

# Fused Gromov-Wasserstein Variance Decomposition with Linear Optimal Transport

Michael Wilson, Tom Needham and Anuj Srivastava

**Abstract**—Wasserstein distances form a family of metrics on spaces of probability measures that have recently seen many applications. However, statistical analysis in these spaces is complex due to the nonlinearity of Wasserstein spaces. One potential solution to this problem is Linear Optimal Transport (LOT). This method allows one to find a Euclidean embedding, called *LOT embedding*, of measures in some Wasserstein spaces, but some information is lost in this embedding. So, to understand whether statistical analysis relying on LOT embeddings can make valid inferences about original data, it is helpful to quantify how well these embeddings describe that data. To answer this question, we present a decomposition of the Fréchet variance of a set of measures in the 2-Wasserstein space, which allows one to compute the percentage of variance explained by LOT embeddings of those measures. We then extend this decomposition to the Fused Gromov-Wasserstein setting. We also present several experiments that explore the relationship between the dimension of the LOT embedding, the percentage of variance explained by the embedding, and the classification accuracy of machine learning classifiers built on the embedded data. We use the MNIST handwritten digits dataset, IMDB-50000 dataset, and Diffusion Tensor MRI images for these experiments. Our results illustrate the effectiveness of low dimensional LOT embeddings in terms of the percentage of variance explained and the classification accuracy of models built on the embedded data.

## I. INTRODUCTION

Probability measures are objects of fundamental importance in statistics, and Optimal Transport (OT) theory [52] provides a powerful framework for studying them. Recently, computational methods for optimal transport [18], [42] have found diverse applications in machine learning and data science [13], [39]. These include applications in computer vision [8], as well as in generative modeling [5], [29], [44] and domain adaptation [27]. Variants of OT have also been applied to diffeomorphic image registration with Wasserstein-Fisher Rao metrics [14], [24], [25], comparison of Gaussian mixture measures [20], [54], and shape or graph matching with Gromov-Wasserstein and Fused Gromov-Wasserstein distances [15]–[17], [36], [37], [50], [51].

This paper focuses on the statistical analysis of datasets where each observation is a point cloud or graph, represented as an empirical probability measure in a metric space. A significant challenge in this approach is that such data objects have no canonical Euclidean representation and thus do not lend themselves to statistical techniques designed for Euclidean data. One potential solution to this problem, called *Linear Optimal Transport* (LOT) [53], offers an embedding of measures supported on  $\mathbb{R}^d$  (and even some Riemannian

manifolds [45]) into a given Euclidean space. However, given that the original data are not Euclidean, one should expect some information to be lost in the embedding process. Thus, to determine whether LOT should be used on a particular dataset, it would be helpful to quantify how much information is lost, how this relates to the dimension of the embedding, and what consequences this has on subsequent statistical analysis.

To address these questions, we present a decomposition of *Fréchet variance* [22] in the 2-Wasserstein space, which allows one to compute the percentage of variance captured by Linear Optimal Transport embeddings. This is similar in spirit to the non-Euclidean PCA (e.g., tangent space PCA [28]) where one evaluates a PCA representation using the fraction of total variance captured by the linear approximation. We then proceed to show that similar decompositions of Fréchet variance can be found for the Gromov-Wasserstein (GW) and Fused Gromov-Wasserstein (FGW) distances.

This paper first recalls the Wasserstein distance, Wasserstein barycenters, and the concept of Fréchet variance, focusing on empirical probability measures. It then develops a principled way to decompose Fréchet variance in Wasserstein space and demonstrates this idea practically using several experiments on real data. We also recall the concept of (Fused) Gromov-Wasserstein distance and show how our results on Fréchet variance and our numerical framework extend to this setting. Our experiments study the relationships between the dimension of the LOT embedding, the percentage of variance explained by the embedding, and the accuracy of machine learning classifiers built on the embedded data. We apply these experiments to the MNIST handwritten digits dataset [33], IMDB-50000 sentiment analysis dataset [35], and Diffusion Tensor MRI (DTMRI) data from the “Human Connectome Project - Young Adult” (HCP-YA) dataset [48]. Our results show that LOT can be used to find low-dimensional Euclidean embeddings of unstructured data while maintaining relatively high classification accuracy, especially in applications to computer vision.

The specific contributions of this work are as follows.

- 1) Using a novel interpretation of a decomposition of squared 2-Wasserstein distances previously noted in [2], [6], we express the 2-Wasserstein Fréchet variance as a sum of deterministic and probabilistic components.
- 2) We introduce a similar decomposition of the 2-Gromov-Wasserstein Fréchet variance into deterministic and probabilistic components, including a proof that the deterministic component in the decomposition corresponds to a 2-Gromov-Wasserstein Fréchet variance. Using these results, we show that one can also define a

decomposition of Fused 2-Gromov-Wasserstein Fréchet variance.

- 3) We present a novel extension of the  $F$ -statistic that allows one to test for equality of  $n$  support barycentric projection for point clouds, with accompanying examples that show how to conduct permutation tests using the statistic on simulated data.
- 4) We present experiments that show how LOT and the variance decomposition can be used for dimensionality reduction of point clouds and graphs and for parameter selection in the Fused Gromov-Wasserstein setting. Our results illustrate that low dimensional LOT embeddings can capture much of the Fréchet variation and provide relatively good classification results in real-world datasets.

The rest of this paper is organized as follows; in Section II, we present relevant background on Optimal Transport and Linear Optimal Transport. In Section III, we present a decomposition of the 2-Wasserstein Fréchet variance for datasets of empirical probability measures supported on  $\mathbb{R}^d$  into deterministic and probabilistic components, and then extend this decomposition to the Gromov-Wasserstein and Fused Gromov-Wasserstein settings. In Section IV, we present the results of numerical experiments that explore the variance decompositions using the MNIST handwritten digits dataset, natural language data from the IMDB-50000 dataset, and DTMRI data from the HCP-YA dataset. In Section V, we conclude by discussing directions for future work.

## II. BACKGROUND MATERIAL

### A. Wasserstein Distances and Empirical Distributions

We start by defining the Wasserstein distance between probability measures supported on an arbitrary metric space. Our main reference for classical optimal transport is [52].

*Definition 1:* Let  $(\Omega, d)$  be a Polish metric space. For  $p \geq 1$ , let

$$\mathcal{P}_p(\Omega) := \{\mu : \int_{\Omega} d(x, x_0)^p d\mu(x) < \infty \forall x_0 \in \Omega\}$$

denote the set of Borel probability measures on  $\Omega$  with finite  $p$ -th moment. For  $\nu, \mu \in \mathcal{P}_p(\Omega)$ , define  $\Pi(\nu, \mu)$  to be the set of *couplings* of  $\nu$  and  $\mu$ , that is, the set of joint probability measures  $\gamma$  on  $\Omega \times \Omega$  with marginals  $\nu$  and  $\mu$ , respectively. Then

$$W_p(\nu, \mu) := \left( \inf_{\gamma \in \Pi(\nu, \mu)} \int_{\Omega \times \Omega} d(x, y)^p d\gamma(x, y) \right)^{\frac{1}{p}} \quad (1)$$

defines a distance on  $\mathcal{P}_p(\Omega)$  called the *p-Wasserstein distance*. We let  $W_2^2(\nu, \mu) := (W_2(\nu, \mu))^2$  denote the squared 2-Wasserstein distance. Elements of  $\Pi(\nu, \mu)$  are also referred to as *transport plans*, and a  $\gamma$  that achieves the infimum of (1) is called an *optimal transport plan*.

For arbitrary  $\nu, \mu \in \mathcal{P}_p(\Omega)$ , it is not possible to calculate  $W_p(\nu, \mu)$  explicitly. However, one can always compute Wasserstein distances between *empirical probability measures*, i.e., discrete probability measures with a finite number of

support points. For  $n \geq 1$ , define the *probability simplex* to be the set

$$\Sigma_n := \{u \in \mathbb{R}_+^n : \|u\|_1 = 1\}.$$

One can represent an empirical probability measure with *weights*  $a = (a_1, a_2, \dots, a_n) \in \Sigma_n$  and *support points*  $X = (x_1, x_2, \dots, x_n) \in \Omega^n$  as  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$ , where  $\delta_x$  denotes the Dirac probability measure at  $x \in \Omega$ . We let  $\text{supp}(\nu)$  denote the *support* of a measure  $\nu$ .

In order to calculate Wasserstein distances between empirical probability measures, one can use the solution to a standard form linear program known as the *transportation problem* [7]. For empirical probability measures  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$ ,  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$ , the set of couplings  $\Pi(\nu, \mu)$  can be identified with the set of matrices

$$U(a, b) := \{\gamma \in \mathbb{R}_+^{n \times m} : \sum_{i=1}^n \gamma_{ij} = b_j, \sum_{j=1}^m \gamma_{ij} = a_i \forall i, j\}.$$

Letting  $D_{ij} = d(x_i, y_j)^2$ , one can compute the squared 2-Wasserstein distance between  $\nu$  and  $\mu$  by using linear programming to find a matrix  $\gamma$  that solves

$$W_2^2(\nu, \mu) = \min_{\gamma \in U(a, b)} \text{trace}(\gamma^T D).$$

Throughout the paper, we mildly abuse terminology and conflate couplings (measures) with elements of  $U(a, b)$  (matrices), when the meaning is clear from context.

We now recall some useful terminology. Let  $(X, \nu)$  and  $(Y, \mu)$  be probability spaces, and let  $T : X \rightarrow Y$  be a measurable map. The *pushforward* of  $\nu$  by  $T$  is the measure  $T_{\#}\nu$  on  $Y$  defined by  $(T_{\#}\nu)(A) = \nu(T^{-1}(A))$ . The map  $T$  is called *measure-preserving* if  $T_{\#}\nu = \mu$ ; in the context of optimal transport, such a map is called a *transport map*. Observe that a transport map always leads to a coupling of  $\nu$  and  $\mu$ , via  $(\text{id}_X \times T)_{\#}\nu$ , where  $\text{id}_X \times T : X \rightarrow X \times Y$  is the map  $x \mapsto (x, T(x))$ . If a transport plan (coupling)  $\gamma$  is induced by a transport map in this manner, we say that  $\gamma$  is a *deterministic coupling*. This is to emphasize the point that a general transport plan is *probabilistic*, in the sense that mass at a single point in  $\text{supp}(\nu)$  can be assigned to multiple points in  $\text{supp}(\mu)$ , whereas this behavior does not occur when the plan is induced by a map.

### B. Wasserstein Barycenters and Fréchet Variance

Using the metric structure of  $\mathcal{P}_p(\Omega)$  one can define Fréchet means, also referred to as *Wasserstein barycenters* [1] in the OT literature.

