

STATISTICAL DATA ANALYSIS IN THE WASSERSTEIN SPACE *

JÉRÉMIE BIGOT¹

Abstract. This paper is concerned by statistical inference problems from a data set whose elements may be modeled as random probability measures such as multiple histograms or point clouds. We propose to review recent contributions in statistics on the use of Wasserstein distances and tools from optimal transport to analyse such data. In particular, we highlight the benefits of using the notions of barycenter and geodesic PCA in the Wasserstein space for the purpose of learning the principal modes of geometric variation in a dataset. In this setting, we discuss existing works and we present some research perspectives related to the emerging field of statistical optimal transport.

Résumé. Cet article est motivé par des problèmes d'inférence statistique à partir d'un ensemble de données dont les éléments peuvent être modélisés comme des mesures de probabilités tels que des histogrammes ou des nuages de points. Il est proposé de donner un aperçu des contributions récentes en statistique sur l'utilisation des distances de Wasserstein et des outils issus du transport optimal pour l'analyse de telles données. Nous insistons en particulier sur l'intérêt des notions de barycentre et d'ACP géodésique dans l'espace de Wasserstein dans le but d'apprendre les principaux modes de variation géométrique de ce type de données. Dans ce contexte, nous discutons des travaux existants et présentons quelques perspectives de recherche dans le domaine émergent du transport optimal pour la statistique.

1. THE EMERGING FIELD OF STATISTICAL OPTIMAL TRANSPORT

In many fields of interest (e.g. in signal and image processing or bio-informatics), one records data in the form of high-dimensional vectors or matrices. For the purpose of learning information from such data sets, a standard statistical analysis consists in considering that the observations are realizations of random variables belonging to a linear space endowed with an Euclidean distance. However, it has now been widely recognized that a significant gain in statistical inference can be achieved by the use of non-Euclidean distances to better capture the geometry of the sets to which the data (in the form of histograms, curves, images or points clouds) truly belong. Such a choice of non-Euclidean distances comes from the observation that the data, in many applications, tend to vary along geodesics that are not straight lines. Therefore, the use of non-Euclidean distances is often meaningful to extract geometric information on non-linear sources of variability. Concepts such as nonlinearity and an appropriate modeling of the location of signals or images intensities are essential tools to solve high-dimensional regression problems for the purpose of estimation and classification. The use of Wasserstein distances based on the mathematics of optimal mass transport [63] allows to incorporate both

* Jérémie Bigot is a member of Institut Universitaire de France (IUF), and this work has been carried out with financial support from the IUF. I would like to thank Hervé Cardot and Pierre Calka for proposing to write this review, as well as the two anonymous referees for their comments and suggestions of improvement.

¹ Institut de Mathématiques de Bordeaux et CNRS (UMR 5251), Université de Bordeaux

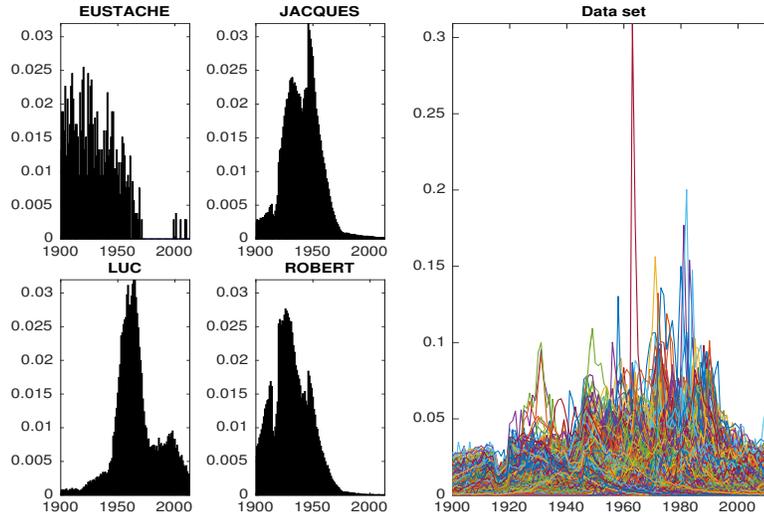


FIGURE 1. Children’s first name at birth - this dataset has been provided by the INSEE (French Institute of Statistics and Economic Studies). A subset of 4 histograms representing the distribution of children born with that name per year in France, and the whole dataset of $n = 1060$ histograms (right), displayed as pdf over the interval [1900, 2013]

spatial and intensity informations when comparing elements in a dataset that can be modeled as probability distributions. This approach leads to an intuitive geometric interpretation of the variability in such datasets.

Examples of such data are numerous. As an illustration, one may consider histograms that represent, for a list of first names, the distribution of children born with that name per year in France between years 1900 and 2013. In Figure 1, we display the histograms of four different names, as well as the whole dataset. Another example arises from the statistical analysis of data acquired from flow cytometry which is a high-throughput technique in biotechnology that can measure a large number of surface and intracellular markers of single cell in a biological sample. With this technique, one can assess individual characteristics (in the form of multivariate data) at a cellular level to determine the type of cell, their functionality and the way they differ. The development of this technology now leads to datasets made of multiple measurements (e.g. up to 18) of millions of individuals cells from different subjects. This leads to statistical inference problems from multiple points clouds such as those displayed in Figure 2 that can be modeled as discrete probability measures supported on \mathbb{R}^d .

Techniques based on optimal transport for data science have thus recently received an increasing interest in mathematical and computational statistics [8, 10–15, 28, 29, 44, 45, 47, 49, 51, 54, 58, 62, 66, 67], machine learning [4, 6, 22, 23, 32, 36–38, 40, 55, 57], image processing and computer vision [7, 17, 24, 30, 31, 41, 52, 59, 60] or computational biology [56].

This paper, whose content is biased by the point of view of a statistician, is an introduction to some recent developments of optimal transport for statistical data analysis. It is focused on the notions of barycenter and geodesic principal component analysis (PCA) in the Wasserstein space and their connection to nonparametric statistics and functional data analysis. We also report the results of numerical experiments to highlight, in an intuitive manner, the potential benefits of such tools for statistical inference. A recent review on the statistical aspects of Wasserstein distances with a rich bibliography on the mathematics of optimal transport can also be found in [50]. More generally, computational optimal transport for data science is a rapidly growing research field. For a detailed introduction and examples of application beyond the prism of statistics we refer to [25].

In Section 2, we present the notions of Wasserstein space and regularized optimal transport. Section 3 is focused on Wasserstein barycenters and statistical inference problems related to this concept. Finally, geodesic principal component analysis (GPCA) in the Wasserstein space is briefly presented in Section 4. Note that the

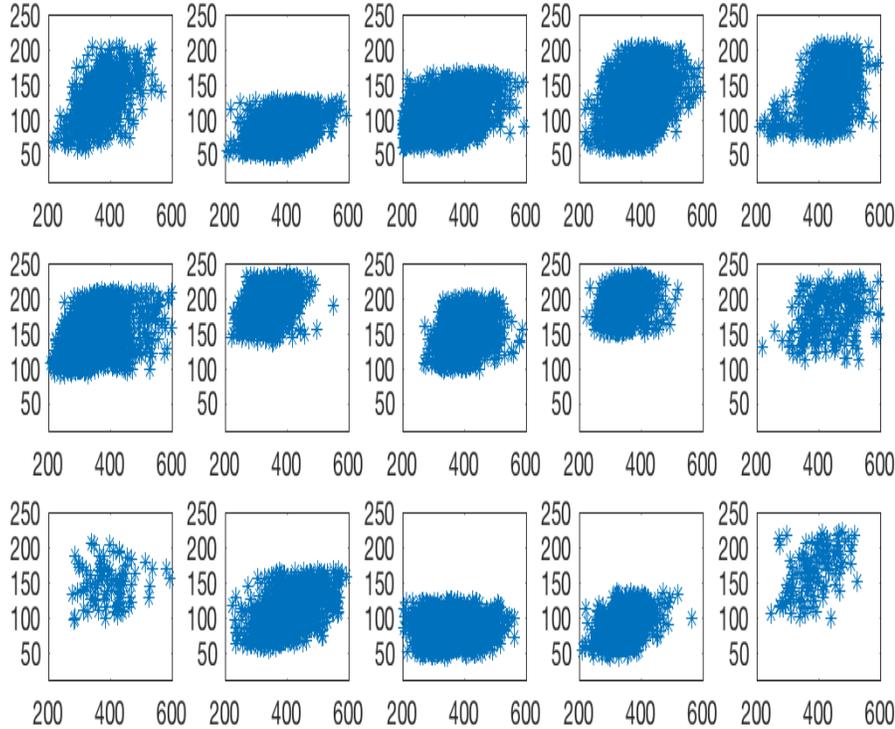


FIGURE 2. Example of flow cytometry data [42] obtained from a renal transplant retrospective study conducted by the Immune Tolerance Network (ITN) and measured from $n = 15$ patients (restricted to a bivariate projection). The horizontal axis - resp. vertical axis represents the values of the forward-scattered light (FSC) - resp. side-scattered light (SSC) cell marker, after an appropriate scaling through an arcsinh transformation and an initial gating on total lymphocytes to remove artefacts. The number of considered cells by patient varies from 88 to 2185.

presentation of the paper is rather informal on the mathematical aspects, and we mainly give intuition of the benefits of statistical optimal transport through the presentation of various numerical examples. Throughout the paper, we also discuss some research perspectives and open problems related to the statistical aspects of barycenters and GPCA in the Wasserstein space.

2. WASSERSTEIN SPACE AND REGULARIZED OPTIMAL TRANSPORT

2.1. The Wasserstein metric

We denote by $\mathcal{P}_2(\Omega)$ the space of probability measures with finite second moment and supports included in a convex domain $\Omega \subset \mathbb{R}^d$. This set can be endowed with the Wasserstein metric W_2 defined as

$$W_2(\mu, \nu) = \inf_{\pi \in \Gamma(\mu, \nu)} \left(\iint_{\Omega^2} |x - y|^2 d\pi(x, y) \right)^{1/2}, \quad \text{for } \mu, \nu \in \mathcal{P}_2(\Omega), \quad (1)$$

where $\Gamma(\mu, \nu)$ is the set of probability measures (also called transport plans or coupling measures) on the product space $\Omega \times \Omega$ with respective marginals μ and ν , and $|\cdot|$ denotes the usual Euclidean norm on \mathbb{R}^d . The metric space $(\mathcal{P}_2(\Omega), W_2)$ is called the (2-)Wasserstein space.

Remark 2.1. The minimisation problem (1) corresponds to the so-called Kantorovich formulation of optimal transport. Other cost functions than the quadratic one $c(x, y) = |x - y|^2$ may be used, but we restrict the discussion to this setting to simplify the presentation. For further details on the formulation of optimal transport for probability measures supported on general metric space we refer to [63].

