
Continuous Hierarchical Representations with Poincaré Variational Auto-Encoders

Emile Mathieu[†]

emile.mathieu@stats.ox.ac.uk

Charline Le Lan[†]

charline.lelan@stats.ox.ac.uk

Chris J. Maddison^{†*}

cmaddis@stats.ox.ac.uk

Ryota Tomioka[‡]

ryoto@microsoft.com

Yee Whye Teh^{†*}

y.w.teh@stats.ox.ac.uk

[†] Department of Statistics, University of Oxford, United Kingdom

* DeepMind, London, United Kingdom

[‡] Microsoft Research, Cambridge, United Kingdom

Abstract

The variational auto-encoder (VAE) is a popular method for learning a generative model and embeddings of the data. Many real datasets are hierarchically structured. However, traditional VAEs map data in a Euclidean latent space which cannot efficiently embed tree-like structures. Hyperbolic spaces with negative curvature can. We therefore endow VAEs with a Poincaré ball model of hyperbolic geometry as a latent space and rigorously derive the necessary methods to work with two main Gaussian generalisations on that space. We empirically show better generalisation to unseen data than the Euclidean counterpart, and can qualitatively and quantitatively better recover hierarchical structures.

1 Introduction

Learning useful representations from unlabelled raw sensory observations, which are often high-dimensional, is a problem of significant importance in machine learning. Variational auto-encoders (VAEs) (Kingma and Welling, 2014; Rezende et al., 2014) are a popular approach to this: they are probabilistic generative models composed of an *encoder* stochastically embedding observations in a low dimensional latent space \mathcal{Z} , and a *decoder* generating observations $x \in \mathcal{X}$ from encodings $z \in \mathcal{Z}$. After training, the encodings constitute a low-dimensional representation of the original raw observations, which can be used as features for a downstream task (e.g. Huang and LeCun, 2006; Coates et al., 2011) or be interpretable for their own sake. VAEs are therefore of interest for representation learning (Bengio et al., 2013), a field which aims to learn *good representations*, e.g. interpretable representations, ones yielding better generalisation, or ones useful for downstream tasks.

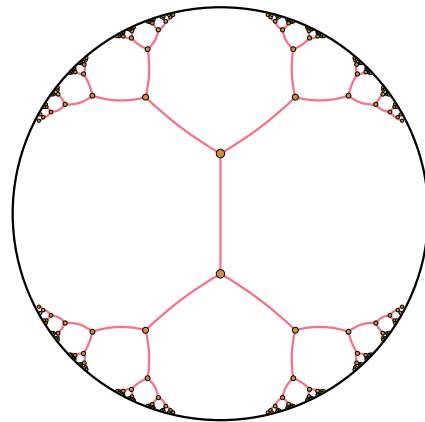


Figure 1: A regular tree isometrically embedded in the Poincaré disc. Red curves are same length *geodesics*, i.e. "straight lines".

It can be argued that in many domains data should be represented hierarchically. For example, in cognitive science, it is widely accepted that human beings use a hierarchy to organise object categories (e.g. Roy et al., 2006; Collins and Quillian, 1969; Keil, 1979). In biology, the theory of evolution (Darwin, 1859) implies that features of living organisms are related in a hierarchical manner given by the evolutionary tree. Explicitly incorporating hierarchical structure in probabilistic models has unsurprisingly been a long-running research topic (e.g. Duda et al., 2000; Heller and Ghahramani, 2005).

Earlier work in this direction tended to use trees as data structures to represent hierarchies. Recently, hyperbolic spaces have been proposed as an alternative continuous approach to learn hierarchical representations from textual and graph-structured data (Nickel and Kiela, 2017; Tifrea et al., 2019). Hyperbolic spaces can be thought of as continuous versions of trees, and vice versa, as illustrated in Figure 1. Trees can be embedded with arbitrarily low error into the Poincaré disc model of hyperbolic geometry (Sarkar, 2012). The exponential growth of the Poincaré surface area with respect to its radius is analogous to the exponential growth of the number of leaves in a tree with respect to its depth. Further, these spaces are smooth, enabling the use of deep learning approaches which rely on differentiability.

We show that replacing VAEs latent space components, which traditionally assume a Euclidean metric over the latent space, by their hyperbolic generalisation helps to represent and discover hierarchies. Our goals are twofold: (a) learn a latent representation that is interpretable in terms of hierarchical relationships among the observations, (b) learn a more efficient representation which generalises better to unseen data that is hierarchically structured. Our main contributions are as follows:

1. We propose efficient and reparametrisable sampling schemes, and calculate the probability density functions, for two canonical Gaussian generalisations defined on the Poincaré ball, namely the maximum-entropy and wrapped normal distributions. These are the ingredients required to train our VAEs.
2. We introduce a decoder architecture that explicitly takes into account the hyperbolic geometry, which we empirically show to be crucial.
3. We empirically demonstrate that endowing a VAE with a Poincaré ball latent space can be beneficial in terms of model generalisation and can yield more interpretable representations.

Our work fits well with a surge of interest in combining hyperbolic geometry and VAEs. Of these, it relates most strongly to the concurrent works of Ovinnikov (2018); Grattarola et al. (2019); Nagano et al. (2019). In contrast to these approaches, we introduce a decoder that takes into account the geometry of the hyperbolic latent space. Along with the *wrapped normal* generalisation used in the latter two articles, we give a thorough treatment of the *maximum entropy normal* generalisation and a rigorous analysis of the difference between the two. Additionally, we train our model by maximising a lower bound on the marginal likelihood, as opposed to Ovinnikov (2018); Grattarola et al. (2019) which consider a Wasserstein and an adversarial auto-encoder setting, respectively. We discuss these works in more detail in Section 4.

2 The Poincaré Ball model of hyperbolic geometry

2.1 Review of Riemannian geometry

Throughout the paper we denote the Euclidean norm and inner product by $\|\cdot\|$ and $\langle \cdot, \cdot \rangle$ respectively. 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 . For each point z of the manifold, the *metric tensor* $\mathfrak{g}(z)$ defines an inner product on the associated tangent space: $\mathfrak{g}(z) = \langle \cdot, \cdot \rangle_z : \mathcal{T}_z\mathcal{M} \times \mathcal{T}_z\mathcal{M} \rightarrow \mathbb{R}$. The *matrix representation of the Riemannian metric* $G(z)$, is defined such that $\forall \mathbf{u}, \mathbf{v} \in \mathcal{T}_z\mathcal{M} \times \mathcal{T}_z\mathcal{M}, \langle \mathbf{u}, \mathbf{v} \rangle_z = \mathfrak{g}(z)(\mathbf{u}, \mathbf{v}) = \mathbf{u}^T G(z) \mathbf{v}$. A *Riemannian manifold* is then defined as a tuple $(\mathcal{M}, \mathfrak{g})$ (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 $d\mathcal{M}(z) = \sqrt{|G(z)|} dz$ on the manifold, with dz being the Lebesgue measure.

