
A Wrapped Normal Distribution on Hyperbolic Space for Gradient-Based Learning

Yoshihiro Nagano¹ Shoichiro Yamaguchi² Yasuhiro Fujita² Masanori Koyama²

Abstract

Hyperbolic space is a geometry that is known to be well-suited for representation learning of data with an underlying hierarchical structure. In this paper, we present a novel hyperbolic distribution called *hyperbolic wrapped distribution*, a wrapped normal distribution on hyperbolic space whose density can be evaluated analytically and differentiated with respect to the parameters. Our distribution enables the gradient-based learning of the probabilistic models on hyperbolic space that could never have been considered before. Also, we can sample from this hyperbolic probability distribution without resorting to auxiliary means like rejection sampling. As applications of our distribution, we develop a hyperbolic-analog of variational autoencoder and a method of probabilistic word embedding on hyperbolic space. We demonstrate the efficacy of our distribution on various datasets including MNIST, Atari 2600 Breakout, and WordNet.

1. Introduction

Recently, hyperbolic geometry is drawing attention as a powerful geometry to assist deep networks in capturing fundamental structural properties of data such as a hierarchy. Hyperbolic attention network (Gülçehre et al., 2019) improved the generalization performance of neural networks on various tasks including machine translation by imposing the hyperbolic geometry on several parts of neural networks. Poincaré embeddings (Nickel & Kiela, 2017) succeeded in learning a parsimonious representation of symbolic data by embedding the dataset into Poincaré balls.

¹Department of Complexity Science and Engineering, The University of Tokyo, Japan ²Preferred Networks, Inc., Japan. Correspondence to: Yoshihiro Nagano <nagano@mns.k.u-tokyo.ac.jp>.

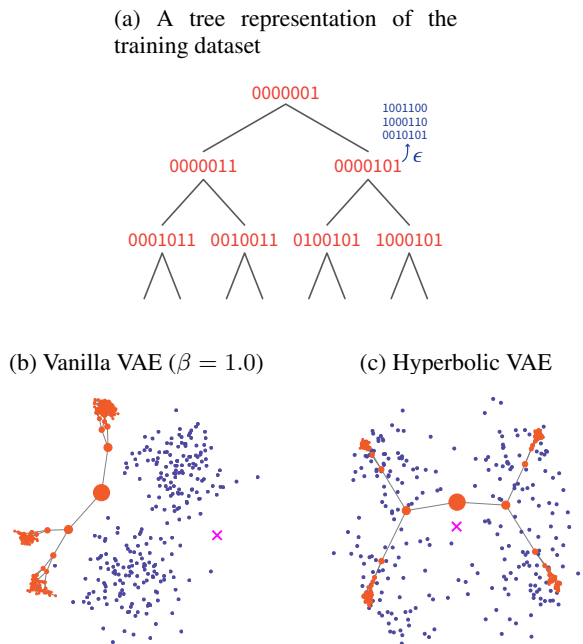


Figure 1: The visual results of Hyperbolic VAE applied to an artificial dataset generated by applying random perturbations to a binary tree. The visualization is being done on the Poincaré ball. The red points are the embeddings of the original tree, and the blue points are the embeddings of noisy observations generated from the tree. The pink \times represents the origin of the hyperbolic space. The VAE was trained without the prior knowledge of the tree structure. Please see 6.1 for experimental details

In the task of data embedding, the choice of the target space determines the properties of the dataset that can be learned from the embedding. For the dataset with a hierarchical structure, in particular, the number of relevant features can grow exponentially with the depth of the hierarchy. Euclidean space is often inadequate for capturing the structural information (Figure 1). If the choice of the target space of the embedding is limited to Euclidean space, one might have to prepare extremely high dimensional space as the target space to guarantee small distortion. However, the same embedding can be done remarkably well if the destination is hyperbolic space (Sarkar, 2012; Sala et al., 2018).

Now, the next natural question is; “how can we extend these works to *probabilistic* inference problems on hyperbolic space?” When we know in advance that there is a hierarchical structure in the dataset, a prior distribution on hyperbolic space might serve as a good *informative* prior. We might also want to make Bayesian inference on a dataset with hierarchical structure by training a variational autoencoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014) with latent variables defined on hyperbolic space. We might also want to conduct probabilistic word embedding into hyperbolic space while taking into account the uncertainty that arises from the underlying hierarchical relationship among words. Finally, it would be best if we can compare different probabilistic models on hyperbolic space based on popular statistical measures like divergence that requires the explicit form of the probability density function.

The endeavors we mentioned in the previous paragraph all require *probability distributions on hyperbolic space* that admit a parametrization of the density function that can be **computed analytically** and **differentiated** with respect to the parameter. Also, we want to be able to **sample from the distribution efficiently**; that is, we do not want to resort to auxiliary methods like rejection sampling.

In this study, we present a novel hyperbolic distribution called *hyperbolic wrapped distribution*, a wrapped normal distribution on hyperbolic space that resolves all these problems. We construct this distribution by defining Gaussian distribution on the tangent space at the origin of the hyperbolic space and projecting the distribution onto hyperbolic space after transporting the tangent space to a desired location in the space. This operation can be formalized by a combination of the parallel transport and the exponential map for the Lorentz model of hyperbolic space.

We can use our hyperbolic wrapped distribution to construct a probabilistic model on hyperbolic space that can be trained with gradient-based learning. For example, our distribution can be used as a prior of a VAE (Figure 1, Figure 6). It is also possible to extend the existing probabilistic embedding method to hyperbolic space using our distribution, such as probabilistic word embedding. We will demonstrate the utility of our method through the experiments of probabilistic hyperbolic models on benchmark datasets including MNIST, Atari 2600 Breakout, and WordNet.

