



Variational Autoencoders with Riemannian Brownian Motion Priors

Kalatzis, Dimitris; Eklund, David; Arvanitidis, Georgios; Hauberg, Søren

Published in:

Proceedings of the 37th International Conference on Machine Learning

Publication date:

2020

Document Version

Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

Citation (APA):

Kalatzis, D., Eklund, D., Arvanitidis, G., & Hauberg, S. (2020). Variational Autoencoders with Riemannian Brownian Motion Priors. In H. Daume, & A. Singh (Eds.), *Proceedings of the 37th International Conference on Machine Learning* (Vol. 119, pp. 5020-5033). International Machine Learning Society (IMLS).

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

Variational Autoencoders with Riemannian Brownian Motion Priors

Dimitris Kalatzis¹ David Eklund² Georgios Arvanitidis³ Søren Hauberg¹

Abstract

Variational Autoencoders (VAEs) represent the given data in a low-dimensional latent space, which is generally assumed to be Euclidean. This assumption naturally leads to the common choice of a standard Gaussian prior over continuous latent variables. Recent work has, however, shown that this prior has a detrimental effect on model capacity, leading to subpar performance. We propose that the Euclidean assumption lies at the heart of this failure mode. To counter this, we assume a Riemannian structure over the latent space, which constitutes a more principled geometric view of the latent codes, and replace the standard Gaussian prior with a Riemannian Brownian motion prior. We propose an efficient inference scheme that does not rely on the unknown normalizing factor of this prior. Finally, we demonstrate that this prior significantly increases model capacity using only one additional scalar parameter.

1. Introduction

Variational autoencoders (VAEs) (Kingma & Welling, 2014; Rezende et al., 2014) simultaneously learn a conditional density $p(\mathbf{x}|\mathbf{z})$ of high dimensional observations and low dimensional representations \mathbf{z} giving rise to these observations. In VAEs, a prior distribution $p(\mathbf{z})$ is assigned to the latent variables which is typically a standard Gaussian. It has, unfortunately, turned out that this choice of distribution is limiting the modelling capacity of VAEs and richer priors have been proposed instead (Tomczak & Welling, 2017; van den Oord et al., 2017; Bauer & Mnih, 2018; Klushyn et al., 2019). In contrast to this popular view, we will argue that the limitations of the prior are not due to *lack of*




Figure 1. The latent space priors of two VAEs trained on the digit 1 from MNIST. *Left*: Using a unit Gaussian prior. *Right*: Using a Riemannian Brownian motion (ours) with trainable (scalar) variance.

capacity, but rather *lack of principle*.

Informally, the Gaussian prior has two key problems.

1. The Euclidean representation is arbitrary. Behind the Gaussian prior lies the assumption that the latent space \mathcal{Z} is Euclidean. However, if the decoder $p_\theta(\mathbf{x}|\mathbf{z})$ is of sufficiently high capacity, then it is always possible to reparameterize the latent space from \mathbf{z} to $h(\mathbf{z})$, $h : \mathcal{Z} \rightarrow \mathcal{Z}$, and then let the decoder invert this reparameterization as part of its decoding process (Arvanitidis et al., 2018; Hauberg, 2018b). This implies that we cannot assign any meaning to specific instantiations of the latent variables, and that Euclidean distances carry limited meaning in \mathcal{Z} . This is an *identifiability problem* and it is well-known that even the most elementary latent variable models are subject to such. For example, Gaussian mixtures can be reparameterized by permuting cluster indices, and principal components can be arbitrarily rotated (Bishop, 2006).

2. Latent manifolds are mismatched onto \mathcal{Z} . In all but the simplest cases, the latent manifold \mathcal{M} giving rise to data observations is embedded in \mathcal{Z} . An encoder with adequate capacity will always recover some smoothed form of \mathcal{M} , which will either result in the latent space containing “holes” of low density or, in \mathcal{M} being mapped to the whole of \mathcal{Z} under the influence of the prior. Both cases will lead to bad samples or convergence problems. This problem is called *manifold mismatch* (Davidson et al., 2018; Falorsi et al., 2018) and is closely related to *distribution mismatch* (Hoffman & Johnson, 2016; Bauer & Mnih, 2018; Rosca et al., 2018) where the prior samples from regions to which the variational posterior (or encoder) does not assign any density. A graphical illustration of this situation can be

¹Section for Cognitive Systems, Department of Applied Mathematics and Computer Science, Technical University of Denmark
²Research Institutes of Sweden, Isafjordsgatan 22, 164 40 Kista, Sweden
³Empirical Inference Department, Max Planck Institute for Intelligent Systems, Tübingen, Germany. Correspondence to: Dimitris Kalatzis <dika@dtu.dk>.

seen on the left panel of Fig. 1, where a VAE is trained on the 1-digits of MNIST under the Gaussian prior. The prior assigns density where there is none.

In this paper, we consider an alternative prior, which is shown in the right panel of Fig. 1. This is a Riemannian Brownian motion model defined over the manifold immersed by the decoder. The Riemannian structure solves the identifiability problem and gives a meaningful representation that is invariant to reparametrizations and at the same time restricts the prior to sample only from the image of \mathcal{M} onto \mathcal{Z} . The prior generalizes the Gaussian to the Riemannian setting. It only has a single scalar variance parameter, yet it is able to capture intrinsic complexities in the data.

2. Background

2.1. Variational autoencoders

VAEs learn a generative model $p_\theta(\mathbf{x}, \mathbf{z})$ by specifying a likelihood of observations conditioned on latent variables $p_\theta(\mathbf{x}|\mathbf{z})$ and a prior over the latent variables $p(\mathbf{z})$. The marginal likelihood of the observations $p_\theta(\mathbf{x}) = \int p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z})d\mathbf{z}$ is intractable. As such, VAEs are trained by maximizing the variational *Evidence Lower Bound* (ELBO) on the marginal likelihood :

$$\mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})), \quad (2.1)$$

where $q_\phi(\mathbf{z}|\mathbf{x})$ denotes the variational family. Kingma & Welling (2014); Rezende et al. (2014) proposed a low variance estimator of stochastic gradients of the ELBO, known as *reparameterization trick*.

In the VAE framework, both the variational family $q_\phi(\mathbf{z}|\mathbf{x})$ and the conditional likelihood $p_\theta(\mathbf{x}|\mathbf{z})$ are parameterized by neural networks with variational parameters ϕ and generative parameters θ . In the language of autoencoders, these networks are often called *encoder* and *decoder* parameterizing the variational family and the generative model respectively. From an autoencoder perspective, Eq. 2.1 can be seen as a loss function involving a data reconstruction term (the generative model) and a regularization term (the KL divergence between the variational family and the prior distribution over the latent variables).

2.2. A primer on Riemannian geometry

The standard Gaussian prior relies on the usual Lebesgue measure which in turn, assumes a Euclidean structure over the latent space \mathcal{Z} . Recently, it has been noted (Arvanitidis et al., 2018; Hauberg, 2018b) that this assumption is mathematically questionable, and that, empirically, Euclidean latent space distances carry little information about the relationship between data points. Rather, a Riemannian interpretation of the latent space appears more promising. Hence we give a short review of Riemannian geometry.