The length of a curve $\gamma : t \mapsto \gamma(t) \in \mathcal{M}$ is given by $L(\gamma) = \int_0^1 \|\gamma'(t)\|_{\gamma(t)}^{1/2} dt$. The concept of straight lines can then be generalised to *geodesics*, which are constant speed curves giving the shortest path between pairs of points \mathbf{z}, \mathbf{y} of the manifold: $\gamma^* = \arg \min L(\gamma)$ with $\gamma(0) = \mathbf{z}$, $\gamma(1) = \mathbf{y}$ and $\|\gamma'(t)\|_{\gamma(t)} = 1$. A *global distance* is thus induced on \mathcal{M} given by $d_{\mathcal{M}}(\mathbf{z}, \mathbf{y}) = \inf L(\gamma)$. 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 $\gamma(0) = \mathbf{z}$ with initial tangent vector $\gamma'(0) = \mathbf{v}$. The corresponding exponential map is then defined by $\exp_{\mathbf{z}}(\mathbf{v}) = \gamma(1)$, as illustrated on Figure 2. The *logarithm map* is the inverse $\log_{\mathbf{z}} = \exp_{\mathbf{z}}^{-1} : \mathcal{M} \rightarrow \mathcal{T}_{\mathbf{z}}\mathcal{M}$. For geodesically complete manifolds, such as the Poincaré ball, $\exp_{\mathbf{z}}$ is well-defined on the full tangent space $\mathcal{T}_{\mathbf{z}}\mathcal{M}$ for all $\mathbf{z} \in \mathcal{M}$.

2.2 The Poincaré ball model of hyperbolic geometry

A d -dimensional hyperbolic space, denoted \mathbb{H}^d , is a complete, simply connected, d -dimensional Riemannian manifold with constant negative curvature c . In contrast with the Euclidean space \mathbb{R}^d , \mathbb{H}^d can be constructed using various isomorphic models (none of which is prevalent), including the hyperboloid model, the Beltrami-Klein model, the Poincaré half-plane model and the Poincaré ball \mathcal{B}_c^d (Beltrami, 1868). The Poincaré ball model is formally defined as the Riemannian manifold $\mathbb{B}_c^d = (\mathcal{B}_c^d, \mathfrak{g}_p^c)$, where \mathcal{B}_c^d is the open ball of radius $1/\sqrt{c}$, and \mathfrak{g}_p^c its *metric tensor*, which along with its induced distance are given by

$$\mathfrak{g}_p^c(\mathbf{z}) = (\lambda_{\mathbf{z}}^c)^2 \mathfrak{g}_e(\mathbf{z}), \quad d_p^c(\mathbf{z}, \mathbf{y}) = \frac{1}{\sqrt{c}} \cosh^{-1} \left(1 + 2c \frac{\|\mathbf{z} - \mathbf{y}\|^2}{(1 - c\|\mathbf{z}\|^2)(1 - c\|\mathbf{y}\|^2)} \right),$$

where $\lambda_{\mathbf{z}}^c = \frac{2}{1 - c\|\mathbf{z}\|^2}$ and \mathfrak{g}_e denotes the Euclidean metric tensor, i.e. the usual dot product. The *Möbius addition* (Ungar, 2008) of \mathbf{z} and \mathbf{y} in \mathbb{B}_c^d is defined as

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

One recovers the Euclidean addition of two vectors in \mathbb{R}^d as $c \rightarrow 0$. Building on that framework, Ganea et al. (2018) derived closed-form formulations for the *exponential map* (illustrated in Figure 2)

$$\exp_{\mathbf{z}}^c(\mathbf{v}) = \mathbf{z} \oplus_c \left(\tanh \left(\sqrt{c} \frac{\lambda_{\mathbf{z}}^c \|\mathbf{v}\|}{2} \right) \frac{\mathbf{v}}{\sqrt{c}\|\mathbf{v}\|} \right)$$

and its inverse, the *logarithm map*

$$\log_{\mathbf{z}}^c(\mathbf{y}) = \frac{2}{\sqrt{c}\lambda_{\mathbf{z}}^c} \tanh^{-1} \left(\sqrt{c} \|\mathbf{y} - \mathbf{z} \oplus_c \mathbf{y}\| \right) \frac{-\mathbf{z} \oplus_c \mathbf{y}}{\|\mathbf{y} - \mathbf{z} \oplus_c \mathbf{y}\|}.$$

3 The Poincaré VAE

We consider the problem of mapping an empirical distribution of observations to a lower dimensional Poincaré ball \mathbb{B}_c^d , as well as learning a map from this latent space $\mathcal{Z} = \mathbb{B}_c^d$ to the observation space \mathcal{X} . Building on the VAE framework, this *Poincaré-VAE* model, or \mathcal{P}^c -VAE for short, differs by the choice of prior and posterior distributions being defined on \mathbb{B}_c^d , and by the encoder g_{ϕ} and decoder f_{θ} maps which take into account the latent space geometry. Their parameters $\{\theta, \phi\}$ are learned by maximising the *evidence lower bound* (ELBO). Our model can be seen as a generalisation of a classical Euclidean VAE (Kingma and Welling, 2014; Rezende et al., 2014) that we denote by \mathcal{N} -VAE, i.e. \mathcal{P}^c -VAE $\xrightarrow{c \rightarrow 0} \mathcal{N}$ -VAE.

3.1 Prior and variational posterior distributions

In order to parametrise distributions on the Poincaré ball, we consider two canonical generalisations of normal distributions on that space. A more detailed review of Gaussian generalisations on manifolds can be found in Appendix B.1.

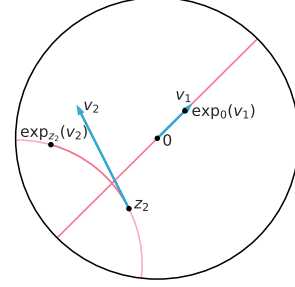


Figure 2: Geodesics and exponential maps in the Poincaré disc.

