 3: Noise stability results under the setting of Theorem 15 with  $\xi_i$ 's and  $\zeta_i$ 's following  $N(0, \sigma^2 I_d)$ , where  $d = 10$  and  $n = 100$ . The plots show the probability  $\mathbb{P}(\hat{\pi} = \pi^*)$ , where  $\hat{\pi}$  corresponds to the output of Algorithm 2 for the blue lines (distance profile) and to (18) for the red lines (standard OT); here, the probability is estimated by repeating each setting 100 times. For (a), the rigid transformation is based on a rotation matrix that only affects the first two coordinates, while (b) is based on a random rotation matrix changing all  $d$  coordinates.

## 4.2 Applications

### 4.2.1 SHAPE CORRESPONDENCE

We demonstrate how Algorithm 1 in practice using stylized data examples. We apply it to the shape correspondence using CAPOD data (Papadakis, 2014). Figure 4(a) and Figure 4(d) show the two 3D images representing dolphins with different poses, from which we obtain point clouds  $X_1, \dots, X_n$  and  $Y_1, \dots, Y_m$  as shown in Figure 4(b) and Figure 4(e), respectively, where  $n = m = 285$  in this example. We want to match the points corresponding to the same part of the dolphins, rather than finding a one-to-one correspondence. As

explained earlier, it is reasonable to match similar points to the same point. For instance, there are many points concentrated around the nose—called the rostrum—of the source dolphin, which are expected to be matched to any points concentrated around the rostrum of the target dolphin. Hence, we first divide the dolphins into several components. Here, we use k-means clustering with  $k = 7$ , which indeed cluster points based on the body parts.

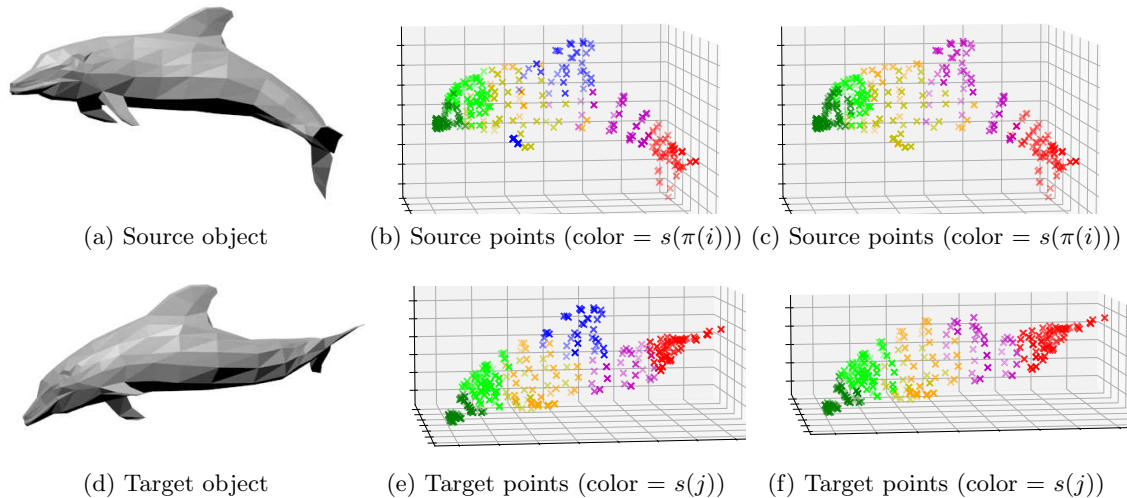


Figure 4: 3D images of dolphins and the corresponding point clouds. (a) and (d) are the source and target 3D images, from which we obtain point clouds  $X_1, \dots, X_n$  and  $Y_1, \dots, Y_m$ , with  $n = m = 285$ , as shown in (b) and (e), respectively. We apply k-means clustering to both point clouds with  $k = 7$  to cluster them based on the body parts. (e) shows the target points  $Y_j$ 's colored by  $s(j)$ 's, where  $s(j) \in \{1, \dots, 7\}$  denotes the index of the component that  $Y_j$  belongs to. (b) shows the source points  $X_i$ 's colored by  $s(\pi(i))$ 's, where  $\pi$  is obtained by applying Algorithm 1. (c) and (f) repeat (b) and (e), respectively, after removing 16 points around the fin of the target dolphin in (e) and clustering again with  $k = 6$ .

Figure 4(e) shows the target points  $Y_j$ 's colored by  $s(j)$ 's, where  $s(j) \in \{1, \dots, 7\}$  is the index of the component that  $Y_j$  belongs to. Now, we run Algorithm 1 and obtain the output  $\pi: [n] \rightarrow [m]$ . For each  $X_i$ , we verify if the part that it belongs to coincides with the part that its match  $Y_{\pi(i)}$  belongs to, which is indexed by  $s(\pi(i))$ . To this end, Figure 4(b) shows the source points  $X_i$ 's colored by  $s(\pi(i))$ 's. We can observe that the proposed procedure recovers the corresponding regions mostly correctly. For instance, for  $X_i$ 's corresponding to the rostrum of the source dolphin, their colors  $s(\pi(i))$ 's (green) indicate that they are matched to the points corresponding to the rostrum of the target dolphin. Similarly, the points corresponding to the head (lime), tail (red), dorsal fin (blue), and peduncle (magenta) are mostly matched correctly, while there are some mismatches, for instance, among the points around the pectoral fin. Figure 4(c) and Figure 4(f) repeat the same procedure after removing 16 points around the fin of the target dolphin in Figure 4(e) and clustering again with  $k = 6$ . Figure 4(c) shows that the matching result is similar to the previous one. In this case, though the points around the fin of the source dolphin do not have corresponding

points in the target dolphin and thus we may ignore matching them, the proposed procedure still matches them to the points around the peduncle (magenta) of the target dolphin, which is a cluster adjacent to the removed fin.

Figure 5 repeats the same procedure for the 3D images of horses. Again, we can verify that the proposed procedure matches the corresponding body parts mostly correctly. Figure 5(c) and Figure 5(f) shows the results after removing 62 points around the two legs on the target horse’s left side in Figure 5(e) and clustering again with  $k = 6$ . We again observe that the points in the common parts are mostly matched correctly.

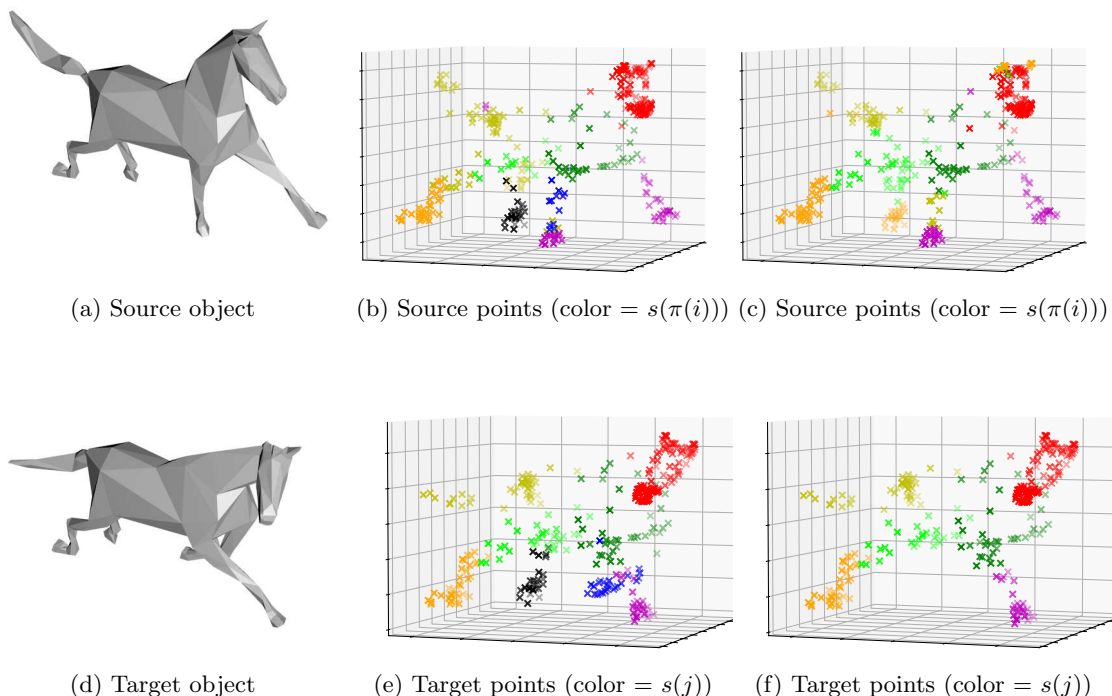


Figure 5: 3D images of horses and the corresponding point clouds. (a) and (d) are the source and target 3D images, from which we obtain point clouds  $X_1, \dots, X_n$  and  $Y_1, \dots, Y_m$ , with  $n = m = 346$ , as shown in (b) and (e), respectively. We apply k-means clustering to both point clouds with  $k = 8$  to cluster them based on the body parts. (e) shows the target points  $Y_j$ ’s colored by  $s(j)$ ’s, where  $s(j) \in \{1, \dots, 8\}$  denotes the index of the component that  $Y_j$  belongs to. (b) shows the source points  $X_i$ ’s colored by  $s(\pi(i))$ ’s, where  $\pi$  is obtained by applying Algorithm 1. (c) and (f) repeat (b) and (e), respectively, after removing 62 points around the two legs on the target horse’s left side in (e) and clustering again with  $k = 6$ .

#### 4.2.2 PARTIAL MATCHING FOR ALIGNMENT OF CRYO-EM STRUCTURES

We apply the developed partial matching procedure to align cryogenic electron microscopy (cryo-EM) structures of human  $\gamma$ -secretase. Cryo-EM has gained significant attention in the




Figure 6: Human  $\gamma$ -secretase. (a) shows the 3D model of human  $\gamma$ -secretase, and (b)-(d) show the projections of the cryo-EM data along the  $x$ ,  $y$ , and  $z$  axes, respectively.

past decade as a powerful tool for unveiling the structures of biomolecules at near-atomic resolution. However, many computational challenges arise; see Bendory et al. (2020) for a review. Aligning cryo-EM structures is particularly challenging due to high noise levels and different conformations. Figure 6(a) shows the 3D model of human  $\gamma$ -secretase, which is thought to be involved in the development of Alzheimer’s disease. The cryo-EM data<sup>3</sup> of human  $\gamma$ -secretase (Bai et al., 2015) essentially represents a probability density function of the intensities of the object, stored as voxels on a 3D grid. Figures 6(b)-(d) show the projections of the cryo-EM data along the  $x$ ,  $y$ , and  $z$  axes, respectively, which clearly demonstrate high noise levels.