For probability measures supported on the real (that is when $d = 1$), the Wasserstein metric has a closed form expression

$$W_2^2(\mu, \nu) = \int_0^1 (F_\mu^-(u) - F_\nu^-(u))^2 du, \quad (2)$$

for any $\mu, \nu \in \mathcal{P}_2(\mathbb{R})$, where F_μ^- (resp. F_ν^-) denotes the quantile function of the measure μ (resp. ν). Hence, in the one-dimensional setting, the use of the Wasserstein metric amounts to perform data analysis in the space of quantile functions (supported on $[0, 1]$) endowed with the usual Hilbert \mathbb{L}_2 metric. When μ and ν are absolutely continuous (a.c.) measures with square integrable probability density functions (pdf) f_μ and f_ν , this remark allows to easily illustrate the difference between choosing either the Wasserstein metric $W_2(\mu, \nu)$ or the Hilbert metric $\|f_\mu - f_\nu\|_2 := \left(\int_0^1 |f_\mu(x) - f_\nu(x)|^2 dx \right)^{1/2}$ for statistical analysis. Indeed, this choice implies two different constructions of the modes of variation in a dataset as illustrated in Figure 3, where we display the geodesics between two Gaussian distributions in the Hilbert space of densities and in the Wasserstein space of probability measure. Recall that, in the one-dimensional setting, the value of the geodesic at time $t \in [0, 1]$ between μ and ν (for the Wasserstein metric) is the measure μ_t whose quantile function is $F_{\mu_t}^- = (1 - t)F_\mu^- + tF_\nu^-$. Therefore, the measure μ_t remains Gaussian along the geodesic path between two Gaussian measures μ and ν . This illustrates the fact that performing statistical data analysis the Wasserstein space allows to simultaneously take into account variations in location and intensity (or mass) in a data set of densities (or histograms). A similar behavior holds in higher-dimensions ($d \geq 2$).

2.2. Regularized optimal transport

A serious limitation for the use of the Wasserstein metric for data analysis is its computational cost between probability measures supported in higher-dimensional spaces (that is beyond the case $d = 1$). Indeed, the computational cost to evaluate a Wasserstein metric between two discrete probability distributions with supports of equal size N is generally of order $\mathcal{O}(N^3 \log N)$. In this discrete setting, computing a Wasserstein distance amounts to solve the linear program (1) whose solution is constrained to belong to the convex set $\Gamma(\mu, \nu)$. Hence, the cost of this convex minimization becomes prohibitive for moderate to large values of N . To reduce the complexity of a linear program, a classical approach in optimization [65] is to add a regularizing entropy term. This is the approach followed in [22, 23] by adding entropy regularization to the transport transport plan in (1), which yields the strictly convex problem (3) below.

To simplify its presentation, we consider discrete probability measures supported on the same finite space $\Omega_N = \{x_1, \dots, x_N\} \in (\mathbb{R}^d)^N$ of cardinal N . A discrete probability distribution μ (with fixed support included in Ω_N) is identified by a vector of weights in the simplex $\Sigma_N = \{r = (r_1, \dots, r_N) \in \mathbb{R}_+^N \text{ with } \sum_{k=1}^N r_k = 1\}$ such that $\mu = \sum_{k=1}^N r_k \delta_{x_k}$ where δ_x is the Dirac distribution at x . The Wasserstein distance between two discrete distributions $\mu = \sum_{k=1}^N r_k \delta_{x_k}$ and $\nu = \sum_{k=1}^N q_k \delta_{x_k}$ (identified to vectors r and q in Σ_N) then becomes

$$W_2(r, q) := \min_{U \in U(r, q)} \langle C, U \rangle^{1/2},$$

where the set of coupling matrices is defined as $U(r, q) := \{U \in \mathbb{R}_+^{N \times N} \text{ such that } U \mathbf{1}_N = r, U^T \mathbf{1}_N = q\}$ with $\mathbf{1}_N$ the N dimensional vector with all entries equal to 1 and C the cost matrix given by $C_{ml} = |x_m - x_l|^2$, for all

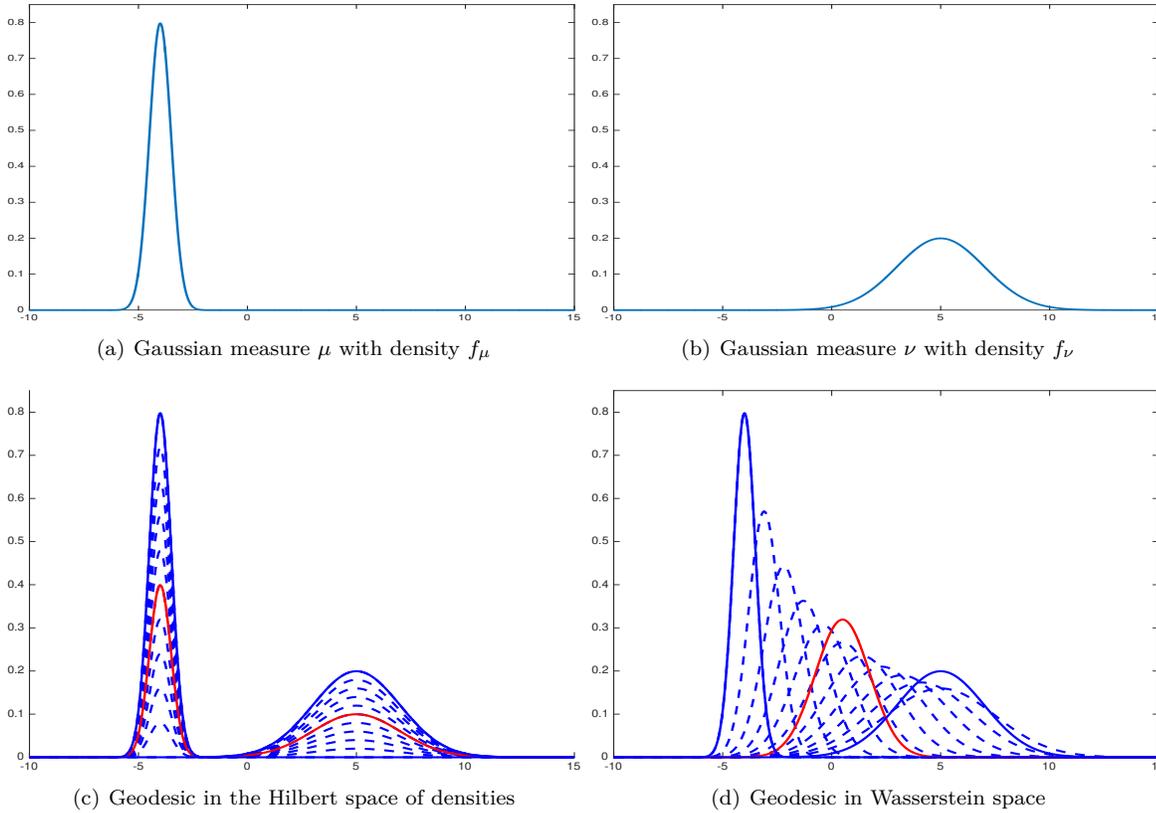


FIGURE 3. First row: densities of two Gaussian measures on \mathbb{R} . Second row: geodesics between μ and ν with respect to the Hilbert metric and the Wasserstein metric. The dotted curves in (c) correspond to the densities $f_t = (1-t)f_\mu + tf_\nu$, and in (d) they correspond to the densities of the measures with quantile functions $F_t^- = (1-t)F_\mu^- + tF_\nu^-$ for $0 < t < 1$. The red curves correspond to $t = 1/2$.

$m, l \in \{1, \dots, N\}$. Then, the notion of entropy regularized optimal transportation [19, 22] leads to the so-called Sinkhorn divergence $W_{2,\varepsilon}(r, q)$ between two discrete measures $r, q \in \Sigma_N$ which is defined as

$$W_{2,\varepsilon}^2(r, q) := \min_{U \in U(r, q)} \langle C, U \rangle - \varepsilon h(U). \quad (3)$$

where $\varepsilon > 0$ is a regularization parameter, and the discrete (negative) entropy for a given coupling matrix $U \in U(r, q)$ is given by $h(U) := -\sum_{m,l} U_{ml} \log U_{ml}$. The introduction of Sinkhorn divergence has been the starting point of the development of computational optimal transport [25] for machine learning as it makes feasible (from a computational point of view) the use of smoothed Wasserstein metrics for data analysis. Moreover, there now exists various toolbox and libraries to carry out data analysis using (possibly regularized) optimal transport in Matlab - <https://optimaltransport.github.io/>, Python - <https://pot.readthedocs.io/en/stable/> and the statistical computing environment R - <https://cran.r-project.org/package=Barycenter>.

3. WASSERSTEIN BARYCENTERS

The most basic statistical inference task from a data set is certainly to compute first order statistics that is to estimate an average location of the data. In this paper, we focus on the estimation of a population mean measure (or density function) from a dataset whose elements are probability measures. The notion of averaging depends on the metric that is chosen to compare elements in a given data set. In this paper, we consider the Wasserstein distance W_2 associated to the quadratic cost for the comparison of probability measures. As introduced in [2], an empirical Wasserstein barycenter $\bar{\mu}_n$ of set of n probability measures ν_1, \dots, ν_n (not necessarily random) in $\mathcal{P}_2(\Omega)$ is defined as a (possibly not unique) minimizer of

$$\mu \mapsto \frac{1}{n} \sum_{i=1}^n W_2^2(\mu, \nu_i), \text{ over } \mu \in \mathcal{P}_2(\Omega). \quad (4)$$

The Wasserstein barycenter corresponds to the notion of empirical Fréchet mean [35] that is an extension of the usual Euclidean barycenter to nonlinear metric spaces. For the case $d = 1$, the quantile formula (2) implies that $\bar{\mu}_n$ is the measure with quantile function $F_{\bar{\mu}_n}^- = \frac{1}{n} \sum_{i=1}^n F_{\nu_i}^-$. For $n = 2$, the Wasserstein barycenter $\bar{\mu}_2$ is the probability measure located at half distance (that is at “time $t = 1/2$ ”) along the geodesic between two absolutely continuous (a.c.) measures ν_1 and ν_2 . As illustrative examples, we display Wasserstein barycenters $\bar{\mu}_2$ between two probability measures in dimension $d = 1$ in Figure 3 and $d = 2$ in Figure 4.

