
Riemannian Continuous Normalizing Flows

Emile Mathieu^{†*}, Maximilian Nickel[‡]

emile.mathieu@stats.ox.ac.uk, maxn@fb.com

[†] Department of Statistics, University of Oxford, UK

[‡] Facebook Artificial Intelligence Research, New York, USA

Abstract

Normalizing flows have shown great promise for modelling flexible probability distributions in a computationally tractable way. However, whilst data is often naturally described on Riemannian manifolds such as spheres, tori, and hyperbolic spaces, most normalizing flows implicitly assume a flat geometry, making them either misspecified or ill-suited in these situations. To overcome this problem, we introduce *Riemannian continuous normalizing flows*, a model which admits the parametrization of flexible probability measures on smooth manifolds by defining flows as the solution to ordinary differential equations. We show that this approach can lead to substantial improvements on both synthetic and real-world data when compared to standard flows or previously introduced projected flows.

1 Introduction

Learning well-specified probabilistic models is at the heart of many problems in machine learning and statistics. Much focus has therefore been placed on developing methods for modelling and inferring expressive probability distributions. Normalizing flows (Rezende and Mohamed, 2016) have shown great promise for this task as they provide a general and extensible framework for modelling highly complex and multimodal distributions (Papamakarios et al., 2019).

An orthogonal but equally important aspect of well-specified models is to correctly characterize the geometry which describes the proximity of data points. Riemannian manifolds provide a general framework for this purpose and are a natural approach to model tasks in many scientific fields ranging from earth and climate science to biology and computer vision. For instance, storm trajectories may be modelled as paths on the sphere (Karpatne et al., 2017), the shape of proteins can be parametrized using tori (Hamelryck et al., 2006), cell developmental processes can be described through paths in hyperbolic space (Klimovskaia et al., 2019), and human actions can be recognized in video using matrix manifolds (Lui, 2012). If appropriately chosen, manifold-informed methods can lead to improved sample complexity and generalization, improved fit in the low parameter regime, and guide inference methods to interpretable models. They can also be understood as a geometric prior that encodes a practitioner’s assumption about the data and imposes an inductive bias.

However, conventional normalizing flows are not readily applicable to such manifold-valued data since their implicit Euclidean assumption makes them unaware of the underlying geometry or borders of the manifold. As a result they would yield distributions having some or all of their mass lying




Figure 1: Trajectories generated on the sphere to model volcano eruptions. Note that these converge to the known *Ring of Fire*.

*Work done while at Facebook AI research.

outside the manifold, rendering them ill-suited or even misspecified so that central concepts like the reverse Kullback-Leibler (KL) divergence would not even be defined.

In this work, we propose a principled way to combine both of these aspects and parametrize flexible probability distributions on Riemannian manifolds. Specifically, we introduce *Riemmanian continuous normalizing flows* in which flows are defined via vector fields on manifolds and computed as the solution to the associated ordinary differential equation (ODE) (see Figure 1 for an illustration). Intuitively, our method operates by first parametrizing a vector field on the manifold with a neural network, then sampling particles from a base distribution, and finally approximating their flow along the vector field using a numerical solver. Both the neural network and the solver are aware of the underlying geometry which ensures that the flow is always located on the manifold – yielding a *Riemannian* method.

This approach allows us to combine multiple important advantages: One major challenge of normalizing flows lies in designing transformations that enable efficient sampling and density computation. By basing our approach on *continuous normalizing flows* (CNFs) (Chen et al., 2019; Grathwohl et al., 2018; Salman et al., 2018) we avoid strong structural constraints to be imposed on the flow, as is the case for most *discrete* normalizing flows. Such unconstrained *free-form* flows have empirically been shown to be highly expressive (Chen et al., 2020; Grathwohl et al., 2018). Moreover, *projected* methods require a differentiable mapping from a Euclidean space to the manifold, yet such a function cannot be bijective, which in turn leads to numerical challenges. By taking a Riemannian approach, our method is more versatile since it does not rely on an ad-hoc projection map and simultaneously reduces numerical artefacts that interfere with training. To the best of our knowledge, our method is the first to combine these properties as existing methods for normalizing flows on manifolds are either discrete (Bose et al., 2020; Rezende et al., 2020), projected (Gemici et al., 2016; Falorsi et al., 2019; Bose et al., 2020) or manifold-specific (Sei, 2011; Bose et al., 2020; Rezende et al., 2020).

We empirically demonstrate the advantages of our method on constant curvature manifolds – i.e., the Poincaré disk and the sphere – and show the benefits of the proposed approach compared to non-Riemannian and projected methods for maximum likelihood estimation and reverse KL minimization. We also apply our method to density estimation on earth-sciences data (e.g., locations of earthquakes, floods and wildfires) and show that it yields better generalization performance and faster convergence.

2 Continuous Normalizing Flows on Riemannian Manifolds

Normalizing flows operate by pushing a simple base distribution through a series of parametrized invertible maps, referred as the *flow*. This can yield a highly complex and multimodal distribution which is typically assumed to live in a Euclidean vector space. Here, we propose a principled approach to extend normalizing flows to manifold-valued data, i.e. *Riemmanian continuous normalizing flows* (RCNFs). Following CNFs (Chen et al., 2019; Grathwohl et al., 2018; Salman et al., 2018) we define manifold flows as the solutions to ODEs. The high-level idea is to parametrize flows through the time-evolution of manifold-valued particles z – in particular via their velocity $\dot{z}(t) = f_\theta(z(t), t)$ where f_θ denotes a *vector field*. Particles are first sampled from a simple base distribution, and then their evolution is integrated by a manifold-aware numerical solver, yielding a new complex multimodal distribution of the particles. This *Riemannian and continuous* approach has the advantages of allowing almost *free-form* neural networks and of not requiring any mapping from a Euclidean space which would potentially lead to numerical challenges.

For practical purposes, we focus our theoretical and experimental discussion on *constant curvature manifolds* (see Table 1). In addition to being widely used in the literature (Nickel and Kiela, 2017; Davidson et al., 2018; Mardia and Jupp, 2000; Hasnat et al., 2017), these manifolds are convenient to work with since most related geometrical quantities are available in closed-form. However, our

Table 1: Summary of d-dimensional continuous constant (sectional) curvature manifolds.

Geometry	Model	Curvature	Coordinates	$\sqrt{\det g} = d \text{ Vol} / d \text{ Leb}_{\mathbb{R}^d}$	Compact	
Euclidean	\mathbb{R}^d	Real vector space	$K = 0$	Cartesian z	1	No
Hyperbolic	\mathbb{B}_K^d	Poincaré ball	$K < 0$	Cartesian z	$\left(2 / (1 + K \ z\ ^2)\right)^d$	No
Elliptic	\mathbb{S}_K^d	Hypersphere	$K > 0$	n-spherical φ	$K^{-\frac{d-1}{2}} \prod_{i=1}^{d-2} \sin(\varphi_i)^{d-i-1}$	Yes

proposed approach is generic and could be used on a broad class of manifolds such as product and matrix manifolds like tori and Grassmannians. For a brief overview of relevant concepts in Riemannian geometry please see Appendix A.1 or Lee (2003) for a more thorough introduction.

In the following, we develop the key components which allow us to define continuous normalizing flows that are aware of the underlying Riemannian geometry: flow, likelihood, and vector field.

Vector flows Flows in conventional normalizing flows are defined as smooth mappings $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^d$ which transform a base distribution $z \sim P_0$ into a complex distribution P_θ . For normalizing flows to be well-behaved and convenient to work with, the flow is required to be *bijective* and *differentiable* which introduces significant structural constraints on ϕ . Continuous normalizing flows overcome this issue by defining the flow $\phi : \mathbb{R}^d \times \mathbb{R} \rightarrow \mathbb{R}^d$ generated by an ordinary differential equation, allowing for unrestricted neural network architectures. Here we show how *vector fields* can be used to define similar flows $\phi : \mathcal{M} \times \mathbb{R} \rightarrow \mathcal{M}$ on general *Riemannian manifolds*.

Consider the temporal evolution of a particle $z(t)$ lying on a d -dimensional manifold \mathcal{M} , whose velocity is given by a *vector field* $f_\theta(z(t), t)$. Intuitively, $f_\theta(z(t), t)$ indicates the direction and speed along which the particle is moving on the *manifold's surface*. Classic examples for such vector fields include weathercocks giving wind direction and compasses pointing toward the magnetic north pole of the earth. Formally, let $\mathcal{T}_z \mathcal{M}$ denote the *tangent space* at z and $\mathcal{T} \mathcal{M} = \bigcap_{z \in \mathcal{M}} \mathcal{T}_z \mathcal{M}$ the associated *tangent bundle*. Furthermore, let $f_\theta : \mathcal{M} \times \mathbb{R} \mapsto \mathcal{T} \mathcal{M}$ denote a vector field on \mathcal{M} . The particle's time-evolution according to f_θ is then given by the following ODE

$$\frac{dz(t)}{dt} = f_\theta(z(t), t). \quad (1)$$

To transform a base distribution using this vector field, we are then interested in a particle's position after time t . When starting at an initial position $z(0) = z_0$, the *flow operator* $\phi : \mathcal{M} \times \mathbb{R} \mapsto \mathcal{M}$ gives the particle's position at any time t as $z(t) = \phi(z_0, t)$. Leveraging the fundamental theorem of flows (Lee, 2003), we can show that under mild conditions, this flow is *bijective* and *differentiable*. We write C^1 for the set of differentiable functions whose derivative are continuous.

Proposition 1 (Vector flows). *Let \mathcal{M} be a smooth complete manifold. Furthermore, let f_θ be a C^1 -bounded time-dependent vector field. Then there exists a global flow $\phi : \mathcal{M} \times \mathbb{R} \mapsto \mathcal{M}$ such that for each $t \in \mathbb{R}$, the map $\phi(\cdot, t) : \mathcal{M} \mapsto \mathcal{M}$ is a C^1 -diffeomorphism (i.e. C^1 bijection with C^1 inverse).*

Proof. See Appendix E.1 for a detailed derivation. \square

Note that scaling the vector field as $f_\theta^\alpha \triangleq \alpha f_\theta$ results in a time-scaled flow $\phi^\alpha(z, t) = \phi(z, \alpha t)$. The integration duration t is therefore arbitrary. Without loss of generality we set $t = 1$ and write $\phi \triangleq \phi(\cdot, 1)$. Concerning the evaluation of the flow ϕ , it generally does not accept a closed-form solution and thus requires to be approximated numerically. To this extent we rely on an explicit and adaptive Runge-Kutta (RK) integrator of order 4 (Dormand and Prince, 1980). However, standard integrators used in CNFs generally do not preserve manifold constraints (Hairer, 2006). To overcome this issue we rely on a *projective solver* (Hairer, 2011). This solver works by conveniently solving the ODE in the ambient Cartesian coordinates and projecting each step onto the manifold. Projections onto \mathbb{S}^d are computationally cheap since they amount to l^2 norm divisions. No projection is required for the Poincaré ball.

Likelihood Having a flow at hand, we are now interested in evaluating the likelihood of our *pushforward* model $P_\theta = \phi_\# P_0$. Here, the *pushforward* operator $\#$ indicates that one obtains samples $z \sim \phi_\# P_0$ as $z = \phi(z_0)$ with $z_0 \sim P_0$. For this purpose, we derive in the following the change in density in terms of the geometry of the manifold and show how to efficiently estimate the likelihood.