Alignment problems in cryo-EM are often tackled by using the density functions (Singer and Yang, 2024; Harpaz and Shkolnisky, 2023) with suitable transformations or by obtaining the point clouds from the density functions (Riahi et al., 2023, 2025), say, utilizing the topology representing network algorithm (Martinetz and Schulten, 1994). The proposed algorithm can be applied to the latter setting. Here, we sample two point clouds from the cryo-EM data, one from the original density function and the other from the density function after applying a rigid transformation; see Figure 7(a). The goal is to recover the rigid transformation from the point clouds.

The proposed partial matching—Algorithm 1—provides correspondences between the points. Once we have the correspondences, we can estimate the rigid transformation using the standard orthogonal Procrustes problem—after centering—between matrices  $\mathbb{X}, \mathbb{Y} \in \mathbb{R}^{n \times d}$  whose rows are the point clouds  $X_1, \dots, X_n$  and  $Y_{\pi(1)}, \dots, Y_{\pi(n)}$ , respectively, where  $\pi$  is the output of Algorithm 1. Figure 7(b) shows the aligned point clouds, which seems reasonable. However, for better results, it is beneficial to subset the good matches by using the threshold  $\rho$  in Algorithm 1. The crux is that we only need to partially match the points for alignment, and focusing on good matches can improve the alignment results. Figures 7(c) and 7(d) apply Algorithm 1 with  $\rho$  being the 0.5 and 0.1 quantiles of  $W(1, \pi(1)), \dots, W(n, \pi(n))$  to use the best 50% and 10% of the correspondences, respectively. In these cases, we obtain the set of index  $I$  from the output of Algorithm 1 and

3. Listed as 5A63 at the Protein Data Bank (PDB) <https://www.rcsb.org/structure/5A63> and as EMD-3061 at the Electron Microscopy Data Bank (EMDB) <https://www.ebi.ac.uk/emdb/EMD-3061>.

estimate the rigid transformation using the submatrices  $\mathbb{X}_I, \mathbb{Y}_I$  of  $\mathbb{X}, \mathbb{Y}$ , respectively, corresponding to the rows indexed by  $I$ . Figures 7(c) and 7(d) indeed look better than using the full correspondences, which we can quantify by computing the Wasserstein-2 distance between the aligned point cloud and the target point cloud. The Wasserstein-2 distance between the aligned point cloud and the target point cloud is 10.68 for the full correspondences, 9.04 for the best 50%, and 8.31 for the best 10%, which confirms that the partial matching improves the alignment results.

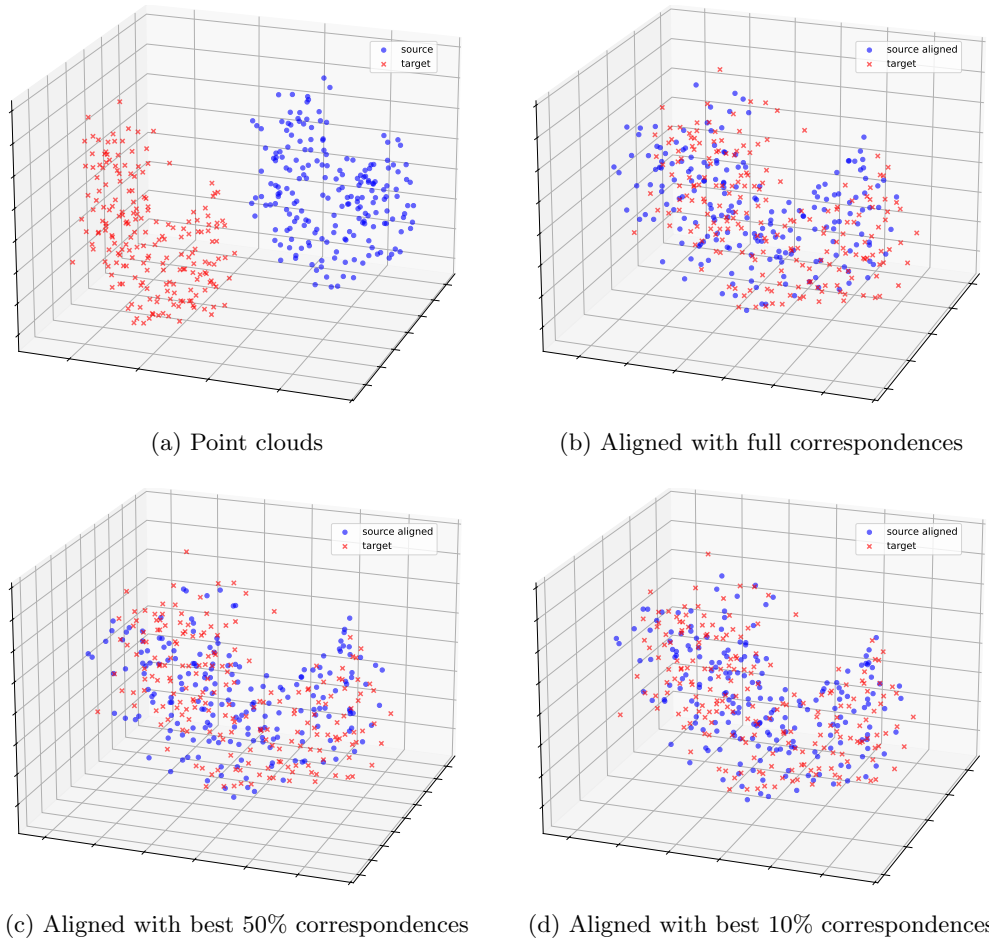


Figure 7: Point clouds from the cryo-EM data of human  $\gamma$ -secretase and alignment results. (a) shows the point clouds sampled from the original density function and the density function after applying a rigid transformation, shown as source and target, respectively. For (b)-(c), we first apply Algorithm 1 to obtain the correspondences and then estimate the rigid transformation using the orthogonal Procrustes problem. (b) shows the aligned point clouds using all the correspondences, while (c) and (d) use the best 50% and 10% of the correspondences, respectively, based on the threshold  $\rho$  in Algorithm 1.




Figure 8: Rabbit  $\text{Ca}_v1.1$ . (a) shows the 3D model of  $\text{Ca}_v1.1$ , and (b)-(d) show the projections of the cryo-EM data along the  $x$ ,  $y$ , and  $z$  axes, respectively.

We repeat the same process for the cryo-EM data of the voltage-gated calcium channel ( $\text{Ca}_v1.1$ ) of rabbit (Wu et al., 2016), which plays a role as the sensor for excitation-contraction coupling of skeletal muscles. Figure 8 shows the 3D model and projections of the cryo-EM data. Figure 9 shows the alignment results. Again, we obtain two point clouds from the cryo-EM data, one from the original density function and the other from the density function after applying a rigid transformation; see Figure 9(a). Figures 9(b)-(d) show the aligned point clouds using all the correspondences, the best 50%, and the best 10%, respectively. The Wasserstein-2 distance between the aligned point cloud and the target point cloud is 11.10 for the full correspondences, 8.23 for the best 50%, and 9.17 for the best 10%. In this case, taking the best 50% correspondences yields the best partial matching result among the three settings.

## 5. Conclusion

We have proposed matching procedures based on distance profiles, which are easily implementable without having to solve nontrivial optimization problems over permutations. Most importantly, we have shown that the proposed method is provably robust to outlier components and noise under some model assumptions that we believe capture essential aspects of empirical applications. To the best of our knowledge, this is the first work that provides theoretical guarantees for the robustness of matching procedures based on distance profiles in the context of object matching. Also, we have derived the new statistical complexity result of the distance profile distance—the lower bound to the Gromov-Wasserstein distance—that has provided insights into how to use the distance profile matching procedures in practice. We have also demonstrated the empirical performance of the proposed methods through simulations and applications to real data.

We suggest several directions for future research. On the theoretical side, it would be interesting to consider other types of matching guarantees in Theorem 14. As mentioned in Section 4, the matching accuracy (19) may be practically more useful. Similarly, for Theorem 15, it would be interesting to consider the probability that  $\hat{\pi}$  is close to  $\pi^*$  under a certain discrepancy, such as the Hamming distance, instead of the probability of  $\hat{\pi} = \pi^*$ .

Also, one can explore a generalization based on mixtures beyond the Euclidean space, say, a general metric or Banach space. On the practical side, it would be interesting to delve deeper into the procedure of finding inliers, that is, the points from the common part of two objects. For interested readers, we discuss in Appendix C.2 how the threshold  $\rho$  of Algorithm 1 can be chosen to identify the points from the common part at least in theory. However, developing a practical method applicable to real data is very challenging in general and needs further investigation. We leave these directions for future research.

