Expected path length on random manifolds

David Eklund[*]   Søren Hauberg[*]

daek@kth.se   sohau@dtu.dk

August 21, 2019

Abstract

Manifold learning seeks a low dimensional representation that faithfully captures the essence of data. Current methods can successfully learn such representations, but do not provide a meaningful set of operations that are associated with the representation. Working towards *operational representation learning*, we endow the latent space of a large class of generative models with a random Riemannian metric, which provides us with elementary operators. As computational tools are unavailable for random Riemannian manifolds, we study deterministic approximations and derive tight error bounds on expected distances.

Key words: random fields, random metrics on manifolds, Gaussian process latent variable models, expected curve length.

1   Introduction

Manifold learning is one of the cornerstones of unsupervised learning. Classical methods such as *Isomap* [31], *Locally linear embeddings* [29], *Laplacian eigenmaps* [4] and more [30, 11] all seek a low dimensional embedding of high dimensional data that preserves prespecified aspects of data. Probabilistic methods often view the data manifold as governed by a latent variable along with a generative model that describes how the latent manifold is to be embedded in the data space. The common theme is the quest for a low dimensional representation that faithfully captures the data.

Ideally, we want an *operational representation*, that is we want to be able to make mathematically meaningful calculations with respect to the learned representation. It has been argued [17] that a good representation should at least support the following:

- **Interpolation:** given two points, a natural unique interpolating curve that follows the manifold should exist.

- **Distances:** the distance between two points should be well defined and informally reflect the amount of energy required to transform one point to another.

---

[*]DTU Compute, Richard Petersens Plads, Building 321, DK-2800 Kgs. Lyngby, Denmark.
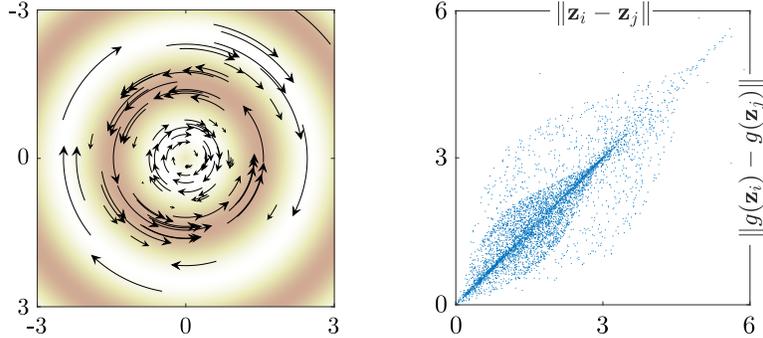
Figure 1: Reparametrizations illustrated. *Left:* A "swirling" transformation of the latent space Z. *Right:* Pair wise distances between points before and after reparametrization; clearly the Euclidean distances change with reparametrizations.

- **Measure:** the representation should be equipped with a measure under which integration is well defined for all points on the manifold.

These are elementary requirements of a representation, but most nonlinear manifold learning schemes do not imply or provide such operations.

In the sequel we will use the following notation. We denote by Z the d-dimensional *representation* or *latent space*, which is learned from data in the *observation space* X. Latent points are denoted $z_i \in Z$, while corresponding observations are $x_i \in X$.

Embedding methods seek a low dimensional embedding $z_{1:N} = \{z_1, \ldots, z_N\} \subset Z$ of the data $x_{1:N} \subset X$. These methods fundamentally only describe the data manifold at the points where data is observed and nowhere else. As such, the low dimensional embedding space is only well defined at $z_{1:N}$. It is common to treat the low dimensional embedding space as equipped with the Euclidean metric, but this is generally a *post hoc* assumption with limited grounding in the embedding method. Fundamentally, the learned representation space is a discrete space that does not lend itself to continuous interpolations. Likewise, the most natural measure will only assign mass to the points $z_{1:N}$, and any associated distribution will be discrete. It is not clear how this can naturally lead to an operational representation.

Generative models estimate a set of low dimensional latent variables $z_{1:N}$ along with a suitable mapping $f : Z \to X$ such that $f(z) \approx x$. It is, again, common to treat the latent space Z together with the Euclidean metric. However, this assumption is unwarranted. As an example, consider the *variational autoencoder (VAE)* [22, 28], which seeks a representation in which $z_{1:N}$ follow a unit Gaussian distribution. Now consider a 2-dimensional latent space and the transformation $g(z) = R_\theta z$, where R_θ is a linear transformation that rotates points by $\theta(z) = \sin(\pi\|z\|)$. This is a smooth invertible transformation with the property that $z \sim \mathcal{N}(0, I) \Rightarrow g(z) \sim \mathcal{N}(0, I)$; see Fig. 1. If the latent variables $z_{1:N}$ and the mapping f is an optimal VAE, then $g(z_{1:N})$ and $f \circ g^{-1}$ is equally optimal. Yet, the latent spaces Z and $g(Z)$ are quite different; Fig. 1 shows the Euclidean distances between pairs of points of the latent space before

and after applying g, for samples drawn from a unit Gaussian. Clearly, the transformed latent space is significantly different from the original space. As the VAE provides no guarantees as to which latent space is recovered, we must be careful when relying on the Euclidean latent space: distances between points are effectively arbitrary, as are straight line interpolations. Ideally, we want a representation that is invariant under such transformations, but current models do not have such properties.

In this paper, we consider probabilistic latent variable models on the form $x = f(z)$ where f is a smooth stochastic process. The latent space can then be endowed with a random Riemannian metric to ensure that the learned latent representation is operational as defined above. We consider a deterministic approximation to the random Riemannian metric, and provide tight approximation bounds for expected distances (Proposition 4.6). The approximation is good when the data is high dimensional, which is often the case in machine learning applications. The analysis justifies the use of deterministic approximations, which in turn lead to computationally tractable algorithms.

The paper is structured to first provide a short primer on (deterministic) Riemannian geometry (Sec. 2). We then extend this class of geometries to the stochastic setting (Sec. 3), and provide our main theoretical contributions (Sec. 4) that analyze to which extend stochastic manifolds are well approximated by deterministic ones. Our analysis holds for any smooth stochastic generative process, which we exemplify (Sec. 5) with Gaussian process latent variable models [24].

2   Riemannian manifolds

A d-dimensional manifold $\mathcal{M}$ embedded in $\mathbb{R}^n$ with $d < n$ is a topological space in which each point $x \in \mathcal{M}$ has a neighborhood that is homeomorphic to $\mathbb{R}^d$ [15]. We may think of $\mathcal{M}$ as a smooth (nonlinear) surface in space that does not self intersect or change dimensionality. At each point $x \in \mathcal{M} \subseteq \mathbb{R}^n$ we have the tangent space $T_x\mathcal{M}$ of $\mathcal{M}$ at x which may be seen as a linear approximation of $\mathcal{M}$ near x. In Fig. 2 a 2-dimensional manifold embedded in $\mathbb{R}^3$ is shown together with part of a tangent space.

Let $f : Z \to \mathcal{M} \subseteq \mathbb{R}^n$ be a parametrization of an open subset $f(Z) \subseteq \mathcal{M}$ defined on some open subset $Z \subseteq \mathbb{R}^d$. A Riemannian metric on Z is an inner product $\langle \cdot, \cdot \rangle_z$ on the tangent spaces $T_zZ \cong T_z\mathbb{R}^d$ which varies smoothly from point to point. Here, *smooth* means infinitely differentiable. Such a metric may be given by a positive definite $(d \times d)$-matrix M_z which depends smoothly on $z \in Z$. The induced inner product is then $\langle v, w \rangle_z = v^T M_z w$ for $v, w \in T_z\mathbb{R}^d$ seen as column vectors.

