



CENTRE FOR **STOCHASTIC GEOMETRY**
AND ADVANCED **BIOIMAGING**



Stefan Sommer

An Infinitesimal Probabilistic Model for Principal Component Analysis of Manifold Valued Data

No. 02, February 2018

An Infinitesimal Probabilistic Model for Principal Component Analysis of Manifold Valued Data

Stefan Sommer

Department of Computer Science (DIKU), University of Copenhagen

Abstract

We provide a probabilistic and infinitesimal view of how the principal component analysis procedure (PCA) can be generalized to analysis of nonlinear manifold valued data. Starting with the probabilistic PCA interpretation of the Euclidean PCA procedure, we show how PCA can be generalized to manifolds in an intrinsic way that does not resort to linearization of the data space. The underlying probability model is constructed by mapping a Euclidean stochastic process to the manifold using stochastic development of Euclidean semimartingales. The construction uses a connection and bundles of covariant tensors to allow global transport of principal eigenvectors, and the model is thereby an example of how principal fiber bundles can be used to handle the lack of global coordinate system and orientations that characterizes manifold valued statistics. We show how curvature implies non-integrability of the equivalent of Euclidean principal subspaces, and how the stochastic flows provide an alternative to explicit construction of such subspaces. We describe two estimation procedures for inference of parameters and prediction of principal components, and we give examples of properties of the model on surfaces embedded in \mathbb{R}^3 .

Keywords: principal component analysis, manifold valued statistics, stochastic development, probabilistic PCA, anisotropic normal distributions, frame bundle

1 Introduction

A central problem in the formulation of statistical methods for analysis of data in nonlinear spaces is the lack of global coordinate systems and global orientation fields. As an example, consider generalizing the notion of covariance matrix to manifold valued random variables: While the Euclidean definition takes the expectation $E[X^i - E[X]^i]E[X^j - E[X]^j]$ of the product of the coordinate components X^i of the centered random variable $X - E[X]$, the coordinate components are not meaningful in the nonlinear situation as the coordinates themselves are not defined. This fact

fundamentally questions what constitutes a natural generalization of covariance. As a second example, consider a standard Euclidean linear latent variable model

$$y = m + Wx + \epsilon \tag{1.1}$$

on \mathbb{R}^d with mean m , coefficient matrix W , latent variables x , and noise ϵ . The columns of W can be seen as encoding the direction in the Euclidean space connected to a change of each element of x . However, on a manifold M , W has a priori only meaning for infinitesimal changes dx in the tangent space $T_m M$, and the lack of global orientation prevents a direct translation between such infinitesimal changes, finite perturbations of m , and global directions on M .

The aim of this paper is to construct a nonlinear manifold generalization of the inherently linear principal component analysis (PCA) procedure, a generalization that is intrinsically based on the geometry of the manifold M and does not resort to a linear approximation of the geometry. The model is based on the Euclidean probabilistic principal component analysis procedure (PPCA, [26]) that interprets PCA as a latent variable model (1.1) with W having low rank $k \leq d$. We use PPCA with a probability model based on a notion of infinitesimal covariance and thereby avoid linearizing the nonlinear data space while incorporating the effect of data anisotropy, here difference in the principal eigenvalues, intrinsically. The model is related to the probabilistic principal geodesic analysis (PPGA, [27]) procedure, however using the probability model and normal distributions defined in [20, 24]. This construction in particular emphasizes the role of the connection on the manifold in linking infinitesimally close tangent spaces. We refer to the method as being *infinitesimal probabilistic* because the connection allows sequences of random, infinitesimal steps to generate the data probability model.

As a second aim, we wish to exemplify how the use of fiber bundle structures provides a way around the lack of coordinates and global orientations on M . The construction in [20, 24] essentially enlarges the manifold by equipping it with a structure group at each point and hence a principal fiber bundle structure. An example of this is the frame bundle FM , viewed as the bundle of invertible linear maps $\text{GL}(\mathbb{R}^d, TM)$, but we will also encounter lower-rank versions $F^k M$ of FM , and the quotient bundle Sym^+ of symmetric positive tensors on TM . Elements of these bundles are here used to model the local anisotropy and covariance of the data. The incorporation of nontrivial covariance couples with curvature leading to families of paths that extend geodesics as being, in a certain sense, most probable paths between data points [22].

The probabilistic construction naturally gives a formulation of inference procedures as maximum likelihood or maximum a posteriori fits to data. Using intrinsically defined probability distributions on the manifold thereby avoids some of the complexities that makes non-probabilistic parametric constructions on manifolds inherently complex to define. For regression, a similar approach has been pursued in [13]. The present paper is partly based on and extends the Oberwolfach abstract [19]. While the analogy to PPCA is mentioned in [20], the focus of that paper is on defining normal-like distributions and not to generalize PPCA as is the focus here.

1.1 Outline

We start by a short review of related work on generalizing PCA to manifolds. This is followed by a description of the fiber bundle geometry necessary for representing and transporting orientations structures over the manifold. We then define and discuss the generalized probabilistic principal component analysis procedure before outlining inference methods. The paper ends with simple numerical experiments and concluding remarks.

2 Background

Extending Euclidean statistical notions, tools, and inference procedures to the nonlinear manifold situation has been treated in multiple works in recent literature. We focus here on PCA-like statistical analysis of data y_1, \dots, y_N , $y_i \in M$ with M being a nonlinear manifold with a priori known structure, for example arising directly from the data, e.g. angular measurements or position measurements on the surface of the earth, or from modeling constraints. We assume the dimension d of M is finite. Note that the setting is different than manifold learning where the objective is to infer the manifold structure from the data. Manifolds lack vector space structure and therefore also a global coordinate system and globally consistent orientations. Instead of inner product structure on Euclidean vectors, we instead often assume the existence of a Riemannian metric g that defines local inner products on infinitesimal variations, vectors in the tangent bundle TM . For the construction in focus in this paper, we however mainly need a connection ∇ and a fixed base measure μ_0 . If M has a Riemannian metric, ∇ can be the Levi-Civita connection of g and μ_0 the Riemannian volume form vol_g .

2.1 Parametric Subspace Constructions

Perhaps the most immediate way to handle the lack of coordinate system is to use tangent spaces to linearize the manifold and thereby implicitly define a local sense of linear coordinate system on the nonlinear space. This approach is used for generalizations of the principal component analysis procedure in tangent space PCA (tPCA). The principal geodesic analysis (PGA, [5]) procedure also uses a tangent space linearization but minimizes the residual distances to the data using manifold distances induced from a Riemannian metric. The central idea in tangent space based procedures is to find a suitable zero-dimensional representation m of the data, often a Frechét mean [6], and subsequently map the data from the manifold to the linear tangent space $T_m M$. Given a Riemannian structure on M , this can be achieved from the geodesic endpoint map Exp_m and its locally defined inverse Log_m . This construction is however not faithful to the geometry: The effective linearization of the manifold is only locally around m a proper view of the geometry as encoded in the Riemannian metric. The curvature of the manifold will distort the linearized view of the data when significant data mass is observed far from m . The fact that the linear view is only one-to-one up to the cut locus of m further emphasizes the approximation in the tangent space linearization.