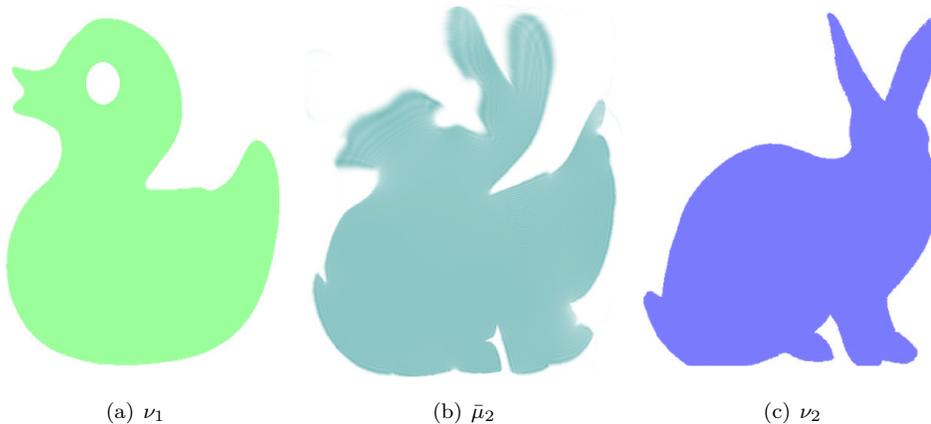


FIGURE 4. A numerical example taken from [17]. Wasserstein barycenter $\bar{\mu}_2$ in (b) between two probability measures ν_1 and ν_2 on \mathbb{R}^2 assumed to a.c. and uniform over their support displayed in (a) and (c).

Computing a Wasserstein barycenter between $n \geq 3$ probability measures is generally a delicate optimisation problem in dimension $d \geq 2$ that we do not discuss here. We refer to [7, 24, 25] for further details and references. In what follows, we focus on the mathematical statistical aspects of Wasserstein barycenters of random probability measures and the problem of their estimation from a nonparametric point of view.

3.1. Statistical models of random probability measures

Let us first introduce the notion of a population Wasserstein barycenter, as well as a two-sampling model of random probability measures that is appropriate to consider the problem of its nonparametric estimation. To this end, let \mathbb{P} be a distribution on the space of probability measures $(\mathcal{P}_2(\Omega), \mathcal{B}(\mathcal{P}_2(\Omega)))$, where $\mathcal{B}(\mathcal{P}_2(\Omega))$ is the Borel σ -algebra generated by the topology induced by W_2 . Throughout the paper, it is always assumed that \mathbb{P}

satisfies the square-integrability condition

$$\int_{\mathcal{P}_2(\Omega)} W_2^2(\mu, \nu) d\mathbb{P}(\nu) < +\infty, \text{ for some } \mu \in \mathcal{P}_2(\Omega).$$

Definition 3.1. A **population Wasserstein barycenter** $\mu_{\mathbb{P}}$ of a distribution \mathbb{P} over $\mathcal{P}_2(\Omega)$ is defined as a (possibly not unique) minimizer of

$$\mu \mapsto \int_{\mathcal{P}_2(\Omega)} W_2^2(\mu, \nu) d\mathbb{P}(\nu), \text{ over } \mu \in \mathcal{P}_2(\Omega). \quad (5)$$

For a discrete distribution $\mathbb{P}_n = \frac{1}{n} \sum \delta_{\nu_i}$ on $\mathcal{P}_2(\Omega)$, one has that $\mu_{\mathbb{P}_n}$ corresponds to the empirical Wasserstein barycenter defined as a minimizer of (4). The existence and unicity of Wasserstein barycenters has been discussed in various works [2,3,15,43,47]. As argued in [3], a sufficient condition for the unicity of $\mu_{\mathbb{P}}$ is that the distribution \mathbb{P} gives a strictly positive mass to the subset $\mathcal{P}_2^{ac}(\Omega)$ of a.c. measures in $\mathcal{P}_2(\Omega)$. Moreover, if \mathbb{P} is supported on the set of measures $\mathcal{P}_2(\Omega) \cap \mathbb{L}_q(\Omega)$ for some $q \in (1, +\infty)$ (i.e. measures in $\mathcal{P}_2^{ac}(\Omega)$ with $\mathbb{L}_q(\Omega)$ densities), it follows that $\mu_{\mathbb{P}}$ admits a density in $\mathbb{L}_q(\Omega)$.

To define an appropriate sampling model, let us now assume that ν_1, \dots, ν_n are independent and identically distributed a.c. random measures sampled from a distribution \mathbb{P} satisfying $\mathbb{P}(\mathcal{P}_2^{ac}(\Omega) \cap \mathbb{L}_2(\Omega)) = 1$. Hence, under such assumptions, the population Wasserstein barycenter $\mu_{\mathbb{P}}$ is unique and absolutely continuous. As raw data in the form of densities are generally not directly available, we focus on the setting where one has access to a set of random vectors $(\mathbf{X}_{i,j})_{1 \leq j \leq p_i; 1 \leq i \leq n}$ in \mathbb{R}^d organized in the form of n experiment units, such that, conditionally on ν_i , one has that $\mathbf{X}_{i,1}, \dots, \mathbf{X}_{i,p_i}$ are iid observations sampled from the measure ν_i for each $1 \leq i \leq n$. The observed data are thus in the form of multiple point clouds $\mathbf{X} = (\mathbf{X}_{i,j})_{1 \leq i \leq n; 1 \leq j \leq p_i}$. Elements in such a set can be modeled as n discrete probability measures

$$\nu_{p_i} = \frac{1}{p_i} \sum_{j=1}^{p_i} \delta_{\mathbf{X}_{i,j}}, \quad (6)$$

from which one may construct another empirical Wasserstein barycenter $\hat{\mu}_{n,p}$ defined as a (not necessarily unique) minimizer of

$$\mu \mapsto \frac{1}{n} \sum_{i=1}^n W_2^2(\mu, \nu_{p_i}), \text{ over } \mu \in \mathcal{P}_2(\Omega), \quad (7)$$

with the convention that $p = \min_{1 \leq i \leq n} p_i$. We shall now discuss the issue of estimating the population Wasserstein barycenter $\mu_{\mathbb{P}}$ using either $\mu_{\mathbb{P}_n}$ (with $\mathbb{P}_n = \frac{1}{n} \sum \delta_{\nu_i}$) or $\hat{\mu}_{n,p}$ as $n \rightarrow +\infty$ and $p \rightarrow +\infty$.

3.2. Convergence of empirical Wasserstein barycenters

3.2.1. Law of large numbers

By the law of large numbers, we refer to the question of showing that $W_2(\mu_{\mathbb{P}_n}, \mu_{\mathbb{P}})$ almost surely converges to zero as $n \rightarrow +\infty$. This problem has been addressed in [15] for a parametric model of random distributions \mathbb{P} over $\mathcal{P}_2(\Omega)$ with Ω assumed to be compact. General results for arbitrary distributions can be found in [47]. Beyond the law of large numbers, it is natural to ask for a (functional) central limit theorem for the empirical Wasserstein barycenter. Nevertheless, this notion is more delicate to formulate as the Wasserstein space is not a Hilbert space. For recent contributions in this direction (but in specific parametric settings) we refer to [3,45].

3.2.2. Rate of convergence in the one dimensional case

Let us now restrict to the one dimensional case of probability measure in $\mathcal{P}_2(\mathbb{R})$ and assume that $p_1 = \dots = p_n = p$. Thanks to the quantile formula (2) for the expression of the Wasserstein metric when $d = 1$, it follows

that the empirical Wasserstein barycenter $\hat{\boldsymbol{\mu}}_{n,p}$ minimizing (7) is the unique discrete measure

$$\hat{\boldsymbol{\mu}}_{n,p} = \frac{1}{p} \sum_{j=1}^p \delta_{\bar{\mathbf{X}}_j^*}, \quad (8)$$

where $\mathbf{X}_{i,1}^* \leq \mathbf{X}_{i,2}^* \leq \dots \leq \mathbf{X}_{i,p}^*$ are the order statistics the i -th sample, and $\bar{\mathbf{X}}_j^* = \frac{1}{n} \sum_{i=1}^n X_{i,j}^*$ for all $1 \leq j \leq p$. Moreover, as shown in [13], the population Wasserstein barycenter $\mu_{\mathbb{P}}$ is the unique a.c. measure with quantile function

$$F_0^-(\alpha) = \mathbb{E}(\mathbf{F}^-(\alpha)) \text{ where } \mathbf{F}^- \text{ is the quantile function of the random measure } \boldsymbol{\nu} \text{ with distribution } \mathbb{P}.$$

In this setting, the following upper bound on the expected quadratic risk $\mathbb{E}(W_2^2(\hat{\boldsymbol{\mu}}_{n,p}, \mu_{\mathbb{P}}))$ has been obtained in [13].

Theorem 3.2. *Let Y_1, \dots, Y_p be iid variables sampled from $\mu_{\mathbb{P}}$, and $\boldsymbol{\mu}_p = \frac{1}{p} \sum_{k=1}^p \delta_{Y_k}$. Then, the estimator $\hat{\boldsymbol{\mu}}_{n,p}$ satisfies*

$$\mathbb{E}(W_2^2(\hat{\boldsymbol{\mu}}_{n,p}, \mu_{\mathbb{P}})) \leq \frac{1}{n} \int_0^1 \text{Var}(\mathbf{F}^-(\alpha)) d\alpha + \mathbb{E}(W_2^2(\boldsymbol{\mu}_p, \mu_{\mathbb{P}})). \quad (9)$$

Hence, the rate of convergence of $\hat{\boldsymbol{\mu}}_{n,p}$ is decomposed into the sum of two terms. The first one goes to zero at the parametric rate $\mathcal{O}(\frac{1}{n})$, while the second one depends on the rate of convergence of the empirical measure $\boldsymbol{\mu}_p$ towards its population counterpart in expected squared Wasserstein distance. For $d = 1$, the question of deriving upper and lower bound for $\mathbb{E}(W_2^2(\boldsymbol{\mu}_p, \mu_{\mathbb{P}}))$ has been thoroughly studied in [16]. In arbitrary dimension d , there is a rich literature in probability on studying the convergence rate of an empirical measure towards its population version in Wasserstein distance. We refer to [9, 16, 27, 34, 64] for extensive references and the most recent contributions on this topic. For example, using results from [16], one has that

$$\mathbb{E}(W_2^2(\boldsymbol{\mu}_p, \mu_{\mathbb{P}})) \leq \frac{2}{p+1} J_2(\mu_{\mathbb{P}}), \text{ where } J_2(\mu_{\mathbb{P}}) = \int_{\Omega} \frac{F_0(x)(1-F_0(x))}{f_0(x)} dx,$$

where F_0 is the cumulative distribution function of $\mu_{\mathbb{P}}$ and f_0 denotes its density. Hence, if $J_2(\mu_{\mathbb{P}}) < +\infty$, it follows that $\hat{\boldsymbol{\mu}}_{n,p}$ converges to zero at the parametric rate $\mathcal{O}(\frac{1}{n} + \frac{1}{p})$. A detailed discussion on the derivation of upper bounds on the rate of convergence of $\hat{\boldsymbol{\mu}}_{n,p}$ is proposed in [13] with some results on minimax optimality of such bounds.