2. Background

2.1. Hyperbolic Geometry

Hyperbolic geometry is a non-Euclidean geometry with a constant negative Gaussian curvature, and it can be visualized as the forward sheet of the two-sheeted hyper-

boloid. There are four common equivalent models used for the hyperbolic geometry: the Klein model, Poincaré disk model, and Lorentz (hyperboloid/Minkowski) model, and Poincaré half-plane model. Many applications of hyperbolic space to machine learning to date have adopted the Poincaré disk model as the subject of study (Nickel & Kiela, 2017; Ganea et al., 2018a;b; Sala et al., 2018). In this study, however, we will use the Lorentz model that, as claimed in Nickel & Kiela (2018), comes with a simpler closed form of the geodesics and does not suffer from the numerical instabilities in approximating the distance. We will also exploit the fact that both exponential map and parallel transportation have a clean closed form in the Lorentz model.

Lorentz model \mathbb{H}^n (Figure 2(a)) can be represented as a set of points $\mathbf{z} \in \mathbb{R}^{n+1}$ with $z_0 > 0$ such that its Lorentzian product (negative Minkowski bilinear form)

$$\langle \mathbf{z}, \mathbf{z}' \rangle_{\mathcal{L}} = -z_0 z'_0 + \sum_{i=1}^n z_i z'_i,$$

with itself is -1 . That is,

$$\mathbb{H}^n = \{ \mathbf{z} \in \mathbb{R}^{n+1} : \langle \mathbf{z}, \mathbf{z} \rangle_{\mathcal{L}} = -1, z_0 > 0 \}. \quad (1)$$

Lorentzian inner product also functions as the metric tensor on hyperbolic space. We will refer to the one-hot vector $\boldsymbol{\mu}_0 = [1, 0, 0, \dots, 0] \in \mathbb{H}^n \subset \mathbb{R}^{n+1}$ as the *origin* of the hyperbolic space. Also, the distance between two points \mathbf{z}, \mathbf{z}' on \mathbb{H}^n is given by $d_{\ell}(\mathbf{z}, \mathbf{z}') = \operatorname{arccosh}(-\langle \mathbf{z}, \mathbf{z}' \rangle_{\mathcal{L}})$, which is also the length of the geodesic that connects \mathbf{z} and \mathbf{z}' .

2.2. Parallel Transport and Exponential Map

The rough explanation of our strategy for the construction of hyperbolic wrapped distribution $\mathcal{G}(\boldsymbol{\mu}, \Sigma)$ with $\boldsymbol{\mu} \in \mathbb{H}^n$ and a positive definite matrix Σ is as follows. We (1) sample a vector from $\mathcal{N}(\mathbf{0}, \Sigma)$, (2) transport the vector from $\boldsymbol{\mu}_0$ to $\boldsymbol{\mu}$ along the geodesic, and (3) project the vector onto the surface. To formalize this sequence of operations, we need to define the tangent space on hyperbolic space as well as the way to transport the tangent space and the way to project a vector in the tangent space to the surface. The transportation of the tangent vector requires *parallel transport*, and the projection of the tangent vector to the surface requires the definition of *exponential map*.

Tangent space of hyperbolic space

Let us use $T_{\boldsymbol{\mu}}\mathbb{H}^n$ to denote the tangent space of \mathbb{H}^n at $\boldsymbol{\mu}$ (Figure 2(a)). Representing $T_{\boldsymbol{\mu}}\mathbb{H}^n$ as a set of vectors in the same ambient space \mathbb{R}^{n+1} into which \mathbb{H}^n is embedded, $T_{\boldsymbol{\mu}}\mathbb{H}^n$ can be characterized as the set of points satisfying the orthogonality relation with respect to the Lorentzian




Figure 2: (a) One-dimensional Lorentz model \mathbb{H}^1 (red) and its tangent space $T_{\mu}\mathbb{H}^1$ (blue). (b) Parallel transport carries $v \in T_{\mu_0}$ (green) to $u \in T_{\mu}$ (blue) while preserving $\|\cdot\|_{\mathcal{L}}$. (c) Exponential map projects the $u \in T_{\mu}$ (blue) to $z \in \mathbb{H}^1$ (red). The distance between μ and $\exp_{\mu}(u)$ which is measured on the surface of \mathbb{H}^n coincides with $\|u\|_{\mathcal{L}}$.

product:

$$T_{\mu}\mathbb{H}^n := \{u: \langle u, \mu \rangle_{\mathcal{L}} = 0\}. \quad (2)$$

$T_{\mu}\mathbb{H}^n$ set can be literally thought of as the tangent space of the forward hyperboloid sheet at μ . Note that $T_{\mu_0}\mathbb{H}^n$ consists of $v \in \mathbb{R}^{n+1}$ with $v_0 = 0$, and $\|v\|_{\mathcal{L}} := \sqrt{\langle v, v \rangle_{\mathcal{L}}} = \|v\|_2$.

Parallel transport and inverse parallel transport

Next, for an arbitrary pair of point $\mu, \nu \in \mathbb{H}^n$, the parallel transport from ν to μ is defined as a map $\text{PT}_{\nu \rightarrow \mu}$ from $T_{\nu}\mathbb{H}^n$ to $T_{\mu}\mathbb{H}^n$ that carries a vector in $T_{\nu}\mathbb{H}^n$ along the geodesic from ν to μ in a parallel manner without changing its metric tensor. In other words, if PT is the parallel transport on hyperbolic space, then $\langle \text{PT}_{\nu \rightarrow \mu}(v), \text{PT}_{\nu \rightarrow \mu}(v') \rangle_{\mathcal{L}} = \langle v, v' \rangle_{\mathcal{L}}$.