When using tangent PCA or similar tangent-space based procedures, principal subspaces found as linear subspaces of $T_m M$ are projected to subspaces of M using Exp_m , i.e. as sprays of geodesics originating at m . Such subspaces are generally only geodesic at m itself unlike the Euclidean situation where a linear subspace always contains straight lines between all of its points. Multiple methods aim at improving this situation by either using particular properties of the data space or by defining other constructions of geometrically natural subspaces: Principal nested spheres (PNS) and composite PNS (CPNS, [10]) use the particular fact that sub-spheres of spheres are totally geodesic. Horizontal component analysis (HCA, [18]) uses parallel transport and Fermi coordinates. Torus PCA (TPCA, [3]) performs surgery on tori to deform them into spheres on which CPNS can be used. Geodesic PCA (GPCA, [8]) finds principal geodesics passing m that minimize residual errors. Recently barycentric subspace analysis (BSA, [16]) uses subspaces arising as affine spans between data points. Common to these approaches is the explicit construction of low-dimensional subspaces that, focusing on different aspects, are as faithful to the nonlinear geometry as possible.

In contrast to the Euclidean situation, subspaces maximizing captured variance are not equivalent to subspaces minimizing residual errors. Even for distributions with local support, recurring and dense geodesics can make subspace projections minimizing residual errors undefined [9]. Projections to subspaces in general have complex behavior, and finding subspaces maximizing variance of projections can also lead to counter-intuitive results or non-existing maxima.

2.2 PCA from a Probabilistic View

Instead of focusing on PCA as a fit of linear subspaces, we take a probabilistic approach and generalize the Euclidean probabilistic PCA (PPCA, [26]) formulation of PCA to the manifold setting. PPCA interprets PCA as a maximum likelihood fit of the factor model (1.1) when restricting W to be of rank $k \leq d$ and setting the covariance matrix for the noise ϵ to be diagonal $\sigma^2 I$. Because x is assumed normally distributed with unit variance, the marginal distribution of y is normal as well. i.e., PPCA assumes

$$y|x \sim N(Wx + m, \sigma^2 I) \quad (2.1)$$

with the latent variables x normally distributed $\mathcal{N}(0, I)$ and i.i.d. isotropic noise $\epsilon \sim \mathcal{N}(0, \sigma^2 I)$. This implies

$$y \sim N(m, \Sigma) \quad (2.2)$$

with $\Sigma = WW^T + \sigma^2 I$.

Assuming σ , m and W are already estimated, the latent variable $x|y_i$ conditioned on the data y_i takes the role of the ordinary principal components of the data y_i in PCA. To get a single data descriptor for y_i , one can take the expectation of $x|y_i$ which has the explicit expression $x_i := E[x|y_i] = (W^T W + \sigma^2 I)^{-1} W^T (y_i - m)$. We here loosely denote x_i as principal components for PPCA.

From (2.2), it is immediate that the log-likelihood of the data y is

$$\mathcal{L}(y; W, \sigma, m) = -\frac{1}{2}(d \ln(2\pi) + \ln |\Sigma| + (y - m)^T \Sigma^{-1} (y - m)) \quad (2.3)$$

and the maximum likelihood estimate for W is up to rotation given by $W_{ML} = U_k(\Lambda - \sigma^2 I)^{1/2}$, $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_k)$ where U_k contains the first k principal eigenvectors of the sample covariance matrix of y_i in the columns, and $\lambda_1, \dots, \lambda_k$ are corresponding eigenvalues.

For both the ML estimates of m and W , and for the principal components x_i , the usual non-probabilistic PCA solution is recovered in the zero noise limit $\sigma^2 \rightarrow 0$. Note that a similar interpretation of PCA can be found in [17]. Here the case $\sigma > 0$ is denoted sensible PCA (SPCA).

Turning to the manifold situation, the probabilistic view implies that the fundamental problem in generalizing PPCA is not to define low-dimensional subspaces as sought by the approaches described in section 2.1 but instead to define a natural generalization of the Euclidean normal distribution to manifolds. PPCA has previously been generalized to manifolds with the probabilistic principal geodesic analysis (PPGA, [27]) procedure. The probability model is for PPGA a Riemannian normal distribution defined via its density

$$p(y; m, \tau) = \frac{1}{C(m, \tau)} e^{-\frac{\tau}{2} d_g(m, y)^2} \quad (2.4)$$

with $C(m, \tau)$ a normalization constant and d_g the distance induced by a Riemannian metric g on M . This distribution is a function of the squared Riemannian distance to m , it is isotropic and closely connected to geodesic distances and least-squares.

We aim in this paper to generalize PPCA using a different probability model. In [20, 24], a generalization of the Euclidean normal distribution is constructed that intrinsically takes anisotropy of the data into account. Instead of using the metric g and squared d_g distances to define a density as for the Riemannian normal distribution in PPGA, a connection ∇ on M is used as the fundamental tool for moving vectors and tensors between infinitesimally close tangent spaces and thereby to distribute orientations on the manifold without linearizing the manifold. This distribution of directions enables stochastic processes to be mapped onto M . With a fixed background measure, the resulting transition distributions have densities and the data thereby likelihoods. We will make this precise in section 4 after describing the necessary fiber bundle geometry in the next section.

3 Fiber Bundle Geometry

We here review aspects of fiber bundle geometry focusing on the concepts necessary for the intrinsic construction of normal-like probability distributions. We therefore omit many geometric details. More information can for example be found in the papers [24, 21], the book [7] and, for a comprehensive exposition of fiber bundle geometry, the book [11].

The focus is to handle the absence of global orientation fields on the manifold. Indeed, we can choose an ordered basis, a frame, for a single tangent space $T_x M$ providing a reference orientation for vectors in $T_x M$. This however does not provide us with information about vectors in $T_y M$, $x \neq y$. When M is equipped with a connection, the parallel transport along a curve can be used to link the tangent spaces $T_x M$ and $T_y M$. However, the parallel transport is dependent on the curve, and

the holonomy of a manifold with non-zero curvature implies that different choices of curves give different parallel transport. This problem is elegantly handled by the Eells-Elworthy-Malliavin construction of Brownian motion that uses a Euclidean martingale and the orthonormal frame bundle to lift the problem to a distribution of orientations over M . The outline below is based on and inspired by this idea.

3.1 The Frame Bundle

The frame bundle is the set of points $x \in M$ and ordered bases ν for $T_x M$. For an element $u = (x, \nu) \in FM$, the frame part ν consists of d basis vectors $\nu_i \in T_x M$. Splitting u in the parts x and ν technically requires a local trivialization of FM . Instead, we let π be the projection $FM \rightarrow M$ that just drops the frame from a frame bundle elements and thus sends u to x , and we write just u and u_i for the frame and basis vectors. The frame bundle can equivalently be defined as the principal bundle $GL(\mathbb{R}^d, TM)$ of invertible linear maps between \mathbb{R}^d and the tangent bundle TM . An element $u \in GL(\mathbb{R}^d, TM)$ assigns to a vector $v \in \mathbb{R}^d$ an element $uv \in T_x M$. The d basis vectors $u_i \in T_x M$ in this view appear as the images ue_1, \dots, ue_d with e_1, \dots, e_d the standard basis for \mathbb{R}^d .

If the manifold is equipped with a connection ∇ , each of the basis vectors u_i can be parallel transported along a curve γ on M passing $x = \pi(u)$. We write the parallel transport of a vector $v \in T_{\gamma_0} M$ along γ as $P_{\gamma,t}(v)$ giving a vector in $T_{\gamma_t} M$. Performing this operation for all u_i gives a transport along γ of the entire frame u . We can thus lift the parallel transport operation from working on vectors in the tangent bundle TM to transporting frames in FM .