3.2.3. Open problems in dimension $d \geq 2$

Beyond the one-dimensional case $d = 1$, it appears that deriving the rate of convergence for empirical Wasserstein barycenters, e.g. for the expected quadratic risk $\mathbb{E}(W_2^2(\hat{\boldsymbol{\mu}}_{n,p}, \mu_{\mathbb{P}}))$, is a much more involved problem. A few works [1, 45] exist on the convergence rate of $\mu_{\mathbb{P}_n}$ to $\mu_{\mathbb{P}}$. However, the more general question of deriving the convergence rate of a Wasserstein barycenter $\hat{\boldsymbol{\mu}}_{n,p}$ (7) obtained from samples of random measure is open for $d \geq 2$ (to the best of our knowledge).

3.3. Regularization of Wasserstein barycenters

An empirical Wasserstein barycenter $\hat{\boldsymbol{\mu}}_{n,p}$ of the discrete measures $\boldsymbol{\nu}_{p_1}, \dots, \boldsymbol{\nu}_{p_n}$, defined by (6), is generally irregular (and even not unique). Irregularity has to be understood in the sense that $\hat{\boldsymbol{\mu}}_{n,p}$ is, in most situations, a discrete probability measure as shown by its expression (8) when $d = 1$. Therefore, as a discrete measure, it poorly represents the smoothness of the population Wasserstein barycenter of a random a.c. measure $\boldsymbol{\nu}$ with distribution \mathbb{P} , as $\mu_{\mathbb{P}}$ is absolutely continuous in this setting as discussed previously. Hence, we now explain how regularizing the Wasserstein barycenter $\hat{\boldsymbol{\mu}}_{n,p}$ in order to construct a.c. and consistent estimators of an a.c. population barycenter in the asymptotic setting where both n and $p = \min_{1 \leq i \leq n} p_i$ tend to infinity. To this end, one may consider two ways of regularizing an empirical barycenter.

3.3.1. Penalized Wasserstein barycenters

A first possibility is to follow the approach proposed in [12]. Thus, we now introduce the penalized Wasserstein barycenter associated to the distribution \mathbb{P} that is defined as a solution of the minimization problem

$$\min_{\mu \in \mathcal{P}_2(\Omega)} \int_{\mathcal{P}_2(\Omega)} W_2^2(\mu, \nu) d\mathbb{P}(\nu) + \gamma E(\mu)$$

where $\gamma > 0$ is a penalization parameter and the penalty function writes

$$E(\mu) = \begin{cases} \|f\|_{H^k(\Omega)}^2, & \text{if } f = \frac{d\mu}{dx} \text{ and } f \geq \alpha, \\ +\infty & \text{otherwise.} \end{cases} \quad (10)$$

where $\|\cdot\|_{H^k(\Omega)}$ denotes the Sobolev norm associated to the $\mathbb{L}^2(\Omega)$ space, $\alpha > 0$ is arbitrarily small and $k > d-1$. The choice (10) for the penalty function E is motivated by the discussion in Section 5 of [12]. Its main advantages are to impose that the penalized Wasserstein barycenter is absolutely continuous with a smooth density. Other examples of penalty functions are discussed in [12] including the class of relative G-functional described in Section 9.4 in [5].

Definition 3.3. For $\gamma > 0$, we consider the Wasserstein barycenters:

$$\mu_{\mathbb{P}}^{\gamma} = \arg \min_{\mu \in \mathcal{P}_2(\Omega)} \int_{\mathcal{P}_2(\Omega)} W_2^2(\mu, \nu) d\mathbb{P}(\nu) + \gamma E(\mu) \quad (11)$$

$$\hat{\mu}_{n,p}^{\gamma} = \arg \min_{\mu \in \mathcal{P}_2(\Omega)} \frac{1}{n} \sum_{i=1}^n W_2^2(\mu, \nu_{p_i}) + \gamma E(\mu) \quad (12)$$

called respectively penalized population barycenter (11) and penalized empirical barycenter (12).

Existence and uniqueness of these barycenters are discussed in [12] for a large class of penalty functions enforcing them to be a.c. One can thus let $\hat{f}_{n,p}^{\gamma}$ (resp. $f_{\mathbb{P}}^{\gamma}$) be the densities associated to $\hat{\mu}_{n,p}^{\gamma}$ (resp. $\mu_{\mathbb{P}}^{\gamma}$). The choice (10) for the penalty function E imposes that the densities of the penalized Wasserstein barycenters are bounded from below by a small constant $\alpha > 0$ over their support that is equal to Ω . In this setting, the following result (see Section 5 in [12]) gives the convergence rate in expected squared $\mathbb{L}_2(\Omega)$ -distance between these densities for a fixed value of $\gamma > 0$.

Theorem 3.4 (Section 5 in [12]). *Assume that $\Omega \subset \mathbb{R}^d$ is compact. Let $\hat{f}_{n,p}^{\gamma}$ and $f_{\mathbb{P}}^{\gamma}$ denote the density functions of $\hat{\mu}_{n,p}^{\gamma}$ and $\mu_{\mathbb{P}}^{\gamma}$, induced by the choice (10) of the penalty function E . Then, provided that $d < 4$, there exists a constant $c > 0$ depending only on Ω such that*

$$\mathbb{E} \left(\|\hat{f}_{n,p}^{\gamma} - f_{\mathbb{P}}^{\gamma}\|_{\mathbb{L}_2(\Omega)}^2 \right) \leq c \left(\frac{1}{\gamma p^{1/4}} + \frac{1}{\gamma n^{1/2}} \right) \quad (13)$$

where $p = \min_{1 \leq i \leq n} p_i$.

Letting $f_{\mathbb{P}}^0$ be the density of the population Wasserstein barycenter $\mu_{\mathbb{P}}$, the convergence of the approximation error term $\|f_{\mathbb{P}}^{\gamma} - f_{\mathbb{P}}^0\|_{\mathbb{L}_2(\Omega)}^2$ is also discussed in [12], where it is shown that it converges to zero as $\gamma \rightarrow 0$. However, no rate of convergence (as a function of γ) has been derived in [12] for this approximation term, and deriving such a result remains an open problem for regularized Wasserstein barycenters. Therefore, only the consistency of $\hat{f}_{n,p}^{\gamma}$ has been derived in [12] in the sense that the quadratic risk $\mathbb{E} \left(\|\hat{f}_{n,p}^{\gamma} - f_{\mathbb{P}}^0\|_{\mathbb{L}_2(\Omega)}^2 \right) \rightarrow 0$ under the assumption that $\gamma = \gamma_{n,p}$ decays to zero as $\min(n, p) \rightarrow +\infty$ at a rate which guarantees that the right hand size of inequality (13) converges to zero. In Theorem 3.4, the condition $d < 4$ may be relaxed at the price of obtaining a slower rate of convergence as $p \rightarrow +\infty$ (see Section 5 in [12]).

Remark 3.5. Consistency of regularized empirical Wasserstein barycenter is discussed in [12] for the Bregman divergence associated to the penalization function E . The specific choice (10) for E finally yields to the use of the $\mathbb{L}_2(\Omega)$ norm to compare the densities of $\hat{\mathbf{f}}_{n,p}^\gamma$, $f_{\mathbb{P}}^\gamma$ and $f_{\mathbb{P}}^0$. The question of deriving consistency results for $\hat{\boldsymbol{\mu}}_{n,p}^\gamma$ using the Wasserstein metric (that is in the space of probability measures and not densities) has not been considered in [12]

3.3.2. Sinkhorn barycenters

Another way to regularize an empirical Wasserstein barycenter is to use the Sinkhorn divergence (3) to compare probability measures. As this divergence is defined for discrete measures supported on the same points, one first needs to choose a fixed grid $\Omega_N = \{x_1, \dots, x_N\}$ made of N points $x_k \in \mathbb{R}^d$ (bin locations that are typically equally spaced). Then, one performs a binning of the data (6) on this grid, leading to a dataset of discrete measures (with supports included in Ω_N) that we denote

$$\tilde{\mathbf{q}}_i^{p_i} = \frac{1}{p_i} \sum_{j=1}^{p_i} \delta_{\tilde{X}_{i,j}}, \quad \text{where } \tilde{X}_{i,j} = \arg \min_{x \in \Omega_N} |x - X_{i,j}|, \quad (14)$$

for $1 \leq i \leq n$. Binning (i.e. choosing the grid Ω_N) surely incorporates some sort of additional regularization, and a discussion on the influence of the grid size N on the smoothness of the barycenter can be found in [11]. The size N of the grid is also guided by numerical issues on the computational cost of the algorithms used to approximate a Sinkhorn divergence.

Then, another possibility to introduce regularization in the computation of Wasserstein barycenters [11,23,24] is to consider the following estimator

$$\hat{\mathbf{r}}_{n,p}^\varepsilon = \arg \min_{r \in \Sigma_N} \frac{1}{n} \sum_{i=1}^n W_{2,\varepsilon}^2(r, \tilde{\mathbf{q}}_i^{p_i}), \quad (15)$$

that can be interpreted as a Fréchet mean with respect to a Sinkhorn divergence. We call $\hat{\mathbf{r}}_{n,p}^\varepsilon$ a Sinkhorn barycenter. From the point of view of statistics, it can be interpreted as a M-estimator defined over the compact set Σ_N . Two key properties of the Sinkhorn divergence have been used in [11] to study the convergence properties of such barycenters:

- (i): for any $q \in \Sigma_N$, the function $r \mapsto W_{2,\varepsilon}^2(r, q)$ is ε -strongly convex for the Euclidean 2-norm
- (ii): let $q \in \Sigma_N$ and $0 < \rho < 1/N$ an arbitrarily small constant. Then, one has that $r \mapsto W_{2,\varepsilon}^2(r, q)$ is $L_{\rho,\varepsilon}$ -Lipschitz on

$$\Sigma_N^\rho = \left\{ r \in \Sigma_N : \min_{1 \leq \ell \leq N} r_\ell \geq \rho \right\},$$

with

$$L_{\rho,\varepsilon} = \left(\sum_{1 \leq m \leq N} \left(2\varepsilon \log(N) + \sup_{1 \leq \ell, k \leq N} |C_{m\ell} - C_{k\ell}| - 2\varepsilon \log(\rho) \right)^2 \right)^{1/2}$$

For a discussion on the behavior of the Lipschitz constant $L_{\rho,\varepsilon}$ as ρ and ε tend to zero, we refer to Section 3.1 in [11]. For example, if $\rho = \varepsilon^\kappa$ for some $\kappa > 0$, then

$$\lim_{\varepsilon \rightarrow 0} L_{\rho,\varepsilon} = \left(\sum_{1 \leq m \leq N} \left(\sup_{1 \leq \ell, k \leq N} |C_{m\ell} - C_{k\ell}| \right)^2 \right)^{1/2}.$$

To guarantee the Lipschitz continuity of the mapping $r \mapsto W_{2,\varepsilon}^2(r, q)$, the analysis of the statistical properties of Sinkhorn barycenters in [11] is restricted to discrete measures belonging to the convex set Σ_N^ρ (thus having non-vanishing entries). Therefore, one has to slightly modify the data at hand so that the observations belong to

Σ_N^ρ . This yields to the following definitions of sampling model, and empirical/population Sinkhorn barycenters in Σ_N^ρ .