Consider the standard inner product between points in $\mathbb{R}^n$, $\langle x, x' \rangle = \sum_{i=1}^n x_i x_i'$ where $x = (x_1, \ldots, x_n)$ and $x' = (x_1', \ldots, x_n')$. Let $z \in Z$ and let $\Delta_1, \Delta_2 \in U \subset \mathbb{R}^d$ where U is an open ball centered at the origin such that $z + U \subseteq Z$. Then we can compute the inner product of Δ_1 and Δ_2 at z using the Taylor expansion of f. Consider the scalar product $\langle f(z + \Delta_1) - f(z), f(z + \Delta_2) - f(z) \rangle$, and the linear part of it in Δ_1 and Δ_2
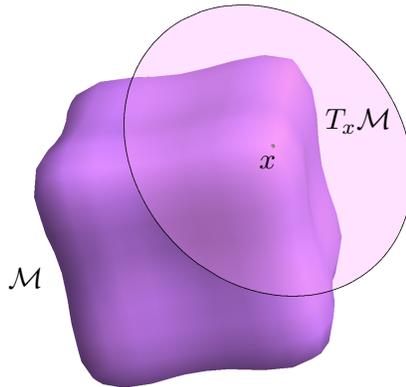
3

Figure 2: Image of an embedded manifold $\mathcal{M}$ and the tangent space $T_x\mathcal{M}$ at a point $x \in \mathcal{M}$.

which we denote by $\langle f(z + \Delta_1) - f(z), f(z + \Delta_2) - f(z) \rangle_0$. Then

$$\langle f(z + \Delta_1) - f(z), f(z + \Delta_2) - f(z) \rangle_0 =$$
$$\langle f(z) + J_f(z)\Delta_1 - f(z), f(z) + J_f(z)\Delta_2 - f(z) \rangle =$$
$$\langle J_f(z)\Delta_1, J_f(z)\Delta_2 \rangle = \Delta_1^\top J_f(z)^\top J_f(z)\Delta_2,$$

where $J_f(z)$ is the $n \times d$ Jacobian matrix of f at z. The $d \times d$ symmetric positive definite matrix $M_z = J_f(z)^\top J_f(z)$ defines a Riemannian metric on Z induced by f which is called the *pullback metric*. Note that the pullback metric corresponds to the Riemannian metric on $\mathcal{M}$ induced by the inner product in the ambient space $\mathbb{R}^n$, which does not depend on the choice of parametrization f. In this way, the pullback metric avoids the parametrization issue discussed in the opening section.

Distances & interpolants. The length of a smooth curve $c : [a, b] \to Z$ under the local inner product is

$$\mathcal{L}(c) = \int_a^b \sqrt{\dot{c}_t^\top M_{c_t} \dot{c}_t} \, dt,$$

where $c_t = c(t)$ and $\dot{c}_t = \partial_t c(t)$ is the curve and its derivative, respectively. Natural interpolants (geodesics) can then be defined as length minimizing curves connecting two points. The length of such a curve is a natural distance measure along the manifold. Unfortunately, minimizing curve length gives rise to a poorly determined optimization problem as the length of a curve is independent of the parametrization. The following proposition provides remedy [15]:

Proposition 2.1. *Let $c : [a, b] \to Z$ be a smooth curve that (locally) minimizes the curve energy*

$$\mathcal{E}(c) = \frac{1}{2} \int_a^b \dot{c}_t^\top M_{c_t} \dot{c}_t \, dt. \tag{1}$$

Then c has constant velocity and is locally length-minimizing.

4

Integration. Given a function $h : \mathbb{R}^n \to \mathbb{R}$ and an open subset $\Omega \subseteq Z$ we can integrate h over $f(\Omega)$ as [26]

$$\int_{f(\Omega)} h(x)\mathrm{d}x = \int_{\Omega} h(f(z))\sqrt{\det(M_z)}\mathrm{d}z.$$

The quantity $\sqrt{\det(M_z)}$ is known as the *Riemannian volume measure* and is akin to the Jacobian determinant in the *change of variables theorem*.

Relation to latent variable models. As stated in the introduction, we are concerned with latent variable models $x = f(z)$, where f is a smooth stochastic process. The above constructions assume that f is deterministic. In this case, Riemannian geometry provides us with the tools to make the representation *operational* as defined in the introduction. This paper is concerned with extending Riemannian geometry to the stochastic domain in order to provide operational representations in latent variable models. To make the resulting constructions practical, we will further show how the stochastic geometry can be approximated well by deterministic geometries that lend themselves to computations.

3 Stochastic Riemannian geometry

In this section we establish the basic definitions in the area of random manifolds. To the best of our knowledge this is a fairly unexplored topic and in our opinion it deserves more attention. Related work has been done in the area of random fields, see for instance [1]. To illustrate the concepts the definitions are followed by some elementary examples. We start by defining random metrics.

Definition 3.1. Let $Z \subseteq \mathbb{R}^d$ be an open subset. A random (or stochastic) Riemannian metric is a matrix-valued random field on Z whose sample paths are Riemannian metrics. We also refer to Z equipped with the random metric as a random manifold.

The stochastic Riemannian metrics considered in this paper are induced by a stochastic process $f : Z \to \mathbb{R}^n$ for some n, where $Z \subseteq \mathbb{R}^d$ is an open subset. More precisely, f is a random field in the sense of [1]. In our examples we will have $Z = \mathbb{R}^d$ but it is useful to keep in mind that the parametrization may not be defined on the whole of $\mathbb{R}^d$. Now, in order for f to induce a random metric it must satisfy some differentiability conditions. One option is to require that any sample path from f is smooth. In that case we say that f is smooth. If in addition the sample paths $s : Z \to \mathbb{R}^n$ are such that the Jacobian matrix J_s has full rank everywhere, f is called a *stochastic immersion*. This implies that s is locally injective and that $s(Z) \subseteq \mathbb{R}^n$ is an immersed submanifold. A stochastic immersion whose sample paths are injective is called a *stochastic embedding*. In this case the sample paths $s : Z \to \mathbb{R}^n$ are embedded submanifolds $s(Z) \subseteq \mathbb{R}^n$. As such, $s(Z)$ has an induced Riemannian metric from the ambient space $\mathbb{R}^n$ and this defines a stochastic Riemannian metric on Z. Another point of view of the basic objects is thus as random embedded submanifolds of $\mathbb{R}^n$. These submanifolds all have the same constant topology and smooth structure induced by Z and hence from an intrinsic point

of view, the random aspect of these manifolds is confined to the Riemannian metric. If the conditions to be a stochastic immersion or embedding are only true with probability 1 we can modify the process by restricting to the measure 1 subset of the probability space where the requirements are fulfilled.

As in the deterministic case, the metric on Z induced by f is given by the pullback metric

$$M = J_f^T J_f,$$

where J_f is the Jacobian of f. In accordance with the definition of stochastic Riemannian metrics, M is a matrix-valued stochastic process parametrized by Z.

If it exists, we may also consider the *expected metric* on Z which is given by the mean value $\mathbb{E}(M) = \mathbb{E}(J_f^T J_f)$.

Definition 3.2. Let $f : Z \to \mathbb{R}^n$ be a stochastic immersion such that all the entries of the matrix $\mathbb{E}(M) = \mathbb{E}(J_f^T J_f)$ are finite smooth functions. We then refer to $\mathbb{E}(M)$ as the expected metric. It defines a Riemannian metric on Z making it into a Riemannian manifold which is called the mean manifold.