The infinitesimal $t \rightarrow 0$ limit of the parallel transport of u along γ gives an infinitesimal variation in FM , i.e. a vector in the tangent bundle TFM of the frame bundle. The span of the tangent vectors arising from such infinitesimal parallel transports, i.e. from choosing curves γ on M with different velocities $\dot{\gamma}_0$, defines a linear subbundle of TFM denoted the horizontal subbundle. Another subbundle of TFM is the vertical subbundle VFM , and we can write TFM as a direct sum $TFM = HFM \oplus VFM$ thanks to the connection. Elements in the vertical bundle are variations of u that keep $x = \pi(u)$ fixed varying only the frame part in the fiber $\pi^{-1}(x)$ above x . Conversely, infinitesimal variations in the horizontal subspace moves x while keeping the frame part of u as fixed as possible as measured by the connection or, equivalently, the parallel transport. HFM variations are thus zero-acceleration as measured by the connection.

An important property of the horizontal bundle HFM is that the pushforward $\pi_* : TFM \rightarrow TM$ of the projection π is a linear isomorphism when restricted to the horizontal space for a given $u \in FM$, i.e. $\pi_*|_{H_u FM} : H_u FM \rightarrow T_{\pi(u)} M$ is invertible. The inverse is called the horizontal lift, here denoted $h_u : T_{\pi(u)} M \rightarrow H_u FM$. That is, we can relate vectors in $T_{\pi(u)} M$ and vectors in $H_u FM$ in a one-to-one fashion. An important consequence, in particular for our purposes, is the fact that the horizontal lift gives a basis of globally defined vector fields $(H_1(u), \dots, H_d(u))$ for HFM . This is very much in contrast to the situation on the base manifold M where topology generally prohibits globally defined non-zero vector fields. We get this basis by, for each basis element $e_i \in \mathbb{R}^d$, using the horizontal lift $u \mapsto h_u(ue_i)$ to get the

HFM valued vector field on FM denoted $H_i(u)$. Moreover, if M has a Riemannian metric, the basis is globally orthonormal for $u \in OM$, OM being the subbundle of FM consisting of orthonormal frames, in the sense that $(\pi_* H_1(u), \dots, \pi_* H_d(u))$ constitutes an orthonormal basis at each point $\pi(u)$.

3.2 Sub-Riemannian Structure

Recall that the density of the Euclidean normal distribution with covariance Σ is a function of the weighed quadratic form $x^T \Sigma^{-1} x$. If we let W be a square root $WW^T = \Sigma$, we can write this as $(W^{-1}x)^T (W^{-1}x)$ using the usual \mathbb{R}^d dot product $v^T v$ of the preimage $W^{-1}x$ regarding W as a linear map $\mathbb{R}^d \rightarrow \mathbb{R}^d$. This construction can be naturally extended to give a sub-Riemannian structure on FM that is then, by definition, related to the density of the normal distribution. Because $u \in FM$ can be regarded a linear map $\text{GL}(\mathbb{R}^d, T_{\pi(u)}M)$, we can take $(u^{-1}v)^T u^{-1}v$ for $v \in T_{\pi(u)}M$. We thus informally regard u^{-1} a square root of the precision matrix Σ^{-1} in $T_x M$, or, conversely, u is a square root of the covariance matrix Σ that is then a matrix on \mathbb{R}^d . To be precise, we define the inner product

$$\Sigma^{-1}(u)(v, w) = \langle u^{-1}v, u^{-1}w \rangle_{\mathbb{R}^n} = (u^{-1}v)^T u^{-1}w, \quad v, w \in T_{\pi(u)}M. \quad (3.1)$$

The Σ^{-1} notation indicates that the inner product should be seen as encoding the precision matrix corresponding to the term $x^T \Sigma^{-1} x$ in the Euclidean normal distribution density. The inner product on $T_{\pi(u)}M$ lifts to an inner product on $T_u FM$

$$\Sigma^{-1}(u)(v_u, w_u) = \langle u^{-1}(\pi_*(v_u)), u^{-1}(\pi_*(w_u)) \rangle_{\mathbb{R}^n}, \quad v_u, w_u \in T_u FM. \quad (3.2)$$

Thanks to $\pi_*|_{H_u FM}$ onto $T_{\pi(u)}M$ being an isomorphism, this product is positive definite on $H_u FM$. However, it degenerates on VFM and therefore does not define a Riemannian structure on FM . It does however defines a sub-Riemannian structure. The sub-Riemannian metric can also be viewed as a map $g_\Sigma : TFM^* \rightarrow HFM$ defined by $\xi_u(v_u) = \Sigma^{-1}(u) \langle v_u, g_\Sigma(\xi_u) \rangle$ for $\xi_u \in T_u FM^*$. This in addition defines a cometric, also denoted g_Σ , in the form of an inner product on T^*FM by $g_\Sigma(\xi_u, \eta_u) = \eta_u(g_\Sigma(\xi_u))$, $\xi_u, \eta_u \in T_u FM^*$. Being inverse to the metric which is modeled after precision matrix, the cometric can be seen as encoding covariance.

3.3 Bundles of Symmetric Positive Definite Tensors

The quadratic form Σ^{-1} is an element of the bundle $\text{Sym}^+ M$ of covariant 2-tensors on M . The projection $\pi : FM \rightarrow M$ can be factored through this bundle giving a map q such that $FM \xrightarrow{\Sigma^{-1}} \text{Sym}^+ M \xrightarrow{q} M$ with $\pi = q \circ \Sigma^{-1}$. It is natural modeling covariance structure using $\text{Sym}^+ M$ since the bundle omits the implicit rotation that a representation of Σ by a square root $u \in FM$ imply. Indeed, $\text{Sym}^+ M$ can be viewed as the quotient $FM/O(d)$ where the orthogonal group $O(d)$ acts on the right by $R.u = uR$ for $R \in O(d)$ and $u \in \text{GL}(\mathbb{R}^d, TM)$.

As shown in [24], the horizontal/vertical splitting of TFM and the sub-Riemannian structure on FM descend to corresponding structures on $\text{Sym}^+ M$. We can therefore work on the two bundles interchangeably in the same way as one shifts between a

square root covariance W and the covariance matrix $\Sigma = WW^T$ in Euclidean statistics. As we will see below, the frame bundle supports the development construction for mapping Euclidean semimartingales to the manifold. We therefore often work on FM keeping in mind that the generated covariance structures can be seen as element of the quotient bundle Sym^+ .

3.4 Development and Stochastic Development

Let x_t be a Euclidean semimartingale on \mathbb{R}^d defined as the solution to the Stratonovich SDE

$$dx_t = b(t, x_t) + W(t, x_t) \circ dB_t \quad (3.3)$$

where B_t is a standard Brownian motion on \mathbb{R}^k , $k \leq d$, and \circ denotes Stratonovich multiplication. We let \mathbb{P}_{x_t} denote its law and, when $k = d$, $p_{x_t}(v; x, T)$ denotes the time T transition density of the process evaluated at $v \in \mathbb{R}^d$ when started with initial conditions $x_0 = x$. Note that when the drift b is zero and W is a time- and spatially stationary full-rank matrix, x_T is normally distributed with covariance $\Sigma = WW^T$ and density $p_{x_t}(v; x, T) = (2\pi T|\Sigma|)^{-\frac{d}{2}} e^{-\frac{1}{2}(v-x)^T \Sigma^{-1}(v-x)}$. This view of the normal distribution arising as the combined effects of an continuous sequence of infinitesimal random steps with covariance $\Sigma = WW^T$ is particularly well-suited for generalizing to the manifold situation.

