

Wasserstein Covariance for Multiple Random Densities

BY ALEXANDER PETERSEN

*Department of Statistics and Applied Probability, University of California, Santa Barbara,
California 93106, U.S.A.
petersen@pstat.ucsb.edu*

5

AND HANS-GEORG MÜLLER

*Department of Statistics, University of California, Davis, California 95616, U.S.A.
hgmueLLer@ucdavis.edu*

SUMMARY

A common feature of methods for analyzing samples of probability density functions is that they respect the geometry inherent to the space of densities. Once a metric is specified for this space, the Fréchet mean is typically used to quantify and visualize the average density from the sample. For one-dimensional densities, the Wasserstein metric is popular due to its theoretical appeal and interpretive value as an optimal transport metric, leading to the Wasserstein–Fréchet mean or barycenter as the mean density. We extend the existing methodology for samples of densities in two key directions. First, motivated by applications in neuroimaging, we consider dependent density data, where a p -vector of univariate random densities is observed for each sampling unit. Second, introduce Wasserstein covariance and propose intuitively appealing estimators for both fixed and diverging p , where the latter corresponds to continuously-indexed densities. We also give theory demonstrating consistency and asymptotic normality, while accounting for errors introduced in the unavoidable preparatory density estimation step. The utility of the Wasserstein covariance matrix is demonstrated through applications to functional connectivity in the brain using functional magnetic resonance imaging data and to the secular evolution of mortality for various countries.

Some key words: Barycenter; Fréchet Mean; Fréchet Variance; Functional Connectivity; Mortality; Random Density.

1. INTRODUCTION

The analysis of samples of density functions or distributions is an important and challenging problem for modern statistical practice (Delicado, 2011). Examples include distributions of age at death or mortality for different countries, warping functions, or distributions of voxel-to-voxel correlations of functional magnetic resonance imaging signals for a sample of subjects. While functional principal component analysis using cross-sectional averaging can be directly applied for samples of density functions (Kneip & Utikal, 2001), more recently techniques have been developed that incorporate the geometric constraints inherent to the space of density functions. A popular metric for data where each data atom corresponds to a randomly sampled distribution or density is the Wasserstein metric, both for its theoretical appeal and its convincing empirical performance in various applications (Bolstad et al., 2003; Zhang & Müller, 2011; Bigot et al., 2016, 2017; Panaretos & Zemel, 2016).

We consider samples of density data where multiple random densities per subject are observed, that is, repeated realizations of a stochastic process defined on \mathcal{D}^p , where \mathcal{D} is a space of one-dimensional density functions. For $p = 1$, a variety of methods have been proposed, focusing on concepts of mean, modes of variation and dimensionality reduction (Delicado, 2011; Petersen & Müller, 2016; Hron et al., 2016). However, generalizations of these methods for $p > 1$, where one has a vector of densities, have not yet been developed, even though such data arise in various applications. An essential first step is the extension of the mean and variance concepts for samples of densities to a vector of means and covariance matrix. We demonstrate here the usefulness of these concepts for applications. The type of data we consider is different from data where one observes, for each subject, a random sample of identically distributed p -vectors, for which the joint distribution is of interest. Instead, we focus on the joint modeling of p univariate densities.

Using the Wasserstein geometry of optimal transport on the space of densities, we extend the concepts of Fréchet mean and variance to a measure of Wasserstein covariance between component densities for a p -variate density-valued process. As one does not observe the densities themselves but rather samples of univariate data that they generate, a preliminary density estimation step is necessary and is taken into account in our analysis. Our theoretical arguments show that the population Wasserstein covariance can be estimated consistently, with a limiting Gaussian distribution under sufficient conditions for the estimation error associated with the preliminary estimation step to be negligible. Motivated by one of the applications, we also consider continuously varying densities, where one observes a discretized version of a continuously evolving density-valued process, similar to densely observed repeated functional data (Park & Staicu, 2015; Chen et al., 2017). For this situation we target a Wasserstein covariance surface and develop theory for its consistent estimation while still accounting for errors in density estimation.

The utility of the proposed methodology is demonstrated through the analysis of functional magnetic resonance imaging and mortality data. In the application to brain imaging, we investigate differences in intra-regional functional connectivity between Alzheimer's patients, cognitively normal subjects, and a third group of individuals diagnosed with mild cognitive impairment, a sign of increased risk for developing dementia. Here, intra-regional functional connectivity is quantified by the distribution of voxel-to-voxel correlations within a neighborhood, called a functional hub, corresponding to a small region in the brain. Multiple densities per subject are obtained for a number of hubs simultaneously, leading to a random vector of densities obtained for each subject in the sample. In the mortality application, we consider distributions of age at death over a range of calendar years, and compare the resulting Wasserstein covariance surfaces between a group of Eastern European countries and a second group of other nations.

2. WASSERSTEIN COVARIANCE

2.1. The Wasserstein Metric and Geometry

Let \mathcal{D} be a class of one-dimensional densities such that $\int_{\mathbb{R}} u^2 f(u) du < \infty$ for all $f \in \mathcal{D}$. For $f, g \in \mathcal{D}$, suppose $Y \sim f$ and consider the collection of nondecreasing maps $T^* : \mathbb{R} \rightarrow \mathbb{R}$, known as transports, such that $T^*(Y) \sim g$. The optimal transport problem that leads to the Wasserstein metric is to find the transport that minimizes

$$\int_{\mathbb{R}} \{T^*(u) - u\}^2 f(u) du, \quad (1)$$

and the solution is known to be $T = G^{-1} \circ F$, where F and G are the distribution functions of f and g , respectively (Ambrosio et al., 2008). The resulting squared Wasserstein distance is

$$d_W^2(f, g) = \int_{\mathbb{R}} \{T(u) - u\}^2 f(u) du. \quad (2)$$

This metric arises from a local inner product (Ambrosio et al., 2008). If $g_j \in \mathcal{D}$, $j = 1, 2$, have distribution functions G_j , the optimal transports $T_j = G_j^{-1} \circ F$ for given F reside in the tangent space of f , and for each f an inner product between T_1 and T_2 can be defined by

$$\langle T_1, T_2 \rangle_f = \int_{\mathbb{R}} \{T_1(u) - u\} \{T_2(u) - u\} f(u) du, \quad (3)$$

so that $d_W(f, g_j)^2 = \langle T_j, T_j \rangle_f = \langle T_j^{-1}, T_j^{-1} \rangle_{g_j}$.