Note that the mean manifold is typically *not* given by the mean value of f, that is the map $\mathbb{E}(f) : Z \to \mathbb{R}^n : z \mapsto \mathbb{E}(f(z))$. The expected metric has been previously studied for the *Gaussian process latent variable model* [32], and the *variational autoencoder* [2].

Remark 3.3. The requirement that a stochastic process $f : Z \to \mathbb{R}^n$ has smooth sample paths with probability 1 is not always straightforward to verify. An alternative assumption that is relatively easy to check is that f is differentiable in mean square, see [1]. We say that f is mean square smooth if f has mean square derivatives of any order. Let J_f denote the mean square Jacobian of a mean square smooth process with smooth covariance function. The expected metric may be considered in this setting as well, assuming that $\mathbb{E}(M) = \mathbb{E}(J_f^T J_f)$ has full rank. The main result of this paper, Proposition 4.6, compares expected length on random manifolds to length in the expected metric. This result holds both in the smooth and mean square smooth setting.

Example 3.4. A significant special case is a Gaussian process or Gaussian random field

$$f : \mathbb{R}^d \to \mathbb{R}^n : q \mapsto (f_1(q), \ldots, f_n(q)),$$

with $d \leq n$. Here, the components $f_i : \mathbb{R}^d \to \mathbb{R}$ are Gaussian processes, which means that the vector $(f_i(q_1), \ldots, f_i(q_r))$ is Gaussian for any $q_1, \ldots, q_r \in \mathbb{R}^d$. In fact, that $f : \mathbb{R}^d \to \mathbb{R}^n$ is a Gaussian process means that $\sum_{i=1}^n \alpha_i f_i$ is a Gaussian process for all $\alpha_1, \ldots, \alpha_n \in \mathbb{R}$. We will be concerned with the case where the random vectors $(f_i(q_1), \ldots, f_i(q_r))$ and $(f_j(q_1), \ldots, f_j(q_r))$ are independent for $i \neq j$ and all $q_1, \ldots, q_r \in \mathbb{R}^d$. The distribution of $(f_i(q_1), \ldots, f_i(q_r))$ is determined by the mean function $\mu_i : \mathbb{R}^d \to \mathbb{R} : p \mapsto \mathbb{E}(f_i(p))$ and the covariance function $k_i : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R} : (p, q) \mapsto \mathbb{E}((f_i(p) - \mu_i(p))(f_i(q) - \mu_i(q)))$. If k_i and μ_i are real analytic functions, then f_i is mean

square smooth. This can be seen from the criterion (1.4.9) in [1]. Moreover, for any column vector $v \in \mathbb{R}^d$,

$$v^T \mathbb{E}(J_f^T J_f) v = \mathbb{E}(v^T J_f^T J_f v) = \mathbb{E}(||J_f v||^2) = \sum_{i=1}^{n} \mathbb{E}(\nabla f_i \cdot v)^2 + \sum_{i=1}^{n} \text{var}(\nabla f_i \cdot v),$$

and $\mathbb{E}(\nabla f_i \cdot v) = \nabla \mu_i \cdot v$. Hence $\mathbb{E}(J_f^T J_f)$ has full rank for example if $\mu = (\mu_1, \ldots, \mu_n)$ is locally injective or ∇f_i has non-degenerate covariance matrix for some i.

For Gaussian processes in machine learning, see [24, 27, 32] and in particular the Gaussian process latent variable model.

Example 3.5. Another special case to keep in mind which is particularly simple is a stochastic embedding $f(x) = L(g(x))$, where $g : \mathbb{R}^d \to \mathbb{R}^m$ a deterministic smooth injection and $L : \mathbb{R}^m \to \mathbb{R}^n$ is a random matrix with $n \leq m$. Consider the case where L is given by a distribution on $\text{SO}(\mathbb{R}^m, \mathbb{R}^n)$, that is any sample of L has orthonormal rows. Let W be the span of the rows of L and fix coordinates on W given by the rows of L as a basis. In this way, f can be seen as a random projection of the deterministic manifold $g(\mathbb{R}^d) \subseteq \mathbb{R}^m$ into W. Note that if $n < m$, this can be viewed as random dimensionality reduction. See [7] for a survey on the topic of random projections in machine learning. In this context it is relevant to recall the Johnson-Lindenstrauss lemma. Roughly speaking, given a finite set $E \subset \mathbb{R}^m$, distances in E are well preserved by a random projection $\mathbb{R}^m \to \mathbb{R}^n$ if n is big enough. See for example [10] for background on the Johnson-Lindenstrauss lemma. In the same spirit one can show that distances along the submanifold $g(\mathbb{R}^d) \subseteq \mathbb{R}^m$ are well preserved under similar circumstances [3, 8, 14, 19, 23, 33].

Example 3.6. Consider a simple random manifold given by a stochastic embedding

$$f : \mathbb{R}^d \to \mathbb{R}^n : z \mapsto (f_1(z), \ldots, f_n(z)),$$

where $f_i = \mu_i + \epsilon_i \sigma_i$ for smooth functions $\mu_i, \sigma_i : \mathbb{R}^d \to \mathbb{R}$ and random variables ϵ_i with $\mathbb{E}(\epsilon_i) = 0$ and $\mathbb{E}(\epsilon_i^2) = 1$. Let $\mu : \mathbb{R}^d \to \mathbb{R}^n : z \mapsto (\mu_1(z), \ldots, \mu_n(z))$, $\sigma : \mathbb{R}^d \to \mathbb{R}^n : z \mapsto (\sigma_1(z), \ldots, \sigma_n(z))$ and define

$$e = \begin{pmatrix} \epsilon_1 & 0 & 0 & \ldots & 0 \\ 0 & \epsilon_2 & 0 & \ldots & 0 \\ & & \vdots & & \\ 0 & 0 & 0 & \ldots & \epsilon_n \end{pmatrix}.$$

Then $f = \mu + e\sigma$, where the product $e\sigma$ is matrix multiplication and μ and σ are viewed as column vectors. Note that $\mathbb{E}(e) = 0$ and $\mathbb{E}(e^2) = \text{Id}_{\mathbb{R}^n}$, that is $\mathbb{E}(e)$ is the zero matrix and $\mathbb{E}(e^2)$ the identity matrix on $\mathbb{R}^n$. For the Jacobian of f we get $J_f = J_\mu + eJ_\sigma$ and hence

$$M = J_f^T J_f = (J_\mu + eJ_\sigma)^T (J_\mu + eJ_\sigma) = J_\mu^T J_\mu + J_\mu^T eJ_\sigma + J_\sigma^T eJ_\mu + J_\sigma^T e^2 J_\sigma.$$

Therefore $\mathbb{E}(J_f^T J_f) = J_\mu^T J_\mu + J_\mu^T \mathbb{E}(e) J_\sigma + J_\sigma^T \mathbb{E}(e) J_\mu + J_\sigma^T \mathbb{E}(e^2) J_\sigma = J_\mu^T J_\mu + J_\sigma^T J_\sigma$. This is also the metric on $\mathbb{R}^d$ induced by the embedding $g : \mathbb{R}^d \to \mathbb{R}^{2n} = \mathbb{R}^n \times \mathbb{R}^n$ given by

$$g : \mathbb{R}^d \to \mathbb{R}^n \times \mathbb{R}^n : z \mapsto (\mu(z), \sigma(z)).$$

The metric induced by g has been studied empirically for variational autoencoders [2]. We see in this example how the mean manifold $g(\mathbb{R}^d) \subseteq \mathbb{R}^{2n}$ depends on both the mean value $\mu = \mathbb{E}(f)$ and the standard deviation map σ.