We achieve this generalization using the stochastic development construction, see e.g. [7]. Recall above the existence of a globally defined basis (H_1, \dots, H_d) for the horizontal bundle HFM . This can be used to define an FM valued process from the semimartingale x_t via the SDE

$$du_t = H_i(u_t) \circ dx_t^i. \quad (3.4)$$

Note the Einstein summation convention implies a summation over the components dx_t^i and the horizontal basis fields. In the deterministic case ($W = 0$ in (3.3)), the ODE is denoted just development or “rolling-without-slipping” due to the fact that the frame represented by a solution u_t is parallel transported, or rolled, along the manifold. This is a consequence of H_i representing infinitesimal parallel transport. In the stochastic case, when $u_0 = u$ is an orthonormal frame with respect to a Riemannian metric, i.e. an element of the orthonormal frame bundle OM , the construction is the basis for the Eells-Elworthy-Malliavin construction of Brownian motion [4]. In the following, we denote by $\phi_u(x_t)$ the solution of (3.4) of a path $x_t \in \mathbb{R}^d$, deterministic or stochastic, started at $u \in FM$. The inverse of ϕ_u is denote anti-development.

3.5 Anisotropic Normal Distributions

In [20, 24], the stochastic development construction is used with x_t a Euclidean Brownian motion to map from a starting frame u to a distribution $y_T = \pi(u_T)$ on M with density $p_{y_t}(v; u, T)$ with respect to a fixed base measure (e.g. the Riemannian volume form vol_g). The orthonormality condition on u in the Eells-Elworthy-Malliavin construction of Brownian motion is thus relaxed. The result is the anisotropic distribution y_T that has nontrivial covariance in the sense that the infinitesimal stochastic

displacements of the process have covariance given by the frame u . The base point $m = \pi(u)$ is thus interpreted as the mean of the distribution, and the frame u itself models the square root covariance. The precision matrix is the inner product $\Sigma^{-1}(u)$ on $T_{\pi(u)}M$. The orientation problem that usually prevents us from defining globally non-zero vector fields on M with special properties, e.g. orthonormality, is thus handled by spreading the orientations stochastically in FM with parallel transport and taking the time T distribution before projecting the resulting distribution to M . The density of the generated distributions is visualized in Figure 2.

Note that the process u_t is actually a semi-elliptic FM -valued Brownian motion with respect to the sub-Riemannian metric g_Σ on FM . The semi-ellipticity arise because the diffusion is generated only in the subspace HFM of TFM . The curvature of M is exactly the non-integrability of the horizontal fields H_i , and non-zero curvature therefore implies that the process will diffuse out of the horizontal bundle and generate a larger subspace of FM . It does however not satisfy the Hörmander condition on TFM , and the diffusion will not fill all of FM .

3.6 Euclideanization

In manifold statistics, the concept of Euclideanizations of the data, mappings from the nonlinear manifold to a Euclidean space that preserve important features of the data, is often appearing. This is for example the case in tangent space PCA and related models where the map Log_m maps the data from M to a linear tangent space T_mM thus providing a linear view of the data. This is however at the cost of flattening the space and thus losing curvature information.

The development procedure provides a Euclideanization that is qualitatively different: Development is a one-to-one map from the Wiener space of paths $W([0, T], \mathbb{R}^d)$ to the corresponding space $W([0, T], M)$ on M . It thus provides a Euclidean view of paths to data instead of the data itself. As we will see below, we can marginalize over all paths $x_t \in W([0, T], \mathbb{R}^d)$ reaching a point $y \in M$ in the sense $\phi_u(x_t)_T = y$. This for example provides a Euclidean view of the data as the expectation over paths x_t whose development reach y , or the expectation over endpoints x_T of paths with development reaching y .

4 Probabilistic Principal Component Analysis on Manifolds

The probabilistic formulation of PPCA makes no reference to linear subspaces which suggests a generalization to nonlinear manifolds can be obtained without constructing submanifolds that in general cannot satisfy the usual properties of linear subspaces of Euclidean spaces, e.g. being totally geodesic. The main difficulty in generalizing PPCA instead lies in the latent variable model being additive and using normal distributions, neither which are directly transferable to manifolds. Instead, we propose to use the anisotropic normal distributions described in section 3.5 to take the place of the normal distributions in Euclidean PPCA.¹

¹In the abstract [19], the procedure was denoted DPCA for diffusion PCA due to the use of semi-elliptic diffusions to generate the probability model but we avoid using this term here.

4.1 The Nonlinear Probability Model

Consider the map $\phi_{B_t, T} : FM \rightarrow \text{Prob}(M)$ that by stochastic development sends $u \in FM$ to $\pi(u_T)$ where the FM diffusion $du_t = H_i(u_t) \circ dB_t^i$ is started at time $t = 0$ at u , and $B_t \in \mathbb{R}^d$ is a Brownian motion. The stopping time T can without loss of generality be assumed $T = 1$. Recall from the discussion earlier in the paper that u represents the mean $m = \pi(u)$ and the frame u the square root covariance of the distribution $\mu_u = \pi(u_T)$. The precision matrix is the inner product $\Sigma^{-1}(u)$ given by u . We let $\Gamma \subset \text{Prob}(M)$ be the image of $\phi_{B_t, T}$, i.e. the set of distributions $\mu_u = \phi_{B_t, T}(u)$ resulting from point-sourced diffusions in FM stopped at time T . We then assume the observed data is distributed according to $\mu_u \in \Gamma$ so that $y \sim \mu_u = \pi(u_T)$ for a diffusion $u_T \in FM$ started at u .

Let μ_0 be a fixed measure on M , e.g. a Riemannian volume form vol_g . For each distribution $\mu \in \Gamma$, we write p_μ for the density satisfying $\mu = p\mu_0 \in \Gamma$. We can then define the log-likelihood

$$\ln \mathcal{L}(y; u) = \ln \mathcal{L}(y; \mu_u) = \ln p_{\mu_u}(y) \quad (4.1)$$

for a sample $y \in M$. Now for samples y_1, \dots, y_N , let $u_{ML} \in FM$ be a maximum for $\ln \mathcal{L}(y_1, \dots, y_N; \mu_u) = \prod_{i=1}^N \ln \mathcal{L}(y_i; u)$. Then u_{ML} contains the parameters of a maximum likelihood fit to the data y_1, \dots, y_N of the parameters of the model in u .

