

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 ON THE WASSERSTEIN GEODESIC PRINCIPAL COMPO- NENT ANALYSIS OF PROBABILITY MEASURES

005 **Anonymous authors**

006 Paper under double-blind review

ABSTRACT

011 This paper focuses on Geodesic Principal Component Analysis (GPCA) on a
012 collection of probability distributions using the Otto-Wasserstein geometry. The
013 goal is to identify geodesic curves in the space of probability measures that best
014 capture the modes of variation of the underlying dataset. We first address the case
015 of a collection of Gaussian distributions, and show how to lift the computations
016 to the space of invertible linear maps. For the more general setting of absolutely
017 continuous probability measures, we leverage a novel approach to parameterizing
018 geodesics in Wasserstein space with neural networks. Finally, we compare to
019 classical tangent PCA through various examples and provide illustrations on real-
020 world datasets.

1 INTRODUCTION

024 In this paper, we are interested in computing the main modes of variation of a dataset of absolutely
025 continuous (a.c.) probability measures supported in \mathbb{R}^d . For data points living in an arbitrary Hilbert
026 space, the classical approach defined by Principal Component Analysis (PCA) consists in finding a
027 sequence of nested affine subspaces on which the projected data retain a maximal part of the variance
028 of the original dataset, or equivalently, yield best lower-dimensional approximations. When dealing
029 with a set of a.c. probability distributions, a natural choice is to identify the probability measures
030 with their probability density functions and to perform PCA on these using the L^2 Hilbert metric.
031 Unfortunately, as highlighted in Cazelles et al. (2018), the components computed in this manner fail
032 to capture the intrinsic structure of the dataset: **the projections onto the components most likely result**
033 **in non-positive and un-normalized functions**. Using the Wasserstein metric W_2 instead has proven to
034 overcome these limitations by taking into account the geometry of the space of distributions.

035 The Wasserstein metric endows the space of probability distributions with a *Riemannian-like* structure,
036 framing the problem as PCA on a (positively) curved Riemannian manifold. A first approach to
037 solve this task, known as *Tangent PCA* (TPCA), consists in embedding the data into the tangent
038 space at a reference point, and applying classical PCA in this flat space, as in Fletcher et al. (2003).
039 TPCA is computationally advantageous but can generically induce distortion in the embedded data,
040 depending on the curvature of the manifold at the reference point and the dispersion of the data. A
041 more geometrically coherent approach is *Geodesic PCA* (GPCA) proposed for Riemannian manifolds
042 in Huckemann et al. (2010); Huckemann & Ziezold (2006), where principal modes of variations are
043 geodesics that minimize the variance of the projection residuals. Following this approach, the first
044 geodesic component of a set of probability measures ν_1, \dots, ν_n in the Wasserstein space solves

$$045 \inf_{t \mapsto \mu(t) \text{ geodesic}} \sum_{i=1}^n \inf_{t_i} W_2^2(\mu(t_i), \nu_i). \quad (1)$$

046 Interestingly, unlike in the Hilbert setting, this criterion is *not* equivalent to maximizing the variance
047 of the projections, which leads to a different notion of PCA on Riemannian manifolds (see Sommer
048 et al. (2010; 2014)).

049 **Related works** TPCA in the Wasserstein space was considered by Wang et al. (2013) through the
050 use of the linearized Wasserstein distance. In a similar approach, Boissard et al. (2015) restrict to
051 distributions that can be obtained by deforming a single template measure. For one-dimensional
052 probability measures, Bigot et al. (2017) have shown that GPCA and its linearized approximation

TPCA coincide, as the embedding into a tangent space is then an isometry when constrained to a convex set. An algorithm in this case has been proposed in Cazelles et al. (2018), with an approximate extension in dimension 2. For higher-dimensional measures, Seguy & Cuturi (2015) solve an approximate version of GPCA, replacing geodesics by generalized geodesics as defined in Ambrosio et al. (2008). Despite all this, a method to solve the *exact* GPCA problem described in equation 1 is still missing for \mathbb{R}^d -valued probability measures. The goal of this paper is to fill this gap.

Main contributions In this paper, we introduce two algorithms to solve the exact GPCA problem in the Wasserstein space of (1) centered Gaussian distributions and (2) a.c. probability measures on \mathbb{R}^d . Our methods are exact in the sense that they do not rely on a linearization of the Wasserstein space, and the components are true geodesics that minimize the cost in equation 1. In the Gaussian case, we leverage the Bures-Wasserstein geometry to lift the computations to the flat space of invertible matrices. We show an example where GPCA and TPCA differ significantly, and relate this effect to curvature. In the general case of a.c. probability distributions, we lift the probability distributions to the space of (non necessarily optimal) maps that pushforward a given reference measure, as described by Otto (2001). This approach is independent of the chosen reference measure and yields a convenient way to parametrize geodesic components and define orthogonality with respect to the Wasserstein metric. In practice, we parametrize geodesic components using multilayer perceptrons (MLPs), trained to minimize the cost in equation 1. We show illustrations on images and 3D point clouds. Along the way, we prove that for univariate Gaussian distributions, GPCA yields the same results whether it is performed in the space of a.c. distributions or restricted to the Gaussian submanifold.

Organization of the paper In Section 2, we present the Wasserstein metric and its restriction to Gaussian distributions, as well as the related Bures-Wasserstein and Otto-Wasserstein geometries. We present GPCA for centered Gaussian distributions in Section 3, and the general case of a.c. probability measures in Section 4. Experiments are presented in Section 5, and the paper ends with a discussion in Section 6. All the proofs and additional experiments are deferred in the appendices.

2 BACKGROUND

The Wasserstein distance Optimal transport is about finding the optimal way to transport mass from one distribution μ on \mathbb{R}^d to another ν with respect to a ground cost, say the Euclidean squared distance. The total transport cost defines the *Wasserstein distance* W_2 between a.c. measures μ, ν with moment of order 2, whose Monge (1781) formulation is given by

$$W_2^2(\mu, \nu) = \int_{\mathbb{R}^d} \|x - T_\mu^\nu(x)\|^2 d\mu(x), \quad (2)$$

and where the map T_μ^ν is the μ -a.s. unique gradient of a convex function verifying $T_\mu^\nu \# \mu = \nu$, as proven by Brenier (1991). When the distributions μ and ν are centered (non-degenerate) Gaussian distributions, they can be identified with their covariance matrices Σ_μ, Σ_ν and equation 2 is referred to as the *Bures-Wasserstein distance* BW_2 on the manifold S_d^{++} of symmetric positive definite (SPD) matrices (see e.g. Modin (2017); Bhatia et al. (2019)):

$$BW_2^2(\Sigma_\mu, \Sigma_\nu) = \text{tr} \left[\Sigma_\mu + \Sigma_\nu - 2(\Sigma_\mu^{1/2} \Sigma_\nu \Sigma_\mu^{1/2})^{1/2} \right]. \quad (3)$$

Both distances can be induced by a Riemannian metric on their respective manifolds, i.e. the space of a.c. distributions and S_d^{++} , as we will see in the following. For more details, see Appendix B.

Bures-Wasserstein geometry of centered Gaussian distributions The set of centered non-degenerate Gaussian distributions on \mathbb{R}^d is identified with the manifold S_d^{++} of SPD matrices. The Riemannian geometry of the Bures-Wasserstein metric in equation 3 can be described by considering S_d^{++} as the quotient of the manifold GL_d of invertible matrices by the right action of the orthogonal group O_d . In this geometry, GL_d is decomposed into equivalence classes called *fibers*. The fiber over $\Sigma \in S_d^{++}$ is defined to be the pre-image of Σ under the projection

$$\pi: A \in GL_d \mapsto AA^\top \in S_d^{++}, \quad (4)$$

and can be obtained as the result of the action of O_d on a representative, e.g. $\Sigma^{1/2}$ the only SPD square root of Σ : $\pi^{-1}(\Sigma) = \{A \in GL_d, AA^\top = \Sigma\} = \Sigma^{1/2}O_d$.

108 Tangent vectors to GL_d are said to be *horizontal* if they
 109 are orthogonal to the fibers with respect to the Frobenius
 110 metric, i.e. if they belong to the space

$$111 \quad \text{Hor}_A := \{X \in \mathbb{R}^{d \times d}, X^\top A - A^\top X = 0\}, \quad (5)$$

113 for a given point $A \in GL_d$. Then the projection π in
 114 equation 4 defines an isometry between the horizontal
 115 subspace Hor_A equipped with the Frobenius inner product
 116 $\langle X, Y \rangle := \text{tr}(XY^\top)$, and S_d^{++} equipped with a
 117 Riemannian metric that induces the Bures-Wasserstein
 118 distance (equation 3) as the geodesic distance. In partic-
 119 ular, this means that moving horizontally along straight
 120 lines in the top space GL_d is equivalent to moving along
 121 geodesics in the bottom space S_d^{++} (see Figure 1), as
 122 recalled in the following proposition.

123 **Proposition 1** (Takatsu (2011); Malagò et al. (2018); Bhatia et al. (2019)). *Any geodesic $t \mapsto \Sigma(t)$ in*
 124 S_d^{++} *for the Bures-Wasserstein metric in equation 3 is the π -projection of a horizontal line segment*
 125 *in GL_d , that is*

$$127 \quad \Sigma(t) = \pi(A + tX) = (A + tX)(A + tX)^\top, \quad A \in GL_d, X \in \text{Hor}_A, \quad (6)$$

128 where t is defined in a certain time interval (t_{\min}, t_{\max}) . Also, the Bures-Wasserstein distance between
 129 two covariance matrices $\Sigma_1, \Sigma_2 \in S_d^{++}$ is given by the minimal distance between their fibers

$$130 \quad BW_2(\Sigma_1, \Sigma_2) = \inf_{Q_1, Q_2 \in O_d} \|\Sigma_1^{1/2} Q_1 - \Sigma_2^{1/2} Q_2\| = \inf_{Q \in SO_d} \|\Sigma_1^{1/2} - \Sigma_2^{1/2} Q\|, \quad (7)$$

132 where $\|\cdot\|$ is the Frobenius norm and SO_d is the special orthogonal group.

133 It is essential to note that the geodesic equation 6 cannot be extended for all time $t \in \mathbb{R}$ (the only
 134 geodesic lines are those obtained by translation (Kloeckner, 2010, Proposition 3.6)). Therefore,
 135 equation 6 is only defined on a time interval (t_{\min}, t_{\max}) that depends on the eigenvalues of XA^{-1}
 136 (see Appendix B.3). More details on the Bures-Wasserstein geometry can be found in Appendix B.2.

138 **Otto-Wasserstein geometry of a.c. probability measures** The Riemannian structure described for
 139 Gaussian distributions is a special case of Otto (2001)’s more general construction : the bottom space
 140 becomes the space $\text{Prob}(\Omega)$ of a.c. distributions supported in a compact set $\Omega \subset \mathbb{R}^d$ while the top
 141 space is the space of diffeomorphisms $\text{Diff}(\Omega)$ endowed with the L^2 metric with respect to a fixed
 142 reference measure ρ (see Figure 19 in Appendix B). The *fibers* of $\text{Diff}(\Omega)$ are then defined to be the
 143 pre-images under the projection

$$144 \quad \pi: \varphi \in \text{Diff}(\Omega) \mapsto \pi(\varphi) = \varphi\# \rho \in \text{Prob}(\Omega). \quad (8)$$

145 In this setting, *horizontal* displacements in $\text{Diff}(\Omega)$ are along vector fields that are gradients of
 146 functions. The projection π defines an isometry between the horizontal subspace equipped with the
 147 $L^2(\rho)$ -inner product and $\text{Prob}(\Omega)$ equipped with a Riemannian metric that induces the Wasserstein
 148 distance as the geodesic distance. In particular, we have the following result.

149 **Proposition 2** (Otto (2001)). *Any geodesic $t \mapsto \mu(t)$ for the Wasserstein metric given in equation 2 is*
 150 *the π -projection of a line segment in $\text{Diff}(\Omega)$ going through a diffeomorphism φ at horizontal speed*
 151 *$\nabla f \circ \varphi$ for some smooth function $f \in \mathcal{C}(\mathbb{R}^d)$. That is, for t defined in a certain interval (t_{\min}, t_{\max}) ,*

$$152 \quad \mu(t) = \pi(\varphi + t\nabla f \circ \varphi) = (\text{id} + t\nabla f)\#(\varphi\# \rho). \quad (9)$$

153 Another geodesic $\tilde{\mu}(t) = \pi(\varphi + t\nabla \tilde{f} \circ \varphi)$ is orthogonal to $\mu(t)$ at $t = 0$ for the Riemannian metric
 154 inducing the Wasserstein distance if and only if $\langle \nabla f \circ \varphi, \nabla \tilde{f} \circ \varphi \rangle_{L^2(\rho)} = 0$.

156 We emphasize that f need not be convex in equation 9, unlike in the more classical parametrization
 157 of geodesics due to McCann (1997) between two distributions μ_0 and $\mu_1 = \nabla u\# \mu_0$:

$$158 \quad \mu(t) = (\text{id} + t(\nabla u - \text{id}))\# \mu_0, \text{ with } t \in [0, 1] \text{ and } u \text{ a convex function.} \quad (10)$$

159 Note that equation 9 parametrizes geodesics provided that $\text{id} + t\nabla f$ is a diffeomorphism, and thus it
 160 is defined on a time interval that depends on the eigenvalues of the Hessian of f . On the other hand,
 161 the convexity condition on the function u in the parametrization of equation 10 ensures that time t is
 162 defined on $[0, 1]$. Both are completely equivalent (see Appendix B.3 for details).

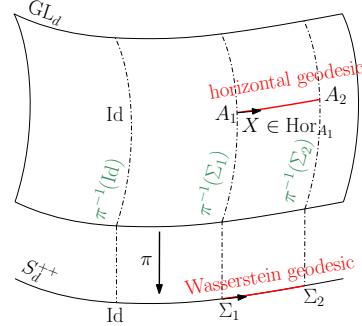


Figure 1: The Bures-Wasserstein geometry of centered Gaussian distributions, inspired by Khesin et al. (2021).

162 3 GEODESIC PCA ON CENTERED GAUSSIAN DISTRIBUTIONS

164 In this section, we consider the exact GPCA problem for the Bures-Wasserstein metric in equation 3.
 165 The data are n centered Gaussian distributions identified with their covariance matrices $\Sigma_1, \dots, \Sigma_n \in$
 166 S_d^{++} . Following Huckemann et al. (2010), we define the first component as the geodesic $t \mapsto \Sigma(t) \in$
 167 S_d^{++} that minimizes the sum of squared residuals of the BW_2 -projections of the data:
 168

$$169 \inf_{t \mapsto \Sigma(t) \text{ geodesic}} \sum_{i=1}^n \inf_{t_i} BW_2^2(\Sigma(t_i), \Sigma_i). \quad (11)$$

171 The second principal component is defined to be the geodesic that minimizes the same cost function,
 172 with the constraint of intersecting the previous component orthogonally. The subsequent principal
 173 components have the additional constraint of going through the intersection of the first two principal
 174 geodesics. This definition does not impose that the geodesic components go through the Wasserstein
 175 barycenter (see Aguech & Carlier (2011)), and in Section 5 we show an example where this is indeed
 176 not verified. This gives an observation of the phenomenon already described in Huckemann & Ziezold
 177 (2006) for spherical geometry. The proofs of this section are deferred to Appendix D.
 178

179 **Learning the geodesic components** Following Proposition 1, we lift the GPCA problem in equa-
 180 tion 11 to the total space GL_d of Otto’s fiber bundle. This has several advantages: the Bures-
 181 Wasserstein distance in the cost function of equation 11 is replaced by the Frobenius norm $\|\cdot\|$, the
 182 geodesic is replaced by a horizontal line segment, and the projection times t_i become explicit. The
 183 price to pay is an optimization over variables $(Q_i)_{i=1}^n$ in SO_d , needed to represent the covariance
 184 matrices Σ_i by invertible matrices $\Sigma_i^{1/2}Q_i$ in their respective fibers.
 185

Proposition 3. *Let $\pi: GL_d \rightarrow S_d^{++}$, $A \mapsto AA^\top$ and $(A_1, X_1, (Q_i)_{i=1}^n)$ be a solution of*

$$187 \inf F(A_1, X_1, (Q_i)_{i=1}^n) := \sum_{i=1}^n \|A_1 + p_{A_1, X_1}(t_i)X_1 - \Sigma_i^{1/2}Q_i\|^2, \quad (12)$$

189 *subject to $A_1 \in GL_d$, $X_1 \in \text{Hor}_{A_1}$, $\|X_1\|^2 = 1$, $Q_1, \dots, Q_n \in SO_d$.*

191 *Then there exist $t_{\min}, t_{\max} \in \mathbb{R}$ such that the geodesic $\Sigma: t \in [t_{\min}, t_{\max}] \mapsto \pi(A_1 + tX_1)$ in S_d^{++}
 192 minimizes equation 11.*

193 Here the t_i are projection times given by $t_i = \langle \Sigma_i^{1/2}Q_i - A_1, X_1 \rangle$, and $p_{A, X}$ is a projection operator
 194 that clips any $t \in \mathbb{R}$ onto a closed interval $[t_{\min}, t_{\max}]$ depending on A and X , such that $A + p_{A, X}(t)X$
 195 is invertible for any t in this interval (see Appendix B.3). Clipping the time parameter of the line
 196 segment is necessary to ensure it remains within GL_d and projects onto a geodesic in S_d^{++} .
 197

198 The second component is a geodesic of
 199 S_d^{++} that orthogonally intersects the first
 200 component. Lifting again the problem
 201 to GL_d , this boils down to searching for
 202 a horizontal line $t \mapsto A_2 + tX_2$ where
 203 $A_2 = (A_1 + t^*X_1)R^*$ for a rotation
 204 matrix R^* , a time $t^* \in [t_{\min}, t_{\max}]$ and a
 205 horizontal vector $X_2 \in \text{Hor}_{A_2}$ such that
 206 $\langle X_2, X_1 R^* \rangle = 0$. The equation for A_2
 207 ensures that the π -projections of the first
 208 two horizontal lines intersect, while the
 209 condition on X_2 ensures that they intersect
 210 orthogonally (since $X_1 R^*$ is horizontal at
 211 A_2 as can easily be checked). See Figure 2.