The explicit formula for the parallel transport on the Lorentz model (Figure 2(b)) is given by:

$$\text{PT}_{\nu \rightarrow \mu}(v) = v + \frac{\langle \mu - \alpha\nu, v \rangle_{\mathcal{L}}}{\alpha + 1}(\nu + \mu), \quad (3)$$

where $\alpha = -\langle \nu, \mu \rangle_{\mathcal{L}}$. The inverse parallel transport $\text{PT}_{\nu \rightarrow \mu}^{-1}$ simply carries the vector in $T_{\mu}\mathbb{H}^n$ back to $T_{\nu}\mathbb{H}^n$ along the geodesic. That is,

$$v = \text{PT}_{\nu \rightarrow \mu}^{-1}(u) = \text{PT}_{\mu \rightarrow \nu}(u). \quad (4)$$

Exponential map and inverse exponential map

Finally, we will describe a function that maps a vector in a tangent space to its surface.

According to the basic theory of differential geometry, every $u \in T_{\mu}\mathbb{H}^n$ determines a unique maximal geodesic $\gamma_{\mu}: [0, 1] \rightarrow \mathbb{H}^n$ with $\gamma_{\mu}(0) = \mu$ and $\dot{\gamma}_{\mu}(0) = u$. Exponential map $\exp_{\mu}: T_{\mu}\mathbb{H}^n \rightarrow \mathbb{H}^n$ is a map defined by $\exp_{\mu}(u) = \gamma_{\mu}(1)$, and we can use this map to project a

vector v in $T_{\mu}\mathbb{H}^n$ onto \mathbb{H}^n in a way that the distance from μ to destination of the map coincides with $\|v\|_{\mathcal{L}}$, the metric norm of v . For hyperbolic space, this map (Figure 2(c)) is given by

$$z = \exp_{\mu}(u) = \cosh(\|u\|_{\mathcal{L}})\mu + \sinh(\|u\|_{\mathcal{L}})\frac{u}{\|u\|_{\mathcal{L}}}. \quad (5)$$

As we can confirm with straightforward computation, this exponential map is norm preserving in the sense that $d_{\ell}(\mu, \exp_{\mu}(u)) = \text{arccosh}(-\langle \mu, \exp_{\mu}(u) \rangle_{\mathcal{L}}) = \|u\|_{\mathcal{L}}$. Now, in order to evaluate the density of a point on hyperbolic space, we need to be able to map the point back to the tangent space, on which the distribution is initially defined. We, therefore, need to be able to compute the inverse of the exponential map, which is also called logarithm map, as well.

Solving eq. (5) for u , we can obtain the inverse exponential map as

$$u = \exp_{\mu}^{-1}(z) = \frac{\text{arccosh}(\alpha)}{\sqrt{\alpha^2 - 1}}(z - \alpha\mu), \quad (6)$$

where $\alpha = -\langle \mu, z \rangle_{\mathcal{L}}$. See Appendix A.1 and A.2 for further details.

3. Hyperbolic Wrapped Distribution

3.1. Construction

Finally, we are ready to provide the construction of our hyperbolic wrapped distribution $\mathcal{G}(\mu, \Sigma)$ on hyperbolic space with $\mu \in \mathbb{H}^n$ and positive definite Σ (Figure 3).

In the language of the differential geometry, our strategy can be re-described as follows:

1. Sample a vector \tilde{v} from the Gaussian distribution $\mathcal{N}(\mathbf{0}, \Sigma)$ defined over \mathbb{R}^n .

Algorithm 1 Sampling on hyperbolic space

Input: parameter $\mu \in \mathbb{H}^n, \Sigma$
Output: $z \in \mathbb{H}^n$
Require: $\mu_0 = (1, 0, \dots, 0)^\top \in \mathbb{H}^n$
Sample $\tilde{v} \sim \mathcal{N}(\mathbf{0}, \Sigma) \in \mathbb{R}^n$
 $v = [0, \tilde{v}] \in T_{\mu_0}\mathbb{H}^n$
Move v to $u = \text{PT}_{\mu_0 \rightarrow \mu}(v) \in T_\mu\mathbb{H}^n$ by eq. (3)
Map u to $z = \exp_\mu(u) \in \mathbb{H}^n$ by eq. (5)

2. Interpret \tilde{v} as an element of $T_{\mu_0}\mathbb{H}^n \subset \mathbb{R}^{n+1}$ by rewriting \tilde{v} as $v = [0, \tilde{v}]$.
3. Parallel transport the vector v to $u \in T_\mu\mathbb{H}^n \subset \mathbb{R}^{n+1}$ along the geodesic from μ_0 to μ .
4. Map u to \mathbb{H}^n by \exp_μ .

Algorithm 1 is an algorithmic description of the sampling procedure based on our construction.

3.2. Probability Density Function

Note that both $\text{PT}_{\mu_0 \rightarrow \mu}$ and \exp_μ are differentiable functions that can be evaluated analytically. Thus, by the construction of $\mathcal{G}(\mu, \Sigma)$, we can compute the probability density of $\mathcal{G}(\mu, \Sigma)$ at $z \in \mathbb{H}^n$ using a composition of differentiable functions, $\text{PT}_{\mu_0 \rightarrow \mu}$ and \exp_μ . Let $\text{proj}_\mu := \exp_\mu \circ \text{PT}_{\mu_0 \rightarrow \mu}$.