Another point of view is that $f = L \circ g$ where

$$L = (\mathrm{Id}_{\mathbb{R}^n} \mid e)$$

is the random $(n \times 2n)$-matrix given by stacking the columns of the matrices $\mathrm{Id}_{\mathbb{R}^n}$ and e side by side. This is a special case of the random manifolds considered in Ex. 3.5. In connection with Ex. 3.7 below we note that $\mathbb{E}(L^T L) = \mathrm{Id}_{\mathbb{R}^{2n}}$.

Example 3.7. Recall Ex. 3.5 and random manifolds given by $f = L \circ g$ where $g : \mathbb{R}^d \to \mathbb{R}^m$ is an embedding and $L : \mathbb{R}^m \to \mathbb{R}^n$ is a random matrix. Assume that $\mathbb{E}(L^T L)$ has finite entries and full rank. We shall see that, similarly to Ex. 3.6, the mean manifold is essentially the manifold $g(\mathbb{R}^d) \subseteq \mathbb{R}^m$. Since $\mathbb{E}(L^T L)$ is symmetric we may diagonalize $\mathbb{E}(L^T L)$ with an orthogonal $(m \times m)$-matrix P, that is $\mathbb{E}(L^T L) = P^T D P$ where D is diagonal. Since $\mathbb{E}(L^T L)$ is positive definite we can take the square roots of the eigenvalues on the diagonal of D to obtain a diagonal matrix $\sqrt{D}$. In this case the stochastic metric is $M = J_f^T J_f = J_g^T L^T L J_g$ and the mean metric is $\mathbb{E}(M) = J_g^T \mathbb{E}(L^T L) J_g$. Hence the mean manifold is induced by the embedding $\mathbb{R}^d \to \mathbb{R}^m : z \mapsto \sqrt{D} P g(z)$, which is just the manifold $g(\mathbb{R}^d) \subseteq \mathbb{R}^m$ up to a change of coordinates. In contrast consider $\mathbb{E}(f) : \mathbb{R}^d \to \mathbb{R}^n$, and note that $\mathbb{E}(f) = \mathbb{E}(L) \circ g$. The image of $\mathbb{E}(f)$ is thus the image of the embedded mean manifold $\sqrt{D} P g(\mathbb{R}^d) \subseteq \mathbb{R}^m$ under the linear map $\mathbb{E}(L)(\sqrt{D} P)^{-1} : \mathbb{R}^m \to \mathbb{R}^n$.

3.1 Expected length versus the expected metric

Let $Z \subseteq \mathbb{R}^d$ be an open subset. We will now explore the topic of shortest paths on a random manifold induced by a stochastic immersion $f : Z \to \mathbb{R}^n$ with expected metric $\mathbb{E}(M)$. In order to talk about expected length of curves on the random manifold we need that f is measurable when viewed as a map $\Omega \times Z \to \mathbb{R}^n$, where Ω is a probability space. Here, Z and $\mathbb{R}^n$ are endowed with the Lebesgue measure and $\Omega \times Z$ with the product measure. Now let $c : [a, b] \to Z$ denote a smooth immersed curve and consider its stochastic immersion $f \circ c$ in $\mathbb{R}^n$. We stress that c is a deterministic curve in Z, while $f \circ c$ is a random curve in $\mathbb{R}^n$. The energy of c, defined as in Eq. 1, is a random quantity and it is natural to consider its expectation with respect to the random metric. Since the energy integrand is positive, Tonelli's Theorem tells us that the expected energy

$\epsilon(c)$ is given by

$$\epsilon(c) = \mathbb{E}(\mathcal{E}(c)) = \frac{1}{2}\mathbb{E}\left(\int_a^b \dot{c}_t^\top M_{c_t}\dot{c}_t \mathrm{d}t\right)$$
$$= \frac{1}{2}\int_a^b \dot{c}_t^\top \mathbb{E}(M_{c_t})\dot{c}_t \mathrm{d}t.$$

This implies that a curve c with minimal expected energy over the stochastic manifold is a geodesic under the deterministic Riemannian metric $\mathbb{E}(M)$.

We can understand a curve with minimal expected energy in more explicit terms as follows. Let $u_t = \mathbb{E}(\|\dot{c}_t\|)$ and $v_t = 1$ denote two functions over the interval $[a, b]$; here we use the shorthand notation $\|\dot{c}_t\| = \sqrt{\dot{c}_t^\top M_{c_t}\dot{c}_t}$. The Cauchy-Schwartz inequality then gives

$$\left(\int_a^b \mathbb{E}(\|\dot{c}_t\|)\mathrm{d}t\right)^2 \leq \int_a^b \mathbb{E}(\|\dot{c}_t\|)^2\mathrm{d}t \cdot \int_a^b \mathrm{d}t$$
$$= (b-a)\int_a^b \mathbb{E}(\|\dot{c}_t\|)^2\mathrm{d}t.$$

Let $l(c) = \mathbb{E}(\mathcal{L}(c))$ denote the expected length of c. Then

$$\int_a^b \mathbb{E}(\|\dot{c}_t\|)^2\mathrm{d}t \geq \frac{l^2(c)}{b-a}. \tag{2}$$

Equality is achieved when u_t and v_t are parallel, that is when $\mathbb{E}(\|\dot{c}_t\|)$ is constant. If the curve is *regular* in the sense that the expected speed $\mathbb{E}(\|\dot{c}_t\|)$ is non-zero for all t, then we can always reparametrize c_t to have constant expected speed and achieve equality in Eq. 2. Since $\mathrm{var}(\|\dot{c}_t\|) = \mathbb{E}(\|\dot{c}_t\|^2) - \mathbb{E}(\|\dot{c}_t\|)^2$, we see that

$$\int_a^b \mathbb{E}(\|\dot{c}_t\|)^2\mathrm{d}t = \int_a^b \mathbb{E}(\|\dot{c}_t\|^2)\mathrm{d}t - \int_a^b \mathrm{var}(\|\dot{c}_t\|)\mathrm{d}t$$
$$= 2\epsilon(c) - \int_a^b \mathrm{var}(\|\dot{c}_t\|)\mathrm{d}t.$$

Assuming that the curve has constant expected speed, we then get

$$\epsilon(c) = \frac{l^2(c)}{2(b-a)} + \frac{1}{2}\int_a^b \mathrm{var}(\|\dot{c}_t\|)\mathrm{d}t.$$

Minimizing expected curve energy, thus, does not always minimize the expected curve length. Rather, this balances the minimization of expected curve length and the minimization of curve variance.

9

4 Expected length in high codimension

Let $Z \subseteq \mathbb{R}^d$ be an open subset and $f : Z \to \mathbb{R}^n$ a stochastic immersion with expected metric $\mathbb{E}(M)$. Any smooth curve $c : [0,1] \to Z$ gives rise to a random curve $f \circ c : [0,1] \to \mathbb{R}^n$. In this section we continue to examine the relationship between the expected length of random curves and the length with respect to the expected metric. More precisely, we show that length with respect to the expected metric is a good approximation for expected length in high ambient dimension n. Assuming independence of the components of $f \circ c$ we could apply a version of the central limit theorem such as the Berry-Esseen theorem to this problem, see for example [12]. We found it more convenient to take a direct approach via the Taylor expansion of the norm of velocity vectors. See also [21, 5] and Chapter 27 of [9] for approximation results in the same vein.