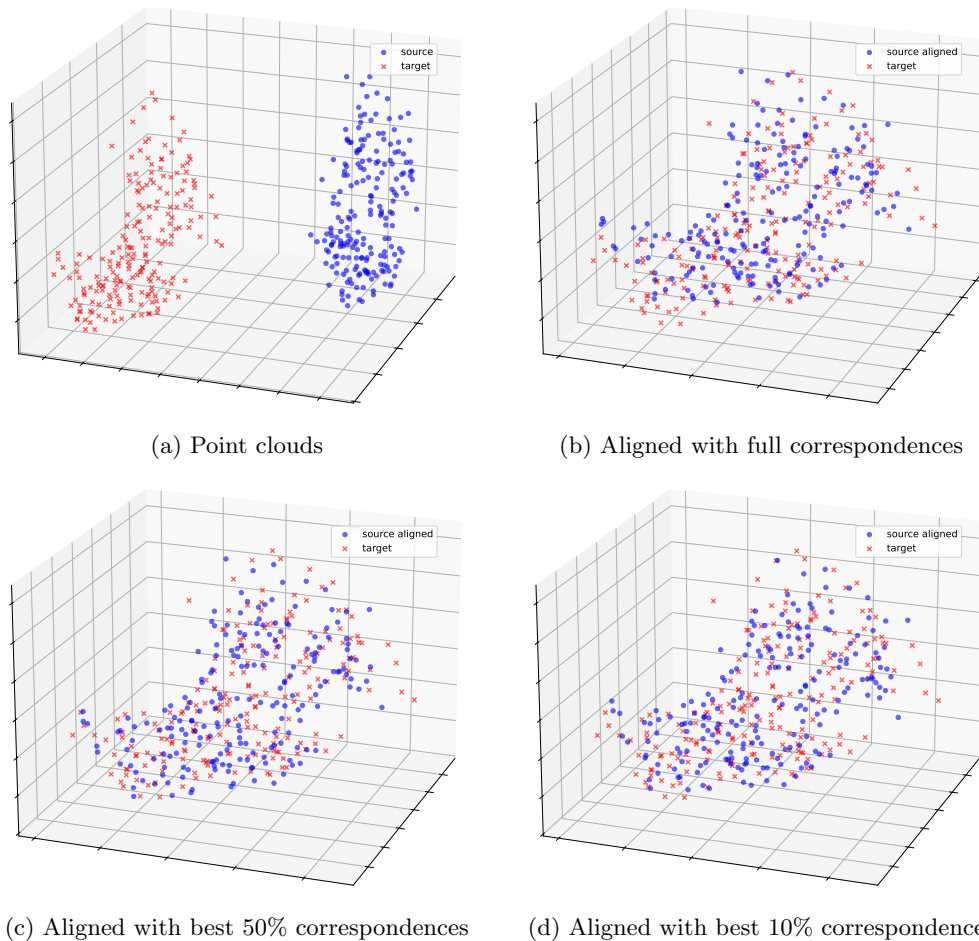


Figure 9: Point clouds from the cryo-EM data of rabbit  $\text{Ca}_v1.1$  and alignment results. (a) shows the point clouds sampled from the original density function and the density function after applying a rigid transformation, shown as source and target, respectively. For (b)-(c), we first apply Algorithm 1 to obtain the correspondences and then estimate the rigid transformation using the orthogonal Procrustes problem. (b) shows the aligned point clouds using all the correspondences, while (c) and (d) use the best 50% and 10% of the correspondences, respectively, based on the threshold  $\rho$  in Algorithm 1.

## Appendix A. Proofs

### A.1 Proof of Theorem 7

**Proof** For  $x \in \mathcal{X}$  and  $r \in \mathbb{R}$ , let  $\phi_{x,r}(\cdot) = 1\{d_{\mathcal{X}}(x, \cdot) \leq r\}$ , which is a function from  $\mathcal{X}$  to  $\{0, 1\}$ . Let  $F_x$  be the cumulative distribution function of  $\mu_x$ . Then,  $F_x(r) = \mu\phi_{x,r}$  by definition. Similarly, letting  $\hat{F}_x$  be the cumulative distribution function of  $\hat{\mu}_x$ , we have  $\hat{F}_x(r) = \hat{\mu}\phi_{x,r}$ . Recalling that  $W_1$  between two probability measures on  $\mathbb{R}$  is the  $L_1$  distance between their cumulative distribution functions, we have

$$W_1(\mu_x, \hat{\mu}_x) = \int_0^\infty |F_x(r) - \hat{F}_x(r)| dr = \int_0^\infty |\mu\phi_{x,r} - \hat{\mu}\phi_{x,r}| dr \leq \sup_{\phi \in \mathcal{F}_x} |\mu\phi - \hat{\mu}\phi|,$$

where  $\mathcal{F}_x := \{\phi_{x,r} : r \geq 0\}$  for any  $x \in \mathcal{X}$ . Therefore,

$$\sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) \leq \sup_{x \in \mathcal{X}} \sup_{\phi \in \mathcal{F}_x} |\mu\phi - \hat{\mu}\phi| = \sup_{\phi \in \mathcal{F}_{\mathcal{X}}} |\mu\phi - \hat{\mu}\phi|,$$

where  $\mathcal{F}_{\mathcal{X}} := \bigcup_{x \in \mathcal{X}} \mathcal{F}_x = \{\phi_{x,r} : x \in \mathcal{X}, r \geq 0\}$ . By Theorem 8.23 of Vershynin (2018),

$$\mathbb{E} \left[ \sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) \right] \leq \mathbb{E} \left[ \sup_{\phi \in \mathcal{F}_{\mathcal{X}}} |\mu\phi - \hat{\mu}\phi| \right] \leq C \sqrt{\frac{\text{vc}(\mathcal{F}_{\mathcal{X}})}{n}}. \quad (20)$$

where  $C$  is an absolute constant that is independent of  $x$ ,  $n$ ,  $\mathcal{X}$ , and  $\mu$ . Here,  $\text{vc}(\mathcal{F}_{\mathcal{X}})$  is the Vapnik-Chervonenkis (VC) dimension of the function class  $\mathcal{F}_{\mathcal{X}}$ . As  $\mathcal{X} \subset \mathbb{R}^d$ , this is upper bounded by the VC dimension of the set of all closed Euclidean balls in  $\mathbb{R}^d$ , which is  $d + 1$ ; see Dudley (1979). Hence, we obtain (7).  $\blacksquare$

### A.2 Proof of Proposition 9

**Proof** Define a map  $S: \mathcal{X} \rightarrow \mathbb{R} \times \mathbb{R}$  by letting  $S(\cdot) = (d_{\mathcal{X}}(x, \cdot), d_{\mathcal{X}}(x', \cdot))$ . Then, one can verify that  $S_{\#}\mu$  is a coupling of  $\mu_x$  and  $\mu_{x'}$ . Therefore,

$$\begin{aligned} W_1(\mu_x, \mu_{x'}) &= \inf_{\rho \in \Pi(\mu_x, \mu_{x'})} \int_{\mathbb{R} \times \mathbb{R}} |z_1 - z_2| d\rho(z_1, z_2) \\ &\leq \int_{\mathbb{R} \times \mathbb{R}} |z_1 - z_2| dS_{\#}\mu(z_1, z_2) \\ &= \int_{\mathcal{X}} |d_{\mathcal{X}}(x, \bar{x}) - d_{\mathcal{X}}(x', \bar{x})| d\mu(\bar{x}) \\ &\leq d_{\mathcal{X}}(x, x'), \end{aligned}$$

where the last inequality follows from the triangle inequality of  $d_{\mathcal{X}}$ .  $\blacksquare$

### A.3 Proof of Proposition 10

**Proof** By definition, one can check that  $\mu_x$  is also given as the mixture of the distance profiles  $(\mu_k)_x$ 's, namely,  $\mu_x = \sum_{k=1}^t p_k (\mu_k)_x$ . Then,

$$W_1(\mu_x, \mu_{x'}) = W_1 \left( \sum_{k=1}^t p_k (\mu_k)_x, \sum_{k=1}^t p_k (\mu_k)_{x'} \right) \leq \sum_{k=1}^t p_k W_1((\mu_k)_x, (\mu_k)_{x'}), \quad (21)$$

where the inequality follows from the fact that if  $\gamma_k \in \Pi((\mu_k)_x, (\mu_k)_{x'})$  is a coupling between  $(\mu_k)_x$  and  $(\mu_k)_{x'}$ , the convex combination  $\sum_{k=1}^t p_k \gamma_k$  is a valid coupling between  $\mu_x$  and  $\mu_{x'}$ . As  $\mu_k$  is rotationally invariant with respect to the center  $\theta_k$ , we can deduce that  $(\mu_k)_x = (\mu_k)_{\theta_k + \|x - \theta_k\|_2 u}$  for any  $u$  on the unit sphere. Hence,

$$\begin{aligned} W_1((\mu_k)_x, (\mu_k)_{x'}) &= W_1((\mu_k)_{\theta_k + \|x - \theta_k\|_2 u}, (\mu_k)_{\theta_k + \|x' - \theta_k\|_2 u}) \\ &\leq \|\theta_k + \|x - \theta_k\|_2 u - \theta_k - \|x' - \theta_k\|_2 u\|_2 \\ &= \left| \|x - \theta_k\|_2 - \|x' - \theta_k\|_2 \right|, \end{aligned} \quad (22)$$

where the inequality uses Proposition 9. Combining (21) and (22), we obtain (8).  $\blacksquare$

### A.4 Proof of Theorem 11

**Proof** By (5), (6), and Theorem 7, it suffices to show that  $\mathbb{E}(\ast) \leq O(n^{-1/t})$ . Define  $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}^t$  by letting  $\Phi(x) = (\sqrt{p_1} \|x - \theta_1\|_2, \dots, \sqrt{p_t} \|x - \theta_t\|_2)$ . From Proposition 10, we have

$$W_1(\mu_x, \mu_{x'}) \leq \sum_{k=1}^t p_k \left| \|x - \theta_k\|_2 - \|x' - \theta_k\|_2 \right| \leq \|\Phi(x) - \Phi(x')\|_2,$$

where the last inequality follows from the Cauchy-Schwarz inequality. For any coupling  $\gamma \in \Pi(\mu, \hat{\mu})$ , note that  $(\Phi \times \Phi)_{\#} \gamma \in \Pi(\Phi_{\#} \mu, \Phi_{\#} \hat{\mu})$ , where  $\Phi \times \Phi: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}^t \times \mathbb{R}^t$  is a map defined by  $(\Phi \times \Phi)(x, x') = (\Phi(x), \Phi(x'))$ . Conversely, for any coupling  $\rho \in \Pi(\Phi_{\#} \mu, \Phi_{\#} \hat{\mu})$ , we can find a coupling  $\gamma \in \Pi(\mu, \hat{\mu})$  such that  $\rho = (\Phi \times \Phi)_{\#} \gamma$  by Lemma 3.12 of Ambrosio and Gigli (2013). In summary,

$$\Pi(\Phi_{\#} \mu, \Phi_{\#} \hat{\mu}) = \{(\Phi \times \Phi)_{\#} \gamma : \gamma \in \Pi(\mu, \hat{\mu})\}. \quad (23)$$