Change in density In normalizing flows, we can compute the likelihood of a sample via the change in density from the base distribution to the pushforward. Applying the chain rule we get

$$\log p_\theta(z) - \log p_0(z_0) = \log \left| \det \frac{\partial \phi^{-1}(z)}{\partial z} \right| = - \log \left| \det \frac{\partial \phi(z_0)}{\partial z} \right|. \quad (2)$$

In general, computing the Jacobian's determinant of the flow is challenging since it requires d reverse-mode automatic differentiations to obtain the full Jacobian matrix, and $O(d^3)$ operations to

compute its determinant. CNFs side step direct computation of the determinant by leveraging the time-continuity of the flow and re-expressing Equation 2 as the integral of the *instantaneous change* in log density $\int_0^t \frac{\partial \log p_\theta(z(t))}{\partial t} dt$. However, standard CNFs make an implicit Euclidean assumption to compute this quantity which is violated for general Riemannian manifolds. To overcome this issue we express the instantaneous change in log-density in terms of the *Riemannian metric*. In particular, let $G(z)$ denote the *matrix representation of the Riemannian metric* for a given manifold \mathcal{M} , then $G(z)$ endows tangent spaces $\mathcal{T}_z \mathcal{M}$ with an inner product. For instance in the Poincaré ball \mathbb{B}^d , it holds that $G(z) = (2 / (1 + \|z\|^2))^2 I_d$, while in Euclidean space \mathbb{R}^d we have $G(z) = I_d$, where I_d denotes the identity matrix. Using the Liouville equation, we can then show that the instantaneous change in variable is defined as follows.

Proposition 2 (Instantaneous change of variables). *Let $z(t)$ be a continuous manifold-valued random variable given in local coordinates, which is described by the ODE from Equation 1 with probability density $p_\theta(z(t))$. The change in log-probability then also follows a differential equation given by*

$$\frac{\partial \log p_\theta(z(t))}{\partial t} = -\text{div}(f_\theta(z(t), t)) = -|G(z(t))|^{-\frac{1}{2}} \text{tr} \left(\frac{\partial \sqrt{|G(z(t))|} f_\theta(z(t), t)}{\partial z} \right) \quad (3)$$

$$= -\text{tr} \left(\frac{\partial f_\theta(z(t), t)}{\partial z} \right) - |G(z(t))|^{-\frac{1}{2}} \left\langle f_\theta(z(t), t), \frac{\partial}{\partial z} \sqrt{|G(z(t))|} \right\rangle. \quad (4)$$

Proof. For a detailed derivation of Equation 3 see Appendix C. \square

Note that in the Euclidean setting $\sqrt{|G(z)|} = 1$ thus the second term of Equation 4 vanishes and we recover the formula from Grathwohl et al. (2018); Chen et al. (2019).

Estimating the divergence Even though the determinant of Equation 2 has been replaced in Equation 3 by a trace operator with lower computational complexity, we still need to compute the full Jacobian matrix of f_θ . Similarly to Grathwohl et al. (2018); Salman et al. (2018), we make use of Hutchinson’s trace estimator to compute the Jacobian efficiently. In particular, Hutchinson (1990) showed that $\text{tr}(A) = \mathbb{E}_{p(\epsilon)}[\epsilon^\top A \epsilon]$ with $p(\epsilon)$ being a d -dimensional random vector such that $\mathbb{E}[\epsilon] = 0$ and $\text{Cov}(\epsilon) = I_d$. Leveraging this trace estimator to approximate the divergence in Equation 3 yields

$$\text{div}(f_\theta(z(t), t)) = |G(z(t))|^{-\frac{1}{2}} \mathbb{E}_{p(\epsilon)} \left[\epsilon^\top \frac{\partial \sqrt{|G(z(t))|} f_\theta(z(t), t)}{\partial z} \epsilon \right]. \quad (5)$$

We note that the variance of this estimator can potentially be high since it scales with the inverse of the determinant term $\sqrt{|G(z(t))|}$ (see Appendix D.2). By integrating Equation 3 over time with the stochastic divergence estimator from Equation 5, we get the following total change in log-density between the manifold-valued random variables z and z_0

$$\log \left(\frac{p_\theta(z)}{p_\theta(z_0)} \right) = - \int_0^1 \text{div}(f_\theta(z(t), t)) dt = -\mathbb{E}_{p(\epsilon)} \left[\int_0^1 |G(z(t))|^{-\frac{1}{2}} \epsilon^\top \frac{\partial \sqrt{|G(z(t))|} f_\theta(z(t), t)}{\partial z} \epsilon dt \right]. \quad (6)$$

It can be seen that Equation 6 accounts again for the underlying geometry through the metric $G(z(t))$. Table 1 lists closed-form solutions of its determinant for constant curvature manifolds. Furthermore, the vector-Jacobian product can be computed through backward auto-differentiation with linear complexity, avoiding the quadratic cost of computing the full Jacobian matrix. Additionally, the integral is approximated via the discretization of the flow returned by the solver.

Choice of base distribution P_0 The closer the initial base distribution is to the target distribution, the easier the learning task should be. However, it is challenging in practice to incorporate such prior knowledge. We consequently use a uniform distribution on \mathbb{S}^d since it is the most "uncertain" distribution. For the Poincaré ball \mathbb{B}^d , we rely on a standard wrapped Gaussian distribution \mathcal{N}^W (Nagano et al., 2019; Mathieu et al., 2019) because it is convenient to work with.

Vector field Finally, we discuss the form of the *vector field* $f_\theta : \mathcal{M} \times \mathbb{R} \rightarrow \mathcal{T}\mathcal{M}$ which generates the flow ϕ used to pushforward samples. We parametrize f_θ via a feed-forward neural network which takes as input manifold-valued particles, and outputs their velocities. The architecture of the vector field has a direct impact on the expressiveness of the distribution and is thus crucially important. In order to take into account these geometrical properties we make use of specific input and output layers that we describe below. The rest of the architecture is based on a multilayer perceptron.

Input layer To inform the neural network about the geometry of the manifold \mathcal{M} , we use as first layer a *geodesic distance layer* (Ganea et al., 2018; Mathieu et al., 2019) which generalizes *linear layers* to manifolds, and can be seen as computing distances to *decision boundaries* on \mathcal{M} . These boundaries are parametrized by *geodesic hyperplanes* H_w , and the associated *neurons* $h_w(z) \propto d_{\mathcal{M}}(z, H_w)$, with $d_{\mathcal{M}}$ being the geodesic distance. Horizontally stacking several of these *neurons* makes a *geodesic distance layer*. We refer to Appendix E.2 for more details.

Output layer To constrain the neural net to \mathcal{TM} , we output vectors in \mathbb{R}^{d+1} when $\mathcal{M} = \mathbb{S}^d$, before projecting them to the tangent space i.e. $\mathbf{f}_{\theta}(z) = \text{proj}_{\mathcal{T}_z\mathcal{M}} \text{neural_net}(z)$. This is not necessary in \mathbb{B}^d since the ambient space is of equal dimension. Yet, velocities scale as $\|\mathbf{f}_{\theta}(z)\|_z = |G(z)|^{1/2} \|\text{neural_net}(z)\|_2$, hence we scale the *neural_net* by $|G(z)|^{-1/2}$ s.t. $\|\mathbf{f}_{\theta}(z)\|_z = \|\text{neural_net}(z)\|_2$.

Regularity For the flow to be bijective, the vector field \mathbf{f}_{θ} is required to be C^1 and bounded (cf Proposition 1). The boundness and smoothness conditions can be satisfied by relying on bounded smooth non-linearities in \mathbf{f}_{θ} such as tanh, along with bounded weight and bias at the last layer.

Training In density estimation and inference tasks, one aims to learn a model P_{θ} with parameters θ by minimising a *divergence* $\mathcal{L}(\theta) = D(P_{\mathcal{D}} \| P_{\theta})$ w.r.t. a target distribution $P_{\mathcal{D}}$. In our case, the parameters θ refer to the parameters of the vector field \mathbf{f}_{θ} . We minimize the loss $\mathcal{L}(\theta)$ using first-order stochastic optimization, which requires Monte Carlo estimates of loss gradients $\nabla_{\theta} \mathcal{L}(\theta)$. We back-propagate gradients through the explicit solver with $O(1/h)$ memory cost, h being the step size. When the loss $\mathcal{L}(\theta)$ is expressed as an expectation over the model P_{θ} , as in the *reverse* KL divergence, we rely on the *reparametrization trick* (Kingma and Welling, 2014; Rezende et al., 2014). In our experiments we will consider both the negative log-likelihood and reverse KL objectives

$$\mathcal{L}^{\text{Like}}(\theta) = -\mathbb{E}_{z \sim P_{\mathcal{D}}} [\log p_{\theta}(z)] \text{ and } \mathcal{L}^{\text{KL}}(\theta) = D_{\text{KL}}(P_{\theta} \| P_{\mathcal{D}}) = \mathbb{E}_{z \sim P_{\theta}} [\log p_{\theta}(z) - \log p_{\mathcal{D}}(z)]. \quad (7)$$

Additionally, regularization terms can be added in the hope of improving training and generalization. See Appendix D for a discussion and connections to the dynamical formulation of optimal transport.

3 Related work

Here we discuss previous work that introduced normalizing flows on manifolds. For clarity we split these into *projected* vs *Riemannian* methods which we describe below.

Projected methods These methods consist in parametrizing a normalizing flow on \mathbb{R}^d and then pushing-forward the resulting distribution along an invertible map $\psi : \mathbb{R}^d \rightarrow \mathcal{M}$. Yet, the existence of such an *invertible* map is equivalent to \mathcal{M} being homeomorphic to \mathbb{R}^d (e.g. being "flat"), hence limiting the scope of that approach. Moreover there is no principled way to choose such a map, and different choices lead to different numerical or computational challenges which we discuss below.

Exponential map The first generic *projected* map that comes to mind in this setting is the exponential map $\exp_{\mu} : T_{\mu}\mathcal{M} \cong \mathbb{R}^d \rightarrow \mathcal{M}$, which parameterizes geodesics starting from μ with velocity $v \in T_{\mu}\mathcal{M}$. This leads to so called *wrapped* distributions $P_{\theta}^W = \exp_{\mu_{\theta}} P$, with P a probability measure on \mathbb{R}^d . This approach has been taken by Falorsi et al. (2019) to parametrize probability distributions on Lie groups. Yet, in compact manifolds – such as spheres or the $\text{SO}(3)$ group – computing the density of *wrapped* distributions requires an infinite summation, which in practice needs to be truncated. This is not the case however on hyperbolic spaces (like the Poincaré ball) since the exponential map is bijective on these manifolds. This approach has been proposed in Bose et al. (2020) where they extend Real-NVP (Dinh et al., 2017) to the hyperboloid model of hyperbolic geometry. In addition to this *wrapped* Real-NVP, they also introduced a hybrid coupling model which is empirically shown to be more expressive. We note however that the exponential map is believed to be "badly behaved" away from the origin (Dooley and Wildberger, 1993; Al-Mohy and Higham, 2010).

Stereographic map Alternatively to the exponential map, Gemici et al. (2016) proposed to parametrize probability distributions on \mathbb{S}^d via the *stereographic projection* defined as $\rho(z) = z_{2:d} / (1 + z_1)$ with *projection point* $-\{\mu_0\} = (-1, 0, \dots, 0)$. Gemici et al. then push a probability measure P defined on \mathbb{R}^d along the inverse of the *stereographic* map ρ , yielding $P_{\theta}^S = \rho_{\theta}^{-1} P$.




Figure 2: Probability densities on \mathbb{B}^2 . Models have been trained by maximum likelihood to fit $\mathcal{N}^W(\exp_0(2\partial x), \Sigma)$. The black semi-circle indicates the disk’s border. The best run out of twelve trainings is shown for each model.




Figure 3: Ablation study of the vector field architecture for the *Riemannian* model. Models have been trained to fit a $\mathcal{N}^W(\exp_0(\partial x), \Sigma)$.




Figure 4: Negative Log-likelihood of CNFs trained to fit a $\mathcal{N}^W(\exp_0(\alpha \partial x), \Sigma)$ target on \mathbb{B}^2 .

However, the stereographic map ρ is not injective, and projects $-\mu_0$ to ∞ . This implies that spherical points close to the projection point $-\{\mu_0\}$ are mapped far away from the origin of the plane. Modelling probability distributions with mass close to $-\{\mu_0\}$ may consequently be numerically challenging since the norm of the Euclidean flow would explode. Similarly, [Rezende et al. \(2020\)](#) introduced flows on hyperspheres and tori by using the inverse tangent function. Although this method is empirically shown to perform well, it similarly suffers from numerical instabilities near singularity points.