A *smooth manifold* \mathcal{M} is a topological manifold endowed with a smooth structure. That is to say \mathcal{M} is locally homeomorphic to Euclidean space and we are able to do calculus on it. For a point $p \in \mathcal{M}$, the *tangent space* $T_p\mathcal{M}$ is a vector space centered on p which contains all tangent vectors to \mathcal{M} passing through point p . With this we can give a formal definition of the Riemannian metric tensor which is of central importance to any analysis involving Riemannian geometry.

Definition 1. (Riemannian metric) (do Carmo, 1992) Given a smooth manifold \mathcal{M} , a Riemannian metric on \mathcal{M} assigns on each point $p \in \mathcal{M}$ an inner product (i.e. a symmetric, positive definite, bilinear form) $\langle \cdot, \cdot \rangle_p$ in the tangent space $T_p\mathcal{M}$ which varies smoothly in the following sense: if $\mathbf{x} : \mathbb{R}^n \supset U \rightarrow \mathcal{M}$ is a local coordinate chart centered at p and $\frac{\partial}{\partial x_i}(q) = \mathbf{dx}_q(0, \dots, 1, \dots, 0)$ for $q \in U$, then $\langle \frac{\partial}{\partial x_i}(q), \frac{\partial}{\partial x_j}(q) \rangle_{\mathbf{x}(q)} = g_{ij}(q)$ is a smooth function on U .

By generalizing the inner product to Riemannian manifolds, the metric tensor gives meaning to length, angle and volume on manifolds. Central to distributions defined on a Riemannian manifold, the volume measure over an infinitesimal region centered at point p is defined as $d\mathcal{M}_p = \sqrt{\det \mathbf{G}_p} dp$, where \mathbf{G}_p is the matrix representation of the metric tensor evaluated at point p . Shortest paths on manifolds are represented by geodesic curves, which generalize straight lines in Euclidean space. A geodesic is a constant speed curve and its length can be computed by integrating the norm of its velocity vector under the metric, in other words $\mathcal{L} = \int_0^1 \|\frac{d\gamma}{dt}\|_g dt$. For $p \in \mathcal{M}$ there is a useful map defined on a neighborhood of the origin of $T_p\mathcal{M}$ called the exponential map. More precisely, the exponential map is a diffeomorphism, i.e. a bijection with a smooth inverse, between an open subset $\mathcal{U} \subset T_p\mathcal{M}$ and an open subset $\mathcal{U}' \subset \mathcal{M}$. Given $p \in \mathcal{M}$ and $v \in \mathcal{U}$, there is a unique geodesic $\gamma : [0, 1] \rightarrow \mathcal{M}$ with $\gamma(0) = p$ and $\frac{d\gamma}{dt}(0) = v$. The exponential map is given by $\text{exp}_p(v) = \gamma(1)$. Note that $\text{exp}_p(0) = p$. The inverse map (from \mathcal{U}' to \mathcal{U}) $\text{exp}_p^{-1} = \text{log}_p$ is called the *logarithmic map*.

Let $\mathcal{M} \subseteq \mathbb{R}^M$ be an embedded n -dimensional manifold and consider local coordinates $\phi : U \rightarrow \mathcal{M}$ with $U \subseteq \mathbb{R}^N$ an open subset. The Euclidean metric on \mathbb{R}^M induces a Riemannian metric on \mathcal{M} . Expressed in terms of the coordinates given by ϕ , this metric is known as the *pull-back metric* on U under ϕ . For $\mathbf{u} \in U$, the pull-back metric $\mathbf{G}_\mathbf{u}$ at \mathbf{u} is given by

$$\mathbf{G}_\mathbf{u} = \mathbf{J}_\phi^\top(\mathbf{u})\mathbf{J}_\phi(\mathbf{u}), \quad (2.2)$$

where \mathbf{J}_ϕ denotes the Jacobian matrix of ϕ .

2.3. VAE decoders as immersions

We will dedicate this subsection to showing that, under certain architectural choices, VAE decoders induce Riemannian

metrics in the latent space. That is to say, they belong to a certain class of maps, called *smooth immersions*, which give rise to *immersed submanifolds*. In other words, we will formally describe our intuition about VAEs mapping the latent space back to data space, using the language of smooth manifolds and Riemannian geometry.

The generative and variational distributions can be seen as families of parameterized mappings $g_\phi : \mathcal{X} \rightarrow \mathcal{Z}$ and $f_\theta : \mathcal{Z} \rightarrow \mathbb{R}^M$, $\mathcal{Z} \subset \mathbb{R}^N$ and $M > N$ and parameters ϕ and θ respectively. The family defined by the generative model is of particular interest. To make the subsequent exposition clearer we will assume a Gaussian generative model and rewrite it in the following form:

$$f_\theta(\mathbf{z}) = \mu_\theta(\mathbf{z}) + \sigma_\theta(\mathbf{z}) \odot \epsilon, \quad \epsilon \sim \mathcal{N}(\mathbf{0}, \mathbb{I}_M) \quad (2.3)$$

with $\mu_\theta : \mathcal{Z} \rightarrow \mathbb{R}^M$, $\sigma_\theta : \mathcal{Z} \rightarrow \mathbb{R}_+^M$, denoting the mean and standard deviation of the generative model parameterized by neural networks with parameters θ and \odot denoting the Hadamard or element-wise product.

Definition 2. (Smooth immersions) Given smooth manifolds \mathcal{M} and \mathcal{M}' with $\dim(\mathcal{M}) < \dim(\mathcal{M}')$, a mapping $f : \mathcal{M} \rightarrow \mathcal{M}'$, a point $p \in \mathcal{M}$ and its image $f(p) \in \mathcal{M}'$, the mapping f is called an immersion if its differential $df_p : T_p\mathcal{M} \rightarrow T_{f(p)}\mathcal{M}'$ is injective for all $p \in \mathcal{M}$.

We will consider a particular Riemannian metric on \mathcal{Z} induced by μ_θ and σ_θ . The architectures of μ_θ and σ_θ are such that these maps are immersions. Consider now the diagonal immersion

$$f : \mathcal{Z} \rightarrow \mathbb{R}^M \times \mathbb{R}_+^M : \mathbf{z} \mapsto (\mu_\theta(\mathbf{z}), \sigma_\theta(\mathbf{z})), \quad (2.4)$$

whose geometry encodes both mean and variance. The random map f_θ is a random projection given by ϵ of the diagonal immersion. Sampling using the decoder can therefore be seen as first sampling the image of this immersion and then randomly projecting down to \mathcal{X} (Eklund & Hauberg, 2019). Taking the pull-back metric \mathbf{G}_z of f to \mathcal{Z} we obtain

$$\mathbf{G}_z = \mathbf{J}_\mu(\mathbf{z})^\top \mathbf{J}_\mu(\mathbf{z}) + \mathbf{J}_\sigma(\mathbf{z})^\top \mathbf{J}_\sigma(\mathbf{z}), \quad (2.5)$$

where \mathbf{J}_μ and \mathbf{J}_σ are the Jacobian matrices of μ_θ and σ_θ .