2.2. Wasserstein Mean, Variance and Covariance

For a random density \mathfrak{F} in \mathcal{D} , its Fréchet mean and Fréchet variance (Fréchet, 1948) are natural tools for relating the distributional properties of \mathfrak{F} to the Wasserstein geometry of \mathcal{D} in terms of first and second order behaviour. When the space \mathcal{D} is endowed with the metric d_W , the Wasserstein–Fréchet, or simply Wasserstein, mean and variance of \mathfrak{F} are

$$f_{\oplus} = \operatorname{argmin}_{f \in \mathcal{D}} E \{d_W(\mathfrak{F}, f)^2\}, \quad \operatorname{var}_{\oplus}(\mathfrak{F}) = E \{d_W(\mathfrak{F}, f_{\oplus})^2\}. \quad (4)$$

For a single density process \mathfrak{F} , the theoretical and practical properties of the Wasserstein mean have been thoroughly investigated (Bolstad et al., 2003; Zhang & Müller, 2011; Panaretos & Zemel, 2016; Bigot et al., 2017), and recently the Wasserstein variance was adopted to quantify variability explained when performing dimension reduction for densities (Petersen & Müller, 2016). To quantify the dependence between two random densities, we propose here the extension of these concepts to a Wasserstein covariance measure.

For a bivariate density process $(\mathfrak{F}_1, \mathfrak{F}_2)$, where the \mathfrak{F}_j are random elements of \mathcal{D} , $j = 1, 2$, denote by $F_{\oplus, j}$ the distribution function of the Wasserstein mean of \mathfrak{F}_j (4) and the corresponding density by $f_{\oplus, j}$. The random optimal transport from $f_{\oplus, j}$ to \mathfrak{F}_j is $T_j = F_j^{-1} \circ F_{\oplus, j}$, where F_j is the distribution function of \mathfrak{F}_j , so that the Wasserstein variances are

$$\operatorname{var}_{\oplus}(\mathfrak{F}_j) = E \{d_W(\mathfrak{F}_j, f_{\oplus, j})^2\} = E \left[\int_{\mathbb{R}} \{T_j(u) - u\}^2 f_{\oplus, j}(u) du \right] = E \left(\langle T_j, T_j \rangle_{f_{\oplus, j}} \right).$$

This suggests defining a Wasserstein covariance measure as an expected inner product between T_1 and T_2 . Since in general $f_{\oplus, 1} \neq f_{\oplus, 2}$, these transports reside in different tangent spaces. Adopting a common technique for manifold-valued data (Yuan et al., 2012), we push T_1 to a new transport map \tilde{T}_1 in the tangent space of $f_{\oplus, 2}$ by a parallel transport map. Intuitively, as T_1 is transported to \tilde{T}_1 along the geodesic connecting $f_{\oplus, 1}$ to $f_{\oplus, 2}$, its angle with the geodesic is preserved. In the Wasserstein geometry, setting $T_{\oplus, 12} = F_{\oplus, 1}^{-1} \circ F_{\oplus, 2}$, one obtains $\tilde{T}_1 = T_1 \circ T_{\oplus, 12} - T_{\oplus, 12} + \operatorname{id}$, where id is the identity map. Due to symmetry of this operation, if $\tilde{T}_2 = T_2 \circ T_{\oplus, 12}^{-1} - T_{\oplus, 12}^{-1} + \operatorname{id}$, then $\langle \tilde{T}_1, \tilde{T}_2 \rangle_{f_{\oplus, 2}} = \langle T_1, T_2 \rangle_{f_{\oplus, 1}}$. Thus, the Wasserstein geometry motivates

$$\operatorname{cov}_{\oplus}(\mathfrak{F}_1, \mathfrak{F}_2) = E \left(\langle \tilde{T}_1, \tilde{T}_2 \rangle_{f_{\oplus, 2}} \right) = E \left(\langle T_1, T_2 \rangle_{f_{\oplus, 1}} \right) = \operatorname{cov}_{\oplus}(\mathfrak{F}_2, \mathfrak{F}_1) \quad (5)$$

as the Wasserstein covariance between the two random densities.