Definition 3.6. Let $0 < \rho < 1/N$, and \mathbb{P} be a probability distribution on Σ_N^ρ . Let $\mathbf{q}_1, \dots, \mathbf{q}_n \in \Sigma_N^\rho$ be an iid sample drawn from the distribution \mathbb{P} . Assume that $(\tilde{X}_{i,j})_{1 \leq j \leq p_i}$ are iid random variables sampled from \mathbf{q}_i , for each $1 \leq i \leq n$ and consider the discrete measures $\hat{\mathbf{q}}_i^{p_i} = \frac{1}{p_i} \sum_{j=1}^{p_i} \delta_{\tilde{X}_{i,j}}$ and $\tilde{\mathbf{q}}_i^{p_i} = (1 - \rho N) \hat{\mathbf{q}}_i^{p_i} + \rho \mathbb{1}_N$, where $\mathbb{1}_N$ is the vector of \mathbb{R}^N with all entries equal to one. Then, we define

$$r^\varepsilon = \arg \min_{r \in \Sigma_N^\rho} \mathbb{E}_{\mathbf{q} \sim \mathbb{P}} [W_{2,\varepsilon}^2(r, \mathbf{q})] \quad \text{the population Sinkhorn barycenter} \quad (16)$$

$$\hat{r}_{n,p}^\varepsilon = \arg \min_{r \in \Sigma_N^\rho} \frac{1}{n} \sum_{i=1}^n W_{2,\varepsilon}^2(r, \hat{\mathbf{q}}_i^{p_i}) \quad \text{the empirical Sinkhorn barycenter} \quad (17)$$

Theorem 3.7 (Section 3 in [11]). *Let $\varepsilon > 0$ and $p = \min_{1 \leq i \leq n} p_i$. Then, one has that*

$$\mathbb{E}(\|r^\varepsilon - \hat{r}_{n,p}^\varepsilon\|_{\mathbb{R}^N}^2) \leq \frac{32L_{\rho,\varepsilon}^2}{\varepsilon^2 n} + \frac{2L_{\rho,\varepsilon}}{\varepsilon} \left(\sqrt{\frac{N}{p}} + 2\rho(N + \sqrt{N}) \right). \quad (18)$$

The upper bound (18) allows to shed some light on how the expected squared distance between the empirical Sinkhorn barycenter and its population version is influenced by the number n of observed measures, the minimal number p of samples per measure and the size N of the grid. However, this upper bound does not converge to zero as $\min(n, p) \rightarrow +\infty$ if ρ remains fixed. The main advantage of this upper bound is to allow the derivation of a data-driven strategy to select ε that is presented in the following subsection below. Finally, one should remark that the behavior of the population Sinkhorn barycenter r^ε as $\varepsilon \rightarrow 0$ has not been considered in [11].

Remark 3.8. A third way to regularize an empirical Wasserstein barycenter (that might look more natural for statisticians) is to perform a preliminary smoothing step of the discrete measures $\nu_{p_1}, \dots, \nu_{p_n}$. For example, one may first smooth the data using standard kernel density estimation $\hat{f}_i^{h_i}(x) = \frac{1}{p_i h_i^d} \sum_{j=1}^{p_i} K\left(\frac{x - \mathbf{X}_{i,j}}{h_i}\right)$ and then define a smoothed Wasserstein barycenter as

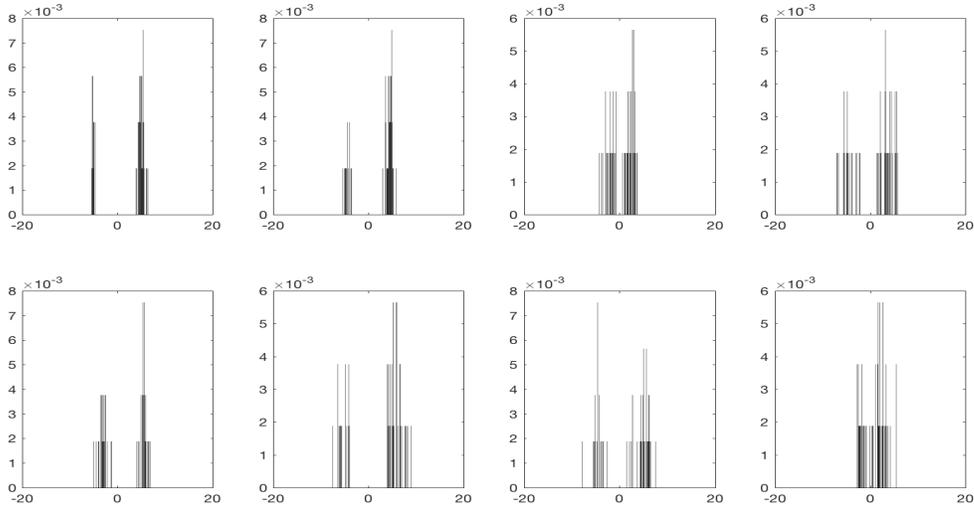
$$\hat{\mu}_{n,p}^h = \arg \min_{\mu \in \mathcal{P}_2(\Omega)} \frac{1}{n} \sum_{i=1}^n W_2^2(\hat{\nu}_i^{h_i}, \mu),$$

where $d\hat{\nu}_i^{h_i}(x) = \hat{f}_i^{h_i}(x)dx$, and $h = \max_{1 \leq i \leq n} h_i$. The consistency of this approach (as $h \rightarrow 0$) has been studied in details by [49, 50] using multiple point processes for the sampling procedure, and it is also discussed in [13].

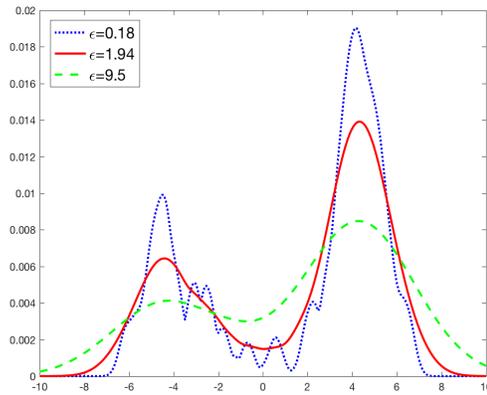
3.4. Data-driven regularization in computational optimal transport

In this paper, we report results on numerical experiments from [11] that are based on the approach proposed in [24] to compute Sinkhorn barycenters on a grid Ω_N of equi-spaced points. We do not report results on the computational aspects of penalized Wasserstein barycenters, and we refer to [11] for illustrations on their numerical performances.

As any nonparametric method, the use of regularized Wasserstein barycenters involves the delicate calibration of a regularization parameter. Choosing the value of ε for the Sinkhorn barycenter $\hat{r}_{n,p}^\varepsilon$ might be guided by the classical bias and variance tradeoff principle as illustrated by the synthetic example displayed in Figure 5. For small values of ε the Sinkhorn barycenter shows many oscillations while it is much smoother for larger values of ε . In some sense, the regularization parameters ε can be interpreted as the usual bandwidth parameter in kernel density estimation and its choice greatly influences the shape of $\hat{r}_{n,p}^\varepsilon$.



(a)



(b)

FIGURE 5. (a) A subset of 8 histograms (out of $n = 15$) obtained with random variables sampled from one-dimensional Gaussian mixtures distributions ν_i (with random means and variances). The population Wasserstein of ν_1, \dots, ν_n is expected to be a mixture of two Gaussian distributions. Histograms are constructed by binning the data $(X_{i,j})_{1 \leq i \leq n; 1 \leq j \leq p}$ on a grid Ω_N of size $N = 2^8$ with $p_1 = \dots = p_n = p = 50$. (b) Three Sinkhorn barycenters $\hat{r}_{n,p}^\varepsilon$ associated to the parameters $\varepsilon = 0.18, 1.94, 9.5$.

Let us now briefly discuss the question of choosing ε in an automatic way. In [11], the following (un-regularized) population Wasserstein is considered

$$r^0 \in \arg \min_{r \in \Sigma_N^p} \mathbb{E}_{q \sim \mathbb{P}} [W_{2,0}^2(r, q)] \quad \text{with} \quad W_{2,0}^2(r, q) := \min_{U \in U(r, q)} \langle C, U \rangle. \quad (19)$$

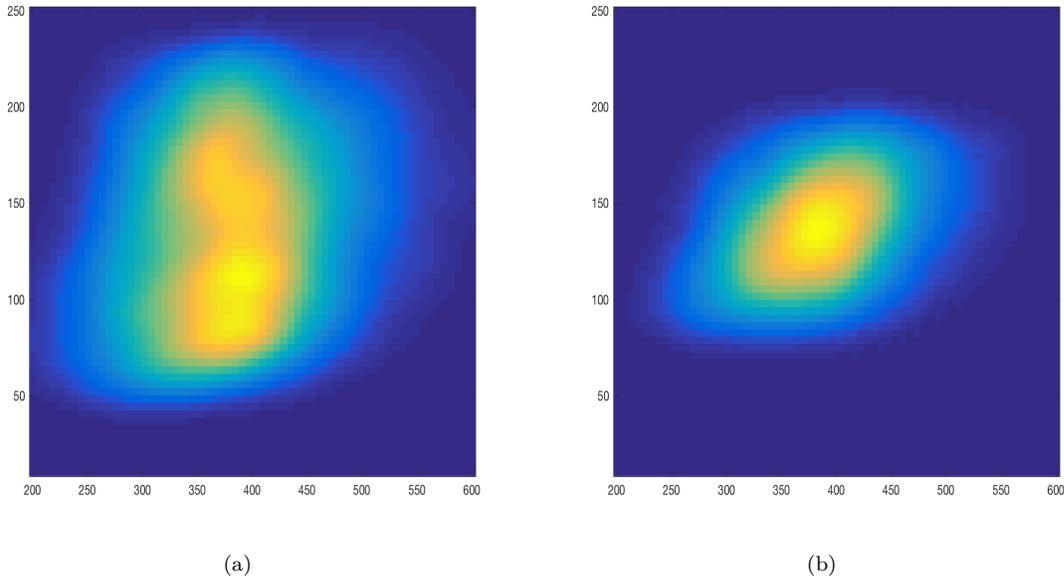


FIGURE 6. Two dimensional flow cytometry dataset. (a) Euclidean mean of the data (after smoothing but without a preliminary step to correct misalignment effects in the raw data), (b) Sinkhorn barycenter $\hat{r}_{n,p}^{\hat{\varepsilon}}$ associated to the parameter $\hat{\varepsilon} = 2.6$.