Riemannian methods In contrast to *projected* methods which rely on mapping the manifold to a Euclidean space, *Riemannian* methods do not. As a consequence they side-step any artefact or numerical instability arising from the manifold’s projection. Early work ([Sei, 2011](#)) proposed transformations along geodesics on the hypersphere by evaluating the exponential map at the gradient of a scalar manifold function. Recently, [Rezende et al. \(2020\)](#) introduced ad-hoc *discrete Riemannian* flows for hyperspheres and tori based on Möbius transformations and spherical splines. We contribute to this line of work by introducing *continuous* flows on general Riemannian manifolds. In contrast to *discrete* flows (e.g. [Bose et al., 2020](#); [Rezende et al., 2020](#)), *time-continuous* flows as ours alleviate strong structural constraints on the flow by implicitly parametrizing it as the solution to an ODE ([Grathwohl et al., 2018](#)). Additionally, recent and concurrent work ([Lou et al., 2020](#); [Falorsi and Forré, 2020](#)) proposed to extend neural ODEs to smooth manifolds.

4 Experimental results

We evaluate the empirical performance of the above-mentioned models on hyperbolic and spherical geometry. We will first discuss experiments on two synthetic datasets where we highlight specific pathologies of the naive and projected methods via unimodal distributions at the point (or the limit) of the pathology. This removes additional modelling artefacts that would be introduced through more complex distributions and allows to demonstrate advantages of our approach on the respective manifolds. We further show that these advantages also translate to substantial gains on highly multi-modal real world datasets.

For all *projected* models (e.g. stereographic and wrapped cf Section 3), the vector field’s architecture is chosen to be a multilayer perceptron as in [Grathwohl et al. \(2018\)](#), whilst the architecture described in Section 2 is used for our *Riemannian* (continuous normalizing flow) model. For fair comparisons,

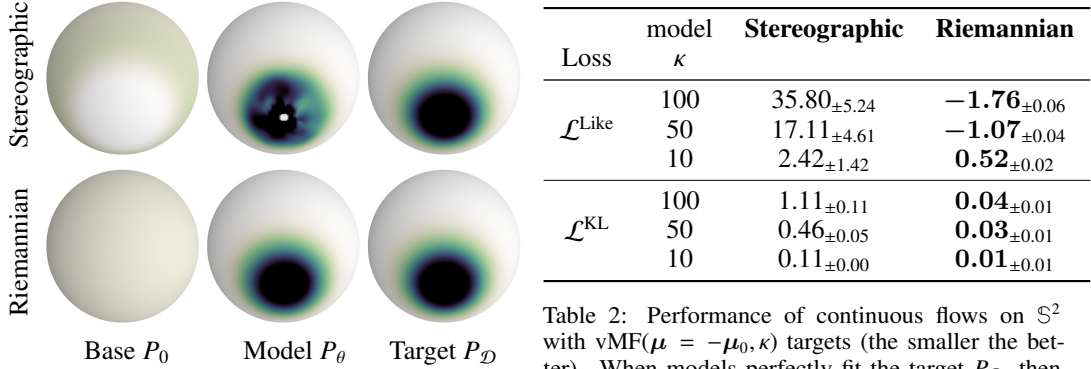


Figure 5: Probability distributions on \mathbb{S}^2 . Models trained to fit a $\text{vMF}(\mu = -\mu_0, \kappa = 10)$.

Table 2: Performance of continuous flows on \mathbb{S}^2 with $\text{vMF}(\mu = -\mu_0, \kappa)$ targets (the smaller the better). When models perfectly fit the target $P_{\mathcal{D}}$, then $\mathcal{L}^{\text{Like}} = \mathbb{H}[P_{\mathcal{D}}]$ which decreases with κ , explaining $\mathcal{L}^{\text{Like}}$'s results for the *Riemannian* model.

we also parametrize *projected* models with a CNF. Also, all models are chosen to have approximately the same number of parameters. All models were implemented in PyTorch (Paszke et al., 2017) and trained by stochastic optimization with Adam (Kingma and Ba, 2015). All 95% confidence intervals are computed over 12 runs. Please refer to Appendix G for full experimental details. We open-source our code for reproducibility ².

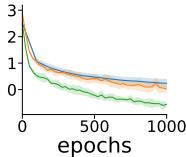
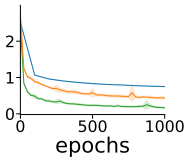
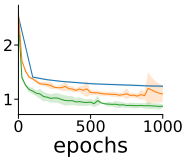
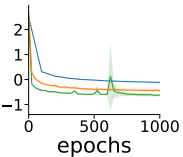
Hyperbolic geometry and limits of conventional and wrapped methods First, we aim to show that conventional normalizing flows are ill-suited for modelling target manifold distributions. These are blind to the geometry, so we expect them to behave poorly when the target is located where the manifold behaves most differently from a Euclidean space. We refer to such models as *naive* and discuss their properties in more detail in Appendix B.1. Second, we wish to inspect the behaviour of *wrapped* models (see Section 3) when the target is away from the exponential map origin.

To this extent we parametrize a wrapped Gaussian target distribution $\mathcal{N}^{\text{W}}(\exp_0(\alpha \partial x), \Sigma) = \exp_{\mu\#} \mathcal{N}(\alpha \partial x, \Sigma)$ defined on the Poincaré disk \mathbb{B}^2 (Nagano et al., 2019; Mathieu et al., 2019). The scalar parameter α allows us to locate the target closer or further away from the origin of the disk. We put three CNFs models on the benchmark; our *Riemannian* (from Section 2), a conventional *naive* and a *wrapped* model. The base distribution P_0 is a standard Gaussian for the *naive* and *wrapped* models, and a standard wrapped Gaussian for the *Riemannian* model. Models are trained by maximum likelihood until convergence. Throughout training, the Hutchinson’s estimator is used to approximate the divergence as in Equation 5. It can be seen from Figure 4 that the *Riemannian* model indeed outperforms the *naive* and *wrapped* models as we increase the values of α – i.e., the closer we move to the boundary of the disk. Figure 2 shows that qualitatively the *naive* and *wrapped* models seem to indeed fail to properly fit the target when it is located far from the origin. Additionally, we assess the architectural choice of the vector field used in our *Riemannian* model. In particular, we conduct an ablation study on the rescaling of the output layer, by training for 10 iterations a *rescaled* and an *unscaled* version of our model. Figure 3 shows that the number of function evaluations (NFE) tends to be large and sometimes even dramatically diverges when the vector field’s output is *unscaled*. In addition to increasing the computational cost, this in turns appears to worsen the convergence’s speed of the model. This further illustrates the benefits of our vector field parameterization.

Spherical geometry and limits of the stereographic projection model Next, we evaluate the ability of our model and the stereographic projection model from Section 3 to approximate distributions on the sphere which are located around the projection point $-\mu_0$. We empirically assess this phenomenon by choosing the target distribution to be a Von-Mises Fisher (Downs, 1972) distribution $\text{vMF}(\mu, \kappa)$ located at $\mu = -\mu_0$, and with concentration κ (which decreases with the variance). Along with the *stereographic* projection method, we also consider our *Riemannian* model from Section 2. We neither included the *naive* model since it is misspecified here (leading to an undefined reverse KL divergence), nor the *wrapped* model as computing its density requires an infinite summation (see Section 3). The base distribution P_0 is chosen to be a standard Gaussian on \mathbb{R}^2 for the *stereographic*

²https://github.com/facebookresearch/riemannian_cnf

Table 3: Negative test log-likelihood of continuous normalizing flows on \mathbb{S}^2 datasets.

	Volcano	Earthquake	Flood	Fire
Mixture vMF ■	$-0.38_{\pm 0.19}$	$0.60_{\pm 0.01}$	$1.10_{\pm 0.02}$	$-0.21_{\pm 0.01}$
Stereographic ■	$-0.07_{\pm 0.17}$	$0.43_{\pm 0.04}$	$1.03_{\pm 0.05}$	$-0.45_{\pm 0.06}$
Riemannian ■	$-0.67_{\pm 0.11}$	$0.16_{\pm 0.03}$	$0.86_{\pm 0.05}$	$-0.64_{\pm 0.03}$
Learning curves				
Data size	829	6124	4877	12810

model and a uniform distribution on \mathbb{S}^2 for the *Riemannian* model. Models are trained by computing the exact divergence. The performance of these two models are quantitatively assessed on both the negative log-likelihood and reverse KL criteria.

Figure 5 shows densities of the target distribution along with the base and learned distributions. We observe that the stereographic model fails to push mass close enough to the singularity point $-\mu_0$, as opposed to the Riemannian model which perfectly fits the target. Table 2 shows the negative log-likelihood and reverse KL losses of both models when varying the concentration parameter κ of the vMF target. The larger the concentration κ is, the closer to the singularity point $-\mu_0$ the target’s mass gets. We observe that the *Riemannian* model outperforms the *stereographic* one to fit the target for both objectives, although this performance gap shrinks as the concentration gets smaller. Also, we believe that the gap in performance is particularly large for the log-likelihood objective because it heavily penalizes models that fail to cover the support of the target. When the vMF target is located away from the singularity point, we noted that both models were performing similarly well.

Density estimation of spherical data Finally, we aim to measure the expressiveness and modelling capabilities of our method on real world datasets. To this extent, we gathered four earth location datasets, representing respectively volcano eruptions (NOAA, 2020b), earthquakes (NOAA, 2020a), floods (Brakenridge, 2017) and wild fires (EOSDIS, 2020). We approximate the earth’s surface (and thus also these data points) as a perfect sphere. Along our *Riemannian* CNF, we also assess the fitting capacity of a mixture of von Mises-Fisher (vMF) distributions and a *stereographic* projected CNF. The locations of the vMF components are learned via stochastic Riemannian optimization (Bonnabel, 2013; Bécigneul and Ganea, 2019). The learning rate and number of components are selected by hyperparameter grid search. In our experiments, we split datasets randomly into training and testing datasets, and fit the models by maximum likelihood estimation on the training dataset. CNF models are trained by computing the exact divergence.

We observe from Table 3 that for all datasets, the *Riemannian* model outperforms its *stereographic* counterpart and the mixture of vMF distributions by a large margin. It can also be seen from the learning curves that the *Riemannian* model converges faster. Figure 6 shows the learned spherical distributions along with the training and testing datasets. We note that qualitatively the *stereographic* distribution is generally more diffuse than its *Riemannian* counterpart. It also appears to allocate some of its mass outside the target support, and to cover less of the data points. Additional figures are shown in Appendix H.

Limitations In the spherical setting, the stochastic estimator to approximate the divergence from Equation 5 exhibits high variance. Its variance scales with the inverse of $\sqrt{|G(z)|} = \sin(\theta)$, which becomes too large around the north pole and thus requires the use of the exact estimator. On large-scale datasets, where the stochastic estimator has important runtime advantages, this issue could be alleviated by choosing a different vector field basis than the one induced by some local coordinates (e.g. Falorsi and Forré, 2020). For the Poincaré ball, no such variance behavior of the stochastic estimator exists and it can readily be applied to large-scale data.