2.3. Expression in Terms of Quantile Functions

The well known fact that for one-dimensional densities, the Wasserstein geometry is closely related to quantile functions (Villani, 2003; Panaretos & Zemel, 2016; Petersen & Müller, 2016) leads to a second characterization of Wasserstein covariance, which is useful for practice. The change of variable $t = F(u)$ applied to (3) gives the alternative expression

$$\langle T_1, T_2 \rangle_f = \int_0^1 \{F^{-1}(t) - G_1^{-1}(t)\} \{F^{-1}(t) - G_2^{-1}(t)\} dt, \quad (6)$$

so that $d_W^2(f, g) = \int_0^1 \{F^{-1}(t) - G^{-1}(t)\}^2 dt$. We can then express the quantities in (4) and (5) in terms of the random quantile functions F_j^{-1} . The Wasserstein means $f_{\oplus, j}$ are characterized by their quantile functions $F_{\oplus, j}^{-1}(t) = E\{F_j^{-1}(t)\}$, $0 \leq t \leq 1$, and for $j, k = 1, 2$, the Wasserstein variances and covariances are

$$\text{cov}_{\oplus}(\mathfrak{F}_j, \mathfrak{F}_k) = E \left[\int_0^1 \{F_j^{-1}(t) - F_{\oplus, j}^{-1}(t)\} \{F_k^{-1}(t) - F_{\oplus, k}^{-1}(t)\} dt \right]. \quad (7)$$

Expressions (6), (7) reveal a connection to functional data analysis. Viewing (F_1^{-1}, F_2^{-1}) as bivariate functional data, key objects are the mean functions $F_{\oplus, j}^{-1}(t)$ and covariance surfaces

$$\mathcal{C}_{jk}(s, t) = \text{cov} \left\{ F_j^{-1}(s), F_k^{-1}(t) \right\}, \quad 1 \leq j, k \leq 2, \quad 0 \leq s, t \leq 1,$$

that characterize the first- and second-order behavior of the processes (Li & Hsing, 2010). Writing \mathcal{M}_{jk} for the integral operator with kernel \mathcal{C}_{jk} , Fubini's theorem implies that

$$\text{cov}_{\oplus}(\mathfrak{F}_j, \mathfrak{F}_k) = \int_0^1 \mathcal{C}_{jk}(t, t) dt = \text{Tr}(\mathcal{M}_{jk}),$$

where $\text{Tr}(\cdot)$ is the operator trace. Accordingly, the Wasserstein variance of each component distribution can be interpreted as a summary of the variability in the quantile process, and Wasserstein covariance as a summary of covariability.

Using quantile functions has two major advantages when multiple densities are observed per subject. First, the derived notions of mean, variance, and covariance have geometric interpretations in the manifold induced by the Wasserstein metric in the space of distributions, as in (4) and (5). Second, quantile functions always have the same support $[0, 1]$ regardless of the distributional supports of the \mathfrak{F}_j , so that the Wasserstein covariance remains well-defined even when the latter differ. In contrast, attempts to define similar covariance summaries based on cross-covariance operators of density or compositional representations are bound to fail when distributional supports differ. For example, if \mathcal{G}_{jk} is the ordinary cross-covariance operator between densities \mathfrak{F}_j and \mathfrak{F}_k , the operator trace is well-defined only when the supports coincide. Even when they are the same, taking $\text{Tr}(\mathcal{G}_{jk})$ as a summary covariance measure has no intuitive geometric meaning.

2.4. Wasserstein Covariance Matrices and Kernels

Considering a p -variate density process $\mathfrak{F} = (\mathfrak{F}_1, \dots, \mathfrak{F}_p)$ with component Wasserstein means and variances $f_{\oplus, j}$ and $\text{var}_{\oplus}(\mathfrak{F}_j)$, $j = 1, \dots, p$, we initially assume that p is fixed, as in the brain imaging example in Section 4.1. The Wasserstein covariance matrix for \mathfrak{F} is the $p \times p$ matrix with elements

$$(\Sigma_{\oplus})_{jk} = \text{cov}_{\oplus}(\mathfrak{F}_j, \mathfrak{F}_k), \quad (8)$$

which is easily seen to be a valid covariance matrix.

Motivated by the mortality example in Section 4.2, suppose the components of \mathfrak{F} are indexed by a continuous variable $y_j \in [0, 1]$, which might represent time. To model the setting of repeatedly observed densities that are measured densely in time, we allow $p \rightarrow \infty$. Then \mathfrak{F} can be thought of as a discretized version of an unobservable process $\mathcal{F}(y)$, $0 \leq y \leq 1$, with $\mathfrak{F}_j = \mathcal{F}(y_j)$. We target the Wasserstein mean surface $f_{\oplus}(\cdot; y)$ and covariance kernel

$$\Sigma_{\oplus}(y, z) = \text{cov}_{\oplus} \{ \mathcal{F}(y), \mathcal{F}(z) \}, \quad 0 \leq y, z \leq 1. \quad (9)$$

3. ESTIMATION OF WASSERSTEIN COVARIANCE OBJECTS

3.1. Density Estimation

While we defined the Wasserstein covariance for fully observed densities, in reality these densities are rarely if ever directly observed. Rather, the data actually available are collections of scalar random variables W_{ijr} ($i = 1, \dots, n$; $j = 1, \dots, p$; $r = 1, \dots, N_{ij}$), where n is the number of subjects i , p is the number of densities or distributions j per subject, and N_{ij} is the number of independent observations r distributed according to the j -th density that are available for the i -th subject and may vary across i, j . The observed data can be viewed as resulting from two independent random mechanisms, where the first random mechanism generates the independent vectors of densities $f_i = \{f_{i1}, \dots, f_{ip}\}$, $i = 1, \dots, n$, while the second generates the observations that are sampled from these distributions, $W_{ijr} \sim f_{ij}$. The W_{ijr} are all independent and for each fixed (i, j) are also identically distributed.

Obtaining density estimates \hat{f}_{ij} from the observed data W_{ijr} and using these as proxies for the f_{ij} , for the asymptotic analysis we need to address the challenge that these estimates are noisy and deviate from the true density targets. Since the targets Σ_{\oplus} can be expressed as integrated moments of the multivariate quantile process, an obvious route would be to estimate the empirical quantile functions and proceed by averaging. However, this has some practical drawbacks and a preferred approach is to first construct a sample of smooth density estimates \hat{f}_{ij} , then obtaining smooth distribution functions by integration, quantile functions as inverse functions and the target quantities from the estimated quantile functions.

