

## A Evidence Lower Bound

The ELBO can readily be extended for Riemannian latent spaces by applying Jensen’s inequality w.r.t. the metric induced measure  $d\mathcal{M}$  which yield

$$\begin{aligned}
\ln p(\mathbf{x}) &= \ln \int_{\mathcal{Z}=\mathcal{M}} p_\theta(\mathbf{x}, \mathbf{z}) d\mathcal{M}(\mathbf{z}) = \ln \int_{\mathcal{M}} p_\theta(\mathbf{x}|\mathbf{z}) p(\mathbf{z}) d\mathcal{M}(\mathbf{z}) \\
&= \ln \int_{\mathcal{M}} \frac{p_\theta(\mathbf{x}|\mathbf{z}) p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x})} q_\phi(\mathbf{z}|\mathbf{x}) d\mathcal{M}(\mathbf{z}) \\
&\geq \int_{\mathcal{M}} \ln \frac{p_\theta(\mathbf{x}|\mathbf{z}) p(\mathbf{z})}{q_\phi(\mathbf{z}|\mathbf{x})} q_\phi(\mathbf{z}|\mathbf{x}) d\mathcal{M}(\mathbf{z}) \\
&= \int_{\mathcal{M}} [\ln p_\theta(\mathbf{x}|\mathbf{z}) - \ln p(\mathbf{z}) - \ln q_\phi(\mathbf{z}|\mathbf{x})] q_\phi(\mathbf{z}|\mathbf{x}) d\mathcal{M}(\mathbf{z}) \\
&= \mathbb{E}_{\mathbf{z} \sim q_\phi(\cdot|\mathbf{x})} [\ln p_\theta(\mathbf{x}|\mathbf{z}) + \ln p(\mathbf{z}) - \ln q_\phi(\mathbf{z}|\mathbf{x})] \\
&\triangleq \mathcal{L}_{\mathcal{M}}(\mathbf{x}; \theta, \phi) \\
&\approx \sum_k \ln p_\theta(\mathbf{x}|\mathbf{z}^k) + \ln p(\mathbf{z}^k) - \ln q_\phi(\mathbf{z}^k|\mathbf{x}), \quad \mathbf{z}^k \sim q_\phi(\cdot|\mathbf{x}) \sqrt{|G(\cdot)|}
\end{aligned}$$

## B Hyperbolic normal distributions

In this section, we first review some canonical generalisation of the normal distributions to Riemannian manifolds, and then introduce in more details the *Riemannian* and *wrapped* normal distributions on the Poincaré ball. Finally, we give architecture and training details about the conducted experiments.

### B.1 Probability measures on Riemannian manifolds

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

$$d\mathcal{M}(\mathbf{z}) = \sqrt{|G(\mathbf{z})|} d\mathbf{z}, \quad (4)$$

with  $d\mathbf{z}$  being the Lebesgue measure. Random variables  $\mathbf{z} \in \mathcal{M}$  would naturally be characterised by the Radon-Nikodym derivative of a measure  $\nu$  w.r.t. the Riemannian measure  $d\mathcal{M}(\cdot)$  (assuming absolute continuity)

$$f(\mathbf{z}) = \frac{d\nu(\mathbf{z})}{d\mathcal{M}(\mathbf{z})}.$$

Since the normal distribution plays such a canonical role in statistics, generalising it to manifold is of interest. Given a Fréchet expectation  $\boldsymbol{\mu} \in \mathcal{M}$  – defined as minimisers of  $\int_{\mathcal{M}} d_{\mathcal{M}}(\boldsymbol{\mu}, \mathbf{z})^2 p(\mathbf{z}) d\mathcal{M}(\mathbf{z})$  – and a dispersion parameter  $\sigma > 0$  (generally not equal to the standard deviation), several properties ought to be verified by such generalised normal distributions. Such a distribution should tend towards a delta function at  $\boldsymbol{\mu}$  when  $\sigma \rightarrow 0$  and to an (improper for non-compact) uniform distribution when  $\sigma \rightarrow \infty$ . Also, as the curvature tends to 0, one should recover the vanilla normal distribution. Hereby, we review canonical generalisations of the normal distribution, which have different theoretical and computational advantages.

**Maximum entropy normal** The property that Penne (2006) takes for granted is the maximization of the entropy given a mean and a covariance matrix, yielding in the isotropic setting

$$\frac{d\nu^{\mathbb{R}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2)}{d\mathcal{M}(\mathbf{z})} = \mathcal{N}_{\mathcal{M}}^{\mathbb{R}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2) = \frac{1}{Z^{\mathbb{R}}} \exp\left(-\frac{d_{\mathcal{M}}(\boldsymbol{\mu}, \mathbf{z})^2}{2\sigma^2}\right), \quad (5)$$

with  $d_{\mathcal{M}}$  being the Riemannian distance on the manifold induced by the tensor metric. Such a formulation – sometimes referred as *Riemannian Normal* distribution – is used by Said et al. (2014) in the Poincaré half-plane, or by Hauberg (2018) in the hypersphere  $\mathbb{S}^d$ . Sampling from such distributions and computing the normalising constant – especially in the anisotropic setting – is usually challenging.

**Wrapped normal** Another generalisation is defined by taking the image by the exponential map of a Gaussian distribution on the tangent space centered at the mean value. Such a distribution has been referred in literature as *wrapped*, *push-forward*, *exp-map* or *tangential* normal distribution. Sampling is therefore straightforward. The pdf is then readily available through the change of variable formula if one can compute the Jacobian of the exponential map (or its inverse). Hence such a distribution is attractive from a computational perspective. Grattarola et al. (2019) and Nagano et al. (2019) rely on such a distribution defined on the hyperboloid model. Wrapped distributions are often encountered in the *directional* statistics (Ley and Verdebout, 2017; Hauberg, 2018).

**Restricted normal** What is more, for sub-manifolds of  $\mathbb{R}^n$ , one can consider the restriction of a normal distribution pdf to the manifold. This yields the Von Mises distribution on  $\mathbb{S}^1$  and the Von Mises-Fisher distribution on  $\mathbb{S}^d$  (Hauberg, 2018) and the Stiefel manifold.

**Diffusion normal** Yet another generalisation arises by defining the normal pdf through the *heat kernel*, or fundamental solution of the heat equation,  $K : \mathbb{R}^+ \times \mathcal{M} \times \mathcal{M} \rightarrow \mathcal{M}$ ,

$$\mathcal{N}_{\mathcal{M}}^{\Delta}(z|\mu, \sigma^2) = K(\sigma^2/2, \mu, z). \quad (6)$$

See for instance Hsu (2008) for an introduction of Brownian motion on Riemannian manifolds and Paeng (2011) for conditions on existence and uniqueness of the kernel. Sampling amounts to simulating a Brownian motion, which may be challenging for non sub-manifolds of  $\mathbb{R}^n$ . Closed form solutions of the heat kernel is available for some manifolds such as spheres or flat tori, otherwise numerical approximations can be used. Such a distribution has been used in a VAE setting (Rey et al., 2019; Li et al., 2019).