Therefore,

$$\begin{aligned} \inf_{\gamma \in \Pi(\mu, \hat{\mu})} \int_{\mathbb{R}^d \times \mathbb{R}^d} W_1(\mu_x, \mu_{x'}) d\gamma(x, x') &\leq \inf_{\gamma \in \Pi(\mu, \hat{\mu})} \int_{\mathbb{R}^d \times \mathbb{R}^d} \|\Phi(x) - \Phi(x')\|_2 d\gamma(x, x') \\ &= \inf_{\gamma \in \Pi(\mu, \hat{\mu})} \int_{\mathbb{R}^t \times \mathbb{R}^t} \|z - z'\|_2 d(\Phi \times \Phi)_{\#} \gamma(z, z') \\ &= \inf_{\rho \in \Pi(\Phi_{\#} \mu, \Phi_{\#} \hat{\mu})} \int_{\mathbb{R}^t \times \mathbb{R}^t} \|z - z'\|_2 d\rho(z, z') \\ &= W_1(\Phi_{\#} \mu, \Phi_{\#} \hat{\mu}), \end{aligned}$$

where the first equality follows from the change of variables, and the second equality is due to (23). As  $\Phi_{\#}\hat{\mu}$  is equivalent to the empirical measure of  $\Phi_{\#}\mu$  which is defined on  $\mathbb{R}^t$ , we conclude from the standard results (Dudley, 1969; Fournier and Guillin, 2015) that  $\mathbb{E}(\ast) \leq W_1(\Phi_{\#}\mu, \Phi_{\#}\hat{\mu}) \leq O(n^{-1/t})$ .  $\blacksquare$

### A.5 Proof of Theorem 14

**Proof** We analyze the deviation of (9) from (13), namely, for any pair  $(i, j) \in [n] \times [m]$ , we derive an upper bound on the deviation  $|W(i, j) - \overline{W}(t(i), s(j))|$ . To this end, we define an intermediate quantity  $W'(i, j)$  as follows: for any  $(i, j) \in [n] \times [m]$ ,

$$\begin{aligned} W'(i, j) &:= W_1\left(\frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_{t(i)} - \theta_{t(\ell)}\|_2}, \frac{1}{m} \sum_{\ell=1}^m \delta_{\|\eta_{s(j)} - \eta_{s(\ell)}\|_2}\right) \\ &= W_1\left(\sum_{k=1}^t \frac{n_k}{n} \delta_{\|\theta_{t(i)} - \theta_k\|_2}, \sum_{k=1}^s \frac{m_k}{m} \delta_{\|\eta_{s(j)} - \eta_k\|_2}\right), \end{aligned}$$

where we define  $n_k = \sum_{\ell=1}^n 1\{t(\ell) = k\}$  and  $m_k = \sum_{\ell=1}^m 1\{s(\ell) = k\}$ . We first derive an upper bound on  $|W'(i, j) - \overline{W}(t(i), s(j))|$  for any  $(i, j) \in [n] \times [m]$ . By the triangle inequality of  $W_1$ , we have

$$\begin{aligned} &|W'(i, j) - \overline{W}(t(i), s(j))| \\ &\leq \underbrace{W_1\left(\sum_{k=1}^t \frac{n_k}{n} \delta_{\|\theta_{t(i)} - \theta_k\|_2}, \sum_{k=1}^t p_k \delta_{\|\theta_{t(i)} - \theta_k\|_2}\right)}_{=:(\ast)} + \underbrace{W_1\left(\sum_{k=1}^s \frac{m_k}{m} \delta_{\|\eta_{s(j)} - \eta_k\|_2}, \sum_{k=1}^s q_k \delta_{\|\eta_{s(j)} - \eta_k\|_2}\right)}_{=:(\ast\ast)}. \end{aligned}$$

For each  $\alpha \in [K]$ , order  $\theta_1, \dots, \theta_t$  by the distance to  $\theta_\alpha$ , namely, find  $k_\alpha : [t] \rightarrow [t]$  such that  $\|\theta_\alpha - \theta_{k_\alpha(1)}\|_2 \leq \dots \leq \|\theta_\alpha - \theta_{k_\alpha(t)}\|_2$ . Then, by Lemma 18, we have

$$\begin{aligned} (\ast) &\leq \max_{\alpha \in [K]} W_1\left(\sum_{k=1}^t \frac{n_k}{n} \delta_{\|\theta_\alpha - \theta_k\|_2}, \sum_{k=1}^t p_k \delta_{\|\theta_\alpha - \theta_k\|_2}\right) \\ &\leq \max_{\alpha \in [K]} \|\theta_\alpha - \theta_{k_\alpha(t)}\|_2 \cdot \max_{j \in [t-1]} \left| \sum_{k=1}^j \left(\frac{n_{k_\alpha(k)}}{n} - p_{k_\alpha(k)}\right) \right| \\ &\leq R \cdot \max_{\alpha \in [K]} \max_{j \in [t-1]} \left| \sum_{k=1}^j \left(\frac{n_{k_\alpha(k)}}{n} - p_{k_\alpha(k)}\right) \right| =: R \cdot Z_1. \end{aligned}$$

Similarly, for each  $\beta \in [K]$ , order  $\eta_1, \dots, \eta_s$  by the distance to  $\eta_\beta$ , namely, find  $\ell_\beta: [s] \rightarrow [s]$  such that  $\|\eta_\beta - \eta_{\ell_\beta(1)}\|_2 \leq \dots \leq \|\eta_\beta - \eta_{\ell_\beta(s)}\|_2$ . Then, by Lemma 18, we have

$$\begin{aligned}
 (**) &\leq \max_{\beta \in [K]} W_1 \left( \sum_{k=1}^s \frac{m_k}{m} \delta_{\|\eta_\beta - \eta_k\|_2}, \sum_{k=1}^s q_k \delta_{\|\eta_\beta - \eta_k\|_2} \right) \\
 &\leq \max_{\beta \in [K]} \|\eta_\beta - \eta_{\ell_\beta(s)}\|_2 \cdot \max_{j \in [s-1]} \left| \sum_{k=1}^j \left( \frac{m_{\ell_\beta(k)}}{m} - q_{\ell_\beta(k)} \right) \right| \\
 &\leq R \cdot \max_{\beta \in [K]} \max_{j \in [s-1]} \left| \sum_{k=1}^j \left( \frac{m_{\ell_\beta(k)}}{m} - q_{\ell_\beta(k)} \right) \right| =: R \cdot Z_2.
 \end{aligned}$$

Hence,  $|W'(i, j) - \overline{W}(t(i), s(j))| \leq R \cdot (Z_1 + Z_2)$ . Next, we bound  $|W(i, j) - W'(i, j)|$  for any  $(i, j) \in [n] \times [m]$ . By the triangle inequality again, we have

$$\begin{aligned}
 &|W(i, j) - W'(i, j)| \\
 &\leq W_1 \left( \frac{1}{n} \sum_{\ell=1}^n \delta_{\|X_i - X_\ell\|_2}, \frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_{t(i)} - \theta_{t(\ell)}\|_2} \right) + W_1 \left( \frac{1}{m} \sum_{\ell=1}^m \delta_{\|Y_j - Y_\ell\|_2}, \frac{1}{m} \sum_{\ell=1}^m \delta_{\|\eta_{s(j)} - \eta_{s(\ell)}\|_2} \right).
 \end{aligned}$$

Meanwhile, notice that  $X_i = \theta_{t(i)} + \xi_i$  and  $Y_j = \eta_{s(j)} + \zeta_j$  for any  $(i, j) \in [n] \times [m]$ , where  $\xi_i | t(i)$  and  $\zeta_j | s(j)$  are independent sub-Gaussian random vectors. Recall that the Wasserstein-1 distance between two sets of points in  $\mathbb{R}$  with the same cardinality satisfies the following:

$$W_1 \left( \frac{1}{n} \sum_{\ell=1}^n \delta_{\|X_i - X_\ell\|_2}, \frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_{t(i)} - \theta_{t(\ell)}\|_2} \right) \leq \frac{1}{n} \sum_{\ell=1}^n \left| \|X_i - X_\ell\|_2 - \|\theta_{t(i)} - \theta_{t(\ell)}\|_2 \right|.$$

From  $\left| \|X_i - X_\ell\|_2 - \|\theta_{t(i)} - \theta_{t(\ell)}\|_2 \right| = \left| \|\theta_{t(i)} - \theta_{t(\ell)} + \xi_i - \xi_\ell\|_2 - \|\theta_{t(i)} - \theta_{t(\ell)}\|_2 \right| \leq \|\xi_i - \xi_\ell\|_2$ , we have

$$W_1 \left( \frac{1}{n} \sum_{\ell=1}^n \delta_{\|X_i - X_\ell\|_2}, \frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_{t(i)} - \theta_{t(\ell)}\|_2} \right) \leq \frac{1}{n} \sum_{\ell=1}^n \|\xi_i - \xi_\ell\|_2 \leq \max_{i, \ell \in [n]} \|\xi_i - \xi_\ell\|_2 =: U.$$

Similarly,

$$W_1 \left( \frac{1}{m} \sum_{\ell=1}^m \delta_{\|Y_j - Y_\ell\|_2}, \frac{1}{m} \sum_{\ell=1}^m \delta_{\|\eta_{s(j)} - \eta_{s(\ell)}\|_2} \right) \leq \max_{j, \ell \in [m]} \|\zeta_j - \zeta_\ell\|_2 =: V.$$