In the PPCA model (2.2), the coefficient matrix W was assumed of rank $k \leq d$. A similar rank k model in the nonlinear setting can be constructed by instead of modelling u directly, letting W be an element of the bundle $F^k M$ of rank k linear maps $\mathbb{R}^k \rightarrow TM$. In addition, we need to represent the isotropic iid. noise ϵ with variance σ^2 . Generators of isotropic noise are elements of the orthonormal frame bundle OM with respect to a Riemannian metric g on M , confer the Eells-Elworthy-Malliavin construction of Brownian motion. We denote such an element by $R \in O_{\pi(u)} M$ to emphasize its pure rotation, no scaling nature. We then set

$$\begin{aligned} dW &= H_i(W_t) \circ dX_t^i + \sigma H_i(R_t) \circ d\epsilon_t^i, \\ dR &= h_{R_t}(\pi_*(dW)). \end{aligned} \quad (4.2)$$

Here ϵ_t is a Brownian motion on \mathbb{R}^d modeling increments of the iid. noise while X_t is now a Brownian motion on \mathbb{R}^k modeling the latent variables. This is a direct extension of the PPCA model (2.2) and the latent variable model (1.1) using the stochastic development construction (3.4). The horizontal fields and stochastic development are defined on $F^k M$ in a similar way as on FM . However, in the no noise situation $\sigma = 0$, the generated distributions μ_W would not have strictly positive density similarly to the Euclidean PPCA case in the limit $\sigma \rightarrow 0$. This is a consequence of the generated process not being full-rank on $HF^k M$ if $W \in F^k M$, $k < d$, and the isotropic noise $d\epsilon_t$ is not added.

Note that the system (4.2) could equivalently be formulated as

$$dR = H_i(R_t) \circ ((W dB_t)^i + \sigma d\epsilon_t^i), \quad (4.3)$$

i.e. by multiplying the Brownian motion X_t with a fixed matrix W and adding noise ϵ_t before multiply on H_i . In practice, with (4.2), we need only simulate the

W_t evolution in $F^k M$ since the isotropic part R_t can be obtained up to rotation by lifting any g orthonormal basis to an element of OM . For high dimensional systems, simulating on F^k , $k \ll d$ can be computationally much more tractable than on FM (or OM). The system (4.3) lives on OM and therefore does not have a similar reduction property. Finally, the system (4.2) separates the latent process $X_t \in \mathbb{R}^k$ from the geometric flow of the coefficient matrix $W_t \in F^k M$. This view emphasizes the role of the fiber bundle $F^k M$ in modeling the flow over M of the coefficient matrix while the latent process X_t is Euclidean.

We saw earlier that u defined a sub-Riemannian structure on FM . The addition in (4.2) can also be seen as a sub-Riemannian metric on FM on the form $g_W + \sigma^2 \tilde{g}$ [21]. Here \tilde{g} is a lift $\tilde{g}(\xi_u, \eta_u) = g(\pi_* \xi_u, \pi_* \eta_u)$ to TFM^* of the Riemannian metric g on M , and g_W is a rank k inner product on $HF^k M$ defined from the map W .

In the following, to simplify notation, we mostly refer to the stochastic process as just u_t without distinguishing between the FM valued full rank version u_t solution to (3.4) and the low-rank version W_t in $F^k M$ solution to (4.2).

4.2 The Principal Components

Euclidean Probabilistic PCA reduces the dimensionality of the data by considering the latent variables conditioned on the observed data $x|y_i$. This random variable converges to the principal components as $\sigma^2 \rightarrow 0$. In the proposed model (4.2), the latent process X_t takes the place of x . With non-zero curvature, the latent process cannot directly summarize the observations in single vectors: sample paths $X_t(\omega)$ generating paths $\pi(W_t(\omega))$ hitting the same endpoint y_i on M will in general not have the same endpoint $X_T(\omega)$ in \mathbb{R}^k , see Figure 4 in the experiments, section 6. However, we can still consider the conditioned latent variable process $X_t|\pi(W_T) = y_i$. Since the latent process lives in \mathbb{R}^k where we can take expectation (in contrast to on M), and summarize by the mean of latent sample paths reaching y_i :

$$\bar{X}_{i,t} = E[X_t|\pi(W_T) = y_i] \quad (4.4)$$

Thus $\bar{X}_{i,t}$ take the role of the latent variables in PPCA. Note that given the source $W \in F^k M$, the sample paths can be equivalently viewed as paths $\pi(W_t(\omega))$ on M or as paths $X_t(\omega)$ in \mathbb{R}^k . Examples of mean paths are illustrated in Figure 3. In \mathbb{R}^k , the data can be further summarized by integrating out the time dependence from \bar{X}_i summarizing the data y_i only in the latent endpoint $X_{i,T}$ with mean $\bar{X}_{i,t}$. The conditioned latent variables in this way provide a Euclideanization of the data similar to those provided by the manifold PCA methods PGA, GPCA, HCA, and BSA. Because of the process nature of the latent variable, the linearization will be quite different from the linearizations provided by these methods.

4.3 Zero Noise Limit

In the $\sigma \rightarrow 0$ limit, PPCA recovers the original PCA formulation with projections to the latent space that either minimize residual error or maximize variance of the projected data. In the nonlinear case (4.2), as σ tends to 0, we get an FM diffusion that progressively concentrates its infinitesimal displacements around

$H_{W_t} F^k M$. The small-time limit of the likelihood behaves Gaussian-like [24] in the sense, $\lim_{t \rightarrow 0} 2t \log p_{\mu_u, t}(y) = -d_{\Sigma^{-1}}(u_0, \pi^{-1}(y))^2$ with $d_{g_{\Sigma^{-1}}}$ the distance on FM induced by the sub-Riemannian metric g_{Σ} . We can then conjecture that, in the limit, we recover projections to the latent space in a similar sense. If we let $Q(u)$ denote the subspace of FM reachable by horizontal paths starting at $u \in FM$, a natural limit notion of the principal components would be

$$\operatorname{argmin}_{\tilde{u} \in Q(u)} d_{g_{\Sigma^{-1}}}(\tilde{u}, \pi^{-1}(x))^2. \quad (4.5)$$

We return to the space Q briefly below, and leave the question if the actual $\sigma \rightarrow 0$ limit of the principal components take a form similar to (4.5) to future work.

4.4 Rotations and Subbundles

As discussed in section 3.3, the above construction is over specified in the sense of an arbitrary rotation being present in the representation u of the covariance similarly to the Euclidean case of specifying covariance with a square root W instead of the actual covariance matrix $\Sigma = WW^T$. We can handle this by quotienting out $O(d)$, instead specifying the construction on Sym^+ , see section 3.3. This has however little influence in practice where the rotation implicit in the matrix u can just be ignored.

4.5 Curvature and Nonintegrability

While principal subspaces in the Euclidean case are linear subspaces of \mathbb{R}^d , the space $Q(W)$ of endpoints of curves starting at $W \in F^k M$, staying horizontal, and generated by the flow equation (3.4) is not in general a k -dimensional submanifold of $F^k M$. The geometric reason is that curvature is equivalent to non-integrability of the horizontal distribution of the vector horizontal fields H_1, \dots, H_k on $F^k M$, i.e. the $VF^k M$ valued Lie brackets $[H_i, H_j]$, $1 \leq i, j \leq k$ are non-zero for some i, j . Thus, the Frobenius theorem tells us that the span doesn't integrate to a k -dimensional submanifold. This can be seen as the key consequence of curvature for PCA like constructions defined via infinitesimal flows. In the present case, we do not need to truncate the non-integrable span to obtain a k -dimensional submanifold as is done when e.g. considering geodesic sprays starting at $\pi(u)$. Instead, we simply model the data as being distributed according to the development of horizontal stochastic flows and thus avoid referring to subspaces in the PCA construction. Note that we can still extract principal components as discussed above.