212 The second component is thus defined by $\Sigma_2(t) = \pi(A_2 + tX_2)$, found by solving:

$$213 \inf F(A_2, X_2, (Q_i)_{i=1}^n) \\ 214 \text{subject to } A_2 = (A_1 + t^*X_1)R^*, R^* \in SO_d, t^* \in [t_{\min}, t_{\max}] \\ 215 X_2 \in \text{Hor}_{A_2}, \|X_2\|^2 = 1, \langle X_2, X_1 R^* \rangle = 0, Q_1, \dots, Q_n \in SO_d. \quad (13)$$

Note that this step requires to find new rotation matrices $(Q_i)_{i=1}^n$. The first two components fix the intersection point $\pi(A_2)$ through which all other geodesic components will pass, see Figure 2. For every higher order component, we search for a velocity vector X_k that is horizontal at some point in the fiber over $\pi(A_2)$ and orthogonal to the lifts of the velocity vectors of the previous components. Details on the implementation of these components are given in Appendix D.2.

Quantifying the difference between TPCA and GPCA In the following, we quantify the distortion induced by linearization in the case of covariances matrices with same eigenvalues.

Proposition 4. *Let $\Sigma \in S_2^{++}$ with eigenvalues a^2, b^2 and $\Sigma' = P_\theta \Sigma P_\theta^\top$ where P_θ is the rotation matrix of angle θ . Then, denoting $\bar{\Sigma} = ((a+b)/2)^2 I$, we have*

$$\frac{BW_2^2(\Sigma, \Sigma')}{BW_{2, \bar{\Sigma}}^2(\Sigma, \Sigma')} = 1 - \left(\frac{a-b}{a+b} \right)^2 \cos^2 \theta + O((a-b)^4), \quad (14)$$

where $BW_{2, \bar{\Sigma}}$ is the linearized Bures-Wasserstein distance at $\bar{\Sigma}$ recalled in equation 29.

For a given θ , equation 14 shows that the distortion is most important for $\frac{|a-b|}{|a+b|}$ close to 1, which corresponds to matrices that are close to the border of the cone, as illustrated in Section 5.1.

On the restriction to the space of Gaussian distributions Geodesic PCA can also be defined in the more general space of a.c. probability distributions, as presented in Section 4. A natural question that arises is whether performing GPCA in the whole space of probability distributions gives the same result as restricting to the space of Gaussian distributions, which is totally geodesic. To our knowledge, the answer to this question is not known in general, although it is true in one dimension.

Proposition 5. *Let $\nu_i = \mathcal{N}(m_i, \sigma_i^2)$ for $i = 1, \dots, n$, be n univariate Gaussian distributions. The first principal geodesic component $t \in [0, 1] \mapsto \mu(t)$ solving equation 1 remains in the space of Gaussian distributions for all $t \in [0, 1]$.*

4 GEODESIC PCA ON A.C. PROBABILITY MEASURES: GPCAGEN

We now tackle the task of performing GPCA on a set of a.c. probability measures ν_1, \dots, ν_n using the Otto-Wasserstein geometry. We propose a parameterization of the geodesic principal components based on Otto’s formulation, leveraging neural networks. Additionally, we introduce a dedicated cost function to optimize the different geodesic components.

Parameterizing geodesics Following Proposition 2 and equation 9, any geodesic $t \mapsto \mu(t)$ in the Wasserstein space $(\text{Prob}(\Omega), W_2)$ can be expressed as $\mu(t) = (\varphi + t\nabla f \circ \varphi)_{\#}\rho$, for t in some interval $[t_{\min}, t_{\max}]$, $\varphi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ a diffeomorphism, $f : \mathbb{R}^d \rightarrow \mathbb{R}$ a smooth function, and ρ a fixed reference measure, taken to be the standard Gaussian distribution in this work. Using multilayer perceptrons (MLPs) to parametrize the functions φ and f , denoted φ_θ and f_ψ , respectively, the curve

$$t \mapsto \mu_{\theta, \psi}(t) = (\text{id} + t\nabla f_\psi)_{\#}(\varphi_\theta \# \rho)$$

is a geodesic for $t \in [t_{\min}, t_{\max}]$, provided that $\text{id} + t\nabla f_\psi \in \text{Diff}(\Omega)$ for all t in this interval. Equivalently, this condition holds if the Hessian matrix $I_d + tH_{f_\psi}(x)$ is positive definite for all $x \in \mathbb{R}^d$ and $t \in [t_{\min}, t_{\max}]$, where $H_{f_\psi}(x)$ denotes the Hessian of f_ψ at x . In practice, we enforce this constraint by monitoring the eigenvalues of $I_d + tH_{f_\psi}(x)$ (see Appendix B.3) and either clipping t or adjusting the interval $[t_{\min}, t_{\max}]$ to ensure that all eigenvalues remain positive. This representation enables to sample from the distributions along the geodesic. Specifically, given the learned vector field φ_θ and function f_ψ , one can sample from $\mu_{\theta, \psi}(t)$ by first drawing $x \sim \rho$ and then applying the transformations φ_θ and $\text{id} + t\nabla f_\psi$ sequentially as $\varphi_\theta(x) + t\nabla f_\psi(\varphi_\theta(x)) \sim \mu_{\theta, \psi}(t)$.

Learning the geodesic components The first principal component in GPCA minimizes the objective in equation 1. The scalar variables t_i specify the projection time of each distribution ν_i onto the geodesic $t \mapsto \mu(t)$. Leveraging the explicit form of Otto’s geodesic, equation 1 can be reformulated as:

$$\inf_{\substack{f \in \mathcal{C}(\mathbb{R}^d), \varphi \in \text{Diff}(\Omega) \\ t_1, \dots, t_n \in [t_{\min}, t_{\max}]}} \mathcal{L}(f, \varphi, t_1, \dots, t_n) := \sum_{i=1}^n W_2^2((\text{id} + t_i \nabla f)_{\#}(\varphi \# \rho), \nu_i). \quad (15)$$

270 We jointly learn the parameters t_i together with the neural networks φ_θ and f_ψ to minimize the
 271 objective in equation 15. In practice, we use the Sinkhorn divergence S_ε that has been proven to be a
 272 differentiable and computationally efficient approximation of the squared Wasserstein distance W_2^2 ,
 273 see Frogner et al. (2015); Genevay et al. (2018); Chizat et al. (2020), and represent the distributions ρ
 274 and ν_i using batches of m samples $x_k \sim \rho$ and $y_j \sim \nu_i$. The optimization proceeds by updating the
 275 parameters based on a single distribution ν_i sampled at each iteration, as detailed in Algorithm 1. To
 276 compute t_{\min} and t_{\max} on line 5 of Algorithm 1, we approximate the extremal eigenvalues of H_{f_ψ} by
 277 evaluating the largest and smallest eigenvalues over the finite set $\{H_{f_\psi}(x_k)\}_{k=1}^m$, and substitute these
 278 estimates into the theoretical bounds from Appendix B.3.

279 **Algorithm 1** Geodesic PCA algorithm for a.c. measures: GPCAGEN

280 1: Initialize φ_θ , f_ψ and the t_i for $1 \leq i \leq n$
 281 2: **while** not converged **do**
 282 3: **for** $i = 1$ to n **do**
 283 4: Draw m i.i.d samples $y_j^{(i)} \sim \nu_i$ and draw m i.i.d samples $x_k \sim \rho$ $1 \leq j, k \leq m$
 284 5: Estimate t_{\min} , t_{\max} with $\{H_{f_\psi}(x_k)\}_{k=1}^m$ and set $t'_i = \min(\max(t_i, t_{\min}), t_{\max})$
 285 6: $z_k^{(i)} \leftarrow (\text{id} + t'_i \nabla f_\psi) \circ (\varphi_\theta)(x_k)$ for $1 \leq k \leq m$
 286 7: $\mathcal{L}_{\theta, \psi, t_i} \leftarrow S_\varepsilon \left(\frac{1}{m} \sum_{k=1}^m \delta_{z_k^{(i)}}, \frac{1}{m} \sum_{j=1}^m \delta_{y_j^{(i)}} \right)$
 287 8: Update φ_θ , f_ψ and the t_i with $\nabla \mathcal{L}_{\theta, \psi, t_i}$
 288 9: **end for**
 289 10: **end while**

292
 293 The second principal component minimizes the objective in equation 1 subject to the constraint that it
 294 intersects the first component orthogonally. Similar to the first component, we use two MLPs, f_{ψ_2}
 295 and φ_{θ_2} , to parameterize the geodesic $t \mapsto \mu_{\theta_2, \psi_2}(t)$, along with n scalar variables t_i^2 , to optimize
 296 the objective in equation 15. We also introduce two additional scalar variables, t_{inter}^1 and t_{inter}^2 , which
 297 define the intersection times of the two geodesics, along with the regularization terms:

298
$$\mathcal{I}(\mu_1, \mu_2, t_{\text{inter}}^1, t_{\text{inter}}^2) = W_2^2(\mu_1(t_{\text{inter}}^1), \mu_2(t_{\text{inter}}^2)) \quad \text{and} \quad \mathcal{O}(g, h) = \frac{\langle g, h \rangle_{L^2(\rho)}^2}{\|g\|_{L^2(\rho)}^2 \|h\|_{L^2(\rho)}^2},$$

300 where \mathcal{I} enforces the geodesics $\mu_1 = \mu_{\theta, \psi}$ and $\mu_2 = \mu_{\theta_2, \psi_2}$ to intersect at the respective times
 301 t_{inter}^1 and t_{inter}^2 , and $\mathcal{O}(g, h)$ ensures orthogonality between the corresponding horizontal vector fields
 302 $g = \nabla f_\psi(\varphi_\theta)$ and $h = \nabla f_{\psi_2}(\varphi_{\theta_2})$ in $L^2(\rho)$. The total objective used to optimize the second
 303 principal component incorporates these regularization terms and is given by:

304
$$\mathcal{L}(f_{\psi_2}, \varphi_{\theta_2}, t_1^2, \dots, t_n^2) + \lambda_I \mathcal{I}(\mu_{\theta, \psi}, \mu_{\theta_2, \psi_2}, t_{\text{inter}}^1, t_{\text{inter}}^2) + \lambda_O \mathcal{O}(\nabla f_\psi(\varphi_\theta), \nabla f_{\psi_2}(\varphi_{\theta_2}))$$

305 where λ_I and λ_O are the regularization parameters controlling the trade-off between the intersection
 306 and orthogonality regularization terms, respectively. Note that in virtue of Proposition 2, the $L^2(\rho)$
 307 inner product in the regularization term \mathcal{O} truly enforces orthogonality of the geodesic components
 308 with respect to the Riemannian metric associated to the Wasserstein distance.

309 The training algorithm used to optimize the second principal component follows the same structure as
 310 Algorithm 1, except for the seventh line, where the regularization terms, estimated using the minibatch
 311 $x_k \sim \rho$, are added to the loss function. Higher-order components can be computed similarly.

312 **5 EXPERIMENTS**

313 **5.1 EXPERIMENTS ON CENTERED GAUSSIAN DISTRIBUTIONS**

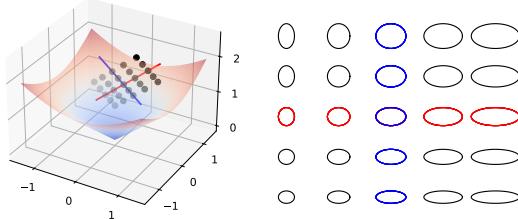
314 In this section, we consider toy examples in S_2^{++} and compare GPCA to its widely used linearized
 315 approximation, TPCA (see Appendix C). We use two coordinate systems for matrices in S_2^{++} : the
 316 first comes from the spectral decomposition, and the second maps any SPD matrix to a point in the
 317 interior of the cone $\mathcal{C} = \{(x, y, z) \in \mathbb{R}^3, z > 0, z^2 < x^2 + y^2\}$

318
$$\Sigma = P_\theta \begin{pmatrix} a^2 & 0 \\ 0 & b^2 \end{pmatrix} P_\theta^\top = \begin{pmatrix} z + y & x \\ x & z - y \end{pmatrix}, \quad (a, b, \theta) \in \mathbb{R}_+^* \times \mathbb{R}_+^* \times \mathbb{R}, \quad (x, y, z) \in \mathcal{C}, \quad (16)$$

324 where P_θ is the rotation matrix of angle θ . Generically, GPCA and TPCA yield very similar results:
 325 for sets of $n = 50$ covariance matrices randomly generated using a uniform distribution on the
 326 parameters (a, b, θ) , GPCA reduces the objective in equation 11 of less than 1% w.r.t. TPCA, on
 327 average for 100 trials. This suggests that TPCA is generally a very good approximation of GPCA.
 328 Two extreme cases are described below: (i) GPCA and TPCA are equivalent and (ii) GPCA and
 329 TPCA drastically differ.

330
 331 **Matrices with same orientation** If we consider a set of covariance matrices that live in the subspace
 332 $\theta = \text{constant}$ in notations of equation 16, then both GPCA and TPCA yield exactly the same results,

333 namely that of linear PCA in the
 334 (a, b) -coordinates. This is because
 335 any such subspace has zero curvature
 336 for the Wasserstein metric, and
 337 geodesics are straight lines in the
 338 (a, b) -coordinates (Appendix D.1).
 339 Figure 3 shows the geodesic compo-
 340 nents obtained for a set of matrices
 341 in the subspace $\theta = 0$ that form a
 342 regular rectangular grid in the (a, b)
 343 coordinates, i.e. $\Sigma_{ij} = \text{diag}(a_i^2, b_j^2)$
 344 where the a_i 's and b_j 's are equally
 345 spaced. They are indeed straight
 346 lines that capture the variations in
 347 a and b respectively.



348
 349 Figure 3: GPCA on a set of diagonal covariance matrices
 350 Σ_{ij} with varying eigenvalues $1 \leq a_i^2 \leq 3$, $1 \leq b_j^2 \leq 2$.
 351 The matrices form a planar grid inside the cone \mathcal{C} of SPD
 352 matrices in equation 16 (**left**), and correspond to ellipses
 353 of varying width and height (**right**). The first component
 354 (red) captures the variation in a , while the second component
 355 (blue) captures the variation in b .

356 **Matrices with same eigenvalues** Now we consider covariance matrices that all have the same
 357 eigenvalues but different orientations. Specifically, we choose $\Sigma_i = P_{\theta_i} \text{diag}(a^2, b^2) P_{\theta_i}^\top$, for positive
 358 reals $a > b$, $\theta_i = i\pi/n$ for $i = 0, \dots, n-1$ and an even number n . In the (x, y, z) coordinates
 359 (equation 16), the covariance matrices are displayed on a circle of equation $x = \text{cst}$ (constant trace)
 360 and $y^2 + z^2 = \text{cst}$ (constant determinant), as shown in Figure 4 (in practice, we choose a slightly open
 361 circle to break the symmetry). Then the Bures-Wasserstein barycenter of the covariance matrices
 362 $\Sigma_1, \dots, \Sigma_n$ is given by $\bar{\Sigma} = (a+b)^2/4 I$ (see Proposition 15 in Appendix D.1). When performing
 363 TPCA on $\Sigma_1, \dots, \Sigma_n$ at the barycenter $\bar{\Sigma}$, the radial distances between $\bar{\Sigma}$ and Σ_i are preserved, but
 364 not the pairwise distances between the Σ_i 's. **Proposition 4 evaluates the level of this distortion. Note**
 365 **that since $(a-b)^2/(a+b)^2 = (x^2 + y^2)/z^2$, the distortion is most important when covariance**
 366 **matrices are close to the border of the cone, see Figure 4 (**left**)**. Indeed, in that case, the results of
 367 GPCA can be very different from those of TPCA and the first component may not even go through
 368 the Wasserstein barycenter $\bar{\Sigma}$, see Figure 4 (**middle**) and Figure 8 in Appendix A. In that case GPCA
 369 may be seen as worse-behaved as TPCA, as some of the Gaussian distributions will project onto
 370 the first geodesic component boundaries, yielding a poor separation. Figure 4 (**right**) shows the
 371 percentage of improvement of the cost in equation 11 (in terms of minimization) of GPCA with
 372 respect to TPCA, in the setting previously described for different values of the ratio $|a-b|/|a+b|$.
 373 on average for 10 runs per value of the ratio. The blue strip indicates standard deviation.

374
 375 **Weather dataset** In this paragraph, we use the Weather CORGIS Dataset to illustrate GPCA based
 376 on empirical covariance matrices. The dataset provides weekly measures of precipitation and wind
 377 speed recorded from March 2016 to January across the 50 U.S. states and the territory of Puerto
 378 Rico. From these measures, we construct two histograms for each state: one for precipitation and
 379 one for wind speed. We then compute the 51 empirical covariances from these histograms. We show
 380 in Figure 14 the projection of each state onto the two first GPCA components computed from the
 381 empirical covariance matrices. We can clearly identify clusters of different weather behavior among
 382 the states.