Riemannian normal One generalisation is the distribution maximising entropy given an expectation and variance (Said et al., 2014; Penneç, 2006; Hauberg, 2018), often called the *Riemannian normal* distribution, which has a density w.r.t. the metric induced measure $d\mathcal{M}$ given by

$$\mathcal{N}_{\mathbb{B}_c^d}^R(z|\boldsymbol{\mu}, \sigma^2) = \frac{d\nu^R(z|\boldsymbol{\mu}, \sigma^2)}{d\mathcal{M}(z)} = \frac{1}{Z^R} \exp\left(-\frac{d_p^c(\boldsymbol{\mu}, z)^2}{2\sigma^2}\right), \quad (1)$$

where $\sigma > 0$ is a dispersion parameter, $\boldsymbol{\mu} \in \mathbb{B}_c^d$ is the Fréchet mean, and Z^R is the normalising constant derived in Appendix B.4.3

Wrapped normal An alternative is to consider the push-forward measure obtained by mapping a normal distribution along the *exponential map* $\exp_{\boldsymbol{\mu}}$. That probability measure is often referred to as the *wrapped normal* distribution, and has been used in auto-encoder frameworks with other manifolds (Grattarola et al., 2019; Nagano et al., 2019; Falorsi et al., 2018). Samples $z \in \mathbb{B}_c^d$ are obtained as $z = \exp_{\boldsymbol{\mu}}^c(v/\lambda_{\boldsymbol{\mu}}^c)$ with $v \sim \mathcal{N}(\cdot|\mathbf{0}, \Sigma)$ and its density is given by (details given in Appendix B.3)

$$\begin{aligned} \mathcal{N}_{\mathbb{B}_c^d}^W(z|\boldsymbol{\mu}, \Sigma) &= \frac{d\nu^W(z|\boldsymbol{\mu}, \Sigma)}{d\mathcal{M}(z)} \\ &= \mathcal{N}(\lambda_{\boldsymbol{\mu}}^c \log_{\boldsymbol{\mu}}(z)|\mathbf{0}, \Sigma) \left(\frac{\sqrt{c} d_p^c(\boldsymbol{\mu}, z)}{\sinh(\sqrt{c} d_p^c(\boldsymbol{\mu}, z))} \right)^{d-1}. \end{aligned} \quad (2)$$

The (usual) normal distribution is recovered for both generalisations as $c \rightarrow 0$. We discuss the benefits and drawbacks of those two distributions in Appendix B.1. We refer to both as *hyperbolic normal* distributions with pdf $\mathcal{N}_{\mathbb{B}_c^d}(z|\boldsymbol{\mu}, \sigma^2)$. Figure 8 shows several probability densities for both distributions.

The prior distribution defined on \mathcal{Z} is chosen to be a hyperbolic normal distribution with mean zero, $p(z) = \mathcal{N}_{\mathbb{B}_c^d}(\cdot|\mathbf{0}, \sigma_0^2)$, and the variational family is chosen to be parametrised as $\mathcal{Q} = \{\mathcal{N}_{\mathbb{B}_c^d}(\cdot|\boldsymbol{\mu}, \sigma^2) \mid \boldsymbol{\mu} \in \mathbb{B}_c^d, \sigma \in \mathbb{R}_*^+\}$.

3.2 Encoder and decoder

We make use of two neural networks, a *decoder* f_{θ} and an *encoder* g_{ϕ} , to parametrise the likelihood $p(\cdot|f_{\theta}(z))$ and the variational posterior $q(\cdot|g_{\phi}(x))$ respectively. The input of f_{θ} and the output of g_{ϕ} need to respect the hyperbolic geometry of \mathcal{Z} . In the following we describe appropriate choices for the first layer of the decoder and the last layer of the encoder.

Decoder In the Euclidean case, an affine transformation can be written in the form $f_{\mathbf{a}, \mathbf{p}}(z) = \langle \mathbf{a}, z - \mathbf{p} \rangle$, with orientation and offset parameters $\mathbf{a}, \mathbf{p} \in \mathbb{R}^d$. This can be rewritten in the form

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

where $H_{\mathbf{a}, \mathbf{p}} = \{z \in \mathbb{R}^p \mid \langle \mathbf{a}, z - \mathbf{p} \rangle = 0\} = \mathbf{p} + \{\mathbf{a}\}^{\perp}$ is the decision hyperplane. The third term is the distance between z and the decision hyperplane $H_{\mathbf{a}, \mathbf{p}}^c$ and the first term refers to the side of $H_{\mathbf{a}, \mathbf{p}}^c$ where z lies. Ganea et al. (2018) analogously introduced an operator $f_{\mathbf{a}, \mathbf{p}}^c: \mathbb{B}_c^d \rightarrow \mathbb{R}^p$ on the Poincaré ball,

$$f_{\mathbf{a}, \mathbf{p}}^c(z) = \text{sign}(\langle \mathbf{a}, \log_{\mathbf{p}}^c(z) \rangle_{\mathbf{p}}) \|\mathbf{a}\|_{\mathbf{p}} d_p^c(z, H_{\mathbf{a}, \mathbf{p}}^c)$$

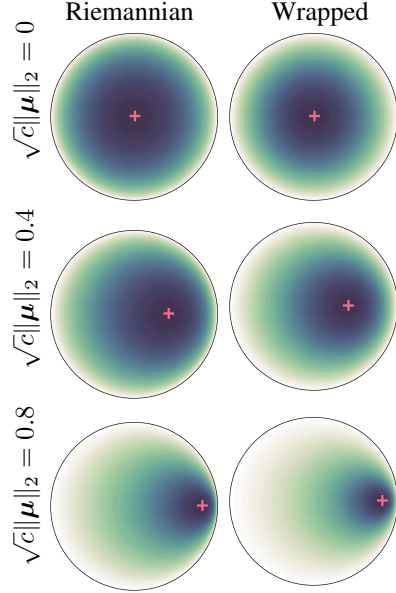


Figure 3: Hyperbolic normal probability density for different Fréchet mean, same standard deviation and $c = 10$. The *Riemannian* hyperbolic radius has a slightly larger mode.