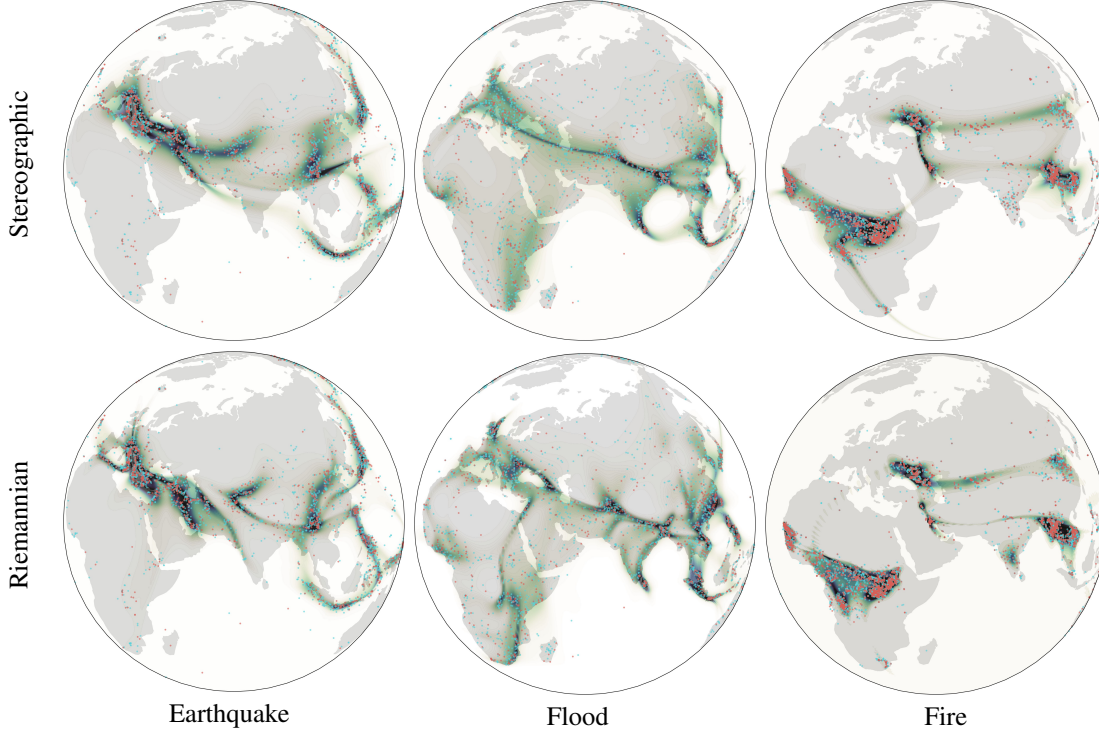


Figure 6: Density estimation for earth sciences data. Blue and red dots represent training and testing datapoints, respectively. Heatmaps depict the log-likelihood of the trained models.

Compared to a well-optimized linear layer, the use of the geodesic distance layer (see Section 2) induces an extra computational cost as shown in Figure 7. Empirically, the geodesic layer helps to improve performance in the hyperbolic setting but had less of an effect in the spherical setting. As such, the geodesic layer can be regarded as an optional component that can improve the quality of the model at an additional computational cost.

5 Discussion

In this paper we proposed a principled way to parametrize expressive probability distributions on Riemannian manifolds. Specifically, we introduced *Riemannian continuous normalizing flows* in which flows are defined via vector fields on manifolds and computed as the solution to the associated ODE. We empirically demonstrated that this method can yield substantial improvements when modelling data on constant curvature manifolds compared to conventional or projected flows.

Broader impact

The work presented in this paper focuses on the learning of well-specified probabilistic models for manifold-valued data. Consequently, its applications are especially promising to advance scientific understanding in fields such as earth and climate science, computational biology, and computer vision. As a foundational method, our work inherits the broader ethical aspects and future societal consequences of machine learning in general.

Acknowledgments

We are grateful to Adam Foster, Yann Dubois, Laura Ruis, Anthony Caterini, Adam Golinski, Chris Maddison, Salem Said, Alessandro Barp, Tom Rainforth and Yee Whye Teh for fruitful discussions and support. EM research leading to these results received funding from the European Research Council under the European Union’s Seventh Framework Programme (FP7/2007- 2013) ERC grant

agreement no. 617071 and he acknowledges Microsoft Research and EPSRC for funding EM's studentship.

References

- Al-Mohy, A. H. and Higham, N. J. (2010). A New Scaling and Squaring Algorithm for the Matrix Exponential. *SIAM Journal on Matrix Analysis and Applications*, 31(3):970–989.
- Ambrosio, L. (2003). Optimal transport maps in Monge-Kantorovich problem.
- Avron, H. and Toledo, S. (2011). Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *Journal of the ACM*, 58(2).
- Bécigneul, G. and Ganeva, O.-E. (2019). Riemannian Adaptive Optimization Methods. *arXiv:1810.00760 [cs, stat]*.
- Benamou, J.-D. and Brenier, Y. (2000). A computational fluid mechanics solution to the Monge-Kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393.
- Blumenson, L. E. (1960). A derivation of n-Dimensional spherical coordinates. *The American Mathematical Monthly*, 67(1):63–66.
- Bonnabel, S. (2013). Stochastic gradient descent on Riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229.
- Bose, A. J., Smofsky, A., Liao, R., Panangaden, P., and Hamilton, W. L. (2020). Latent Variable Modelling with Hyperbolic Normalizing Flows. *arXiv:2002.06336 [cs, stat]*.
- Brakenridge, G. (2017). Global active archive of large flood events. <http://floodobservatory.colorado.edu/Archives/index.html>. Dartmouth Flood Observatory, University of Colorado,.
- Chen, R. T. Q., Behrmann, J., Duvenaud, D., and Jacobsen, J.-H. (2020). Residual Flows for Invertible Generative Modeling. *arXiv:1906.02735 [cs, stat]*.
- Chen, R. T. Q., Rubanova, Y., Bettencourt, J., and Duvenaud, D. (2019). Neural Ordinary Differential Equations. *arXiv:1806.07366 [cs, stat]*.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. (2018). Hyperspherical Variational Auto-Encoders. *arXiv:1804.00891 [cs, stat]*.
- Dinh, L., Sohl-Dickstein, J., and Bengio, S. (2017). Density estimation using Real NVP. *arXiv:1605.08803 [cs, stat]*.
- Dooley, A. and Wildberger, N. (1993). Harmonic analysis and the global exponential map for compact Lie groups. *Functional Analysis and Its Applications*, 27(1):21–27.
- Dormand, R. J. and Prince, J. P. (1980). A family of embedded Runge-Kutta formulae. *Journal of Computational and Applied Mathematics*, pages 19–26.
- Downs (1972). Orientational statistics. *Biometrika*, 59.
- Drucker, H. and Cun, Y. L. (1992). Improving Generalization Performance Using Double Backpropagation. *IEEE Transactions on Neural Networks and Learning Systems*, 3(6):991–997.
- EOSDIS (2020). Active fire data. <https://earthdata.nasa.gov/earth-observation-data/near-real-time/firms/active-fire-data>. Land, Atmosphere Near real-time Capability for EOS (LANCE) system operated by NASA's Earth Science Data and Information System (ESDIS).
- Falorsi, L., de Haan, P., Davidson, T. R., and Forré, P. (2019). Reparameterizing Distributions on Lie Groups. *arXiv:1903.02958 [cs, math, stat]*.
- Falorsi, L. and Forré, P. (2020). Neural Ordinary Differential Equations on Manifolds. *arXiv:2006.06663 [cs, stat]*.

- Finlay, C., Jacobsen, J.-H., Nurbekyan, L., and Oberman, A. M. (2020). How to train your neural ODE. *arXiv:2002.02798 [cs, stat]*.
- Ganea, O.-E., Bécigneul, G., and Hofmann, T. (2018). Hyperbolic Neural Networks. *arXiv:1805.09112 [cs, stat]*.
- Gemici, M. C., Rezende, D., and Mohamed, S. (2016). Normalizing Flows on Riemannian Manifolds. *arXiv:1611.02304 [cs, math, stat]*.
- Grathwohl, W., Chen, R. T. Q., Bettencourt, J., Sutskever, I., and Duvenaud, D. (2018). FFIORD: Free-form Continuous Dynamics for Scalable Reversible Generative Models. *arXiv:1810.01367 [cs, stat]*.
- Hairer, E. (2006). *Geometric Numerical Integration : Structure-Preserving Algorithms for Ordinary Differential Equations*. Springer Series in Computational Mathematics ; 31. Springer, Berlin, 2nd ed. edition.
- Hairer, E. (2011). Solving Differential Equations on Manifolds. page 55.
- Hamelryck, T., Kent, J. T., and Krogh, A. (2006). Sampling Realistic Protein Conformations Using Local Structural Bias. *PLoS Computational Biology*, 2(9).
- Hasnat, M. A., Bohné, J., Milgram, J., Gentric, S., and Chen, L. (2017). Von Mises-Fisher Mixture Model-based Deep learning: Application to Face Verification. *arXiv:1706.04264 [cs]*.
- Hutchinson (1990). A stochastic estimator of the trace of the influence matrix for laplacian smoothing splines. *Communications in Statistics - Simulation and Computation*, 19(2):433–450.
- Karpatne, A., Ebert-Uphoff, I., Ravela, S., Babaie, H. A., and Kumar, V. (2017). Machine Learning for the Geosciences: Challenges and Opportunities. *arXiv:1711.04708 [physics]*.
- Kingma, D. P. and Ba, J. (2015). Adam: A Method for Stochastic Optimization. *arXiv:1412.6980 [cs]*.
- Kingma, D. P. and Welling, M. (2014). Auto-Encoding Variational Bayes. *arXiv:1312.6114 [cs, stat]*.
- Klimovskaia, A., Lopez-Paz, D., Bottou, L., and Nickel, M. (2019). Poincaré Maps for Analyzing Complex Hierarchies in Single-Cell Data. Preprint, Bioinformatics.
- Lee, J. M. (2003). *Introduction to Smooth Manifolds*. Number 218 in Graduate Texts in Mathematics. Springer, New York.
- Lou, A., Lim, D., Katsman, I., Huang, L., Jiang, Q., Lim, S.-N., and De Sa, C. (2020). Neural manifold ordinary differential equations.
- Lui, Y. M. (2012). Advances in matrix manifolds for computer vision. *Image and Vision Computing*, 30(6):380–388.
- Mardia, K. V. and Jupp, P. E. (2000). *Directional Statistics*. Wiley Series in Probability and Statistics. J. Wiley, Chichester ; New York.
- Mathieu, E., Lan, C. L., Maddison, C. J., Tomioka, R., and Teh, Y. W. (2019). Continuous Hierarchical Representations with Poincaré Variational Auto-Encoders.
- Nagano, Y., Yamaguchi, S., Fujita, Y., and Koyama, M. (2019). A Wrapped Normal Distribution on Hyperbolic Space for Gradient-Based Learning. *arXiv:1902.02992 [cs, stat]*.
- Nickel, M. and Kiela, D. (2017). Poincaré Embeddings for Learning Hierarchical Representations. *arXiv:1705.08039 [cs, stat]*.
- NOAA (2020a). Global significant earthquake database. <https://data.nodc.noaa.gov/cgi-bin/iso?id=gov.noaa.ngdc.mgg.hazards:G012153>. National Geophysical Data Center / World Data Service (NGDC/WDS): NCEI/WDS Global Significant Earthquake Database. NOAA National Centers for Environmental Information.