The metric \mathbf{G}_z was studied by Arvanitidis et al. (2018) and is known to yield geodesics that follow high density regions in latent space. As an example, Fig. 2 shows geodesics of a VAE trained on 1-digits from MNIST, which follow the data due to the variance term of the metric, which penalizes geodesics going through low density regions of the latent space.

3. Geometric latent priors

It is evident that the geometric structure over the latent space carries significant information about data density that the




Figure 2. Example geodesics under the pull-back metric (2.5). The associated VAE is the same as in Fig. 1.

traditional Euclidean interpretation foregoes. With this in mind, we propose that the prior should be defined with respect to the geometric structure. We could opt for a *Riemannian normal distribution*, which is well-studied (Oller, 1993; Mardia & Jupp, 2000; Pennec, 2006; Arvanitidis et al., 2016; Hauberg, 2018a). Unfortunately, computing its normalization constant is expensive and involves Monte Carlo integration. Furthermore, it is equally hard to sample from this distribution, since it generally requires rejection sampling with non-trivial proposal distributions.

Instead we consider a cheap and flexible alternative, namely the heat kernel of a Brownian motion process (Hsu, 2002). A Brownian motion X_t on an immersed Riemannian manifold $\mathcal{M} \subseteq \mathbb{R}^M$ can be defined through a stochastic differential equation on Stratonovich form:

$$dX_t = \sum_{\alpha=1}^M P_\alpha(X_t) \circ dW_t^\alpha. \quad (3.1)$$

Here $W_t = (W_t^1, \dots, W_t^M)$ is a Brownian motion in \mathbb{R}^M and $P_1(X_t), \dots, P_M(X_t)$ denotes the projection of the standard basis of \mathbb{R}^M onto the tangent space of \mathcal{M} at X_t . This way, a Brownian motion on \mathcal{M} is driven by a Euclidean Brownian motion W_t projected to the tangent space. Fixing an initial point $\mu \in \mathcal{M}$ and a time $t > 0$, Brownian motion starting at μ running for time t gives rise to a random variable on \mathcal{M} . Its density function is the *transition density* $p(\mathbf{x})$. An alternative description of Brownian motion on \mathcal{M} is that $p(\mathbf{x})$ is the heat kernel associated to the Laplace-Beltrami operator of a scalar function h on \mathcal{M} :

$$\Delta h = d\mathcal{M}^{-1} \partial_i (d\mathcal{M} g^{ij} \partial_j h) \quad (3.2)$$

where $d\mathcal{M}$ is the volume measure of the immersed submanifold \mathcal{M} , g^{ij} are the components of the inverse metric tensor and $\partial_i := \frac{\partial}{\partial x^i}$, $\partial_j := \frac{\partial}{\partial x^j}$ are the basis vectors at the tangent space $T_p\mathcal{M}$. We will express the transition density in terms of local coordinates $\mathcal{Z} \rightarrow \mathcal{M}$ on \mathcal{M} . Conveniently, we may approximate the transition density by a so-called Parametrix expansion in a power series (Hsu, 2002). In this paper we will use the zeroth order approximation which




Figure 3. Inferred latent space for a toy data set, embedded via a non-linear function in \mathbb{R}^{100} . The background color, with blue representing lower and red representing higher values, from left to right, show: the (log) standard deviation estimated by a typical neural network; the associated (log) volume measure; the RBF (log) standard deviation estimate; and the associated (log) volume measure. Best viewed in color.

gives rise to the following expression for $p(\mathbf{z})$ with $\mathbf{z} \in \mathcal{Z}$:

$$p(\mathbf{z}) \approx (2\pi t)^{-d/2} H_0 \exp\left(-\frac{l^2(\mathbf{z}, \boldsymbol{\mu})}{2t}\right), \quad (3.3)$$

where:

- $t \in \mathbb{R}$, denotes the duration of the Brownian motion, and corresponds to *variance* on Euclidean manifolds.
- d is the dimensionality of \mathbf{z} .
- $\boldsymbol{\mu} \in \mathcal{Z}$ is the center of the Brownian motion.
- $l(\cdot, \cdot)$ is the geodesic distance on the manifold.
- $H_0 = \left(\frac{\det \mathbf{G}_{\mathbf{z}}}{\det \mathbf{G}_{\boldsymbol{\mu}}}\right)^{1/2}$ is the ratio of the Riemannian volume measure evaluated at points \mathbf{z} and $\boldsymbol{\mu}$ respectively.

Equation 3.3 can be evaluated reasonably fast as no Monte Carlo integration is required. The most expensive computation is the evaluation of the geodesic distance for which several efficient algorithms exist (Hennig & Hauberg, 2014; Arvanitidis et al., 2019). Here we parameterize the geodesic as a cubic spline and perform direct energy minimization.

3.1. Inference

Since we use the heat kernel density function for the prior $p(\mathbf{z})$, we need the variational family $q_\phi(\mathbf{z}|\mathbf{x})$ to be defined with respect to the same Riemannian measure. We therefore also use the heat kernel density function for the variational family, which is parameterized by the encoder network with variational parameters ϕ . The parameter t of the prior is learned through optimization. The ELBO can be derived with respect to the volume measure $d\mathcal{M}$:

$$\begin{aligned} \log p(\mathbf{x}) &\geq \mathcal{L}_{\mathcal{M}}(\mathbf{x}; \theta, \phi) \\ &\triangleq \int_{\mathcal{M}} \log\left(\frac{p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x})}\right) q_\phi(\mathbf{z}|\mathbf{x}) d\mathcal{M}_{\mathbf{z}} \\ &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log p_\theta(\mathbf{x}|\mathbf{z})] - \text{KL}(q_\phi(\mathbf{z}|\mathbf{x})||p(\mathbf{z})). \end{aligned} \quad (3.4)$$

This ELBO can be estimated using Monte Carlo samples from the variational posterior. With no analytical solution

to the KL divergence we resort to Monte Carlo integration:

$$\begin{aligned} \text{KL}(q||p) &= \int_{\mathcal{M}} \log \frac{q_\phi(\mathbf{z}|\mathbf{x})}{p(\mathbf{z})} q_\phi(\mathbf{z}|\mathbf{x}) d\mathcal{M}_{\mathbf{z}} \\ &= \mathbb{E}_{q(\mathbf{z}|\mathbf{x})}[\log q(\mathbf{z}|\mathbf{x}) - \log p(\mathbf{z})] \\ &\approx \frac{1}{N} \sum_{i=1}^N (\log q(\mathbf{z}_i|\mathbf{x}) - \log p(\mathbf{z}_i)) \end{aligned} \quad (3.5)$$

with:

$$\log q_\phi(\mathbf{z}|\mathbf{x}) = -\frac{d}{2} \log(2\pi t_q) + \log H_{0,q} - \frac{l_q^2}{2t_q} \quad (3.6)$$

$$\log p(\mathbf{z}) = -\frac{d}{2} \log(2\pi t_p) + \log H_{0,p} - \frac{l_p^2}{2t_p} \quad (3.7)$$

where $l_q^2 = l^2(\mathbf{z}, \boldsymbol{\mu}_q)$, $l_p^2 = l^2(\mathbf{z}, \boldsymbol{\mu}_p)$.