383 5.2 EXPERIMENTS ON ABSOLUTELY CONTINUOUS DISTRIBUTIONS

384
 385 We conduct a preliminary experiment on a synthetic dataset with known geodesics to verify that
 386 our algorithm, GPCAGEN (Section 4), accurately recovers the first two principal components. We

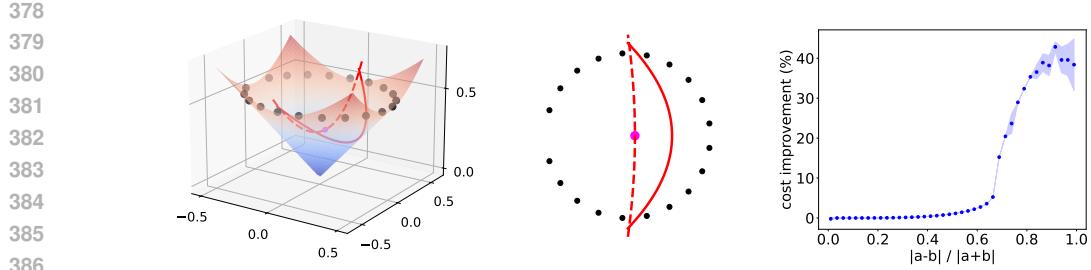


Figure 4: Comparison between tangent and geodesic PCA on a set of $n = 20$ covariance matrices with same eigenvalues a^2, b^2 and different orientations θ . **(left)** They are equally spaced on an (open) circle in a horizontal plane inside the cone of SPD matrices. The first component of TPCA (dashed red line) goes through the Fréchet mean $\bar{\Sigma}$ (magenta dot), a multiple of the identity, while the component of GPCA (solid red line) does not. Here $|a - b|/|a + b| \approx 0.8$. **(middle)** Representation of the left figure in the (x, y) coordinates. **(right)** Evolution of the first component cost improvement (in the sense of minimization) of GPCA with respect to TPCA, as a function of the ratio $|a - b|/|a + b|$.

then apply GPCAGEN to 3D point clouds from the ModelNet40 dataset (Wu et al. (2015)) and to color distributions of images from the Landscape Pictures dataset (Rougetet (2020)). **An additional experiment in Appendix A.3 demonstrate how GPCA can be used for outlier detection.** For these experiments, f_ψ and φ_θ are MLPs with four hidden layers of size 128 and an output layer of size 1 and d respectively. We found that setting the regularization coefficients λ_I and λ_O to 1.0 ensures the algorithm works as expected in all experiments. A discussion of the regularization coefficients, along with details on the architecture and hyperparameters, is provided in Appendix E.

MNIST geodesics. We represent each image from the MNIST dataset (LeCun et al. (2010)) as a probability measure over \mathbb{R}^4 . The grayscale pixel intensities define a normalized density over spatial coordinates $(x, y) \in \mathbb{R}^2$, and we further assign each pixel two additional values corresponding to red and blue color channels. We construct two orthogonal geodesics: the first one interpolates between a digit "1" and a digit "2", both assigned a fixed purple by setting the color channels to 0.5. The second one is defined from the midpoint of the first, by linearly interpolating the color from red to blue. As shown in Figures 5 and 9, GPCAGEN successfully recovers the two geodesics intersecting orthogonally. A second experiment on the MNIST dataset is displayed in Appendix A.

3D point cloud. We use the ModelNet40 3D point cloud dataset (Wu et al. (2015)) and apply GPCA to a subset of 100 randomly selected lamp point clouds. Figure 6 (**middle row**) and Figure 7 (**left**) demonstrate that the first principal component captures the distinction between hanging lamps (chandeliers) and standing lamps (floor lamps), while the second component reflects variations in the thickness of the lamp structure. We conduct a similar experiment on 100 point clouds from ModelNet40 representing different chairs. As shown in Figure 6 (**top row**) and Figure 10, the first principal component captures the height of the seat, while the second component distinguishes between chairs and armchairs.

Landscape images. We took 39 images from the Landscape Pictures dataset (Rougetet (2020)) and use GPCAGEN on the corresponding point clouds, where each point cloud represents color

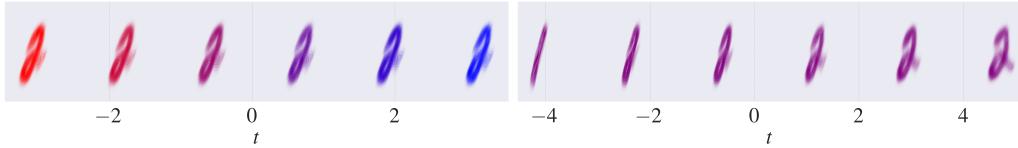


Figure 5: Densities of probability distributions uniformly sampled along the first and second principal geodesics components. GPCAGEN successfully recovers the two orthogonally intersecting geodesics constructed from MNIST data. The first component (**left**) captures variation in color space, while the second component (**right**) recovers the interpolation from the digit "1" to the digit "2".

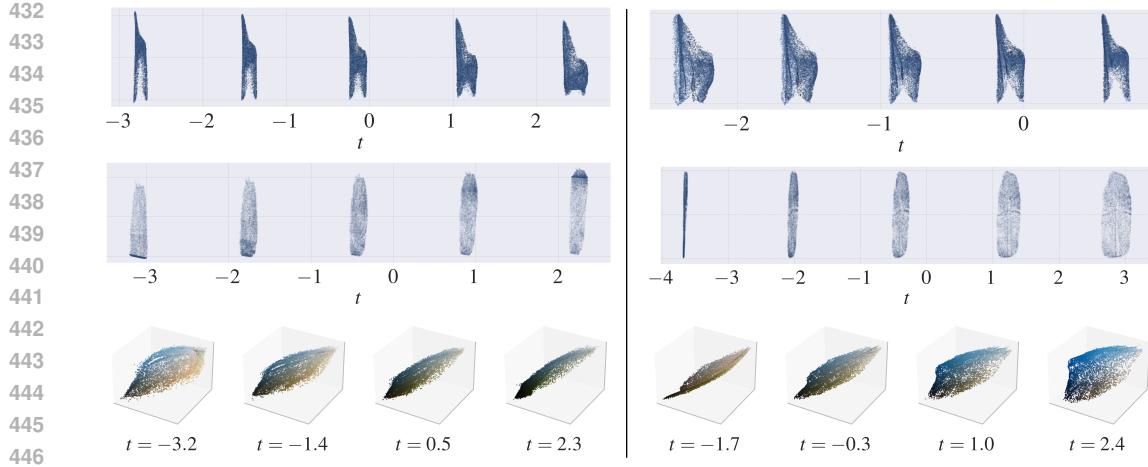


Figure 6: Empirical distributions uniformly sampled along the geodesics corresponding to the first (left) and second (right) principal components, as computed by GPCAGEN in the 3D point cloud of chairs experiment (top row), the 3D point cloud of lamps experiment (middle row) and the Landscape images experiment (bottom row).

distribution in the image. Figure 6 (bottom row) and Figure 7 (right) show that the first component captures variations in overall brightness, ranging from bright to dark images, while the second component separates mostly green images from mostly blue ones.

Baselines An obvious baseline for GPCAGEN is TPCA. Unlike GPCAGEN, which learns continuous geodesics from empirical distributions of absolutely continuous measures, TPCA acts on discrete measures. A direct numerical comparison between the two methods is therefore not meaningful. However, we include in Appendix A.2 the two principal components returned by TPCA on the 3D point cloud experiments. We observe in Figure 16 that the discrete nature of TPCA produces artifacts, including holes in certain regions, excessive mass concentration in others, and intermediate distributions that no longer resemble valid objects.

Another natural baseline consists in embedding point clouds into a latent space of dimension d then performing standard PCA on the resulting latent vectors. This approach, in addition to being computationally expensive, does not produce meaningful modes of variation, as shown in Section A.2 of the appendices.

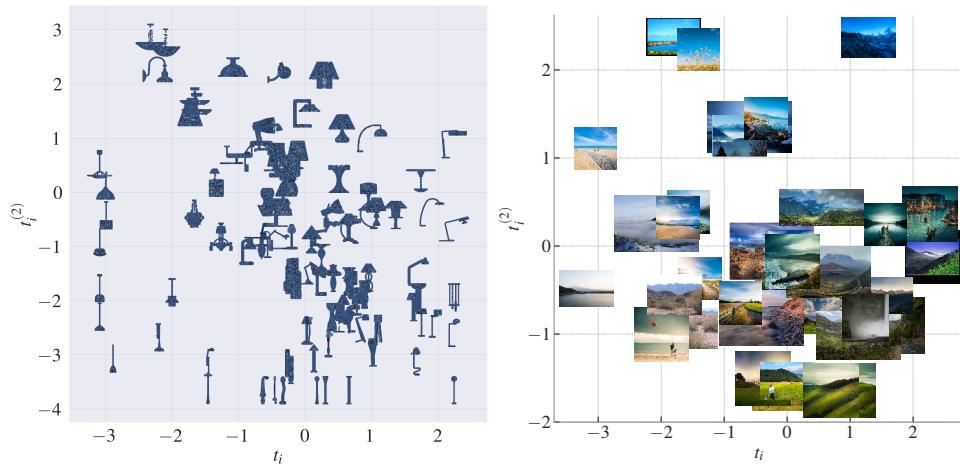


Figure 7: Each lamp point cloud (left) and each image (right) is embedded in the plane according to its projection times onto the first and second principal components computed by GPCAGEN.

486 6 DISCUSSION
487488 GPCA is a statistical approach for learning the main modes of variations of a set of probability
489 distributions. The first components capture meaningful structure for data lying on a curved space,
490 which then enables downstream tasks such as classification, clustering, and outlier detection. In
491 this work, we have proposed two methods for computing exact GPCA : one tailored for Gaussian
492 distributions and the other for the more general case of a.c. probability distributions. In the Gaussian
493 case, our experiments suggest that GPCA and TPCA generically yield very similar results, except for
494 distributions with covariance matrices that are close to the boundary of the SPD cone, for which GPCA
495 can yield undesirable effects as suggested by the pathological example of Figure 4. In the general case
496 of a.c. probability measures, a key advantage of our approach is that it operates directly on continuous
497 distributions, avoiding the need for empirical approximations of the ν_i , which would require equal
498 sample sizes and can introduce discretization artifacts in the recovered components. Additionally, our
499 method enables sampling from any point along the geodesic components—something not possible
500 with discrete approximations commonly used in TPCA. Otto’s parametrization also allowed us to
501 avoid relying on input convex neural networks (ICNNs) by not requiring convex functions, with the
502 trade-off being the need to estimate the eigenvalues of the Hessian of f . This perspective opens new
503 directions for parametrizing convex functions without imposing hard architectural constraints.
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539

540 REPRODUCIBILITY STATEMENT
541

542 All implementation details of our proposed method, including model architectures, training proce-
543 dures, and hyperparameter settings, are provided in Section 5 of the main paper and in Appendix E
544 and D.2. Original theoretical results are presented with complete proofs in Appendix D. The datasets
545 used in our experiments are publicly available. We will release the source code to reproduce all
546 experiments associated with this paper at a later stage.

548 REFERENCES
549

550 Martial Aguech and Guillaume Carlier. Barycenters in the Wasserstein space. *SIAM Journal on*
551 *Mathematical Analysis*, 43(2):904–924, 2011.

552 Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the*
553 *space of probability measures*. Springer Science & Business Media, 2008.

554 Luigi Ambrosio, Alberto Bressan, Dirk Helbing, Axel Klar, Enrique Zuazua, Luigi Ambrosio, and
555 Nicola Gigli. A user’s guide to optimal transport. *Modelling and Optimisation of Flows on*
556 *Networks: Cetraro, Italy 2009*, Editors: Benedetto Piccoli, Michel Rascle, pp. 1–155, 2013.

557 Rajendra Bhatia, Tanvi Jain, and Yongdo Lim. On the Bures–Wasserstein distance between positive
558 definite matrices. *Expositiones Mathematicae*, 37(2):165–191, 2019.

559 Jérémie Bigot, Raúl Gouet, Thierry Klein, and Alfredo López. Geodesic PCA in the Wasserstein
560 space by convex PCA. *Annales de l’Institut Henri Poincaré - Probabilités et Statistiques*, 53(1):
561 1–26, 2017.

562 Emmanuel Boissard, Thibaut Le Gouic, and Jean-Michel Loubes. Distribution’s template estimate
563 with Wasserstein metrics. *Bernoulli*, 21(2):740–759, 2015.

564 Nicolas Boumal. *An introduction to optimization on smooth manifolds*. Cambridge University Press,
565 2023.

566 Yann Brenier. Polar factorization and monotone rearrangement of vector-valued functions. *Commu-*
567 *nications on pure and applied mathematics*, 44(4):375–417, 1991.

568 Anna Calissano, Aasa Feragen, and Simone Vantini. Populations of unlabelled networks: Graph
569 space geometry and generalized geodesic principal components. *Biometrika*, pp. 147–170, 2024.

570 Elsa Cazelles, Vivien Seguy, Jérémie Bigot, Marco Cuturi, and Nicolas Papadakis. Geodesic PCA
571 versus log-PCA of histograms in the Wasserstein space. *SIAM Journal on Scientific Computing*,
572 40(2):B429–B456, 2018.

573 Lenaic Chizat, Pierre Roussillon, Flavien Léger, François-Xavier Vialard, and Gabriel Peyré. Faster
574 wasserstein distance estimation with the sinkhorn divergence. *Advances in neural information*
575 *processing systems*, 33:2257–2269, 2020.

576 Jed A Duersch, Meiyue Shao, Chao Yang, and Ming Gu. A robust and efficient implementation of
577 LOBPCG. *SIAM Journal on Scientific Computing*, 40(5):C655–C676, 2018.

578 David G Ebin and Jerrold Marsden. Groups of diffeomorphisms and the motion of an incompressible
579 fluid. *Annals of Mathematics*, 92(1):102–163, 1970.

580 P. Thomas Fletcher, Conglin Lu, and Sarang Joshi. Statistics of shape via principal geodesic analysis
581 on Lie groups. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern*
582 *Recognition, 2003. Proceedings.*, volume 1, pp. I–I. IEEE, 2003.

583 Charlie Frogner, Chiyuan Zhang, Hossein Mobahi, Mauricio Araya, and Tomaso A Poggio. Learning
584 with a wasserstein loss. *Advances in neural information processing systems*, 28, 2015.

585 Aude Genevay, Gabriel Peyré, and Marco Cuturi. Learning generative models with sinkhorn di-
586 vergences. In *International Conference on Artificial Intelligence and Statistics*, pp. 1608–1617.
587 PMLR, 2018.

594 Philipp Harms, Elodie Maignant, and Stefan Schlager. Approximation of Riemannian distances and
 595 applications to distance-based learning on manifolds. *arXiv preprint :1904.11860*, 2019.
 596

597 Wen Huang and Ke Wei. Riemannian proximal gradient methods. *Mathematical Programming*, 194
 598 (1):371–413, 2022.

599 Stephan Huckemann and Herbert Ziezold. Principal component analysis for Riemannian manifolds,
 600 with an application to triangular shape spaces. *Advances in Applied Probability*, 38(2):299–319,
 601 2006.

602 Stephan Huckemann, Thomas Hotz, and Axel Munk. Intrinsic shape analysis: Geodesic PCA for
 603 Riemannian manifolds modulo isometric Lie group actions. *Statistica Sinica*, pp. 1–58, 2010.

604

605 Boris Khesin, Gerard Misiołek, and Klas Modin. Geometric hydrodynamics and infinite-dimensional
 606 Newton’s equations. *Bulletin of the American Mathematical Society*, 58(3):377–442, 2021.

607

608 Benoit Kloeckner. A geometric study of Wasserstein spaces: Euclidean spaces. *Annali della Scuola
 609 Normale Superiore di Pisa-Classe di Scienze*, 9(2):297–323, 2010.

610 Yann LeCun, Corinna Cortes, and CJ Burges. MNIST handwritten digit database. *ATT Labs [Online]*.
 611 Available: <http://yann.lecun.com/exdb/mnist>, 2, 2010.

612 Cyril Letrouit and Quentin Mérigot. Gluing methods for quantitative stability of optimal transport
 613 maps. *arXiv preprint arXiv:2411.04908*, 2024.

614

615 Chi-Kwong Li and Fuzhen Zhang. Eigenvalue continuity and Gersgorin’s theorem. *Electronic
 616 Journal of Linear Algebra*, 2019.

617 Yingjie Ma, Xi Gao, Chao Liu, and Jie Li. Improved SQP and SLSQP algorithms for feasible
 618 path-based process optimisation. *Computers & Chemical Engineering*, pp. 108751, 2024.

619

620 Luigi Malagò, Luigi Montrucchio, and Giovanni Pistone. Wasserstein Riemannian geometry of
 621 Gaussian densities. *Information Geometry*, 1:137–179, 2018.

622

623 Robert J McCann. A convexity principle for interacting gases. *Advances in mathematics*, 128(1):
 624 153–179, 1997.

625

626 Nina Miolane, Nicolas Guigui, Alice Le Brigant, Johan Mathe, Benjamin Hou, Yann Thanwerdas,
 627 Stefan Heyder, Olivier Peltre, Niklas Koep, Hadi Zaatiti, Hatem Hajri, Yann Cabanes, Thomas
 628 Gerald, Paul Chauchat, Christian Shewmake, Daniel Brooks, Bernhard Kainz, Claire Donnat,
 629 Susan Holmes, and Xavier Pennec. Geomstats: A Python package for Riemannian geometry
 630 in machine learning. *Journal of Machine Learning Research*, 21(223):1–9, 2020. URL <http://jmlr.org/papers/v21/19-027.html>.

631

632 Klas Modin. Geometry of matrix decompositions seen through optimal transport and information
 633 geometry. *Journal of Geometric Mechanics*, 9(3):335–390, 2017.

634

635 Gaspard Monge. Mémoire sur la théorie des déblais et des remblais. *Mem. Math. Phys. Acad. Royale
 636 Sci.*, pp. 666–704, 1781.

637

638 Felix Otto. The geometry of dissipative evolution equations: the porous medium equation. *Communications in Partial Differential Equations*, 2001.

639

640 Michael J.D. Powell. On search directions for minimization algorithms. *Mathematical programming*,
 4:193–201, 1973.

641

642 Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets
 643 for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision
 and pattern recognition*, pp. 652–660, 2017.

644

645 James O. Ramsay and Bernard W. Silverman. *Applied functional data analysis: methods and case
 646 studies*. Springer, 2002.

647 Arnaud Rougetet. Landscape pictures dataset. <https://www.kaggle.com/datasets/arnaud58/landscape-pictures>, 2020.

648 Klaus D Schmidt. On inequalities for moments and the covariance of monotone functions. *Insurance: Mathematics and Economics*, 55:91–95, 2014.

649

650

651 Vivien Seguy and Marco Cuturi. Principal geodesic analysis for probability measures under the optimal transport metric. *Advances in Neural Information Processing Systems*, 28, 2015.

652

653 Stefan Sommer, François Lauze, and Mads Nielsen. The differential of the exponential map, Jacobi fields and exact principal geodesic analysis. *CoRR, abs/1008.1902*, 2010.