Hence,  $|W(i, j) - W'(i, j)| \leq U + V$  and  $|W(i, j) - \bar{W}(t(i), s(j))| \leq R(Z_1 + Z_2) + U + V$ . Now, fix  $i \in [n]$  such that  $t(i) \in [K]$ . Then,

$$\begin{aligned}
 & \min_{\substack{j \in [m] \\ s(j) \neq t(i)}} W(i, j) - \min_{\substack{j \in [m] \\ s(j) = t(i)}} W(i, j) \\
 & \geq \min_{\substack{j \in [m] \\ s(j) \neq t(i)}} \bar{W}(t(i), s(j)) - \min_{\substack{j \in [m] \\ s(j) = t(i)}} \bar{W}(t(i), s(j)) - 2R(Z_1 + Z_2) - 2(U + V) \\
 & \geq \min_{\beta \in [s] \setminus \{t(i)\}} \bar{W}(t(i), \beta) - \bar{W}(t(i), t(i)) - 2R(Z_1 + Z_2) - 2(U + V) \tag{24} \\
 & \geq \underbrace{\min_{\substack{\alpha \in [K] \\ \beta \in [s] \setminus \{\alpha\}}} \bar{W}(\alpha, \beta) - R \cdot \left(1 - \sum_{k \in [K]} p_k\right)}_{=: \Omega} - 2R(Z_1 + Z_2) - 2(U + V),
 \end{aligned}$$

where the last inequality follows from Lemma 17. Therefore,  $2R(Z_1 + Z_2) + 2(U + V) < \Omega$  implies that  $s(\pi(i)) = t(i)$  for any  $i \in [n]$  such that  $t(i) \in [K]$ . We show that  $\mathbb{P}(E) \leq \delta$ , where  $E$  is the event  $(2R(Z_1 + Z_2) + 2(U + V) \geq \Omega)$ . By Lemma 19, we have

$$\begin{aligned}
 \mathbb{P}\left(Z_1 \geq \frac{\Omega}{8R}\right) & \leq 2K(t-1) \exp\left(-\frac{2n\Omega^2}{(8R)^2}\right), \\
 \mathbb{P}\left(Z_2 \geq \frac{\Omega}{8R}\right) & \leq 2K(s-1) \exp\left(-\frac{2m\Omega^2}{(8R)^2}\right).
 \end{aligned}$$

Meanwhile, for  $i \neq \ell$ , note that  $\xi_i - \xi_\ell | t(i), t(\ell)$  is a sub-Gaussian random vector with variance proxy  $\sigma_{t(i)}^2 + \sigma_{t(\ell)}^2$ . By Lemma 20, we have

$$\mathbb{P}\left(\max_{i, \ell \in [n]} \|\xi_i - \xi_\ell\|_2 \geq \frac{\Omega}{8}\right) \leq \sum_{i \neq \ell} \mathbb{P}\left(\|\xi_i - \xi_\ell\|_2 \geq \frac{\Omega}{8}\right) \leq n^2 e^{2d} \exp\left(-\frac{(\Omega/8)^2}{16\sigma^2}\right).$$

Similarly, we have

$$\mathbb{P}\left(\max_{j, \ell \in [m]} \|\zeta_j - \zeta_\ell\|_2 \geq \frac{\Omega}{8}\right) \leq m^2 e^{2d} \exp\left(-\frac{(\Omega/8)^2}{16\tau^2}\right).$$

Putting all together, we have

$$\begin{aligned}
 \mathbb{P}(E) & \leq \mathbb{P}\left(Z_1 \geq \frac{\Omega}{4R}\right) + \mathbb{P}\left(Z_2 \geq \frac{\Omega}{4R}\right) + \mathbb{P}\left(U \geq \frac{\Omega}{8}\right) + \mathbb{P}\left(V \geq \frac{\Omega}{8}\right) \\
 & \leq 2K(t+s-2) \exp\left(-\frac{2(n \wedge m)\Omega^2}{(8R)^2}\right) + 2(n \vee m)^2 e^{2d} \exp\left(-\frac{(\Omega/8)^2}{16\Gamma^2}\right) \\
 & \leq \frac{\delta}{2} + \frac{\delta}{2} = \delta,
 \end{aligned}$$

where the last inequality follows from (12). Hence, we have (11).  $\blacksquare$

**Remark 16** From the proof of Theorem 14, one may have noticed that we can replace the left-hand side of (12) with

$$\Omega_o := \min_{\alpha \in [K]} \left( \min_{\beta \in [s] \setminus \{\alpha\}} \overline{W}(\alpha, \beta) - \overline{W}(\alpha, \alpha) \right).$$

To see this, it suffices to observe that we can replace  $\Omega$  in the last line of (24) with  $\Omega_o$ , and the rest of the proof follows the same with  $\Omega_o$  in place of  $\Omega$ . Since  $\Omega_o \geq \Omega$ , the condition (12) is a stronger condition than requiring  $\Omega_o$  to be at least the right-hand side of (12). We chose to present (12) as it is written in terms of the outlier proportion which is more interpretable.

Lastly, we provide the lemmas used in the proof of Theorem 14.

**Lemma 17** For the centers  $\theta_1, \dots, \theta_t$  and  $\eta_1, \dots, \eta_s$  in Assumption 1, define  $\overline{W}$  as in (13). Then, for any  $k \in [K]$ , we have

$$\overline{W}(\alpha, \alpha) \leq \left( 1 - \sum_{k \in [K]} p_k \right) \cdot \max_{(\alpha, \beta) \in [t] \times [s]} \|\theta_\alpha - \eta_\beta\|_2. \quad (25)$$

**Proof** The Wasserstein-1 distance on  $\mathcal{P}(\mathbb{R})$  can be written as the  $L^1$  distance between cumulative distribution functions. For  $\alpha \in [K]$ , we have

$$\begin{aligned} \overline{W}(\alpha, \alpha) &= \int_{\mathbb{R}} \left| \sum_{k=1}^t p_k 1\{\|\theta_\alpha - \theta_k\|_2 \leq z\} - \sum_{k=1}^s q_k 1\{\|\eta_\alpha - \eta_k\|_2 \leq z\} \right| dz \\ &= \int_{\mathbb{R}} \left| \sum_{k \in [t] \setminus [K]} p_k 1\{\|\theta_\alpha - \theta_k\|_2 \leq z\} - \sum_{k \in [s] \setminus [K]} q_k 1\{\|\theta_\alpha - \eta_k\|_2 \leq z\} \right| dz \\ &= (1 - \lambda) W_1 \left( \sum_{k \in [t] \setminus [K]} \frac{p_k}{1 - \lambda} \delta_{\|\theta_\alpha - \theta_k\|_2}, \sum_{k \in [s] \setminus [K]} \frac{q_k}{1 - \lambda} \delta_{\|\theta_\alpha - \eta_k\|_2} \right), \end{aligned}$$

where  $\lambda = \sum_{k \in [K]} p_k$ . Since  $\|\theta_\alpha - \theta_k\|_2 - \|\theta_\alpha - \eta_\ell\|_2 \leq \|\theta_k - \eta_\ell\|_2$  for any  $k \in [t] \setminus [K]$  and  $\ell \in [s] \setminus [K]$ , one can deduce that (25) holds.  $\blacksquare$

**Lemma 18** Let  $z_1 \leq \dots \leq z_t$  be any real numbers. For any  $(p_1, \dots, p_t), (q_1, \dots, q_t) \in \Delta_t$ , we have

$$W_1 \left( \sum_{k=1}^t p_k \delta_{z_k}, \sum_{k=1}^t q_k \delta_{z_k} \right) \leq (z_t - z_1) \cdot \max_{j \in [t-1]} \left| \sum_{k=1}^j (p_k - q_k) \right|. \quad (26)$$

**Proof** Notice that

$$W_1 \left( \sum_{k=1}^t p_k \delta_{z_k}, \sum_{k=1}^t q_k \delta_{z_k} \right) = \int_{\mathbb{R}} \left| \sum_{k=1}^t (p_k - q_k) 1\{z_k \leq z\} \right| dz = \sum_{j=1}^{t-1} \left| \sum_{k=1}^j (p_k - q_k) \right| \cdot (z_{j+1} - z_j),$$

where the last equality follows from

$$\left| \sum_{k=1}^t (p_k - q_k) 1\{z_k \leq z\} \right| = \sum_{j=1}^{t-1} \left| \sum_{k=1}^j (p_k - q_k) \right| \cdot 1\{z_j \leq z < z_{j+1}\} \quad \forall z \in \mathbb{R}.$$

Therefore, we have (26). ■

### A.6 Proof of Theorem 15

**Proof** For the proof, we may assume that  $\pi^*$  is the identity without loss of generality. First, letting  $Q(\pi) = \sum_{i=1}^n W_1(\hat{\mu}_i, \hat{\nu}_{\pi(i)})$ , observe that

$$(\hat{\pi} \neq \pi^*) \subset \bigcup_{\pi \in \mathcal{S}_n \setminus \{\pi^*\}} (Q(\pi^*) \geq Q(\pi)).$$

The event  $Q(\pi^*) \geq Q(\pi)$  is equivalent to

$$\sum_{i=1}^n (W_1(\hat{\mu}_i, \hat{\nu}_i) - W_1(\hat{\mu}_i, \hat{\nu}_{\pi(i)})) = \sum_{i: \pi(i) \neq i} (W_1(\hat{\mu}_i, \hat{\nu}_i) - W_1(\hat{\mu}_i, \hat{\nu}_{\pi(i)})) \geq 0,$$

which is contained in the following events:

$$\bigcup_{i: \pi(i) \neq i} (W_1(\hat{\mu}_i, \hat{\nu}_i) \geq W_1(\hat{\mu}_i, \hat{\nu}_{\pi(i)})) \subset \bigcup_{i=1}^n \bigcup_{j \neq i} (W_1(\hat{\mu}_i, \hat{\nu}_i) \geq W_1(\hat{\mu}_i, \hat{\nu}_j))$$

Therefore,

$$(\hat{\pi} \neq \pi^*) \subset \bigcup_{i=1}^n \bigcup_{j \neq i} (W_1(\hat{\mu}_i, \hat{\nu}_i) \geq W_1(\hat{\mu}_i, \hat{\nu}_j)). \quad (27)$$