In general, if \mathbf{X} is a random variable endowed with the probability density function $p(x)$, the log likelihood of $Y = f(\mathbf{X})$ at y can be expressed as

$$\log p(y) = \log p(x) - \log \det \left(\frac{\partial f}{\partial x} \right)$$

where f is an invertible and continuous map. Thus, all we need in order to evaluate the probability density of $\mathcal{G}(\mu, \Sigma)$ at $z = \text{proj}_\mu(v)$ is the way to evaluate $\det(\partial \text{proj}_\mu(v)/\partial v)$:

$$\log p(z) = \log p(v) - \log \det \left(\frac{\partial \text{proj}_\mu(v)}{\partial v} \right). \quad (7)$$

Algorithm 2 is an algorithmic description for the computation of the pdf.

For the implementation of algorithm 1 and algorithm 2, we would need to be able to evaluate not only $\exp_\mu(u)$, $\text{PT}_{\mu_0 \rightarrow \mu}(v)$ and their inverses, but also need to evaluate the determinant. We provide an analytic solution to each one of them below.

Log-determinant

For the evaluation of (7), we need to compute the log determinant of the projection function that maps a vector in the tangent space $T_{\mu_0}(\mathbb{H}^n)$ at origin to the tangent space $T_\mu(\mathbb{H}^n)$ at an arbitrary point μ in the hyperbolic space.

Algorithm 2 Calculate log-pdf

Input: sample $z \in \mathbb{H}^n$, parameter $\mu \in \mathbb{H}^n, \Sigma$
Output: $\log p(z)$
Require: $\mu_0 = (1, 0, \dots, 0)^\top \in \mathbb{H}^n$
Map z to $u = \exp_\mu^{-1}(z) \in T_\mu\mathbb{H}^n$ by eq. (6)
Move u to $v = \text{PT}_{\mu_0 \rightarrow \mu}^{-1}(u) \in T_{\mu_0}\mathbb{H}^n$ by eq. (4)
Calculate $\log p(z)$ by eq. (7)

Appealing to the chain-rule and the property of determinant, we can decompose the expression into two components:

$$\begin{aligned} & \det \left(\frac{\partial \text{proj}_\mu(v)}{\partial v} \right) \\ &= \det \left(\frac{\partial \exp_\mu(u)}{\partial u} \right) \cdot \det \left(\frac{\partial \text{PT}_{\mu_0 \rightarrow \mu}(v)}{\partial v} \right). \quad (8) \end{aligned}$$

We evaluate each piece one by one. First, let us recall that $\partial \exp_\mu(u)/\partial u$ is a map that sends an element in $T_u(T_\mu(\mathbb{H}^n)) = T_\mu(\mathbb{H}^n)$ to an element in $T_v(\mathbb{H}^n)$, where $v = \exp_\mu(u)$. We have a freedom in choosing a basis to evaluate the determinant of this expression. For convenience, let us choose an orthonormal basis of $T_\mu(\mathbb{H}^n)$ that contains $\bar{u} = u/\|u\|_{\mathcal{L}}$:

$$\{\bar{u}, u'_1, u'_2, \dots, u'_{n-1}\}$$

Now, computing the directional derivative of \exp_μ with respect to each basis element, we get

$$d \exp_\mu(\bar{u}) = \sinh(r)\mu + \cosh(r)\bar{u}, \quad (9)$$

$$d \exp_\mu(u'_k) = \frac{\sinh r}{r} u'_k, \quad (10)$$

where $r = \|u\|_{\mathcal{L}}$. Because the directional derivative in the direction of \bar{u} has magnitude 1 and because each component is orthogonal to each other, the norm of the directional derivative is given by

$$\det \left(\frac{\partial \exp_\mu(u)}{\partial u} \right) = \left(\frac{\sinh r}{r} \right)^{n-1}, \quad (11)$$

which yields the value of the desired determinant. Let us next compute the determinant of the derivative of parallel transport. Evaluating the determinant with an orthogonal basis of $T_v(T_{\mu_0}(\mathbb{H}^n))$, we get

$$d \text{PT}_{\mu_0 \rightarrow \mu}(\xi) = \text{PT}_{\mu_0 \rightarrow \mu}(\xi).$$

Because parallel transport is itself a norm preserving map, this means that $\det(\partial \text{PT}_{\mu_0 \rightarrow \mu}(v)/\partial v) = 1$.

All together, we get

$$\det \left(\frac{\partial \text{proj}_\mu(v)}{\partial v} \right) = \left(\frac{\sinh r}{r} \right)^{n-1}.$$

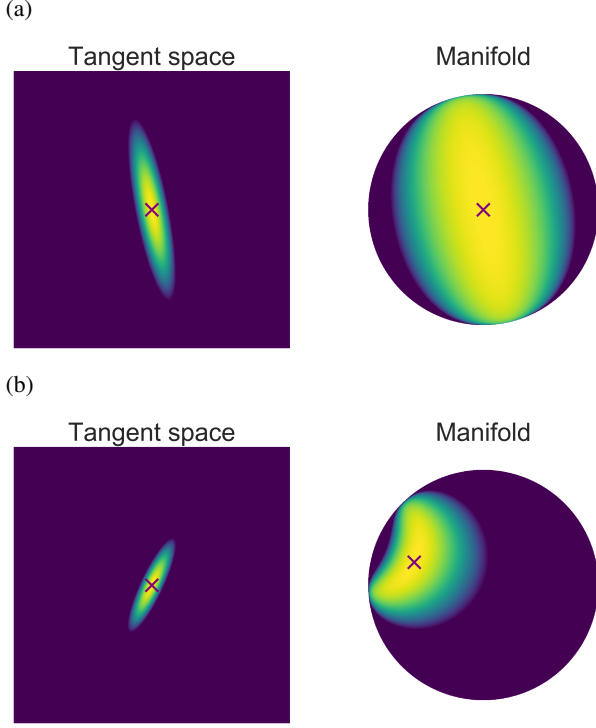


Figure 3: The heatmaps of log-likelihood of the hyperbolic wrapped distribution with various μ and Σ . We designate the origin of hyperbolic space by the \times mark. See Appendix B for further details.