**Other than normal distributions** Of course one needs not to restrict itself to generalisations of the normal distribution. For instance, one could consider a wrapped spherical Student-t as  $z \sim \exp_{\mu\#} S_t(0, \nu)$  or a Riemannian Student-t with density proportional to  $(1 + d_{\mathcal{M}}(z, \mu)^2/\nu)^{(-\nu+1)/2}$  (by making sure that this density is  $d_{\mathcal{M}}$ -integrable).

## B.2 Hyperbolic polar coordinates

In this subsection, we review the hyperbolic polar change of coordinates allowing us to reparametrise hyperbolic normal distributions in a similar fashion than the Box-Muller transform (Box and Muller, 1958).

**Polar coordinates** Euclidean polar coordinates, express points  $z \in \mathbb{R}^d$  through a radius  $r \geq 0$  and a direction  $\alpha \in \mathbb{S}^{d-1}$  such that  $z = r\alpha$ . Yet, one could choose another *pole* (or *reference point*)  $\mu \neq \mathbf{0}$  such that  $z = \mu + r\alpha$ . Consequently,  $r = d_E(\mu, z)$ . An analogous change of variables can also be constructed in Riemannian manifolds relying on the exponential map instead of the addition operator. Given a *pole*  $\mu \in \mathbb{B}_c^d$ , the point of hyperbolic polar coordinates  $z = (r, \alpha)$  is defined as  $z = \gamma(r)$ , with  $r = d_p^c(\mu, z)$  and  $\gamma : \mathbb{R}^+ \rightarrow \mathbb{B}_c^d$  a curve such that  $\gamma'(0) = \alpha \in \mathbb{S}^{d-1}$ . Hence  $z = \exp_{\mu}^c\left(\frac{r}{\lambda_{\mu}^c}\alpha\right)$  since  $d_p^c(\mu, z) = \|\ln_{\mu}^c(x)\|_{\mu} = \|\frac{r}{\lambda_{\mu}^c}\alpha\|_{\mu} = r$ .

**Tensor metric** We derive below the expression of the Poincaré ball metric in such hyperbolic polar coordinate, for the specific setting where  $\mu = \mathbf{0}$ :  $z = \exp_{\mathbf{0}}^c(\frac{r}{2}\alpha)$ . Switching to Euclidean polar coordinate we get

$$\begin{aligned} ds_{\mathbb{B}_c^d}^2 &= (\lambda_z^c)^2 (dz_1^2 + \dots + dz_d^2) = \frac{4}{(1 - c\|x\|^2)^2} dz^2 \\ &= \frac{4}{(1 - c\rho^2)^2} (d\rho^2 + \rho^2 ds_{\mathbb{S}^{d-1}}^2). \end{aligned} \quad (7)$$

Let's define  $r = d_p^c(\mathbf{0}, z) = L(\gamma)$ , with  $\gamma$  being the geodesic joining  $\mathbf{0}$  and  $z$ . Since such a geodesic is the segment  $[\mathbf{0}, z]$ , we have

$$r = \int_0^{\rho} \lambda_t^c dt = \int_0^{\rho} \frac{2}{1 - ct^2} dt = \int_0^{\sqrt{c}\rho} \frac{2}{1 - t^2} \frac{dt}{\sqrt{c}} = \frac{2}{\sqrt{c}} \tanh^{-1}(\sqrt{c}\rho).$$

Plugging  $\rho = \frac{1}{\sqrt{c}} \tanh(\sqrt{c}\frac{r}{2})$  (and  $d\rho = (1 - c\rho^2)/2dr$ ) into Eq 7 yields

$$\begin{aligned} ds_{\mathbb{B}_c^d}^2 &= \frac{4}{(1 - c\rho^2)^2} \frac{1}{4} (1 - c\rho^2)^2 dr^2 + \left(2 \frac{\rho}{1 - c\rho^2}\right)^2 ds_{\mathbb{S}^{d-1}}^2 \\ &= dr^2 + \left(2 \frac{\frac{1}{\sqrt{c}} \tanh(\sqrt{c}\frac{r}{2})}{1 - c \left(\frac{1}{\sqrt{c}} \tanh(\sqrt{c}\frac{r}{2})\right)^2}\right)^2 ds_{\mathbb{S}^{d-1}}^2 \\ &= dr^2 + \left(\frac{1}{\sqrt{c}} \sinh(\sqrt{c}r)\right)^2 ds_{\mathbb{S}^{d-1}}^2. \end{aligned} \quad (8)$$

The Euclidean line element is recovered when  $c \rightarrow 0$

$$ds_{\mathbb{R}^d}^2 = dr^2 + r^2 ds_{\mathbb{S}^{d-1}}^2. \quad (9)$$

In an appropriate orthonormal basis of  $\mathcal{T}_{\boldsymbol{\mu}} \mathbb{B}_c^d$  the hyperbolic polar coordinate leads to the following expression of the matrix of the metric

$$G(\mathbf{z}) = \begin{pmatrix} 1 & 0 \\ 0 & \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}r}\right)^2 \mathbf{I}_{d-1} \end{pmatrix}. \quad (10)$$

Hence, the density of the Riemannian measure with respect to the image of the Lebesgue measure of  $\mathcal{T}_{\boldsymbol{\mu}} \mathbb{B}_c^d$  by  $\exp_{\boldsymbol{\mu}}^c$  is given by

$$\sqrt{|G(\mathbf{z})|} = \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}r}\right)^{d-1}. \quad (11)$$

This result holds for any *reference point*  $\boldsymbol{\mu} \in \mathbb{B}_c^d$ , with  $r = d_p^c(\boldsymbol{\mu}, \mathbf{z})$ , since the metric induced measure is invariant under the isometries of the manifold (i.e. Möbius transformations). This result can also be found in Chevallier et al. (2015); Said et al. (2014). Also, the fact that the line element  $ds_{\mathbb{B}_c^d}^2$  and equivalently the metric  $G$  only depends on the radius in hyperbolic polar coordinate, is a consequence of the hyperbolic space's isotropy.