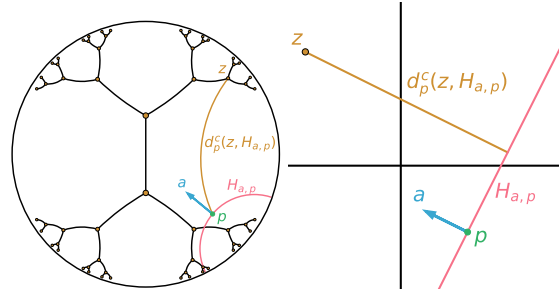


Figure 4: Illustration of an orthogonal projection on a hyperplane in a Poincaré disc (Left) and an Euclidean plane (Right).

with $H_{\mathbf{a},\mathbf{p}}^c = \{\mathbf{z} \in \mathbb{B}_c^d \mid \langle \mathbf{a}, \log_p^c(\mathbf{z}) \rangle = 0\} = \exp_p^c(\{\mathbf{a}\}^\perp)$. A closed-form expression for the distance $d_p^c(\mathbf{z}, H_{\mathbf{a},\mathbf{p}}^c)$ was also derived, $d_p^c(\mathbf{z}, H_{\mathbf{a},\mathbf{p}}^c) = \frac{1}{\sqrt{c}} \sinh^{-1} \left(\frac{2\sqrt{c} | \langle -\mathbf{p} \oplus_c \mathbf{z}, \mathbf{a} \rangle |}{(1-c\|\mathbf{p} \oplus_c \mathbf{z}\|^2)\|\mathbf{a}\|} \right)$. The hyperplane decision boundary $H_{\mathbf{a},\mathbf{p}}^c$ is called *gyroplane* and is a semi-hypersphere orthogonal to the Poincaré ball’s boundary as illustrated on Figure 4. The decoder’s first layer, called *gyroplane* layer, is chosen to be a concatenation of such operators, which are then composed with a standard feed-forward neural network.

Encoder The encoder g_ϕ outputs a Fréchet mean $\boldsymbol{\mu} \in \mathbb{B}_c^d$ and a distortion $\sigma \in \mathbb{R}_*^+$ which parametrise the hyperbolic variational posterior. The Fréchet mean $\boldsymbol{\mu}$ is obtained as the image of the exponential map \exp_0^c , and the distortion σ through a *softplus* function.

3.3 Training

We follow a standard variational approach by deriving a lower bound on the marginal likelihood. The ELBO is optimised via an unbiased Monte Carlo (MC) estimator thanks to the reparametrisable sampling schemes that we introduce for both hyperbolic normal distributions.

Objective The *evidence lower bound* (ELBO) can readily be extended to Riemannian latent spaces by applying Jensen’s inequality w.r.t. $d\mathcal{M}$ (see Appendix A)

$$\log p(\mathbf{x}) \geq \mathcal{L}_{\mathcal{M}}(\mathbf{x}; \theta, \phi) \triangleq \int_{\mathcal{M}} \ln \left(\frac{p_\theta(\mathbf{x}|\mathbf{z})p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x})} \right) q_\phi(\mathbf{z}|\mathbf{x}) d\mathcal{M}(\mathbf{z}).$$

Densities have been introduced earlier in Equations 1 and 2.

Reparametrisation In the Euclidean setting, by working in polar coordinates, an isotropic normal distribution centred at $\boldsymbol{\mu}$ can be described by a directional vector $\boldsymbol{\alpha}$ uniformly distributed on the hypersphere and a univariate radius $r = d_E(\boldsymbol{\mu}, \mathbf{z})$ following a χ -distribution. In the Poincaré ball we can rely on a similar representation, through a *hyperbolic polar* change of coordinates, given by

$$\mathbf{z} = \exp_\mu^c \left(G(\boldsymbol{\mu})^{-\frac{1}{2}} \mathbf{v} \right) = \exp_\mu^c \left(\frac{r}{\lambda_\mu^c} \boldsymbol{\alpha} \right) \quad (3)$$

with $\mathbf{v} = r\boldsymbol{\alpha}$ and $r = d_p^c(\boldsymbol{\mu}, \mathbf{z})$. The direction $\boldsymbol{\alpha}$ is still uniformly distributed on the hypersphere and for the *wrapped normal*, the radius r is still χ -distributed, while for the *Riemannian normal* its density $\rho^R(r)$ is given by (derived in Appendix B.4.1)

$$\rho^W(r) \propto \mathbb{1}_{\mathbb{R}_+}(r) e^{-\frac{r^2}{2\sigma^2}} r^{d-1}, \quad \rho^R(r) \propto \mathbb{1}_{\mathbb{R}_+}(r) e^{-\frac{r^2}{2\sigma^2}} \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}} \right)^{d-1}.$$

The latter density $\rho^R(r)$ can efficiently be sampled via rejection sampling with a piecewise exponential distribution proposal. This makes use of its log-concavity. The *Riemannian normal* sampling scheme is not directly affected by dimensionality since the radius is a *one-dimensional* random variable. Full sampling schemes are described in Algorithm 1, and in Appendices B.4.1 and B.4.2.

Gradients Gradients $\nabla_\mu \mathbf{z}$ can straightforwardly be computed thanks to the exponential map reparametrisation (Eq 3), and gradients w.r.t. the dispersion $\nabla_\sigma \mathbf{z}$ are readily available for the *wrapped normal*. For the *Riemannian normal*, we additionally rely on an implicit reparametrisation (Figurnov et al., 2018) of ρ^R via its cdf $F^R(r; \sigma)$.

Optimisation Parameters of the model living in the Poincaré ball are parametrised via the exponential mapping: $\phi_i = \exp_0^c(\phi_i^0)$ with $\phi_i^0 \in \mathbb{R}^m$, so we can make use of usual optimisation schemes. Alternatively, one could directly optimise such manifold parameters with manifold gradient descent schemes (Bonnabel, 2013).

Algorithm 1 Hyperbolic normal sampling scheme

Require: $\boldsymbol{\mu}, \sigma^2$, dimension d , curvature c
if Wrapped normal **then** $\mathbf{v} \sim \mathcal{N}(\mathbf{0}_d, \sigma^2)$
else if Riemannian normal **then**
 Let g be a piecewise exponential proposal
 while sample r not accepted **do**
 Propose $r \sim g(\cdot), u \sim \mathcal{U}([0, 1])$
 if $u < \frac{\rho^R(r)}{g(r)}$ **then** Accept sample r
 Sample direction $\boldsymbol{\alpha} \sim \mathcal{U}(\mathbb{S}^{d-1})$
 $\mathbf{v} \leftarrow r\boldsymbol{\alpha}$
 Return $\mathbf{z} = \exp_\mu^c(\mathbf{v}/\lambda_\mu^c)$

4 Related work

Hierarchical models The Bayesian Nonparametric literature has a rich history of explicitly modelling the hierarchical structure of data (Teh et al., 2008; Heller and Ghahramani, 2005; Griffiths et al., 2004; Ghahramani et al., 2010; Larsen et al., 2001; Salakhutdinov et al., 2011). The discrete nature of trees used in such models makes learning difficult, whereas performing optimisation in a continuous hyperbolic space is an attractive alternative. Such an approach has been empirically and theoretically shown to be useful for graphs and word embeddings (Nickel and Kiela, 2017, 2018; Chamberlain et al., 2017; Sala et al., 2018; Tifrea et al., 2019).