Thus, the final form of the Monte Carlo evaluation of the KL divergence is:

$$\begin{aligned} \text{KL}(q||p) &\approx \frac{1}{2} \left[\frac{1}{N} \sum_{i=1}^N \left(\log \det \mathbf{G}_{\boldsymbol{\mu}_p}(\mathbf{z}_i) - \right. \right. \\ &\log \det \mathbf{G}_{\boldsymbol{\mu}_q}(\mathbf{z}_i) + \frac{l^2(\mathbf{z}_i, \boldsymbol{\mu}_p)}{t_p} - \frac{l^2(\mathbf{z}_i, \boldsymbol{\mu}_q)}{t_q} \left. \right) \\ &\left. + d(\log t_p - \log t_q) \right] \end{aligned} \quad (3.8)$$

3.2. Sampling

In the previous section we mentioned that a Brownian motion (BM) on the manifold can be derived by projecting each BM step X_t on the tangent space at t . However we will take each step directly in the latent space and avoid having to evaluate the exponential map. Given a manifold \mathcal{M} with dimension N , the immersion $f : \mathcal{M} \rightarrow \mathbb{R}^M$, a point $\mathbf{a} \in \mathcal{M}$ and its image under f , $\mathbf{A} \in \mathbb{R}^M$ we take a random step from \mathbf{A} :

$$\Delta \sim \mathcal{N}(\mathbf{0}, \Sigma_M). \quad (3.9)$$

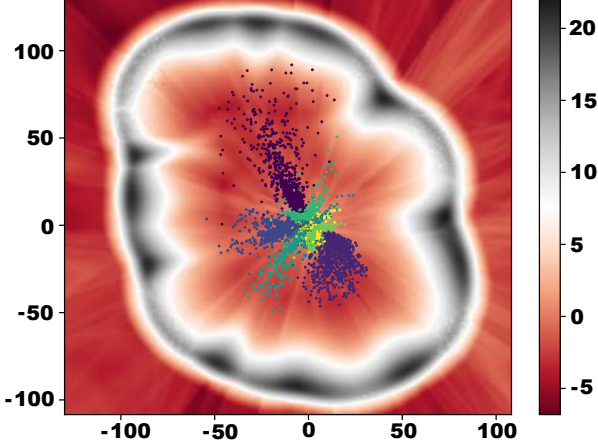


Figure 4. Latent space of an \mathcal{R} -VAE, plotted against the Riemannian volume measure dM . Once again note the “borders” created by the metric roughly demarcating the latent code support. The latent codes are colored according to label. Best viewed in color.

Applying a Taylor expansion we have:

$$f(\mathbf{a} + \epsilon) = f(\mathbf{a}) + \mathbf{J}_a \epsilon + \mathcal{O}(\epsilon^2). \quad (3.10)$$

With $\Delta = f(\mathbf{a} + \epsilon) - f(\mathbf{a})$ we have:

$$\Delta = \mathbf{J}_a \epsilon + \mathcal{O}(\epsilon^2). \quad (3.11)$$

For small ϵ an approximation to taking a step directly in the latent space is then $\mathbf{b} = \mathbf{a} + \epsilon$ with $\epsilon \approx \mathbf{J}_a^+ \Delta$ and $\mathbf{J}_a^+ = (\mathbf{J}_a^\top \mathbf{J}_a)^{-1} \mathbf{J}_a^\top \in \mathbb{R}^{N \times M}$ the pseudoinverse of \mathbf{J}_a . Since $\Delta \sim \mathcal{N}(\mathbf{0}, \Sigma_M)$ the step ϵ can be written:

$$\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{J}_a^+ \Sigma_M (\mathbf{J}_a^+)^{\top}). \quad (3.12)$$

We consider an isotropic heat kernel so in our case $\Sigma_M = \sigma^2 \mathbf{I}$. Furthermore:

$$\begin{aligned} \mathbf{J}_a^+ \Sigma_M (\mathbf{J}_a^+)^{\top} &= (\mathbf{J}_a^\top \mathbf{J}_a)^{-1} \mathbf{J}_a^\top \Sigma_M \mathbf{J}_a (\mathbf{J}_a^\top \mathbf{J}_a)^{-\top} \\ &= \sigma^2 (\mathbf{J}_a^\top \mathbf{J}_a)^{-1} \mathbf{J}_a^\top \mathbf{J}_a (\mathbf{J}_a^\top \mathbf{J}_a)^{-\top} \\ &= \sigma^2 (\mathbf{J}_a^\top \mathbf{J}_a)^{-\top} = \sigma^2 (\mathbf{J}_a^\top \mathbf{J}_a)^{-1}. \end{aligned} \quad (3.13)$$

This implies that

$$\epsilon \sim \mathcal{N}(\mathbf{0}, \sigma^2 (\mathbf{J}_a^\top \mathbf{J}_a)^{-1}). \quad (3.14)$$

Thus, to sample from the prior we simply need to run Brownian motion for $t = 1, \dots, T$:

$$\mathbf{z}_t \sim \mathcal{N}\left(\mathbf{z}_{t-1}, \frac{\sigma^2}{T} (\mathbf{J}_{\mathbf{z}_{t-1}}^\top \mathbf{J}_{\mathbf{z}_{t-1}})^{-1}\right) \quad (3.15)$$

An obvious concern regarding the computational cost of sampling is the inverting of the metric tensor. While this is

Table 1. Results on MNIST (mean & std deviation over 10 runs). *Rec* denotes the negative conditional likelihood.

Model	Neg. ELBO	Rec	KL
VAE			
d = 2	-1030.38 \pm 5.34	-1033.06 \pm 5.48	2.68 \pm .14
d = 5	-1076.64 \pm 4.48	-1078.91 \pm 4.44	2.27 \pm .04
d = 10	-1110.79 \pm 1.17	-1113.01 \pm 1.13	2.22 \pm .03
VAE-VampPrior			
d = 2	-1045.03 \pm 5.22	-1047.34 \pm 5.20	2.30 \pm .03
d = 5	-1109.74 \pm 4.87	-1111.63 \pm 4.87	1.88 \pm .01
d = 10	-1116.58 \pm 4.23	-1118.27 \pm 4.20	1.69 \pm .02
\mathcal{R} -VAE			
d = 2	-1047.29 \pm 2.77	-1053.70 \pm 2.75	14.33 \pm .01
d = 5	-1141.06 \pm 7.09	-1177.86 \pm 3.39	28.00 \pm .2
d = 10	-1170.03 \pm 18.52	-1280.94 \pm 14.67	57.76 \pm 3.85

a valid concern for large latent dimensionalities, in practice and for the typical number of latent dimensions found in generative modelling literature the sampling cost is bearable, considering that the operation can be parallelized for K samples. We further note that from a practical standpoint for small diffusion times the number of discretized steps can be small. The time complexity of the sampling operation is

$$\mathcal{O}(KHM + KMN^2 + N^3) \quad (3.16)$$

where K is the number of samples, N is the latent space dimensionality, M is the input space dimensionality and H is the decoder hidden layer size.