**Integration** We now make use of the aforementioned hyperbolic polar coordinates to integrate functions following Said et al. (2014). The integral of a function  $f : \mathbb{B}_c^d \rightarrow \mathbb{R}$  can be computed by using polar coordinates,

$$\begin{aligned} \int_{\mathbb{B}_c^d} f(\mathbf{z}) d\mathcal{M}(\mathbf{z}) &= \int_{\mathbb{B}_c^d} f(\mathbf{z}) \sqrt{|G(\mathbf{z})|} d\mathbf{z} \\ &= \int_{\mathcal{T}_{\boldsymbol{\mu}} \mathbb{B}_c^d \cong \mathbb{R}^d} f(\mathbf{v}) \sqrt{|G(\mathbf{v})|} d\mathbf{v} \end{aligned} \quad (12)$$

$$\begin{aligned} &= \int_{\mathbb{R}_+} \int_{\mathbb{S}^{d-1}} f(r) \sqrt{|G(r)|} dr r^{d-1} ds_{\mathbb{S}^{d-1}} \\ &= \int_{\mathbb{R}_+} \int_{\mathbb{S}^{d-1}} f(r) \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}r}\right)^{d-1} dr r^{d-1} ds_{\mathbb{S}^{d-1}} \\ &= \int_{\mathbb{R}_+} \int_{\mathbb{S}^{d-1}} f(r) \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}}\right)^{d-1} dr ds_{\mathbb{S}^{d-1}}. \end{aligned} \quad (13)$$

### B.3 Wrapped hyperbolic normal distribution on $\mathbb{B}_c^d$

**Anisotropic** The *wrapped normal* distribution considers a normal distribution in the tangent space  $\mathcal{T}_{\boldsymbol{\mu}} \mathbb{B}_c^d$  being pushed-forward along the exponential map. One can obtain sampled as follow

$$\mathbf{z} = \exp_{\boldsymbol{\mu}}^c \left( G(\boldsymbol{\mu})^{-\frac{1}{2}} \mathbf{v} \right) = \exp_{\boldsymbol{\mu}}^c \left( \frac{\mathbf{v}}{\lambda_{\boldsymbol{\mu}}^c} \right), \text{ with } \mathbf{v} \sim \mathcal{N}(\cdot | \mathbf{0}, \Sigma). \quad (14)$$

Then, its density is given by

$$\begin{aligned} \mathcal{N}_{\mathbb{B}_c^d}^{\mathbb{W}}(\mathbf{z}|\boldsymbol{\mu}, \Sigma) &= \mathcal{N}\left(G(\boldsymbol{\mu})^{1/2} \log_{\boldsymbol{\mu}}(\mathbf{z}) \mid \mathbf{0}, \Sigma\right) \left(\frac{\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z})}{\sinh(\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z}))}\right)^{d-1} \\ &= \mathcal{N}\left(\lambda_{\boldsymbol{\mu}}^c \log_{\boldsymbol{\mu}}(\mathbf{z}) \mid \mathbf{0}, \Sigma\right) \left(\frac{\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z})}{\sinh(\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z}))}\right)^{d-1} \end{aligned} \quad (15)$$

with  $G(\boldsymbol{\mu})^{1/2}$  the unique square-root matrix of  $G(\boldsymbol{\mu})$  (thanks to the positive definiteness of the metric tensor). This can be shown by plugging this density as  $f$  in Equation (12) with  $\mathbf{v} = r\boldsymbol{\alpha} = \lambda_{\boldsymbol{\mu}}^c \log_{\boldsymbol{\mu}}(\mathbf{z})$ , we get

$$\begin{aligned} \int_{\mathbb{B}_c^d} \mathcal{N}_{\mathbb{B}_c^d}^{\mathbb{W}}(\mathbf{z}|\boldsymbol{\mu}, \Sigma) d\mathcal{M}(\mathbf{z}) &= \int_{\mathcal{T}_{\boldsymbol{\mu}} \mathbb{B}_c^d \cong \mathbb{R}^d} \mathcal{N}(\mathbf{v} \mid \mathbf{0}, \Sigma) \left(\frac{\sqrt{c} \|\mathbf{v}\|_2}{\sinh(\sqrt{c} \|\mathbf{v}\|_2)}\right)^{d-1} \sqrt{|G(\mathbf{v})|} d\mathbf{v} \\ &= \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{v} \mid \mathbf{0}, \Sigma) \left(\frac{\sqrt{c} \|\mathbf{v}\|_2}{\sinh(\sqrt{c} \|\mathbf{v}\|_2)}\right)^{d-1} \left(\frac{\sinh(\sqrt{c} \|\mathbf{v}\|_2)}{\sqrt{c} \|\mathbf{v}\|_2}\right)^{d-1} d\mathbf{v} \\ &= \int_{\mathbb{R}^d} \mathcal{N}(\mathbf{v} \mid \mathbf{0}, \Sigma) d\mathbf{v}. \end{aligned}$$

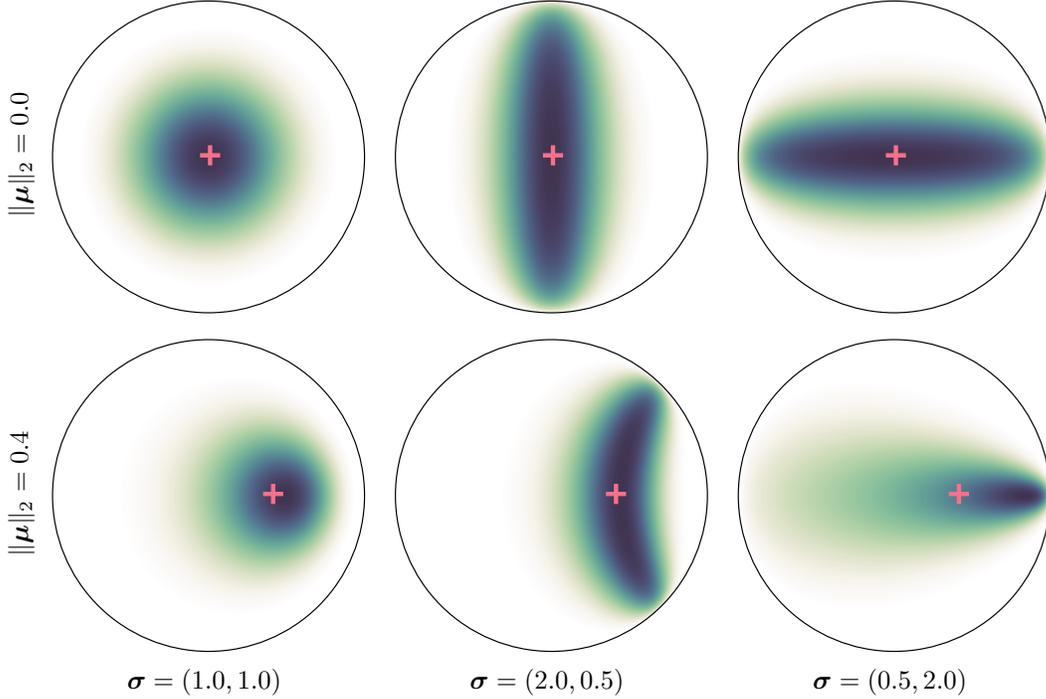


Figure 8: Anisotropic wrapped normal probability measures for Fréchet means  $\boldsymbol{\mu}$  (red +), concentrations  $\Sigma = \text{diag}(\boldsymbol{\sigma})$  and  $c = 1$ .