Distributions on manifold Probability measures defined on manifolds are of interest to model uncertainty of data living (either intrinsically or assumed to) on such spaces, e.g. directional statistics (Ley and Verdebout, 2017; Mardia and Jupp, 2009). Pennec (2006) introduced a maximum entropy generalisation of the normal distribution, often referred to as *Riemannian normal*, which has been used for maximum likelihood estimation in the Poincaré half-plane (Said et al., 2014) and on the hypersphere (Hauberg, 2018). Another class of manifold probability measures are *wrapped* distributions, i.e. push-forward of distributions defined on a tangent space, often along the exponential map. They have recently been used in auto-encoder frameworks on the *hyperboloid* model (of hyperbolic geometry) (Grattarola et al., 2019; Nagano et al., 2019) and on Lie groups (Falorsi et al., 2018). Rey et al. (2019); Li et al. (2019) proposed to parametrise a variational family through a Brownian motion on manifolds such as spheres, tori, projective spaces and $SO(3)$.

VAEs with Riemannian latent manifold VAEs with non Euclidean latent space have been recently introduced, such as Davidson et al. (2018) making use of hyperspherical geometry and Falorsi et al. (2018) endowing the latent space with a $SO(3)$ group structure. Concurrent work considers endowing auto-encoders (AEs) with a hyperbolic latent space. Grattarola et al. (2019) introduces a constant curvature manifold (CCM) (i.e. hyperspherical, Euclidean and hyperboloid) latent space within an adversarial auto-encoder framework. However, the encoder and decoder are not designed to explicitly take into account the latent space geometry. Ovinnikov (2018) recently proposed to endow a VAE latent space with a Poincaré ball model. They choose a Wasserstein Auto-Encoder framework (Tolstikhin et al., 2018) because they could not derive a closed-form solution of the ELBO’s entropy term. We instead rely on a MC estimate of the ELBO by introducing a novel reparametrisation of the Riemannian normal. They discuss the *Riemannian* normal distribution, yet they make a number of heuristic approximations for sampling and reparametrisation. Also, Nagano et al. (2019) propose using a *wrapped* normal distribution to model uncertainty on the *hyperboloid* model of hyperbolic space. They derive its density and a reparametrisable sampling scheme, allowing such a distribution to be used in a variational learning framework. They apply this *wrapped* normal distribution to stochastically embed graphs and to parametrise the variational family in VAEs. Ovinnikov (2018) and Nagano et al. (2019) rely on a standard feed-forward decoder architecture, which does not take into account the hyperbolic geometry.

5 Experiments

We implemented our model and ran our experiments within the automatic differentiation framework PyTorch (Paszke et al., 2017). We open-source our code for reproducibility and to benefit the community ¹. Experimental details are fully described in Appendix C.

5.1 Branching diffusion process

We assess our modelling assumption on data generated from a branching diffusion process which explicitly incorporate hierarchical structure. Nodes $y_i \in \mathbb{R}^n$ are normally distributed with mean given by their parent and with unit variance. Models are trained on a noisy vector representations (x_1, \dots, x_N) , hence do not have access to the true hierarchical representation. We train several \mathcal{P}^c -VAEs with increasing curvatures, along with a vanilla \mathcal{N} -VAE as a baseline. Table 1 shows that the \mathcal{P}^c -VAE outperforms its Euclidean counterpart in terms of test marginal likelihood. As expected, we observe that the performance of the \mathcal{N} -VAE is recovered as the curvature c tends to zero. Also, we notice that increasing the prior distribution distortion σ_0 helps embeddings lie closer to the border,

¹<https://github.com/emilemathieu/pvae>

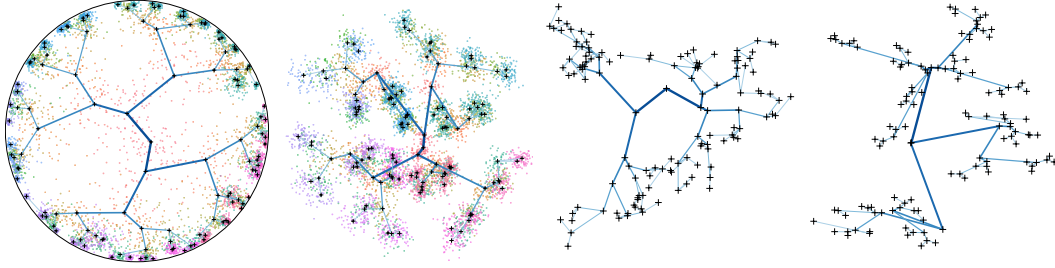


Figure 5: Latent representations learned by \mathcal{P}^1 -VAE (Leftmost), \mathcal{N} -VAE (Center-Left), PCA (Center-Right) and GPLVM (Rightmost) trained on synthetic dataset. Embeddings are represented by black crosses, and colour dots are posterior samples. Blue lines represent true hierarchy.

and as a consequence improved generalisation performance. Figure 5 represents latent embeddings for \mathcal{P}^1 -VAE and \mathcal{N} -VAE, along with two embedding baselines: principal component analysis (PCA) and a Gaussian process latent variable model (GPLVM). A hierarchical structure is somewhat learned by all models, yet \mathcal{P}^c -VAE’s latent representation is the least distorted.

Table 1: Negative test marginal likelihood estimates $\mathcal{L}_{\text{IWAE}}$ (Burda et al., 2015) (computed with 5000 samples) on the synthetic dataset. 95% confidence intervals are computed over 20 trainings.