The term $\mathbb{E}(|r^\varepsilon - \hat{r}_{n,p}^\varepsilon|^2)$ can then be interpreted as “variance term”, while $|r^\varepsilon - r^0|$ is referred to as a bias term. Ideally, one would like to select a value of ε which gives a good tradeoff between these two terms whose values change in opposite directions as ε varies. However, the decay of the bias term (as $\varepsilon \rightarrow 0$) is typically unknown and only an estimation of the variance term is available. In [11], it is thus proposed to use the upper bound (18) to automatically calibrate the parameter $\varepsilon > 0$ by following Goldenshluger-Lepski (GL) principle [39] as formulated in [46] for standard kernel density estimation. In [11], an adaptation of the GL’s principle is used to estimate a bias-variance trade-off functional that is minimized over a finite set of regularization parameters to obtain an automatic selection of ε for Sinkhorn barycenters. The method consists in comparing estimators pairwise (for a given range of regularization parameters) with respect to a loss function, and we refer to [11] for further details. As an illustration of the usefulness of the GL’s principle in statistical optimal transport, we report in Figure 6 numerical experiments from [11] on the computation of a Sinkhorn barycenter (with a data-driven choice of ε) in dimension $d = 2$ for the flow cytometry dataset displayed in Figure 2(b). This regularized barycenter clearly corrects some mis-alignment issues in these measurements. We also display in Figure 6(a) the Euclidean mean of this dataset (after a preliminary kernel smoothing step). The support of this Euclidean mean is more spread out due to the presence of a strong variance in spatial location in the data from one individual to another.

4. GEOMETRY OF THE WASSERSTEIN SPACE AND GEODESIC PCA

Principal Component Analysis (PCA) is certainly the most widely used approach to reduce the dimension of datasets whose elements are high-dimensional vectors or matrices. When such elements can be modeled as probability measures, we explain how using the Wasserstein metric to compute their modes of variation accordingly around their barycenter in the Wasserstein space. We briefly describe the setting where the data

can be modeled as histograms supported on the real line, and we refer to extensions in the case of probability measures supported on \mathbb{R}^d for $d \geq 2$.

4.1. Geometry of the Wasserstein space

The space $\mathcal{P}_2(\Omega)$ endowed with the 2-Wasserstein distance is not a Hilbert space. As a consequence, standard PCA in a linear space based on the diagonalisation of a covariance matrix cannot be applied directly to compute principal modes of variations with respect to the Wasserstein metric. Nevertheless, a meaningful notion of PCA [14, 20, 58] can still be defined by relying on the pseudo-Riemannian structure of the Wasserstein space which was extensively studied in [5].

This geometric structure allows to define a notion of geodesic PCA in the Wasserstein space which shares similarities with analogs of PCA for data belonging to a Riemannian manifolds. Principal Geodesic Analysis (PGA) [33, 61] is a geometric method for statistical analysis that extends the notion of classical PCA in Hilbert spaces [26, 53] to curved Riemannian manifolds by replacing Euclidean concepts of vector means, lines and orthogonality by the more general notions of Fréchet mean, geodesics in Riemannian manifolds and orthogonality in tangent spaces. For $d = 1$, the geometric structure of the Wasserstein space is relatively easy to describe, and we refer to [14, 20] for a detailed presentation. The setting $d \geq 2$ is more involved. Some elements on the geometry of the Wasserstein space can be found in Section 4.4 in [50], and an extensive study can be found in [5].

4.2. Geodesic PCA in the one-dimensional case

Following these principles of statistical analysis in manifolds, a framework for geodesic PCA (GPCA) of probability measures supported on a interval $\Omega \subset \mathbb{R}$ has been introduced in [14]. GPCA is defined as the problem of estimating a principal geodesic subspace (of a given dimension) which maximizes the variance of the projection of the data to that subspace whose base point is the Wasserstein barycenter of the observations. In [14], a detailed study is proposed on the existence and consistency of GPCA. It is also shown that GPCA in the Wasserstein space is equivalent to map the data in the tangent space at the empirical Wasserstein barycenter (assumed to be a.c.) and then to perform a PCA that is constrained to lie in a convex and closed subset of functions. Mapping the data to this tangent space is not difficult in the one-dimensional case as it amounts to computing a set of optimal maps between the data and their Wasserstein barycenter, for which a closed form is available using their quantile functions. In particular, optimal maps belong to the set of non-decreasing functions. This convex constrained PCA can also be understood from the quantile formula (2) for the Wasserstein metric when $d = 1$. Indeed, GPCA for probability measures supported on $\Omega \subset \mathbb{R}$ corresponds to a PCA of quantile functions that is constrained to lie in the convex subset of non-decreasing function in $\mathbb{L}_2([0, 1])$. Nevertheless, such a convex PCA leads to the minimization of a non-convex and non-differentiable objective function. This is a delicate optimisation problem (even for $d = 1$), and numerical methods to compute such a PCA have been proposed in [20].

To shed some lights on the benefits of GPCA for data analysis, we report numerical experiments on synthetic data in Figure 7 and for the dataset on children's first name at birth in Figure 8. When the size or locations of significant bins in observed histograms significantly vary from one histogram to another, standard PCA with respect to the Euclidean metric (in the space of densities) fail in recovering geometric variations in location. To the contrary, GPCA in the Wasserstein space is able to capture meaningful sources of variations in a dataset by separating variability in location and scaling in the observed histograms along the first and second geodesic principal components.

4.3. Geodesic PCA in dimension $d \geq 2$

The study of Geodesic PCA of histograms in the Wasserstein space is much more involved in dimension $d \geq 2$, both on the theoretical and numerical aspects. Indeed, for $d \geq 2$, one could always define GPCA as the construction of a principal geodesic subspace maximizing the variance of the projection of the data to that subspace. However, a rigorous analysis of such a construction in dimension $d \geq 2$ has not been proposed so

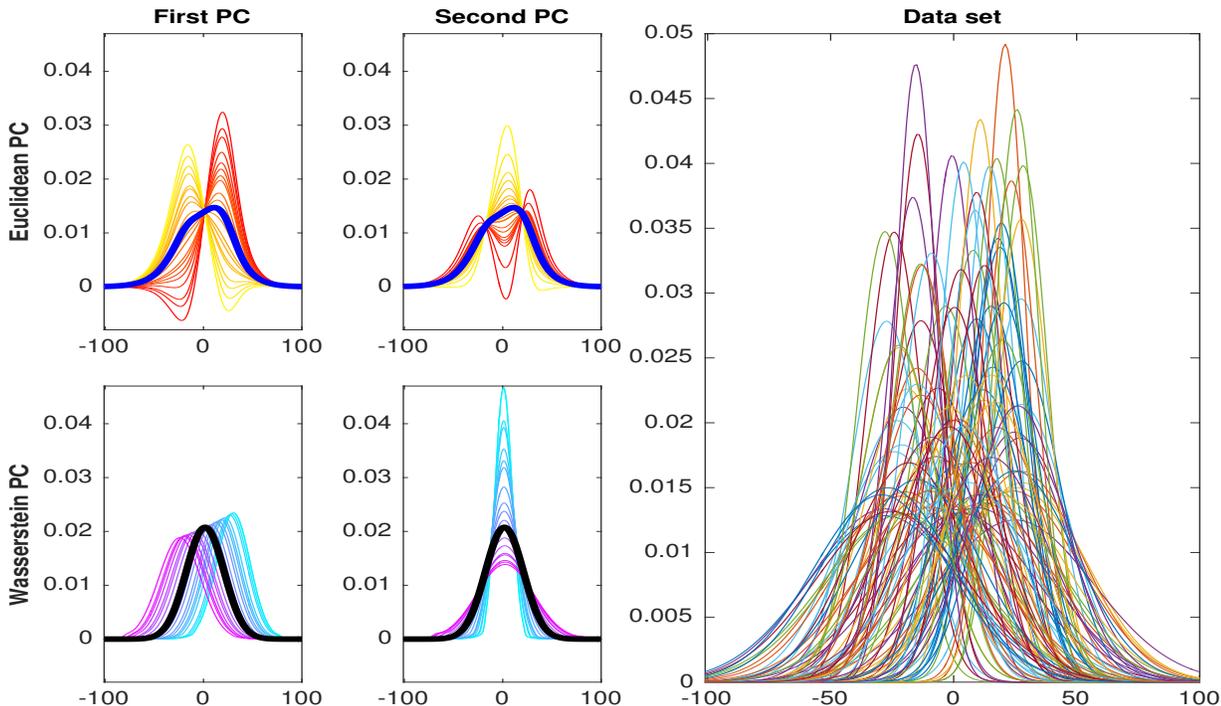


FIGURE 7. Synthetic data set: Gaussian probability densities with random mean and variance. **Euclidean PC**: standard PCA with respect to the Euclidean metric between densities - first and second principal components (PC) of the data around their Euclidean mean (blue curve). The yellow to red curves encode the progression of the first two main directions of variation in the Hilbert space of square integrable functions. **Wasserstein PC**: geodesic PCA in the Wasserstein space - first and second principal components (PC) of the data around their Wasserstein barycenter (black curve). The curves from light blue to violet encode the progression of the pdf of the first two main geodesics directions of variation in $\mathcal{P}_2(\Omega)$.

far. There also exist two main differences with the one-dimensional case. First, GPCA cannot be equivalent to map the data in the tangent space at their barycenter (assumed to be a.c.) and then to perform a standard PCA in that linear space as the Wasserstein space is a curved manifold for $d \geq 2$. Secondly, optimal maps sending a set of histograms to such a tangent space are more difficult to compute for $d \geq 2$. Moreover, it is well known [18, 21] that such maps are equal to the gradient of convex functions. This implies that an adaptation of existing algorithms for $d = 1$ leads to optimisation problems with more difficult constraints to handle than a PCA constrained to lie in the convex set of non-decreasing functions.

Nevertheless, numerical approximation to the construction of GPCA in dimension $d = 2$ have been proposed in [20, 58]. As an illustrative example, we report numerical experiments from [20] on the analysis of data from the MNIST database [48] which contains grayscale images of handwritten digits. All grayscale images in this dataset are of size 28×28 pixels. We normalize them so that the sum of pixel grayscale values sum to one. In this manner, the images are histograms supported on a 2D grid of size 28×28 . The ground metric for the Wasserstein distance is then the squared Euclidean distance between the locations of bins in this 2D grid. In Figure 9, we display the first principal components computed from 1000 images of each digit using the algorithmic approach proposed in [20]. It can be seen that GPCA captures very well the main source of variability of each digit such as changes in the shape of the '0' or the presence or absence of the lower loop of the '2'.