4.1 Expected norm of high dimensional vectors

We will consider a sequence W_n of random vectors in $\mathbb{R}^n$. This means that for each integer $1 \leq n$ we have an $\mathbb{R}^n$-valued random variable W_n. Consider the norm $w_n = ||W_n||$. Assume that w_n^2 has first and second moments and put $m_n = \sqrt{\mathbb{E}(w_n^2)}$ and $\Sigma_n = \sqrt{n \cdot \mathrm{var}(w_n^2)}$. We say that m_n is bounded away from 0 if there is a constant $0 < b$ such that $b < m_n$ for all n. For $1 \leq k$, let $\mu_k(w_n^2) = \mathbb{E}((w_n^2 - m_n^2)^k)$ denote the k-th central moment of w_n^2.

Definition 4.1. We call a sequence W_n as above balanced if m_n is bounded away from 0 and m_n, $n^2 \mu_3(w_n^2)$ and $n^2 \mu_4(w_n^2)$ are bounded sequences.

Suppose that W_n is balanced. We shall see that Σ_n is a bounded sequence. Let $Z_n = n(w_n^2 - m_n^2)^2$ and let $\mathcal{X}_n$ be the indicator function for the event $\{Z_n < 1\}$. We know that $\mathbb{E}(Z_n^2) = n^2 \mu_4(w_n^2)$ is bounded and need to show that $n\mu_2(w_n^2) = \mathbb{E}(Z_n)$ is bounded. This follows from $\mathbb{E}(Z_n) \leq 1 + \mathbb{E}(Z_n(1 - \mathcal{X}_n)) \leq 1 + \mathbb{E}(Z_n^2)$.

Remark 4.2. Consider the cubic Taylor polynomial of $\sqrt{x}$ around $x = 1$, $P(x) = 1 + \frac{1}{2}(x-1) - \frac{1}{8}(x-1)^2 + \frac{1}{16}(x-1)^3$. We will show that $|\sqrt{x} - P(x)| \leq \frac{5}{16}(x-1)^4$ for all $x \geq 0$. Let $y = \sqrt{x}$ and put $Q(y) = P(y^2) - y$. Since $16Q(y) = (y-1)^4(y^2 + 4y + 5)$, we have that $Q(y) \geq 0$ for $y \geq 0$. Hence $P(x) - \sqrt{x} \geq 0$ for $x \geq 0$. Let $R(y) = \frac{5}{16}(y^2-1)^4 - P(y^2) + y$. Since $16R(y) = y(y-1)^4(5y^3 + 20y^2 + 29y + 16) \geq 0$ for $y \geq 0$, we have that $\frac{5}{16}(x-1)^4 \geq P(x) - \sqrt{x}$ for $x \geq 0$.

Proposition 4.3. *Let W_n be a balanced sequence of vectors. Then*

$$\mathbb{E}(w_n) = m_n - \frac{\Sigma_n^2}{8n \cdot m_n^3} + \mathcal{O}(n^{-2}).$$

Proof. We will first prove the statement assuming that $m_n = 1$ for all n. In this case we need to show that $\mathbb{E}(w_n) - (1 - \Sigma_n^2/8n) = \mathcal{O}(n^{-2})$. Let $P(x) = 1 + \frac{1}{2}(x-1) - \frac{1}{8}(x-1)^2 + \frac{1}{16}(x-1)^3$ be the cubic Taylor polynomial of $\sqrt{x}$ around $x = 1$ and note

that $\mathbb{E}(P(w_n^2)) = 1 - \Sigma_n^2/8n + \frac{1}{16}\mu_3(w_n^2)$. Since by assumption $\mu_3(w_n^2) = \mathcal{O}(n^{-2})$, it is enough to show that $\mathbb{E}(w_n) - \mathbb{E}(P(w_n^2)) = \mathcal{O}(n^{-2})$. Note that

$$|\mathbb{E}(w_n) - \mathbb{E}(P(w_n^2))| = |\mathbb{E}(w_n - P(w_n^2))| \leq \mathbb{E}(|w_n - P(w_n^2)|).$$

Also, $|\sqrt{x} - P(x)| \leq \frac{5}{16}(x-1)^4$ for all $x \geq 0$ by Remark 4.2. Hence $|w_n - P(w_n^2)| \leq (w_n^2 - 1)^4$ and by assumption we have $\mathbb{E}((w_n^2 - 1)^4) = \mathcal{O}(n^{-2})$. It follows that $\mathbb{E}(|w_n - P(w_n^2)|) \leq \mathbb{E}((w_n^2 - 1)^4) = \mathcal{O}(n^{-2})$.

Now consider the general case of a balanced sequence W_n and let m_n and Σ_n denote the corresponding sequences associated to W_n. Since W_n/m_n is balanced, $\mathbb{E}(||W_n/m_n||^2) = 1$ for all n and $n \cdot \text{var}(||W_n/m_n||^2) = \Sigma_n^2/m_n^4$, we have by above that

$$\mathbb{E}(w_n/m_n) = 1 - \frac{\Sigma_n^2/m_n^4}{8n} + \mathcal{O}(n^{-2}).$$

But m_n is bounded and so the claim follows by multiplying by m_n. $\qquad\square$

Remark 4.4. Suppose that W_n is balanced. Since Σ_n is bounded and m_n bounded away from 0 we have that $\limsup_{n\to\infty} \Sigma_n < \infty$ and $\liminf_{n\to\infty} m_n > 0$. Let $A, b \in \mathbb{R}$ be such that $A > \limsup_{n\to\infty} \Sigma_n$ and $0 < b < \liminf_{n\to\infty} m_n$. Then, by Proposition 4.3,

$$0 \leq m_n - \mathbb{E}(w_n) \leq \frac{A^2}{8n \cdot b^3}, \tag{3}$$

for large enough n. In particular, if $\Sigma_n \to \Sigma$ and $m_n \to m$ with $\Sigma, m \in \mathbb{R}$, then Eq. 3 holds for any $A > \Sigma$ and $0 < b < m$. In this case we also have that

$$\lim_{n\to\infty} \mathbb{E}(w_n) = \lim_{n\to\infty} m_n \tag{4}$$

and Proposition 4.3 gives some additional information concerning the rate of convergence. Note also that $\text{var}(w_n) = m_n^2 - \mathbb{E}(w_n)^2 \to 0$ for balanced W_n.

4.2 Normed sequences of independent random variables

For $k \geq 1$ and a random variable Y with first moment $\mathbb{E}(Y)$, let $\mu_k(Y) = \mathbb{E}((Y - \mathbb{E}(Y))^k)$ denote the k-th central moment. Now consider a sequence $X_1, X_2, \ldots$ of independent random variables and let

$$W_n = (X_1/\sqrt{n}, \ldots, X_n/\sqrt{n})$$

be the corresponding normalized sequence of vectors. We say that the sequence $X_1, X_2, \ldots$ has *bounded moments* if the moments $\mathbb{E}(X_i^k)$ form a bounded sequence for any $k \geq 1$. This implies that the central moments $\mu_k(X_i)$ are bounded as well. Let $w_n = ||W_n||$. The definition of a balanced sequence of vectors W_n is motivated by this setup since, for instance, in this case $m_n^2 = \mathbb{E}(w_n^2) = \frac{1}{n}\sum_{i=1}^n \mathbb{E}(X_i^2)$ is bounded if the sequence $X_1, X_2, \ldots$ has bounded moments. Similarly, $\Sigma_n^2 = n \cdot \text{var}(w_n^2) = n \cdot \frac{1}{n^2}\sum_{i=1}^n \text{var}(X_i^2)$ is bounded in this case.