*Definition 2:* A *p-Wasserstein barycenter* of measures  $\mu_\ell \in \mathcal{P}_p(\Omega)$ ,  $\ell = 1, \dots, N$ , is any measure  $\nu$  satisfying

$$\nu \in \text{argmin}_{\mu \in \mathcal{P}_p(\Omega)} \left( \sum_{\ell=1}^N W_p^p(\mu, \mu_\ell) \right). \quad (2)$$

While there exist methods to calculate Wasserstein barycenters for sets of empirical probability measures [4], they are often computationally intractable for large datasets. For the applications presented in Section IV, we use *free support* Wasserstein barycenter algorithms [19], [43] to compute approximations of Wasserstein barycenters for sets of empirical

probability measures. The free support barycenter algorithm ([19], Algorithm 2) seeks to minimize Eq. (2) with respect to the support points of the barycenter (for some fixed weight vector  $a$ ) by minimizing a local quadratic approximation of the gradient of the transport cost. Given an estimate of the barycenter  $\nu^t$  with weights  $a$  and locations  $X^t$ , and letting  $Y^\ell$  be a matrix containing the locations of a measure  $\mu_\ell$  with weights  $b^\ell$ ,  $\ell = 1, \dots, N$ , with  $\gamma_t \in U(a, b^\ell)$  an optimal coupling of  $\nu^t$  and  $\mu_\ell$ , one uses the update rule [19]

$$X^{t+1} = \frac{1}{N} \sum_{\ell=1}^N \text{diag}(a)^{-1} \gamma^t Y^\ell.$$

Given a set of measures  $\mu_\ell$ ,  $\ell = 1, \dots, N$ , and a 2-Wasserstein barycenter of those measures  $\nu$ , the sample *Fréchet variance* of the  $\mu_\ell$ s is given by

$$\widehat{Var}_W := \frac{1}{N} \sum_{\ell=1}^N W_2^2(\nu, \mu_\ell).$$

The sample Fréchet variance provides a way to quantify the dispersion of data in metric spaces, where standard definitions of variance/covariance cannot necessarily be applied. Figure 1 shows a schematic illustration of Wasserstein barycenter of empirical measures  $\mu_\ell$ ,  $\ell = 1, \dots, N$ .

### C. Linear Optimal Transport

Many standard statistical methods (e.g., Regression, Principal Components Analysis (PCA), Support Vector Machines (SVM), etc.) are naturally designed for Euclidean data. When working with non-Euclidean data, the typical approach is to *linearize* the data and then apply standard statistical methods. In the Riemannian setting, a standard pipeline used to linearize data is to (1) calculate a Fréchet mean of the data, and (2) apply the logarithmic map (calculated with respect to the Fréchet mean) to get tangent vector representations of the data, at which point statistical methods for Euclidean data can be applied (for examples, see [30], [46]).

Linear Optimal Transport (LOT) [53] is an analogous pipeline for linearizing data in a Wasserstein space with respect to a template measure (for which a natural choice is the Wasserstein barycenter). In order to linearize measure data, LOT uses *barycentric projection*.

*Definition 3:* (See [3, Def. 5.4.2]) Let  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$ ,  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$  be empirical probability measures on  $\mathbb{R}^d$ , and let  $\gamma$  be an optimal coupling of  $\nu$  and  $\mu$ . Then the *barycentric projection map induced by  $\gamma$* ,  $T : \text{supp}(\nu) \rightarrow \mathbb{R}^d$ , is defined to be

$$T(x_i) := \sum_j \frac{\gamma_{ij}}{a_i} y_j, \quad i = 1, \dots, n,$$

and the *barycentric projection of  $\mu$  w.r.t.  $\nu$*  is given by

$$T_\# \nu = \sum_{i=1}^n a_i \delta_{T(x_i)}. \quad (3)$$

One can use barycentric projection maps to define a vector field on  $\text{supp}(\nu)$  as follows. We first identify  $\text{supp}(\nu)$  with

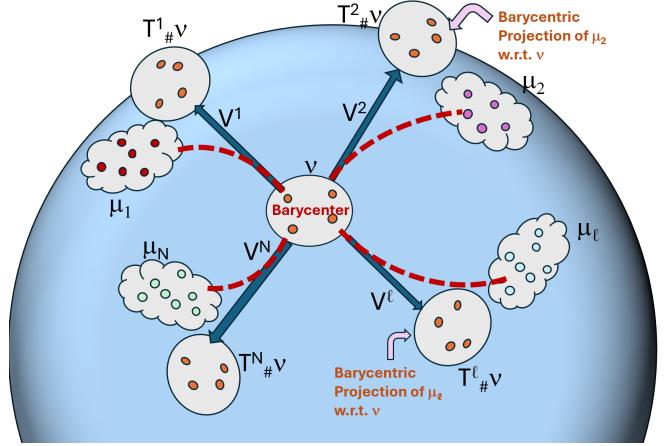


Fig. 1. Illustration of empirical measures  $\mu_1, \mu_2, \dots, \mu_l$  and their barycenter  $\nu$ . The measures  $T_\# \nu^\ell$  denote the barycentric projections of  $\nu$  towards  $\mu_\ell$ . The vectors  $V^\ell$  represent the Euclidean translations from  $\nu$  to  $T_\# \nu^\ell$

a matrix  $x \in \mathbb{R}^{n \times d}$ . Then the vector associated with the  $i$ th support point of  $\nu$  is given by

$$V_i = T(x_i) - x_i \in \mathbb{R}^d, \quad i = 1, 2, \dots, n. \quad (4)$$

Importantly, for a template measure  $\nu$  with  $n$  support points, the vector fields calculated for any  $\mu_\ell = \sum_{j=1}^{m_\ell} b_j \delta_{y_j}$  will be of fixed dimension; i.e., the size of  $V^\ell \in \mathbb{R}^{n \times d}$  will not depend on  $m_\ell$ . Thus, one can vectorize the  $V^\ell$ 's and apply standard statistical methods to these vector representations of the measures. We refer to the  $V^\ell$ 's as LOT embeddings of the  $\mu_\ell$ s.

The fact that Wasserstein spaces are not Riemannian manifolds leads to some issues with the analogy to the pipeline for linearizing data in the Riemannian setting. As noted in [12], barycentric projection with respect to an empirical probability measure is a lossy process in general. Thus, a natural question is - “How well do LOT embeddings describe the original data?”. In the next section, we show that, for measures supported on Euclidean spaces, one can quantify exactly the percentage of Fréchet variance explained by LOT embeddings for three popular OT distances.

### III. FRÉCHET VARIANCE DECOMPOSITIONS WITH LOT

In this section, we present conditions under which Linear Optimal Transport can be used to decompose Fréchet variance with respect to the Wasserstein, Gromov-Wasserstein, and Fused Gromov-Wasserstein distances.

#### A. Decomposition of Wasserstein Distance

We begin with a proposition first presented in [2].

*Proposition 1:* Let  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$  and  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$  be empirical probability measures supported on  $\mathbb{R}^d$ , let  $\gamma \in U(a, b)$  be an optimal coupling of  $\nu$  and  $\mu$ , and let  $T$  be the barycentric projection map induced by  $\gamma$  as defined in Eq. 3. Then

$$W_2^2(\nu, \mu) = W_2^2(\nu, T_\# \nu) + \sum_{ij} \gamma_{ij} \|T(x_i) - y_j\|^2. \quad (5)$$

Although this result was derived in [2], for the sake of completeness, and due to its correspondence with the proof of Theorem 1, we include a proof of Proposition 1 in Appendix ???. Because the terms on the RHS of Eq. (5) correspond to the transport cost of a deterministic coupling (of  $\nu$  and  $T_{\#}\nu$ ) and a probabilistic coupling (of  $T_{\#}\nu$  and  $\mu$ ), we refer to these as the *deterministic* and *probabilistic components* of the decomposition, respectively. More on that later. The proof of the proposition shows that one can find a similar decomposition into deterministic and probabilistic components for any coupling  $\gamma \in U(a, b)$ ; however, for non-optimal couplings, the first term in (5) will no longer be a squared Wasserstein distance.

We also provide here an interpretation of the decomposition which is novel, to the best of our knowledge. Given an arbitrary coupling  $\gamma \in U(a, b)$ , one obtains a collection of  $m$  vector fields on  $\text{supp}(\nu)$ : the vectors  $W_i^1, \dots, W_i^m$  based at the point  $x_i$  are given by

$$W_i^j = \gamma_{ij} \cdot (y_j - x_i).$$

We say that the coupling  $\gamma$  is *deterministic* if there is exactly one nonzero vector based at each  $x_i$ . On the opposite end of the spectrum, we say that  $\gamma$  is *purely probabilistic* if, for each  $i$ ,

$$\sum_{j=1}^m W_i^j = 0.$$

Intuitively, a deterministic coupling acts on the density  $\nu$  by moving its points around in the ambient space—in fact, it is clear from the definition that a coupling is deterministic if and only if it is induced by a transport map. On the other hand, a purely probabilistic coupling acts on  $\nu$  by completely redistributing its mass; it has no component which acts by moving any point in the ambient space. The Eq. (5) decomposes an optimal coupling  $\gamma$  into a deterministic coupling between  $\nu$  and  $T_{\#}\nu$  and a purely probabilistic coupling  $\gamma$  between  $T_{\#}\nu$  and  $\mu$ . To see the latter point, observe that

$$\begin{aligned} \sum_{j=1}^m W_i^j &= \sum_j \gamma_{ij} (y_j - T(x_i)) = \sum_j \gamma_{ij} y_j - T(x_i) \sum_j \gamma_{ij} \\ &= \sum_j \gamma_{ij} y_j - \left( \sum_j \frac{\gamma_{ij}}{a_i} y_j \right) a_i = 0. \end{aligned}$$

### B. Wasserstein Variance Decomposition

Note that, since Eq. (5) applies to a single squared 2-Wasserstein distance, it also works for averages of squared 2-Wasserstein distances. In particular, the sample Fréchet variance of a set of measures  $\mu_\ell$ ,  $\ell = 1, \dots, N$  calculated with respect to a free support barycenter  $\nu^n$  with  $n$  support points, denoted  $\widehat{Var}_W^n$ , can be decomposed into deterministic and probabilistic components.

**Definition 4:** Let  $\mu_\ell = \sum_{j=1}^{m_\ell} b_j^\ell \delta_{y_j^\ell}$ ,  $\ell = 1, \dots, N$  be empirical probability measures with  $\nu^n = \sum_{i=1}^n a_i \delta_{x_i}$  a free support barycenter of the  $\mu_\ell$ s with  $n$  support points,  $\gamma^\ell$  an optimal coupling of  $\nu^n$  and  $\mu_\ell$ , and  $T^\ell$  the barycentric projection map induced by  $\gamma^\ell$ . Then

$$\begin{aligned} \widehat{Var}_W^n &:= \frac{1}{N} \sum_{\ell=1}^N W_2^2(\nu^n, \mu_\ell) \\ &= \frac{1}{N} \sum_{\ell=1}^N W_2^2(\nu^n, T_\#^\ell \nu^n) + \frac{1}{N} \sum_{\ell=1}^N \sum_{ij} \gamma_{ij}^\ell \|T^\ell(x_i) - y_j^\ell\|^2. \end{aligned} \quad (6)$$

Note that the deterministic component (the first term on the RHS of Eq. (6)) corresponds to the sample Fréchet variance of the barycentric projections of the  $\mu_\ell$ s w.r.t.  $\nu^n$ , and thus can be used to quantify the percentage of the variance in the  $\mu_\ell$ s that is explained by the LOT embedding.