- NOAA (2020b). Global significant volcanic eruptions database. <https://data.nodc.noaa.gov/cgi-bin/iso?id=gov.noaa.ngdc.mgg.hazards:G10147>. National Geophysical Data Center / World Data Service (NGDC/WDS): NCEI/WDS Global Significant Volcanic Eruptions Database. NOAA National Centers for Environmental Information.
- Novak, R., Bahri, Y., Abolafia, D. A., Pennington, J., and Sohl-Dickstein, J. (2018). Sensitivity and Generalization in Neural Networks: An Empirical Study.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2019). Normalizing Flows for Probabilistic Modeling and Inference. *arXiv:1912.02762 [cs, stat]*.
- Papamakarios, G., Pavlakou, T., and Murray, I. (2018). Masked Autoregressive Flow for Density Estimation. *arXiv:1705.07057 [cs, stat]*.
- Paszke, A., Gross, S., Chintala, S., Chanan, G., Yang, E., DeVito, Z., Lin, Z., Desmaison, A., Antiga, L., and Lerer, A. (2017). Automatic differentiation in PyTorch. In *NIPS-W*.
- Pennec, X. (2006). Intrinsic Statistics on Riemannian Manifolds: Basic Tools for Geometric Measurements. *Journal of Mathematical Imaging and Vision*, 25(1):127–154.
- Petersen, P. (2006). *Riemannian Geometry*. Springer-Verlag New York.
- Rezende, D. J. and Mohamed, S. (2016). Variational Inference with Normalizing Flows. *arXiv:1505.05770 [cs, stat]*.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic Backpropagation and Approximate Inference in Deep Generative Models. *arXiv:1401.4082 [cs, stat]*.
- Rezende, D. J., Papamakarios, G., Racanière, S., Albergo, M. S., Kanwar, G., Shanahan, P. E., and Cranmer, K. (2020). Normalizing Flows on Tori and Spheres. *arXiv:2002.02428 [cs, stat]*.
- Salman, H., Yadollahpour, P., Fletcher, T., and Batmanghelich, K. (2018). Deep Diffeomorphic Normalizing Flows. *arXiv:1810.03256 [cs, stat]*.
- Sei, T. (2011). A Jacobian Inequality for Gradient Maps on the Sphere and Its Application to Directional Statistics. *Communications in Statistics - Theory and Methods*, 42(14):2525–2542.
- Skopek, O., Ganea, O.-E., and Bécigneul, G. (2019). Mixed-curvature Variational Autoencoders.
- Ungar, A. A. (2008). A gyrovector space approach to hyperbolic geometry. *Synthesis Lectures on Mathematics and Statistics*, 1(1):1–194.

Appendix for

Riemannian Continuous Normalizing Flows

A Constant curvature manifolds

In the following, we provide a brief overview of Riemannian geometry and constant curvature manifolds, specifically the Poincaré ball and the hypersphere models. We will use $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ to denote the Euclidean norm and inner product. For norms and inner products on tangent spaces $\mathcal{T}_z\mathcal{M}$, we write $\|\cdot\|_z$ and $\langle \cdot, \cdot \rangle_z$ where $z \in \mathcal{M}$.

A.1 Review of Riemannian geometry

A real, smooth *manifold* \mathcal{M} is a set of points z , which is "locally similar" to a linear space. For every point z of the manifold \mathcal{M} is attached a real vector space of the same dimensionality as \mathcal{M} called the *tangent space* $\mathcal{T}_z\mathcal{M}$. Intuitively, it contains all the possible directions in which one can tangentially pass through z . Taking the disjoint union of all tangent spaces yields the *tangent bundle* $\mathcal{TM} = \bigcup_{z \in \mathcal{M}} \mathcal{T}_z\mathcal{M}$. For each point z of the manifold, the *metric tensor* $g(z)$ defines an inner product on the associated tangent space as $g(z) = \langle \cdot, \cdot \rangle_z : \mathcal{T}_z\mathcal{M} \times \mathcal{T}_z\mathcal{M} \rightarrow \mathbb{R}$. The *matrix representation of the Riemannian metric* $G(z)$, is defined such that

$$\forall u, v \in \mathcal{T}_z\mathcal{M} \times \mathcal{T}_z\mathcal{M}, \langle u, v \rangle_z = g(z)(u, v) = u^T G(z) v.$$

A *Riemannian manifold* is then given as a tuple (\mathcal{M}, g) (Petersen, 2006). The metric tensor gives a *local* notion of angle, length of curves, surface area and volume, from which *global* quantities can be derived by integrating local contributions. A norm is induced by the inner product on $\mathcal{T}_z\mathcal{M}$: $\|\cdot\|_z = \sqrt{\langle \cdot, \cdot \rangle_z}$. An infinitesimal volume element is induced on each tangent space $\mathcal{T}_z\mathcal{M}$, and thus a measure $d\text{Vol}(z) = \sqrt{|G(z)|} d\text{Leb}(z)$ on the manifold, with $\text{Leb}(z)$ being the Lebesgue measure. The length of a curve $\gamma : t \mapsto \gamma(t) \in \mathcal{M}$ is given by $L(\gamma) = \int_0^1 \|\gamma'(t)\|_{\gamma(t)} dt$. The concept of straight lines can then be generalized to *geodesics*, which are constant speed curves giving the shortest path between pairs of points z, y of the manifold: $\gamma^* = \arg \min L(\gamma)$ with $\gamma(0) = z$, $\gamma(1) = y$ and $\|\gamma'(t)\|_{\gamma(t)} = 1$. A *global* distance is thus induced on \mathcal{M} given by

$$d_{\mathcal{M}}(z, y) = \inf L(\gamma).$$

Endowing \mathcal{M} with that distance consequently defines a metric space $(\mathcal{M}, d_{\mathcal{M}})$. The concept of moving along a "straight" curve with constant velocity is given by the *exponential map*. In particular, there is a unique unit speed *geodesic* γ satisfying $\gamma(0) = z$ with initial tangent vector $\gamma'(0) = v$. The corresponding exponential map is then defined by $\exp_z(v) = \gamma(1)$. The *logarithm map* is the inverse $\log_z = \exp_z^{-1} : \mathcal{M} \rightarrow \mathcal{T}_z\mathcal{M}$. The \exp_z map is well-defined on the full tangent space $\mathcal{T}_z\mathcal{M}$ for all $z \in \mathcal{M}$ if and only if \mathcal{M} is geodesically complete, i.e. if all geodesics can "run" indefinitely. This is the case for the Poincaré ball and hypersphere.

A.2 The Poincaré ball model of hyperbolic geometry

In the following, we provide a brief overview of key concepts related to hyperbolic geometry. A d -dimensional *hyperbolic* space is a complete, simply connected, d -dimensional Riemannian manifold with *constant negative curvature* K . The *Poincaré ball* is one model of this geometry, and is formally defined as the Riemannian manifold $\mathbb{B}_K^d = (\mathcal{B}_K^d, g_K)$. Here \mathcal{B}_K^d denotes the open ball of radius $1/\sqrt{|K|}$, and g_K the *metric tensor* $g_K(z) = (\lambda_z^K)^2 g^e(z)$, where $\lambda_z^K = \frac{2}{1+K\|z\|^2}$ and g^e denotes the Euclidean metric tensor, i.e. the usual dot product. The induced *invariant measure* Vol is absolutely continuous with respect to the Lebesgue measure Leb , and its density is given by $\frac{d\text{Vol}}{d\text{Leb}}(z) = \sqrt{|G(z)|} = (\lambda_z^K)^d$ for all $z \in \mathbb{B}_K^d$. As motivated by Skopek et al. (2019), the Poincaré ball \mathbb{B}_K^d can conveniently be described through the formalism of *gyrovector spaces* (Ungar, 2008). These can be seen as an analogy to the way vector spaces are used in Euclidean geometry, but in the non-Euclidean geometry setting. In

particular, the *Möbius addition* \oplus_K of \mathbf{z}, \mathbf{y} in \mathbb{B}_K^d is defined as

$$\mathbf{z} \oplus_K \mathbf{y} = \frac{(1 - 2K \langle \mathbf{z}, \mathbf{y} \rangle - K \|\mathbf{y}\|^2) \mathbf{z} + (1 + K \|\mathbf{z}\|^2) \mathbf{y}}{1 - 2K \langle \mathbf{z}, \mathbf{y} \rangle + K^2 \|\mathbf{z}\|^2 \|\mathbf{y}\|^2}.$$

Then the *exponential map* can be expressed via this *Möbius addition* as

$$\exp_z^K(\mathbf{v}) = \mathbf{z} \oplus_K \left(\tanh \left(\sqrt{-K} \frac{\lambda_z^K \|\mathbf{v}\|}{2} \right) \frac{\mathbf{v}}{\sqrt{-K} \|\mathbf{v}\|} \right)$$

where $\mathbf{x} = -\mathbf{z} \oplus_K \mathbf{y}$ for all $\mathbf{x}, \mathbf{y} \in \mathbb{B}_K^d$.

A.3 The hypersphere model of elliptic geometry

In the following, we discuss key concepts related to positively curved spaces known as *elliptic* spaces, and in particular to the *hypersphere* model. The d -sphere, or *hyperpsphere*, is a compact submanifold of \mathbb{R}^{d+1} with positive constant curvature K whose support is defined by $\mathcal{S}_K^d = \{\mathbf{z} \in \mathbb{R}^{d+1} \mid \langle \mathbf{z}, \mathbf{z} \rangle = 1/K\}$. It is endowed with the pull-back metric of the ambient Euclidean space.

Sphere In the two-dimensional setting $d = 2$, we rely on polar coordinates to parametrize the sphere \mathbb{S}^2 . These coordinates consist of polar $\theta \in [0, \pi]$ and azimuth $\varphi \in [0, 2\pi)$ angles. The ambient Cartesian coordinates are then given by $\mathbf{r}(\theta, \varphi) = (\sin(\theta) \cos(\varphi), \sin(\theta) \sin(\varphi), \cos(\theta))$. We have $\sqrt{|G(\theta, \varphi)|} = \sin(\theta)$. Applying the generic divergence formula (see Equation 11) yields the celebrated spherical divergence formula

$$\operatorname{div}(\mathbf{g}) = \frac{1}{\sin(\theta)} \frac{\partial}{\partial \theta} (\sin(\theta) g^\theta(\theta, \varphi)) + \frac{1}{\sin(\theta)} \frac{\partial}{\partial \varphi} (g^\varphi(\theta, \varphi)).$$

Hypersphere For higher dimensions, we can rely on the n -spherical coordinate system in which the coordinates consist of $d - 1$ angular coordinates $\varphi_1, \dots, \varphi_{d-2} \in [0, \pi]$ and $\varphi_{d-2} \in [0, 2\pi)$ (Blumenson, 1960). Then we have $\sqrt{|G(\varphi)|} = \sin^{d-2}(\varphi_1) \sin^{d-3}(\varphi_2) \dots \sin(\varphi_{d-2})$.

Using the ambient cartesian coordinates, the *exponential map* is given by

$$\exp_\mu^c(\mathbf{v}) = \cos(\sqrt{K} \|\mathbf{v}\|) \boldsymbol{\mu} + \sin(\sqrt{K} \|\mathbf{v}\|) \frac{\mathbf{v}}{\sqrt{K} \|\mathbf{v}\|}$$

for all $\mathbf{z} \in \mathbb{S}_K^d$ and $\mathbf{v} \in \mathcal{T}_z \mathbb{S}_K^d$.

B Probability measures on Riemannian manifolds

In what follows, we discuss core concepts of probability measures on Riemannian manifolds and show how naive methods lead to ill- and mis-specified models on manifolds.

Probability measures and random vectors can intrinsically be defined on Riemannian manifolds so as to model uncertainty on non-flat spaces (Pennec, 2006). The Riemannian metric $G(\mathbf{z})$ induces an infinitesimal volume element on each tangent space $\mathcal{T}_z \mathcal{M}$, and thus a measure on the manifold,

$$d \operatorname{Vol}(\mathbf{z}) = \sqrt{|G(\mathbf{z})|} d \operatorname{Leb},$$

with Leb the Lebesgue measure. Manifold-valued random variables would naturally be characterized by the *Radon-Nikodym derivative* of a measure ν w.r.t. the Riemannian measure Vol (assuming absolute continuity)

$$p(\mathbf{z}) = \frac{d\nu}{d \operatorname{Vol}}(\mathbf{z}).$$

B.1 Ambient Euclidean probability distributions

Unfortunately, conventional probabilistic models implicitly assume a *flat* geometry. This in turn cause these models to either be misspecified or ill-suited to fit manifold distributions. Below we discuss the reasons why.