4. Meaningful variance estimation

We now turn to the problem of restricting our prior to sample from the image of our manifold in \mathcal{Z} . Since typically the geometry of the data is not known a priori, we adopt the Bayesian approach and relate uncertainty estimation in the generative model to the geometry of the latent manifold. Specifically, since the generative model parameterizes $f_\theta : \mathcal{Z} \rightarrow \mathcal{X}$ we construct it such that the pull-back metric will acquire high values away from the data support and thereby restrict prior samples to high density regions of the latent manifold.

In Sec. 2.3 we described the metric tensor arising from the diagonal immersion f . By the form of the metric, it is clear that both $\mu_\theta(\mathbf{z})$ and $\sigma_\theta(\mathbf{z})$ contribute to the manifold geometry. In recent works (Arvanitidis et al., 2018; Hauberg, 2018b; Detlefsen et al., 2019) it was shown that neural network variance estimates are typically poor in regions away from the training data, due to poor extrapolation properties. Thus, neural networks cannot be trusted to properly estimate the variance of the generative model “off-the-shelf” when the functional form of the immersion (and thus the geometry of the data) is not known a priori. By extension, this leads to poor estimates of latent manifold geometry and latent densities. Arvanitidis et al. (2018) propose to use a

Table 2. Results on FashionMNIST (mean & std deviation over 10 runs). *Rec* denotes the negative conditional likelihood.

Model	Neg. ELBO	Rec	KL
VAE			
d = 2	-443.13 \pm 10.67	-447.44 \pm 10.8	4.31 \pm .14
d = 5	-511.65 \pm 3.70	-517.41 \pm 3.84	5.76 \pm .21
d = 10	-525.05 \pm 5.87	-530.86 \pm 5.9	5.81 \pm .05
VAE-VampPrior			
d = 2	-705.90 \pm 17.3	-708.45 \pm 17.29	2.54 \pm .01
d = 5	-769.27 \pm 5.	-770.1 \pm 5.02	0.83 \pm .09
d = 10	-774.17 \pm 10.83	-777.75 \pm 10.78	3.57 \pm .06
\mathcal{R} -VAE			
d = 2	-708.77 \pm 6.93	-722.41 \pm 5.736	13.64 \pm 1.51
d = 5	-889.62 \pm 3.44	-913.61 \pm 3.38	23.83 \pm .8
d = 10	-959.2 \pm 5.37	-1001.4 \pm 4.08	40.35 \pm .8

radial basis function (RBF) network (Que & Belkin, 2016) to estimate precision, rather than variance. We adopt this approach due to its simplicity and relative numerical stability, however we note that similar approaches for principled variance estimation exist (Detlefsen et al., 2019; ?).

The influence of the RBF network can be seen in Fig. 3, where it is compared with a usual neural network variance estimate. Note that the metric creates “borders” demarcating the regions to which the latent codes have been mapped by the encoder. This makes interpolations and random walks generally follow the trend of the latent points instead of wondering off the support. Thus, this regularization scheme restricts prior sampling to such high density regions. A similar effect is not observed in the usual Gaussian VAE, where the prior samples from regions to which the variational posterior has not necessarily placed probability density (Hoffman & Johnson, 2016; Rosca et al., 2018).

5. Experiments

5.1. Generative modelling

For our first experiment we train a VAE with a Riemannian Brownian motion prior (\mathcal{R} -VAE) for different latent dimensions and compare it to a VAE with a standard Normal prior and a VAE with a VampPrior. Tables 1 & 2 show the results. \mathcal{R} -VAE achieves a better lower bound than both its Euclidean counterparts. The Brownian motion prior adapts to the latent code support and as such yields more expressive representations. On the other hand, with only a single parameter it results in a model that generalizes better than VAEs with a VampPrior.

5.2. Classification

We next assess the usefulness of the latent representations of \mathcal{R} -VAE. Fig. 4 shows the latent code clusters. \mathcal{R} -VAE has produced more separable clusters in the latent space due to the prior adapting to the latent codes, which results in a less

regularized clustering. We quantitatively measured the utility of the \mathcal{R} -VAE latent codes in different dimensionalities by training a classifier to predict digit labels and measuring the average overall and per-digit F1 score. Table 3 shows the results when comparing against the same classifier trained on latent codes derived by a VAE. \mathcal{R} -VAE has a significant advantage in low dimensions. As dimensionality increases this advantage becomes non-existent. An explanation for this is that due to the KL annealing of the Euclidean VAE, its representations have become more informative.

5.3. Qualitative results

Finally we explore the geometric properties of a \mathcal{R} -VAE with a 2-dimensional latent space. Fig 4 shows the learned manifold. As in Fig. 3, the influence of the variance network on the metric can be seen in the “borders” surrounding the latent code support.

We begin by investigating the behavior of distances on the induced manifold. Fig. 5 shows the geodesic curves between two pairs of random points on the manifold, compared against their Euclidean counterpart. The geodesic interpolation is influenced by the metric tensor, which makes sure that shortest paths will generally avoid areas of low density. This can easily be seen in top left Fig. 5, where the geodesic curve follows a path along a high density region. Contrast this to the Euclidean straight line between the two points traversing a lower density region. Reconstructed images along the curves can be seen in the middle and bottom rows. Even in less apparent cases (top right Fig. 5), reconstructions of latent codes along geodesic curves generally provide smoother transitions between the curve endpoints as can be seen by comparing the middle right and bottom right sections of the figure.

Next, we investigate sampling from \mathcal{R} -VAE. In Sec. 4 we claimed that a Brownian motion prior coupled with the RBF regularization of the decoder variance network would yield samples that mostly avoid low density regions of the latent space. To empirically prove this, we executed two sets of multiple sampling runs on the latent manifold. In the first set we ran Brownian motion with the learned prior parameters. These runs and the resulting images are displayed in Fig. 6. The random walks generally stay within high density regions of the manifold. Cases where they explore low density regions do exist but they are rare. The samples generally seem clear although sometimes their quality drops, especially when the sampler is transitioning between classes, where variance estimates are higher. This could potentially be rectified with a less aggressive deterministic warm-up scheme, which would result in more concentrated densities with thinner tails, although between-class variance estimates would likely still be higher compared to within-class ones. For the second set of the sampling runs, we increased the

Table 3. Per digit and average F1 score for a classifier trained on the learned latent codes of VAE and \mathcal{R} -VAE. Results are averaged over 5 classifier training runs.