It is worth noting that the Law of Total Variance [11] can be seen as a special case of this decomposition, where  $d = 1$ ,  $n = 1$  and  $N \geq 2$ , or in other words, a ‘Law of Total Fréchet Variance’ defines a Pythagorean-like theorem in the 2-Wasserstein space for empirical probability measures supported on  $\mathbb{R}^d$ . This decomposition thus also relates to the  $F$ -statistic used in One-Way Analysis of Variance [9], [26]: observe that if the  $\mu_\ell \in \mathcal{P}(\mathbb{R})$  (treated as sample data rather than measures) all have uniform weights, and we take  $n = 1$ , one can calculate an  $F$ -statistic as

$$F^{n,d} = \frac{(\sum_{\ell=1}^N m_\ell) - N}{N-1} \frac{\sum_{\ell=1}^N W_2^2(\nu^n, T_\#^\ell \nu^n)}{\sum_{\ell=1}^N \sum_{ij} \gamma_{ij}^\ell \|T^\ell(x_i) - y_j^\ell\|^2}. \quad (7)$$

This trivially extends to a weighted ANOVA if the  $\mu_\ell$ s have non-uniform weights. However, this formula can also be used for  $n > 1$  (testing for equality of  $n$ -support barycentric projection) and/or  $d > 1$  (multivariate data). In an intuitive sense, the statistic can be seen as testing for equality of cluster means across data sets; the numerator captures the average ‘between group’ variation of cluster means  $T^\ell(x_i)$  about  $\frac{1}{N} \sum_{\ell=1}^N T^\ell(x_i)$ , while the denominator captures the average ‘within group’ variance of all clusters across all datasets. In this interpretation, standard one-way ANOVA would be a special case where the datasets are all assumed to have 1 cluster. To be a bit more rigorous, note that  $F^{n,d}$  will equal zero if all of the couplings  $\gamma^\ell$  are purely probabilistic and will increase without bound as the percent of Fréchet Variance explained by the deterministic component increases, being undefined if all of the couplings are deterministic.

Although the question of how the values  $n$  and  $d$  affect the distribution of the statistic is quite interesting, deriving the distribution of  $F^{n,d}$  for  $n > 1$  and/or  $d > 1$  is beyond the scope of this work. Instead, in section IV-A, we present examples of how one can conduct a permutation test for equality of  $n$ -support barycentric projection using  $F^{n,d}$ .

In the following sections, we extend the decomposition of Fréchet Variance presented in (6) to two other popular OT distances.

### C. Gromov-Wasserstein Variance Decomposition

The Gromov-Wasserstein distance [15], [36], [37] is an optimal transport-based distance between *measure networks*, as we now recall.

**Definition 5:** A *measure network* is a triple  $(X, \omega_X, \nu)$ , where  $X$  is a Polish space,  $\nu$  is a Borel probability measure on  $X$ , and  $\omega_X$  is a square integrable function on  $X \times X$  (w.r.t  $\nu \otimes \nu$ ). The *2-Gromov-Wasserstein distance* is given by

$$\begin{aligned} GW_2^2(\mathcal{X}, \mathcal{Y}) &= \inf_{\gamma \in \Pi(\nu, \mu)} \int_{(X \times Y)^2} (\omega_X(x, x') - \omega_Y(y, y'))^2 \\ &\quad \gamma(dx, dy) \gamma(dx', dy') \end{aligned}$$

Gromov-Wasserstein optimal couplings always exist [15, Theorem 2.2], so that the infimum in the definition is actually a minimum. We are particularly interested in finite measure networks, where we assume that  $X = \{x_1, \dots, x_n\}$ ,  $Y = \{y_1, \dots, y_m\}$  and that the measures are given by  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$ ,  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$ . We adopt the notation  $A_{ik} = \omega_X(x_i, x_k)$  and  $B_{jl} = \omega_Y(y_j, y_l)$ , so that  $GW_2^2(\mathcal{X}, \mathcal{Y})$  can be written as

$$GW_2^2(\mathcal{X}, \mathcal{Y}) = \min_{\gamma \in U(a, b)} \sum_{ijkl} \gamma_{ij} \gamma_{kl} |A_{ik} - B_{jl}|^2$$

One can define a natural notion of barycentric projection in the GW setting as follows (see [40, Definition 1]).

**Definition 6:** Let  $\mathcal{X} = (X, \omega_X, \nu)$  and  $\mathcal{Y} = (Y, \omega_Y, \mu)$  be two finite measure networks. Let  $\gamma$  be a Gromov-Wasserstein optimal coupling of  $\nu$  and  $\mu$ , and define

$$\omega_C(x_i, x_k) := \sum_{jl} \frac{\gamma_{ij}}{a_i} \frac{\gamma_{kl}}{a_k} B_{jl}.$$

Then the *GW barycentric projection* of  $\mathcal{Y}$  with respect to  $\mathcal{X}$  is given by

$$\mathcal{T} := (X, \omega_C, \nu).$$

Using this notion of barycentric projection, one can extend the decomposition in Eq. (5) to the Gromov-Wasserstein setting.

**Theorem 1:** Let  $\mathcal{X} = (X, \omega_X, \nu)$  and  $\mathcal{Y} = (Y, \omega_Y, \mu)$  be two finite measure networks,  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$ , and  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$ . Let  $\gamma$  be a Gromov-Wasserstein optimal coupling of  $\mathcal{X}$  and  $\mathcal{Y}$ , define  $B_{jl} = \omega_Y(y_j, y_l)$  and  $C_{ik} = \sum_{jl} \frac{\gamma_{ij}}{a_i} \frac{\gamma_{kl}}{a_k} B_{jl}$ , and let  $\mathcal{T}$  denote the GW barycentric projection of  $\mathcal{Y}$  with respect to  $\mathcal{X}$ . Then

$$\begin{aligned} GW_2^2(\mathcal{X}, \mathcal{Y}) &= GW_2^2(\mathcal{X}, \mathcal{T}) + \sum_{ijkl} \gamma_{ij} \gamma_{kl} |C_{ik} - B_{jl}|^2 \\ &= GW_2^2(\mathcal{X}, \mathcal{T}) + (\text{diam}_2(\mathcal{Y})^2 - \text{diam}_2(\mathcal{T})^2), \end{aligned} \quad (8)$$

where

$$\text{diam}_2(\mathcal{X}) := \left( \iint_{X \times X} \omega_X(x, x')^2 \nu(dx) \nu(dx') \right)^{1/2}. \quad (9)$$

Note that  $C_{ik}$  is the evaluation of  $\omega_C$  on the pairs  $(x_i, x_k)$ .

The *2-diameter* of a measure network, as defined in (9), gives a natural notion of its size, and appears in certain estimates of GW distances (see [15], [37]). The proof of the theorem is provided in Appendix ???. The proof shows that an optimal coupling of  $\mathcal{X}$  and  $\mathcal{T}$  is given by the identity coupling  $(\text{id}_X \times \text{id}_X) \# \nu$ . Additionally, it shows that *any* coupling  $\gamma$  separates into deterministic and probabilistic components, though, again, the first term will only necessarily be  $GW_2^2$  if  $\gamma$  is optimal.

#### D. Fused Gromov-Wasserstein Variance Decomposition

The Fused Gromov-Wasserstein distance [51] is an optimal transport distance between ‘structured data objects’ - measure networks whose nodes are elements of the same metric space. For the purposes of defining a variance decomposition, we are only interested in the special case of the Fused 2-Gromov-Wasserstein distance.

**Definition 7:** Let  $(\Omega, d)$  be a metric space,  $\mathcal{X} = (X, \omega_X, \nu)$  and  $\mathcal{Y} = (Y, \omega_Y, \mu)$  be two measure networks with  $X, Y \subset \Omega$ , and  $\alpha \in [0, 1]$ . Then the *squared  $\alpha$ -Fused 2-Gromov-Wasserstein distance* is given by

$$\begin{aligned} FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{Y}) &= \inf_{\gamma \in \Pi(\nu, \mu)} \int_{X \times Y \times X \times Y} [\alpha d(x, y)^2 \\ &\quad + (1 - \alpha)(\omega_X(x, x') - \omega_Y(y, y'))^2] \gamma(dx, dy) \gamma(dx', dy'). \end{aligned}$$

Using Definitions 3 and 6, it is straightforward to define a notion of barycentric projection in the FGW setting, for measure networks supported on  $\mathbb{R}^d$  (see [40, Definition 1]).

**Definition 8:** Let  $\mathcal{X} = (X, \omega_X, \nu)$  and  $\mathcal{Y} = (Y, \omega_Y, \mu)$  be two finite measure networks with  $X = \{x_i \in \mathbb{R}^d, i = 1, \dots, n\}$ ,  $Y = \{y_j \in \mathbb{R}^d, j = 1, \dots, m\}$ ,  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$ , and  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$ . Let  $\gamma$  be an  $\alpha$ -Fused 2-Gromov-Wasserstein optimal coupling of  $\mathcal{X}$  and  $\mathcal{Y}$ , define  $T(x_i) = \sum_j \frac{\gamma_{ij}}{a_i} y_j$ ,  $T(X) \in \mathbb{R}^{n \times d}$  such that  $T(X)_i = T(x_i)$ ,  $B_{jl} = \omega_Y(y_j, y_l)$  and  $\omega_C(x_i, x_k) = \sum_{jl} \frac{\gamma_{ij}}{a_i} \frac{\gamma_{kl}}{a_k} B_{jl}$ . Then the FGW barycentric projection of  $\mathcal{Y}$  with respect to  $\mathcal{X}$  is given by

$$\mathcal{T} := (T(X), \omega_C, T_{\#} \nu)$$

Given equations (5) and (8), it is easy to see that one can define a decomposition of the squared 2-FGW distance for structured data objects supported on  $\mathbb{R}^d$  as well. For notational simplicity, we define transport cost functions as follows.

**Definition 9:** Let  $\mathcal{X} = (X, \omega_X, \nu)$  and  $\mathcal{Y} = (Y, \omega_Y, \mu)$  be two measure networks with  $X, Y \subset \mathbb{R}^d$ ,  $\nu = \sum_{i=1}^n a_i \delta_{x_i}$  and  $\mu = \sum_{j=1}^m b_j \delta_{y_j}$ . Let  $\gamma$  be any coupling of  $\nu$  and  $\mu$ . Then the *squared 2-Wasserstein transport cost* of  $\gamma \in \Pi(\nu, \mu)$  is defined to be

$$C_W(\gamma) := \sum_{ij} \gamma_{ij} \|x_i - y_j\|^2$$

the *squared 2-Gromov-Wasserstein transport cost* of  $\gamma$  is defined to be

$$C_{GW}(\gamma) := \sum_{ijkl} \gamma_{ij} \gamma_{kl} (\omega_X(x_i, x_k) - \omega_Y(y_j, y_l))^2$$

and the *squared  $\alpha$ -Fused 2-Gromov-Wasserstein transport cost* of  $\gamma$  is defined to be

$$C_{FGW}^{\alpha}(\gamma) := \alpha C_W(\gamma) + (1 - \alpha) C_{GW}(\gamma)$$

Due to the fact that the separability of the decompositions given in Eq. (5) and Eq. (8) do not depend on the optimality of  $\gamma$  (see the comments following each result), we get the following corollary.