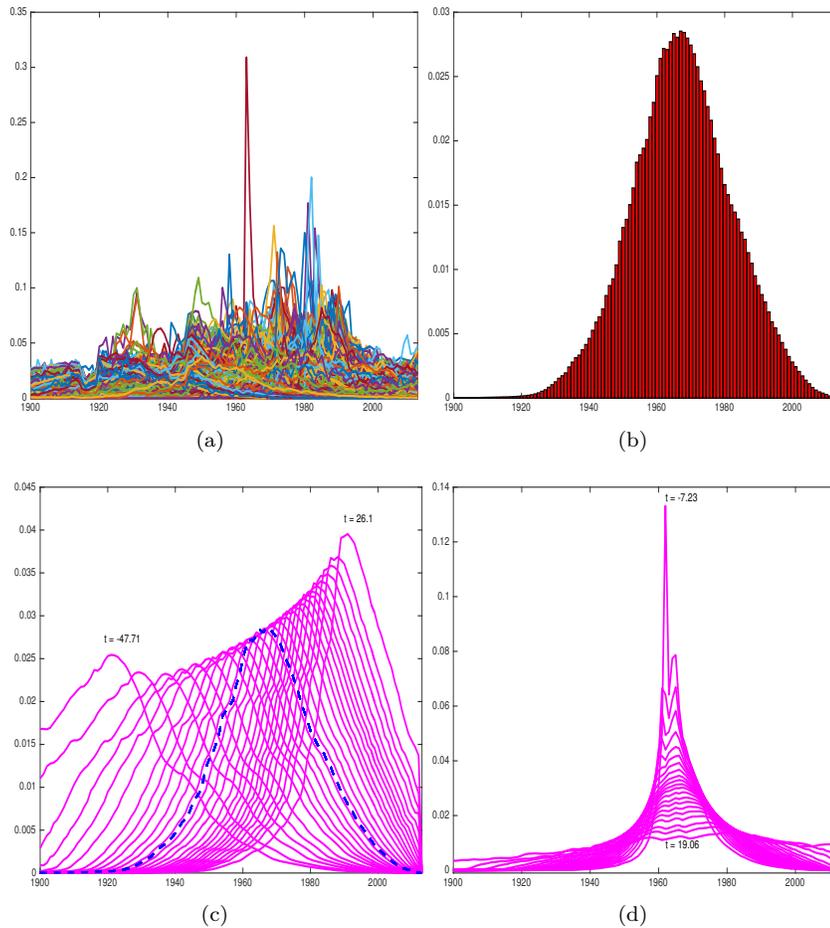


FIGURE 8. Geodesic PCA of the data on children’s first name at birth. (a) The whole dataset made of $n = 1060$ histograms. (b) Density of the Wasserstein barycenter. (c) First direction and (d) second direction of geodesic variation in $\mathcal{P}_2(\Omega)$ (at the level of pdf).

REFERENCES

- [1] A.-C. A., T. LE GOUIC, V. SPOKOINY, AND Q. PARIS, *Convergence rates for empirical barycenters in metric spaces: curvature, convexity and extendible geodesics*, arXiv preprint arXiv:1806.02740, (2018).
- [2] M. AGUEH AND G. CARLIER, *Barycenters in the Wasserstein space*, SIAM Journal on Mathematical Analysis, 43 (2011), pp. 904–924.
- [3] ———, *Vers un théorème de la limite centrale dans l’espace de Wasserstein?*, Comptes Rendus Mathématique, 355 (2017), pp. 812–818.
- [4] J. ALTSCHULER, J. WEED, AND P. RIGOLLET, *Near-linear time approximation algorithms for optimal transport via sinkhorn iteration*, in Advances in Neural Information Processing Systems 30, I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, eds., Curran Associates, Inc., 2017, pp. 1964–1974.
- [5] L. AMBROSIO, N. GIGLI, AND G. SAVARÉ, *Gradient flows: in metric spaces and in the space of probability measures*, Springer Science & Business Media, 2008.
- [6] M. ARJOVSKY, S. CHINTALA, AND L. BOTTOU, *Wasserstein generative adversarial networks*, in Proceedings of the 34th International Conference on Machine Learning, D. Precup and Y. W. Teh, eds., vol. 70 of Proceedings of Machine Learning Research, International Convention Centre, Sydney, Australia, 06–11 Aug 2017, PMLR, pp. 214–223.

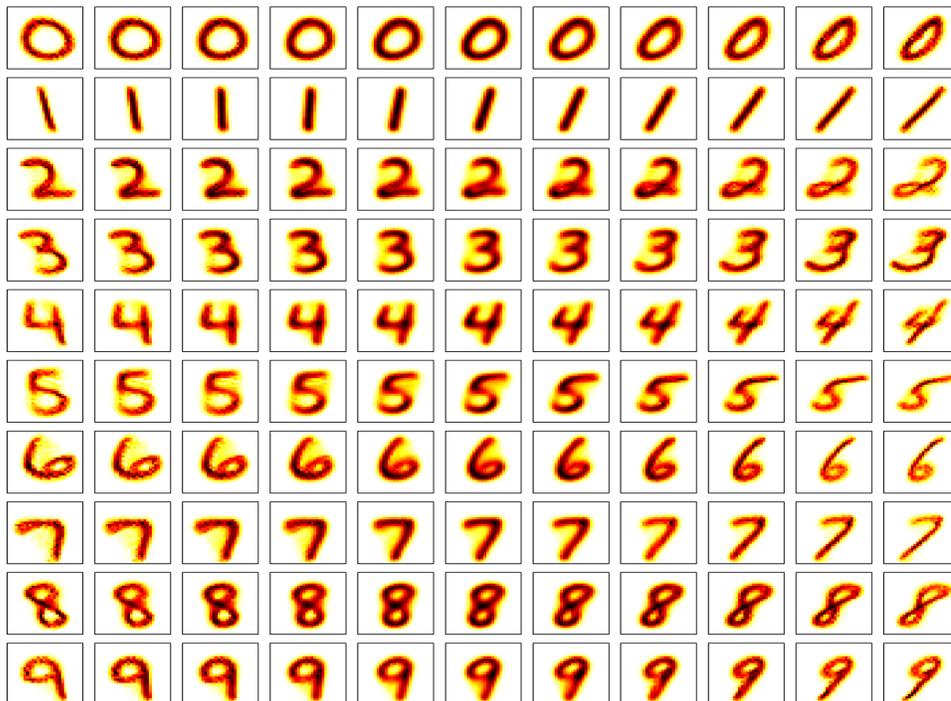


FIGURE 9. First principal geodesics for 1000 images of each digit from the MNIST database [48]. The middle column represents the pdf of the Wasserstein barycenter of each digit (computed from 1000 images whose pixels are normalized to form a set of histograms). Each line allows to visualize the first geodesic of variation of each digit around their barycenter in the sense of optimal transport. The cost is quadratic and it corresponds to the squared distance between pixel locations.

- [7] J.-D. BENAMOU, G. CARLIER, M. CUTURI, L. NENNA, AND G. PEYRÉ, *Iterative Bregman projections for regularized transportation problems*, SIAM Journal on Scientific Computing, 37 (2015), pp. A1111–A1138.
- [8] B. BERCU AND BIGOT, *Asymptotic distribution and convergence rates of stochastic algorithms for entropic optimal transportation between probability measures*, arXiv preprint arXiv:1812.09150, (2019).
- [9] Q. BERTHET AND J. WEED, *Estimation of smooth densities in wasserstein distance*, arXiv preprint arXiv:1902.01778, (2019).
- [10] J. BIGOT, E. CAZELLES, AND N. PAPADAKIS, *Central limit theorems for sinkhorn divergence between probability distributions on finite spaces and statistical applications*, arXiv preprint arXiv:1711.08947, (2017).
- [11] ———, *Data-driven regularization of wasserstein barycenters with an application to multivariate density registration*, arXiv preprint arXiv:1804.08962, (2018).
- [12] ———, *Penalization of barycenters in the Wasserstein space*, SIAM J. Math. Analysis, To be published (2019).
- [13] J. BIGOT, R. GOUET, T. KLEIN, AND A. LOPEZ, *Upper and lower risk bounds for estimating the Wasserstein barycenter of random measures on the real line*, Electronic Journal of Statistics, (2018).
- [14] J. BIGOT, R. GOUET, T. KLEIN, A. LÓPEZ, ET AL., *Geodesic PCA in the Wasserstein space by convex PCA*, Annales de l’Institut Henri Poincaré, Probabilités et Statistiques, 53 (2017), pp. 1–26.
- [15] BIGOT, JÉRÉMIE AND KLEIN, THIERRY, *Characterization of barycenters in the wasserstein space by averaging optimal transport maps*, ESAIM: PS, 22 (2018), pp. 35–57.
- [16] S. BOBKOV AND M. LEDOUX, *One-dimensional empirical measures, order statistics and Kantorovich transport distances*, To appear in the Memoirs of the American Mathematical Society, 2014.
- [17] N. BONNEEL, J. RABIN, G. PEYRÉ, AND H. PFISTER, *Sliced and radon Wasserstein barycenters of measures*, Journal of Mathematical Imaging and Vision, 51 (2015), pp. 22–45.
- [18] Y. BRENIER, *Polar factorization and monotone rearrangement of vector-valued functions*, Comm. Pure Appl. Math., 44 (1991), pp. 375–417.