Digits	0	1	2	3	4	5	6	7	8	9	Avg
VAE											
d = 2	0.94	0.95	0.88	0.67	0.55	0.42	0.86	0.68	0.61	0.53	$0.72_{\pm.002}$
d = 5	0.95	0.97	0.94	0.90	0.90	0.89	0.95	0.93	0.88	0.87	$0.92_{\pm.001}$
d = 10	0.98	0.99	0.97	0.94	0.96	0.95	0.98	0.97	0.93	0.94	$0.96_{\pm.001}$
\mathcal{R} -VAE											
d = 2	0.95	0.97	0.89	0.68	0.64	0.56	0.88	0.85	0.71	0.64	$0.78_{\pm.002}$
d = 5	0.95	0.98	0.94	0.91	0.94	0.88	0.95	0.93	0.90	0.89	$0.93_{\pm.0008}$
d = 10	0.98	0.98	0.96	0.95	0.96	0.95	0.97	0.97	0.93	0.94	$0.96_{\pm.001}$




Figure 5. Top: Interpolations plotted in the latent space of \mathcal{R} -VAE. Black indicates a geodesic interpolant, red indicates a Euclidean interpolant. Middle: Images reconstructed along the geodesic interpolation. Bottom: Images reconstructed along the Euclidean interpolation. The latent codes are color-coded according to label. Best viewed in color.

duration of the Brownian motion. These runs are displayed along with the sampled images in Fig. 7. The influence of the variance estimates on the metric tensor is clearly shown here. As the sampler is moving farther away from the latent code support, evaluations of the metric tensor increase making these regions harder to traverse. As a result the random walk either oscillates with decreased speed and stops close to the boundary (as in Figures 7a and 7b) or returns to higher density regions of the manifold. This clearly shows that \mathcal{R} -VAE mostly avoids the manifold mismatch problem.

6. Related work

Learned priors. In recent literature many works have identified the adverse effects of the KL divergence regular-

ization when the prior is chosen to be a standard Gaussian. As such, there have been many approaches of learning a more flexible prior. Chen et al. (2016) propose learning an autoregressive prior by applying an Inverse Autoregressive transformation (Kingma et al., 2016) to a simple prior. Nalisnick & Smyth (2016) propose a non-parametric stick-breaking prior. (Tomczak & Welling, 2017) propose learning the prior as a mixture of variational posteriors. More recently, Bauer & Mnih (2018) present a rejection sampling approach with a learned acceptance function, while Klushyn et al. (2019) proposed a hierarchical prior through an alternative formulation of the objective.

Non-Euclidean latent space. Arvanitidis et al. (2018) was one of the first to analyze the latent space of a VAE




Figure 6. Top: Brownian motion runs on the learned latent manifold. Bottom: Corresponding sampled images. The sampler mostly stays in high density regions of the latent manifold. Best viewed in color.




Figure 7. Brownian motion runs with artificially increased t (diffusion) parameter beyond the learned value. Note that the borders created by the metric tensor stop the sampler from exploring low density regions any further - the sampler either stops (a and b) or returns to regions of higher density (c). This effect is observed in the sampled images. Best viewed in color.

from a non-Euclidean perspective. This work was inspired by Tosi et al. (2014) that studied the Riemannian geometry of the Gaussian process latent variable model (Lawrence, 2005). Arvanitidis et al. (2018) train a Euclidean VAE and fit a latent Riemannian LAND distribution (Arvanitidis et al., 2016) and show that this view of the latent space leads to more accurate statistical estimates, as well as better sample quality.

Since then, a number of other works have appeared in literature that propose learning non-Euclidean latent manifolds. Xu & Durrett (2018) and Davidson et al. (2018) learn a VAE with a von Mises-Fisher latent distribution, which samples codes on the unit hypersphere. Similarly, Mathieu et al. (2019) and Nagano et al. (2019) extend VAEs to hyperbolic spaces. Mathieu et al. (2019) assume a Poincaré ball model as a latent space and present 2 generalizations of the Euclidean Gaussian distribution - a wrapped Normal and the Riemannian Normal distributions, of which only the latter is a maximum entropy generalization. In practice, they perform similarly. Nagano et al. (2019) assume a Lorentz hyperbolic model as a latent space and also present a wrapped Normal generalization of the Gaussian. While these works have correctly identified the problem of the standard Gaussian not being a truly uninformative prior, due to the Euclidean assumption, they have proposed approaches which are designed for observations with known geometries. Most of the time, however, this information is not available and a more general framework for learning geometrically informed VAEs is needed. In response to this, Skopek et al. (2019) propose VAEs with the latent space modelled as a product of constant curvature manifolds, where each component curvature is learned. While more general than a model with a fixed curvature latent manifold, this framework still requires the specification of number of component manifolds along with the sign of their respective curvature. Finally, similar to our approach, Li et al. (2019) and Rey et al. (2019) both propose the heat kernel as a variational family representing a Brownian motion process on a Riemannian manifold. They test their approaches on a priori chosen manifolds.

7. Conclusion

In this paper we presented VAEs with Riemannian manifolds as latent spaces and proposed a Riemannian generalization of the Gaussian along with an efficient sampling scheme. We show that the pull-back metric informs distances in the latent space, remaining invariant to reparameterizations. We further make explicit the relationship between uncertainty estimation and proper latent geometry and qualitatively show that geometrically informed priors avoid manifold mismatch by drawing samples from the image of the manifold in the latent space. Quantitatively, we

show that our approach outperforms Euclidean VAEs both in an unsupervised learning task and a classification task, especially in low latent space dimensions.

Acknowledgements

SH and DE were supported by a research grant (15334) from VILLUM FONDEN. This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement n° 757360). We gratefully acknowledge the support of the NVIDIA Corporation with the donation of GPU hardware.

References

- Arvanitidis, G., Hansen, L. K., and Hauberg, S. A locally adaptive normal distribution. In *Advances in Neural Information Processing Systems*, pp. 4251–4259, 2016.
- Arvanitidis, G., Hansen, L. K., and Hauberg, S. Latent space oddity: on the curvature of deep generative models. In *International Conference on Learning Representations (ICLR)*, 2018.
- Arvanitidis, G., Hauberg, S., Hennig, P., and Schober, M. Fast and robust shortest paths on manifolds learned from data. In *Proceedings of the 19th international Conference on Artificial Intelligence and Statistics (AISTATS)*, 2019.
- Bauer, M. and Mnih, A. Resampled priors for variational autoencoders. *arXiv preprint arXiv:1810.11428*, 2018.
- Bishop, C. M. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2006.
- Chen, X., Kingma, D. P., Salimans, T., Duan, Y., Dhariwal, P., Schulman, J., Sutskever, I., and Abbeel, P. Variational lossy autoencoder. *arXiv preprint arXiv:1611.02731*, 2016.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891*, 2018.
- Detlefsen, N. S., Jørgensen, M., and Hauberg, S. Reliable training and estimation of variance networks. *arXiv preprint arXiv:1906.03260*, 2019.
- do Carmo, M. *Riemannian Geometry*. Mathematics (Boston, Mass.). Birkhäuser, 1992.
- Eklund, D. and Hauberg, S. Expected path length on random manifolds. *arXiv preprint arXiv:1908.07377*, 2019.
- Falorsi, L., de Haan, P., Davidson, T. R., De Cao, N., Weiler, M., Forré, P., and Cohen, T. S. Explorations in