Now, recall that we can write a rigid transformation  $T: \mathbb{R}^d \rightarrow \mathbb{R}^d$  as  $T(\theta) = R\theta + b$  for some orthonormal matrix  $R \in \mathbb{R}^{d \times d}$  and some vector  $b \in \mathbb{R}^d$ . Hence, we have

$$\begin{aligned} \hat{\mu}_i &= \frac{1}{n} \sum_{\ell=1}^n \delta_{\|X_i - X_\ell\|_2} = \frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_i - \theta_\ell + \xi_i - \xi_\ell\|_2} = \frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_i - \theta_\ell + \xi_{i\ell}\|_2}, \\ \hat{\nu}_i &= \frac{1}{n} \sum_{\ell=1}^n \delta_{\|Y_i - Y_\ell\|_2} = \frac{1}{n} \sum_{\ell=1}^n \delta_{\|R(\theta_i - \theta_\ell) + \zeta_i - \zeta_\ell\|_2} = \frac{1}{n} \sum_{\ell=1}^n \delta_{\|\theta_i - \theta_\ell + \zeta_{i\ell}\|_2}, \end{aligned}$$

where we let  $\xi_{i\ell} = \xi_i - \xi_\ell$  and  $\zeta_{i\ell} = R^\top(\zeta_i - \zeta_\ell)$ . Next, recall that the Wasserstein-1 distance between two sets of points in  $\mathbb{R}$  with the same cardinality satisfies the following:

$$W_1(\hat{\mu}_i, \hat{\nu}_j) = \min_{\pi \in \mathcal{S}_n} \frac{1}{n} \sum_{\ell=1}^n \left| \|\theta_i - \theta_\ell + \xi_{i\ell}\|_2 - \|\theta_j - \theta_{\pi(\ell)} + \zeta_{j\pi(\ell)}\|_2 \right|.$$

Note that

$$\begin{aligned}
 & \left| \|\theta_i - \theta_\ell + \xi_{i\ell}\|_2 - \|\theta_j - \theta_{\pi(\ell)} + \zeta_{j\pi(\ell)}\|_2 \right| - \left| \|\theta_i - \theta_\ell\|_2 - \|\theta_j - \theta_{\pi(\ell)}\|_2 \right| \\
 & \leq \left| \|\theta_i - \theta_\ell + \xi_{i\ell}\|_2 - \|\theta_i - \theta_\ell\|_2 \right| + \left| \|\theta_j - \theta_{\pi(\ell)}\|_2 - \|\theta_j - \theta_{\pi(\ell)} + \zeta_{j\pi(\ell)}\|_2 \right| \\
 & \leq \|\xi_{i\ell}\|_2 + \|\zeta_{j\pi(\ell)}\|_2 \\
 & \leq (U + V),
 \end{aligned}$$

where we define  $U = \max_{i,\ell \in [n]} \|\xi_{i\ell}\|_2$  and  $V = \max_{j,\ell \in [n]} \|\zeta_{j\ell}\|_2$ . As

$$W_1(\mu_i, \mu_j) = \min_{\pi \in \mathcal{S}_n} \frac{1}{n} \sum_{\ell=1}^n \left| \|\theta_i - \theta_\ell\|_2^2 - \|\theta_j - \theta_{\pi(\ell)}\|_2^2 \right|,$$

we conclude that

$$|W_1(\hat{\mu}_i, \hat{\nu}_j) - W_1(\mu_i, \mu_j)| \leq (U + V) \quad \forall i, j \in [n].$$

Therefore, if we can find  $i, j \in [n]$  such that  $i \neq j$  and  $W_1(\hat{\mu}_i, \hat{\nu}_i) \geq W_1(\hat{\mu}_i, \hat{\nu}_j)$ , we must have

$$(U + V) \geq W_1(\hat{\mu}_i, \hat{\nu}_i) \geq W_1(\hat{\mu}_i, \hat{\nu}_j) \geq W_1(\mu_i, \mu_j) - (U + V) \geq \Phi - (U + V).$$

Accordingly, by (27), we have

$$\mathbb{P}(\hat{\pi} \neq \pi^*) \leq \mathbb{P}(2(U + V) \geq \Phi) \leq \mathbb{P}\left(U \geq \frac{\Phi}{4}\right) + \mathbb{P}\left(V \geq \frac{\Phi}{4}\right).$$

Since  $\xi_{i\ell} = \xi_i - \xi_\ell$  is a sub-Gaussian random vector with variance proxy  $\sigma_i^2 + \sigma_\ell^2$  for  $i \neq \ell$ , one can deduce from Lemma 20 and (17) that

$$\mathbb{P}\left(\max_{i,\ell \in [n]} \|\xi_{i\ell}\|_2 \geq \frac{\Phi}{4}\right) \leq \sum_{i \neq \ell} \mathbb{P}\left(\|\xi_{i\ell}\|_2 \geq \frac{\Phi}{4}\right) \leq n^2 e^{2d} \exp\left(-\frac{(\Phi/4)^2}{16\sigma^2}\right) \leq \frac{\delta}{2},$$

where  $\sigma = \max\{\sigma_1, \dots, \sigma_n\}$ . Similarly, we have

$$\mathbb{P}\left(\max_{j,\ell \in [n]} \|\zeta_{j\ell}\|_2 \geq \frac{\Phi}{4}\right) \leq \frac{\delta}{2}.$$

Therefore, we have  $\mathbb{P}(\hat{\pi} \neq \pi^*) \leq \delta$ . ■

## Appendix B. Lemmas

We have used the following lemmas on the concentration of the multinomial distribution and the sub-Gaussian distribution in the proofs of Theorems 14 and 15. The former is an application of Proposition A.6.6 of Vaart and Wellner (1996), while the latter is a standard result that can be deduced from the Chernoff bound; we provide the proof for completeness.

**Lemma 19** Suppose  $(z_1, \dots, z_t)$  is from the multinomial distribution with the number of trials  $n$  and the probability vector  $(p_1, \dots, p_t) \in \Delta_t$ . Let  $\mathcal{A} \subset 2^{[t]}$  be a collection of subsets of  $[t]$ . Then, for any  $\varepsilon \in (0, 1)$ , we have

$$\mathbb{P} \left( \max_{A \in \mathcal{A}} \left| \sum_{k \in A} \left( \frac{z_k}{n} - p_k \right) \right| \geq \varepsilon \right) \leq 2|\mathcal{A}| \cdot e^{-2n\varepsilon^2}.$$

**Lemma 20** We say a probability measure  $\mu \in \mathcal{P}(\mathbb{R}^d)$  is sub-Gaussian with variance proxy  $\sigma^2 > 0$  if the following holds for  $X \sim \mu$ :

$$\mathbb{E} e^{u^\top (X - \mathbb{E}X)} \leq \exp \left( \frac{\sigma^2 \|u\|_2^2}{2} \right) \quad \forall u \in \mathbb{R}^d.$$

In this case, for any  $t \geq 0$ , we have

$$\mathbb{P} (\|X - \mathbb{E}X\|_2 \geq t) \leq e^{2d} \exp \left( -\frac{t^2}{8\sigma^2} \right). \quad (28)$$

**Proof** Without loss of generality, we may assume  $\mathbb{E}X = 0$ . Let  $B_r = \{x \in \mathbb{R}^d : \|x\|_2 \leq r\}$  be the closed ball of radius  $r$  centered at the origin, and recall that  $\|x\|_2 = r \cdot \sup_{u \in B_r} u^\top x$  for any  $x \in \mathbb{R}^d$ . Now, for  $\varepsilon \in (0, 1)$ , let  $N_\varepsilon$  be a minimal  $\varepsilon$ -covering of  $B_1$ . For any  $u \in B_1$ , we can find  $z_u \in N_\varepsilon$  such that  $\|u - z_u\|_2 \leq \varepsilon$ . Hence, for any  $u \in B_1$ , we have

$$u^\top x \leq (u - z_u)^\top x + z_u^\top x \leq \sup_{z \in B_\varepsilon} z^\top x + \max_{z \in N_\varepsilon} z^\top x \leq \varepsilon \|x\|_2 + \max_{z \in N_\varepsilon} z^\top x,$$

which implies

$$\|x\|_2 \leq \frac{1}{1 - \varepsilon} \max_{z \in N_\varepsilon} z^\top x.$$

For any  $z \in N_\varepsilon \subset B_1$ , note that  $z^\top X$  is a sub-Gaussian random variable with variance proxy  $\sigma^2$ . Hence, for any  $t \geq 0$ , we have

$$\mathbb{P}(\|X\|_2 \geq t) \leq \mathbb{P}(\max_{z \in N_\varepsilon} z^\top X \geq (1 - \varepsilon)t) \leq |N_\varepsilon| \cdot \exp \left( -\frac{(1 - \varepsilon)^2 t^2}{2\sigma^2} \right).$$

By taking  $\varepsilon = 1/2$  and using  $|N_\varepsilon| \leq (1 + 2/\varepsilon)^d$ , we have (28); we use  $|N_{1/2}| \leq 5^d \leq e^{2d}$ . ■

## Appendix C. Extensions

### C.1 Extension of Theorem 7 to General Metric Measure Spaces

From (20) in the proof of Theorem 7, one can deduce that Theorem 7 can be extended to any metric measure space  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  by replacing  $d + 1$  in (7) with the VC dimension of  $\mathcal{F}_{\mathcal{X}}$  defined in the proof, which is the VC dimension of the set of all balls in  $\mathcal{X}$ . Unfortunately, controlling the VC dimension of the set of all balls in a general metric measure space is highly nontrivial. A more straightforward quantity is the  $\varepsilon$ -covering number  $\mathcal{N}(\varepsilon, \mathcal{X}, d_{\mathcal{X}})$ , namely, the smallest possible number of  $\varepsilon$ -balls to cover  $\mathcal{X}$ . The following theorem derives an upper bound using the covering number for a general bounded metric measure space  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$ .

**Theorem 21** Let  $(\mathcal{X}, d_{\mathcal{X}}, \mu)$  be a metric measure space and  $\hat{\mu} = \frac{1}{n} \sum_{i=1}^n \delta_{X_i}$  be the empirical measure based on  $X_1, \dots, X_n$  that are i.i.d. from  $\mu$ . Suppose  $\Delta = \sup_{x, x' \in \mathcal{X}} d_{\mathcal{X}}(x, x') < \infty$  and the support of  $\mu$  is exactly  $\mathcal{X}$ . Let  $\mathcal{N}(\varepsilon, \mathcal{X}, d_{\mathcal{X}})$  be the  $\varepsilon$ -covering number of  $\mathcal{X}$  under  $d_{\mathcal{X}}$ . Then,

$$\mathbb{E} \left[ \sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) \right] \leq \frac{C}{\sqrt{n}} \int_0^{\Delta/4} \sqrt{\frac{\Delta}{\varepsilon} + \log \mathcal{N}(\varepsilon, \mathcal{X}, d_{\mathcal{X}})} d\varepsilon, \quad (29)$$

where  $C$  is an absolute constant that is independent of  $n$ ,  $\mathcal{X}$ , and  $\mu$ .