In some cases, we can say more about the structure of $Q(W)$ or $Q(u)$. If $k = d$ and the bracket span $\operatorname{Lie}(H_1, \dots, H_d)$ of the horizontal fields has constant dimension for any point at M , there exists a subbundle of FM on which H_1, \dots, H_d satisfies the Hörmander condition. In this case, the reachable set $Q(u)$ is this subbundle of FM , and $\pi(U)$ is a submanifold of M . Note that nonzero curvature implies that the dimension $\dim(Q(u))$ is greater than d . An example of this case can be seen for the sphere \mathbb{S}^2 where the bracket span has rank 3 and $Q(u) \simeq O\mathbb{S}^2$ for any orthonormal u . However, if the sphere is deformed to be locally flat in a neighborhood of $\pi(u)$, the rank of the bracket span is lowered to 2 in this neighborhood and the constant rank condition fails.

4.6 Extensions

As noted in [26], the probabilistic formulation has advantages beyond the theoretical insight and the ability to perform estimation with MLE. This includes extension to mixed models where data are assumed distributed according to a sum of multiple latent models of the form (1.1), in effect allowing different centers m_1, \dots, m_j or more complicated shaped distributions. Similar flexibility is present in the manifold situation. The stochastic process u_t or W_t can be started at multiple points $u_1, \dots, u_j \in FM$ and the resulting densities averaged.

5 Inference and Predictions

We here describe two estimation approaches. The first is based on the estimators described in [20, 24] that use the anisotropically weighted energy of the most probable paths as surrogates for the log data likelihood. The second approach outlines a Monte Carlo method for estimating transition densities from which the likelihood can be optimized.

5.1 Most Probable Paths

In [24], the small time asymptotic limit of $p_{\mu_u, t}(\cdot)$ is used to suggest the estimator

$$\operatorname{argmin}_{u \in FM} \sum_{i=1}^N \left(d_{\Sigma^{-1}}(u, \pi_u^{-1}(y_i))^2 - N \log(\det(u)_g) \right) \quad (5.1)$$

for the maximum likelihood fit of an anisotropic normal distribution to data points $y_1, \dots, y_N \in M$. The FM distance $d_{\Sigma^{-1}}$ is dependent on u , and a minimizer for (5.1) can be found by iterative optimization. Because the small time asymptotic limit is used, the estimator is reasonable for data with limited variation around $\pi(u)$. The $d_{\Sigma^{-1}}$ distances are realized by most probable paths on FM [22], a family of paths that generalizes geodesics when u is not orthonormal.

5.2 Bridge Simulation

We generally do not wish to restrict to cases where the data variation is small. As data variation and curvature increases, the estimator (5.1) will provide a progressively less precise approximation of the optimal likelihood (4.1). Instead, we here describe a bridge simulation scheme based on the conditioned diffusion bridge simulation method of [2] and the maximum likelihood estimation in [23].

In [2], simulation of diffusion processes $x_t \in \mathbb{R}^d$ given by the Itô SDE

$$dx_t = b(t, x_t)dt + W(t, x_t)dB_t \quad (5.2)$$

conditioned on hitting a point $v \in \mathbb{R}^d$ at time T is considered based on the idea of adding a drift term that guides the diffusion towards the target v . The resulting modified SDE takes the form

$$d\tilde{x}_t = b(t, \tilde{x}_t)dt - \frac{\tilde{x}_t - v}{T - t}dt + W(t, \tilde{x}_t)dB_t. \quad (5.3)$$

Under reasonable assumptions, including that W is invertible for all t, x , [2] shows that $E_{x_t|v}[f(x_t)] = E_{\tilde{x}_t}[f(\tilde{x}_t)\varphi(\tilde{x}_t)]$ for measurable maps f on $W([0, T], \mathbb{R}^d)$. Here φ is a correction factor that takes into account the difference of the laws of the process $x_t|v$ that is conditioned on hitting v at time T , and the modified process \tilde{x}_t . Note that \tilde{x}_t by construction will hit v a.s. The density of x_T can be recovered from this construction as

$$p(v; x, T) = \left(\frac{|\Sigma^{-1}(v)|}{2\pi T} \right)^{\frac{d}{2}} e^{-\frac{\|W(x)^{-1}(x-v)\|^2}{2T}} \mathbb{E}_{\tilde{x}_t}[\varphi(\tilde{x}_t)] \quad (5.4)$$

with $\Sigma = WW^T$.

We now suggest to use a similar approach for estimating the likelihood of the data under the proposed manifold PPCA model. Because we do not have a diffusion process with invertible diffusion field W as above (the process u_t (or W_t) is only semi-elliptic on FM), we will not here give a rigorous argument for the convergence of the procedure. We will instead sketch an approach that uses the fact that the data in the model is only observed at M while the process u_t lives in FM . This situation is related to the case of partial observations treated in [14], see also the semi-elliptic phase-space flows in [1]. We assume we have a chart that covers M except for a set of measure zero, and we use this to write the process in coordinates. This is for example the case in the experiments we present in the next section for data on surfaces in \mathbb{R}^3 .

The idea is now to take coordinates on FM as in [15] and write the u_t diffusion as a process in coordinates. The coordinates imply a trivialization of FM so we can write $u_t = (x_t, \nu_t)$ with $x_t \in M$ and ν_t frames. We write the system in short form as

$$\begin{pmatrix} dx_t \\ d\nu_t \end{pmatrix} = \begin{pmatrix} b_x \\ b_\nu \end{pmatrix} dt + \begin{pmatrix} W_x \\ W_\nu \end{pmatrix} \circ dW. \quad (5.5)$$

The W_x part is actually just ν by construction of the process. We then make a modified process $\tilde{u}_t = (\tilde{x}_t, \tilde{\nu}_t)$

$$\begin{pmatrix} d\tilde{x}_t \\ d\tilde{\nu}_t \end{pmatrix} = \begin{pmatrix} b_x \\ b_\nu \end{pmatrix} dt - \begin{pmatrix} \tilde{\nu} \\ W_\nu \end{pmatrix} \frac{\tilde{\nu}^{-1}(\tilde{x}_t - v)}{T - t} dt + \begin{pmatrix} \tilde{\nu} \\ W_\nu \end{pmatrix} \circ dW. \quad (5.6)$$

Intuitively, the use of $\tilde{\nu}^{-1}$ in the drift term produces a correction that after multiplication on $(\tilde{\nu}, W_\nu)^T$ points in the direction $\tilde{x}_t - v$ on M while staying horizontal on TFM .

Without arguing for convergence here, we aim for \tilde{x}_t to hit v at time T because of the added drift term. We then find the correction term φ as in [2], and arrive at the expression

$$p(v; u, T) = (2\pi T |u|_g^2)^{-\frac{d}{2}} e^{-\frac{\|u^{-1}(x-v)\|^2}{2T}} \mathbb{E}_{\tilde{u}_t}[\varphi(\tilde{u}_t)] \quad (5.7)$$

for the density of the process with correction term φ .

We can now write the density with respect to μ_0 , e.g. vol_g with g a Riemannian metric, and sample from $E_{\tilde{u}_t}[\varphi(\tilde{u}_t)]$ with a Monte Carlo scheme. We do this with Hamiltonian updates to keep the acceptance rate high. We can then optimize for $u = (x, \nu)$ either by directly taking gradients with respect to u of the sample approximation of the likelihood, or by an EM-approach.

6 Experiments

We aim here to visualize the effect of the method and the influence of curvature on two low-dimensional manifolds, the sphere \mathbb{S}^2 and a non-spherical ellipsoid. While curvature effects are visible in both cases, the non-symmetrical nature of the ellipsoid emphasizes the differences to the Euclidean situation. For both manifolds, we illustrate samples from the model with fixed mean and covariance encoded in the frame bundle element u . We then for optimal u illustrate how the non-linearity affects the principal components (4.4). After this, we illustrate iterations of a direct optimization of the approximate data likelihood from the density expression (5.7) using Monte Carlo sampling.