11

Lemma 4.5. *Let $X_1, X_2, \ldots$ be independent with bounded moments. If m_n is bounded away from 0 then $W_n = (X_1/\sqrt{n}, \ldots, X_n/\sqrt{n})$ is balanced.*

Proof. We need to show that $n^2 \mu_3(w_n^2)$ and $n^2 \mu_4(w_n^2)$ are bounded. Note that $n^2 \mu_3(w_n^2) = n^2 \mu_3(\frac{1}{n} \sum_{i=1}^{n} X_i^2) = \frac{1}{n} \sum_{i=1}^{n} \mu_3(X_i^2)$, which is bounded. Similarly,

$$n^2 \mu_4(w_n^2) = \frac{1}{n^2} \mu_4(\sum_{i=1}^{n} X_i^2) = \frac{1}{n^2} \sum_{i=1}^{n} \mu_4(X_i^2) + \frac{6}{n^2} \sum_{i<j} \mu_2(X_i^2)\mu_2(X_j^2),$$

which is bounded as well. $\qquad\square$

4.3 Expected length of curves

Consider a sequence of stochastic processes $f_1, f_2, \ldots$ defined on $[0,1]$ such that for any $t \in [0,1]$, $f_1'(t), f_2'(t), \ldots$ are independent. Let

$$\phi_n : [0,1] \to \mathbb{R}^n : t \mapsto (f_1(t)/\sqrt{n}, \ldots, f_n(t)/\sqrt{n})$$

and assume that ϕ_n is a stochastic immersion or a mean square smooth process on $(0,1)$. As in Sec. 3.1, we also assume that ϕ_n has an expected metric and is measurable when seen as a map $\Omega \times [0,1] \to \mathbb{R}^n$, where Ω is a probability space. Furthermore, suppose that the sequence $f_1', f_2', \ldots$ has *uniformly bounded moments* in the sense that for any $k \geq 1$, there is a constant C_k such that $|\mathbb{E}(f_i'(t)^k)| \leq C_k$ for all i and $t \in [0,1]$. Let $w_n(t) = ||\phi_n'(t)||$, $m_n(t) = \sqrt{\mathbb{E}(w_n^2(t))}$ and $\Sigma_n(t) = \sqrt{n \cdot \mathrm{var}(w_n^2(t))}$. Then $\sup_{n,t} \Sigma_n(t) < \infty$. Furthermore, we assume that m_n is uniformly bounded away from 0, meaning that $0 < \inf_{n,t} m_n(t)$. Let L_n denote the length of ϕ_n in the expected metric and l_n the expected length of ϕ_n. In other words

$$L_n = \int_0^1 \mathbb{E}(||\phi_n'(t)||^2)^{1/2} \, dt, \quad l_n = \mathbb{E}\left(\int_0^1 ||\phi_n'(t)|| \, dt\right) = \int_0^1 \mathbb{E}(||\phi_n'(t)||) \, dt. \quad (5)$$

Let $\sup_{n,t} \Sigma_n(t) < A$ and $0 < b < \inf_{n,t} m_n(t)$. By Lemma 4.5 and Eq. 3, for any $t \in [0,1]$ there is a $N_t > 0$ such that $m_n(t) - \mathbb{E}(w_n(t)) \leq A^2/(8n \cdot b^3)$ for all $n > N_t$. In fact, due to the uniform bounds on the moments $\mathbb{E}(f_i'(t)^k)$ we have that for large enough n, $m_n(t) - \mathbb{E}(w_n(t)) \leq A^2/(8n \cdot b^3)$ for all $t \in [0,1]$.

Proposition 4.6. *With $\sup_{n,t} \Sigma_n(t) < A$ and $0 < b < \inf_{n,t} m_n(t)$,*

$$0 \leq \frac{L_n - l_n}{L_n} \leq \frac{A^2}{8nb^4},$$

for large enough n.

Proof. Since for large enough n, $0 \leq m_n(t) - \mathbb{E}(w_n(t)) \leq A^2/(8n \cdot b^3)$ for all $t \in [0,1]$, integrating both sides over $[0,1]$ gives $0 \leq L_n - l_n \leq A^2/(8nb^3)$. Divide by b and note that $b < L_n$ for all n. $\qquad\square$

This result implies that in high codimension we can minimize expected energy instead of minimizing expected length. A curve minimizing expected energy can be found by computing the expected metric and using standard tools from differential geometry to recover a geodesic associated with this metric. It it interesting to note that the length in the expected metric bounds the expected length from above. Consequently, by minimizing expected energy we minimize an upper bound on expected length. Such notions are standard in *variational inference* [6].

Remark 4.7. Expected speed and expected length of random curves are quite natural quantities to consider. For instance, minimal expected length is an interesting candidate for a distance measure along random manifolds. What about simply taking the expectation value $\mathbb{E}(\phi_n)$ component wise and considering the length of this curve? This seems natural enough but is not enough to capture the notion of expected length. The velocity of the curve $\mathbb{E}(\phi_n)$ at $t \in [0,1]$ is $||\mathbb{E}(\phi_n'(t))||$. Let $t \in [0,1]$ and assume for simplicity that $\mathbb{E}(||\phi_n'(t)||)$ converges as $n \to \infty$. By Remark 4.4, $\mathbb{E}(||\phi_n'(t)||^2)$ converges as well and

$$\lim_{n \to \infty} \mathbb{E}(||\phi_n'(t)||^2) = \lim_{n \to \infty} \mathbb{E}(||\phi_n'(t)||)^2.$$

Also, $\mathbb{E}(||\phi_n'(t)||^2) = \frac{1}{n} \sum_i \mathbb{E}(f_i'^2(t)) = \frac{1}{n} \sum_i \operatorname{var}(f_i'(t)) + ||\mathbb{E}(\phi_n'(t))||^2$. Thus,

$$\lim_{n \to \infty} ||\mathbb{E}(\phi_n'(t))|| \neq \lim_{n \to \infty} \mathbb{E}(||\phi_n'(t)||)$$

unless $\frac{1}{n} \sum_i \operatorname{var}(f_i'(t)) \to 0$ as $n \to \infty$.

5 Gaussian processes

We will now have a closer look at the case of Gaussian processes [27].

5.1 Definitions

For a smooth function $h : \mathbb{R}^d \to \mathbb{R} : (p_1, \ldots, p_d) \mapsto h(p_1, \ldots, p_d)$ we will use the notation

$$h_{p_{i_1}, \ldots, p_{i_j}} = \frac{\partial^j h}{\partial p_{i_1} \ldots \partial p_{i_j}},$$

where $i_1, \ldots, i_j \in \{1, \ldots, d\}$. A Gaussian process $f : \mathbb{R}^d \to \mathbb{R}$ is a stochastic process such that for any $q_1, \ldots, q_r \in \mathbb{R}^d$, $(f(q_1), \ldots, f(q_r))$ is a Gaussian vector. The distribution of the vector $(f(q_1), \ldots, f(q_r))$ is determined by the mean function $\mu = \mathbb{E}(f) : \mathbb{R}^d \to \mathbb{R} : p \mapsto \mathbb{E}(f(p))$ and covariance function

$$k : \mathbb{R}^d \times \mathbb{R}^d \to \mathbb{R} : (p, q) \mapsto \operatorname{covar}(f(p), f(q)).$$

The function k is also known as the kernel of f. We will assume that μ and k are real analytic functions. As explained in Ex. 3.4, this implies that f is mean square smooth. For such a Gaussian process $f : \mathbb{R}^d \to \mathbb{R}$, the partial derivative f_{p_i} with respect to p_i

for $i \in \{1, \dots, d\}$ is a Gaussian process with mean function μ_{p_i} and kernel k_{p_i, q_i}. More generally, for $p, q \in \mathbb{R}^d$, $f_{p_i}(p)$ and $f_{p_j}(q)$ have covariance

$$\operatorname{covar}(f_{p_i}(p), f_{p_j}(q)) = k_{p_i, q_j}(p, q),$$

see [27]. Gaussian processes $f_1, \dots, f_m : \mathbb{R}^d \to \mathbb{R}$ are *independent* if

$$(f_1(q_1), \dots, f_1(q_r)), \dots, (f_m(q_1), \dots, f_m(q_r))$$

are independent for all $q_1, \dots, q_r \in \mathbb{R}^d$. A Gaussian process $f : \mathbb{R}^d \to \mathbb{R}^m : p \mapsto (f_1(p), \dots, f_m(p))$ is a stochastic process such that $\sum_{i=1}^m \alpha_i f_i$ is a Gaussian process for all $\alpha_1, \dots, \alpha_m \in \mathbb{R}$. We will call such a process *symmetric* if $f_1, \dots, f_m$ are independent and all have the same kernel.