**Proof** By the Kantorovich-Rubinstein theorem, e.g., Theorem 1.14 of Villani (2003), we have

$$W_1(\mu_x, \hat{\mu}_x) = \sup_{f \in \text{Lip}_1} (\mu_x f - \hat{\mu}_x f),$$

where  $\text{Lip}_1 = \{f: [0, \Delta] \rightarrow \mathbb{R} \mid f(0) = 0 \text{ and } f \text{ is 1-Lipschitz}\}$ . Here, we have used the fact that the supports of  $\mu_x$  and  $\hat{\mu}_x$  are contained in the interval  $[0, \Delta]$  for any  $x \in \mathcal{X}$ . Hence,

$$\sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) = \sup_{x \in \mathcal{X}} \sup_{f \in \text{Lip}_1} (\mu_x f - \hat{\mu}_x f) = \sup_{\phi \in \mathcal{F}} (\mu \phi - \hat{\mu} \phi),$$

where  $\mathcal{F} = \{f \circ d_{\mathcal{X}, x} : f \in \text{Lip}_1, x \in \mathcal{X}\}$  which is a function class from  $\mathcal{X}$  to  $\mathbb{R}$ . Here,  $d_{\mathcal{X}, x}: \mathcal{X} \rightarrow \mathbb{R}$  is defined by  $d_{\mathcal{X}, x}(\cdot) = d_{\mathcal{X}}(x, \cdot)$ . Using symmetrization and Dudley's chaining argument (see Chapter 8 of Vershynin (2018)), one can derive the following:

$$\mathbb{E} \left[ \sup_{x \in \mathcal{X}} W_1(\mu_x, \hat{\mu}_x) \right] = \mathbb{E} \left[ \sup_{\phi \in \mathcal{F}} (\mu \phi - \hat{\mu} \phi) \right] \leq \frac{C}{\sqrt{n}} \int_0^{\Delta/2} \sqrt{\log \mathcal{N}(\varepsilon, \mathcal{F}, \|\cdot\|_{\mathcal{X}})} d\varepsilon, \quad (30)$$

where  $\mathcal{N}(\varepsilon, \mathcal{F}, \|\cdot\|_{\mathcal{X}})$  denotes the  $\varepsilon$ -covering number of  $\mathcal{F}$  with respect to  $\|\cdot\|_{\mathcal{X}}$ , namely, the sup norm defined by  $\|\phi\|_{\mathcal{X}} = \sup_{x \in \mathcal{X}} |\phi(x)|$  for any  $\phi: \mathcal{X} \rightarrow \mathbb{R}$ . Here, we are using the fact that  $\sup_{\phi \in \mathcal{F}} \|\phi\|_{\mathcal{X}} \leq \Delta$  since  $\sup_{z \in [0, \Delta]} |f(z)| \leq \Delta$  for any  $f \in \text{Lip}_1$  by Lipschitzness. Now, noticing that  $\mathcal{F}$  is a function class consisting of composite functions, we upper bound  $\mathcal{N}(\varepsilon, \mathcal{F}, \|\cdot\|_{\mathcal{X}})$  by the product of the  $\varepsilon$ -covering number of  $\text{Lip}_1$  with respect to the sup norm  $\|\cdot\|_{\infty}$  on  $[0, \Delta]$ , that is,  $\|f\|_{\infty} = \sup_{z \in [0, \Delta]} |f(z)|$ , and the  $\varepsilon$ -covering number of  $\mathcal{X}$  with respect to  $d_{\mathcal{X}}$ . Concretely, we claim

$$\mathcal{N}(2\varepsilon, \mathcal{F}, \|\cdot\|_{\mathcal{X}}) \leq \mathcal{N}(\varepsilon, \text{Lip}_1, \|\cdot\|_{\infty}) \times \mathcal{N}(\varepsilon, \mathcal{X}, d_{\mathcal{X}}). \quad (31)$$

To see this, let  $\text{Lip}_1^{\varepsilon} \subset \text{Lip}_1$  and  $\mathcal{X}^{\varepsilon} \subset \mathcal{X}$  be  $\varepsilon$ -coverings of  $\text{Lip}_1$  and  $\mathcal{X}$ , respectively. For  $f \in \text{Lip}_1$  and  $x \in \mathcal{X}$ , we can find  $f' \in \text{Lip}_1^{\varepsilon}$  and  $x' \in \mathcal{X}^{\varepsilon}$  such that  $\|f - f'\|_{\infty} \leq \varepsilon$  and  $d_{\mathcal{X}}(x, x') \leq \varepsilon$ . Then,

$$\begin{aligned} \|f \circ d_{\mathcal{X}, x} - f' \circ d_{\mathcal{X}, x'}\|_{\mathcal{X}} &\leq \|f \circ d_{\mathcal{X}, x} - f \circ d_{\mathcal{X}, x'}\|_{\mathcal{X}} + \|f \circ d_{\mathcal{X}, x'} - f' \circ d_{\mathcal{X}, x'}\|_{\mathcal{X}} \\ &\leq \|d_{\mathcal{X}, x} - d_{\mathcal{X}, x'}\|_{\mathcal{X}} + \|f - f'\|_{\infty} \\ &\leq d_{\mathcal{X}}(x, x') + \varepsilon \\ &\leq 2\varepsilon. \end{aligned}$$

Hence, we obtain (31). The standard upper bound on the covering number of  $\text{Lip}_1$  (see Exercise 8.2.7 of Vershynin (2018)) tells

$$\log \mathcal{N}(\varepsilon, \text{Lip}_1, \|\cdot\|_{\infty}) \leq \frac{c\Delta}{\varepsilon}, \quad (32)$$

where  $c$  is some absolute constant. Combining (30), (31), and (32), we obtain (29).  $\blacksquare$

**Remark 22** *If one assumes that  $\mathcal{X}$  is a compact subset of  $\mathbb{R}^d$  and  $d_{\mathcal{X}}$  is the standard Euclidean distance in Theorem 21, the bounds coincide with those in Theorem 7. To see this, it suffices to notice that  $\mathcal{N}(\varepsilon, \mathcal{X}, d_{\mathcal{X}})$  scales as  $(1 + 2/\varepsilon)^d$ ; see Section 4.2 of Vershynin (2018). The key difference, however, is that Theorem 21 requires compactness of the support of  $\mu$ , while Theorem 7 does not.*

## C.2 Extension of Theorem 14: Identifying Points from the Common Part

Section 3.2 focuses on the matching of the points from the common part, namely, the first  $K$  components. The matching guarantee ensures that all  $X_i$ 's such that  $t(i) \in [K]$  are correctly matched, while ignoring the points from the outlier components. If we assume a stronger condition by replacing the left-hand side of (12) with the following quantity

$$\min_{\substack{\alpha \in [K] \\ \beta \in [s] \setminus \{\alpha\}}} \overline{W}(\alpha, \beta) \wedge \min_{\substack{\alpha \in [t] \setminus [K] \\ \beta \in [s]}} \overline{W}(\alpha, \beta) - R \cdot \left(1 - \sum_{k \in [K]} p_k\right), \quad (33)$$

we can show that

$$\max_{\substack{i \in [n] \\ t(i) \in [K]}} W(i, \pi(i)) < \min_{\substack{i \in [n] \\ t(i) \notin [K]}} W(i, \pi(i)) \quad (34)$$

holds with probability at least  $1 - \delta$ . This means that if the threshold  $\rho$  in Algorithm 1 is chosen between the two quantities in (34), the output  $I$  of the algorithm exactly coincides with the set  $\{i \in [n] : t(i) \in [K]\}$ , thereby identifying the points from the common part.

Let us prove (34). To this end, first recall from the proof of Theorem 14 that the event  $C = (2R(Z_1 + Z_2) + 2(U + V) < \Omega)$  implies  $s(\pi(i)) = t(i)$  for any  $i \in [n]$  such that  $t(i) \in [K]$ . Hence, on the event  $C$ , one can deduce that for such  $i \in [n]$ , we have

$$W(i, \pi(i)) \leq \overline{W}(t(i), t(i)) + R(Z_1 + Z_2) + U + V,$$

which follows from  $|W(i, j) - \overline{W}(t(i), s(j))| \leq R(Z_1 + Z_2) + (U + V)$  for any  $i, j \in [n]$ . Therefore, invoking Lemma 17, we have

$$\max_{\substack{i \in [n] \\ t(i) \in [K]}} W(i, \pi(i)) \leq R \cdot \left(1 - \sum_{k \in [K]} p_k\right) + R(Z_1 + Z_2) + (U + V).$$

Meanwhile,

$$\begin{aligned} \min_{\substack{i \in [n] \\ t(i) \notin [K]}} W(i, \pi(i)) &\geq \min_{\substack{i \in [n] \\ t(i) \notin [K]}} \overline{W}(t(i), s(\pi(i))) - R(Z_1 + Z_2) - (U + V) \\ &\geq \min_{\substack{\alpha \in [t] \setminus [K] \\ \beta \in [s]}} \overline{W}(\alpha, \beta) - R(Z_1 + Z_2) - (U + V). \end{aligned}$$

Now, let  $C'$  be the event  $(2R(Z_1 + Z_2) + 2(U + V) < \Omega')$ , where  $\Omega'$  represents the quantity in (33). Clearly, we have  $\Omega' \leq \Omega$  and thus  $C' \subset C$ . Therefore, on the event  $C'$ , we have

$$\min_{\substack{i \in [n] \\ t(i) \notin [K]}} W(i, \pi(i)) - \max_{\substack{i \in [n] \\ t(i) \in [K]}} W(i, \pi(i)) \geq \Omega' - 2R(Z_1 + Z_2) - 2(U + V) > 0.$$

Now, we have  $\mathbb{P}(C') \geq 1 - \delta$  by repeating the proof of Theorem 14 with the left-hand side of (12), which is  $\Omega$ , replaced by  $\Omega'$ . Therefore, we have (34) with probability at least  $1 - \delta$ .