	σ_0	Models					
		\mathcal{N} -VAE	$\mathcal{P}^{0.1}$ -VAE	$\mathcal{P}^{0.3}$ -VAE	$\mathcal{P}^{0.8}$ -VAE	$\mathcal{P}^{1.0}$ -VAE	$\mathcal{P}^{1.2}$ -VAE
$\mathcal{L}_{\text{IWAE}}$	1	57.1 \pm 0.2	57.1 \pm 0.2	57.2 \pm 0.2	56.9 \pm 0.2	56.7 \pm 0.2	56.6 \pm 0.2
$\mathcal{L}_{\text{IWAE}}$	1.7	57.0 \pm 0.2	56.8 \pm 0.2	56.6 \pm 0.2	55.9 \pm 0.2	55.7 \pm 0.2	55.6 \pm 0.2

5.2 Mnist digits

The MNIST (LeCun and Cortes, 2010) dataset has been used in the literature for hierarchical modelling (Salakhutdinov et al., 2011; Saha et al., 2018). One can view the natural clustering in MNIST images as a hierarchy with each of the 10 classes being internal nodes of the hierarchy. We empirically assess whether our model can take advantage of such simple underlying hierarchical structure, first by measuring its generalisation capacity via the test marginal log-likelihood. Table 2 shows that our model outperforms its Euclidean counterpart, especially for low latent dimension. This can be interpreted through an information bottleneck perspective; as the latent dimensionality increases, the pressure on the embeddings quality decreases, hence the gain from the hyperbolic geometry is reduced (as observed by Nickel and Kiela (2017)). Also, by using the *Riemannian normal* distribution, we achieve slightly better results than with the *wrapped normal*.

Table 2: Negative test marginal likelihood estimates computed with 5000 samples. 95% confidence intervals are computed over 10 runs. * indicates numerically unstable settings.

	c	Dimensionality			
		2	5	10	20
\mathcal{N} -VAE	(0)	144.5 \pm 0.4	114.7 \pm 0.1	100.2 \pm 0.1	97.6 \pm 0.1
\mathcal{P} -VAE (Wrapped)	0.1	143.9 \pm 0.5	115.5 \pm 0.3	100.2 \pm 0.1	97.2 \pm 0.1
	0.2	144.2 \pm 0.5	115.3 \pm 0.3	100.0 \pm 0.1	97.1 \pm 0.1
	0.7	143.8 \pm 0.6	115.1 \pm 0.3	100.2 \pm 0.1	97.5 \pm 0.1
	1.4	144.0 \pm 0.6	114.7 \pm 0.1	100.7 \pm 0.1	98.0 \pm 0.1
\mathcal{P} -VAE (Riemannian)	0.1	143.7 \pm 0.6	115.2 \pm 0.2	99.9 \pm 0.1	97.0 \pm 0.1
	0.2	143.8 \pm 0.4	114.7 \pm 0.3	99.7 \pm 0.1	97.4 \pm 0.1
	0.7	143.1 \pm 0.4	114.1 \pm 0.2	101.2 \pm 0.2	*
	1.4	142.5 \pm 0.4	115.5 \pm 0.3	*	*

We conduct an ablation study to assess the usefulness of the *gyroplane* layer introduced in Section 3.2. To do so we estimate the test marginal log-likelihood for different choices of decoder. We select a multi-layer perceptron (MLP) to be the baseline decoder. We additionally compare to a MLP pre-composed by \log_0 , which can be seen as a linearisation of the space around the centre of the ball. Figure 6 shows the relative performance improvement of decoders over the MLP baseline w.r.t. the latent space dimension. We observe that *linearising* the input of a MLP through the logarithm map slightly improves generalisation, and that using a *gyroplane* layer as the first layer of the decoder additionally improves generalisation. Yet, these performance gains appear to decrease as the latent dimensionality increases.

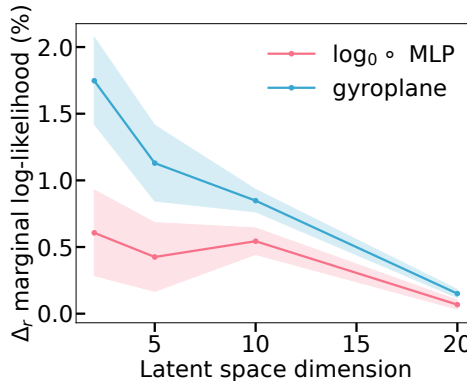


Figure 6: Decoder ablation study on MNIST with wrapped normal \mathcal{P}^1 -VAE. Baseline decoder is a MLP.

Second, we explore the learned latent representations of the trained \mathcal{P} -VAE and \mathcal{N} -VAE models shown in Figure 7. Qualitatively our \mathcal{P} -VAE produces a clearer partitioning of the digits, in groupings of $\{4, 7, 9\}$, $\{0, 6\}$, $\{2, 3, 5, 8\}$ and $\{1\}$, with right-slanting $\{5, 8\}$ being placed separately from the non-slanting ones. Recall that distances increase towards the edge of the Poincaré ball. We quantitatively assess the quality of the embeddings by training a classifier predicting labels. Table 3 shows that the embeddings learned by our \mathcal{P} -VAE model yield on average an 2% increase in accuracy over the digits. The full confusion matrices are shown in Figure 12 in Appendix.

Table 3: Per digit accuracy of a classifier trained on the learned latent 2-d embeddings. Results are averaged over 10 sets of embeddings and 5 classifier trainings.

Digits	0	1	2	3	4	5	6	7	8	9	Avg
\mathcal{N} -VAE	89	97	81	75	59	43	89	78	68	57	73.6
$\mathcal{P}^{1.4}$ -VAE	94	97	82	79	69	47	90	77	68	53	75.6




Figure 7: MNIST Posteriors mean (Left) sub-sample of digit images associated with posteriors mean (Middle) Model samples (Right) – for $\mathcal{P}^{1.4}$ -VAE (Top) and \mathcal{N} -VAE (Bottom).

5.3 Graph embeddings

We evaluate the performance of a variational graph auto-encoder (VGAE) (Kipf and Welling, 2016) with Poincaré ball latent space for link prediction in networks. Edges in complex networks can typically be explained by a latent hierarchy over the nodes (Clauset et al., 2008). We believe the Poincaré ball latent space should help in terms of generalisation. We demonstrate these capabilities on three network datasets: a graph of Ph.D. advisor-advisee relationships (Nooy et al., 2011), a phylogenetic tree expressing genetic heritage (Hofbauer et al., 2016; Sanderson and Eriksson, 1994) and a biological set representing disease relationships (Goh et al., 2007; Rossi and Ahmed, 2015).