**Isotropic** In the isotropic setting, we therefore get

$$\int_{\mathbb{B}_c^d} \mathcal{N}_{\mathbb{B}_c^d}^{\mathbb{W}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2) d\mathcal{M}(\mathbf{z}) = \int_{\mathbb{R}_+} \int_{\mathbb{S}^{d-1}} \frac{1}{Z^{\mathbb{R}}} e^{-\frac{r^2}{2\sigma^2}} r^{d-1} dr d\mathbb{S}^{d-1}. \quad (16)$$

The hyperbolic radius  $r = d_p^c(\boldsymbol{\mu}, \mathbf{z})$  consequently follows the usual  $\chi$  distribution with density

$$\rho^{\mathbb{W}}(r) \propto \mathbb{1}_{\mathbb{R}_+}(r) e^{-\frac{r^2}{2\sigma^2}} r^{d-1}, \quad (17)$$

and the density of the *wrapped* normal given by

$$\mathcal{N}_{\mathbb{B}_c^d}^{\mathbb{W}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2) = \frac{d\nu^{\mathbb{W}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2)}{d\mathcal{M}(\mathbf{z})} = (2\pi\sigma^2)^{-d/2} \exp\left(-\frac{d_p^c(\boldsymbol{\mu}, \mathbf{z})^2}{2\sigma^2}\right) \left(\frac{\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z})}{\sinh(\sqrt{c} d_p^c(\boldsymbol{\mu}, \mathbf{z}))}\right)^{d-1}.$$

## B.4 Maximum entropy hyperbolic normal distribution on $\mathbb{B}_c^d$

Alternatively, by considering the maximum entropy generalisation of the normal distribution one gets (Pennec, 2006)

$$\mathcal{N}_{\mathbb{B}_c^d}^{\mathbb{R}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2) = \frac{d\nu^{\mathbb{R}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2)}{d\mathcal{M}(\mathbf{z})} = \frac{1}{Z^{\mathbb{R}}} \exp\left(-\frac{d_p^c(\boldsymbol{\mu}, \mathbf{z})^2}{2\sigma^2}\right). \quad (18)$$

Such a pdf can be computed pointwise once  $Z^{\mathbb{R}}$  is known, which we derive in Appendix B.4.3. Also, we observe that as  $c$  and  $\sigma$  get smaller (resp. bigger), the *Riemannian* normal pdf gets closer (resp. further) to the *wrapped* normal pdf.

### B.4.1 Reparametrisation

Plugging the *Riemannian* normal density as  $f$  in Equation (13), with  $r = d_p^c(\boldsymbol{\mu}, \mathbf{z})$ , we have

$$\begin{aligned} \int_{\mathbb{B}_c^d} \mathcal{N}_{\mathbb{B}_c^d}^{\mathbb{R}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2) d\mathcal{M}(\mathbf{z}) &= \int_{\mathbb{R}_+} \int_{\mathbb{S}^{d-1}} \frac{1}{Z^{\mathbb{R}}} e^{-\frac{r^2}{2\sigma^2}} \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}}\right)^{d-1} dr ds_{\mathbb{S}^{d-1}} \\ &= \frac{1}{Z^{\mathbb{R}}} \left( \int_{\mathbb{R}_+} e^{-\frac{r^2}{2\sigma^2}} \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}}\right)^{d-1} dr \right) \left( \int_{\mathbb{S}^{d-1}} ds_{\mathbb{S}^{d-1}} \right) \end{aligned} \quad (19)$$

Hence, samples  $\mathbf{z} \sim \mathcal{N}_{\mathcal{M}}^{\mathbb{R}}(\mathbf{z}|\boldsymbol{\mu}, \sigma^2)d\mathcal{M}(\mathbf{z})$  can be reparametrised as

$$\mathbf{z} = \exp_{\boldsymbol{\mu}}^c \left( \frac{r}{\lambda_{\boldsymbol{\mu}}^c} \boldsymbol{\alpha} \right) \quad (20)$$

with the direction  $\boldsymbol{\alpha}$  being uniformly distributed on the hypersphere  $\mathbb{S}^{d-1}$ , i.e.

$$\boldsymbol{\alpha} \sim \mathcal{U}(\mathbb{S}^{d-1})$$

and the hyperbolic radius  $r = d_p^c(\boldsymbol{\mu}, \mathbf{z})$  distributed according to the following density (w.r.t the Lebesgue measure)

$$\rho^{\mathbb{R}}(r) = \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} e^{-\frac{r^2}{2\sigma^2}} \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}}\right)^{d-1}. \quad (21)$$

**Developed expression** By expanding the sinh term using the binomial formula, we get

$$\begin{aligned} \rho^{\mathbb{R}}(r) &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} e^{-\frac{r^2}{2\sigma^2}} \left(\frac{\sinh(\sqrt{c}r)}{\sqrt{c}}\right)^{d-1} \\ &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} e^{-\frac{r^2}{2\sigma^2}} \left(\frac{e^{\sqrt{c}r} - e^{-\sqrt{c}r}}{2\sqrt{c}}\right)^{d-1} \\ &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} e^{-\frac{r^2}{2\sigma^2}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} \binom{d-1}{k} (e^{\sqrt{c}r})^{d-1-k} (-e^{-\sqrt{c}r})^k \\ &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} e^{-\frac{r^2}{2\sigma^2}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{(d-1-2k)\sqrt{c}r} \\ &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{-\frac{r^2}{2\sigma^2} + (d-1-2k)\sqrt{c}r} \\ &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c \sigma^2} e^{-\frac{1}{2\sigma^2} [r - (d-1-2k)\sqrt{c}\sigma^2]^2}. \end{aligned} \quad (22)$$

## B.4.2 Sampling

In this section we detail the sampling scheme that we use for the Riemannian normal distribution  $\mathcal{N}_{\mathbb{B}_c^d}^R(\cdot|\boldsymbol{\mu}, \sigma^2)$ , along with a reparametrisation which allows to compute gradients with respect to the parameters  $\boldsymbol{\mu}$  and  $\sigma$ .

**Sampling challenges due to the hyperbolic geometry** Several properties of the Euclidean space do not generalise to the hyperbolic setting, unfortunately hardening the task of obtaining samples from *Riemannian* normal distributions. First, one can factorise a normal density through the space’s dimensions – thanks to the Pythagorean theorem – hence allowing to divide the task on several subspaces and then concatenate the samples. Such a property does not extend to the hyperbolic geometry, thus seemingly preventing us from focusing on 2-dimensional samples. Second, in Euclidean geometry, the polar radius  $r$  is distributed according to  $\rho^W(r) = \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r} e^{-\frac{r^2}{2\sigma^2}} r^{d-1}$ , making it easy by a linear change of variable to take into account different scaling values. The non-linearity of  $\sinh$  prevent us from using such a simple change of variable.