**Corollary 1:** Let  $\mathcal{X} = (X, \omega_X, \nu)$  and  $\mathcal{Y} = (Y, \omega_Y, \mu)$  be two measure networks, as in Definition 8. Let  $\alpha \in [0, 1]$  and let  $\gamma$  be an  $\alpha$ -Fused 2-Gromov-Wasserstein optimal coupling of  $\nu$  and  $\mu$ ,  $T(x_i) = \sum_j \frac{\gamma_{ij}}{a_i} y_j$ , and  $\pi = (T, \text{id}_Y) \# \gamma$ , where

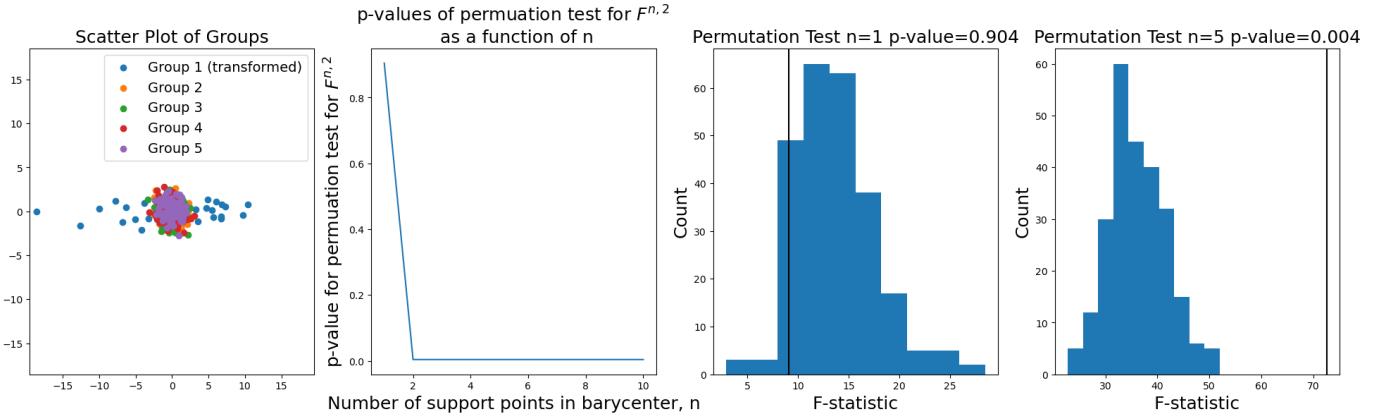


Fig. 2. Left: Sample data from 5 groups; Middle Left: p-values for permutations test for different values of  $n$ ; Middle Right: Permutation distribution and observed F statistic for  $n=1$ ; Right: Permutation distribution and observed F statistic for  $n=5$ .

$(T, \text{id}_Y) : X \times Y \rightarrow T(X) \times Y$  is the map  $(x, y) \mapsto (T(x), y)$ . Then

$$FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{Y}) = FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{T}) + C_{FGW}^\alpha(\pi)$$

*Proof.* We verify this by direct calculation, using the Fréchet variance decompositions Proposition 1 and Theorem 1:

$$\begin{aligned} FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{Y}) &= \alpha C_W(\gamma) + (1 - \alpha) C_{GW}(\gamma) \\ &= \alpha(C_W(T) + C_W(\pi)) + (1 - \alpha)(C_{GW}(T) + C_{GW}(\pi)) \\ &= C_{FGW}^\alpha(T) + C_{FGW}^\alpha(\pi) \\ &= FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{T}) + C_{FGW}^\alpha(\pi), \end{aligned}$$

where the last equality holds by [40, Lemma 1, part 1].  $\square$

Thus, given a data set of structured data objects  $\mu_\ell, \ell = 1, \dots, N$ , the sample  $\alpha$ -Fused 2-Gromov-Wasserstein Fréchet variance can be decomposed as

$$\begin{aligned} \widehat{Var}_{FGW_{2,\alpha}}^n &:= \frac{1}{N} \sum_{\ell=1}^N FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{Y}^\ell) \\ &= \frac{\alpha}{N} \sum_{\ell=1}^N FGW_{2,\alpha}^2(\mathcal{X}, \mathcal{T}^\ell) + \frac{1 - \alpha}{N} \sum_{\ell=1}^N C_{FGW}^\alpha(\pi^\ell) \end{aligned}$$

#### IV. APPLICATIONS TO DATA ANALYSIS

In this section, we demonstrate the theoretical tools developed in section III to represent and analyze datasets of measures and measure networks.

##### A. Illustrating $F^{n,d}$ Using Simulated Data

Here we present results for permutation tests of equality of  $n$ -support barycentric projection, for  $d = 2$ . Specifically, we generate sample data (with sample size  $m = 100$ ) from standard Gaussians (on  $\mathbb{R}^2$ ) for  $N = 5$  groups, and apply the transformation  $T(x, y) = (x^3, y)$  to one of the Gaussian samples. The left panel of Figure 2 presents a scatter plot of the data. We then calculate the test statistic  $F^{n,2}$  for this data, and estimate the permutation distribution of  $F^{n,2}$  with  $k=250$  permutations, for  $n = 1, \dots, 10$ . Importantly, this transformation does not change the expected value of the mean

of the transformed data (i.e.  $E[T(X, Y)] = E[(X, Y)] = \mathbf{0}$ ), and so the expected value of the 1-support barycenter of the sample should be unchanged. The second panel from the left presents the  $p$ -values of a permutation test of the  $F^{n,2}$  statistic on the data, for different values of  $n$ . As expected, the test fails to reject the hypothesis of 1-support barycentric projection, but rejects the hypothesis of equality of  $n$ -support barycentric projection for  $n > 1$ . The third and fourth panels from the left present the estimated permutation distributions of  $F^{1,2}$  and  $F^{5,2}$  calculated on the data.

#### B. Decompositions

In the following sections, we study the behavior of the decompositions presented in Section III, as well as the accuracy of machine learning classifiers applied to LOT embedded data, as a function of the number of support points in a Free Support barycenter. In section IV-C, we present results for digit classification on the MNIST handwritten digits dataset, in section IV-D we present results for sentiment analysis on the IMDB-50000 dataset, and in section IV-E, we present results on gender classification on DTMRI data from the HCP-YA dataset.

#### C. MNIST Data

We begin with experiments studying the variance decomposition using the MNIST handwritten digits dataset. The dataset is a set of grayscale images of handwritten digits, with classes 0, ..., 9. We represent each image as an empirical measure network in  $\mathbb{R}^2$ , using pixel locations as support points/nodes, Euclidean distances between support points as edge weights, and pixel intensities (normalized to sum to 1) as weights. Using these representations, we use the free support Wasserstein barycenter algorithms [19], [43] to calculate barycenters with different numbers of support points, and calculate LOT embeddings with respect to these barycenters, for  $\alpha = 0, 0.25, 0.5, 0.75, 1$ . These LOT embeddings have dimensions that depends on the number of support points/nodes in the free support barycenter, and this allows us to choose the dimension of the LOT embedded data.

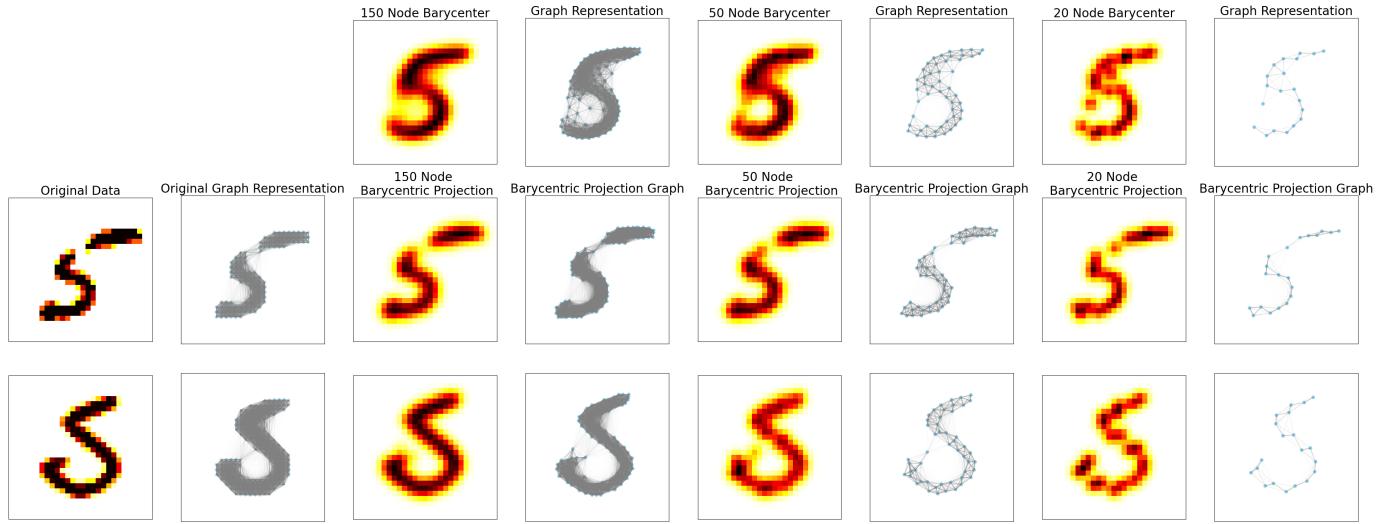


Fig. 3. Top Row: Gaussian Kernel Reconstructions and graph representations for free support Fused Gromov-Wasserstein barycenters ( $\alpha = 0.5$ ) of MNIST 5 digit, with  $n = 150, 50$ , and  $20$  nodes; Bottom Rows: Original Data, Gaussian Kernel Reconstructions of Barycentric projections, and graph representations for two 5 digits from the MNIST data set.

The top row of Fig. 3 presents Gaussian kernel reconstructions (reconstructed using the algorithm described in appendix ??) of free support barycenters and their nodes, for the digit “5”, for  $n = 150, 50, 20$  and  $\alpha = 0.5$ . The bottom two rows of Fig. 3 present the original images, and Gaussian kernel reconstructions of the barycentric projections (and the node locations) for two particular observations from the dataset, calculated with respect to the barycenter in the top row of the corresponding column. This plot motivates the use of free support barycenters with reduced support sizes in the LOT pipeline. For example, we see little degradation of the reconstructions going from a 150-point barycenter to a 50-point barycenter. While there is noticeably more degradation when we use a 20-support barycenter, the barycentric projections still appear to maintain their fundamental structure.