654

655

656 Stefan Sommer, François Lauze, and Mads Nielsen. Optimization over geodesics for exact principal geodesic analysis. *Advances in Computational Mathematics*, 40:283–313, 2014.

657

658 Asuka Takatsu. Wasserstein geometry of Gaussian measures. *Osaka J. Math.*, 48:1005–1026, 2011.

659

660 Yann Thanwerdas. *Riemannian and stratified geometries on covariance and correlation matrices*. PhD thesis, Université Côte d’Azur, 2022.

661

662 Pauli Virtanen, Ralf Gommers, Travis E. Oliphant, Matt Haberland, Tyler Reddy, David Cournapeau, Evgeni Burovski, Pearu Peterson, Warren Weckesser, Jonathan Bright, Stéfan J. van der Walt, Matthew Brett, Joshua Wilson, K. Jarrod Millman, Nikolay Mayorov, Andrew R. J. Nelson, Eric Jones, Robert Kern, Eric Larson, C J Carey, İlhan Polat, Yu Feng, Eric W. Moore, Jake VanderPlas, Denis Laxalde, Josef Perktold, Robert Cimrman, Ian Henriksen, E. A. Quintero, Charles R. Harris, Anne M. Archibald, Antônio H. Ribeiro, Fabian Pedregosa, Paul van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

663

664

665

666

667

668

669

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

671

672

673

674 Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes, 2015. URL <https://arxiv.org/abs/1406.5670>.

675

676

677 Constantin Zalinescu. *Convex analysis in general vector spaces*. World scientific, 2002.

678

679 Juan Zhou, Kangkang Deng, Hongxia Wang, and Zheng Peng. Inexact Riemannian gradient descent method for nonconvex optimization. *arXiv preprint arXiv:2409.11181*, 2024.

680

681

682

683

684

685

686

687

688

689

690

691

692

693

694

695

696

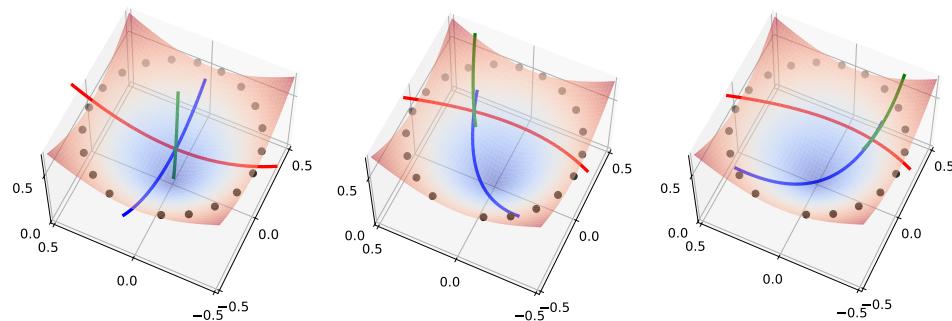
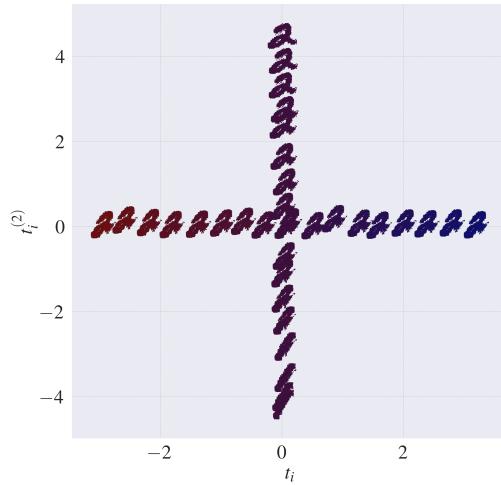
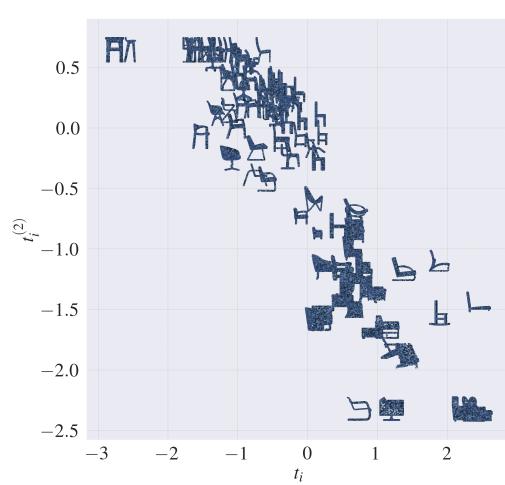
697

698

699

700

701

702 **A ADDITIONAL EXPERIMENTS AND FIGURES**
703704 **A.1 GEODESIC PCA**
705706 Here we present additional figures to further explain the experiments described in the paper. Figure 8
707 concerns the experiment on Gaussian distributions with diagonal covariances described in Section 5.1
708 corresponding to Figure 4. It shows all three principal components found by tangent PCA (**left**) and
709 geodesic PCA, in two equally optimal solutions (**middle, right**).
710721
722 Figure 8: Principal geodesic components of a set of Gaussian distributions whose covariance matrices
723 have same eigenvalues and different orientations, as described in Section 5.1. Tangent PCA yields a
724 unique solution (**left**) where geodesic components cross at the barycenter, while geodesic PCA yields
725 two equally optimal solutions (**middle, right**) where the geodesic components cross at another point.
726 The first geodesic component is shown in red, the second in blue, the third in green.
727727 Figure 9 displays on the plane the two first geodesic components of the MNIST experiment of Section
728 5.2, while Figure 10 shows the planar representation of the 3D point cloud of chairs experiment
729 given by the projection onto the first two geodesic components found by GPCAGEN algorithm and
730 depicted in Figure 6 (**top row**).
731748 Figure 9: Each point cloud, corresponding to
749 a distribution along one of the artificially con-
750 structed geodesics, is embedded in the plane ac-
751 cording to its projection times onto the first and
752 second geodesics returned by the GPCAGEN
753 algorithm. We observe that GPCAGEN suc-
754 cessfully recovers the two orthogonally intersecting
755 geodesics designed from MNIST-based interpo-
lations of digit shape and color.
756757 Figure 10: Each chair point cloud is embedded in
758 the plane according to its projection times onto
759 the first and second geodesics returned by the
760 GPCAGEN algorithm.
761

Finally, we present an additional experiment on the MNIST dataset. We use the same color construction as in the experiment presented in Section 5.2, we then apply GPCAGEN to a dataset of 20 red digits "1", 20 blue digits "1", 20 red digits "2", and 20 blue digits "2" (see Figure 12). As shown in Figures 11 and 12, GPCAGEN again identifies two orthogonal geodesics: the first primarily captures variation in color, while the second captures variation in shape—from digit "2" to digit "1".

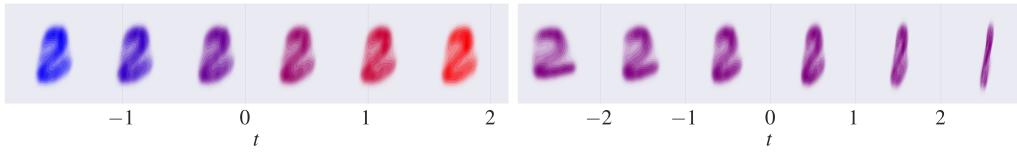


Figure 11: Densities of probability distributions uniformly sampled along the geodesics corresponding to the first and second principal components. The first component (**left**) returned by GPCAGEN captures variation in color space, while the second component (**right**) recovers the interpolation between digit "2" and digit "1".

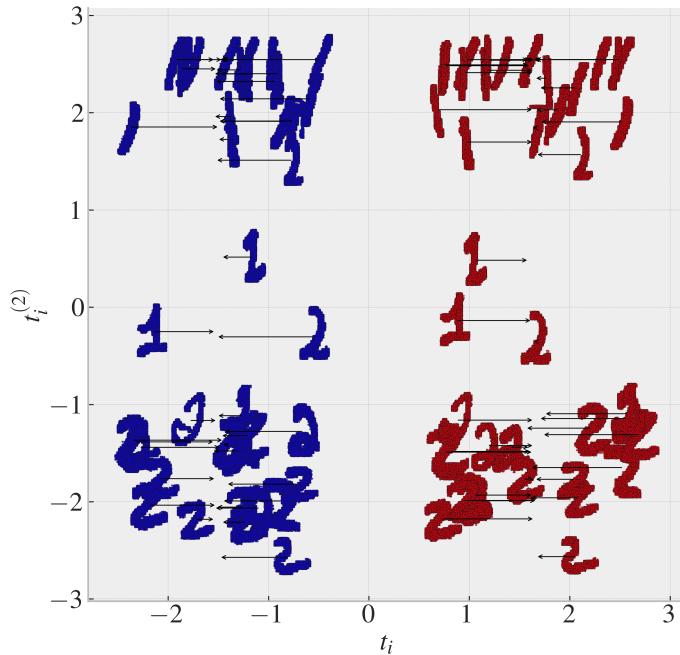


Figure 12: Each MNIST digit is embedded in the plane (the arrows indicate the exact position of each digit) according to its projection times onto the first and second geodesics returned by the GPCAGEN algorithm. We observe that the first principal component recovered by GPCAGEN captures variation in color, while the second component reflects the transformation from digit "2" to digit "1".

A.2 COMPARISON OF GPCA TO RELATED METHODS

Other notions of PCA on Gaussian distributions There exist a wide variety of metrics on the space of symmetric positive definite matrices, such as e.g. the log-Euclidean, Euclidean-Cholesky or affine-invariant metrics (see Thanwerdas (2022) for a comprehensive overview). Each of these metrics could be used to perform PCA on centered Gaussian distributions. However, there is no obvious quantitative way to compare the results. Each method optimizes its own criterion, and any metric that one could think of to compare the methods would rely on a choice of underlying metric on the space of SPD matrices. Comparison of PCA methods with two different metrics thus boils down to comparing the metrics themselves. We illustrate in Figure 13 the behavior of covariances matrices along geodesics for different metrics.

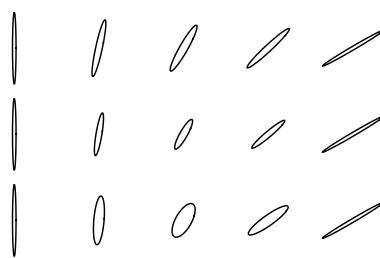


Figure 13: Geodesics on the space of symmetric positive definite matrices from left to right, for **(top)** the Bures-Wasserstein metric, **(middle)** the log-Euclidean metric and **(bottom)** the Euclidean metric on the Cholesky coefficients.

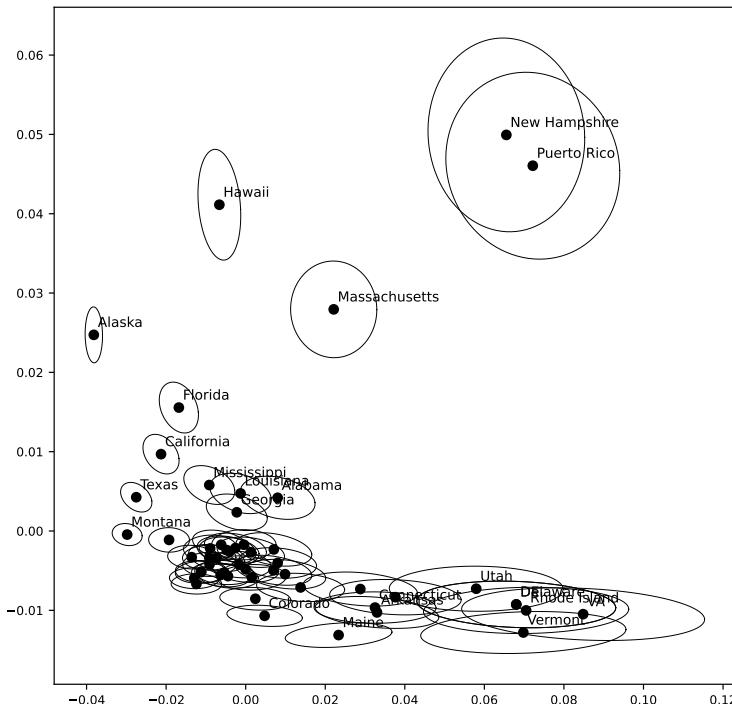


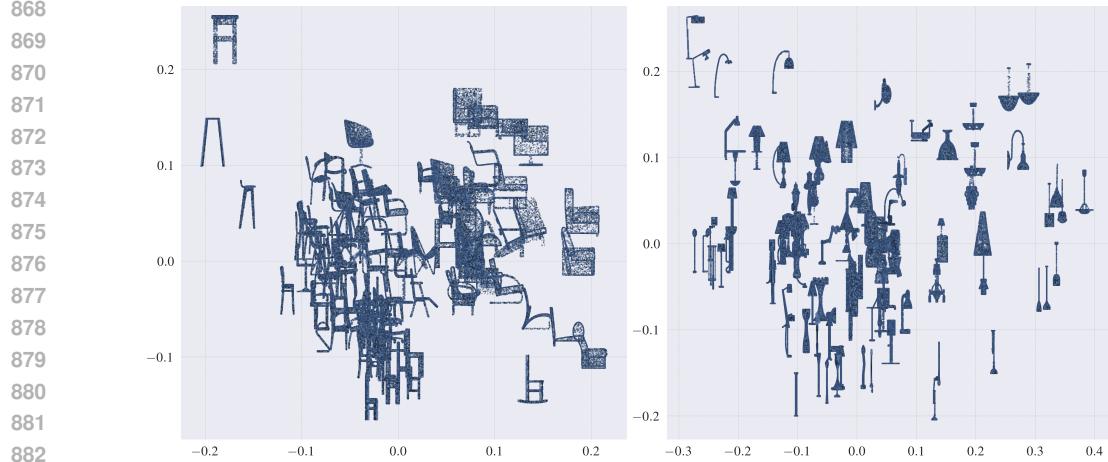
Figure 14: Each dot represents a state projected onto the first two GPCA components computed from the empirical covariance matrices, which are also shown in the figure.

TPCA on 3D Point Cloud Data Here we present the results returned by TPCA on the 3D point cloud experiments, see Figures 16 and 15, and compare them to those obtained by GPCAGEN.

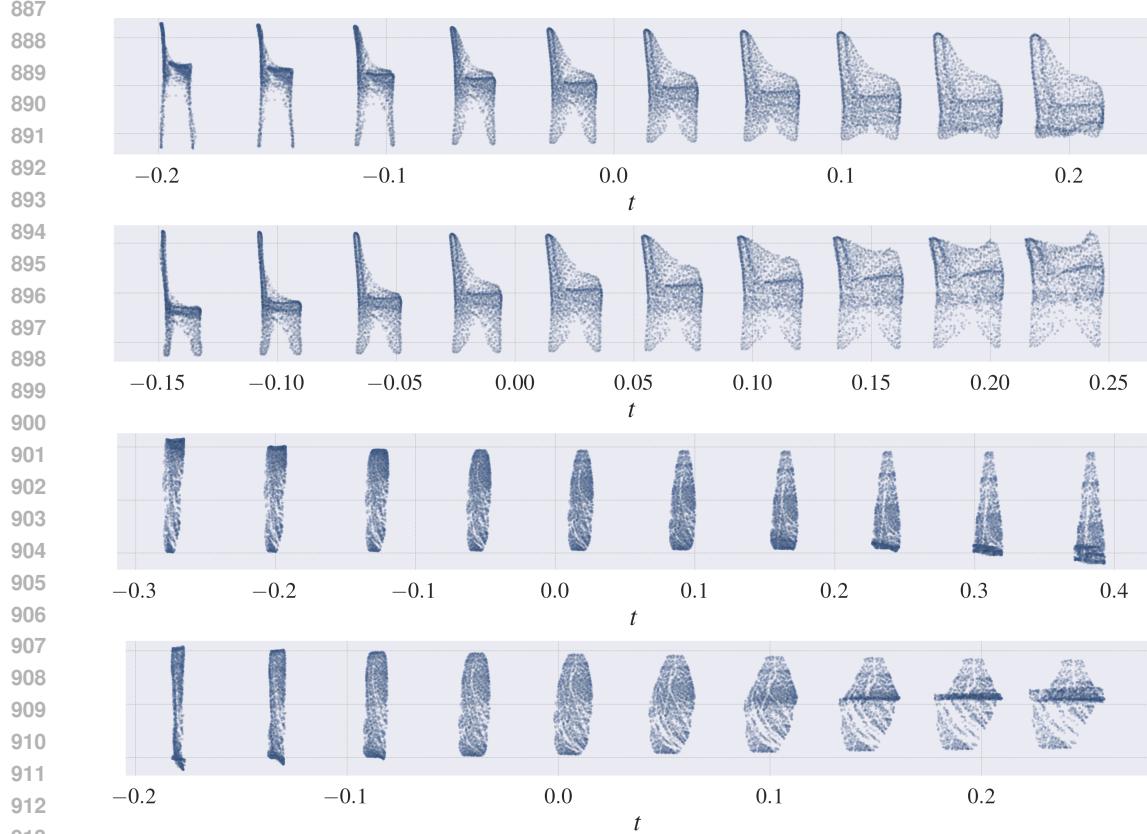
For the lamps dataset, the first component is similar and captures the distinction between hanging and standing lamps. The second component focuses on the object thickness, like the second GPCAGEN component, but also on whether mass is concentrated at the extremities or the middle of the lamp structure.

For the chairs dataset, both geodesics obtained by TPCA differ from those returned by GPCA. The first component interpolates from a thin chair with a high seat to a low-seated armchair. The second component captures a transition from a thin chair with a low seat to a high-seated armchair. The first TPCA component appears to blend the first and second GPCAGEN geodesics, while the second is complementary to the first TPCA component.

864
 865 Finally, due to the discrete nature of the TPCA algorithm, we observe discretization artifacts in the
 866 TPCA components: holes in some parts of the space, mass concentration in others, and intermediate
 867 distributions that do not resemble valid objects.