**Computing gradients with respect to parameters** So as to compute gradients of samples  $\mathbf{z}$  with respect to the parameters  $\boldsymbol{\mu}$  and  $\sigma$  of samples of a hyperbolic distributions, we respectively rely on the reparametrisation given by Eq 20 for  $\nabla_{\boldsymbol{\mu}}\mathbf{z}$ , and on an implicit reparametrisation (Figurnov et al., 2018) of  $r$  for  $\nabla_{\sigma}\mathbf{z}$ . We have  $\mathbf{z} = \exp_{\boldsymbol{\mu}}^c\left(\frac{r}{\lambda_{\boldsymbol{\mu}}^c}\boldsymbol{\alpha}\right)$  with  $\boldsymbol{\alpha} \sim \mathcal{U}(\mathbb{S}^{d-1})$  and  $r \sim \rho^R(\cdot)$ . Hence,

$$\nabla_{\boldsymbol{\mu}}\mathbf{z} = \nabla_{\boldsymbol{\mu}} \exp_{\boldsymbol{\mu}}^c(\mathbf{u}), \quad (23)$$

with  $\mathbf{u} = \frac{r}{\lambda_{\boldsymbol{\mu}}^c}\boldsymbol{\alpha}$  (actually) independent of  $\boldsymbol{\mu}$ , and

$$\nabla_{\sigma}\mathbf{z} = \nabla_{\sigma} \exp_{\boldsymbol{\mu}}^c(\mathbf{u}) = \nabla_{\mathbf{u}} \exp_{\boldsymbol{\mu}}^c(\mathbf{u}) \frac{\boldsymbol{\alpha}}{\lambda_{\boldsymbol{\mu}}^c} \nabla_{\sigma} r, \quad (24)$$

with  $\nabla_{\sigma}(r)$  computed via the implicit reparametrisation given by

$$\begin{aligned} \nabla_{\sigma}(r) &= -(\nabla_r F^R(r, \sigma))^{-1} \nabla_{\sigma} F^R(r, \sigma) \\ &= -(\rho^R(r; \sigma))^{-1} \nabla_{\sigma} F^R(r, \sigma). \end{aligned} \quad (25)$$

**Sampling hyperbolic radii** Unfortunately the density of the hyperbolic radius  $\rho^R(r)$  is not a well-known distribution and its cumulative density function does not seem analytically invertible. We therefore rely on rejection sampling methods.

**Adaptive Rejection Sampling** By making use of the log-concavity of  $\rho^R$ , we can rely on a piecewise exponential distribution proposal from adaptive rejection sampling (ARS) (R. Gilks and Wild, 1992). Such a proposal automatically adapt itself with respect to the parameters  $\sigma$ ,  $c$  and  $d$ . Even though  $\mathcal{N}_{\mathbb{B}_c^d}$  is defined on a  $d$ -dimensional manifold,  $\rho^R$  is a univariate distribution hence the sampling scheme is not directly affected by dimensionality. The difficulty in ARS is to choose the initial set of points to construct the piecewise exponential proposal. To do so, we first compute the mean  $m = \mathbb{E}_{r \sim \rho^R}[r]$  and standard deviation  $s = \mathbb{V}_{r \sim \rho^R}[r]^{1/2}$  of the targeted distribution. Then we choose a grid  $\boldsymbol{\eta} = (\eta_1, \dots, \eta_K) = (\text{linspace}(\eta_{max}, \eta_{min}, K/2), \text{linspace}(\eta_{min}, \eta_{max}, K/2))$ . Eventually, we set the initial points  $(x_1, \dots, x_K)$  to  $x_k = m + \eta_k * \min(s, 0.95 * m/\eta_{max})$ . For our experiments we chose  $\eta_{min} = .1, \eta_{max} = 3, K = 20$ . We do not adapt the proposal within the rejection sampling since we empirically found it unnecessary.

Alternatively, we derived below two non-adaptive proposal distributions along with their rejection rate constants. Yet, we observe that these rates do not scale well the dimensionality  $d$  and distortion  $\sigma$ , making them ill-suited for practical purposes.

**Rejection Sampling with truncated Normal proposal** The developed expression of  $\rho(r)^R$  from Eq 22 highlights the fact that the density can immediately be upper bounded by a truncated normal

density:

$$\begin{aligned}\rho(r)^{\mathbb{R}} &= \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c\sigma^2} e^{-\frac{1}{2\sigma^2} [r-(d-1-2k)\sqrt{c}\sigma^2]^2} \\ &\leq \frac{\mathbb{1}_{\mathbb{R}_+}(r)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{2k=0}^{d-1} \binom{d-1}{2k} e^{\frac{(d-1-4k)^2}{2} c\sigma^2} e^{-\frac{1}{2\sigma^2} [r-(d-1-4k)\sqrt{c}\sigma^2]^2}.\end{aligned}$$

Then we choose our proposal  $g$  to be the truncated normal distribution associated with  $k = 0$ , i.e. with mean  $(d-1)\sqrt{c}\sigma^2$  and variance  $\sigma^2$

$$\begin{aligned}g(r) &= \frac{\mathbb{1}_{r>0}}{\sigma \left(1 - \Phi\left(-\frac{(d-1)\sqrt{c}\sigma^2}{\sigma}\right)\right)} \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2\sigma^2} (r-(d-1)\sqrt{c}\sigma^2)^2} \\ &= \frac{1}{\sqrt{2\pi}} \frac{\mathbb{1}_{r>0}}{\sigma \left(1 - \frac{1}{2} - \frac{1}{2} \operatorname{erf}\left(-\frac{(d-1)\sqrt{c}\sigma}{\sqrt{2}}\right)\right)} e^{-\frac{1}{2\sigma^2} (r-(d-1)\sqrt{c}\sigma^2)^2} \\ &= \sqrt{\frac{2}{\pi}} \frac{\mathbb{1}_{r>0}}{\sigma \left(1 + \operatorname{erf}\left(\frac{(d-1)\sqrt{c}\sigma}{\sqrt{2}}\right)\right)} e^{-\frac{1}{2\sigma^2} (r-(d-1)\sqrt{c}\sigma^2)^2} \\ &= \frac{\mathbb{1}_{r>0}}{Z_g(\sigma)} e^{-\frac{1}{2\sigma^2} (r-(d-1)\sqrt{c}\sigma^2)^2}\end{aligned}\tag{26}$$

with

$$Z_g = \sqrt{\frac{\pi}{2}} \sigma \left(1 + \operatorname{erf}\left(\frac{(d-1)\sqrt{c}\sigma}{\sqrt{2}}\right)\right).\tag{27}$$