Fig. 4 presents a more quantitative exploration of the embedding, based on the decomposition in Corollary 1, for different values of  $n$  and  $\alpha$ . Specifically, we calculate a free support barycenter on a training set of 5000 randomly selected digits for  $n = 1, \dots, 50$ , and then calculate and plot the components of the decomposition, as functions of  $n$ , on a (separate) test set of 20000 randomly selected digits. The leftmost panel presents these results; the  $x$ -axis corresponds to  $n$ , the number of support points in the free support barycenter, and the  $y$ -axis corresponds to the different components of the decomposition. A consequence of Corollary 1 is that one can quantify not just the percentage of variance explained by LOT embeddings, but also the percentage of variance explained by the node and edge components of the embedding separately. The left panel of Figure 6 presents these components of the decomposition as a function of number of support points, for  $\alpha = 0.5$ . The middle panel of Figure 6 presents the percentage of variance explained by the deterministic component, for  $\alpha = 0.0, 0.25, 0.5, 0.75, 1$ . Importantly, for  $\alpha = 0.1$ , we use the Wasserstein and GW barycentric projections - though one could technically use an FGW barycentric projection (i.e., one

could still calculate the edge component of the barycentric projection for  $\alpha = 0$  or the node location component for  $\alpha = 1$ ). Interestingly, the percentage of variance explained is nearly identical for all values of  $\alpha$ . Considering the left panel of Figure 4, we observe that  $\widehat{Var}_{FGW_{2,\alpha}}^n$  decreases as the number of support points increases. This makes sense; more support points should give you more degrees of freedom with which to describe variation, thus resulting in shorter geodesics. The deterministic components, on the other hand, tends to increase with the number of support points. This also makes sense, and mirrors the previous observation; more support points in the barycenter means more of the variation can be described with transport maps.

The middle panel presents an elbow plot of the percentage of  $\widehat{Var}_{FGW,\alpha}^n$  described by the deterministic component; it is the pointwise ratio of the orange curve to the blue curve presented in the left panel. First, note that the percentage of variance explained generally increases with  $n$ , but it's not strictly non-decreasing (in particular, note the spike at  $n = 8, 9$  for the percentage of variance explained for higher values of  $\alpha$  (purple and red curves)). This is likely due (at least in part) to the fact that the barycenters are only approximations, though it is an open question whether the deterministic component should be strictly non-decreasing in  $n$ . We also note that barycentric projections describe more than 87.5% of  $\widehat{Var}_{FGW,0}^{50}$ , with other values of  $\alpha$  slightly higher. Somewhat surprisingly, the percentage of variance explained is almost identical for all of the distances.

In the right panel, we present the results of an experiment studying the average 5-fold cross validation classification accuracy of a multiclass LightGBM [32] gradient boosted tree trained on LOT embeddings (with respect to a free support barycenter built on the training set of 5000 digits) of 20000 digits from the MNIST dataset for different values of  $n$  and  $\alpha$ . In order to linearize our measure networks, we vectorize FGW barycentric projections using the mapping

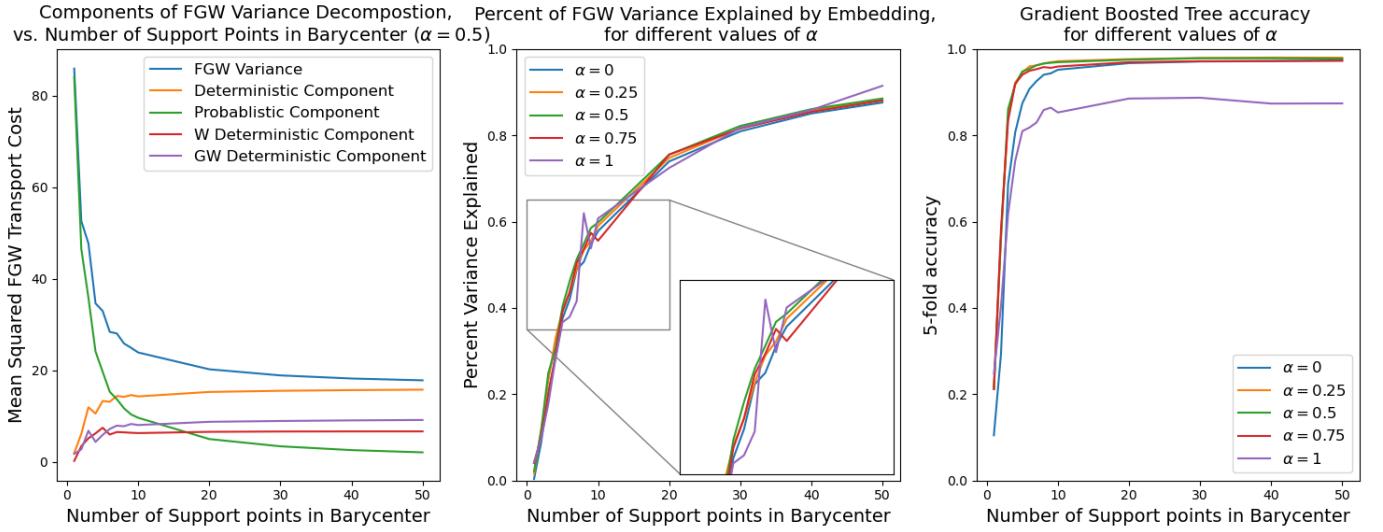


Fig. 4. Left: Components of decomposition calculated with respect to a free support Barycenter with different numbers of support points, for MNIST data; Middle: percentage of 2-Wasserstein variance explained by LOT embeddings with respect to a free support Barycenter with different numbers of support points, for MNIST data; Right: Multiclass classification accuracy of a gradient boosted tree built on LOT embeddings with respect to a free support Barycenter with different numbers of support points

$$(T(X), \|\cdot\|_2, C) \rightarrow \begin{bmatrix} \text{Vec}(T(X)) \\ \text{Vec}(\text{triu}(2C - \text{diag}(C))) \end{bmatrix}$$

where  $T(X)$  and  $C$  are defined as in Definition 8). Note, for  $\alpha = 0$  or 1, we only vectorize the nodes or edges, respectively.

We find that a gradient boosted tree trained on LOT embeddings with  $\alpha = 0.5$ , calculated with respect to a free support barycenter with 7 support points achieves an average of 96.23% test set accuracy, while the performance for 50 support points achieves only a slight improvement of 97.76%. This is a promising result, suggesting that LOT embeddings with respect to a barycenter with a low number of support points can explain a large proportion of information relevant to classification in some datasets.

Despite the fact that there appears to be little difference between the FGW and Wasserstein decompositions in terms of percentage of variance explained, we find that FGW LOT embeddings outperform Wasserstein LOT embeddings in terms of classification accuracy for lower numbers of support points. However, this comparison is a bit misleading, since the dimension of the embeddings depends on  $\alpha$ . For a template measure network with  $n$  nodes on  $\mathbb{R}^d$ , the dimension of the linearizations is given by  $I(\alpha < 1)nd + I(\alpha > 0)(n + n(n - 1)/2)$ .

Thus, to compare performance in terms of the dimension of the embedding rather than the number of support points, one would have to compare, for example, the  $\alpha = 0.5, n = 7$  embedding to the  $\alpha = 0, n = 21$  embedding. In this case, the 5-fold accuracy for the  $\alpha = 0, n = 21$  embeddings is 97.09%, which is slightly better than the  $\alpha = 0.5, n = 7$  embeddings, which achieved just 96.61%. This suggests that LOT embeddings with  $\alpha = 0$  can be more efficient in terms of dimension than embeddings with  $\alpha > 0$ .

#### D. IMDB-50000 Data

Here, we explore an application of LOT to the task of

sentiment analysis in the IMDB-50000 dataset. The dataset is made up of 50000 movie reviews (paragraphs of text, ranging from 10-1200 words) from the International Movie Database website, each associated with either a positive or negative sentiment. After training a Word2Vec embedding on the full dataset (with embedding dimension 100), we represent each review as a point cloud in the embedding space. Treating these points as the nodes of a measure network allows us to apply LOT to the data. Interestingly, in this context, LOT can be seen as a sort of linear document embedding of the reviews. Specifically, we represent each review as an empirical measure network  $(X, B, \mu)$ , where  $X$  contains the locations of the vector embeddings of the words in the review,  $B_{jl} = |l - j|$  captures the number of words separating  $X_j$  and  $X_l$  in the review (note, the indices reflect the order of the words in the review, i.e.,  $y_j^\ell$  corresponds to the vector embedding of the  $j$ th word in review  $\ell$ ), and  $\mu$  corresponds to a uniform measure supported on  $X$ . Similar to the MNIST experiment, we calculate barycenters on a training set of 5000 randomly selected reviews, and then calculate the variance decomposition and conduct 5-fold cross validation with a gradient boosted tree classifier on a test set of 20000 randomly selected reviews.

Figure 5 contains the results of the experiment. Specifically, we calculate the variance decomposition and build a gradient boosted tree on LOT embeddings with respect to barycenters with  $n = 1, 5, 10, 25, 50, 100$  support points, for  $\alpha = 0, 0.000005, 0.000025, 1$ . These values of  $\alpha$  were chosen to balance the percent variance explained by the node versus edge components. The middle panel in Figure 5 shows the percentage of variance explained by the LOT embedding for different values of  $\alpha$ . We see that the two lower values of  $\alpha$  (which weigh the Wasserstein component more heavily) show much slower growth in percent variance explained than the two higher values. This comes down to 1) the difference in

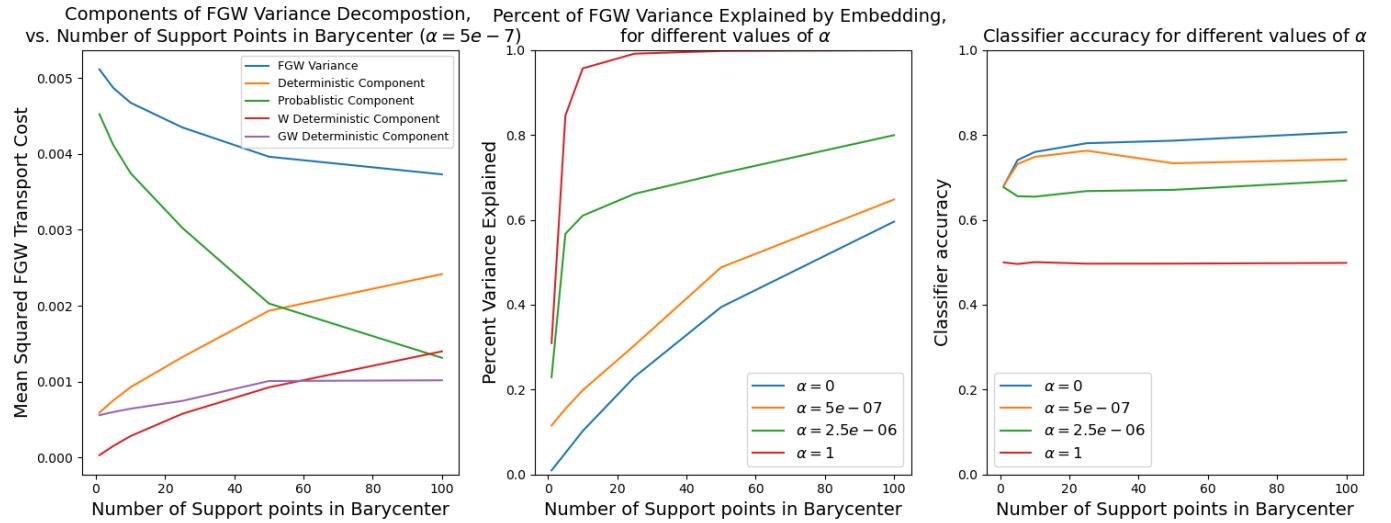


Fig. 5. Left: Components of decomposition calculated with respect to a free support Barycenter with different numbers of support points, for IMDB data; Middle: percentage of 2-Wasserstein variance explained by LOT embeddings with respect to a free support Barycenter with different numbers of support points; Right: Classification accuracy of a gradient boosted tree built on LOT embeddings with respect to a free support Barycenter with different numbers of support points

scale between the node and edge components of our measure networks and 2) the fact that the edge matrices all have fairly similar structure in terms of the GW distance, leading to little GW variance in the data for our choice of edge weights.