5.2 Gaussian process latent variable models

Gaussian processes are used in machine learning in the context of Gaussian process latent variable models (GPLVMs), see [24, 27, 32]. In GPLVMs we consider a Gaussian process prior $F : \mathbb{R}^d \to \mathbb{R}^m$ which is symmetric, has zero mean and whose kernel is of a particular form, depending on finitely many hyper parameters. A common choice is a Radial Basis Function (RBF) kernel

$$K(p, q) = \sigma_0^2 e^{-\frac{1}{2l^2}\|p - q\|^2}$$

with variance σ_0^2 and length scale $l > 0$. Another example is the linear covariance case where the kernel is given by the Euclidean scalar product: $K(p, q) = p^T q$ where $p, q \in \mathbb{R}^d$ are seen as column vectors. In this case, the GPLVM reduces to probabilistic principal component analysis [24]. Typically, the data is assumed to be observed with additive iid Gaussian noise with variance σ_1^2, which introduces one more hyper parameter for the model. Given a finite set of data points $Y \subset \mathbb{R}^m$ we solve a maximum likelihood type problem to compute the hyper parameters of the model as well as a set of latent points $X \subseteq \mathbb{R}^d$ corresponding to the data. This can be done for example using the software package [16]. Combined with Gaussian process regression, the result is a Gaussian process posterior $f : \mathbb{R}^d \to \mathbb{R}^m : p \mapsto f(p) = (f_1(p), \dots, f_m(p))$ which fits the data. The process f is a symmetric Gaussian process with mean and kernel given below.

For matrices $A \in \mathbb{R}^{d \times r}$ and $B \in \mathbb{R}^{d \times s}$ with columns $A = (a_1, \dots, a_r)$ and $B = (b_1, \dots, b_s)$ we will use $K(A, B)$ to denote the matrix given by $\{K(A, B)\}_{i,j} = K(a_i, b_j)$. Assume that the data points are distinct, let $N = |X| = |Y|$ be the number of data points and consider $Y \in \mathbb{R}^{m \times N}$ and $X \in \mathbb{R}^{d \times N}$. Let $R(X) = (K(X, X) + \sigma_1^2 \operatorname{Id}_{\mathbb{R}^N})$ and assume that $R(X)$ is invertible. This is the case if the kernel K is positive definite in the sense that $K(X, X)$ is positive definite, or K is semi-positive definite and $\sigma_1 \neq 0$. The mean μ and kernel k of the posterior process are then given by $\mu(p) = Y R(X)^{-T} K(X, p)$ and $k(p, q) = K(p, q) - K(p, X) R(X)^{-1} K(X, q)$, see [27] for details.

Let $f : \mathbb{R}^d \to \mathbb{R}^m$ be the posterior of a GPLVM. If the kernel of the prior is real analytic then so is the kernel and mean of f and hence f is mean square smooth. Let J_f denote the mean square Jacobian. For $p \in \mathbb{R}^d$, the metric induced by f at p is $J_f(p)^T J_f(p)$, which follows a non-central Wishart distribution, see [25].

Example 5.1. Consider a GPLVM $f : \mathbb{R}^d \to \mathbb{R}^m$ with prior kernel $K(p, q) = p^T q$ where $p, q \in \mathbb{R}^d$ are column vectors. As we shall see, the expected metric is a constant matrix in this case. This means that the mean manifold is flat and that geodesics are straight lines for this choice of kernel for the prior. Let $\mu = \mathbb{E}(f) : \mathbb{R}^d \to \mathbb{R}^m$ and note that $\mu(p) = YR(X)^{-T}X^T p$ with notation as above. Hence $\mathbb{E}(J_f) = J_\mu = YR(X)^{-T}X^T$ is constant. Differentiating the posterior kernel we get that the expected metric is given by

$$\mathbb{E}(J_f^T J_f) = \mathbb{E}(J_f)^T \mathbb{E}(J_f) + m \cdot (\mathrm{Id}_d - XR(X)^{-1}X^T).$$

5.2.1 Empirical illustration

Consider the posterior process $f : \mathbb{R}^d \to \mathbb{R}^m$ of a GPLVM with $f(p) = (f_1(p), \ldots, f_m(p))$ and its projections $\Phi_n : \mathbb{R}^d \to \mathbb{R}^n : p \mapsto (f_1(p), \ldots, f_n(p))$ for $n \leq m$. Given a curve $c : [0, 1] \to \mathbb{R}^d$ we acquire a sequence of Gaussian curves $\phi_n = \Phi_n \circ c : [0, 1] \to \mathbb{R}^n$ up to $n = m$.

Example 5.2. To illustrate the results of Sec. 4.3, consider the 120×120-pixel image of a bird in Fig. 3a. Images of this resolution may be seen as points in $\mathbb{R}^m$, where $m = 120^2 = 14400$. We produce a sequence of N points in $\mathbb{R}^m$ by rotating the image (using interpolation) by an angle $2\pi k/N$ for $k = 0, \ldots, N - 1$, see Fig. 3.

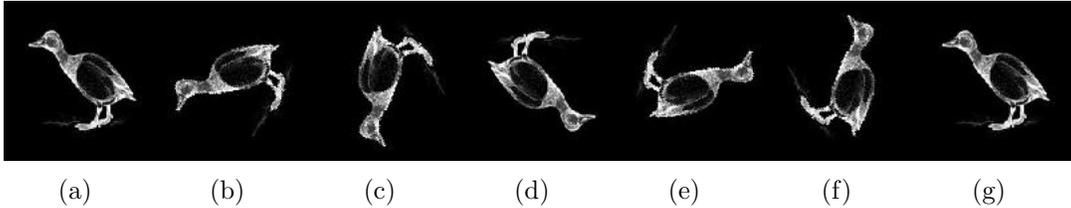| (a) | (b) | (c) | (d) | (e) | (f) | (g) |

Figure 3: A rotated image.

Feeding these to a GPLVM with RBF kernel and d-dimensional latent space we obtain a Gaussian process $f : \mathbb{R}^d \to \mathbb{R}^m$ together with a sequence of N latent points $X \subset \mathbb{R}^d$. For any $n \leq m$ we have a Gaussian process $\Phi_n : \mathbb{R}^d \to \mathbb{R}^n$ given by projection onto the n first coordinates.

Let $c : [0, 1] \to \mathbb{R}^d$ be the line segment joining the first two points of X and put $\phi_n = \Phi_n \circ c$. Let l_n and L_n be given as in Eq. 5. Using [16] with $d = 6$ and $N = 100$ we have estimated l_n and L_n empirically by sampling the posterior process. Fig. 4a displays the relative error $(L_n - l_n)/L_n$ as a function of data dimension n. We have also included the graph of a reference function $h(n) = A^2/(8nb^4)$ for empirical estimates of constants A and b as in Proposition 4.6. The difference between $(L_n - l_n)/L_n$ and $h(n)$ is plotted in Fig. 4b. This illustrates Proposition 4.6 as the theory matches the empirical study.