The theoretical analysis of the Wasserstein covariance estimates in Section 3.2 below requires the following assumption, where a preliminary density estimator is generically denoted by \hat{f} .

Assumption 1. There is a compact interval I such that, for any $f \in \mathcal{D}$, its support D_f is a compact interval contained in I . If W_1, \dots, W_N is a random sample from f , the density estimate \hat{f} based on this sample is a probability density function on D_f such that, for some decreasing sequence $b_N = o(1)$ as $N \rightarrow \infty$,

$$\sup_{f \in \mathcal{D}} E \left\{ d_W(f, \hat{f}) \right\} = O(b_N).$$

A density estimator that satisfies Assumption 1 is described in Petersen & Müller (2016). If the support D_f is known and the condition

$$\sup_{f \in \mathcal{D}} \sup_{u \in D_f} \max \{ f(u), 1/f(u), |f'(u)| \} < \infty$$

is satisfied, then one may take $b_N = N^{-1/3}$; see Proposition 1 in the Supplementary Material. With no assumed uniform lower bound on the random densities, Panaretos & Zemel (2016) proposed a density estimator for which $b_N = N^{-1/4}$.

Lastly, because np densities need to be simultaneously estimated, each from data of varying sample sizes N_{ij} , these need to be tied to the number of independent subjects n , as follows.

Assumption 2. There exists a sequence $N = N(n)$ such that $\min_{i,j} N_{ij} \geq N$ and $N \rightarrow \infty$ as $n \rightarrow \infty$.

Here N is a uniform lower bound on the sample sizes for the np densities to be estimated that must diverge with the number of subjects n to ensure uniform consistency of the densities.

3.2. Wasserstein Covariance Estimation

For fixed p , given densities $f_i = (f_{i1}, \dots, f_{ip})$, $i = 1, \dots, n$, that are independently and identically distributed according to $\mathfrak{F} = (\mathfrak{F}_1, \dots, \mathfrak{F}_p)$, our main goal is to estimate the Wasserstein covariance matrix Σ_{\oplus} , with elements defined in (8). We compute density estimates \hat{f}_{ij} , which are then mapped to their quantile function estimates $\hat{X}_{ij} = \hat{F}_{ij}^{-1}$. Write $\hat{X}_{ij}^c = \hat{X}_{ij} - n^{-1} \sum_{i=1}^n \hat{X}_{ij}$ and $\hat{\mathcal{C}}_{jk}(s, t) = n^{-1} \sum_{i=1}^n \hat{X}_{ij}^c(s) \hat{X}_{ik}^c(t)$. To target the Wasserstein covariance Σ_{\oplus} , the results of Section 2.3 suggest the estimator

$$\left(\hat{\Sigma}_{\oplus} \right)_{jk} = \int_0^1 \hat{\mathcal{C}}_{jk}(t, t) dt. \quad (10)$$

Theorem 1 demonstrates the overall rate of convergence of the Wasserstein covariance estimator, establishing asymptotic normality when N diverges sufficiently fast. While our focus is on the Wasserstein covariance, the same rate of convergence is also obtained for the full covariance surface $\hat{\mathcal{C}}_{jk}(s, t)$ as an estimator of $\mathcal{C}_{jk}(s, t)$; see Theorem 3 in the Supplementary Material, where also auxiliary results and proofs can be found.

THEOREM 1. *Suppose Assumptions 1 and 2 hold, and that $\mathfrak{F}_j \in \mathcal{D}$ almost surely, $j = 1, \dots, p$. Then $\|\Sigma_{\oplus} - \hat{\Sigma}_{\oplus}\|_F = O_p(n^{-1/2} + b_N)$, where $\|\cdot\|_F$ denotes the Frobenius norm. Additionally, if $n^{1/2}b_N$ converges to zero, then there exists a zero-mean $p \times p$ Gaussian matrix \mathfrak{C} such that $n^{1/2}(\hat{\Sigma}_{\oplus} - \Sigma_{\oplus})$ converges weakly to \mathfrak{C} .*

The covariance structure of \mathfrak{C} is a four-dimensional array defined in the Supplementary Material. Under regularity conditions, the density estimator in Petersen & Müller (2016) satisfies $b_N = N^{-1/3}$ so that asymptotic Gaussianity is obtained for $N = n^q$, $q > 3/2$. Faster rates of convergence $b_N = N^{-\rho}$, with $1/3 < \rho < 1/2$, can be obtained under additional smoothness conditions and for suitable density estimators, and then weaker conditions on q will suffice.

In the continuously indexed case, the vectors of densities $f_i = (f_{i1}, \dots, f_{ip})$ correspond to discretized versions of independent and identically distributed realizations $\tilde{f}_i(\cdot; y)$ of a latent dynamic density surface $\mathcal{F}(y)$, $0 \leq y \leq 1$, where f_{ij} are independently distributed for different i as $\mathcal{F}(y_{ij})$. For simplicity, we require

Assumption 3. The number of observation times $p = p(n)$ satisfies $np^{-1} = O(1)$, and these are equidistant and common for all subjects, i.e., $y_{ij} = (j-1)/(p-1)$, $1 \leq j \leq p$.

We then estimate $\Sigma_{\oplus}(y_j, y_k)$ by the sample Wasserstein covariance kernel estimator

$$\hat{\Sigma}_{\oplus}(y_j, y_k) = \int_0^1 \hat{\mathcal{C}}_{jk}(t, t) dt, \quad (11)$$

followed by linearly interpolating these discretized estimates to obtain $\hat{\Sigma}_{\oplus}(y, z)$ for any $0 \leq y, z \leq 1$. For this interpolation to be negligible, we require two additional assumptions.