The rationale behind our choice of edge weights was that the graph components of the barycentric projections might capture information about consistent sequences of words across reviews, that could not be captured in the ‘bag-of-words’ information contained in just the node locations. Our classification results, presented in right panel of Figure 5, suggest that this is not the case, with models built only on node locations ( $\alpha = 0$ ) outperforming models for all other values of  $\alpha$ . Somewhat surprisingly, the 5-fold cross validation accuracy of a model built on LOT embeddings with respect to just 1 support point (corresponding to simply averaging all the points in the point cloud for  $\alpha = 0$ ) achieves greater than 68% accuracy. However, we also find that there is an increase in classification accuracy for higher numbers of support points, for  $\alpha < 1$ . The  $\alpha = 0, n = 100$  LOT embedding performs the best, achieving just over 80% accuracy. Note that the dimension of this embedding is rather high at  $nd = 10000$ . These results suggest that our edge matrices do not contribute much to the classification accuracy; even when we select an  $\alpha$  that balances the W and GW components, we see no benefit in using structured data object representations over simple point clouds - though other edge weights might produce better results. In any case, our results do illustrate that LOT can be used to effectively embed and classify natural language data.

#### E. Diffusion Tensor MRI - Human Connectome Project

DTMRI is an approach to MRI that measures the diffusion of water molecules in tissue, such as white matter fiber tracts in the brain. The most common data representations of DTMRI images of white matter fiber tracts are based on ‘tractography’ methods, which seek to reconstruct white matter fiber tracts

as sets of curves in  $\mathbb{R}^3$ . Tractography datasets can be quite large, and only implicitly contain diffusion information in the tangent directions of the individual curves.

Tract density images (TDI), first proposed in [10], offer another data representation for modeling white matter fiber tracts. Instead of representing the tract as a set of 3-d curves, one counts the number of curves passing through each voxel in the image, and represents the tract by associating with each voxel the number of curves that pass through it - effectively representing the tract as an empirical measure on  $\mathbb{R}^3$ . As the number of curves estimated is often much larger than the number of voxels they pass through, these TDI representations can produce datasets with smaller memory requirements than those produced by tractography methods. Optimal Transport methods for both representations have previously been explored in [25].

An additional benefit of TDI representations is that they offer a natural way to explicitly include diffusion information in their data representations. Different approaches have been developed to add diffusion information to TDIs, such as ‘directionally encoded’ TDIs [41], or ‘Tract Orientation Density Images’ (TODI) [21]. The approach taken here is to associate with each voxel in a TDI a diffusion tensor, i.e., a covariance matrix constructed from the main fiber directions and fractional anisotropy values (as determined by DSI studio [23]) associated with that voxel. We refer to these representations as Covariance TDIs (CTDI). While these CTDIs will contain less information about diffusivity than TODIs, a benefit is that we are able to represent tracts as empirical probability measures on the space  $\mathbb{R}^3 \times \text{Sym}_3^+$  (where  $\text{Sym}_3^+$  corresponds to the set of symmetric positive definite matrices), which allows us to run statistical analyses via LOT. While one could run a similar experiment comparing models built on LOT embeddings with different values of  $\alpha$ , because the DTMRI target measures often have 1000’s of support points, this ends up being very

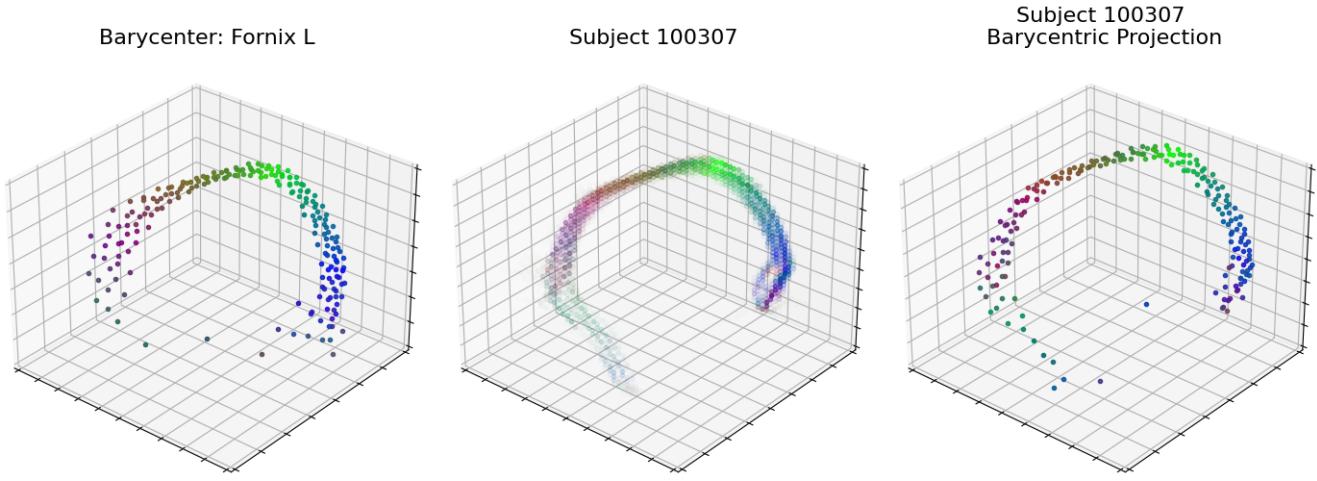


Fig. 6. Points represent locations in  $\mathbb{R}^3$ ,  $(r, g, b)$  values correspond to (normalized) diagonal elements of covariance matrices. Left: free support barycenter of 'Fornix L' in HCP data (200 support points); Middle: Subject 100307's 'Fornix L' in the HCP data (6980 support points); Right: Barycentric Projection of Subject 100307's 'Fornix L' with respect to calculated barycenter (200 support points).

computationally expensive. Thus, we choose to study only the Wasserstein LOT embeddings.

Specifically, we apply our framework to the task of classifying gender in the Human Connectome Project - Young Adult (HCP-YA) dataset. The HCP-YA data set is composed of 1200 young adults (aged 22-35) from 300 families. This means the dataset contains many siblings and twins, including identical twins. Clearly, if subjects  $i$  and  $j$  are identical twins, then  $P(\text{gender}_i = M | \text{gender}_j = M) \approx 1$ . However, by removing 1 of each pair of identical twins from the dataset, we believe it is reasonable to assume that  $P(\text{gender}_i = M | \text{gender}_j = M, \text{Family}_j = \text{Family}_i) \approx .5 \forall i, j$  satisfying  $\text{Family}_j = \text{Family}_i$  (and  $i \neq j$ ), i.e. that the gender of one sibling is independent of the gender of their other siblings. While one might expect dependence of CTDIs between siblings, given the assumption that the gender of siblings is independent, we expect any differences learned by our model to be due to differences in CTDIs alone. Thus, while our results don't necessarily represent the performance of the model on a true i.i.d. sample from the population, the classification accuracy should still depend on our model's ability to measure differences in the CTDIs (and hopefully, therefore, in the underlying white matter microstructure).

1) *Data Representation*: Fiber orientation maps for 1065 subjects from the Human Connectome Project (produced using the pipeline presented in [55]) were obtained from the DSI Studio website<sup>1</sup>. After preprocessing the data (including removing one of each pair of identical twins), we were left with a dataset with 908 subjects. TDIs for each region of interest (ROI) of each subject were produced using DSI studio, and converted to empirical probability measures as follows; given the TDI of a particular ROI for subject  $\ell$ , we represent it as an empirical measure on  $\mathbb{R}^3$

$$\tilde{\mu}_\ell = \sum_{j=1}^{m_\ell} w_j \delta_{x_j},$$

<sup>1</sup>[https://brain.labsolver.org/hcp\\_ya.html](https://brain.labsolver.org/hcp_ya.html), FIB 1.25-mm MNI space

where  $x \in \mathbb{R}^{m_\ell \times 3}$  represents the voxel locations that make up the ROI of subject  $\ell$  (as determined by DSI Studio), and  $w_j$  denotes the number of curves present in the voxel at  $x_j$ . We then produce a covariance matrix  $\Sigma_j$  for each  $x_j$ , using the fiber orientations  $u_{x_j}^r \in \mathbb{R}^{3 \times 1}$  and fractional anisotropy  $\text{FA}_{x_j}^r \in \mathbb{R}^+$  (obtained from the fiber orientation map) associated with that voxel

$$\Sigma_j = \text{Proj}_{\text{Sym}_3^+} \left( \sum_{r=1}^5 \text{FA}_{x_j}^r u_{x_j}^r (u_{x_j}^r)^T \right).$$

Often, these covariances are not full rank, so we project them into the set of symmetric positive definite matrices  $\text{Sym}_3^+$  by replacing any non-positive eigenvalues with a small positive value. Thus, letting  $\theta_j = (x_j, \Sigma_j)$ , and normalizing the TDI weights as  $b_j = w_j / (\sum_{j=1}^{m_\ell} w_j)$ , we represent the ROI of subject  $\ell$  as an empirical probability measure on  $\mathbb{R}^3 \times \text{Sym}_3^+$

$$\mu_\ell = \sum_{j=1}^{m_\ell} b_j \theta_j.$$

For the purposes of calculating the decomposition, we conduct our statistical analysis by equipping  $\mathbb{R}^3$  and  $\text{Sym}_3^+$  with their respective Euclidean metrics. While, at this point, one could choose other Riemannian metrics on  $\mathbb{R}^3$  and  $\text{Sym}_3^+$  [45], we choose to use the Euclidean metrics because it enables us to isometrically embed the locations  $\theta$  into  $\mathbb{R}^9$  (and thus easily calculate the percentage of variance explained by the LOT embedding) using the mapping

$$\left( \begin{bmatrix} v_1 \\ v_2 \\ v_3 \end{bmatrix}, \begin{bmatrix} v_4 & v_5 & v_6 \\ v_5 & v_7 & v_8 \\ v_6 & v_8 & v_9 \end{bmatrix} \right) \rightarrow [v_1, v_2, v_3, v_4, 2v_5, 2v_6, v_7, 2v_8, v_9]^T. \quad (10)$$