Computing the ratio of the densities yield

$$\begin{aligned}\frac{\rho(r)^{\mathbb{R}}}{g(r)} &= \frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c\sigma^2} e^{-\frac{1}{2\sigma^2} [r-(d-1-2k)\sqrt{c}\sigma^2]^2} e^{+\frac{1}{2\sigma^2} [r-(d-1)\sqrt{c}\sigma^2]^2} \\ &= \frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c\sigma^2} e^{2k\sqrt{c}((d-1-k)\sqrt{c}\sigma^2 - r)}.\end{aligned}$$

Hence

$$\rho(r)^{\mathbb{R}}/g(r) \leq M \triangleq \frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} e^{\frac{(d-1)^2 c\sigma^2}{2}}.\tag{28}$$

**Rejection Sampling with Gamma proposal** Now let's consider the following Gamma(2,  $\sigma$ ) density:

$$g(r) = \frac{\mathbb{1}_{r>0}}{Z_g(\sigma)} r e^{-\frac{r}{\sigma}}$$

with

$$Z_g(\sigma) = \Gamma(2)\sigma^2.$$

Then log ratio of the densities can be upper bounded as following:

$$\begin{aligned}\ln\left(\frac{\rho(r)^{\mathbb{R}}}{g(r)}\right) &= \ln\frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} - \frac{r^2}{2\sigma^2} + (d-1) \ln(e^{\sqrt{c}r} - e^{-\sqrt{c}r}) - (d-1) \ln 2 - \ln r + \frac{r}{\sigma} \\ &= \ln\frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} - (d-1) \ln 2 - \underbrace{\frac{r^2}{2\sigma^2} + \left((d-1)\sqrt{c} + \frac{1}{\sigma}\right)r}_{\leq \frac{((d-1)\sqrt{c}\sigma+1)^2}{2}} + \underbrace{(d-1) \ln\left(\frac{1 - e^{-2\sqrt{c}r}}{r}\right)}_{\leq (d-1) \ln(2\sqrt{c})} \\ &\leq \ln\frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} + \frac{((d-1)\sqrt{c}\sigma+1)^2}{2} + (d-1) \ln\sqrt{c}.\end{aligned}$$

Hence

$$\rho(r)^{\mathbb{R}}/g(r) \leq M \triangleq \frac{Z_g(\sigma)}{Z_r^{\mathbb{R}}} c^{\frac{d-1}{2}} e^{\frac{((d-1)\sqrt{c}\sigma+1)^2}{2}}.\tag{29}$$

### B.4.3 Normalisation constant

In order to evaluate the density of the *Riemannian* normal distribution, we need to compute the normalisation constant, which we derive in this subsection.

**Cumulative density function** First let's derive the cumulative density function of the hyperbolic radius. Integrating the expended density of Eq (22) yields

$$\begin{aligned}
F_r^R(r) &= \int_{-\infty}^r \rho^R(r) dr \\
&= \frac{1}{Z_r^R} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c\sigma^2} \\
&\quad \times \int_0^r e^{-\frac{1}{2\sigma^2} [r - (d-1-2k)\sqrt{c}\sigma^2]^2} dr \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c\sigma^2} \\
&\quad \times \left[ \operatorname{erf} \left( \frac{r - (d-1-2k)\sqrt{c}\sigma^2}{\sqrt{2}\sigma} \right) \operatorname{erf} \left( \frac{(d-1-2k)\sqrt{c}\sigma}{\sqrt{2}} \right) \right] \quad (30)
\end{aligned}$$

with  $\Phi : x \mapsto \frac{1}{2} \left( 1 + \operatorname{erf} \left( \frac{x}{\sqrt{2}} \right) \right)$ , the cumulative distribution function of a standard normal distribution.

**Taking the limit**  $F_r^R(r) \xrightarrow[r \rightarrow \infty]{} 1$  in Eq (30) yield

$$Z_r^R = \sqrt{\frac{\pi}{2}} \sigma \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c\sigma^2} \left[ 1 + \operatorname{erf} \left( \frac{(d-1-2k)\sqrt{c}\sigma}{\sqrt{2}} \right) \right]. \quad (31)$$

Note that by the antisymmetry of erf, one can simplify Eq (31) with a sum over  $\lceil d/2 \rceil$  terms (as done in Hauberg (2018)). Also, computing such a sum is much more stable by relying on the *log sum exp* trick. Integrating Equation (19) of Appendix B.2 gives

$$Z^R = Z_r^R Z_\alpha \quad (32)$$

As a reminder, the surface area of the  $d - 1$ -dimensional hypersphere with radius 1 is given by

$$Z_\alpha = A_{\mathbb{S}^{d-1}} = \frac{2\pi^{d/2}}{\Gamma(d/2)}.$$

For the special case of  $c = 1$  and  $d = 2$  we recover the formula given in Said et al. (2014)

$$Z_r^R = \sqrt{\frac{\pi}{2}} \sigma e^{\frac{\sigma^2}{2}} \operatorname{erf} \left( \frac{\sigma}{\sqrt{2}} \right).$$

### B.4.4 Expectation of hyperbolic radii

Computing the expectation of the hyperbolic radius  $r \sim \rho^R$  is of use to choose the initial set of points to construct the piecewise exponential proposal. By integrating the expended density of Eq (22), we

get

$$\begin{aligned}
\mathbb{E}[r] &= \int_{-\infty}^{\infty} r \rho^{\mathbb{R}}(r) dr \\
&= \frac{1}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c \sigma^2} \int_0^{\infty} r e^{-\frac{1}{2\sigma^2} [r - (d-1-2k)\sqrt{c}\sigma^2]^2} dr \\
&= \frac{1}{Z_r^{\mathbb{R}}} \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} e^{\frac{(d-1-2k)^2}{2} c \sigma^2} \\
&\quad \times \left[ \sqrt{\frac{\pi}{2}} (d-1-2k) \sqrt{c} \sigma^2 \left( 1 + \operatorname{erf} \left( \frac{(d-1-2k)\sqrt{c}\sigma}{\sqrt{2}} \right) \right) + \sigma^2 e^{-\frac{(d-1-2k)^2 c \sigma^2}{2}} \right] \\
&= \frac{1}{Z_r^{\mathbb{R}}} \sqrt{\frac{\pi}{2}} \sigma \frac{1}{(2\sqrt{c})^{d-1}} \sum_{k=0}^{d-1} (-1)^k \binom{d-1}{k} \\
&\quad \times \left[ e^{\frac{(d-1-2k)^2}{2} c \sigma^2} (d-1-2k) \sqrt{c} \sigma^2 \left( 1 + \operatorname{erf} \left( \frac{(d-1-2k)\sqrt{c}\sigma}{\sqrt{2}} \right) \right) + \sigma \sqrt{\frac{2}{\pi}} \right]
\end{aligned} \tag{33}$$