We will provide the details of the computation in Appendix A.3. and A.4. The whole evaluation of the log determinant can be computed in $\mathcal{O}(n)$. Figure 3 shows example densities on \mathbb{H}^2 .

Since the metric at the tangent space coincides with the Euclidean metric, we can produce various types of distributions on hyperbolic space by applying our construction strategy to other distributions defined on Euclidean space, such as Laplace and Cauchy distribution.

4. Applications of $\mathcal{G}(\mu, \Sigma)$

4.1. Hyperbolic Variational Autoencoder

As an application of hyperbolic wrapped distribution $\mathcal{G}(\mu, \Sigma)$, we will introduce *hyperbolic variational autoencoder* (Hyperbolic VAE), a variant of the variational autoencoder (VAE) (Kingma & Welling, 2014; Rezende et al., 2014) in which the latent variables are defined on hyperbolic space. Given dataset $\mathcal{D} = \{\mathbf{x}^{(i)}\}_{i=1}^N$, the method of variational autoencoder aims to train a decoder model $p_{\theta}(\mathbf{x}|\mathbf{z})$ that can create from $p_{\theta}(\mathbf{z})$ a dataset that resembles \mathcal{D} . The decoder model is trained together with the encoder model $q_{\phi}(\mathbf{z}|\mathbf{x})$ by maximizing the sum of evidence lower

bound (ELBO) that is defined for each $\mathbf{x}^{(i)}$;

$$\begin{aligned} \log p_{\theta}(\mathbf{x}^{(i)}) &\geq \mathcal{L}(\theta, \phi; \mathbf{x}^{(i)}) = \\ &\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})} [\log p_{\theta}(\mathbf{x}^{(i)}|\mathbf{z})] - D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})||p_{\theta}(\mathbf{z})), \end{aligned} \quad (12)$$

where $q_{\phi}(\mathbf{z}|\mathbf{x}^{(i)})$ is the variational posterior distribution. In classic VAE, the choice of the prior p_{θ} is the standard normal, and the posterior distribution is also variationally approximated by a Gaussian. Hyperbolic VAE is a simple modification of the classic VAE in which $p_{\theta} = \mathcal{G}(\mu_0, I)$ and $q_{\phi} = \mathcal{G}(\mu, \Sigma)$. The model of μ and Σ is often referred to as an encoder. This parametric formulation of q_{ϕ} is called reparametrization trick, and it enables the evaluation of the gradient of the objective function with respect to the network parameters. To compare our method against, we used β -VAE (Higgins et al., 2017), a variant of VAE that applies a scalar weight β to the KL term in the objective function.

In Hyperbolic VAE, we assure that output μ of the encoder is in \mathbb{H}^n by applying \exp_{μ_0} to the final layer of the encoder. That is, if \mathbf{h} is the output, we can simply use

$$\mu = \exp_{\mu_0}(\mathbf{h}) = \left(\cosh(\|\mathbf{h}\|_2), \sinh(\|\mathbf{h}\|_2) \frac{\mathbf{h}}{\|\mathbf{h}\|_2} \right)^{\top}.$$

As stated in the previous sections, our distribution $\mathcal{G}(\mu, \Sigma)$ allows us to evaluate the ELBO exactly and to take the gradient of the objective function. In a way, our distribution of the variational posterior is an hyperbolic-analog of the reparametrization trick.

4.2. Word Embedding

We can use our hyperbolic wrapped distribution \mathcal{G} for probabilistic word embedding. The work of Vilnis & McCallum (2015) attempted to extract the linguistic and contextual properties of words in a dictionary by embedding every word and every context to a Gaussian distribution defined on Euclidean space. We may extend their work by changing the destination of the map to the family of \mathcal{G} . Let us write $a \sim b$ to convey that there is a link between words a and b , and let us use q_s to designate the distribution to be assigned to the word s . The objective function used in Vilnis & McCallum (2015) is given by

$$\mathcal{L}(\theta) = \mathbb{E}_{(s \sim t, s \not\sim t')} [\max(0, m + E(s, t) - E(s, t'))],$$

where $E(s, t)$ represents the measure of similarity between s and t evaluated with $D_{\text{KL}}(q_s||q_t)$. In the original work, q_s and q_t were chosen to be a Gaussian distribution. We can incorporate hyperbolic geometry into this idea by choosing $q_s = \mathcal{G}(\mu(s), \Sigma(s))$.

5. Related Work

The construction of hyperbolic distribution based on projection is not entirely new on its own. For example, CCM-AAE (Grattarola et al., 2018) uses a prior distribution on hyperbolic space centered at origin by projecting a distribution constructed on the tangent space. Wrapped normal distribution (on sphere) also is a creation of similar philosophy. Still yet, as mentioned in the introduction, most studies to date that use hyperbolic space consider only deterministic mappings (Nickel & Kiela, 2017; 2018; Ganea et al., 2018a;b; Gülçehre et al., 2019).

Very recently, Ovinnikov (2019) and Mathieu et al. (2019) proposed an extension of Gaussian distribution on Poincaré ball model and its application to VAEs. Ovinnikov (2019) used Wasserstein Maximum Mean Discrepancy (Gretton et al., 2012), and Mathieu et al. (2019) used Monte Carlo approximation of ELBO for training their models. By construction, however, their method can only create isotropic distribution on Riemannian manifold and they also have to use rejection sampling in their method. Meanwhile, our method can wrap $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ onto \mathbb{H}^n for arbitrary choice of Σ , and we do not have to use rejection sampling. Our distribution $\mathcal{G}(\boldsymbol{\mu}, \Sigma)$ can be defined for any $\boldsymbol{\mu}$ in \mathbb{H}^n and any positive definite matrix $\Sigma \in \mathbb{R}_+^{n \times n}$.