We follow the VGAE model, which maps the adjacency matrix A to node embeddings Z through a graph convolutional network (GCN), and reconstructs A by predicting edge probabilities from the node embeddings. In order to take into account the latent space geometry, we parametrise the probability of an edge by $p(A_{ij} = 1|z_i, z_j) = 1 - \tanh(d_{\mathcal{M}}(z_i, z_j)) \in (0, 1]$ with $d_{\mathcal{M}}$ the latent geodesic metric. We use a *Wrapped* Gaussian prior and variational posterior for the \mathcal{P}^1 -VAE.

We set the latent dimension to 5. We follow the training and evaluation procedures introduced in Kipf and Welling (2016). Models are trained on an incomplete adjacency matrix where some of the edges have randomly been removed. A test set is formed from previously removed edges and an equal number of randomly sampled pairs of unconnected nodes. We report in Table 4 the *area under the ROC curve* (AUC) and *average precision* (AP) evaluated on the test set. It can be observed that the \mathcal{P} -VAE performs better than its Euclidean counterpart in terms of generalisation to unseen edges.

Table 4: Results on network link prediction. 95% confidence intervals are computed over 40 runs.

	Phylogenetic		CS PhDs		Diseases	
	AUC	AP	AUC	AP	AUC	AP
\mathcal{N} -VAE	54.2 \pm 2.2	54.0 \pm 2.1	56.5 \pm 1.1	56.4 \pm 1.1	89.8 \pm 0.7	91.8 \pm 0.7
\mathcal{P} -VAE	59.0\pm1.9	55.5 \pm 1.6	59.8\pm1.2	56.7 \pm 1.2	92.3\pm0.7	93.6\pm0.5

6 Conclusion

In this paper we have explored VAEs with a Poincaré ball latent space. We gave a thorough treatment of two canonical – *wrapped* and *maximum entropy* – normal generalisations on that space, and a rigorous analysis of the difference between the two. We derived the necessary ingredients for training such VAEs, namely efficient and reparametrisable sampling schemes, along with probability density functions for these two distributions. We introduced a decoder architecture explicitly taking into account the hyperbolic geometry, and empirically showed that it is crucial for the hyperbolic latent space to be useful. We empirically demonstrated that endowing a VAE with a Poincaré ball latent space can be beneficial in terms of model generalisation and can yield more interpretable representations if the data has hierarchical structure.

There are a number of interesting future directions. There are many models of hyperbolic geometry, and several have been considered in a gradient-based setting. Yet, it is still unclear which models should be preferred and which of their properties matter. Also, it would be useful to consider principled ways of assessing whether a given dataset has an underlying hierarchical structure, in the same way that topological data analysis (Pascucci et al., 2011) attempts to discover the topologies that underlie datasets.

Acknowledgments We are extremely grateful to Adam Foster, Phillippe Gagnon and Emmanuel Chevallier for their help. EM, YWT’s 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 they acknowledge Microsoft Research and EPSRC for funding EM’s studentship, and EPSRC grant agreement no. EP/N509711/1 for funding CL’s studentship.

References

- Beltrami, E. (1868). *Teoria fondamentale degli spazii di curvatura costante: memoria*. F. Zanetti.
- Bengio, Y., Courville, A., and Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Trans. Pattern Anal. Mach. Intell.*, 35(8):1798–1828.
- Bonnabel, S. (2013). Stochastic gradient descent on riemannian manifolds. *IEEE Transactions on Automatic Control*, 58(9):2217–2229.
- Box, G. E. P. and Muller, M. E. (1958). A note on the generation of random normal deviates. *Ann. Math. Statist.*, 29(2):610–611.
- Burda, Y., Grosse, R., and Salakhutdinov, R. (2015). Importance Weighted Autoencoders. *arXiv.org*.
- Chamberlain, B. P., Clough, J., and Deisenroth, M. P. (2017). Neural Embeddings of Graphs in Hyperbolic Space. *arXiv.org*.
- Chevallier, E., Barbaresco, F., and Angulo, J. (2015). Probability density estimation on the hyperbolic space applied to radar processing. In Nielsen, F. and Barbaresco, F., editors, *Geometric Science of Information*, pages 753–761, Cham. Springer International Publishing.
- Clauset, A., Moore, C., and Newman, M. E. J. (2008). Hierarchical structure and the prediction of missing links in networks. *Nature*, 453:98–101.
- Coates, A., Lee, H., and Ng, A. (2011). An analysis of single-layer networks in unsupervised feature learning. In Gordon, G., Dunson, D., and Dudík, M., editors, *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*, volume 15 of *JMLR Workshop and Conference Proceedings*, pages 215–223. JMLR W&CP.
- Collins, A. and Quillian, M. (1969). Retrieval time from semantic memory. *Journal of Verbal Learning and Verbal Behavior*, 8:240–248.
- Darwin, C. (1859). *On the Origin of Species by Means of Natural Selection*. Murray, London. or the Preservation of Favored Races in the Struggle for Life.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. (2018). Hyperspherical variational auto-encoders. *34th Conference on Uncertainty in Artificial Intelligence (UAI-18)*.
- Duda, R. O., Hart, P. E., and Stork, D. G. (2000). *Pattern Classification (2Nd Edition)*. Wiley-Interscience, New York, NY, USA.
- Falorsi, L., de Haan, P., Davidson, T. R., De Cao, N., Weiler, M., Forré, P., and Cohen, T. S. (2018). Explorations in Homeomorphic Variational Auto-Encoding. *arXiv e-prints*, page arXiv:1807.04689.
- Figurnov, M., Mohamed, S., and Mnih, A. (2018). Implicit reparameterization gradients. In *International Conference on Neural Information Processing Systems*, pages 439–450.
- Ganea, O.-E., Bécigneul, G., and Hofmann, T. (2018). Hyperbolic neural networks. In *International Conference on Neural Information Processing Systems*, pages 5350–5360.
- Ghahramani, Z., Jordan, M. I., and Adams, R. P. (2010). Tree-structured stick breaking for hierarchical data. In Lafferty, J. D., Williams, C. K. I., Shawe-Taylor, J., Zemel, R. S., and Culotta, A., editors, *Advances in Neural Information Processing Systems 23*, pages 19–27. Curran Associates, Inc.
- Goh, K.-I., Cusick, M. E., Valle, D., Childs, B., Vidal, M., and Barabási, A.-L. (2007). The human disease network. *Proceedings of the National Academy of Sciences*, 104(21):8685–8690.