## C Experimental details

In this section we give more details on the datasets, architecture designs and optimisation schemes used for the experimental results given in Section 5

### C.1 Synthetic Branching Diffusion Process

**Generation** Nodes  $(\mathbf{y}_1, \dots, \mathbf{y}_N) \in \mathbb{R}^n$  of the branching diffusion process are sampled as follow

$$\mathbf{y}_i \sim \mathcal{N}(\cdot | \mathbf{y}_{\pi(i)}, \sigma_0^2) \quad \forall i \in 1, \dots, N$$

with  $\pi(i)$  being the index of the  $i$ th node’s ancestor and  $d(i)$  its depth. Then, noisy observations are sampled for each node  $\mathbf{x}_i$ ,

$$\mathbf{x}_{i,j} = \mathbf{y}_i + \boldsymbol{\epsilon}_{i,j}, \quad \boldsymbol{\epsilon}_{i,j} \sim \mathcal{N}(\cdot | \mathbf{0}, \sigma_j^2) \quad \forall i, j.$$

The root  $x_0$  is set to  $\mathbf{0}$  for simplicity. The observation dimension is set to  $n = 50$ . The dataset  $(\mathbf{x}_{i,j})_{i,j}$  is centered and normalised to have unit variance. Thus, the choice of variance  $\sigma_0^2$  does not matter and it is set to  $\sigma_0 = 1$ . The number of noisy observations is set to  $J = 5$ , and its variance to  $\sigma_j^2 = \sigma_0^2/5 = 1/5$ . The depth is set to 6 and the branching factor to 2.

**Architectures** Both  $\mathcal{N}$ -VAE and  $\mathcal{P}^c$ -VAE decoders parametrise the mean of the unit variance Gaussian likelihood  $\mathcal{N}(\cdot | f_{\theta}(\mathbf{z}), 1)$ . Their encoders parametrise the mean and the log-variance of respectively an isotropic normal distribution  $\mathcal{N}(\cdot | g_{\phi}(\mathbf{z}))$  and an isotropic hyperbolic normal distribution  $\mathcal{N}_{\mathbb{B}_c^d}(\cdot | g_{\phi}(\mathbf{z}))$ . The  $\mathcal{N}$ -VAE’s encoder and decoder are composed of 2 Fully-Connected layers with a ReLU activation in between, as summed up in Tables 5 and 6. The  $\mathcal{P}^c$ -VAE’s design is similar, the differences being that the decoder’s output is mapped to manifold via the exponential map  $\exp_0^c$ , and the decoder’s first layer is made of *gyroplane units* presented in Section 3.2, as summarised in Tables 7 and 8. Observations live in  $\mathcal{X} = \mathbb{R}^{50}$  and the latent space dimensionality  $d$  is set to  $d = 2$ .

Table 5: Encoder network for  $\mathcal{N}$ -VAE

Layer	Output dim	Activation
Input	50	Identity
FC	200	ReLU
FC	2, 1	Identity

Table 6: Decoder network for  $\mathcal{N}$ -VAE

Layer	Output dim	Activation
Input	2	Identity
FC	200	ReLU
FC	50	Identity

Table 7: Encoder network for  $\mathcal{P}^c$ -VAE

Layer	Output dim	Activation
Input	50	Identity
FC	200	ReLU
FC	2, 1	$\exp_{\mathbf{0}}^c$ , Identity

Table 8: Decoder network for  $\mathcal{P}^c$ -VAE

Layer	Output dim	Activation
Input	2	Identity
Gyroplane	200	ReLU
FC	50	Identity

The synthetic datasets are generated as described in Section 5, then centred and normalised to unit variance. There are then randomly split into training and testing datasets with a proportion 0.7.

**Optimisation** Gyroplane offset  $\mathbf{p} \in \mathbb{B}_c^d$  are only *implicitly* parametrised to live in the manifold, by projecting a real vector  $\mathbf{p} = \exp_{\mathbf{0}}^c(\mathbf{p}')$ . Hence, all parameters  $\{\boldsymbol{\theta}, \boldsymbol{\phi}\}$  of the model explicitly live in Euclidean spaces which means that usual optimisation schemes can be applied. We therefore rely on Adam optimiser (Kingma and Ba, 2016) with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and a constant learning rate set to  $1e - 3$ . Models are trained with mini-batches of size 64 for 1000 epochs. The ELBO is approximated with a MC estimate with  $K = 1$ .

**Baselines** The principal component analysis (PCA) embeddings are obtained via a singular-value decomposition (SVD) by projecting the dataset on the basis associated with the two highest singular values. The Gaussian process latent variable model (GPLVM) embeddings are obtained by maximising the marginal likelihood of a (non-Bayesian) GPLVM with RBF kernel, and whose latent variables are initialised with PCA.

## C.2 MNIST digits

The MNIST dataset (LeCun and Cortes, 2010) contains 60,000 training and 10,000 test images of ten handwritten digits (zero to nine), with 28x28 pixels.

**Architectures** The architectures used for the encoder and the decoder for Mnist are similar to the ones used for the Synthetic Branching Diffusion Process. They differ by the dimensions of the observation space ( $\mathcal{X} = \mathbb{R}^{28 \times 28}$ ) and hidden space. The output of the first fully connected layer is here equal to 600. The latent space dimensionality  $d$  is set to 2, 5, 10 and 20 respectively. The bias of the decoder’s last layer is set to the average value of digits (for each pixel). The architectures used for the classifier are similar than the decoder architectures, the only difference being the output dimensionality (10 labels). We initialise the classifier’s first layer with decoder’s first layer weights. Then the classifier is trained to minimise the cross entropy for 5 epochs, with mini-batches of size 64 and a constant learning rate of  $1e - 3$ .

**Optimisation** We use  $[0, 1]$  normalised data as targets for the mean of a Bernoulli distribution, using negative cross-entropy for log  $p(\mathbf{x}|\mathbf{z})$ . We set the prior distribution’s distortion to  $\sigma = 1$ . We rely on Adam optimiser with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and a constant learning rate of  $5e - 4$ . Models are trained with mini-batches of size 128 for 80 epochs.

## C.3 Graph embeddings

The PhD advisor-advisee relationships graph (Nooy et al., 2011) contains 344 nodes and 343 edges. The phylogenetic tree expressing genetic heritage (Hofbauer et al., 2016; Sanderson and Eriksson, 1994) contains 1025 nodes and 1043 edges. The biological set representing disease relationships (Goh et al., 2007; Rossi and Ahmed, 2015) contains 516 nodes and 1188 edges. We follow the training and evaluation procedure introduced in Kipf and Welling (2016).