15
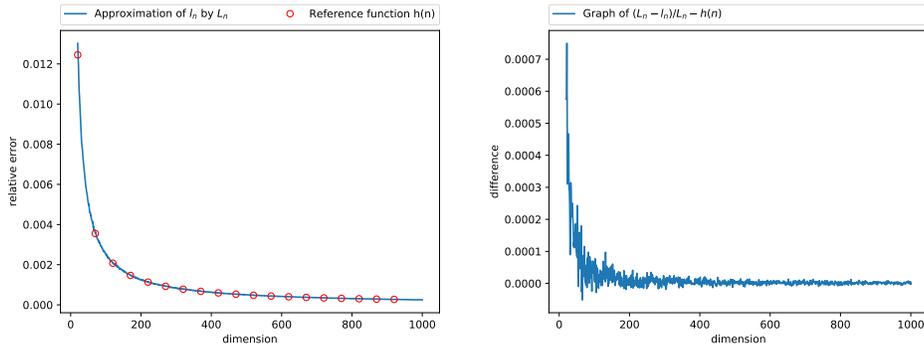
(a) Plots of $(L_n - l_n)/L_n$ and $h(n)$. (b) Plot of $(L_n - l_n)/L_n - h(n)$.

Figure 4

6 Concluding remarks

Starting from the goal of learning an *operational representation*, we have studied a general class of generative latent variable models $x = f(z)$, where f is a smooth stochastic process. The latent space can here be endowed with a random Riemannian metric, such that elementary operations (interpolations, distances, integration, etc.) can be defined in a way that is invariant to reparametrizations of the model.

Mathematically, this is a natural approach, but it does not lend itself easily to computations as computational tools do not exist for working with random Riemannian manifolds. In this paper we have provided a deterministic approximation to this large class of random Riemannian metrics, and provided tight approximation bounds. In particular, it is worth noting that the bound is very tight when data is high dimensional, which is the common case for machine learning. Within this deterministic approximation, we can apply standard tools for computations over Riemannian manifolds [18, 20, 13], and thereby realize the idea of operational representation learning.

Acknowledgments

References

[1] R. J. Adler and J. E. Taylor. *Random Fields and Geometry.* Springer, 2007 edition, 2007.

16

[2] G. Arvanitidis, L. K. Hansen, and S. Hauberg. Latent space oddity: on the curvature of deep generative models. *International Conference on Learning Representations (ICLR)*, 2018.

[3] R. G. Baraniuk and M. B. Wakin. Random projections of smooth manifolds. *Foundations of Computational Mathematics*, 9:51–77, 2009.

[4] Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. *Neural Computation*, 15(6):1373–1396, June 2003.

[5] G. Biau and D. M. Mason. High-dimensional p-norms. *arXiv:1311.0587*, 2013.

[6] David M Blei, Alp Kucukelbir, and Jon D McAuliffe. Variational inference: A review for statisticians. *Journal of the American Statistical Association*, 112(518):859–877, 2017.

[7] Avrim Blum. Random projection, margins, kernels, and feature-selection. In Craig Saunders, Marko Grobelnik, Steve Gunn, and John Shawe-Taylor, editors, *Subspace, Latent Structure and Feature Selection*, pages 52–68. Springer Berlin Heidelberg, 2006.

[8] Kenneth L. Clarkson. Tighter bounds for random projections of manifolds. In *Proceedings of the Twenty-fourth Annual Symposium on Computational Geometry*, SCG '08, pages 39–48, New York, NY, USA, 2008. ACM.

[9] H. Cramér. *Mathematical methods of statistics*. Asia Publishing House, 1962.

[10] Sanjoy Dasgupta and Anupam Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures & Algorithms*, 22(1):60–65, 2003.

[11] David L. Donoho and Carrie Grimes. Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data. *Proceedings of the National Academy of Sciences*, 100(10):5591–5596, 2003.

[12] Carl-Gustav Esseen. On the remainder term in the central limit theorem. *Ark. Mat.*, 8(1):7–15, 11 1969.

[13] Oren Freifeld, Søren Hauberg, and Michael J. Black. Model transport: Towards scalable transfer learning on manifolds. In *Proceedings IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, Columbus, Ohio, USA, June 2014.

[14] Yoav Freund, Sanjoy Dasgupta, Kabra Mayank, and Nakul Verma. Learning the structure of manifolds using random projections. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 473–480. Curran Associates, Inc., 2008.

[15] Sylvestre Gallot, Dominique Hulin, and Jacques Lafontaine. *Riemannian geometry*, volume 3. Springer, 1990.

[16] GPy. GPy: A Gaussian process framework in python. `http://github.com/SheffieldML/GPy`, 2012.

[17] S. Hauberg. Only bayes should learn a manifold (on the estimation of differential geometric structure from data). *arXiv:1806.04994*, 2018.

[18] Søren Hauberg, Oren Freifeld, and Michael J. Black. A geometric take on metric learning. In P. Bartlett, F.C.N. Pereira, C.J.C. Burges, L. Bottou, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems (NIPS) 25*, pages 2033–2041. MIT Press, 2012.

[19] Chinmay Hegde, Michael Wakin, and Richard Baraniuk. Random projections for manifold learning. In J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, editors, *Advances in Neural Information Processing Systems 20*, pages 641–648. Curran Associates, Inc., 2008.

[20] Philipp Hennig and Søren Hauberg. Probabilistic solutions to differential equations and their application to Riemannian statistics. In *Proceedings of the 17th international Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 33, 2014.

[21] R. A. Khan. Approximation for the expectation of a function of the sample mean. *Statistics*, 38:117–122, 2004.

[22] Diederik P Kingma and Max Welling. Auto-Encoding Variational Bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014.

[23] Subhaneil Lahiri, Peiran Gao, and Surya Ganguli. Random projections of random manifolds. *arXiv:1607.04331*, 07 2016.

[24] Neil Lawrence. Probabilistic non-linear principal component analysis with Gaussian process latent variable models. *J. Mach. Learn. Res.*, 6:1783–1816, December 2005.

[25] R. J. Muirhead. *Aspects of Multivariate Statistical Theory*. Wiley, 2005.

[26] Xavier Pennec. Intrinsic Statistics on Riemannian Manifolds: Basic Tools for Geometric Measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154, July 2006.

[27] C. E. Rasmussen and C. K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.

[28] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. Stochastic backpropagation and approximate inference in deep generative models. In Eric P. Xing and Tony Jebara, editors, *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1278–1286. PMLR, 2014.

[29] Sam T Roweis and Lawrence K Saul. Nonlinear dimensionality reduction by locally linear embedding. *Science*, 290(5500):2323–2326, 2000.

[30] Bernhard Schölkopf, Alexander Smola, and Klaus-Robert Mller. Kernel principal component analysis. In *Advances in Kernel Methods - Support Vector Learning*, pages 327–352, 1999.

[31] Joshua B Tenenbaum, Vin De Silva, and John C Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000.

[32] Alessandra Tosi, Søren Hauberg, Alfredo Vellido, and Neil D. Lawrence. Metrics for probabilistic geometries. In *The Conference on Uncertainty in Artificial Intelligence (UAI)*, Quebec, Canada, July 2014.

[33] N. Verma. A note on random projections for preserving paths on a manifold. *UC San Diego, Tech. Report CS2011-0971*, 2011.