- homeomorphic variational auto-encoding. *arXiv preprint arXiv:1807.04689*, 2018.
- Hauberg, S. Directional statistics with the spherical normal distribution. In *2018 21st International Conference on Information Fusion (FUSION)*, pp. 704–711. IEEE, 2018a.
- Hauberg, S. Only bayes should learn a manifold (on the estimation of differential geometric structure from data). *arXiv preprint arXiv:1806.04994*, 2018b.
- He, K., Zhang, X., Ren, S., and Sun, J. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *2015 IEEE International Conference on Computer Vision, ICCV 2015, Santiago, Chile, December 7-13, 2015*, pp. 1026–1034. IEEE Computer Society, 2015. doi: 10.1109/ICCV.2015.123. URL <https://doi.org/10.1109/ICCV.2015.123>.
- Hennig, P. and Hauberg, S. 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.
- Hoffman, M. D. and Johnson, M. J. Elbo surgery: yet another way to carve up the variational evidence lower bound. In *Workshop in Advances in Approximate Bayesian Inference, NIPS*, volume 1, pp. 2, 2016.
- Hsu, E. P. *Stochastic analysis on manifolds*, volume 38. American Mathematical Soc., 2002.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In Bengio, Y. and LeCun, Y. (eds.), *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Kingma, D. P. and Welling, M. Auto-Encoding Variational Bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR)*, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *Advances in neural information processing systems*, pp. 4743–4751, 2016.
- Klushyn, A., Chen, N., Kurle, R., Cseke, B., and van der Smagt, P. Learning hierarchical priors in vaes. In *Advances in Neural Information Processing Systems*, pp. 2866–2875, 2019.
- Lawrence, N. D. Probabilistic non-linear principal component analysis with gaussian process latent variable models. *Journal of machine learning research*, 6(Nov):1783–1816, 2005.
- Li, H., Lindenbaum, O., Cheng, X., and Cloninger, A. Variational diffusion autoencoders with random walk sampling. *arXiv preprint arXiv:1905.12724*, 2019.
- Mardia, K. V. and Jupp, P. E. Basic concepts and models. *Mardia KV, Jupp PE. Directional statistics, 2nd edition. Chichester (UK): John Wiley & Sons*, pp. 25–56, 2000.
- Mathieu, E., Le Lan, C., Maddison, C. J., Tomioka, R., and Teh, Y. W. Continuous hierarchical representations with poincaré variational auto-encoders. In *Advances in neural information processing systems*, pp. 12544–12555, 2019.
- Nagano, Y., Yamaguchi, S., Fujita, Y., and Koyama, M. A differentiable gaussian-like distribution on hyperbolic space for gradient-based learning. *arXiv preprint arXiv:1902.02992*, 2019.
- Nalisnick, E. and Smyth, P. Stick-breaking variational autoencoders. *arXiv preprint arXiv:1605.06197*, 2016.
- Oller, J. M. On an intrinsic analysis of statistical estimation. In *Multivariate Analysis: Future Directions 2*, pp. 421–437. Elsevier, 1993.
- Penneç, X. Intrinsic statistics on riemannian manifolds: Basic tools for geometric measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127, 2006.
- Que, Q. and Belkin, M. Back to the future: Radial basis function networks revisited. In *Artificial Intelligence and Statistics (AISTATS)*, 2016.
- Rey, L. A. P., Menkovski, V., and Portegies, J. W. Diffusion variational autoencoders. *arXiv preprint arXiv:1901.08991*, 2019.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and variational inference in deep latent gaussian models. In *International Conference on Machine Learning*, volume 2, 2014.
- Rosca, M., Lakshminarayanan, B., and Mohamed, S. Distribution matching in variational inference. *arXiv preprint arXiv:1802.06847*, 2018.
- Skopek, O., Ganea, O.-E., and Bécigneul, G. Mixed-curvature variational autoencoders. *arXiv preprint arXiv:1911.08411*, 2019.
- Tomczak, J. M. and Welling, M. Vae with a vampprior. *arXiv preprint arXiv:1705.07120*, 2017.

Tosi, A., Hauberg, S., Vellido, A., and Lawrence, N. D. Metrics for Probabilistic Geometries. In *The Conference on Uncertainty in Artificial Intelligence (UAI)*, July 2014.

van den Oord, A., Vinyals, O., et al. Neural discrete representation learning. In *Advances in Neural Information Processing Systems*, pp. 6306–6315, 2017.

Xu, J. and Durrett, G. Spherical latent spaces for stable variational autoencoders. *arXiv preprint arXiv:1808.10805*, 2018.

A. On neural network-based immersions

For the decoder map 2.3 to be a valid immersion, its differential df needs to be injective for all $p \in \mathcal{M}$ as stated in definition 2. The differential of f is represented by its Jacobian matrix \mathbf{J}_f and for it to be injective for all $p \in \mathcal{M}$, it needs to be full rank. This is ensured if for the MLPs representing the decoder μ_θ and σ_ψ the following are true:

- Each hidden layer in the network has an equal or greater number of units to the previous layer ($n_{L-1} \leq n_L$).
- All weight matrices in the network are full rank.
- The activation functions are at least twice differentiable and strictly monotone.

In our experiments, we opt for the same number of units in each hidden layer of the network and ELU non-linearities. In theory, the ELU activation function could present problems since it has a point of discontinuity at 0, however we did not experience any numerical instability that would arise in such case. All weight matrices are initialized uniformly (He et al., 2015) which practically has zero probability of yielding low rank weight matrices. While theoretically this could change via the gradient updates of the weights, this would once again immediately break experiments because of numerical instabilities, which we did not observe.

B. Geodesic estimation

We estimate geodesic distances by minimizing curve energy. In detail, we represent the geodesic curve with a cubic spline with parameters initialized to form a straight line. These parameters are then optimized via gradient descent by minimizing the curve energy:

$$\begin{aligned} \mathcal{E}(\gamma) &= \frac{1}{2} \int_0^1 \|\dot{\gamma}(t)\|_g^2 dt \\ &= \frac{1}{2} \int_0^1 \dot{\gamma}^T(t) \mathbf{G}_\gamma \dot{\gamma}(t) dt \end{aligned} \quad (\text{B.1})$$

where γ is the geodesic curve, $\dot{\gamma}$ is the first derivative of the curve, i.e. its velocity vector and \mathbf{G}_γ is the matrix representation of the metric tensor evaluated at the curve points. The integral B.1 is computed by numerical approximation, where the partition of the interval can be chosen as a hyperparameter.