Let $P_{\mathcal{D}}$ be a target probability measure that we aim to approximate, and which is defined on a d -dimensional manifold $\mathcal{M} \subseteq \mathbb{R}^D$. Furthermore, we assume it admits a Radon-Nikodym derivative $p_{\mathcal{D}}$ with respect to the manifold invariant measure Vol , denoting $P_{\mathcal{D}} \ll \text{Vol}$ with \ll denoting absolute continuity. Conventional normalizing flows implicitly assume the parametrized probability measure P_{θ} to have support on the ambient space \mathbb{R}^D and to be absolutely continuous with respect to the Lebesgue measure $\text{Leb}_{\mathbb{R}^D}$. We denote its density by p_{θ} .

Next, assume $D = d$, such as for $\mathcal{M} = \mathbb{B}^d \subseteq \mathbb{R}^d$. With z the d -dimensional Cartesian coordinates, we have $\frac{d\text{Vol}}{d\text{Leb}_{\mathbb{R}^d}}(z) = \sqrt{|G(z)|}$. One could then see the manifold-valued target $P_{\mathcal{D}}$ as being a probability measure on \mathbb{R}^d with a density w.r.t. the Lebesgue measure given by

$$\frac{dP_{\mathcal{D}}}{d\text{Leb}_{\mathbb{R}^d}}(z) = p_{\mathcal{D}}(z) \sqrt{|G(z)|} \triangleq \tilde{p}_{\mathcal{D}}(z).$$

In general $P_{\mathcal{D}} \ll P_{\theta}$ which implies that the *forward* Kullback-Leibler divergence, or *negative log-likelihood* up to constants, is defined and given by

$$\mathcal{L}^{\text{Like}}(\theta) = D_{\text{KL}}(P_{\mathcal{D}} \parallel P_{\theta}) + \mathbb{H}(P_{\mathcal{D}}) = \mathbb{E}_{P_{\mathcal{D}}} \left[\log \left(\frac{\tilde{p}_{\mathcal{D}}(z)}{p_{\theta}(z)} \right) \right] + \mathbb{H}(P_{\mathcal{D}}) = -\mathbb{E}_{P_{\mathcal{D}}} \left[\log \frac{p_{\theta}(z)}{\sqrt{|G(z)|}} \right].$$

Minimising $\mathcal{L}^{\text{Like}}(\theta)$ amounts to pushing-forward P_{θ} 's mass so that empirical observations $z_i \sim P_{\mathcal{D}}$ have a positive likelihood under P_{θ} . Yet, in general the model (P_{θ}) has (most of his) mass outside the manifold's support which may cause such a *naive* approach to be ill-suited. More crucially it implies that in general the model's mass is *not* covering the full target's support. In that case, the *reverse* Kullback-Leibler divergence $D_{\text{KL}}(P_{\theta} \parallel P_{\mathcal{D}}) = \mathcal{L}^{\text{KL}}(\theta)$ is not even defined.

Next, consider the case where \mathcal{M} is a submanifold embedded in \mathbb{R}^D with $D > d$, such as $\mathcal{M} = \mathbb{S}^d$ where $D = d + 1$. In this setting the *naive* model P_{θ} is even *misspecified* since it is defined on a different probability space than the target. In the limit $\text{supp}(P_{\theta}) \rightarrow \mathcal{M}$, P_{θ} is not defined because we have that $\int_{\mathbb{R}^D} P_{\theta} \rightarrow \infty$. The target does consequently not belong to the model's class.

C Instantaneous change of variable

In the following we derive the *instantaneous change of density* that a manifold-valued random variable induces when its dynamics are governed by an ODE. We show that in the Riemannian setting this *instantaneous change of density* can be expressed in terms of the manifold's metric tensor.

Proof of Proposition 2

Proof. For a time dependant particles $z(t)$, whose dynamics are given by the following ODE

$$\frac{dz(t)}{dt} = f(z(t), t)$$

the change in density is given by the Liouville equation (or Fokker-Planck equation without the diffusion term); $\forall z \in \mathcal{M}, \forall t \in [0, T]$

$$\begin{aligned} \frac{\partial}{\partial t} p(z, t) &= -\text{div}(p(z, t) f(z, t)) \\ &= -\left\langle \frac{\partial}{\partial z} p(z, t), f(z, t) \right\rangle_z - p(z, t) \text{div}(f(z, t)) \end{aligned}$$

where the last step was obtained by applying the divergence product rule. By introducing the time dependence in $z(t)$ and differentiating with respect to time we get

$$\begin{aligned} \frac{\partial}{\partial t} p(z(t), t) &= \left\langle \frac{\partial}{\partial z} p(z(t), t), \frac{\partial}{\partial t} z(t) \right\rangle_{z(t)} + \frac{\partial}{\partial t} p(z(t), t) \\ &= \left\langle \frac{\partial}{\partial z} p(z(t), t), f(z(t), t) \right\rangle_{z(t)} - \left\langle \frac{\partial}{\partial z} p(z(t), t), f(z(t), t) \right\rangle_{z(t)} - p(z(t), t) \text{div}(f(z(t), t)) \\ &= -p(z(t), t) \text{div}(f(z(t), t)) \end{aligned}$$

Hence the evolution of the log density is given by

$$\frac{\partial}{\partial t} \log p(z(t), t) = -\text{div}(f(z(t), t)). \quad (8)$$

□

Divergence computation For a Riemannian manifold (\mathcal{M}, g) , with local coordinates z , the divergence of a vector field \mathbf{f} is given by

$$\operatorname{div}(\mathbf{f}(z, t)) = \frac{1}{\sqrt{|G(z)|}} \sum_{i=1}^d \frac{\partial}{\partial z^i} \left(\sqrt{|G(z)|} f^i(z, t) \right) \quad (9)$$

$$\begin{aligned} &= \frac{1}{\sqrt{|G(z)|}} \sum_{i=1}^d \left(\sqrt{|G(z)|} \frac{\partial}{\partial z^i} f^i(z, t) + f^i(z, t) \frac{\partial}{\partial z^i} \sqrt{|G(z)|} \right) \\ &= \sum_{i=1}^d \frac{\partial}{\partial z^i} f^i(z, t) + \frac{1}{\sqrt{|G(z)|}} \sum_{i=1}^d f^i(z, t) \frac{\partial}{\partial z^i} \sqrt{|G(z)|} \\ &= \operatorname{tr} \left(\frac{\partial}{\partial z} \mathbf{f}(z, t) \right) + \frac{1}{\sqrt{|G(z)|}} \left\langle \mathbf{f}(z, t), \frac{\partial}{\partial z} \sqrt{|G(z)|} \right\rangle. \end{aligned} \quad (10)$$

We note that in Equation 9, f_i are the components of the vector field \mathbf{f} with respect to the local unnormalized covariant basis $(e_i)_{i=1}^d = \left(\left(\frac{\partial}{\partial z^i} \right)_z \right)_{i=1}^d$. However it is convenient to work with local basis having unit length vectors. If we write \hat{e}_i for this normalized basis, and \hat{f}^i for the components of \mathbf{f} with respect to this normalized basis, we have that

$$\mathbf{f} = \sum_i f^i e_i = \sum_i f^i \|e_i\| \frac{e_i}{\|e_i\|} = \sum_i f^i \sqrt{G_{ii}} \frac{e_i}{\|e_i\|} = \sum_i \hat{f}^i \hat{e}_i$$

using one of the properties of the metric tensor. By dotting both sides of the last equality with the contravariant element \hat{e}_i we get that $\hat{f}^i = f^i \sqrt{G_{ii}}$. Substituting in Equation 9 yields

$$\operatorname{div}(\hat{\mathbf{f}}(z, t)) = \frac{1}{\sqrt{|G(z)|}} \sum_{i=1}^d \frac{\partial}{\partial z^i} \left(\sqrt{\frac{|G(z)|}{G_{ii}(z)}} \hat{f}^i(z, t) \right). \quad (11)$$

Combining Equations 8 and 11 and we finally get

$$\frac{\partial \log p(z(t), t)}{\partial t} = - \frac{1}{\sqrt{|G(z)|}} \sum_{i=1}^d \frac{\partial}{\partial z^i} \left(\sqrt{\frac{|G(z)|}{G_{ii}(z)}} \hat{f}^i(z, t) \right). \quad (12)$$

We rely on this Equation 12 for practical numerical experiments.

D regularization

D.1 l^2 -norm

Henceforth we motivate the use of an l^2 norm regularization in the context of continuous normalizing flows. We do so by highlighting a connection with the dynamical formulation of optimal transport, and by proving that this formulation still holds in the manifold setting.

Monge-Kantorovich mass transfer problem Let $(\mathcal{M}, d_{\mathcal{M}})$ be a metric space, and $c : \mathcal{M} \times \mathcal{M} \rightarrow [0, +\infty)$ a measurable map. Given probability measures p_0 and p_T on \mathcal{M} , Monge's formulation of the optimal transportation problem is to find a transport map $\phi^* : \mathcal{M} \rightarrow \mathcal{M}$ that realizes the infimum

$$\inf_{\phi} \int_{\mathcal{M}} c(\phi(z), z) p_0(dz) \text{ s.t. } p_T \# = p_0.$$

It can be shown that this yields a metric on probability measures, and for $c = d_{\mathcal{M}}^2$, it is called the L^2 Kantorovich (or Wasserstein) distance

$$d_{W^2}(p_0, p_T)^2 = \inf_{\phi} \int_{\mathcal{M}} d_{\mathcal{M}}(\phi(z), z)^2 p_0(dz). \quad (13)$$

By reintroducing the time variable in the L^2 Monge-Kantorovich mass transfer problem, the optimal transport map ϕ^* can be reformulated as the generated flow from an optimal vector field \mathbf{f} .

Proposition 3 (Dynamical formulation from (Benamou and Brenier, 2000)). *Indeed we have*

$$d_{W^2}(p_0, p_T)^2 = \inf \frac{1}{T} \int_0^T \|f\|_{p_t}^2 dt = \inf \frac{1}{T} \int_0^T \int_{\mathcal{M}} \langle f(z, t), f(z, t) \rangle_z p_t(dz) dt \quad (14)$$

where the infimum is taken among all weakly continuous distributional solutions of the continuity equation $\frac{\partial}{\partial t} p_t = -\text{div}(p_t f)$ such that $p(0) = p_0$ and $p(T) = p_T$. Writing $\phi_t^* = \phi^*(\cdot, t)$ the flow generated by the optimal ODE, then the optimal transport map is given by $\phi^* = \phi_T^*$.

The RHS of Equation (14) can then be approximated with no extra-cost with a Monte Carlo estimator given samples from $p_t = \phi_t \# p_0$.

Manifold Setting Let's now focus on the setting where \mathcal{M} is a Riemannian manifold.

Proposition 4 (Optimal map (Ambrosio, 2003)). *Assume that \mathcal{M} is a C^3 , complete Riemannian manifold with no boundary and $d_{\mathcal{M}}$ is the Riemannian distance. If p_0, p_T have finite second order moments and p_0 is absolutely continuous with respect to $\text{vol}_{\mathcal{M}}$, then there exists a unique optimal transport map ϕ for the Monge-Kantorovich problem with cost $c = d_{\mathcal{M}}^2$. Moreover there exists a potential $h : \mathcal{M} \mapsto \mathbb{R}$ such that*

$$\phi^*(z) = \exp_z(-\nabla h(z)) \quad \text{vol}_{\mathcal{M}} - a.e..$$

Proposition 3 has been stated and proved for the case $\mathcal{M} = \mathbb{R}^d$. Below we extend the proof given by Benamou and Brenier (2000) for the manifold setting.