Assumption 4. There exists a constant $L_1 > 0$ such that

$$|\Sigma_{\oplus}(y_1, z_1) - \Sigma_{\oplus}(y_2, z_2)| \leq L_1(|y_1 - y_2| + |z_1 - z_2|), \quad 0 \leq y_1, y_2, z_1, z_2 \leq 1.$$

Assumption 5. With $X_i(t; y)$ denoting the random quantile function corresponding to $\tilde{f}_i(\cdot; y)$ and $X_i^c(t; y) = X_i(t; y) - E\{X_i(t; y)\}$,

$$\int_{[0,1]^4} \text{cov}\{X_1^c(s; y)X_1^c(s; z), X_1^c(t; y)X_1^c(t; z)\} ds dt dy dz < \infty.$$

THEOREM 2. *Suppose Assumptions 1–5 hold and that $\mathcal{F}(y) \in \mathcal{D}$ almost surely. Then*

$$\int_0^1 \int_0^1 \left\{ \Sigma_{\oplus}(y, z) - \hat{\Sigma}_{\oplus}(y, z) \right\}^2 dy dz = O_p(n^{-1} + b_N^2).$$

4. APPLICATIONS

4.1. Functional Connectivity in Brain Imaging

The study of functional connectivity in the brain involves the identification of voxels or regions which exhibit similar behaviour, as quantified by neuroimaging techniques such as electroencephalography and functional magnetic resonance imaging. Of special interest are connections that are present when subjects are in the resting state (Allen et al., 2014). Studies of connections between spatially remote regions revealed the so-called default mode network in the resting brain, including nodes of high centrality that have been characterized as functional connectivity hubs (Buckner et al., 2009). The strength of connections between neighboring voxels, as opposed to those between remote regions, contains important information related to various biological factors (Gao et al., 2016) and neurological diseases (Zalesky et al., 2012). The strength of local connections in a particular brain region has been quantified in various ways (Tomasí & Volkow, 2010; Zang et al., 2004).

Petersen & Müller (2016) demonstrated the utility of a relatively simple approach to quantifying local connectivity within functional connectivity hubs that uses probability density functions formed by smoothing histograms of pairwise temporal correlations between the signals of each voxel within a hub and the signal at its central seed voxel. We implement this approach for the $p = 10$ hubs in Table 3 of Buckner et al. (2009). The resulting 10-dimensional vectors of densities that are obtained for each subject are then analyzed with the proposed Wasserstein covariance. The ten hubs of interest are located in the left and right parietal lobes, medial superior frontal lobe, medial prefrontal lobe, left and right middle frontal lobes, posterior cingulate/precuneus region, right supramarginal lobe and the left and right middle temporal lobes.

The densities were constructed for 171 cognitively normal subjects, 65 subjects with Alzheimer’s disease, and a third group of 86 subjects with mild cognitive impairment, with details about data preprocessing as in Petersen et al. (2016). The densities were then used to compute estimated Wasserstein mean densities $\hat{f}_{\oplus, j}$ for each group separately, which are displayed for each of the 10 hubs in Figure 1. The various groups show remarkable similarities in terms of the average Wasserstein connectivity distributions across these ten functional hubs. In contrast, the Wasserstein covariance and correlation matrices in Figure 2 exhibit clear differences between the groups and provide valuable additional information. Correlations are overall stronger for the normal subjects and those with mild cognitive impairment and exhibit different patterns of dependency across hubs. Complete second order behaviour can be studied through the estimated quantile covariance function estimates \hat{C}_{jk} , where for $p = 10$, there are 45 such functions for each group. We examine the off-diagonal elements for $j < k$ by plotting slices corresponding to the three quartiles $s = 0.25, 0.5, 0.75$; see the Supplementary Material.




Fig. 1. Wasserstein means for ten functional connectivity hubs for normal (left), mild cognitive impairment (middle) and Alzheimer's (right) subjects, displayed as densities.

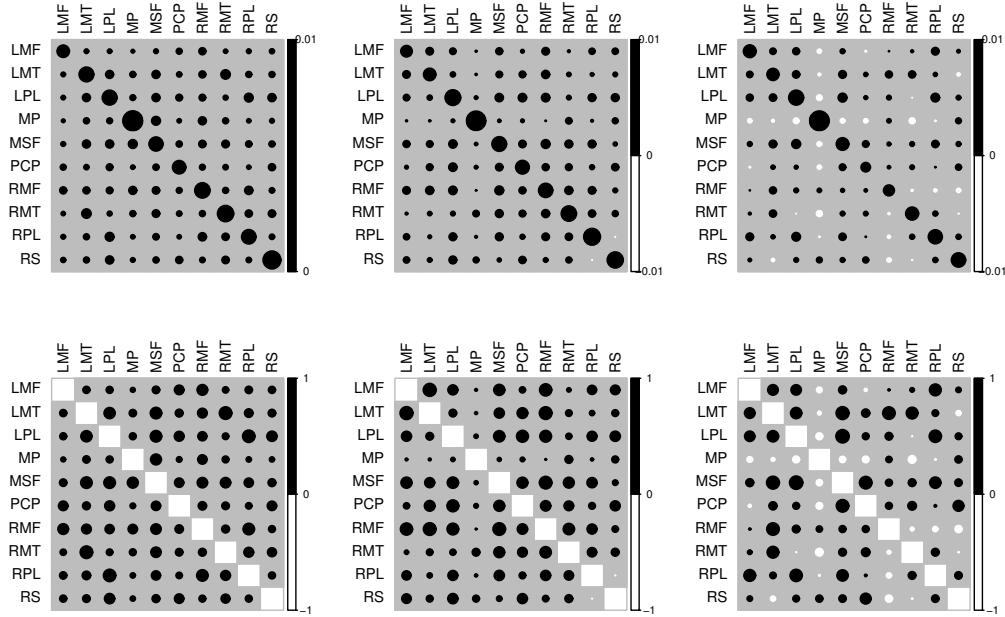


Fig. 2. Estimated Wasserstein covariance (top row) and correlation (bottom row) matrices for normal (left), mild cognitive impairment (middle) and Alzheimer's (right) subjects. Labels: LMF and RMF (left and right middle frontal), LPL and RPL (left and right parietal), LMT and RMT (left and right middle temporal), MSF (medial superior frontal), MP (medial prefrontal), PCP (posterior cingulate/precuneus) and RS (right supramarginal). Positive (negative) values are drawn in black (white) and larger circles correspond to larger absolute values.