- [19] G. CARLIER, V. DUVAL, G. PEYRÉ, AND B. SCHMITZER, *Convergence of entropic schemes for optimal transport and gradient flows*, SIAM J. Math. Analysis, 49 (2017), pp. 1385–1418.
- [20] E. CAZELLES, V. SEGUY, J. BIGOT, M. CUTURI, AND N. PAPADAKIS, *Log-PCA versus geodesic PCA of histograms in the Wasserstein space*, SIAM Journal on Scientific Computing, to be published (2018).
- [21] J. A. CUESTA AND C. MATRAN, *Notes on the Wasserstein metric in Hilbert spaces*, Ann. Probab., 17 (1989), pp. 1264–1276.
- [22] M. CUTURI, *Sinkhorn distances: Lightspeed computation of optimal transport*, in Advances in Neural Information Processing Systems 26, C. J. C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Q. Weinberger, eds., Curran Associates, Inc., 2013, pp. 2292–2300.
- [23] M. CUTURI AND A. DOUCET, *Fast computation of Wasserstein barycenters*, in International Conference on Machine Learning 2014, PMLR W&CP, vol. 32, 2014, pp. 685–693.
- [24] M. CUTURI AND G. PEYRÉ, *A smoothed dual approach for variational Wasserstein problems*, SIAM Journal on Imaging Sciences, 9 (2016), pp. 320–343.
- [25] M. CUTURI AND G. PEYRÉ, *Computational Optimal Transport*, Book available at <https://optimaltransport.github.io/book/>, 2017.
- [26] J. DAUXOIS, A. POUSSE, AND Y. ROMAIN, *Asymptotic theory for the principal component analysis of a vector random function: some applications to statistical inference*, J. Multivariate Anal., 12 (1982), pp. 136–154.
- [27] J. DEDECKER AND F. MERLEVÈDE, *Behavior of the empirical wasserstein distance in \mathbb{R}^d under moment conditions*, Electron. J. Probab., 24 (2019), p. 32 pp.
- [28] E. DEL BARRIO, J. A. CUESTA-ALBERTOS, C. MATRÁN, AND A. MAYO-ÍSCAR, *Robust clustering tools based on optimal transportation*, Statistics and Computing, 29 (2019), pp. 139–160.
- [29] E. DEL BARRIO, P. GORDALIZA, H. LESCORNEL, AND J.-M. LOUBES, *Central limit theorem and bootstrap procedure for Wasserstein variations with an application to structural relationships between distributions*, Journal of Multivariate Analysis, 169 (2019), pp. 341 – 362.
- [30] S. FERRADANS, N. PAPADAKIS, G. PEYRÉ, AND J.-F. AUJOL, *Regularized discrete optimal transport*, SIAM Journal on Imaging Sciences, 7 (2014), pp. 1853–1882.
- [31] J. FEYDY, B. CHARLIER, F.-X. VIALARD, AND G. PEYRÉ, *Optimal transport for diffeomorphic registration*, in Medical Image Computing and Computer Assisted Intervention - MICCAI 2017, Cham, 2017, Springer International Publishing, pp. 291–299.
- [32] R. FLAMARY, M. CUTURI, N. COURTY, AND A. RAKOTOMAMONJY, *Wasserstein discriminant analysis*, Machine Learning, 107 (2018), pp. 1923–1945.
- [33] P. T. FLETCHER, C. LU, S. M. PIZER, AND S. JOSHI, *Principal geodesic analysis for the study of nonlinear statistics of shape.*, IEEE Transactions on Medical Imaging, 23 (2004), pp. 995–1005.
- [34] N. FOURNIER AND A. GUILLIN, *On the rate of convergence in Wasserstein distance of the empirical measure*, Probability Theory and Related Fields, 162 (2015), pp. 707–738.
- [35] M. FRÉCHET, *Les éléments aléatoires de nature quelconque dans un espace distancié*, Annales de l’Institut H.Poincaré, Sect. B, Probabilités et Statistiques, 10 (1948), pp. 235–310.
- [36] C. FROGNER, C. ZHANG, H. MOBAHI, M. ARAYA-POLO, AND T. POGGIO, *Learning with a wasserstein loss*, in Proceedings of the 28th International Conference on Neural Information Processing Systems - Volume 2, NIPS’15, Cambridge, MA, USA, 2015, MIT Press, pp. 2053–2061.
- [37] A. GENEVAY, M. CUTURI, G. PEYRÉ, AND F. BACH, *Stochastic optimization for large-scale optimal transport*, in Advances in Neural Information Processing Systems 29, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, eds., Curran Associates, Inc., 2016, pp. 3440–3448.
- [38] A. GENEVAY, G. PEYRE, AND M. CUTURI, *Learning generative models with sinkhorn divergences*, in Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics, A. Storkey and F. Perez-Cruz, eds., vol. 84 of Proceedings of Machine Learning Research, Playa Blanca, Lanzarote, Canary Islands, 09–11 Apr 2018, PMLR, pp. 1608–1617.
- [39] A. GOLDENSHLUGER, O. LEPSKI, ET AL., *Universal pointwise selection rule in multivariate function estimation*, Bernoulli, 14 (2008), pp. 1150–1190.
- [40] P. GORDALIZA, E. D. BARRIO, G. FABRICE, AND L. JEAN-MICHEL, *Obtaining fairness using optimal transport theory*, in Proceedings of the 36th International Conference on Machine Learning, K. Chaudhuri and R. Salakhutdinov, eds., vol. 97 of Proceedings of Machine Learning Research, Long Beach, California, USA, 09–15 Jun 2019, PMLR, pp. 2357–2365.
- [41] A. GRAMFORT, G. PEYRÉ, AND M. CUTURI, *Fast optimal transport averaging of neuroimaging data*, in International Conference on Information Processing in Medical Imaging, Springer, 2015, pp. 261–272.
- [42] F. HAHNE, A. KHODABAKHSHI, A. BASHASHATI, C.-J. WONG, R. GASCOYNE, A. WENG, V. SEYFERT-MARGOLIS, K. BOURCIER, A. ASARE, T. LUMLEY, R. GENTLEMAN, AND R. BRINKMAN, *Per-channel basis normalization methods for flow cytometry data*, Cytometry Part A, 77 (2010), pp. 121–131.
- [43] Y.-H. KIM AND B. PASS, *Wasserstein barycenters over Riemannian manifolds*, Advances in Mathematics, 307 (2017), pp. 640–683.
- [44] M. KLATT, C. TAMELING, AND A. MUNK, *Empirical regularized optimal transport: Statistical theory and applications*, Preprint, arXiv:1810.09880, (2018).

- [45] A. KROSHNIN, V. SPOKOINY, AND A. SUVORIKOVA, *Statistical inference for Bures-Wasserstein*, arXiv preprint arXiv:1901.00226, (2019).
- [46] C. LACOUR AND P. MASSART, *Minimal penalty for goldenshluger-lepski method*, Stochastic Processes and their Applications, 126 (2016), pp. 3774–3789.
- [47] T. LE GOUIC AND J.-M. LOUBES, *Existence and Consistency of Wasserstein Barycenters*, Probability Theory and Related Fields, 168 (2016), pp. 901–917.
- [48] Y. LECUN, *The mnist database of handwritten digits*, <http://yann.lecun.com/exdb/mnist/>, (1998).
- [49] V. M. PANARETOS AND Y. ZEMEL, *Amplitude and phase variation of point processes*, Annals of Statistics, 44 (2016), pp. 771–812.
- [50] ———, *Statistical aspects of wasserstein distances.*, Annual Reviews of Statistics and its Applications, 6 (2018), pp. 405–431.
- [51] A. PETERSEN AND H.-G. MÜLLER, *Wasserstein covariance for multiple random densities*, arXiv preprint arXiv:1812.07694, (2019).
- [52] J. RABIN AND N. PAPADAKIS, *Convex color image segmentation with optimal transport distances*, in International Conference on Scale Space and Variational Methods in Computer Vision, Springer, 2015, pp. 256–269.
- [53] J. O. RAMSAY AND B. W. SILVERMAN, *Functional data analysis*, Springer Series in Statistics, Springer, New York, second ed., 2005.
- [54] P. RIGOLLET AND J. WEED, *Entropic optimal transport is maximum-likelihood deconvolution*, Comptes Rendus Mathematique, 356 (2018), pp. 1228 – 1235.
- [55] A. ROLET, M. CUTURI, AND G. PEYRÉ, *Fast dictionary learning with a smoothed Wasserstein loss*, in Proc. International Conference on Artificial Intelligence and Statistics (AISTATS), 2016.
- [56] G. SCHIEBINGER, J. SHU, M. TABAKA, B. CLEARY, V. SUBRAMANIAN, A. SOLOMON, J. GOULD, S. LIU, S. LIN, P. BERUBE, L. LEE, J. CHEN, J. BRUMBAUGH, P. RIGOLLET, K. HOCHEDLINGER, R. JAENISCH, A. REGEV, AND E. S. LANDER, *Optimal-transport analysis of single-cell gene expression identifies developmental trajectories in reprogramming*, Cell, 176 (2019), p. 1517.
- [57] M. SCHMITZ, M. HEITZ, N. BONNEEL, F. NGOLÈ, D. COEURJOLLY, M. CUTURI, G. PEYRÉ, AND J. STARCK, *Wasserstein dictionary learning: Optimal transport-based unsupervised nonlinear dictionary learning*, SIAM Journal on Imaging Sciences, 11 (2018), pp. 643–678.
- [58] V. SEGUY AND M. CUTURI, *Principal geodesic analysis for probability measures under the optimal transport metric*, in Advances in Neural Information Processing Systems 28, C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, eds., Curran Associates, Inc., 2015, pp. 3294–3302.
- [59] J. SOLOMON, F. DE GOES, G. PEYRÉ, M. CUTURI, A. BUTSCHER, A. NGUYEN, T. DU, AND L. GUIBAS, *Convolutional wasserstein distances: Efficient optimal transportation on geometric domains*, ACM Trans. Graph., 34 (2015), pp. 66:1–66:11.
- [60] J. SOLOMON, G. PEYRÉ, V. G. KIM, AND S. SRA, *Entropic metric alignment for correspondence problems*, ACM Trans. Graph., 35 (2016), pp. 72:1–72:13.
- [61] J. SOMMER, F. LAUZE, S. HAUBERG, AND M. NIELSEN, *Manifold valued statistics, exact principal geodesic analysis and the effect of linear approximations*, in Computer Vision ECCV 2010, K. Daniilidis, P. Maragos, and N. Paragios, eds., vol. 6316 of Lecture Notes in Computer Science, Springer Berlin Heidelberg, 2010, pp. 43–56.
- [62] M. SOMMERFELD AND A. MUNK, *Inference for empirical Wasserstein distances on finite spaces*, Journal of the Royal Statistical Society: Series B (Statistical Methodology), (2016).
- [63] C. VILLANI, *Topics in optimal transportation*, vol. 58 of Graduate Studies in Mathematics, American Mathematical Society, 2003.
- [64] J. WEED AND F. BACH, *Sharp asymptotic and finite-sample rates of convergence of empirical measures in wasserstein distance*, Bernoulli, To be published (2019).
- [65] A. G. WILSON, *The use of entropy maximising models, in the theory of trip distribution, mode split and route split*, Journal of Transport Economics and Policy, (1969), pp. 108–126.
- [66] Y. ZEMEL AND V. M. PANARETOS, *Fréchet means and procrustes analysis in wasserstein space*, Bernoulli, 25 (2019), pp. 932–976.
- [67] P. C. ÁLVAREZ ESTEBAN, E. DEL BARRIO, J. A. CUESTA-ALBERTOS, AND C. MATRÁN, *Wide consensus aggregation in the wasserstein space. application to location-scatter families*, Bernoulli, 24 (2018), pp. 3147–3179.