Proof of Proposition 3. We follow the same reasoning as the one developed for the Euclidean setting. Let's first upper bound the Wasserstein distance, and then state the optimal flow which yields equality. We have

$$\begin{aligned} \frac{1}{T} \int_0^T \int_{\mathcal{M}} \|f(z, t)\|_z^2 p_t(dz) dt &= \frac{1}{T} \int_0^T \int_{\mathcal{M}} \|f(\phi(z, t), t)\|_z^2 p_0(dz) dt \\ &= \frac{1}{T} \int_0^T \int_{\mathcal{M}} \left\| \frac{\partial}{\partial t} \phi(z, t) \right\|_z^2 p_0(dz) dt \\ &\geq \int_{\mathcal{M}} d_{\mathcal{M}}(\phi(z, T), \phi(z, 0))^2 p_0(dz) dt \\ &= \int_{\mathcal{M}} d_{\mathcal{M}}(\phi(z, T), z)^2 p_0(dz) dt \\ &\geq \int_{\mathcal{M}} d_{\mathcal{M}}(\phi(z), z)^2 p_0(dz) dt \\ &= d_{W^2}(p_0, p_T)^2. \end{aligned}$$

Thus, the optimal choice of flow ϕ is given by

$$\phi(z, t) = \exp_z \left(\frac{t}{T} \log_z(\phi^*(z)) \right), \quad (15)$$

since $\phi(z, 0) = z$, $\phi(z, T) = \phi^*(z)$ and

$$\left\| \frac{\partial}{\partial t} \phi(z, t) \right\|_z = \left\| \frac{\partial}{\partial t} \phi(z, t=0) \right\|_z = \left\| \log_z(\phi^*(z)) \right\|_z = d_{\mathcal{M}}(\phi^*(z), z).$$

□

Note that the optimal flow from Equation 15 yields integral paths $\gamma(t) = \phi(z, t)$ that are *geodesics* and have constant *velocity*.

Motivation Regularizing the vector field with the RHS of Equation 14 would hence tend to make the generated flow ϕ_T closer to the optimal map ψ^* . By doing so, one hopes to increase smoothness of f and consequently lower the solver NFE given a fixed tolerance.

This has been observed in the Euclidean setting by Finlay et al. (2020). They empirically showed that regularizing the loss of a CNF with the vector field's l^2 norm improves training speed. Motivated by the successful use of gradient regularization (Novak et al., 2018; Drucker and Cun, 1992), they showed that additionally regularizing the Frobenius norm of the vector field's Jacobian helps. In the following subsection we remind that this regularization term can also be motivated from an estimator's variance perspective.

D.2 Frobenius norm

Hutchinson’s estimator Hutchinson’s estimator (Hutchinson, 1990) is a simple way to obtain a stochastic estimate of the trace of a matrix. Given a d -dimensional random vector $\epsilon \sim p$ such that $\mathbb{E}[\epsilon] = 0$ and $\text{Cov}(\epsilon) = I_d$, we have

$$\text{tr}(A) = \mathbb{E}_{\epsilon \sim p}[\epsilon^T A \epsilon].$$

Rademacher and Gaussian distributions have been used in practice. For a Rademacher, the variance is given by (Avron and Toledo, 2011)

$$\mathbb{V}_{\epsilon \sim p}[\epsilon^T A \epsilon] = 2 \|A\|_F - 2 \sum_i A_{ii}^2,$$

whereas for a Gaussian it is given by

$$\mathbb{V}_{\epsilon \sim p}[\epsilon^T A \epsilon] = 2 \|A\|_F.$$

Divergence computation As reminded in Appendix C by Equation 10, computing the vector field divergence $\text{div}(f(z, t))$ involves the computation of the trace of vector field’s Jacobian $\text{tr}\left(\frac{\partial}{\partial z} f(z, t)\right)$. As highlighted in Grathwohl et al. (2018); Salman et al. (2018), one can rely on the Hutchinson’s estimator to estimate this trace with $A = \frac{\partial}{\partial z} f(z, t)$.

The variance of this estimator thus depends on the Frobenius norm of the vector’s field Jacobian $\|\frac{\partial}{\partial z} f(z, t)\|_F$, as noted in Grathwohl et al. (2018). Regularizing this Jacobian should then improve training by reducing the variance of the divergence estimator.

E Vector flows and neural architecture

Hereafter we discuss about flows generated by vector fields, and neural architectural choices that we make for their parametrization. Properties of vector fields have direct consequences on the properties of the generated flow and in turn on the associated pushforward probability distributions. In particular we derive sufficient conditions on the flow so that it is *global*, i.e. is a bijection mapping the manifold to itself.

E.1 Existence and uniqueness of a global flow

We start by discussing about vector flows and sufficient conditions on their uniqueness and existence.

Local flow First we remind the Fundamental theorem of flows (Lee, 2003) which gives the existence and uniqueness of a smooth *local* flow.

Proposition 5 (Fundamental theorem of flows). *Let \mathcal{M} be a smooth complete manifold with local coordinates z . Let $f_\theta : \mathcal{M} \times \mathbb{R} \mapsto \mathcal{T}\mathcal{M}$ a C^1 time- dependent vector field and $z_0 \in \mathcal{M}$. Then there exists an open interval I with $0 \in I$, an open subset $U \subseteq \mathcal{M}$ containing z_0 , and a unique smooth map $\phi : I \times U \mapsto \mathcal{M}$ called local flow which satisfies the following properties. We write $\phi_t(z) = \phi(z, t)$.*

1. $\frac{\partial}{\partial t} \phi(z, t) = f_\theta(\phi(z, t), t)$ for all $z, t \in U \times I$, and $\phi_0 = \text{id}_{\mathcal{M}}$.
2. For each $t \in I$, the map $\phi_t : U \mapsto \mathcal{M}$ is a local C^1 -diffeomorphism.

Note that with such assumptions, the existence and uniqueness of flows ϕ_t are only *local*.

Global flow We would like the flow ϕ to be defined for all times and on the whole manifold, i.e. a *global* flow $\phi : \mathcal{M} \times \mathbb{R} \mapsto \mathcal{M}$. Fortunately, if \mathcal{M} is *compact* (such as n -spheres and torii), then the flow is global (Lee, 2003). We show below that another sufficient condition for the flow to be global is that the vector field be bounded.

Proposition 6 (Global Flow). *Let \mathcal{M} be a smooth complete manifold. Let $f_\theta : \mathcal{M} \times \mathbb{R} \mapsto \mathcal{T}\mathcal{M}$ be a C^1 bounded time-dependent vector field. Then the domain of the flow ϕ is $\mathbb{R} \times \mathcal{M}$, i.e. the flow is global.*

Corollary 6.1. *For each $t \in \mathbb{R}$, the map $\phi_t : \mathcal{M} \mapsto \mathcal{M}$ is a C^1 -diffeomorphism.*

Proof of Proposition 6. Let $c > 0$ s.t. $\|\mathbf{f}\| < c$, and $\mathbf{z}_0 \in \mathcal{M}$ be an initial point. Proposition 5 gives the existence of an open interval $I = (a, b)$, a neighbourhood U of \mathbf{z}_0 and a local flow $\phi : (a, b) \times U \mapsto \mathcal{M}$. We write $\gamma = \phi(\mathbf{z}_0, \cdot)$. The maximal interval of γ is (a, b) , which means that γ cannot be extended outside (a, b) . Suppose that $b < \infty$.

The integral path γ is Lipschitz continuous on (a, b) since we have

$$d_{\mathcal{M}}(\gamma(t), \gamma(s)) \leq \int_s^t \|\gamma'(t)\| dt = \int_s^t \|\mathbf{f}(\gamma(t), t)\| dt \leq c |t - s| \quad (16)$$

for all $s < t \in (a, b)$.

Let (t_n) be a sequence in (a, b) that converges to b . Then since (t_n) is a convergent sequence, it must also be a Cauchy sequence. Then $\gamma(t_n)$ is also a Cauchy sequence by Equation 16. Since \mathcal{M} is geodesically complete, it follows by Hopf-Rinow theorem that $(\mathcal{M}, d_{\mathcal{M}})$ is complete, hence that $\gamma(t_n)$ converges to a point $\mathbf{p} \in \mathcal{M}$.

Now suppose that (s_n) is another sequence in (a, b) that converges to b . Then by Equation 16 $\lim_{n \rightarrow \infty} d(\gamma(s_n), \gamma(t_n)) = 0$, thus $\gamma(s_n)$ also converges to $\lim_{n \rightarrow \infty} \gamma(t_n) = \mathbf{p}$. So for every sequence (t_n) in (a, b) that converges to b , we have that $(\gamma(t_n))$ converges to \mathbf{p} . Therefore by the sequential criterion for limits, we have that γ has the limit \mathbf{p} at the point b . Therefore, define $\gamma(b) = \mathbf{p}$ and so γ is continuous at b which is a contradiction. \square

E.2 Geodesic distance layer

The expressiveness of CNFs directly depends on the expressiveness of the vector field and consequently on its architecture. Below we detail and motivate the use of a *geodesic distance* layer, as an input layer for the vector field neural architecture.

Linear layer A linear layer with one *neuron* can be written in the form $h_{\mathbf{a}, \mathbf{p}}(\mathbf{z}) = \langle \mathbf{a}, \mathbf{z} - \mathbf{p} \rangle$, with orientation and offset parameters $\mathbf{a}, \mathbf{p} \in \mathbb{R}^d$. Stacking l such neurons h yields a linear layer with width l . This neuron can be rewritten in the form

$$h_{\mathbf{a}, \mathbf{p}}(\mathbf{z}) = \text{sign}(\langle \mathbf{a}, \mathbf{z} - \mathbf{p} \rangle) \|\mathbf{a}\| d_E(\mathbf{z}, H_{\mathbf{a}, \mathbf{p}}^K)$$

where $H_{\mathbf{a}, \mathbf{p}} = \{\mathbf{z} \in \mathbb{R}^p \mid \langle \mathbf{a}, \mathbf{z} - \mathbf{p} \rangle = 0\} = \mathbf{p} + \{\mathbf{a}\}^\perp$ is the decision hyperplane. The third term is the distance between \mathbf{z} and the decision hyperplane $H_{\mathbf{a}, \mathbf{p}}^K$ and the first term refers to the side of $H_{\mathbf{a}, \mathbf{p}}^K$ where \mathbf{z} lies.

Poincaré ball Ganea et al. (2018) analogously introduced a neuron $f_{\mathbf{a}, \mathbf{p}}^K : \mathbb{B}_K^d \rightarrow \mathbb{R}^p$ on the Poincaré ball,

$$h_{\mathbf{a}, \mathbf{p}}^K(\mathbf{z}) = \text{sign}\left(\left\langle \mathbf{a}, \log_p^K(\mathbf{z}) \right\rangle_p\right) \|\mathbf{a}\|_p d^K(\mathbf{z}, H_{\mathbf{a}, \mathbf{p}}^K) \quad (17)$$

with $H_{\mathbf{a}, \mathbf{p}}^K = \{\mathbf{z} \in \mathbb{B}_K^d \mid \langle \mathbf{a}, \log_p^K(\mathbf{z}) \rangle_p = 0\} = \exp_p^K(\{\mathbf{a}\}^\perp)$. A closed-form expression for the distance $d^K(\mathbf{z}, H_{\mathbf{a}, \mathbf{p}}^K)$ was also derived, $d^K(\mathbf{z}, H_{\mathbf{a}, \mathbf{p}}^K) = \frac{1}{\sqrt{|K|}} \sinh^{-1} \left(\frac{2 \sqrt{|K|} |(-\mathbf{p} \oplus_K \mathbf{z}, \mathbf{a})|}{(1+K \|-\mathbf{p} \oplus_K \mathbf{z}\|^2) \|\mathbf{a}\|} \right)$ in the Poincaré ball. To avoid an over-parametrization of the hyperplane, we set $\mathbf{p} = \exp_0(\mathbf{a}_0)$, and $\mathbf{a} = \Gamma_{0 \rightarrow \mathbf{p}}(\mathbf{a}_0)$ with Γ parallel transport (under Levi-Civita connection). We observed that the term $\|\mathbf{a}\|_p$ from Equation 17 was sometimes causing numerical instabilities, and that when it was not it also did not improve performance. We consequently removed this scaling term. The hyperplane decision boundary $H_{\mathbf{a}, \mathbf{p}}^K$ is called *gyroplane* and is a semi-hypersphere orthogonal to the Poincaré ball's boundary.