We can also test for group differences in Wasserstein covariance using bootstrap samples obtained by centering all quantile functions with respect to their group means. To construct a bootstrap sample under the null hypothesis that all groups have the same Wasserstein covariance matrix, we simply center each multivariate quantile process at its corresponding group Wasser-




Fig. 3. Estimated Wasserstein correlation submatrices corresponding to lateral hub pairs for normal (left) and Alzheimer's (right) subjects, after reordering of hubs by hierarchical clustering using Ward's criterion. Rectangles indicate the groupings when three clusters are used. Labels correspond to those in Figure 2 and further explanations can be found in the caption of Figure 2.

stein mean. By pooling these centered processes, a bootstrap sample is obtained by sampling with replacement and then dividing into three groups of proper size. With estimates $\hat{\Sigma}_{\oplus}^n$, $\hat{\Sigma}_{\oplus}^m$ and $\hat{\Sigma}_{\oplus}^a$ for the groups of normal, mildly cognitively impaired, and Alzheimer's subjects, the test statistic for the global null hypothesis that all groups are equal is

$$S = \left\| \log \hat{\Sigma}_{\oplus}^n - \log \hat{\Sigma}_{\oplus}^m \right\|_F^2 + \left\| \log \hat{\Sigma}_{\oplus}^n - \log \hat{\Sigma}_{\oplus}^a \right\|_F^2 + \left\| \log \hat{\Sigma}_{\oplus}^m - \log \hat{\Sigma}_{\oplus}^a \right\|_F^2.$$

Using 1000 bootstrap samples, the global p -value was found to be $p = 0.015$. An alternative test, obtained by replacing the estimated Wasserstein covariance matrices with the corresponding correlations in the above statistic, resulted in $p = 0.093$.

To further explore the differences between normal and Alzheimer's groups, the matrices were reduced to a subset of rows/columns corresponding to the lateral hub pairs in the middle frontal, middle temporal and parietal regions, then reordered using Ward's hierarchical clustering algorithm with three clusters, visualized in Figure 3. This demonstrates the presence of asymmetry in the Alzheimer's disease group, which is absent in the normal group. Derflinger et al. (2011) and others report findings of similar asymmetries in the brains of Alzheimer's patients.

One can ask how the Wasserstein covariance approach compares to established approaches in brain imaging. A commonly used measure of local connectivity is regional homogeneity (Zang et al., 2004), corresponding to Kendall's coefficient of concordance between the BOLD signals at the seed voxel and those of its immediate neighbors. This scalar measure of regional homogeneity can be computed for each hub and each subject, resulting in a sample of ordinary multivariate data. The group covariance and correlation matrices are shown in the Supplementary Material. Some common patterns are seen, most notably an increased number of negative correlations for the Alzheimer's group. However, the differences are not as stark, and the regional homogeneity covariances do not reveal the asymmetry seen in the Wasserstein covariance analysis.

4.2. Distribution of Age at Death for Period Cohorts

To gain a better understanding of human longevity, the study of the temporal evolution of the distributions of age at death and their dependency structure over calendar time is of interest. The Human Mortality Database provides yearly mortality and population data for 38 countries at

`<www.mortality.org>`, which have been previously analyzed with various functional data analysis techniques (Hyndman & Shang, 2010; Chiou & Müller, 2009).

For a given country and calendar year, the probability distribution for mortality can be represented by its density. Consider a country for which life tables are available for the years y_j ($j = 1, \dots, p$). For integer-valued ages $a = 0, \dots, 110$, the life table provides the size of the population m_a at least a years old, normalized so that $m_0 = 100000$. These life tables were converted to histograms of age-at-death, which we then smoothed, applying the `hades` package, available at `<http://www.stat.ucdavis.edu/hades/>`, with a smoothing bandwidth of $h = 2$. This led to estimated densities of age-at-death on the age domain $[20\text{years}, 110\text{years}]$.

To illustrate the proposed methods for continuously varying densities, we considered 32 countries and identify a subgroup of $n_E = 8$ countries located in Eastern Europe for comparison with the remaining $n_O = 24$ countries. Densities were estimated for every year between 1985 and 2005. Wasserstein means are depicted in the Supplementary Material and are found to be quite similar for the two groups, demonstrating increasing longevity over calendar time. In contrast, the estimated Wasserstein covariance and correlation surfaces of these two groups differ quite drastically, as seen in Figure 4.

Examining the diagonals of the Wasserstein covariance plots, the Eastern European countries are characterized by stagnant Wasserstein variability until 1993 when it increases sharply, followed by a steady increase between 1998 and 2005. For the non-Eastern European countries, the Wasserstein variability is maximal in 1994. Wasserstein correlations reveal that mortality dependencies are high and roughly constant over time for the non-Eastern European countries, while they are weak for Eastern European countries between years before 1990 and those after 1995.