883 Figure 15: For the chair and the lamp experiment, each point cloud is embedded in the plane according
 884 to its projection times onto the first and second principal components computed by TPCA.
 885



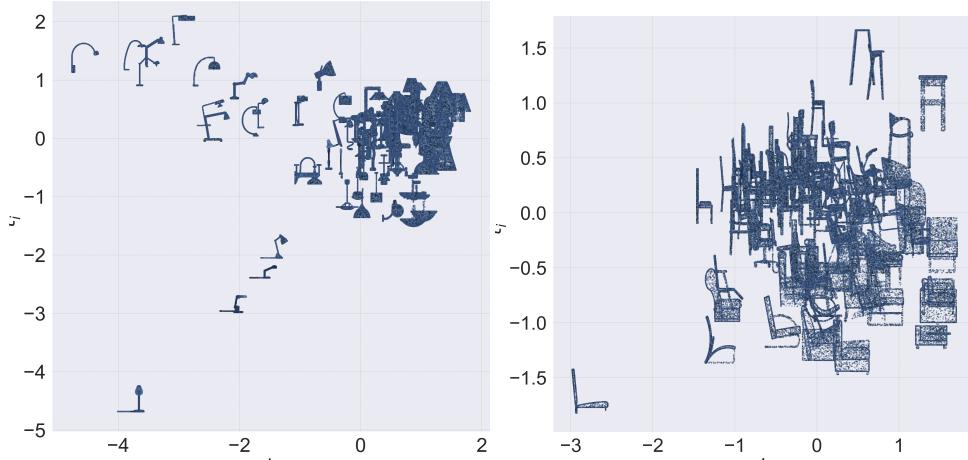
914 Figure 16: Empirical distributions uniformly sampled along the geodesics corresponding to the first
 915 (first line) and second (second line) principal components, as computed by TPCA in the 3D point
 916 cloud of chairs experiment (top rows) and the 3D point cloud of lamps experiment (bottom rows).
 917

918 **PCA computed in the latent space of PointNet.** For the 3D point-cloud datasets, we evaluated
 919 the natural baseline that consists in embedding point clouds into a latent space of dimension d
 920 and then performing standard PCA on the resulting latent vectors. We used a pretrained PointNet
 921 autoencoder (Qi et al., 2017) from the public repository https://github.com/vinits5/pc_autoencoder, trained on ModelNet40, to encode each point cloud (chairs and lamps) into
 922 a d -dimensional latent representation, on which PCA was applied. Figure 17 shows the resulting
 923 2D projections. We observe some clustering of similar objects; for example, large lamps tend to
 924 group together in the lamp dataset, and chairs versus armchairs form distinguishable clusters. The
 925 second principal component for chairs appears to correlate with the height of the seat. Beyond these
 926 observations, however, PCA provides limited separability (especially for lamps), and the recovered
 927 components are difficult to interpret.
 928

929 More generally, this approach presents several important limitations:

930

 931 - Training a point-cloud autoencoder requires a large collection of distributions. In our case
 932 (100 distributions), we need to rely on a pretrained autoencoder trained on related dataset.
 933 - PCA on autoencoder embeddings relies heavily on the geometry learned by the encoder.
 934 The learned geometry is not guaranteed to align with the Wasserstein structure and the
 935 recovered principal components may not reflect meaningful modes of variation (as observed
 936 in the experiments above). Moreover, for a given autoencoder that we wish to train, different
 937 random seeds at initialization can lead to different learned geometries and thus different
 938 PCA components, which is not suitable.



955 Figure 17: For the lamp (left) and the chair (right) experiment, each point cloud is embedded in the
 956 plane according to its projection times onto the first and second principal components computed by
 957 the POINTNET + PCA method.

958 A.3 APPLICATION OF GPCA TO OUTLIER DETECTION

959 In this section, we demonstrate how GPCA can be used for outlier detection. The underlying intuition
 960 is that GPCA components capture the structure of the dataset on which they are trained, and samples
 961 from a different dataset are expected to lie far from the learned components in Wasserstein
 962 distance. In this experiment, we use the ModelNet40 3D point cloud dataset (Wu et al., 2015) and apply
 963 GPCA to a subset of 100 randomly selected chair point clouds to compute the first two components.
 964 For a new point cloud X , we define its score as the sum of the Wasserstein distances between
 965 X and its projections onto the first two learned GPCA components. To compute the Wasserstein
 966 distance between X and a component, we use `ot.emd` from the POT library. Specifically, for each
 967 component, we perform a grid search over 20 equally spaced values of t between t_{\min} and t_{\max} ,
 968 computing the Wasserstein distance between X and 2048 samples drawn from the component at
 969 each t , and select the t that minimizes this distance. We repeat the same procedure for the second
 970 component and sum the two minimal distances to obtain the final score.
 971

We evaluate this approach on 120 point clouds: 60 new chairs (not used for training) and 60 point clouds of cars. The left histogram in Figure 18 shows the resulting scores. We observe that the scores of the chair point clouds (in blue) are lower than those of the car point clouds (in green), indicating that it is possible to detect whether a point cloud is not a chair using this score. We also repeat the experiment with 60 point clouds of planes, shown in the right histogram of Figure 18, and observe that the separation between chair and plane scores is even more pronounced.

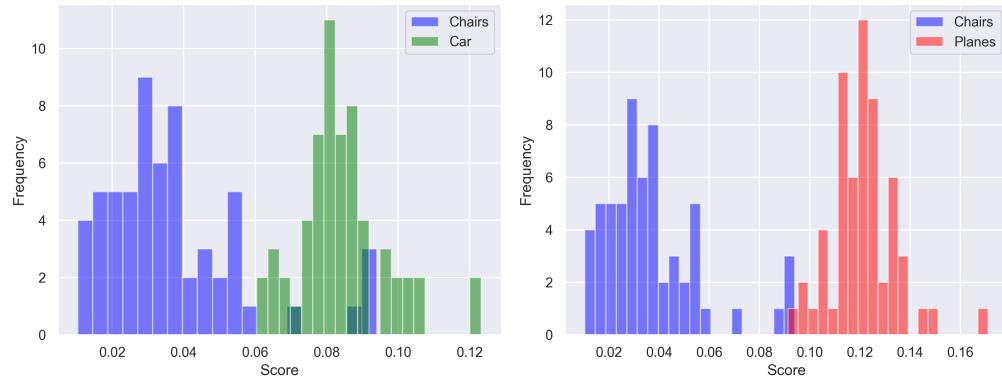


Figure 18: GPCA scores obtained on 60 new point clouds of chairs (never seen during training) and 60 point clouds of cars (left) / planes (right). The separation of the histograms indicates that GPCA can be used for outlier detection.

B THE OTTO-WASSERSTEIN GEOMETRY

In this section, we briefly describe the fiber bundle structure over the Wasserstein space due to Otto (2001), that is behind the Riemannian interpretation of the Wasserstein distance. We then present its restriction to the space of centered non-degenerate Gaussian distributions, which coincides with the Bures-Wasserstein Riemannian geometry on SPD matrices. Finally, we relate Otto’s parametrization of geodesics to McCann’s interpolation.

We present these well-known results without proofs and refer the interested reader to Otto (2001); Khesin et al. (2021) and (Ambrosio et al., 2013, Section 6.1) for more details in the general setting and to Takatsu (2011); Malagò et al. (2018); Bhatia et al. (2019) for details and proofs in the Gaussian setting.

B.1 THE OTTO-WASSERSTEIN GEOMETRY OF A.C. DISTRIBUTIONS

Consider the space $\text{Prob}(\Omega)$ of absolutely continuous probability measures with smooth densities with respect to the Lebesgue measure, and support included in a compact set $\Omega \subset \mathbb{R}^d$, as well as the space $\text{Diff}(\Omega)$ of diffeomorphisms on Ω . These spaces can be equipped with an infinite-dimensional manifold structure, see e.g. Ebin & Marsden (1970), that we will not describe here. The tangent space of $\text{Diff}(\Omega)$ at $\varphi \in \text{Diff}(\Omega)$ is given by

$$T_\varphi \text{Diff}(\Omega) = \{v \circ \varphi, v : \Omega \rightarrow \mathbb{R}^d \text{ vector field}\}.$$

We fix a reference measure $\rho \in \text{Prob}(\Omega)$ and equip $\text{Diff}(\Omega)$ with the L^2 -metric with respect to ρ , defined for any tangent vectors $u \circ \varphi, v \circ \varphi \in T_\varphi \text{Diff}(\Omega)$ as

$$\langle u \circ \varphi, v \circ \varphi \rangle_{L^2(\rho)} := \int (u \circ \varphi) \cdot (v \circ \varphi) d\rho = \int u \cdot v d\mu,$$

where $\mu = \varphi_\# \rho$. Then the space of diffeomorphisms can be decomposed into *fibers*, defined to be equivalence classes under the projection

$$\pi : \text{Diff}(\Omega) \rightarrow \text{Prob}(\Omega), \quad \varphi \mapsto \varphi_\# \rho.$$

1026 Specifically, the fiber over $\mu \in \text{Prob}(\Omega)$ is given by $\pi^{-1}(\mu) = \{\varphi \in \text{Diff}(\Omega), \varphi \# \rho = \mu\}$, see Figure
 1027 19 (right). The tangent space to the fiber $\pi^{-1}(\mu)$ at $\varphi \in \text{Diff}(\Omega)$ and its orthogonal with respect to
 1028 the $L^2(\rho)$ -metric are referred to as the *vertical* and *horizontal* spaces respectively :

$$1029 \text{Ver}_\varphi := \ker d\pi_\varphi, \quad \text{Hor}_\varphi := (\text{Ver}_\varphi)^\perp,$$

1030 where $d\pi_\varphi: T_\varphi \text{Diff}(\Omega) \rightarrow T_{\pi(\varphi)} \text{Prob}(\Omega)$ denotes the differential of π at φ . Moving along vertical
 1031 vectors in $\text{Diff}(\Omega)$ means staying in the same fiber, i.e. projecting always to the same measure μ in
 1032 the bottom space. On the contrary, moving along horizontal vectors means moving orthogonally to
 1033 the fibers, i.e., in the direction that gets fastest away from the fiber. The following proposition gives
 1034 the form of vertical and horizontal vectors.
 1035

1036 **Proposition 6.** *Let $\varphi \in \text{Diff}(\Omega)$. Then*

$$1037 \text{Ver}_\varphi = \{w \circ \varphi, \nabla \cdot (w\mu) = 0\}, \\ 1038 \text{Hor}_\varphi = \{\nabla f \circ \varphi, f \in C^\infty(\Omega)\}.$$

1041 The following results state that line segments and $L^2(\rho)$ -distances in $\text{Diff}(\Omega)$ can be used to compute
 1042 Wasserstein geodesics and distances in the space of probability measures $\text{Prob}(\Omega)$, provided we
 1043 restrict to horizontal displacements.

1044 **Proposition 7.** *The projection $\pi: \text{Diff}(\Omega) \rightarrow \text{Prob}(\Omega)$ is a Riemannian submersion, i.e.
 1045 $d\pi_\varphi: \text{Hor}_\varphi \rightarrow T_{\pi(\varphi)} \text{Prob}(\Omega)$ is an isometry for any $\varphi \in \text{Diff}(\Omega)$.*

1046 This implies the following.

1047 **Proposition 8** (Proposition 2 in main). *Any geodesic $t \mapsto \mu(t)$ for the Wasserstein metric in
 1048 equation 2 is the π -projection of a line segment in $\text{Diff}(\Omega)$ going through a diffeomorphism φ at
 1049 horizontal speed $\nabla f \circ \varphi$ for some smooth function $f \in C(\mathbb{R}^d)$. That is, for t defined in a certain
 1050 interval (t_{\min}, t_{\max}) ,*

$$1051 \mu(t) = \pi(\varphi + t\nabla f \circ \varphi) = (\text{id} + t\nabla f)_\#(\varphi \# \rho). \quad (17)$$

1052 *Another geodesic $\tilde{\mu}(t) = \pi(\varphi + t\nabla \tilde{f} \circ \varphi)$ is orthogonal to $\mu(t)$ at $t = 0$ for the Riemannian metric
 1053 inducing the Wasserstein distance if and only if $\langle \nabla f \circ \varphi, \nabla \tilde{f} \circ \varphi \rangle_{L^2(\rho)} = 0$.*

1054 We comment on the link between this parametrization and McCann's interpolation in Section B.3.

1058 B.2 THE OTTO-WASSERSTEIN GEOMETRY OF GAUSSIAN DISTRIBUTIONS

1059 The Bures-Wasserstein distance in equation 3 on the space S_d^{++} of symmetric positive definite
 1060 (SPD) matrices is the geodesic distance induced by a Riemannian metric g^{BW} , which can be written
 1061 in different ways. Here we use the expression from (Thanwerdas, 2022, Table 4.7), defined for
 1062 $\Sigma = PDP^\top \in S_d^{++}$ and $U = PU'P^\top \in S_d$, by

$$1063 1064 g_\Sigma^{BW}(U, U) = \frac{1}{2} \sum_{1 \leq i, j \leq d} \frac{1}{d_i + d_j} U'_{ij}^2, \quad (18)$$

1065 where the d_i 's are the diagonal elements of D . The associated Riemannian geometry can be described
 1066 by Otto's fiber bundle restricted to the space of centered Gaussian distributions, in the following way.

1067 In this setting, diffeomorphisms are restricted to invertible linear maps $\varphi: u \mapsto Au$ for some invertible
 1068 matrix A , i.e. the space of diffeomorphisms is replaced by the Lie group of invertible matrices GL_d .
 1069 Tangent vectors are then given by linear maps $u \mapsto Xu$ for any matrix $X \in \mathbb{R}^{d \times d}$. Fixing the
 1070 standard normal distribution $\rho = \mathcal{N}(0, \text{Id})$ as reference measure, the L^2 -metric with respect to ρ
 1071 between $u \mapsto Xu$ and $u \mapsto Yu$ is then written, for any $X, Y \in \mathbb{R}^{d \times d}$:

$$1072 1073 \int_{\mathbb{R}^d} \varphi(u)^\top \psi(u) d\rho(u) = \int_{\mathbb{R}^d} \text{tr}(\varphi(u)\psi(u)^\top) d\rho(u) = \text{tr} \left(\int_{\mathbb{R}^d} Xu u^\top Y^\top d\rho(u) \right) = \text{tr}(XY^\top),$$

1074 yielding the standard Frobenius inner product on (the tangent space of) GL_d . We obtain a fibration of
 1075 the top space GL_d over the bottom space S_d^{++} by considering the following projection

$$1076 1077 \pi: GL_d \rightarrow S_d^{++}, \quad A \mapsto AA^\top, \quad (19)$$

1080 see Figure 19 (left). The fiber over $\Sigma \in S_d^{++}$ is
 1081

$$1082 \pi^{-1}(\Sigma) = \{A \in GL_d, AA^\top = \Sigma\} = \Sigma^{1/2}O_d, \quad (20)$$

1083 where O_d denotes the space of orthogonal matrices and $\Sigma^{1/2}$ denotes the only SPD square root of the
 1084 SPD matrix Σ . The differential of the projection $\pi(A) = AA^\top$ is given by
 1085

$$1086 d\pi_A(X) = XA^\top + AX^\top. \quad (21)$$

1087 Therefore, vertical vectors, which are those tangent to the fibers, or equivalently, those belonging to
 1088 the kernel of $d\pi_A(X)$, are given by
 1089

$$\begin{aligned} 1090 \text{Ver}_A &:= \{X \in \mathbb{R}^{d \times d}, XA^\top + AX^\top = 0\} \\ 1091 &= \{X \in \mathbb{R}^{d \times d}, XA^\top \text{ is antisymmetric}\} \\ 1092 &= \{X = K(A^\top)^{-1}, K \in S_d^\perp\} = S_d^\perp(A^\top)^{-1}. \end{aligned}$$

1093 where S_d^\perp denotes the space of antisymmetric matrices of size d . Once again, moving along vertical
 1094 vectors in GL_d means staying in the same fiber, i.e. projecting always to the same SPD matrix in the
 1095 bottom space S_d^{++} . Horizontal vectors are those that are orthogonal to all vertical vectors (for the
 1096 Frobenius metric), i.e. matrices X such that for any antisymmetric matrix K :

$$1097 0 = \langle X, K(A^\top)^{-1} \rangle = \text{tr}(XA^{-1}K^\top)$$

1100 which is equivalent to XA^{-1} symmetric (this can be seen by taking for K the basis elements of S_d^\perp
 1101 in the above equation), yielding

$$\begin{aligned} 1102 \text{Hor}_A &:= \{X \in \mathbb{R}^{d \times d}, (A^\top)^{-1}X^\top = XA^{-1}\} \\ 1103 &= \{X \in \mathbb{R}^{d \times d}, X^\top A - A^\top X = 0\} \\ 1104 &= \{X = KA, K \in S_d\} = S_d A \end{aligned}$$

1105 where S_d denotes the space of symmetric matrices.
 1106

1107 **Proposition 9.** *The projection $\pi: GL_d \rightarrow S_d^{++}$, $A \mapsto AA^\top$ is a Riemannian submersion, i.e. $d\pi_A$
 1108 is an isometry from Hor_A equipped with the Frobenius inner product to $T_{\pi(A)}S_d^{++}$ equipped with the
 1109 inner product $g_{\pi(A)}^{BW}$, for any $A \in GL_d$.*

1110 Just like in the general case, this yields a way to lift the computation of geodesics and distances.
 1111

1112 **Proposition 10** (Propositon 1 in main). *Any geodesic $t \mapsto \Sigma(t)$ in S_d^{++} for the Bures-Wasserstein
 1113 metric in equation 3 is the π -projection of a horizontal line segment in GL_d , that is*

$$1114 \Sigma(t) = \pi(A + tX) = (A + tX)(A + tX)^\top, \quad A \in GL_d, X \in \text{Hor}_A, \quad (22)$$

1115 where t is defined in a certain time interval (t_{\min}, t_{\max}) . Also, the Bures-Wasserstein distance between
 1116 two covariance matrices $\Sigma_1, \Sigma_2 \in S_d^{++}$ is given by the minimal distance between their fibers
 1117

$$1118 1119 BW_2(\Sigma_1, \Sigma_2) = \inf_{Q_1, Q_2 \in O_d} \|\Sigma_1^{1/2}Q_1 - \Sigma_2^{1/2}Q_2\| = \inf_{Q \in SO_d} \|\Sigma_1^{1/2} - \Sigma_2^{1/2}Q\|, \quad (23)$$