For word embedding, several deterministic methods have been proposed to date, including the celebrated Word2Vec (Mikolov et al., 2013). The aforementioned Nickel & Kiela (2017) uses deterministic hyperbolic embedding to exploit the hierarchical relationships among words. The probabilistic word embedding was first proposed by Vilnis & McCallum (2015). As stated in the method section, their method maps each word to a Gaussian distribution on Euclidean space. Their work suggests the importance of investigating the uncertainty of word embedding. In the field of representation learning of word vectors, our work is the first in using hyperbolic probability distribution for word embedding.

On the other hand, the idea to use a noninformative, non-Gaussian prior in VAE is not new. For example, Davidson et al. (2018) proposes the use of von Mises-Fisher prior, and Rolfe (2017); Jang et al. (2017) use discrete distributions as their prior. With the method of Normalizing flow (Rezende & Mohamed, 2015), one can construct even more complex priors as well (Kingma et al., 2016). The appropriate choice of the prior shall depend on the type of dataset. As we will show in the experiment section, our distribution is well suited to the dataset with underlying tree structures. Another choice of the VAE prior that specializes in such dataset has been proposed by Vikram et al. (2018). For the sampling, they use time-marginalized coalescent, a model that samples a random tree structure by a stochastic process. Theoretically, their method can be used in combina-

Model	Correlation	Correlation w/ noise	
Vanilla	$\beta = 0.1$	0.665 \pm .006	0.470 \pm .018
	$\beta = 1.0$	0.644 \pm .007	0.550 \pm .012
	$\beta = 2.0$	0.644 \pm .011	0.537 \pm .012
	$\beta = 3.0$	0.638 \pm .004	0.501 \pm .044
	$\beta = 4.0$	0.217 \pm .143	0.002 \pm .042
Hyperbolic	0.768\pm.003	0.590\pm.018	

Table 1: Results of tree embedding experiments for the Hyperbolic VAE and Vanilla VAEs trained with different weight constants for the KL term. We calculated the mean and the ± 1 SD with five different experiments.

tion with our approach by replacing their Gaussian random walk with a hyperbolic random walk.

6. Experiments

6.1. Synthetic Binary Tree

We trained Hyperbolic VAE for an artificial dataset constructed from a binary tree of depth $d = 8$. To construct the dataset, we first obtained a binary representation for each node in the tree so that the Hamming distance between any pair of nodes is the same as the distance on the graph representation of the tree (Figure 1(a)). Let us call the set of binaries obtained this way by A_0 . We then generated a set of binaries, A , by randomly flipping each coordinate value of A_0 with probability $\epsilon = 0.1$. The binary set A was then embedded into \mathbb{R}^d by mapping $a_1 a_2 \dots a_d$ to $[a_1, a_2, \dots, a_d]$. We used an Multi Layer Perceptron (MLP) of depth 3 and 100 hidden variables at each layer for both encoder and decoder. For activation function we used \tanh .

Table 1 summarizes the quantitative comparison of Vanilla VAE against our Hyperbolic VAE. For each pair of points in the tree, we computed their Hamming distance as well as their distance in the latent space of VAE. That is, we used hyperbolic distance for Hyperbolic VAE, and used Euclidean distance for Vanilla VAE. We used the strength of correlation between the Hamming distances and the distances in the latent space as a measure of performance. Hyperbolic VAE was performing better both on the original tree and on the artificial dataset generated from the tree. Vanilla VAE performed the best with $\beta = 2.0$, and collapsed with $\beta = 3.0$. The difference between Vanilla VAE and Hyperbolic VAE can be observed with much more clarity using the 2-dimensional visualization of the generated dataset on Poincaré Ball (See Figure 1 and Appendix C.1). The red points are the embeddings of A_0 , and the blue points are the embeddings of all other points in A . The pink \times mark designates the origin of hyperbolic space. For the visualization, we used the canonical diffeomorphism between the Lorenz model and the Poincaré ball model.

n	Vanilla VAE	Hyperbolic VAE
2	$-140.45 \pm .47$	-138.61 ± 0.45
5	$-105.78 \pm .51$	-105.38 ± 0.61
10	$-86.25 \pm .52$	-86.40 ± 0.28
20	$-77.89 \pm .36$	-79.23 ± 0.20

Table 2: Quantitative comparison of Hyperbolic VAE against Vanilla VAE on the MNIST dataset in terms of log-likelihood (LL) for several values of latent space dimension n . LL was computed using 500 samples of latent variables. We calculated the mean and the ± 1 SD with five different experiments.

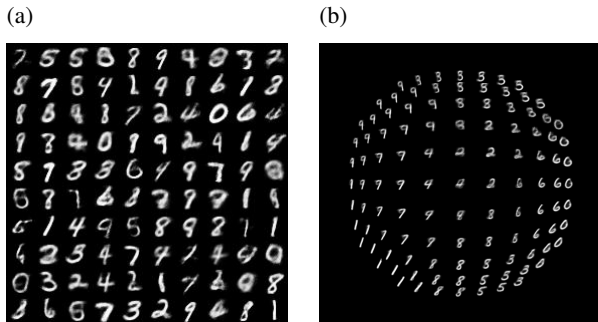


Figure 4: (a) Samples generated from the Hyperbolic VAE trained on MNIST with latent dimension $n = 5$. (b): Interpolation of the MNIST dataset produced by the Hyperbolic VAE with latent dimension $n = 2$, represented on the Poincaré ball.