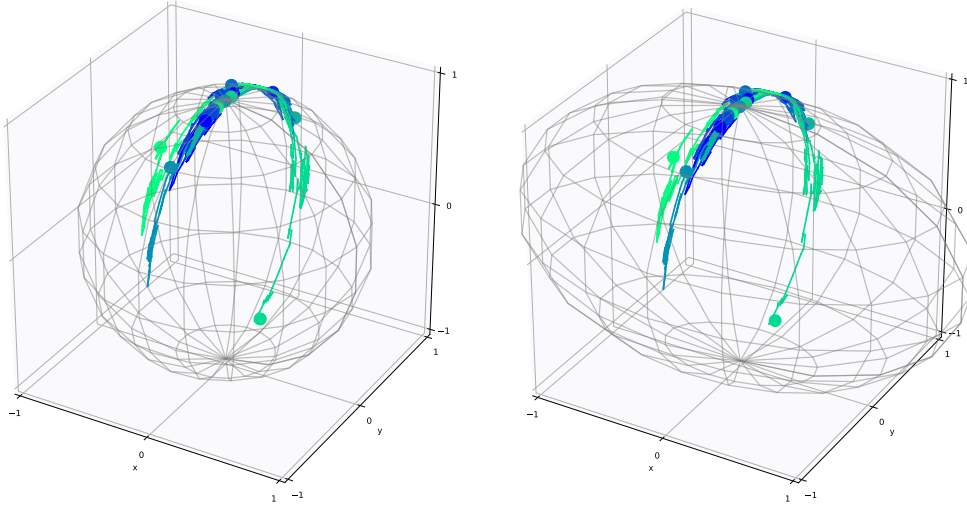


Figure 1: Samples with corresponding trajectories on the sphere \mathbb{S}^2 and an ellipsoid. The variance is 1 in the axis of major variation, and noise with variance $\sigma = .1$ is added in the orthogonal direction.

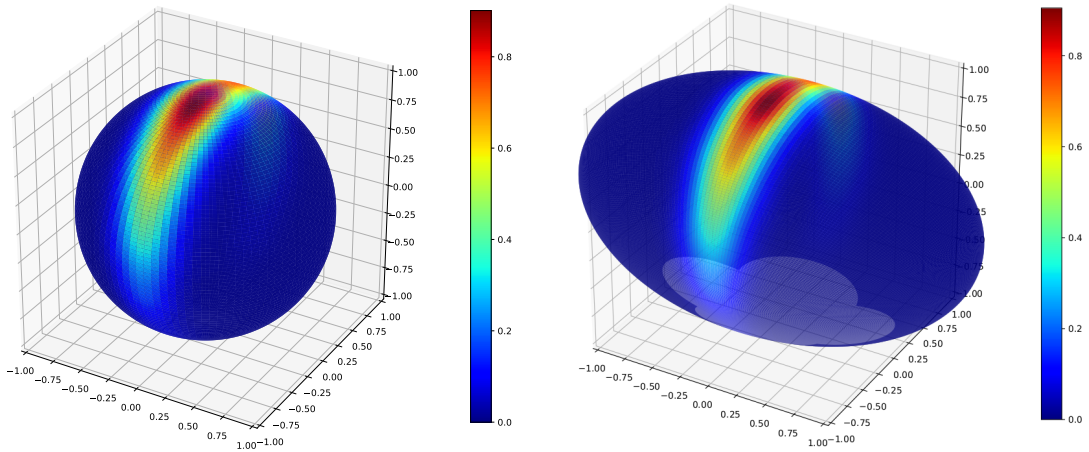


Figure 2: Density plots on the generated distribution on both surfaces.

The experiments are performed using the differential geometry library Theano Geometry² that is based on the Theano framework [25] for symbolic expression, automatic differentiation, and subsequent numerical evaluation. See also [12] for an extended description of the use of automatic differentiation for differential geometric and nonlinear statistical computations. Sampling from the likelihood expression (5.7) with Hamiltonian updates coupled with gradients for u involves very complex expressions with high order derivatives that would be practically infeasible to derive by hand. Fortunately, the use of automatic differentiation removes this complexity.

6.1 Density and Forward Sampling

Figure 1 shows samples from the probability model on the sphere \mathbb{S}^2 and the ellipsoid with variance 1 in one axis, and noise with variance $\sigma = .1$ in the orthogonal axis corresponding to the model (4.2). The starting point of the diffusion $\pi(u)$ corresponding to the mean is on both surfaces the north pole. The trajectories of the anisotropic process leading to the generated samples are visualized along with the endpoints. Figure 2 shows the corresponding density on both surfaces.




Figure 3: Samples from the process (5.6) conditioned on hitting the point v . The mean path (blue) plotted deviates from a geodesic (red) to v because of the coupling between the curvature and the anisotropic covariance.

²<https://bitbucket.com/stefansommer/theanogeometry>




Figure 4: Samples from the latent process X_t corresponding to the samples in Figure 3. The mean latent path is plotted in blue. Even though the process is conditioned on hitting v , the endpoints (black) of the latent path deviates. The mean path is not straight and therefore does not correspond to a geodesic on the surfaces.




Figure 5: Density plot of the trajectories in Figure 4. A straight line (red) corresponding to a geodesic on the surface from the north pole to v is plotted for comparison with the mean path (blue solid).




Figure 6: Samples used for the ML estimation in Figure 7.




Figure 7: (left) Estimated likelihood as a function of iterations of the iterative ML optimization. (right) Evolution of the variance and noise σ during the iterations of the algorithm. Ground truth values in red.

6.2 Principal Components

For a fixed point v on the surfaces, Figure 3 shows sample trajectories from the bridge process (5.6) generated with a Hamiltonian MCMC sampler for approximate evaluation of the expectation in the density expression (5.7). The corresponding samples from the latent process X_t are shown in Figure 4. Notice how the endpoints of the latent process samples vary even though the trajectories on the surfaces always end at v . This effect is a direct consequence of non-zero curvature. Figure 5 shows a density plot of the latent process, still conditioned on v . The mean latent path is plotted in blue in Figure 4 and 5, and the development of the mean path together with the parallel transported frame along the path are plotted on Figure 3. In Euclidean space, the mean path would be a straight line corresponding to the geodesics in Figure 3. While the mean latent paths for both surfaces clearly deviate from straight lines, the effect of the curvature is clearly more emphasized on the non-symmetric ellipsoid.

6.3 Maximum Likelihood

For the samples in Figure 6, we plot in Figure 7 the negative log-likelihood computed from a sample approximation of (5.7), the estimated variance in the axis of major variation, and σ . The horizontal axis shows the evolution of the negative log-likelihood and estimated variance during the evolution of an iterative maximum likelihood optimization. The samples are generated with variance .4 in the major axis and $\sigma = .075$ noise. The algorithm makes repeated sample approximations of (5.7), calculates the gradient and updates the parameters. As can be seen from the figure, the parameter estimates converges approximately to the true values.