1120 where $\|\cdot\|$ is the Frobenius norm and SO_d is the special orthogonal group.
 1121

1122 Formula in equation 22 and the first equality of equation 23 are direct consequences of the fact that
 1123 π is a Riemannian submersion. To obtain the second equality of equation 23, we first notice that
 1124 optimizing on $Q_1, Q_2 \in O_d$ is equivalent to optimizing on a single $Q \in O_d$ thanks to the invariance
 1125 of the Frobenius metric w.r.t. the right action of O_d . And second, that the infimum is attained at (see
 1126 (Bhatia et al., 2019, Equations 3 and 35))

$$1127 1128 Q^* = \Sigma_2^{-1/2}T\Sigma_1^{1/2}, \quad \text{where} \quad T = \Sigma_1^{-1/2}(\Sigma_1^{1/2}\Sigma_2\Sigma_1^{1/2})^{1/2}\Sigma_1^{-1/2}$$

1129 is the Monge map from Σ_1 to Σ_2 (see (Malagò et al., 2018, equation 8)), and so Q^* has positive
 1130 determinant and belongs to SO_d .
 1131

1132 Thus the closest element of the fiber $\pi^{-1}(\Sigma_2)$ to $\Sigma_1^{1/2}$ is given by $\Sigma_2^{1/2}Q^* = T\Sigma_1^{1/2}$, i.e. by left
 1133 multiplying $\Sigma_1^{1/2}$ by the Monge map T . This is more generally true for any representative of Σ_1 :

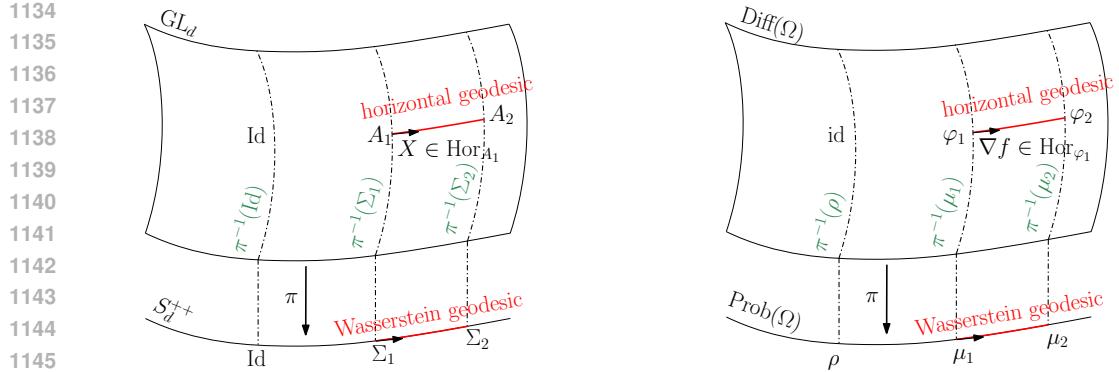


Figure 19: The Otto-Wasserstein geometry of (left) centered Gaussian distributions and (right) a.c. probability distributions. Figures inspired by Khesin et al. (2021).

Proposition 11. Let $\Sigma_1, \Sigma_2 \in S_d^{++}$, T the Monge map from Σ_1 to Σ_2 , $A_1 \in \pi^{-1}(\Sigma_1)$. Then $A_2 := TA_1$ is said to be aligned with respect to A_1 , that is, it is the closest point in $\pi^{-1}(\Sigma_2)$ to A_1 . More precisely, we have

1. $A_2 - A_1 = (T - I)A_1 \in \text{Hor}_{A_1}$
2. $\text{Log}_{\Sigma_1}(\Sigma_2) := d\pi_{A_1}((T - I)A_1) = (T - I)\Sigma_1 + \Sigma_1(T - I)$
3. $BW_2(\Sigma_1, \Sigma_2) = \|\text{Log}_{\Sigma_1} \Sigma_2\|_{\Sigma_1}^{BW} = \|(T - I)A_1\|$

where Log is the Riemannian logarithm map, $\|\cdot\|_{\Sigma}^{BW} = \sqrt{g_{\Sigma}^{BW}(\cdot, \cdot)}$ and $\|\cdot\|$ is the Frobenius norm.

This means that to compute the Bures-Wasserstein distance between two covariance matrices Σ_1 and Σ_2 , one can consider any representative A_1 in the fiber over Σ_1 , compute the representative A_2 of Σ_2 aligned to A_1 (using the Monge map) and finally compute the Frobenius norm of $A_2 - A_1$.

B.3 GEODESIC PARAMETRIZATION

There are two classical parameterizations for Wasserstein geodesics in the space of a.c. probability measures.

McCann's interpolation The first one, due to McCann (1997), is given between two probability distributions μ_0 and μ_1 , and depends on the optimal transport map in equation 2, obtained as the gradient of a convex function u , that is $T_{\mu_0}^{\mu_1} = \nabla u$ and

$$\mu_t = ((1-t)\text{id} + t\nabla u)_{\#}\mu_0 = (\text{id} + t(\nabla u - \text{id}))_{\#}\mu_0, \quad t \in [0, 1]. \quad (24)$$

Otto's geodesic The second one, exploiting Otto's fiber bundle geometry in Otto (2001), consists in writing a geodesic in the Wasserstein space as the projection of a horizontal geodesic in the total space of diffeomorphisms. Such a horizontal geodesic is a line segment going through a diffeomorphism φ with a horizontal speed $\nabla f \circ \varphi$, where f is any smooth function (not necessarily convex). Therefore we get

$$\mu_s = (\varphi + s\nabla f \circ \varphi)_{\#}\rho = (\text{id} + s\nabla f)_{\#}(\varphi_{\#}\rho), \quad s \in (s_0, s_1). \quad (25)$$

In this second expression, the bounds on the time s depends on the function f . Indeed, for μ_s to be a geodesic, $\text{id} + s\nabla f$ needs to remain in the space of diffeomorphisms for a given s , which means that $\text{id} + s\text{Hess } f$ needs to be positive definite. Therefore, we get the following conditions depending on the minimum λ_{\min} and maximum λ_{\max} eigenvalues of $\text{Hess } f$:

$$\begin{cases} s \in (-\infty, -1/\lambda_{\min}) & \text{if } \lambda_{\max} < 0, \\ s \in (-1/\lambda_{\max}, +\infty) & \text{if } \lambda_{\min} > 0, \\ s \in (-1/\lambda_{\max}, -1/\lambda_{\min}) & \text{if } \lambda_{\min} < 0 < \lambda_{\max}. \end{cases} \quad (26)$$

1188 It is clear that equation 24 is a particular case of equation 25, where we choose $\varphi_{\#}\rho = \mu_0$ and
 1189 $\nabla f = \nabla u - \text{id}$. Conversely, one can write equation 25 under the form of equation 24. For a given
 1190 diffeomorphism φ and function f , consider the geodesic given by equation 25, and set $\mu_0 = \varphi_{\#}\rho$.
 1191 Assume that we are in the case where all eigenvalues of $\text{Hess } f$ are negative, then s must be in
 1192 $] -\infty, -1/\lambda_{\min}[$. Consider $s^* \in]0, -1/\lambda_{\min}[$, and define $\mu_1 := \mu_{s^*} = (\text{id} + s^* \nabla f)_{\#}\mu_0$. Setting
 1193 $t = s/s^*$ we have that the geodesic between μ_0 and μ_1 is written

$$\mu_t = (\text{id} + ts^* \nabla f)_{\#}\mu_0 = (\text{id} + t(\nabla u - \text{id}))_{\#}\mu_0, \quad t \in [0, 1].$$

1196 for $u(x) = s^* f + \|x\|^2/2$. Now for any eigenvalue λ_i of H_f the Hessian of f , we have

$$\lambda_i > \lambda_{\min} > -1/s^* \quad \text{i.e.} \quad s^* \lambda_i + 1 > 0.$$

1199 by the interval of definition of s^* . This means that the Hessian $H_u = s^* H_f + \text{id}$ is positive definite,
 1200 which means that u is necessarily convex. The other cases work similarly.

1202 **The Gaussian case** Transposing Otto's formulation in equation 25 to the case of a geodesic between
 1203 Gaussian distributions means that for $A \in GL_d$ and $X \in \text{Hor}_A$ such that $\|X\| = 1$, the interval of
 1204 definition of a geodesic depends on the invertibility of $A + sX$. In turn, the maximal interval of
 1205 definition of $s \in (s_0, s_1)$ is defined from the eigenvalues of XA^{-1} , through the same formula in
 1206 equation 26.

1208 C LINEARIZED OPTIMAL TRANSPORT AND TANGENT PCA

1210 In this section, we provide the definition of linearized Wasserstein distance and details on how to
 1211 perform tangent PCA for both Gaussian distributions and general a.c. distributions. Tangent PCA
 1212 is a widely used approach to compute PCA on the Wasserstein space, that consists in embedding
 1213 probability distributions into the tangent space at some reference measure ρ , and performing PCA in
 1214 the tangent space with respect to the linearized Wasserstein distance.

1216 C.1 THE CASE OF CENTERED GAUSSIAN DISTRIBUTIONS

1218 We consider n covariance matrices $\Sigma_1, \dots, \Sigma_n$ and their Bures-Wasserstein barycenter (or Fréchet
 1219 mean) $\bar{\Sigma}$, that is, the SPD matrix verifying (see Aguech & Carlier (2011)):

$$1220 \bar{\Sigma} = \arg \min_{\Sigma \in S_d^{++}} \sum_{i=1}^n BW_2^2(\Sigma, \Sigma_i). \quad (27)$$

1224 The idea behind tangent PCA is to represent each data point by the corresponding tangent vector,
 1225 given by the Riemannian logarithm map, in the tangent space at the reference point $\bar{\Sigma}$, i.e.

$$1226 \{\text{Log}_{\bar{\Sigma}} \Sigma_i\}_{i=1}^n \subset T_{\bar{\Sigma}} S_d^{++}. \quad (28)$$

1228 Now, one can lift the computations from the tangent space at $\bar{\Sigma}$ to the horizontal space at a point in the
 1229 fiber over $\bar{\Sigma}$, say $A := \bar{\Sigma}^{1/2}$, by aligning all representatives to A , see Proposition 11. The key point
 1230 is that the tangent space at $\bar{\Sigma}$ equipped with the Bures-Wasserstein Riemannian metric is isometric to
 1231 $\text{Hor}_A := S_d A$ equipped with the Frobenius inner product – where we recall that S_d is the space of
 1232 symmetric matrices. This means that instead of performing PCA for the Bures-Wasserstein inner
 1233 product on the tangent vectors in equation 28, we can instead perform linear PCA on their pre-images
 1234 by $d\pi_A$, see Proposition 11:

$$1235 \{(T_i - I)A\}_{i=1}^n \subset \text{Hor}_{A_1}, \quad \text{where} \quad T_i = \Sigma_i^{-1/2}(\Sigma_i^{1/2}\bar{\Sigma}\Sigma_i^{1/2})^{1/2}\Sigma_i^{-1/2}.$$

1237 T_i is the optimal transport map from $\bar{\Sigma}$ to Σ_i , see Section B.2. Now, noticing that

$$1238 \langle K_1 A, K_2 A \rangle = \text{Tr}(K_1 A A^\top K_2^\top) = \text{Tr}(K_1 \bar{\Sigma} K_2^\top), \quad \forall K_1, K_2 \in S_d,$$

1240 we see that the space Hor_A equipped with the Frobenius inner product is itself isometric to S_d
 1241 equipped with the Frobenius inner product weighted by $\bar{\Sigma}$. Therefore, tangent PCA is performed
 through Euclidean PCA on the (centered) vectors $\{T_i - I\}_{i=1}^n$, in the vector space S_d , with respect

to the Frobenius metric weighted by $\bar{\Sigma}$. Another way to see this is by noticing that the linearized Bures-Wasserstein distance $BW_{2,\bar{\Sigma}}$ with respect to $\bar{\Sigma}$ is given by

$$\begin{aligned} BW_{2,\bar{\Sigma}}(\Sigma_1, \Sigma_2) &:= \|\text{Log}_{\bar{\Sigma}} \Sigma_1 - \text{Log}_{\bar{\Sigma}} \Sigma_2\|_{\bar{\Sigma}}^{BW} \\ &= \|d\pi_{\bar{\Sigma}^{1/2}}((T_1 - I)\bar{\Sigma}^{1/2}) - d\pi_{\bar{\Sigma}^{1/2}}((T_2 - I)\bar{\Sigma}^{1/2})\|_{\bar{\Sigma}}^{BW} \\ &= \|(T_1 - I)\bar{\Sigma}^{1/2} - (T_2 - I)\bar{\Sigma}^{1/2}\| \\ &= \|(T_1 - T_2)\bar{\Sigma}^{1/2}\| \end{aligned}$$

where $\|\cdot\|^{BW}$ denotes the norm associated to the Bures Wasserstein Riemannian metric in equation 18, π is Otto's projection in equation 19, and we have used Propositions 9 and 11. Finally,

$$BW_{2,\bar{\Sigma}}(\Sigma_1, \Sigma_2) := \|\text{Log}_{\bar{\Sigma}} \Sigma_1 - \text{Log}_{\bar{\Sigma}} \Sigma_2\|_{\bar{\Sigma}}^{BW} = \|T_1 - T_2\|_{\bar{\Sigma}}, \quad (29)$$

where $\|\cdot\|_{\bar{\Sigma}}$ denotes the Frobenius norm weighted by $\bar{\Sigma}$.

C.2 THE CASE OF A.C. DISTRIBUTIONS

Similarly, one can embed a.c. probability distributions ν_1, \dots, ν_n into the $L^2(\rho)$ space at some a.c. reference measure ρ through the optimal maps $\nu_i \mapsto T_\rho^{\nu_i}$ in the Monge problem in equation 2. Then, the Wasserstein distance can be approximated by the linearized Wasserstein distance in Wang et al. (2013) given by

$$W_{2,\rho}(\nu_1, \nu_2) = \|T_\rho^{\nu_1} - T_\rho^{\nu_2}\|_{L^2(\rho)}. \quad (30)$$

Note that as previously mentioned, this metric induces distortions : while the radial distances from ρ to any μ_i are preserved, that is $\|\text{id} - T_\rho^{\nu_i}\|_{L^2(\rho)} = W_2(\rho, \nu_i)$, other distances are not $\|T_\rho^{\nu_1} - T_\rho^{\nu_2}\|_{L^2(\rho)} \neq W_2(\nu_1, \nu_2)$. A recent paper by Letrouit & Mérigot (2024) proved however, that under some assumptions, $W_{2,\rho}$ is bi-Hölder equivalent to W_2 , which indicates that the distortion effect can be controlled.

Then, denoting $\bar{\nu}_n$ the Wasserstein barycenter as in Aguech & Carlier (2011) of ν_1, \dots, ν_n , that is the solution of

$$\bar{\nu}_n \in \arg \min_{\nu} \sum_{i=1}^n W_2^2(\nu, \nu_i), \quad (31)$$

tangent PCA consists in performing classical PCA, see e.g. Ramsay & Silverman (2002), of $(T_{\bar{\nu}_n}^{\nu_i} - \text{id})_{i=1}^n$ in the Hilbert space $L^2(\bar{\nu}_n)$.

D GEODESIC PCA FOR GAUSSIAN DISTRIBUTIONS

In this section, we present the proofs related to geodesic PCA for Gaussian distributions and the implementation of our algorithm in this case.

D.1 PROOFS RELATED TO GPCA FOR GAUSSIAN DISTRIBUTIONS

We first prove the existence of minimizers for the GPCA problems lifted to Otto's fiber bundle.