2) *Calculating the LOT Embedding*: Given a set of empirical measures  $\mu_\ell, \ell = 1, \dots, N$ , each representing a particular

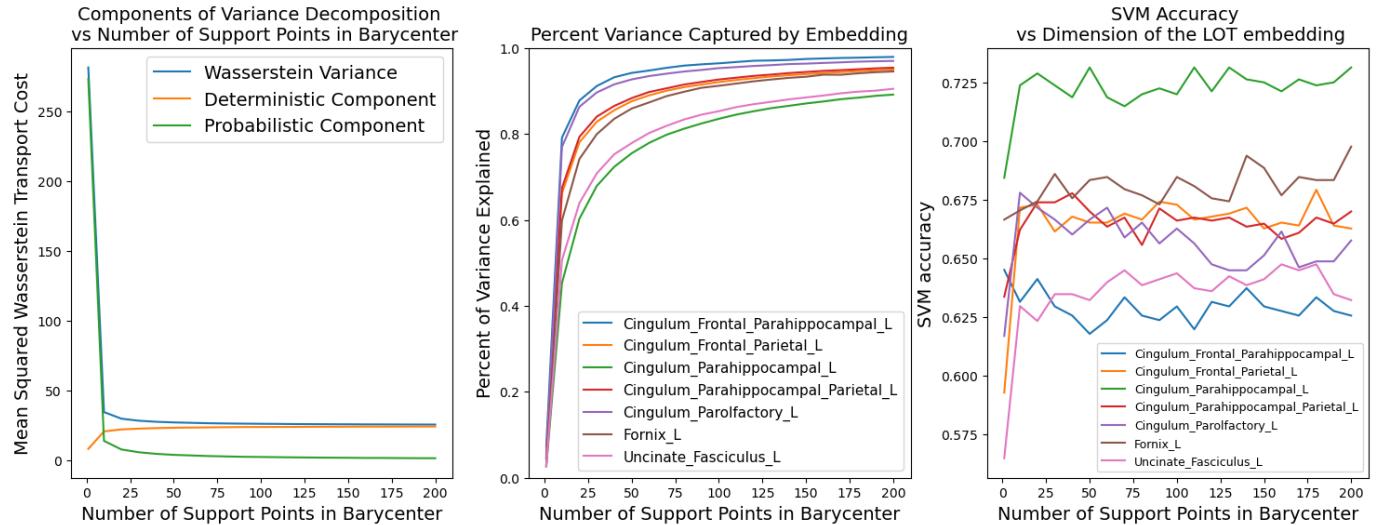


Fig. 7. Left: Fréchet variance and components of decomposition as a function of number of support points in free support barycenter (for  $n=1,10,20,\dots,200$ ) for Fornix L in HCP-YA; Middle: percentage of variance explained by LOT embedding with respect to a free support barycenter with  $n$  support points, for ROI (in left hemisphere) in HCP-YA; Right: SVM classification accuracy on LOT embedded data with respect to number of support points in free support barycenter.

ROI in a subject, we calculate a free support barycenter  $\nu$  using the algorithm described in [19] (initialized with uniform weights, points in  $\mathbb{R}^3$  generated from a Standard Normal distribution, and 3-d covariances generated from a log-uniform distribution [38] over  $\text{Sym}_3^+$ ). The free support barycenter calculated for the tract ‘Fornix L’ (with  $n = 200$  support points) is presented in the left panel of Fig. 6.

Next, given  $\mu_\ell, \ell = 1, \dots, N$ , and the free support barycenter  $\nu = \sum_{i=1}^n a_i \theta_i$ , we calculate matrix representations of the vector fields of each subject  $V^\ell \in \mathbb{R}^{n \times d}$  as in Eq. (3). The  $V^\ell$ ’s can be vectorized as  $Z_\ell = \text{Vec}(V^\ell) \in \mathbb{R}^{nd}$ , at which point, standard statistical methods (i.e. PCA, Regression, SVM) can be applied to the LOT embedded dataset  $Z \in \mathbb{R}^{nd \times N}$ . The original data and the barycentric projection representation for a particular subject (calculated with respect to the barycenter presented in the left panel of Fig. 6) are presented in the middle and right panels of Fig. 6, respectively.

3) *DTMRI Classification Results:* Fig. 7 presents results of numerical experiments applied to the DTMRI data, similar to those presented in Fig. 4. The left panel corresponds to the components of the variance decomposition specifically for the ‘Fornix L.’ The middle panel presents elbow plots for all ROI (in the left hemisphere). Notably, we see that for all regions, the large majority of variation is explained by LOT embeddings calculated with respect to free support barycenters with just 25 – 75 support points. Given that most regions have on the order of 1000’s of support points, this corresponds to a significant reduction of dimensionality. The right panel corresponds to the performance of SVM classifiers built on *only the covariance components* (we include results for classifiers built on other combinations of location and covariance components in appendix ??) of the LOT embeddings with respect to an  $n$  support barycenter, for  $n = 1, 10, 20, 50, 100, 200$ . For most regions, we see a large jump in performance going from 1 support point to 10 support points, with most models hitting

maximum performance (relative to choice of  $n$ ) by about  $n = 50$  support points.

Table I presents the results of Leave-One-Out (LOO) cross validation for an SVM classifier applied to only the covariance components of the vector representations of the LOT embedded measures, for different values of  $n$  (number of support points in barycenter). We also compare our method to a ‘baseline’ model that is often used in DTMRI analysis, which simply involves comparing the average FA value for each region of each subject [31], [47], [49]. Notably, our results suggest that our model outperforms the baseline for all values of  $n$ , and that SVMs built on LOT embedded data achieve near maximum classification accuracy (relative to other choices of parameters) for free support barycenters with just 10 support points, with little to no improvement for larger values of  $n$ , both for individual regions and for a soft voting classifier built on class probabilities of the SVMs built on the individual regions, presented in the bottom row of Table I.

## V. CONCLUSION

This paper presents a decomposition of the Fused 2-Gromov-Wasserstein Fréchet variance of measures supported on Euclidean spaces using Linear Optimal Transport. The decomposition gives data scientists a tool that aids in parameter selection when working with FGW distances and helps them better interpret the quality of LOT embeddings. It also suggests connections to other methods for calculating Free Support barycenters [34]. The decomposition also suggests a generalization of the F-statistic used in ANOVA, which allows one to test a data set of empirical probability measures for equality of  $n$  support barycenter. Furthermore, our methods have provided promising results for the statistical modeling of DTMRI images, using Linear Optimal Transport to naturally fold dimensionality reduction and non-linear image registration into standard machine learning pipelines. This applica-

Region	N	Classification of Gender in HCP-YA					
		Average	FA	n=1	n=10	n=50	n=200
Cingulum Frontal Parahippocampal L	513	63.93	64.52	62.96	62.18	62.57	62.77
Cingulum Frontal Parahippocampal R	476	63.22	65.76	65.55	65.76	64.29	65.13
Cingulum Frontal Parietal L	786	53.06	59.29	66.92	66.28	66.92	65.9
Cingulum Frontal Parietal R	785	59.49	63.69	68.54	68.54	68.15	68.03
Cingulum Parahippocampal L	786	67.17	68.45	71.25	72.52	73.03	72.39
Cingulum Parahippocampal R	786	63.36	64.76	70.1	70.74	69.59	69.85
Cingulum Parahippocampal Parietal L	770	63.38	63.38	65.97	67.4	66.1	66.49
Cingulum Parahippocampal Parietal R	779	57.38	64.57	67.27	67.39	66.5	65.85
Cingulum Parolfactory L	786	56.85	61.7	67.81	67.81	65.27	65.39
Cingulum Parolfactory R	786	57.88	58.27	60.81	62.34	62.72	62.34
Fornix L	771	65.36	66.67	66.93	68.09	69.0	68.35
Fornix R	768	63.16	64.45	67.32	67.71	68.49	68.75
Uncinate Fasciculus L	786	55.99	56.49	63.23	63.61	64.38	64.38
Uncinate Fasciculus R	786	53.80	53.44	63.99	64.89	65.01	64.76
All* Regions Combined	739	64.95	69.14	75.10	75.23	74.96	75.10

\*Cingulum Frontal Parahippocampal L & Cingulum Frontal Parahippocampal R excluded due to lower sample sizes

TABLE I

LEAVE ONE OUT CROSS VALIDATION ACCURACY FOR SVM BUILT ON COVARIANCE COMPONENTS OF LOT EMBEDDINGS OF DTMRI DATA WITH RESPECT TO  $n$ -SUPPORT BARYCENTERS. N DENOTES SAMPLE SIZE - THERE IS MISSINGNESS DUE TO ISSUES IN DATA PREPROCESSING.

tion suggests many directions for future research, such as increasing the complexity of the data representations by using other Riemannian metrics on SPD matrices [45] or ‘orientation distribution functions’ rather than covariances [21], scaling OT with convolutions and/or Sinkhorn [18], [25] to build models on larger regions or the whole brain, and/or implementing partial/unbalanced LOT methods [6].

### A. Acknowledgments

This research was supported in part by NSF grants DMS-2107808, DMS-2324962 to Needham and DMS-1953087, DMS-241378 to Srivastava, and NIH grant R01 MH120299 to Srivastava. DTMRI data were provided [in part] by the Human Connectome Project, WU-Minn Consortium (Principal Investigators: David Van Essen and Kamil Ugurbil; 1U54MH091657) funded by the 16 NIH Institutes and Centers that support the NIH Blueprint for Neuroscience Research; and by the McDonnell Center for Systems Neuroscience at Washington University.