6.2. MNIST

We applied Hyperbolic VAE to a binarized version of MNIST. We used an MLP of depth 3 and 500 hidden units at each layer for both the encoder and the decoder. Table 2 shows the quantitative results of the experiments. Log-likelihood was approximated with an empirical integration of the Bayesian predictor with respect to the latent variables (Burda et al., 2016). Our method outperformed Vanilla VAE with small latent dimension. We also evaluated our method in terms of ELBO. Please see Appendix C.2 for further results. Figure 4(a) are the samples of the Hyperbolic VAE that was trained with 5-dimensional latent variables, and Figure 4(b) are the Poincaré Ball representations of the interpolations produced on \mathbb{H}^2 by the Hyperbolic VAE that was trained with 2-dimensional latent variables.

6.3. Atari 2600 Breakout

In reinforcement learning, the number of possible state-action trajectories grows exponentially with the time horizon. We may say that these trajectories often have a tree-like hierarchical structure that starts from the initial states. We applied our Hyperbolic VAE to a set of trajectories that

were explored by a trained policy during multiple episodes of Breakout in Atari 2600. To collect the trajectories, we used a pretrained Deep Q-Network (Mnih et al., 2015), and used epsilon-greedy with $\epsilon = 0.1$. We amassed a set of trajectories whose total length is 100,000, of which we used 80,000 as the training set, 10,000 as the validation set, and 10,000 as the test set. Each frame in the dataset was gray-scaled and resized to 80×80 . The images in the Figure 5 are samples from the dataset. We used a DCGAN-based architecture (Radford et al., 2016) with latent space dimension $n = 20$. Please see Appendix D for more details.




Figure 5: Examples of the observed screens in Atari 2600 Breakout.

The Figure 6 is a visualization of our results. The top three rows are the samples from Vanilla VAE, and the bottom three rows are the samples from Hyperbolic VAE. Each row consists of samples generated from latent variables of the form $a\tilde{v}/\|\tilde{v}\|_2$ with positive scalar a in range $[1, 10]$. Samples in each row are listed in increasing order of a . For Vanilla VAE, we used $\mathcal{N}(\mathbf{0}, I)$ as the prior. For Hyperbolic VAE, we used $\mathcal{G}(\mu_0, I)$ as the prior. We can see that the number of blocks decreases gradually and consistently in each row for Hyperbolic VAE. Please see Appendix C.3 for more details and more visualizations.

In Breakout, the number of blocks is always finite, and blocks are located only in a specific region. Let's refer to this specific region as R . In order to evaluate each model-output based on the number of blocks, we binarized each pixel in each output based on a prescribed luminance threshold and measured the proportion of the pixels with pixel value 1 in the region R . For each generated image, we used this proportion as the measure of the number blocks contained in the image.

Figure 7 shows the estimated proportions of remaining blocks for Vanilla and Hyperbolic VAEs with different norm of \tilde{v} . For Vanilla VAE, samples generated from \tilde{v} with its norm as large as $\|\tilde{v}\|_2 = 200$ contained considerable amount of blocks. On the other hand, the number of blocks contained in a sample generated by Hyperbolic VAE decreased more consistently with the norm of $\|\tilde{v}\|_2$. This fact suggests that the cumulative reward up to a given state can be approximated well by the norm of Hyperbolic VAE's latent representation. To validate this, we computed latent representation for each state in the test set and measured its correlation with the cumulative reward. The cor-




Figure 6: Samples from Vanilla and Hyperbolic VAEs trained on Atari 2600 Breakout screens. Each row was generated by sweeping the norm of \tilde{v} from 1.0 to 10.0 in a log-scale.

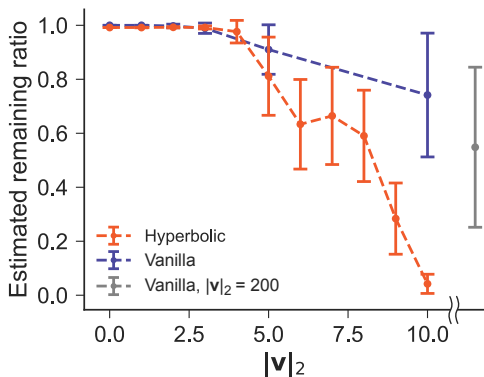


Figure 7: Estimated proportions of remaining blocks for Vanilla and Hyperbolic VAEs trained on Atari 2600 Breakout screens as they vary with the norm of latent variables sampled from a prior.

relation was 0.846 for the Hyperbolic VAE. For the Vanilla VAE, the correlation was 0.712. We emphasize that no information regarding the reward was used during the training of both Vanilla and Hyperbolic VAEs.

6.4. Word Embeddings

Lastly, we applied hyperbolic wrapped distribution to word embedding problem. We trained probabilistic word embedding models with WordNet nouns dataset (Miller, 1998) and evaluated the reconstruction performance of them (Table 3). We followed the procedure of Poincaré embedding (Nickel & Kiela, 2017) and initialized all embeddings in the neighborhood of the origin. In particular, we initialized each weight in the first linear part of the embedding by $\mathcal{N}(0, 0.01)$. We treated the first 50 epochs as a burn-in phase and reduced the learning rate by a factor of 40 after the burn-in phase.