Lemma 1. *The GPCA problem in equation 12 for the first component admits a global minimum.*

Proof. First, let us define the set of normalized matrices $\mathbb{B} := \{X \in \mathbb{R}^{d \times d}, \|X\| = 1\}$. By denoting λ_{\min} (resp. λ_{\max}) the smallest (resp. largest) eigenvalue of XA^{-1} , extending the geodesic $t \mapsto A + tX$ as far as possible (see Section B.3) means that the closed interval $[t_{\min}, t_{\max}]$ is defined for some fixed $\varepsilon > 0$ by

$$\left\{ \begin{array}{ll} (-\infty, -1/\lambda_{\min} - \varepsilon] & \text{if } \lambda_{\max} < 0, \\ [-1/\lambda_{\max} + \varepsilon, +\infty) & \text{if } \lambda_{\min} > 0, \\ [-1/\lambda_{\max} + \varepsilon, -1/\lambda_{\min} - \varepsilon] & \text{if } \lambda_{\min} < 0 < \lambda_{\max}. \end{array} \right. \quad (32)$$

1296 Let us now consider the function
 1297

$$1298 F: GL_d \times \mathbb{B} \times (\mathbb{R}^{d \times d})^n \longrightarrow \mathbb{R}$$

$$1299 1300 1301 (A, X, (Q_i)_{i=1}^n) \longmapsto \sum_{i=1}^n \|A + p_{(A,X)}(t_i)X - \Sigma_i^{1/2}Q_i\|^2 =: \sum_{i=1}^n g_i(A, X, Q_i),$$

1302 where $t_i = \langle \Sigma_i^{1/2}Q_i - A, X \rangle$ and $p_{(A,X)}: \mathbb{R} \rightarrow \mathbb{R}$ is the projection operator that clips a point t into
 1303 $[t_{\min}, t_{\max}]$, which depends on A and X . Then the function F is continuous on $GL_d \times \mathbb{B} \times (\mathbb{R}^{d \times d})^n$
 1304 as composition of linear and continuous functions. Note that the function $(A, X) \mapsto p_{(A,X)}(t_i)$
 1305 is continuous by eigenvalue continuity, see Li & Zhang (2019). Additionally, the function F
 1306 is coercive (see e.g. Zalinescu (2002)) on $GL_d \times \mathbb{B} \times (\mathbb{R}^{d \times d})^n$. Indeed, on a diagonal $\{A =$
 1307 $\Sigma_i^{1/2}Q_i, \text{ for } (A, Q_i) \in GL_d \times \mathbb{R}^{d \times d}\}$ for some $i \in \{1, \dots, n\}$, we have $t_i = 0$, and therefore
 1308 we have either $g_i(A, X, Q_i) = 0$ if $p_{(A,X)}(0) = 0$, or $g_i(A, X, Q_i) = \varepsilon \|X\|^2 = \varepsilon$ otherwise.
 1309 This would imply that $g_i(A, X, Q_i)$ doesn't go to infinity when the norm $\|(A, X, Q_i)\| \rightarrow \infty$.
 1310 However, in this case, we have $g_j(A, X, Q_j) \rightarrow \infty$ when $\|(A, X, Q_j)\| \rightarrow \infty$ for any $j \neq i$.
 1311 Moreover, as $p_{(A,X)}(t_i)$ is a clipping, it won't play a role in the coercivity. We conclude by the
 1312 fact that the function $(A, X) \mapsto X^\top A - A^\top X$ is continuous, implying that the set of constraint
 1313 $\{(A, X) \in GL_d \times \mathbb{R}^{d \times d} : X^\top A - A^\top X = 0\}$ is closed and \mathbb{B} and SO_d are compact. The
 1314 optimization problem in equation 12 thus admits a global minimum. \square
 1315

1316 Note that this result also applies for the second component in equation 13 and the higher order
 1317 components.

1318 **Proposition 12** (Proposition 3 in main). *Let $\pi: GL_d \rightarrow S_d^{++}$, $A \mapsto AA^\top$ and $(A_1, X_1, (Q_i)_{i=1}^n)$
 1319 be a solution of*

$$1321 1322 1323 \inf F(A_1, X_1, (Q_i)_{i=1}^n) := \sum_{i=1}^n \|A_1 + p_{A_1, X_1}(t_i)X_1 - \Sigma_i^{1/2}Q_i\|^2,$$

$$1324 \text{subject to } A_1 \in GL_d, X_1 \in \text{Hor}_{A_1}, \|X_1\|^2 = 1, Q_1, \dots, Q_n \in SO_d.$$

1325 Then there exist $t_{\min}, t_{\max} \in \mathbb{R}$ such that the geodesic $\Sigma: t \in [t_{\min}, t_{\max}] \mapsto \pi(A_1 + tX_1)$ in S_d^{++}
 1326 minimizes equation 11.

1328 *Proof.* A horizontal geodesic in GL_d is a straight line going through a base point $A \in GL_d$ in
 1329 the direction of a horizontal vector $X \in \text{Hor}_A$ (that we consider normalized, ie. $\|X\|^2 = 1$), i.e.
 1330 $t \mapsto A + tX \in GL_d$. Denoting $[t_{\min}, t_{\max}]$ the interval constructed in equation 32 which depends on
 1331 the eigenvalues of XA^{-1} , we have that $(\pi(A + tX))_{t \in [t_{\min}, t_{\max}]}$ is a geodesic in the Bures-Wasserstein
 1332 sense, see Proposition 1, and

$$1333 1334 \min_{t \in [t_{\min}, t_{\max}]} BW_2^2(\pi(A + tX), \Sigma_i) = \min_{t \in [t_{\min}, t_{\max}]} \inf_{Q_i \in SO_d} \|A + tX - \Sigma_i^{1/2}Q_i\|^2$$

$$1335 1336 = \inf_{Q_i \in SO_d} \|A + p_{(A,X)}(t_i)X - \Sigma_i^{1/2}Q_i\|^2,$$

1338 where $t_i = \langle \Sigma_i^{1/2}Q_i - A, X \rangle$ is the (orthogonal) projection time of $\Sigma_i^{1/2}Q_i$ onto the line $t \mapsto A + tX$.

1339 We therefore deduce that a set of solution $(A, X, (Q_i)_{i=1}^n)$ of equation 12 defines a proper geodesic
 1340 $(\pi(A + tX))_{t \in [t_{\min}, t_{\max}]}$, solution of problem in equation 11.

1341 \square

1343 **Proposition 13** (Proposition 5 in main). *Let $\nu_i = \mathcal{N}(m_i, \sigma_i^2)$ for $i = 1, \dots, n$ be n univariate
 1344 Gaussian distributions. The first principal geodesic component $t \in [0, 1] \mapsto \mu(t)$ solving equation 1
 1345 remains in the geodesic space of Gaussian distributions for all $t \in [0, 1]$.*

1346 *Proof.* Let $\text{Prob}_2(\mathbb{R})$ be the set of a.c. probability measures on \mathbb{R} that have finite second moment,
 1347 and \mathcal{Q} the set of corresponding quantile functions :

$$1348 1349 \mathcal{Q} = \{F_\nu^{-1}; \nu \in \text{Prob}_2(\mathbb{R})\}$$

\mathcal{Q} is the set of increasing, left-continuous functions $q : (0, 1) \rightarrow \mathbb{R}$, and a convex cone in $L^2([0, 1])$, the set of square-integrable functions on $[0, 1]$. The mapping

$$\Phi : \nu \mapsto F_\nu^{-1} \quad (33)$$

defines an isometry between $\text{Prob}_2(\mathbb{R})$ equipped with the Wasserstein metric, and \mathcal{Q} equipped with the L^2 metric (see e.g. Bigot et al. (2017)), that is, for any $\mu, \nu \in \text{Prob}_2(\mathbb{R})$,

$$W_2(\mu, \nu) = \|F_\mu^{-1} - F_\nu^{-1}\|_{L^2([0,1])}.$$

The map Φ in equation 33 also defines an isometry from the set of (univariate) Gaussian distributions to the set of all Gaussian quantile functions \mathcal{G} . This space \mathcal{G} is the upper-half of the plane \mathcal{F} spanned by the constant function $\mathbf{1}$ and the quantile function F_0^{-1} of the standard normal distribution:

$$\mathcal{G} = \mathbb{R} \cdot \mathbf{1} + \mathbb{R}_+^* \cdot F_0^{-1} \subset \mathcal{F} := \text{span}(\mathbf{1}, F_0^{-1}).$$

Now, consider n normal distributions ν_1, \dots, ν_n , and $(\mu(t))_{t \in [0,1]}$ the first principal geodesic component found by minimizing equation 1, the sum of squared residuals in $\text{Prob}_2(\mathbb{R})$. Since μ is a Wasserstein geodesic in $\text{Prob}_2(\mathbb{R})$ and Φ is an isometry, the curve $t \mapsto \Phi(\mu)(t) = F_{\mu(t)}^{-1}$ is an $L^2([0, 1])$ -geodesic in \mathcal{Q} , i.e. a line segment

$$t \in [0, 1] \mapsto F_{\mu(t)}^{-1} = (1 - t)F_{\mu(0)}^{-1} + tF_{\mu(1)}^{-1}.$$

Since $\{1, F_0^{-1}\}$ forms an orthonormal basis of \mathcal{F} , the orthogonal projection of this line segment on \mathcal{F} is given by

$$t \in [0, 1] \mapsto \langle F_{\mu(t)}^{-1}, \mathbf{1} \rangle \mathbf{1} + \langle F_{\mu(t)}^{-1}, F_0^{-1} \rangle F_0^{-1},$$

which lies in \mathcal{G} . To see this, we need to show that the following value is positive:

$$\langle F_{\mu(t)}^{-1}, F_0^{-1} \rangle = \int_0^1 F_{\mu(t)}^{-1}(y) F_0^{-1}(y) dy = \int_{\mathbb{R}} x F_0^{-1} \circ F_{\mu(t)}(x) d\mu(t)(x) = \mathbb{E}(XT(X)),$$

where $X \sim \mu(t)$ and $T = F_0^{-1} \circ F_{\mu(t)}$ is the Monge map from $\mu(t)$ to the standard normal distribution. Since T is increasing, we indeed have $\mathbb{E}(XT(X)) > 0$ (see e.g. the proof of Theorem 2.2 in Schmidt (2014)).

Finally, since $\Phi(\mu)$ orthogonally projects from \mathcal{Q} to \mathcal{G} w.r.t the L^2 metric and Φ defines an isometry, we get that the geodesic μ orthogonally projects to a geodesic $\pi(\mu)$ in the space of Gaussian distributions, w.r.t. the Wasserstein metric. By the distance minimizing property of orthogonal projections, we know that the cost function in equation 1 evaluated at $\pi(\mu)$ is no larger than its value at μ . Since μ is optimal, we get that $\mu = \pi(\mu)$ and μ belongs to the space of Gaussian distributions. \square

Proposition 14. *Let Σ_1, Σ_2 two SPD matrices that are diagonalizable in the same orthonormal basis, i.e.*

$$\Sigma_1 = P \begin{pmatrix} a_1^2 & 0 \\ 0 & b_1^2 \end{pmatrix} P^\top \quad \text{and} \quad \Sigma_2 = P \begin{pmatrix} a_2^2 & 0 \\ 0 & b_2^2 \end{pmatrix} P^\top,$$

where P is orthogonal. Then $BW_2^2(\Sigma_1, \Sigma_2) = (a_1 - a_2)^2 + (b_1 - b_2)^2$, and thus the Bures-Wasserstein geodesic between Σ_1 and Σ_2 is given by

$$\Sigma(t) = P \begin{pmatrix} ((1-t)a_1 + tb_1)^2 & 0 \\ 0 & ((1-t)a_2 + tb_2)^2 \end{pmatrix} P^\top, \quad 0 \leq t \leq 1.$$

Proof. This is a straightforward computation using equation 3. \square

Proposition 15. *Let us consider $n = 2p$ covariance matrices $\Sigma_i = \Sigma(a, b, \theta_i)$ as defined in equation 16, where $\theta_i = i\pi/n$ for $i = 0, \dots, n-1$. Then, the Bures-Wasserstein barycenter in equation 27 of these covariance matrices is given by $\bar{\Sigma} = (a+b)^2/4I$.*

Proof. Each pair of covariance matrices

$$\Sigma_i = P_{\theta_i} \begin{pmatrix} a^2 & 0 \\ 0 & b^2 \end{pmatrix} P_{\theta_i}^\top, \quad \text{and} \quad \Sigma_{i+p} = P_{\theta_i + \pi/2} D P_{\theta_i + \pi/2}^\top = P_{\theta_i} \begin{pmatrix} b^2 & 0 \\ 0 & a^2 \end{pmatrix} P_{\theta_i}^\top$$

1404 are diagonalizable in the same basis, and so by Proposition 14, the geodesic from Σ_i to Σ_{i+p} is
 1405

$$1406 \quad \Sigma(t) = P_{\theta_i} \begin{pmatrix} ((1-t)a+tb)^2 & 0 \\ 0 & ((1-t)b+ta)^2 \end{pmatrix} P_{\theta_i}^\top, \quad 0 \leq t \leq 1.$$

1408 In particular, the Fréchet mean is given by $\bar{\Sigma} = \Sigma(1/2) = ((a+b)/2)^2 I$. Since each pair of
 1409 covariance matrices has the same Fréchet mean, the Fréchet mean of the whole set $\Sigma_1, \dots, \Sigma_n$ is
 1410 also given by $\bar{\Sigma}$. \square
 1411

1412 **Proposition 16** (Proposition 4 in main). *Let $\Sigma \in S_2^{++}$ with eigenvalues a^2, b^2 and $\Sigma' = P_\theta \Sigma P_\theta^\top$
 1413 where P_θ is the rotation matrix of angle θ . Then, denoting $\bar{\Sigma} = ((a+b)/2)^2 I$ we have*

$$1414 \quad \frac{BW_2^2(\Sigma, \Sigma')}{BW_{2,\bar{\Sigma}}^2(\Sigma, \Sigma')} = 1 - \left(\frac{a-b}{a+b} \right)^2 \cos^2 \theta + O((a-b)^4). \quad (34)$$

1417 *Proof.* Recall that the linearized Bures-Wasserstein distance at $\bar{\Sigma}$ between Σ and Σ' is given by the
 1418 distance between their images by the Riemannian logarithm map $U := \text{Log}_{\bar{\Sigma}} \Sigma$ and $U' := \text{Log}_{\bar{\Sigma}} \Sigma'$
 1419 in the tangent space at $\bar{\Sigma}$, i.e.
 1420

$$1421 \quad BW_{2,\bar{\Sigma}}(\Sigma, \Sigma') = \|U - U'\|_{\bar{\Sigma}}^{BW},$$

1422 where $\|\cdot\|^{BW}$ denotes the norm associated to the Bures-Wasserstein Riemannian metric in equation 18.
 1423 As in any Riemannian manifold, the true geodesic distance can be approximated by this linearized
 1424 distance in the tangent space, corrected by the curvature (see e.g. Lemma 1 in Harms et al. (2019)):
 1425

$$1426 \quad BW_2^2(\Sigma, \Sigma') = (\|U - U'\|_{\bar{\Sigma}}^{BW})^2 - \frac{1}{3} R_{\bar{\Sigma}}(U, U', U, U') + O(\|U\|_{\bar{\Sigma}}^{BW} + \|U'\|_{\bar{\Sigma}}^{BW})^6, \quad (35)$$

1427 where $R_{\bar{\Sigma}}$ is the curvature tensor.
 1428

1429 Recall from equation 18 that the Bures-Wasserstein norm of a vector U is expressed in an eigenvector
 1430 basis of the base point, here $\bar{\Sigma}$. Since any basis is an eigenvector basis of $\bar{\Sigma}$, it is convenient to
 1431 choose that of Σ , which we can assume without loss of generality to be the canonical basis. Thus
 1432 we write $\Sigma = D$ where $D = \text{diag}(a^2, b^2)$ and $\Sigma' = P_\theta D P_\theta^\top$, and the norm associated to the
 1433 Bures-Wasserstein Riemannian metric is given by

$$1434 \quad \|U\|_{\bar{\Sigma}}^{BW} = \frac{1}{2} \sum_{1 \leq i, j \leq 2} \frac{1}{d_i + d_j} U_{ij}^2$$

1437 where the d_i 's are the eigenvalues of $\bar{\Sigma}$, given here by $d_1 = d_2 = ((a+b)/2)^2$. From Proposition 11
 1438 we have

$$1439 \quad U := \text{Log}_{\bar{\Sigma}} \Sigma = (T - I)\bar{\Sigma} + \bar{\Sigma}(T - I),$$

$$1440 \quad U' := \text{Log}_{\bar{\Sigma}} \Sigma' = (T' - I)\bar{\Sigma} + \bar{\Sigma}(T' - I),$$

1442 where

$$1443 \quad T := \bar{\Sigma}^{-1/2} (\bar{\Sigma}^{1/2} \Sigma \bar{\Sigma}^{1/2})^{1/2} \bar{\Sigma}^{-1/2} = \frac{2}{a+b} D^{1/2},$$

$$1445 \quad T' := \bar{\Sigma}^{-1/2} (\bar{\Sigma}^{1/2} \Sigma' \bar{\Sigma}^{1/2})^{1/2} \bar{\Sigma}^{-1/2} = \frac{2}{a+b} P_\theta D^{1/2} P_\theta^\top,$$

1447 and easily get

$$1448 \quad U = \frac{a^2 - b^2}{2} J, \quad U' = \frac{a^2 - b^2}{2} P_\theta J P_\theta^\top, \quad \text{where} \quad P_\theta J P_\theta^\top = \begin{pmatrix} \cos 2\theta & \sin 2\theta \\ \sin 2\theta & -\cos 2\theta \end{pmatrix}$$

1451 and $J = \text{diag}(1, -1)$. Thus after some computations we obtain

$$1452 \quad \|U\|_{\bar{\Sigma}}^{BW} = \|U'\|_{\bar{\Sigma}}^{BW} = |a-b|/\sqrt{2},$$

$$1453 \quad BW_{2,\bar{\Sigma}}(\Sigma, \Sigma') = \|U - U'\|_{\bar{\Sigma}}^{BW} = \sqrt{2}|(a-b)\sin\theta|. \quad (36)$$

1455 To compute the curvature tensor, we use the following formula from (Thanwerdas, 2022, Table 4.7)

$$1456 \quad R_{\bar{\Sigma}}(U, U', U, U') = \frac{3}{2} \sum_{i,j} \frac{d_i d_j}{d_i + d_j} [U_0, U'_0]_{ij}^2$$

1458 where $[A, B] = AB - BA$ is the Lie bracket of matrices, U_0 and U'_0 are the only symmetric matrices
 1459 verifying the Sylvester equations $U = U_0\bar{\Sigma} + \bar{\Sigma}U_0$ and $U' = U'_0\bar{\Sigma} + \bar{\Sigma}U'_0$ respectively. Since $\bar{\Sigma}$ is a
 1460 multiple of the identity, we easily get

$$1462 \quad U_0 = \frac{a-b}{a+b}J, \quad U'_0 = \frac{a-b}{a+b}P_\theta J P_\theta^\top$$

1464 and straightforward computations yield

$$1466 \quad R_{\bar{\Sigma}}(U, U', U, U') = \frac{3}{2} \frac{(a-b)^4}{(a+b)^2} \sin^2 2\theta. \quad (37)$$

1468 Finally, putting together equation 35, equation 36 and equation 37 and we obtain

$$1470 \quad BW_2^2(\Sigma, \Sigma') = BW_{2,\bar{\Sigma}}^2(\Sigma, \Sigma') - 2 \frac{(a-b)^4}{(a+b)^2} \sin^2 \theta \cos^2 \theta + O((a-b)^6),$$

1472 and dividing by the squared linearized optimal transport distance yields the desired result. \square

1474 D.2 IMPLEMENTATION OF GPCA FOR GAUSSIAN DISTRIBUTIONS

1476 As described in Section 3, the first and second components of geodesic PCA are respectively found
 1477 by solving the minimization problems in equation 12 and equation 13. The geodesic components are
 1478 given by

$$1479 \quad \Sigma_i(t) = (A_i + tX_i)(A_i + tX_i)^\top, \quad \text{for } i = 1, 2,$$

1480 where $A_1 \in GL_d$ and $X_1 \in \text{Hor}_{A_1}$ are minimizers of equation 12, and $A_2 \in GL_d$ and $X_2 \in \text{Hor}_{A_2}$
 1481 minimizers of equation 13. The matrix $\pi(A_2)$ is the crossing point through which all geodesic
 1482 components intersect, see Figure 2. The higher order components are found in a analogous way:
 1483 for the k -th component, we search for a horizontal segment $t \mapsto A_k + tX_k$ where A_k belongs to
 1484 the fiber over the intersection point (we parametrize it w.r.t. the previous position in the fiber, i.e.
 1485 $A_k = A_{k-1}R_{k-1}$ for a certain $R_{k-1} \in SO_d$) and the horizontal velocity vector X_k is orthogonal to
 1486 the lifts of the velocity vectors of the previous component. Thus, the k -th component, $k \geq 3$, solves:

$$1487 \quad \inf F(A_k, X_k, (Q_i)_{i=1}^n)$$

$$1488 \quad \text{subject to } A_k = A_{k-1}R_{k-1}, R_{k-1} \in SO_d, X_k \in \text{Hor}_{A_k}, \|X_k\|^2 = 1,$$

$$1489 \quad \langle X_k, X_{k-\ell}R_{k-\ell} \dots R_{k-1} \rangle = 0, 1 \leq \ell \leq k-1, Q_1, \dots, Q_n \in SO_d.$$

1491 Following Huckemann et al. (2010) and Calissano et al. (2024), we propose an iterative algorithm to
 1492 implement these components, that, for each component, alternates two steps:

1493 (Step 1) minimization of the objective function F (see equation 12) with respect to $(Q_i)_{i=1}^n$ for fixed
 1494 (A, X) ,
 1496 (Step 2) minimization of the objective function F with respect to (A, X) for fixed $(Q_i)_{i=1}^n$.