## References

- Yonathan Aflalo, Alexander Bronstein, and Ron Kimmel. On convex relaxation of graph isomorphism. *Proceedings of the National Academy of Sciences*, 112(10):2942–2947, 2015.
- David Alvarez-Melis and Tommi S Jaakkola. Gromov-Wasserstein alignment of word embedding spaces. *arXiv preprint arXiv:1809.00013*, 2018.
- Luigi Ambrosio and Nicola Gigli. *A User’s Guide to Optimal Transport*, pages 1–155. Springer Berlin Heidelberg, 2013.
- Xiao-chen Bai, Chuangye Yan, Guanghui Yang, Peilong Lu, Dan Ma, Linfeng Sun, Rui Zhou, Sjors HW Scheres, and Yigong Shi. An atomic structure of human  $\gamma$ -secretase. *Nature*, 525(7568):212–217, 2015.
- Serge Belongie, Jitendra Malik, and Jan Puzicha. Shape matching and object recognition using shape contexts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(4):509–522, 2002.
- Tamir Bendory, Alberto Bartschaggi, and Amit Singer. Single-particle cryo-electron microscopy: Mathematical theory, computational challenges, and opportunities. *IEEE Signal Processing Magazine*, 37(2):58–76, 2020.
- P.J. Besl and Neil D. McKay. A method for registration of 3-d shapes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 14(2):239–256, 1992.
- Emmanuel Boissard and Thibaut Le Gouic. On the mean speed of convergence of empirical and occupation measures in Wasserstein distance. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 50(2):539 – 563, 2014.
- Stephen P Borgatti, Ajay Mehra, Daniel J Brass, and Giuseppe Labianca. Network analysis in the social sciences. *Science*, 323(5916):892–895, 2009.
- Claire BréchetEAU. A statistical test of isomorphism between metric-measure spaces using the distance-to-a-measure signature. *Electronic Journal of Statistics*, 13(1):795 – 849, 2019.
- Daniel Brinkman and Peter J Olver. Invariant histograms. *The American Mathematical Monthly*, 119(1):4–24, 2012.

- Charlotte Bunne, David Alvarez-Melis, Andreas Krause, and Stefanie Jegelka. Learning generative models across incomparable spaces. In *International Conference on Machine Learning*, 2019.
- Olivier Collier and Arnak S Dalalyan. Minimax rates in permutation estimation for feature matching. *Journal of Machine Learning Research*, 17(1):162–192, 2016.
- Steffen Dereich, Michael Scheutzow, and Reik Schottstedt. Constructive quantization: Approximation by empirical measures. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 49(4):1183–1203, 2013.
- Jian Ding, Zongming Ma, Yihong Wu, and Jiaming Xu. Efficient random graph matching via degree profiles. *Probability Theory and Related Fields*, 179:29–115, 2021.
- Jian Ding, Yumou Fei, and Yuanzheng Wang. Efficiently matching random inhomogeneous graphs via degree profiles. *arXiv preprint arXiv:2310.10441*, 2023.
- Richard M Dudley. Balls in  $\mathbb{R}^k$  do not cut all subsets of  $k + 2$  points. *Advances in Mathematics*, 31(3):306–308, 1979.
- Richard Mansfield Dudley. The speed of mean Glivenko-Cantelli convergence. *The Annals of Mathematical Statistics*, 40(1):40–50, 1969.
- Nicolas Fournier and Arnaud Guillin. On the rate of convergence in Wasserstein distance of the empirical measure. *Probability Theory and Related Fields*, 162(3-4):707–738, 2015.
- Tigran Galstyan and Arshak Minasyan. Optimality of the least sum of logarithms in the problem of matching map recovery in the presence of noise and outliers. *Armenian Journal of Mathematics*, 15(5):1–9, 2023.
- Tigran Galstyan, Arshak Minasyan, and Arnak S Dalalyan. Optimal detection of the feature matching map in presence of noise and outliers. *Electronic Journal of Statistics*, 16(2): 5720–5750, 2022.
- Steven J Gortler, Louis Theran, and Dylan P Thurston. Generic unlabeled global rigidity. *Forum of Mathematics, Sigma*, 7:e21, 2019.
- Yael Harpaz and Yoel Shkolnisky. Three-dimensional alignment of density maps in cryo-electron microscopy. *Biological Imaging*, 3:e8, 2023.
- YoonHaeng Hur, Wenxuan Guo, and Tengyuan Liang. Reversible Gromov–Monge sampler for simulation-based inference. *SIAM Journal on Mathematics of Data Science*, 6(2): 283–310, 2024.
- Takayuki Iguchi, Dustin G Mixon, Jesse Peterson, and Soledad Villar. Probably certifiably correct k-means clustering. *Mathematical Programming*, 165:605–642, 2017.
- Itay Kezurer, Shahar Z Kovalsky, Ronen Basri, and Yaron Lipman. Tight relaxation of quadratic matching. *Computer Graphics Forum*, 34(5):115–128, 2015.

- Tjalling C Koopmans and Martin Beckmann. Assignment problems and the location of economic activities. *Econometrica*, pages 53–76, 1957.
- Dmitriy Kunisky and Jonathan Niles-Weed. Strong recovery of geometric planted matchings. In *Symposium on Discrete Algorithms*, 2022.
- Marius Leordeanu and Martial Hebert. A spectral technique for correspondence problems using pairwise constraints. In *International Conference on Computer Vision*, 2005.
- Eliane Maria Loiola, Nair Maria Maia De Abreu, Paulo Oswaldo Boaventura-Netto, Peter Hahn, and Tania Querido. A survey for the quadratic assignment problem. *European Journal of Operational Research*, 176(2):657–690, 2007.
- Siddharth Manay, Daniel Cremers, Byung-Woo Hong, Anthony J Yezzi, and Stefano Soatto. Integral invariants for shape matching. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(10):1602–1618, 2006.
- Thomas Martinetz and Klaus Schulten. Topology representing networks. *Neural Networks*, 7(3):507–522, 1994.
- Facundo Mémoli. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of Computational Mathematics*, 11:417–487, 2011.
- Facundo Mémoli and Tom Needham. Distance distributions and inverse problems for metric measure spaces. *Studies in Applied Mathematics*, 149(4):943–1001, 2022.
- Arshak Minasyan, Tigran Galstyan, Sona Hunanyan, and Arnak Dalalyan. Matching map recovery with an unknown number of outliers. In *International Conference on Artificial Intelligence and Statistics*, 2023.
- Abhinav Nellore and Rachel Ward. Recovery guarantees for exemplar-based clustering. *Information and Computation*, 245:165–180, 2015.
- Antonio Norelli, Marco Fumero, Valentino Maiorca, Luca Moschella, Emanuele Rodola, and Francesco Locatello. ASIF: Coupled data turns unimodal models to multimodal without training. In *Advances in Neural Information Processing Systems*, 2023.
- Panagiotis Papadakis. The canonically posed 3d objects dataset. In *Eurographics Workshop on 3D Object Retrieval*, 2014.
- Gabriel Peyré and Marco Cuturi. Computational optimal transport: with applications to data science. *Foundations and Trends® in Machine Learning*, 11(5-6):355–607, 2019.
- Gabriel Peyré, Marco Cuturi, and Justin Solomon. Gromov-Wasserstein averaging of kernel and distance matrices. In *International Conference on Machine Learning*, 2016.
- Anand Rangarajan, Alan Yuille, and Eric Mjolsness. Convergence properties of the softassign quadratic assignment algorithm. *Neural Computation*, 11(6):1455–1474, 1999.

- Aryan Tajmir Riahi, Geoffrey Woollard, Frédéric Poitevin, Anne Condon, and Khanh Dao Duc. Alignot: An optimal transport based algorithm for fast 3d alignment with applications to cryogenic electron microscopy density maps. *IEEE/ACM Transactions on Computational Biology and Bioinformatics*, 20(6):3842–3850, 2023.
- Aryan Tajmir Riahi, Chenwei Zhang, Anne Condon, James Chen, and Khanh Dao Duc. Alignment of partially overlapping cryo-em maps using unbalanced Gromov-Wasserstein divergence. *PRX Life*, 3(2):023003, 2025.
- Mauro R Ruggeri and Dietmar Saupe. Isometry-invariant matching of point set surfaces. In *Eurographics Workshop on 3D Object Retrieval*, 2008.
- Yusuf Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020.
- Amit Singer and Ruiyi Yang. Alignment of density maps in Wasserstein distance. *Biological Imaging*, 4:e5, 2024.
- Justin Solomon, Gabriel Peyré, Vladimir G Kim, and Suvrit Sra. Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (ToG)*, 35(4):1–13, 2016.
- Olaf Sporns, Giulio Tononi, and Rolf Kötter. The human connectome: a structural description of the human brain. *PLoS Computational Biology*, 1(4):e42, 2005.
- Aad W. Vaart and Jon A. Wellner. *Weak Convergence and Empirical Processes: With Applications to Statistics*. Springer, 1996.
- Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. *Computer Graphics Forum*, 30(6):1681–1707, 2011.
- Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge University Press, 2018.
- Cédric Villani. *Topics in Optimal Transportation*. American Mathematical Society, 2003.
- Grace Wahba. A least squares estimate of satellite attitude. *SIAM Review*, 7(3):409–409, 1965.
- Haoyu Wang, Yihong Wu, Jiaming Xu, and Israel Yolou. Random graph matching in geometric models: the case of complete graphs. In *Conference on Learning Theory*, 2022.
- Jonathan Weed and Francis Bach. Sharp asymptotic and finite-sample rates of convergence of empirical measures in Wasserstein distance. *Bernoulli*, 25(4A):2620–2648, 2019.
- Christoph Alexander Weitkamp, Katharina Proksch, Carla Taveling, and Axel Munk. Distribution of distances based object matching: asymptotic inference. *Journal of the American Statistical Association*, pages 1–14, 2022.
- Jianping Wu, Zhen Yan, Zhangqiang Li, Xingyang Qian, Shan Lu, Mengqiu Dong, Qiang Zhou, and Nieng Yan. Structure of the voltage-gated calcium channel Ca<sub>v</sub>1.1 at 3.6 Å resolution. *Nature*, 537(7619):191–196, 2016.

Heng Yang and Luca Carlone. A quaternion-based certifiably optimal solution to the Wahba problem with outliers. In *International Conference on Computer Vision*, 2019.

Qing Zhao, Stefan E Karisch, Franz Rendl, and Henry Wolkowicz. Semidefinite programming relaxations for the quadratic assignment problem. *Journal of Combinatorial Optimization*, 2:71–109, 1998.