7 Conclusion and Outlook

The probabilistic formulation in PPCA allows to generalize the PCA procedure to manifolds with a focus on data likelihoods in contrast to constructions of subspaces. This has previously been pursued with probabilistic PGA [27]. Here, we provide a generalization based on a different probability model using stochastic flows in the frame bundle and related fiber bundles. The main feature of the model is the intrinsic definition that does not refer to a linear tangent space approximation, is infinitesimal in modelling stochastic differential flows, and focuses on the generated likelihood and density instead of squared Riemannian distances. The model uses fiber bundle geometry that reveals important geometric information about the construction. As an example, the non-integrability of the horizontal subbundle is directly related to the curvature of the manifold. Instead of truncating the non-closure of the bracket of the horizontal basis fields to provide a submanifold, the construction allows the diffusion to spread into higher-dimensional subspaces. The data manifold is thereby not linearized and the curvature preserved in the analysis of the data.

The construction is based on the anisotropic normal distributions defined in [20, 24]. In addition to the presented PCA formulation, a regression model based on these distributions has been presented in [13]. We hope in future work to be able to use this and similar geometric constructions that preserve the nonlinear nature of the data space to generalize more statistical procedures to analysis of manifold valued data in intrinsic ways.

Acknowledgments

The work was supported by the Danish Council for Independent Research, and the CSGB Centre for Stochastic Geometry and Advanced Bioimaging funded by a grant from the Villum foundation. The research was partially performed at the Mathematisches Forschungsinstitut Oberwolfach (MFO), 2014 and 2018.

References

- [1] Arnaudon, A., Holm, D.D., Sommer, S.: A Geometric Framework for Stochastic Shape Analysis. submitted, arXiv:1703.09971 [cs, math] (2017)
- [2] Delyon, B., Hu, Y.: Simulation of conditioned diffusion and application to parameter estimation. *Stochastic Processes and their Applications* **116**(11), 1660–1675 (2006). DOI 10.1016/j.spa.2006.04.004
- [3] Eltzner, B., Huckemann, S., Mardia, K.V.: Torus Principal Component Analysis with an Application to RNA Structures. arXiv:1511.04993 [q-bio, stat] (2015). ArXiv: 1511.04993
- [4] Elworthy, D.: Geometric aspects of diffusions on manifolds. In: P.L. Hennequin (ed.) *École d’Été de Probabilités de Saint-Flour XV–XVII, 1985–87*, no. 1362 in *Lecture Notes in Mathematics*, pp. 277–425. Springer Berlin Heidelberg (1988)

- [5] Fletcher, P., Lu, C., Pizer, S., Joshi, S.: Principal geodesic analysis for the study of nonlinear statistics of shape. *Medical Imaging, IEEE Transactions on* (2004). DOI 10.1109/TMI.2004.831793
- [6] Fréchet, M.: Les éléments aléatoires de nature quelconque dans un espace distancié. *Ann. Inst. H. Poincaré* **10**, 215–310 (1948)
- [7] Hsu, E.P.: *Stochastic Analysis on Manifolds*. American Mathematical Soc. (2002)
- [8] Huckemann, S., Hotz, T., Munk, A.: Intrinsic shape analysis: Geodesic PCA for Riemannian manifolds modulo isometric Lie group actions. *Statistica Sinica* **20**(1), 1–100 (2010)
- [9] Huckemann, S., Hotz, T., Munk, A.: Rejoinder to the Discussion of “Intrinsic Shape Analysis: Geodesic Principal Component Analysis for Riemannian Manifolds under Isometric Lie Group Actions”. *Statistica Sinica* **20**(1) (2010)
- [10] Jung, S., Dryden, I.L., Marron, J.S.: Analysis of principal nested spheres. *Biometrika* **99**(3), 551–568 (2012). DOI 10.1093/biomet/ass022
- [11] Kolář, I., Slovák, J., Michor, P.W.: *Natural Operations in Differential Geometry*. Springer Berlin Heidelberg, Berlin, Heidelberg (1993)
- [12] Kuhnelt, L., Arnaudon, A., Sommer, S.: *Deep Learning Numerics for Differential Geometry*. in preparation (2017)
- [13] Kühnel, L., Sommer, S.: Stochastic Development Regression on Non-linear Manifolds. In: *Information Processing in Medical Imaging, Lecture Notes in Computer Science*, pp. 53–64. Springer, Cham (2017). DOI 10.1007/978-3-319-59050-9_5
- [14] Marchand, J.L.: Conditioning diffusions with respect to partial observations. *arXiv:1105.1608 [math]* (2011). ArXiv: 1105.1608
- [15] Mok, K.P.: On the differential geometry of frame bundles of Riemannian manifolds. *Journal Fur Die Reine Und Angewandte Mathematik* **1978**(302), 16–31 (1978). DOI 10.1515/crll.1978.302.16
- [16] Pennec, X.: Barycentric Subspace Analysis on Manifolds. *arXiv:1607.02833 [math, stat]* (2016). ArXiv: 1607.02833
- [17] Roweis, S.: EM Algorithms for PCA and SPCA. In: *Proceedings of the 1997 Conference on Advances in Neural Information Processing Systems 10, NIPS '97*, pp. 626–632. MIT Press, Cambridge, MA, USA (1998)
- [18] Sommer, S.: Horizontal Dimensionality Reduction and Iterated Frame Bundle Development. In: *Geometric Science of Information, LNCS*, pp. 76–83. Springer (2013)
- [19] Sommer, S.: *Diffusion Processes and PCA on Manifolds*. Mathematisches Forschungsinstitut Oberwolfach (2014). https://www.mfo.de/document/1440a/OWR_2014_44.pdf
- [20] Sommer, S.: Anisotropic Distributions on Manifolds: Template Estimation and Most Probable Paths. In: *Information Processing in Medical Imaging, Lecture Notes in Computer Science*, vol. 9123, pp. 193–204. Springer (2015)
- [21] Sommer, S.: Evolution Equations with Anisotropic Distributions and Diffusion PCA. In: F. Nielsen, F. Barbaresco (eds.) *Geometric Science of Information*, no. 9389 in *Lecture Notes in Computer Science*, pp. 3–11. Springer International Publishing (2015). DOI: 10.1007/978-3-319-25040-3_1

- [22] Sommer, S.: Anisotropically Weighted and Nonholonomically Constrained Evolutions on Manifolds. *Entropy* **18**(12), 425 (2016). DOI 10.3390/e18120425
- [23] Sommer, S., Arnaudon, A., Kuhnel, L., Joshi, S.: Bridge Simulation and Metric Estimation on Landmark Manifolds. In: *Graphs in Biomedical Image Analysis, Computational Anatomy and Imaging Genetics, Lecture Notes in Computer Science*, pp. 79–91. Springer, Cham (2017). DOI: 10.1007/978-3-319-67675-3_8
- [24] Sommer, S., Svane, A.M.: Modelling anisotropic covariance using stochastic development and sub-Riemannian frame bundle geometry. *Journal of Geometric Mechanics* **9**(3), 391–410 (2017). DOI 10.3934/jgm.2017015
- [25] Team, T.T.D.: Theano: A Python framework for fast computation of mathematical expressions. *arXiv:1605.02688 [cs]* (2016). ArXiv: 1605.02688
- [26] Tipping, M.E., Bishop, C.M.: Probabilistic Principal Component Analysis. *Journal of the Royal Statistical Society. Series B* **61**(3), 611–622 (1999)
- [27] Zhang, M., Fletcher, P.: Probabilistic Principal Geodesic Analysis. In: *NIPS*, pp. 1178–1186 (2013)