These exploratory findings are not altogether surprising, given the economic and societal upheaval in Eastern Europe beginning around 1990, which is seen to be reflected in the Wasserstein covariance. In addition to exploratory analysis, we implemented a bootstrap test as described in the previous section. With estimated Wasserstein covariance surfaces $\hat{\Sigma}_{\oplus}^E$ and $\hat{\Sigma}_{\oplus}^O$, the test statistic was computed as the square root distance between the associated operators (Pigoli et al., 2014) and implemented by computing the Frobenius distance between the principal square roots of the discretized matrix estimates. Computing this statistic for 1000 bootstrap samples, the p -value for the difference between Wasserstein covariances of Eastern European and other countries was 0.287, while it was 0.007 using Wasserstein correlations, providing some evidence for differences in the distributions of the density processes.

5. DISCUSSION

For studying the covariance structure of vectors of random densities, the Wasserstein approach is preferred over possible alternatives such as the transformation approach (Petersen & Müller, 2016) or compositional methods based on the Aitchison geometry (Egozcue et al., 2006). This is because of its convincing practical behavior for the construction of barycenters (Bolstad et al., 2003; Zhang & Müller, 2011) and the theoretically appealing connections with optimal transport.

Specifically, the transformation approach where densities are mapped to the entire Hilbert space L^2 by means of a suitable transformation such as the log-quantile density transformation, could be applied to the components of the vectors of densities, which would lead to unconstrained multivariate functional data. For such data, any one of numerous available measures of functional covariance and correlation (Leurgans et al., 1993; Dubin & Müller, 2005; Eubank & Hsing, 2008; Yang et al., 2011) could then be harnessed. One would need to choose among many possible covariance measures and transformation maps, none of which is isometric to the Wasserstein distance. The resulting metric distortions make such an approach difficult to interpret.

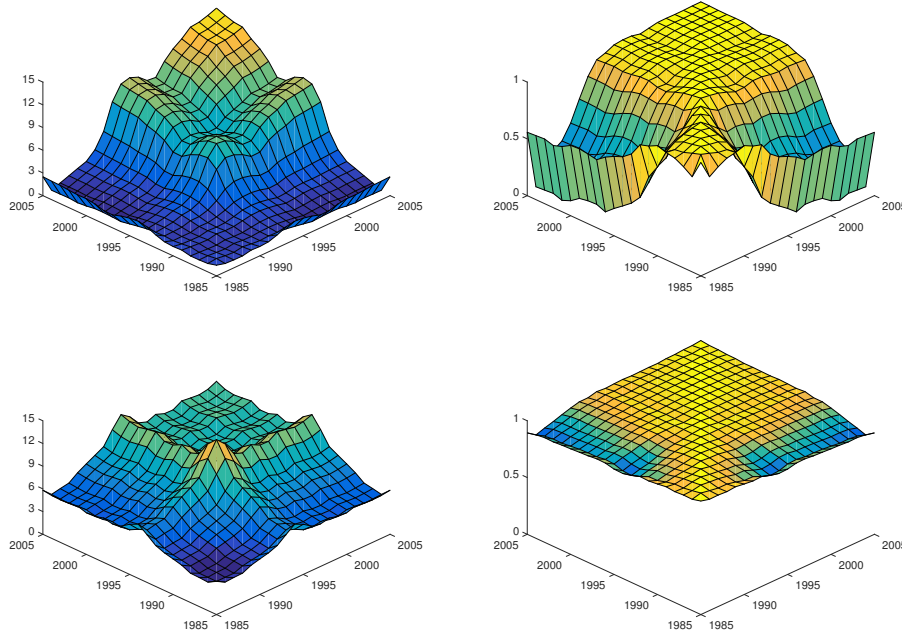


Fig. 4. Estimated Wasserstein covariance (left column) and correlation (right column) matrices for Eastern European (top row) and other (bottom row) countries.

In contrast, the proposed Wasserstein covariance has a canonical interpretation as an expected value of inner products of optimal transport maps. This means that the proposed Wasserstein covariance is similar to the notion of a regular covariance for an appropriate inner product, and can be interpreted as a measure of the degree of synchronization of the movement of probability mass from the Fréchet means to the random components of a bivariate density process. It thus emerges as a natural and compelling extension of the Wasserstein–Fréchet variance. The quantile function representation of the Wasserstein covariance in (6) facilitates the joint Wasserstein analysis of p one-dimensional distributions with Wasserstein covariance matrices and surfaces, enhancing the appeal of the proposed approach for practical applications.

While the geometric notion of covariance in (5) can be extended to the case where the sample densities have a multivariate domain, the implementation via quantile functions cannot be extended for this case, and therefore the asymptotic theory we provide remains limited to the case of one-dimensional densities. An extension to multivariate domains and analogous notions of covariance with respect to other metrics are open problems for future research.

REFERENCES

- ALLEN, E. A., DAMARAJU, E., PLIS, S. M., ERHARDT, E. B., EICHELE, T. & CALHOUN, V. D. (2014). Tracking whole-brain connectivity dynamics in the resting state. *Cerebral Cortex* **24**, 663–676.
- AMBROSIO, L., GIGLI, N. & SAVARÉ, G. (2008). *Gradient Flows in Metric Spaces and in the Space of Probability Measures*. Springer.
- BIGOT, J., CAZELLES, E. & PAPADAKIS, N. (2016). Regularization of barycenters in the Wasserstein space. *arXiv preprint arXiv:1606.01025*.
- BIGOT, J., GOUET, R., KLEIN, T. & LÓPEZ, A. (2017). Geodesic PCA in the Wasserstein space by convex PCA. *Annales de l’Institut Henri Poincaré B: Probability and Statistics* **53**, 1–26.