### REFERENCES

- [1] Martial Aguech and Guillaume Carlier. Barycenters in the Wasserstein space. *SIAM Journal on Mathematical Analysis*, 43(2):904–924, 2011.
- [2] Aurélien Alfonsi and Benjamin Jourdain. Squared quadratic Wasserstein distance: optimal couplings and Lions differentiability. *ESAIM: Probability and Statistics*, 24:703–717, 2020.
- [3] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows in metric spaces and in the space of probability measures*. Springer Science & Business Media, 2008.
- [4] Ethan Anderes, Steffen Borgwardt, and Jacob Miller. Discrete Wasserstein barycenters: Optimal transport for discrete data, 2015.
- [5] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein generative adversarial networks. In *International conference on machine learning*, pages 214–223. PMLR, 2017.
- [6] Yikun Bai, Ivan Vladimir Medri, Rocio Diaz Martin, Rana Shahroz, and Soheil Kolouri. Linear optimal partial transport embedding. In *International Conference on Machine Learning*, pages 1492–1520. PMLR, 2023.
- [7] Dimitris Bertsimas and John N Tsitsiklis. *Introduction to linear optimization*, volume 6. Athena Scientific Belmont, MA, 1997.
- [8] Nicolas Bonneel and Julie Digne. A survey of optimal transport for computer graphics and computer vision. *Computer Graphics Forum*, 42(2):439–460, 2023.
- [9] George EP Box, William H Hunter, Stuart Hunter, et al. *Statistics for experimenters*, volume 664. John Wiley and sons New York, 1978.
- [10] Fernando Calamante, Se-Hong Oh, Jacques-Donald Tournier, Sung-Yeon Park, Young-Don Son, Jun-Young Chung, Je-Geun Chi, Graeme D Jackson, Chan-Woong Park, Young-Bo Kim, et al. Super-resolution track-density imaging of thalamic substructures: Comparison with high-resolution anatomical magnetic resonance imaging at 7.0 t. Technical report, Wiley Online Library, 2013.
- [11] George Casella and Roger Berger. *Statistical Inference*. Duxbury Resource Center, June 2001.
- [12] Elsa Cazelles, Vivien Seguy, Jérémie Bigot, Marco Cuturi, and Nicolas Papadakis. Log-pca versus geodesic pca of histograms in the Wasserstein space. *arXiv preprint arXiv:1708.08143*, 2017.
- [13] Sinho Chewi, Jonathan Niles-Weed, and Philippe Rigollet. Statistical optimal transport. *arXiv preprint arXiv:2407.18163*, 2024.
- [14] Lenaic Chizat, Gabriel Peyré, Bernhard Schmitzer, and François-Xavier Vialard. An interpolating distance between optimal transport and Fisher-Rao metrics. *Foundations of Computational Mathematics*, 18:1–44, 2018.
- [15] Samir Chowdhury and Facundo Mémoli. The Gromov–Wasserstein distance between networks and stable network invariants. *Information and Inference: A Journal of the IMA*, 8(4):757–787, 2019.
- [16] Samir Chowdhury and Tom Needham. Gromov-Wasserstein averaging in a Riemannian framework. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 842–843, 2020.
- [17] Samir Chowdhury and Tom Needham. Generalized spectral clustering via Gromov-Wasserstein learning. In *International Conference on Artificial Intelligence and Statistics*, pages 712–720. PMLR, 2021.
- [18] Marco Cuturi. Sinkhorn distances: Lightspeed computation of optimal transport. *Advances in neural information processing systems*, 26, 2013.
- [19] Marco Cuturi and Arnaud Doucet. Fast computation of Wasserstein barycenters. In *International conference on machine learning*, pages 685–693. PMLR, 2014.
- [20] Julie Delon and Agnès Desolneux. A Wasserstein-type distance in the space of gaussian mixture models. *SIAM Journal on Imaging Sciences*, 13(2):936–970, 2020.
- [21] Thijs Dhollander, Louise Emsell, Wim Van Hecke, Frederik Maes, Stefan Sunaert, and Paul Suetens. Track orientation density imaging (tod)

and track orientation distribution (tod) based tractography. *NeuroImage*, 94:312–336, 2014.

[22] Paromita Dubey and Hans-Georg Müller. Fréchet analysis of variance for random objects. *Biometrika*, 106(4):803–821, 2019.

[23] Frank Yeh et. al. Dsi studio.

[24] Jean Feydy, Benjamin Charlier, François-Xavier Vialard, and Gabriel Peyré. Optimal transport for diffeomorphic registration. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2017: 20th International Conference, Quebec City, QC, Canada, September 11–13, 2017, Proceedings, Part I 20*, pages 291–299. Springer, 2017.

[25] Jean Feydy, Pierre Roussillon, Alain Trouvé, and Pietro Gori. Fast and scalable optimal transport for brain tractograms. In *Medical Image Computing and Computer Assisted Intervention- MICCAI 2019: 22nd International Conference, Shenzhen, China, October 13–17, 2019, Proceedings, Part III 22*, pages 636–644. Springer, 2019.

[26] Ronald Aylmer Fisher. Statistical methods for research workers. In *Breakthroughs in statistics: Methodology and distribution*, pages 66–70. Springer, 1970.

[27] Rémi Flamary, Nicholas Courty, Davis Tuia, and Alain Rakotomamonjy. Optimal transport for domain adaptation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 1(1-40):2, 2016.

[28] P Thomas Fletcher, Conglin Lu, Stephen M Pizer, and Sarang Joshi. Principal geodesic analysis for the study of nonlinear statistics of shape. *IEEE transactions on medical imaging*, 23(8):995–1005, 2004.

[29] Aude Genevay, Gabriel Peyré, and Marco Cuturi. Learning generative models with sinkhorn divergences. In *International Conference on Artificial Intelligence and Statistics*, pages 1608–1617. PMLR, 2018.

[30] Nicolas Guigui, Nina Miolane, Xavier Pennec, et al. Introduction to riemannian geometry and geometric statistics: from basic theory to implementation with geomstats. *Foundations and Trends® in Machine Learning*, 16(3):329–493, 2023.

[31] Richard A Kanaan, Matthew Allin, Marco Picchioni, Gareth J Barker, Eileen Daly, Sukhwinder S Shergill, James Woolley, and Philip K McGuire. Gender differences in white matter microstructure. *PloS one*, 7(6):e38272, 2012.

[32] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30, 2017.

[33] Yann LeCun and Corinna Cortes. MNIST handwritten digit database. <http://yann.lecun.com/exdb/mnist/>, 2010.

[34] Giulia Luise, Saverio Salzo, Massimiliano Pontil, and Carlo Ciliberto. Sinkhorn barycenters with free support via frank-wolfe algorithm. *Advances in neural information processing systems*, 32, 2019.

[35] Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150, 2011.

[36] Facundo Mémoli. On the use of Gromov-Hausdorff Distances for Shape Comparison. In M. Botsch, R. Pajarola, B. Chen, and M. Zwicker, editors, *Eurographics Symposium on Point-Based Graphics*. The Eurographics Association, 2007.

[37] Facundo Mémoli. Gromov–Wasserstein distances and the metric approach to object matching. *Foundations of computational mathematics*, 11:417–487, 2011.

[38] Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas, Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, et al. Geomstats: a python package for Riemannian geometry in machine learning. *The Journal of Machine Learning Research*, 21(1):9203–9211, 2020.

[39] Eduardo Fernandes Montesuma, Fred Ngole Mboula, and Antoine Souloumiac. Recent advances in optimal transport for machine learning. *arXiv preprint arXiv:2306.16156*, 2023.

[40] Dai Hai Nguyen and Koji Tsuda. On a linear fused Gromov-Wasserstein distance for graph structured data. *Pattern Recognition*, 138:109351, 2023.

[41] Sinisa Pajevic and Carlo Pierpaoli. Color schemes to represent the orientation of anisotropic tissues from diffusion tensor data: application to white matter fiber tract mapping in the human brain. *Magnetic Resonance in Medicine: An Official Journal of the International Society for Magnetic Resonance in Medicine*, 42(3):526–540, 1999.

[42] Gabriel Peyré and Marco Cuturi. Computational optimal transport. *Foundations and Trends in Machine Learning*, 11(5–6):355–607, 2019.

[43] Gabriel Peyré, Marco Cuturi, and Justin Solomon. Gromov-Wasserstein averaging of kernel and distance matrices. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2664–2672, New York, New York, USA, 20–22 Jun 2016. PMLR.

[44] Michael E Sander, Pierre Albin, Mathieu Blondel, and Gabriel Peyré. Sinkformers: Transformers with doubly stochastic attention. In *International Conference on Artificial Intelligence and Statistics*, pages 3515–3530. PMLR, 2022.

[45] Clément Sarrazin and Bernhard Schmitzer. Linearized optimal transport on manifolds. *SIAM Journal on Mathematical Analysis*, 56(4):4970–5016, 2024.

[46] A. Srivastava and E. Klassen. *Functional and Shape Data Analysis*. Springer Series in Statistics, New York, 2016.

[47] Hidemasa Takao, Naoto Hayashi, and Kuni Ohtomo. Sex dimorphism in the white matter: fractional anisotropy and brain size. *Journal of Magnetic Resonance Imaging*, 39(4):917–923, 2014.

[48] David C Van Essen, Kamil Ugurbil, Edward Auerbach, Deanna Barch, Timothy EJ Behrens, Richard Bucholz, Acer Chang, Liyong Chen, Maurizio Corbetta, Sandra W Curtiss, et al. The human connectome project: a data acquisition perspective. *Neuroimage*, 62(4):2222–2231, 2012.

[49] J Van Hemmen, IMJ Saris, PT Cohen-Kettenis, DJ Veltman, PJW Pouwels, and Julie Bakker. Sex differences in white matter microstructure in the human brain predominantly reflect differences in sex hormone exposure. *Cerebral Cortex*, 27(5):2994–3001, 2017.

[50] Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Fused Gromov-Wasserstein distance for structured objects: theoretical foundations and mathematical properties, 2018.

[51] Titouan Vayer, Laetitia Chapel, Rémi Flamary, Romain Tavenard, and Nicolas Courty. Optimal transport for structured data with application on graphs, 2019.

[52] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2009.

[53] Wei Wang, Dejan Slepčev, Saurav Basu, John A Ozolek, and Gustavo K Rohde. A linear optimal transportation framework for quantifying and visualizing variations in sets of images. *International journal of computer vision*, 101:254–269, 2013.

[54] Michael Wilson, Tom Needham, Chiwoo Park, Supartee Kundu, and Anuj Srivastava. A Wasserstein-type distance for gaussian mixtures on vector bundles with applications to shape analysis. *SIAM Journal on Imaging Sciences*, 17(3):1433–1466, 2024.

[55] Fang-Cheng Yeh, Jean M Vettel, Aarti Singh, Barnabas Poczos, Scott T Grafton, Kirk I Erickson, Wen-Yih I Tseng, and Timothy D Verstynen. Quantifying differences and similarities in whole-brain white matter architecture using local connectome fingerprints. *PLoS computational biology*, 12(11):e1005203, 2016.



**Michael Wilson** is a Ph.D. candidate at the Florida State University Department of Statistics. His research interests include applications of Optimal Transport and Statistical Shape Analysis to imaging data.



**Tom Needham** is an Assistant Professor of Mathematics at Florida State University. He received his Ph. D. in Mathematics from University of Georgia. He works on applications of geometry and topology to problems in data science, computer vision and signal processing.



**Anuj Srivastava** is a Professor in the Department of Statistics and a Distinguished Research Professor at the Florida State University. His research interests include statistical analysis on nonlinear manifolds, statistical computer vision, and functional and shape analysis. He has graduated 30+ Ph.D. students with placements in academia, industry, and government labs. He has co-authored more than 270 papers in peer-reviewed journals, top-tier conferences, and several books, including the 2016 Springer textbook on "Functional and Shape Data Analysis." He is a fellow of IAPR, IEEE, ASA, IMS, and AAAS.