**Architectures** We also follow the featureless architecture introduced in Kipf and Welling (2016), namely a two-layer GCN with 32 hidden dimensions to parametrise the variational posteriors, and a likelihood which factorises along edges  $p(\mathbf{A}|\mathbf{Z}) = \prod_{i=1}^N \prod_{j=1}^N p(A_{ij}|\mathbf{z}_i, \mathbf{z}_j)$ , with  $\mathbf{A}$  being the adjacency matrix. The probability of an edge is defined through the latent metric by  $p(A_{ij} = 1|\mathbf{z}_i, \mathbf{z}_j) = 1 - \tanh(d_{\mathcal{M}}(\mathbf{z}_i, \mathbf{z}_j))$ . For the Poincaré ball latent space, the encoder output is projected

on the manifold:  $\mu = \exp_0(\text{GCN}_\mu(\mathbf{A}))$ . The latent dimension is set to 5 for the experiments. We use a *Wrapped* Gaussian prior and variational posterior.

**Optimisation** We use the adjacency matrix  $\mathbf{A}$  as target for the mean of a Bernoulli distribution, using negative cross-entropy for  $\log p(\mathbf{A}|\mathbf{Z})$ . We rely on Adam optimiser with parameters  $\beta_1 = 0.9$ ,  $\beta_2 = 0.999$  and a constant learning rate of  $1e^{-2}$ . We perform full-batch gradient descent for 800 epochs and make use of the reparametrisation trick for training.

## D More experimental qualitative results

Figure 9 shows latent representations of  $\mathcal{P}^c$ -VAEs with different curvatures. With "small" curvatures, we observe that embeddings lie close the center of the ball, where the geometry is close to be Euclidean. Similarly as Figure 9, Figure 10 illustrates the learned latent representations of  $\mathcal{P}^c$ -VAE with decreasing curvatures  $c$ , by highlighting the leaned *gyroplanes* of the decoder.

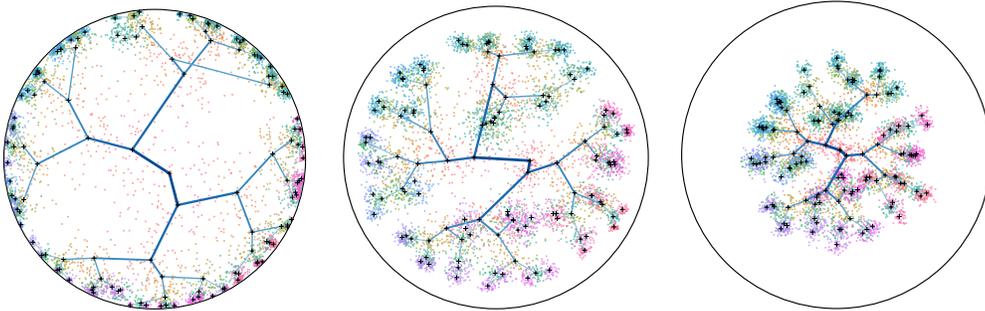


Figure 9: Branching diffusion process latent representations of  $\mathcal{P}^c$ -VAE with decreasing curvatures  $c = 1.2, 0.3, 0.1$  (Left to Right).

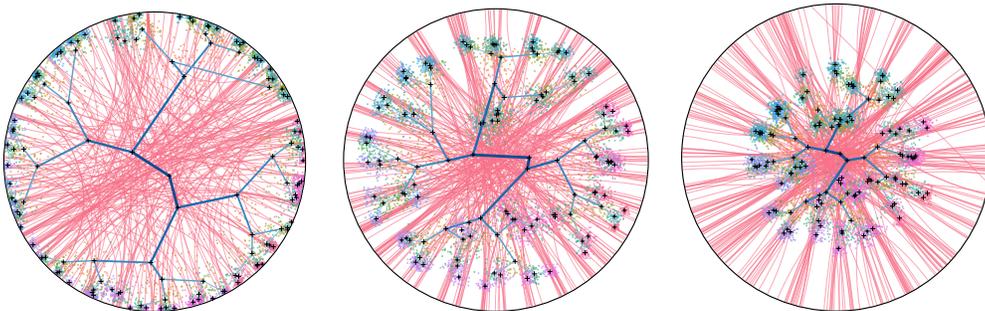


Figure 10: Branching diffusion process latent representations of  $\mathcal{P}^c$ -VAE with decreasing curvatures  $c = 1.2, 0.3, 0.1$  (Left to Right).

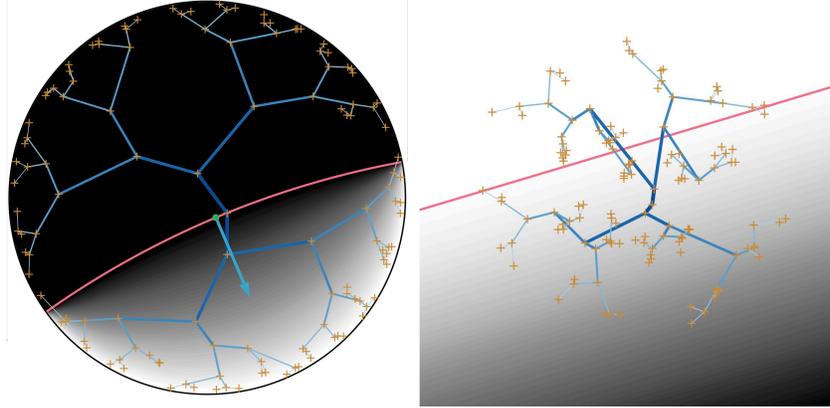


Figure 11: Branching diffusion process latent representation of  $\mathcal{P}^1$ -VAE (Left) and  $\mathcal{N}$ -VAE (Right) with heatmap of the log distance to the hyperplane (in pink).

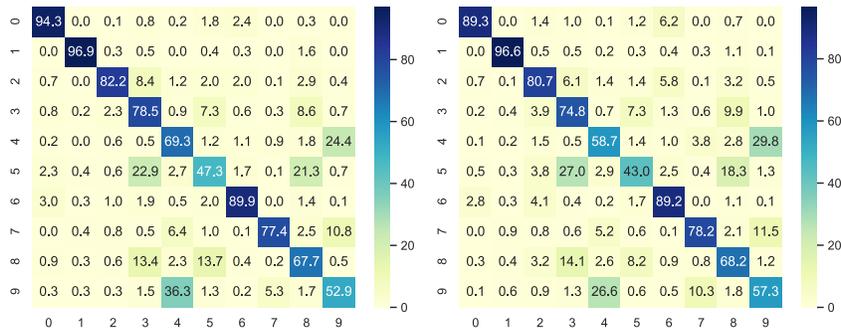


Figure 12: MNIST average confusion matrices of the classifiers trained on embeddings from the  $\mathcal{P}^{1.4}$ -VAE (Left) and  $\mathcal{N}$ -VAE (Right) models.