- BOLSTAD, B. M., IRIZARRY, R., ÅSTRAND, M. & SPEED, T. (2003). A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics* **19**, 185–193.
- 330 BUCKNER, R. L., SEPULCRE, J., TALUKDAR, T., KRIENEN, F. M., LIU, H., HEDDEN, T., ANDREWS-HANNA, J. R., SPERLING, R. A. & JOHNSON, K. A. (2009). Cortical hubs revealed by intrinsic functional connectivity: mapping, assessment of stability, and relation to Alzheimer’s disease. *The Journal of Neuroscience* **29**, 1860–1873.
- CHEN, K., DELICADO, P. & MÜLLER, H.-G. (2017). Modeling function-valued stochastic processes, with applications to fertility dynamics. *Journal of the Royal Statistical Society, Series B (Theory and Methodology)* **79**, 177–196.
- 335 CHIOU, J.-M. & MÜLLER, H.-G. (2009). Modeling hazard rates as functional data for the analysis of cohort lifetables and mortality forecasting. *Journal of the American Statistical Association* **104**, 572–585.
- DELICADO, P. (2011). Dimensionality reduction when data are density functions. *Computational Statistics and Data Analysis* **55**, 401–420.
- 340 DERFLINGER, S., SORG, C., GASER, C., MYERS, N., ARSIC, M., KURZ, A., ZIMMER, C., WOHLSCHLÄGER, A. & MÜHLAU, M. (2011). Grey-matter atrophy in Alzheimer’s disease is asymmetric but not lateralized. *Journal of Alzheimer’s Disease* **25**, 347–357.
- DUBIN, J. A. & MÜLLER, H.-G. (2005). Dynamical correlation for multivariate longitudinal data. *Journal of the American Statistical Association* **100**, 872–881.
- 345 EGOZCUE, J. J., DIAZ-BARRERO, J. L. & PAWLOWSKY-GLAHN, V. (2006). Hilbert space of probability density functions based on Aitchison geometry. *Acta Mathematica Sinica* **22**, 1175–1182.
- EUBANK, R. L. & HSING, T. (2008). Canonical correlation for stochastic processes. *Stochastic Processes and their Applications* **118**, 1634–1661.
- FRÉCHET, M. (1948). Les éléments aléatoires de nature quelconque dans un espace distancié. *Annales de l’Institut Henri Poincaré* **10**, 215–310.
- 350 GAO, Q., XU, F., JIANG, C., CHEN, Z., CHEN, H., LIAO, H. & ZHAO, L. (2016). Decreased functional connectivity density in pain-related brain regions of female migraine patients without aura. *Brain Research* **1632**, 73–81.
- HRON, K., MENAFOGLIO, A., TEMPL, M., HRUZOVÁ, K. & FILZMOSER, P. (2016). Simplicial principal component analysis for density functions in bayes spaces. *Computational Statistics & Data Analysis* **94**, 330–350.
- 355 HYNDMAN, R. J. & SHANG, H. L. (2010). Rainbow plots, bagplots and boxplots for functional data. *Journal of Computational and Graphical Statistics* **19**, 29–45.
- KNEIP, A. & UTIKAL, K. J. (2001). Inference for density families using functional principal component analysis. *Journal of the American Statistical Association* **96**, 519–542.
- LEURGANS, S. E., MOYEED, R. A. & SILVERMAN, B. W. (1993). Canonical correlation analysis when the data are curves. *Journal of the Royal Statistical Society: Series B* **55**, 725–740.
- 360 LI, Y. & HSING, T. (2010). Uniform convergence rates for nonparametric regression and principal component analysis in functional/longitudinal data. *Annals of Statistics* **38**, 3321–3351.
- PANARETOS, V. M. & ZEMEL, Y. (2016). Amplitude and phase variation of point processes. *The Annals of Statistics* **44**, 771–812.
- 365 PARK, S. Y. & STAICU, A.-M. (2015). Longitudinal functional data analysis. *Stat* **4**, 212–226.
- PETERSEN, A. & MÜLLER, H.-G. (2016). Functional data analysis for density functions by transformation to a Hilbert space. *Annals of Statistics* **44**, 183–218.
- PETERSEN, A., ZHAO, J., CARMICHAEL, O. & MÜLLER, H.-G. (2016). Quantifying individual brain connectivity with functional principal component analysis for networks. *Brain Connectivity* **6**, 540–547.
- 370 PIGOLI, D., ASTON, J. A., DRYDEN, I. L. & SECCHI, P. (2014). Distances and inference for covariance operators. *Biometrika* **101**, 409–422.
- TOMASI, D. & VOLKOW, N. D. (2010). Functional connectivity density mapping. *Proceedings of the National Academy of Sciences* **107**, 9885–9890.
- VILLANI, C. (2003). *Topics in Optimal Transportation*. American Mathematical Society.
- 375 YANG, W., MÜLLER, H.-G. & STADTMÜLLER, U. (2011). Functional singular component analysis. *Journal of the Royal Statistical Society: Series B* **73**, 303–324.
- YUAN, Y., ZHU, H., LIN, W. & MARRON, J. (2012). Local polynomial regression for symmetric positive definite matrices. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **74**, 697–719.
- 380 ZALESKY, A., FORNITO, A., EGAN, G. F., PANTELIS, C. & BULLMORE, E. T. (2012). The relationship between regional and inter-regional functional connectivity deficits in Schizophrenia. *Human Brain Mapping* **33**, 2535–2549.
- ZANG, Y., JIANG, T., LU, Y., HE, Y. & TIAN, L. (2004). Regional homogeneity approach to fMRI data analysis. *Neuroimage* **22**, 394–400.
- 385 ZHANG, Z. & MÜLLER, H.-G. (2011). Functional density synchronization. *Computational Statistics and Data Analysis* **55**, 2234–2249.