- Grattarola, D., Livi, L., and Alippi, C. (2019). Adversarial autoencoders with constant-curvature latent manifolds. *Appl. Soft Comput.*, 81.
- Griffiths, T. L., Jordan, M. I., Tenenbaum, J. B., and Blei, D. M. (2004). Hierarchical topic models and the nested chinese restaurant process. In Thrun, S., Saul, L. K., and Schölkopf, B., editors, *Advances in Neural Information Processing Systems 16*, pages 17–24. MIT Press.
- Hauberg, S. (2018). Directional statistics with the spherical normal distribution. In *Proceedings of 2018 21st International Conference on Information Fusion, FUSION 2018*, pages 704–711. IEEE.
- Heller, K. A. and Ghahramani, Z. (2005). Bayesian hierarchical clustering. In *International Conference on Machine Learning, ICML '05*, pages 297–304, New York, NY, USA. ACM.
- Hofbauer, W., Forrest, L., M. Hollingsworth, P., and Hart, M. (2016). Preliminary insights from dna barcoding into the diversity of mosses colonising modern building surfaces. *Bryophyte Diversity and Evolution*, 38:1.
- Hsu, E. P. (2008). A brief introduction to brownian motion on a riemannian manifold.
- Huang, F. J. and LeCun, Y. (2006). Large-scale learning with SVM and convolutional for generic object categorization. In *CVPR (1)*, pages 284–291. IEEE Computer Society.
- Keil, F. (1979). *Semantic and Conceptual Development: An Ontological Perspective*. Cognitive science series. Harvard University Press.
- Kingma, D. P. and Ba, J. (2016). Adam: A Method for Stochastic Optimization. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kingma, D. P. and Welling, M. (2014). Auto-encoding variational bayes. In *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Kipf, T. N. and Welling, M. (2016). Variational graph auto-encoders. *Workshop on Bayesian Deep Learning, NIPS*.
- Larsen, J., Szymkowiak, A., and Hansen, L. K. (2001). Probabilistic hierarchical clustering with labeled and unlabeled data.
- LeCun, Y. and Cortes, C. (2010). MNIST handwritten digit database.
- Ley, C. and Verdebout, T. (2017). *Modern Directional Statistics*. New York: Chapman and Hall/CRC.
- Li, H., Lindenbaum, O., Cheng, X., and Cloninger, A. (2019). Variational Random Walk Autoencoders. *arXiv.org*.
- Mardia, K. and Jupp, P. (2009). *Directional Statistics*. Wiley Series in Probability and Statistics. Wiley.
- Nagano, Y., Yamaguchi, S., Fujita, Y., and Koyama, M. (2019). A Differentiable Gaussian-like Distribution on Hyperbolic Space for Gradient-Based Learning. In *International Conference on Machine Learning (ICML)*.
- Nickel, M. and Kiela, D. (2017). Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems*, pages 6341–6350.
- Nickel, M. and Kiela, D. (2018). Learning Continuous Hierarchies in the Lorentz Model of Hyperbolic Geometry. In *International Conference on Machine Learning (ICML)*.
- Nooy, W. D., Mrvar, A., and Batagelj, V. (2011). *Exploratory Social Network Analysis with Pajek*. Cambridge University Press, New York, NY, USA.
- Ovinnikov, I. (2018). Poincaré Wasserstein Autoencoder. *NeurIPS Workshop on Bayesian Deep Learning*, pages 1–8.
- Paeng, S.-H. (2011). Brownian motion on manifolds with time-dependent metrics and stochastic completeness. *Journal of Geometry and Physics*, 61(5):940 – 946.

- Pascucci, V., Tricoche, X., Hagen, H., and Tierny, J. (2011). *Topological Methods in Data Analysis and Visualization: Theory, Algorithms, and Applications*. Springer Publishing Company, Incorporated, 1st edition.
- 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.
- Petersen, P. (2006). *Riemannian Geometry*. Springer-Verlag New York.
- R. Gilks, W. and Wild, P. (1992). Adaptive rejection sampling for gibbs sampling. 41:337–348.
- Rey, L. A. P., Menkovski, V., and Portegies, J. W. (2019). Diffusion Variational Autoencoders. *CoRR*.
- Rezende, D. J., Mohamed, S., and Wierstra, D. (2014). Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *International Conference on Machine Learning (ICML)*.
- Rossi, R. A. and Ahmed, N. K. (2015). The network data repository with interactive graph analytics and visualization. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, AAAI’15, pages 4292–4293. AAAI Press.
- Roy, D. M., Kemp, C., Mansinghka, V. K., and Tenenbaum, J. B. (2006). Learning annotated hierarchies from relational data. In *NIPS*, pages 1185–1192. MIT Press.
- Saha, S., Varma, G., and Jawahar, C. V. (2018). Class2str: End to end latent hierarchy learning. *International Conference on Pattern Recognition (ICPR)*, pages 1000–1005.
- Said, S., Bombrun, L., and Berthoumieu, Y. (2014). New riemannian priors on the univariate normal model. *Entropy*, 16(7):4015–4031.
- Sala, F., De Sa, C., Gu, A., and Re, C. (2018). Representation tradeoffs for hyperbolic embeddings. In Dy, J. and Krause, A., editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 4460–4469, Stockholmsmässan, Stockholm Sweden. PMLR.
- Salakhutdinov, R. R., Tenenbaum, J. B., and Torralba, A. (2011). One-shot learning with a hierarchical nonparametric bayesian model. In *ICML Unsupervised and Transfer Learning*.
- Sanderson, M. J., M. J. D. W. P. and Eriksson, T. (1994). Treebase: a prototype database of phylogenetic analyses and an interactive tool for browsing the phylogeny of life. *American Journal of Botany*.
- Sarkar, R. (2012). Low distortion delaunay embedding of trees in hyperbolic plane. In van Kreveld, M. and Speckmann, B., editors, *Graph Drawing*, pages 355–366, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Teh, Y. W., Daume III, H., and Roy, D. M. (2008). Bayesian agglomerative clustering with coalescents. In Platt, J. C., Koller, D., Singer, Y., and Roweis, S. T., editors, *Advances in Neural Information Processing Systems 20*, pages 1473–1480. Curran Associates, Inc.
- Tifrea, A., Becigneul, G., and Ganea, O.-E. (2019). Poincare glove: Hyperbolic word embeddings. In *International Conference on Learning Representations (ICLR)*.
- Tolstikhin, I., Bousquet, O., Gelly, S., and Schoelkopf, B. (2018). Wasserstein auto-encoders. In *International Conference on Learning Representations*.
- Ungar, A. A. (2008). A gyrovector space approach to hyperbolic geometry. *Synthesis Lectures on Mathematics and Statistics*, 1(1):1–194.