In Table 3, ‘Euclid’ refers to the word embedding with

Gaussian distribution on Euclidean space (Vilnis & McCallum, 2015), and ‘Hyperbolic’ refers to our proposed method based on hyperbolic wrapped distribution. Our hyperbolic model performed better than Vilnis’ Euclidean counterpart when the latent space is low dimensional. We used diagonal variance for both models above. Please see Appendix C.4 for the full results. We also performed the same experiment with unit variance. The performance difference with small latent dimension was much more remarkable when we use unit variance.

n	Euclid		Hyperbolic	
	MAP	Rank	MAP	Rank
5	0.296±.006	25.09±.80	0.506±.017	20.55±1.34
10	0.778±.007	4.70±.05	0.795±.007	5.07±.12
20	0.894±.002	2.23±.03	0.897±.005	2.54±.20
50	0.942±.003	1.51±.04	0.975±.001	1.19±.01
100	0.953±.002	1.34±.02	0.978±.002	1.15±.01

Table 3: Experimental results of the reconstruction performance on the transitive closure of the WordNet noun hierarchy for several latent space dimension n . We calculated the mean and the ± 1 SD with three different experiments.

7. Conclusion

In this paper, we proposed a novel parametrization for the density of hyperbolic wrapped distribution that can both be differentiated and evaluated analytically. Our experimental results on hyperbolic word embedding and hyperbolic VAE suggest that there is much more room left for the application of hyperbolic space. Our parametrization enables gradient-based training of probabilistic models defined on hyperbolic space and opens the door to the investigation of complex models on hyperbolic space that could not have been explored before.

Acknowledgements

We would like to deeply grateful to Masaki Watanabe for profound discussions and comments to our entire study. We would like to thank Tomohiro Hayase, and Kenta Oono for helpful discussions. We also thank Takeru Miyato and Sosuke Kobayashi for insightful reviews on the paper. This paper is based on results obtained from Nagano’s internship at Preferred Networks, Inc. This work is partially supported by Masason Foundation.

References

- Burda, Y., Grosse, R. B., and Salakhutdinov, R. Importance weighted autoencoders. In *Proceedings of the 4th International Conference on Learning Representations*, 2016.
- Davidson, T. R., Falorsi, L., De Cao, N., Kipf, T., and Tomczak, J. M. Hyperspherical variational auto-encoders. *34th Conference on Uncertainty in Artificial Intelligence*, 2018.
- Ganea, O., Bécigneul, G., and Hofmann, T. Hyperbolic entailment cones for learning hierarchical embeddings. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1632–1641, 2018a.
- Ganea, O., Bécigneul, G., and Hofmann, T. In *Advances in Neural Information Processing Systems 31*, pp. 5350–5360, 2018b.
- Grattarola, D., Livi, L., and Alippi, C. Adversarial autoencoders with constant-curvature latent manifolds. *CoRR*, abs/1812.04314, 2018.
- Gretton, A., Borgwardt, K. M., Rasch, M. J., Schölkopf, B., and Smola, A. A kernel two-sample test. *The Journal of Machine Learning Research*, 13:723–773, March 2012.
- Gülçehre, Ç., Denil, M., Malinowski, M., Razavi, A., Pascanu, R., Hermann, K. M., Battaglia, P., Bapst, V., Raposo, D., Santoro, A., and de Freitas, N. Hyperbolic attention networks. In *Proceedings of the 7th International Conference on Learning Representations*, 2019.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. β -vae: Learning basic visual concepts with a constrained variational framework. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- Jang, E., Gu, S., and Poole, B. Categorical reparameterization with gumbel-softmax. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations*, 2014.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. Improved variational inference with inverse autoregressive flow. In *Advances in Neural Information Processing Systems 29*, pp. 4743–4751. 2016.
- Mathieu, E., Lan, C. L., Maddison, C. J., Tomioka, R., and Teh, Y. W. Hierarchical representations with poincaré variational auto-encoders. *CoRR*, abs/1901.06033, 2019.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems 26*, pp. 3111–3119. 2013.
- Miller, G. *WordNet: An electronic lexical database*. MIT press, 1998.
- Mnih, V., Kavukcuoglu, K., Silver, D., Rusu, A. A., Veness, J., Bellemare, M. G., Graves, A., Riedmiller, M., Fidjeland, A. K., Ostrovski, G., et al. Human-level control through deep reinforcement learning. *Nature*, 518 (7540):529, 2015.
- Nickel, M. and Kiela, D. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems 30*, pp. 6338–6347. 2017.
- Nickel, M. and Kiela, D. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 3776–3785, 2018.
- Ovinnikov, I. Poincaré wasserstein autoencoder. *CoRR*, abs/1901.01427, 2019.
- Radford, A., Metz, L., and Chintala, S. Unsupervised representation learning with deep convolutional generative adversarial networks. In *Proceedings of the 4th International Conference on Learning Representations*, 2016.
- Rezende, D. J. and Mohamed, S. Variational inference with normalizing flows. In *Proceedings of the 32nd International Conference on Machine Learning*, pp. 1530–1538, 2015.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32, pp. 1278–1286, 2014.

- Rolfe, J. T. Discrete variational autoencoders. In *Proceedings of the 5th International Conference on Learning Representations*, 2017.
- Sala, F., De Sa, C., Gu, A., and Re, C. Representation tradeoffs for hyperbolic embeddings. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80, pp. 4460–4469, 2018.
- Sarkar, R. Low distortion delaunay embedding of trees in hyperbolic plane. In *Graph Drawing*, pp. 355–366, 2012.
- Vikram, S., Hoffman, M. D., and Johnson, M. J. The loracs prior for vaes: Letting the trees speak for the data. *CoRR*, abs/1810.06891, 2018.
- Vilnis, L. and McCallum, A. Word representations via gaussian embedding. In *Proceedings of the 3rd International Conference on Learning Representations*, 2015.