1497 In dimension $d = 2$, any rotation matrix Q can be parametrized by a scalar angle θ and both steps
 1498 are solved using the Sequential Least Squares Programming (SLSQP) algorithm (see e.g. Ma et al.
 1499 (2024)) available on the *scipy python library* and given by Virtanen et al. (2020). In higher dimension,
 1500 each minimization with respect to a rotation matrix is performed using Riemannian gradient descent
 1501 on SO_d , relying on the Riemannian geometry of SO_d induced by the standard Frobenius metric of
 1502 the ambient space $\mathbb{R}^{d \times d}$. In particular we use the exponential map implemented in the *Python library*
 1503 *geomstats* developed by Miolane et al. (2020). More details on the Riemannian geometry of SO_d
 1504 and the Riemannian gradient descent procedure can be found e.g. in (Boumal, 2023, Sections 7.4 and
 1505 4.3).

1506 Unfortunately, we cannot ensure the convergence of the iterates of the proposed block alternating
 1507 algorithm, as classical arguments require uniqueness of the minimizer at each iterations as proven
 1508 in Powell (1973). This is unachievable in our problem: the line with base point A and direction
 1509 $X \in \text{Hor}_A$ and the line with base point AQ and direction $XQ \in \text{Hor}_{AQ}$ for $Q \in O_d$ project onto the
 1510 same geodesic in the bottom space. However, regarding (Step 1), and thanks to Theorem 3.7 in Huang
 1511 & Wei (2022), we have for fixed (A, X) that the cost function $f: (Q_1, \dots, Q_n) \mapsto F(A, X, (Q_i)_{i=1}^n)$
 has the Riemannian Kurdyka-Łojasiewicz property at any point of $(O_d)^n$. Finally, we have the

convergence of the iterates towards an accumulation point thanks to Theorem 3.14 in Zhou et al. (2024). The three assumptions in this theorem are verified in our case : Assumption (3.5) (L -Retraction Smoothness) is obtained because $\text{grad } f$ is Lipschitz, and Corollary 10.54 in Boumal (2023); Assumption (3.7) (bounded from below) directly holds because $f \geq 0$; Assumption (3.8) (ndividual Retraction Lipschitzness) is verified thanks to Corollary 10.47 in Boumal (2023).

Scalability of the algorithm Surely, the computational time of our algorithm for Gaussian distributions will increase with the dimension. However, the algorithm can be made less sensitive to the number of input covariance matrices by parallelizing (Step 2) of our algorithm, which consists in updating the orthogonal matrices $(Q_i)_{i=1}^n$. This would significantly reduce the overall computational cost of the algorithm. Also, we currently use the `scipy` toolbox to solve (Step 1), which could also be accelerated using a more powerful optimization toolbox.

E HYPERPARAMETERS

E.1 HYPERPARAMETERS SETTING

Hyperparameter	Value
f_ψ architecture	dense MLP $d \rightarrow 128 \rightarrow 128 \rightarrow 128 \rightarrow 128 \rightarrow 1$ ELU activation functions
f_ψ optimizer	Adam step size = 0.0005 $\beta_1 = 0.9$ $\beta_2 = 0.999$
φ_θ architecture	dense MLP $d \rightarrow 128 \rightarrow 128 \rightarrow 128 \rightarrow 128 \rightarrow d$ RELU activation functions
φ_θ optimizer	Adam step size = 0.0005 $\beta_1 = 0.9$ $\beta_2 = 0.999$
t_i optimizer	Adam step size = 0.001 $\beta_1 = 0.9$ $\beta_2 = 0.999$
batch size	1024
number of gradient steps first component	120,000
number of gradient steps second component	200,000
$\lambda_{\mathcal{O}}$	1.0
$\lambda_{\mathcal{I}}$	1.0

Table 1: Hyperparameters used across all experiments.

All experiments were conducted on a single V100 GPU with 32GB of memory, using a shared set of hyperparameters detailed in Table 1. The same hyperparameters are used for computing both the first and second geodesic components, except for the number of gradient steps (see Table 1), which is increased for the second component. This is likely due to the additional complexity introduced by the intersection and orthogonality constraints enforced through regularization. Both f_ψ and φ_θ are implemented as standard multilayer perceptrons (MLPs) with four hidden layers of width 128. We use ELU activation functions in f_ψ because its gradient is used to parameterize a transport map in our formulation, and ELUs are commonly employed in such settings. The Sinkhorn divergence S_ε is used in the loss function as a surrogate for the squared Wasserstein distance to compute the geodesic components. The regularization parameter ε must be adapted to the scale of the data; we set it as $\varepsilon = 0.01 \mathbb{E}_{x,x' \sim \nu_i} \|x - x'\|^2$, where the expectation is approximated via Monte Carlo using the

1566 current minibatch samples. Note that setting ε this way is the default configuration in the OTT-JAX
 1567 library. For computing the second geodesic component, we fix the regularization coefficients λ_O and
 1568 λ_I to 1.0, which we found to be robust across all experiments. While increasing them (e.g., to 10.0)
 1569 typically yields similar results, excessively large values may degrade performance. Conversely, if
 1570 these regularization terms are too small, the algorithm tends to recover the first component as the
 1571 second, due to its lower cost. In practice, we monitor the regularization terms during optimization
 1572 to ensure they decrease sufficiently relative to their initial values, confirming that the optimization
 1573 effectively optimize the intersection and orthogonality constraints. To determine the hyperparameters
 1574 in Table 1, we performed a grid search over the optimizer learning rate for the t_i in $10^{-4}, 10^{-3}, 10^{-2}$,
 1575 and over the regularization coefficients λ_O and λ_I in 0.1, 1.0, 10.0, 100.0. We found that setting both
 1576 regularization terms to 1.0 consistently yielded good performance across all experiments, [see Section](#)
 1577 [E.2](#).

1578 **Note on φ parameterization.** Note that although φ is theoretically required to be a diffeomorphism
 1579 in Otto’s parameterization of geodesics (equation 9), we parameterize it using a simple MLP. Initially,
 1580 we experimented with normalizing flows to ensure invertibility, but observed that a standard MLP
 1581 yielded similar results. In Otto’s geodesic framework, φ serves to modify the reference measure ρ
 1582 and define the measure at $t = 0$ along the geodesic. If φ is not a diffeomorphism and the pushforward
 1583 $\varphi_\# \rho$ is not absolutely continuous, the resulting geodesic becomes degenerate, which may hinder
 1584 optimization of the loss equation in equation 1. In practice, however, we found that the MLP φ_θ
 1585 reliably produces absolutely continuous measures, which is sufficient for our method.

1587 E.2 IMPACT OF THE REGULARIZATIONS ON GPCA

1589 For the estimation of the second GPCA component, we introduce two regularization terms,
 1590 $\mathcal{I}(\mu_{\theta_1, \psi}, \mu_{\theta_2, \psi_2}, t_{\text{inter}}^1, t_{\text{inter}}^2)$ and $\mathcal{O}(\nabla f_\psi(\varphi_\theta), \nabla f_{\psi_2}(\varphi_{\theta_2}))$, with their associated regularization co-
 1591 efficients λ_I and λ_O . The first term enforces that the two components intersect, while the second
 1592 ensures that the components remain orthogonal. Experimentally, we observe that setting both co-
 1593 efficients to $\lambda_I = \lambda_O = 1.0$ robustly enforces these constraints across all experiments while still
 1594 producing meaningful principal components. Conversely, if these regularization terms are too small,
 1595 the algorithm tends to recover the first component as the second, at it gives the lowest cost. In
 1596 practice, we monitor the regularization terms during optimization to ensure they decrease sufficiently
 1597 relative to their initial values. This permits to confirm that the optimization effectively optimize the
 1598 intersection and orthogonality constraints. This section aims at quantifying the impact of the two
 1599 regularizing coefficients λ_I and λ_O on the computed geodesics. We focus on the 3D point-cloud
 1600 experiments with lamps.

1601 E.2.1 ORTHOGONALITY REGULARIZATION

1603 In this part, we set the regularization term λ_I to 1.0 and compute GPCA for different values of
 1604 λ_O . The resulting second component is shown in Figure 20. The GPCA cost of this component, as
 1605 defined in equation 15, together with the quantity measuring the orthogonality between components,
 1606 $\mathcal{O}(\nabla f_\psi(\varphi_\theta), \nabla f_{\psi_2}(\varphi_{\theta_2}))$, are reported in Table 2. The quantities reported in Table 2 are estimated
 1607 on batches of size 2048. The variance is computed over 100 runs for the orthogonality measure
 1608 and 5 runs for the GPCA cost. Note that each run of the orthogonality estimation already involves
 1609 computing 100 Wasserstein distances, since we have 100 point clouds.

λ_O	Orthogonality: $\mathcal{O}(\nabla f_\psi(\varphi_\theta), \nabla f_{\psi_2}(\varphi_{\theta_2}))$	GPCA cost (second component)
0.001	0.823 ± 0.007	3.22 ± 0.007
0.01	0.722 ± 0.008	3.30 ± 0.01
0.1	$2.72 \times 10^{-3} \pm 7 \times 10^{-5}$	5.08 ± 0.01
1.0	$1.12 \times 10^{-4} \pm 0.07 \times 10^{-4}$	5.13 ± 0.02
10.0	$8.6 \times 10^{-6} \pm 5 \times 10^{-7}$	5.36 ± 0.02
100.0	$4.8 \times 10^{-6} \pm 3 \times 10^{-7}$	5.45 ± 0.02

1618 **Table 2: Orthogonality regularization value and second-component loss for different values of λ_O .**

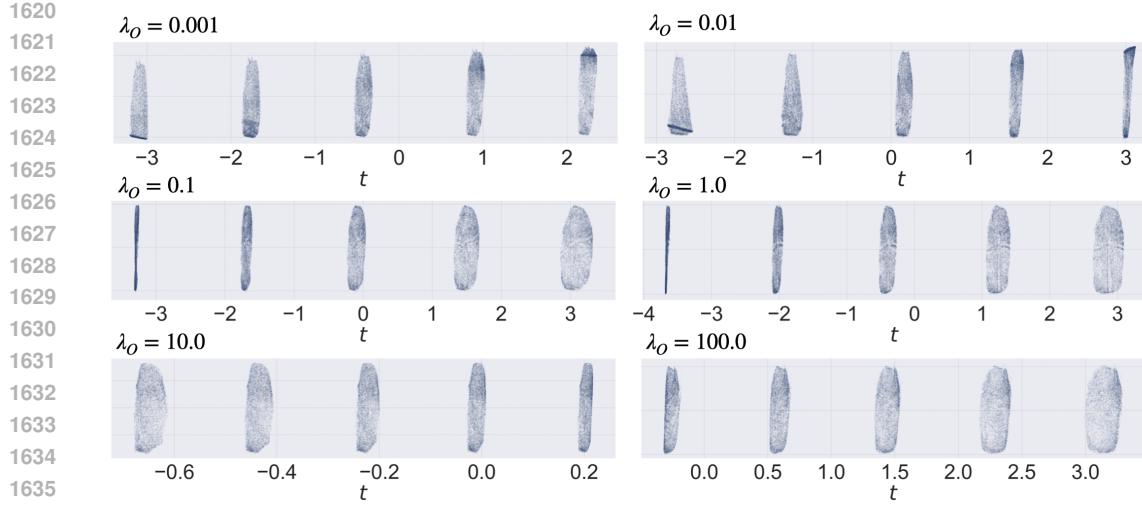


Figure 20: Empirical distributions sampled uniformly along the geodesics associated with the second GPCA principal component for different values of the regularization coefficient λ_O . In all experiments, the other regularization coefficient is fixed at $\lambda_I = 1.0$.

Note that the GPCA cost of the second component should be compared with that of the first component, which is 3.295 ± 0.006 . Table 2 shows that for low values of λ_O (i.e., 0.001 and 0.01), the orthogonality quantity is large, and the recovered "second" component is in fact identical to the first component, as illustrated in Figure 20. This is also reflected in the GPCA cost (see Table 2), which matches the one of the first component. For higher values of λ_O (0.1, 1.0, 10.0, 100.0), the algorithm successfully recovers a distinct second component.

E.2.2 REGULARIZATION ON THE INTERSECTION OF THE GEODESICS

λ_I	Intersection: $\mathcal{I}(\mu_{\theta,\psi}, \mu_{\theta_2,\psi_2}, t_{\text{inter}}^1, t_{\text{inter}}^2)$	GPCA cost (second component)
0.001	$1.9 \times 10^{-3} \pm 2 \times 10^{-4}$	5.16 ± 0.01
1.0	$1.3 \times 10^{-3} \pm 1 \times 10^{-4}$	5.13 ± 0.02
10.0	$1.19 \times 10^{-4} \pm 2 \times 10^{-6}$	5.18 ± 0.02
100.0	$1.48 \times 10^{-5} \pm 5 \times 10^{-7}$	5.08 ± 0.01

Table 3: Wasserstein distance between $\mu_1(t_{\text{inter}}^1)$ and $\mu_2(t_{\text{inter}}^2)$ and second-component loss for different values of λ_I .

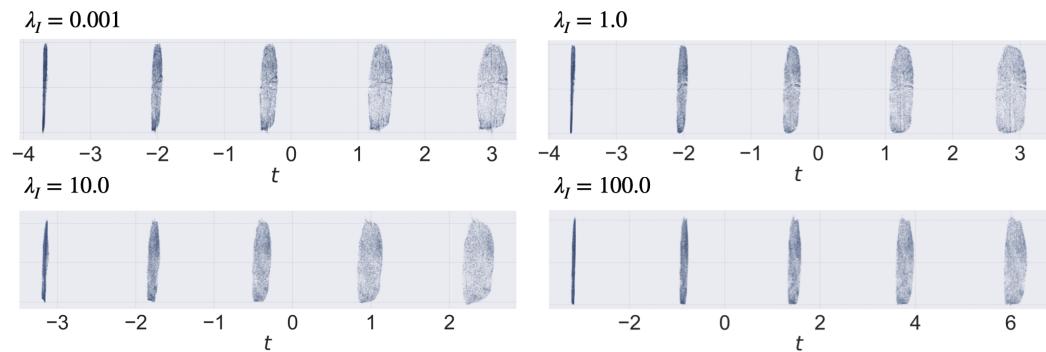


Figure 21: Empirical distributions sampled uniformly along the geodesics associated with the second GPCA principal component for different values of the regularization coefficient λ_I . In all experiments, the other regularization coefficient is fixed at $\lambda_O = 1.0$.

1674 In this part, we set the regularization term λ_O to 1.0 and compute GPCA for different values of λ_I .
 1675 The second component is displayed in Figure 21; the GPCA cost of this component, as well as the
 1676 quantity measuring the intersection of the components, $\mathcal{I}(\mu_{\theta,\psi}, \mu_{\theta_2,\psi_2}, t_{\text{inter}}^1, t_{\text{inter}}^2)$, are reported in
 1677 Table 3. The quantities reported in Table 3 are estimated on batches of size 2048. The variance is
 1678 computed over 100 runs for the intersection measure and 5 runs for the GPCA cost.

1679 We observe from the recovered geodesics in Figure 21 that this regularization term plays a less
 1680 significant role than the orthogonality term. Moreover, Table 3 shows that increasing λ_I does not
 1681 affect negatively the GPCA cost of the recovered component.

1683 E.2.3 SCALABILITY OF OUR GPCAGEN ALGORITHM

1684 For general distributions, there are two types of “scaling” that can affect the algorithm:

- 1687 1. Number of probability measures (n): The number of measures ν_i directly determines the
 1688 iterations of the inner loop in Algorithm 1 (line 3). Consequently, the training time scales
 1689 linearly with n .
- 1690 2. Dimension of the space (d): As the dimension of the space in which the ν_i lies increases, the
 1691 main challenge consists in accurately estimating the maximum and minimum eigenvalues
 1692 that the Hessian of f can take. As discussed with reviewer oUMT, in high dimensions, it
 1693 becomes necessary to use algorithms that avoid computing the full Hessian and instead rely
 1694 on matrix-vector products, such as the LOBPCG algorithm Duersch et al. (2018). Furthermore,
 1695 rather than relying solely on the samples in the training batch, an adversarial approach
 1696 would be needed to track the eigenvectors corresponding to the worst-case eigenvalues.

1697 F USE OF LARGE LANGUAGE MODELS (LLMs)

1700 LLMs were used only to assist with polishing the writing; all research ideas, experiments, and
 1701 analyses were conducted independently by the authors.