Hypersphere In hyperspherical geometry, geodesics are great circles which can be parametrized by a vector $\mathbf{w} \in \mathbb{R}^{d+1}$ as $H_{\mathbf{w}} = \{\mathbf{z} \in \mathbb{S}^d \mid \langle \mathbf{w}, \mathbf{z} \rangle = 0\}$. The geodesic distance between $\mathbf{z} \in \mathbb{S}^d$ and the hyperplane $H_{\mathbf{w}}$ is then given by

$$d(\mathbf{z}, H_{\mathbf{w}}) = \left| \sin^{-1} \left(\frac{\langle \mathbf{w}, \mathbf{z} \rangle}{\sqrt{\langle \mathbf{w}, \mathbf{w} \rangle}} \right) \right|.$$

In a similar fashion, a neuron is now defined by

$$h_{\mathbf{w}}(\mathbf{z}) = \|\mathbf{w}\|_2 \sin^{-1} \left(\frac{\langle \mathbf{w}, \mathbf{z} \rangle}{\sqrt{\langle \mathbf{w}, \mathbf{w} \rangle}} \right).$$

Geodesic distance layer One can then horizontally-stack l neurons to make a *geodesic distance* layer $g : \mathcal{M} \mapsto \mathbb{R}^l$ (Mathieu et al., 2019). Any standard feed-forward neural network can then be vertically-stacked on top of this layer.

F Extensions

F.1 Product of manifolds

Having described CNFs for complete smooth manifolds in Section 2, we extend these for product manifolds $\mathcal{M} = \mathcal{M}_1 \times \cdots \times \mathcal{M}_k$. For instance a d -dimensional torus is defined as $\mathbb{T}^d = \underbrace{\mathbb{S}^1 \times \cdots \times \mathbb{S}^1}_d$.

Any density $p_\theta(z_1, \dots, z_K)$ can be decomposed via the chain rule of probability as

$$p_\theta(z_1, \dots, z_K) = \prod_k p_{\theta_k}(z_k | z_1, \dots, z_{k-1})$$

where each conditional $p_{\theta_k}(z_k | z_1, \dots, z_{k-1})$ is a density on \mathcal{M}_k . As suggested in Rezende et al. (2020), each conditional density can be implemented via a flow $\phi_k : \mathcal{M}_k \mapsto \mathcal{M}_k$ generated by a vector field \mathbf{f}_k , whose parameters θ_k are a function of (z_1, \dots, z_{k-1}) . Such a flow $\phi = \phi_1 \circ \cdots \circ \phi_k$ is called *autoregressive* (Papamakarios et al., 2018) and conveniently has a lower triangular Jacobian, which determinant can be computed efficiently as the product of the diagonal term.

G Experimental details

Below we fully describe the experimental settings used to generate results introduced in Section 4. We open-source our code for reproducibility purposes³.

Architecture The architecture of the *vector field* \mathbf{f}_θ is given by a multilayer perceptron (MLP) with 3 hidden layers and 64 hidden units – as in (Grathwohl et al., 2018) – for *projected* (e.g. stereographic and wrapped cf Section 3) and *naive* (cf Appendix B.1) models. We rely on tanh activation. For our Riemmanian continuous normalizing flow (RCNF), the input layer of the MLP is replaced by a *geodesic distance* layer (Ganea et al., 2018; Mathieu et al., 2019) (see Appendix E.2).

Objectives We consider two objectives, a Monte Carlo (MC) estimator of the negative log-likelihood

$$\hat{\mathcal{L}}^{\text{Like}}(\theta) = - \sum_{i=1}^B \log p_\theta(z_i) \text{ with } z_i \sim P_{\mathcal{D}}$$

and a MC estimator of the reverse KL divergence

$$\hat{\mathcal{L}}^{\text{KL}}(\theta) = \sum_{i=1}^B \log p_\theta(h_\theta(\epsilon_i)) - \log p_{\mathcal{D}}(h_\theta(\epsilon_i))$$

with $z_i \sim P_\theta$ being reparametrized as $z_i = h_\theta(\epsilon_i)$ and $\epsilon_i \sim P$.

Optimization All models are trained by the stochastic optimizer Adam (Kingma and Ba, 2015) with parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, batch-size of 400 data-points and a learning rate set to $1e^{-3}$.

Training We rely on the Dormand-Prince solver (Dormand and Prince, 1980), an adaptive Runge-Kutta 4(5) solver, with absolute and relative tolerance of $1e - 5$ to compute approximate numerical solutions of the ODE. Each solver step is projected onto the manifold.

G.1 Hyperbolic geometry and limits of conventional and wrapped methods

In this experiment the target is set to be a wrapped normal on \mathbb{B}^2 (Nagano et al., 2019; Mathieu et al., 2019) with density $\mathcal{N}^w(\exp_0(\alpha \partial x), \Sigma) = \exp_{\mu^\#} \mathcal{N}(\alpha \partial x, \Sigma)$ with $\Sigma = \text{diag}(0.3, 1.0)$. The scalar parameter α allows us to locate the target closer or further away from the origin of the disk. Through this experiment we consider three CNFs:

³https://github.com/facebookresearch/riemannian_cnf

- *Naive*: $P_\theta^N = \phi_\#^{\mathbb{R}^2} \mathcal{N}(0, 1)$
- *Wrapped*: $P_\theta^W = (\exp_0 \circ \phi^{\mathbb{R}^2})_\# \mathcal{N}(0, 1)$
- *Riemannian*: $P_\theta^R = \phi_\#^{\mathbb{B}^2} \mathcal{N}^W(0, 1)$

with $\phi^{\mathbb{R}^2}$ a conventional CNF on \mathbb{R}^2 , $\phi^{\mathbb{B}^2}$ our RCNF introduced in Section 2, $\mathcal{N}(0, 1)$ the standard Gaussian and $\mathcal{N}^W(0, 1)$ the standard wrapped normal. For the RCNF we scale the vector field as

$$f_\theta(z) = |G(z)|^{-1/2} \text{neural_net}(z) = \left(\frac{1 - \|z\|^2}{2} \right)^2 \text{neural_net}(z).$$

These three models are trained for 1500 iterations, by minimizing the negative log-likelihood (see Figure 4). When training, the divergence is approximated by the (Hutchinson) stochastic estimator from Equation 5.

G.2 Spherical geometry

Through the following spherical experiments we consider the two following models

- *Stereographic*: $P_\theta^S = (\rho^{-1} \circ \phi^{\mathbb{R}^2})_\# \mathcal{N}(0, 1)$
- *Riemannian*: $P_\theta^R = \phi_\#^{\mathbb{S}^2} \mathcal{U}(\mathbb{S}^2)$

with ρ^{-1} the inverse of the stereographic projection, $\phi^{\mathbb{R}^2}$ a conventional CNF on \mathbb{R}^2 , $\phi^{\mathbb{S}^2}$ our RCNF, $\mathcal{N}(0, 1)$ the standard Gaussian and $\mathcal{U}(\mathbb{S}^2)$ the uniform distribution on \mathbb{S}^2 . For the RCNF we project the output layer of the vector field as

$$f_\theta(z) = \text{proj}_{T_z \mathbb{S}^2} \text{neural_net}(z) = \frac{\text{neural_net}(z)}{\|\text{neural_net}(z)\|^2}$$

so as to enforce output vectors to be *tangent*. All spherical experiments were performed using the exact divergence estimator.

Limits of the stereographic projection model In this experiment the target is chosen to be a vMF(μ, κ) located at $\mu = -\mu_0$ with $-\{\mu_0\} = (-1, 0, \dots, 0)$. Both models are trained for 3000 iterations by minimizing the negative log-likelihood and the reverse KL divergence.

Density estimation of spherical data Finally we consider four earth location datasets, representing respectively volcano eruptions (NOAA, 2020b), earthquakes (NOAA, 2020a), floods (Brakenridge, 2017) and wild fires (EOSDIS, 2020). Concerning the CNFs, these models are trained by minimizing the negative log-likelihood for 1000 epochs. We observed that training models with a solver’s tolerance of $1e-5$ was computationally intensive so we lowered this tolerance to $1e-3$ during training, while keeping it to $1e-5$ for evaluation. We additionally, observed that annealing the learning rate such that $\alpha(t) = 0.98^{(t/300)} \alpha_0$ with $\alpha_0 = 1e^{-3}$ helped training convergence.

Concerning the mixture of von Mises-Fisher distributions, the parameters are learned by minimizing the negative log-likelihood with Riemannian Adam (Bécigneul and Ganea, 2019). The number of epochs is set to 10000 for all datasets but for the Volcano eruption one which requires 30000 for the vMF model to converge. The learning rate and number of components are selected by performing a hyperparameter grid search.

H Additional figures




Figure 7: Ablation study of the geodesic layer computational impact for the *Riemannian* model. Negative Log-likelihood of *Riemannian* CNFs trained to fit a $\mathcal{N}^W(\exp_0(\alpha \partial x), \Sigma)$ target on \mathbb{B}^2 .

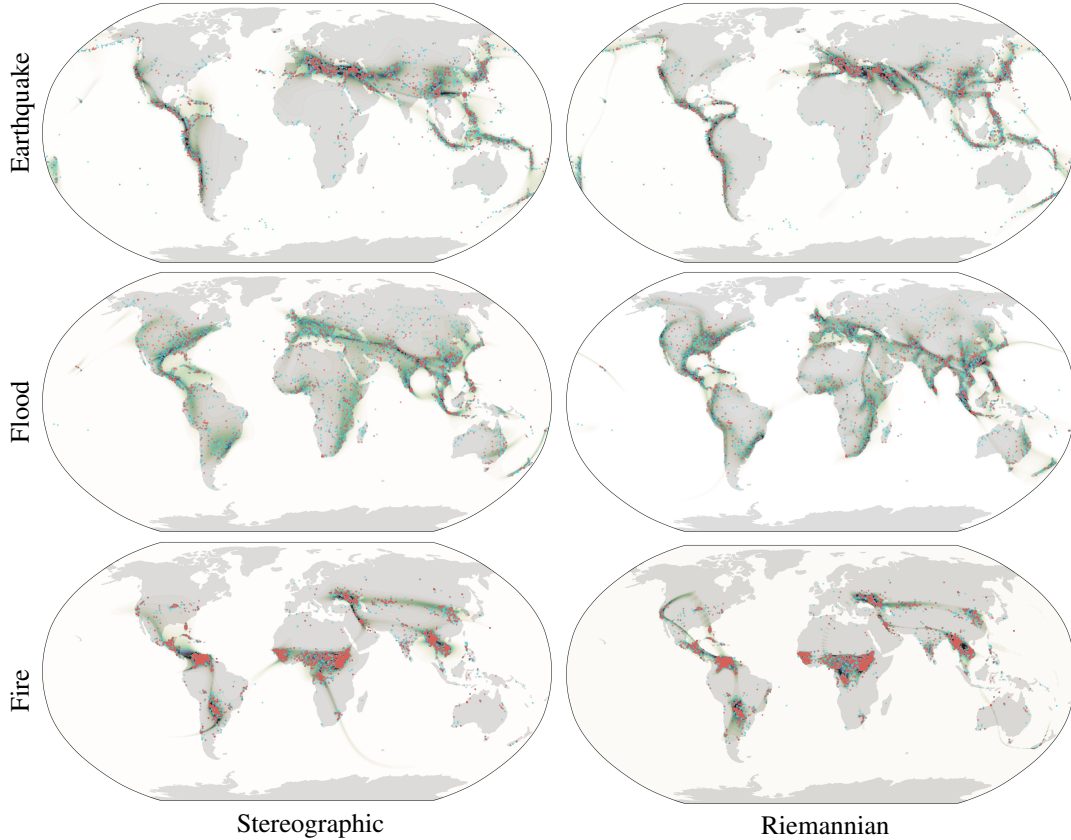


Figure 8: Density estimation for earth sciences data with Robinson projection. Blue and red dots represent training and testing datapoints, respectively. Heatmaps depict the log-likelihood of the trained models.