C. Experimental setup

The architectures of all model variants are shown below in Tables 4 and 5. The encoder mean and variance, as well as the decoder mean are modelled by 2-layer MLPs as shown below. The decoder mean mirrors the encoder mean, while the *precision* β is estimated by the RBF network. The number of the RBF centers is set to 350 and the bandwidth is set to 0.01 in all cases. For a fair comparison, all models share the same underlying architecture for the encoder and decoder. Tables 4 and 5 summarize the architectures, listing the activation function for each layer with the units corresponding to each layer in parentheses.

Table 4. Encoder network architectures.

Network	Layer 1	Layer 2	Output
$\mu_\phi(\mathbf{x})$	ELU (300)	ELU (300)	Linear ($\dim(\mathcal{Z})$)
$\sigma_\phi^2(\mathbf{x})$	ELU (300)	ELU (300)	Softplus ($\dim(\mathcal{Z})$)

Table 5. Decoder network architectures. * denotes strictly positive weights.

Network	Layer 1	Layer 2	Output
$\mu_\theta(\mathbf{z})$	ELU (300)	ELU (300)	Linear ($\dim(\mathcal{X})$)
$\beta_\psi(\mathbf{z})$	RBF ($\mathbb{R}^{\dim(\mathcal{Z}) \times 350}$)	Linear* ($\dim(\mathcal{X})$)	Identity ($\dim(\mathcal{X})$)

C.1. Section 5.1 experiment

Detlefsen et al. (2019) highlighted the importance of optimizing the mean and variance components separately, when training VAEs with Gaussian generative models. Following this paradigm, in all our experiments we first optimize the encoder components (μ_ϕ and σ_ϕ) along with the decoder μ_θ . Then, keeping these fixed, we optimize the decoder σ_ψ . All models were trained for 300 epochs. More specifically, the \mathcal{R} -VAE was trained as an autoencoder (optimizing only the encoder μ_ϕ and σ_ϕ and the decoder μ_θ) for the first 100 epochs and for the remaining 200 epochs the latent prior and the decoder β_ψ were optimized. Similarly for a VAE, it was deterministically warmed up for 100 epochs and for the remaining 200 epochs, the decoder β_ψ was optimized. All experiments were run with the Adam optimizer (Kingma & Ba, 2015) with default parameter settings and a fixed learning rate of 10^{-3} . The batch size was 100 for all models.

C.2. Section 5.2 experiment

The classifier used on this section was a single, 100-unit layer MLP with ReLU non-linearities, trained for 100 epochs with the Adam optimizer with default parameter settings and a learning rate of 10^{-3} . The batch size was set at 64. The architectures of the models giving rise to the latent representations are as in the previous section.

C.3. Runtime comparisons

Below is the wall clock time for every model used in the experiments. The statistics were computed without a fixed seed. The latent space dimensions are denoted by d . In VAE-VampPrior, n denotes the number of mixture components in the latent prior.

Table 6. Per epoch training time for each model. Mean and std deviation in seconds, computed over 100 epochs on MNIST.

Model	$d = 2$	$d = 5$	$d = 10$
VAE	10.64 \pm .51	11.01 \pm .77	11.10 \pm .60
VAE-VampPrior ($n = 128$)	10.66 \pm .6	11.22 \pm .9	11.37 \pm .96
VAE-VampPrior ($n = 256$)	10.72 \pm .3	11.34 \pm 1.21	11.52 \pm .77
VAE-VampPrior ($n = 512$)	10.9 \pm .34	11.38 \pm .93	12.18 \pm 1.12
\mathcal{R} -VAE	55.73 \pm 4.36	59.97 \pm 1.33	60.13 \pm 1.19

C.4. Complete results for VAE-VampPrior

Tables 1 & 2 show the results of the best performing VampPrior model variant. Here we show the complete results of the VAE-VampPrior in all settings. Below n denotes the number of mixture components in the latent prior, while d denotes the latent space dimensions.

Table 7. MNIST results of VAE with VampPrior for varying latent space dimensions and number of mixture components in the latent prior.

Model	Neg. ELBO	Rec	KL
$d = 2$			
VAE-VampPrior ($n = 128$)	-1039.66 \pm 2.56	-1042.13 \pm 2.56	2.46 \pm .01
VAE-VampPrior ($n = 256$)	-1045.04 \pm 5.20	-1047.34 \pm 5.22	2.30 \pm .03
VAE-VampPrior ($n = 512$)	-1040.79 \pm 9.23	-1043.24 \pm 9.25	2.45 \pm .05
$d = 5$			
VAE-VampPrior ($n = 128$)	-1100.77 \pm 4.98	-1102.46 \pm 4.91	1.69 \pm .06
VAE-VampPrior ($n = 256$)	-1103.29 \pm 1.85	-1105.04 \pm 1.79	1.75 \pm .12
VAE-VampPrior ($n = 512$)	-1109.74 \pm 4.87	-1111.63 \pm 4.87	1.88 \pm .01
$d = 10$			
VAE-VampPrior ($n = 128$)	-1110.05 \pm 6.10	-1112.23 \pm 5.82	1.84 \pm .04
VAE-VampPrior ($n = 256$)	-1116.58 \pm 4.23	-1118.28 \pm 4.20	1.69 \pm .02
VAE-VampPrior ($n = 512$)	-1100.64 \pm 2.93	-1102.42 \pm 2.97	1.78 \pm .03

Table 8. FashionMNIST results of VAE with VampPrior for varying latent space dimensions and number of mixture components in the latent prior.

Model	Neg. ELBO	Rec	KL
$d = 2$			
VAE-VampPrior ($n = 128$)	$-694.63_{\pm 8.65}$	$-697.14_{\pm 8.65}$	$2.50_{\pm .01}$
VAE-VampPrior ($n = 256$)	$-702.67_{\pm 17.45}$	$-705.19_{\pm 17.44}$	$2.52_{\pm .04}$
VAE-VampPrior ($n = 512$)	$-705.90_{\pm 21.29}$	$-708.45_{\pm 21.29}$	$2.54_{\pm .01}$
$d = 5$			
VAE-VampPrior ($n = 128$)	$-755.80_{\pm .66}$	$-756.58_{\pm .71}$	$0.77_{\pm .06}$
VAE-VampPrior ($n = 256$)	$-767.54_{\pm 3.22}$	$-768.33_{\pm 3.31}$	$0.78_{\pm .09}$
VAE-VampPrior ($n = 512$)	$-769.27_{\pm 5.0}$	$-770.10_{\pm 5.02}$	$0.83_{\pm .09}$
$d = 10$			
VAE-VampPrior ($n = 128$)	$-754.47_{\pm 6.78}$	$-758.20_{\pm 6.72}$	$3.72_{\pm .06}$
VAE-VampPrior ($n = 256$)	$-756.13_{\pm 5.40}$	$-760.49_{\pm 5.1}$	$3.69_{\pm .06}$
VAE-VampPrior ($n = 512$)	$-774.17_{\pm 10.83}$	$-777.75_{\pm 10.78}$	$3.58_{\pm .06}$