

000
001
002
003

INTRINSIC LORENTZ NEURAL NETWORK

004 **Anonymous authors**
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053
ABSTRACT

Real-world data frequently exhibit latent hierarchical structures, which can be naturally represented by hyperbolic geometry. Although recent hyperbolic neural networks have demonstrated promising results, many existing architectures remain partially intrinsic, mixing Euclidean operations with hyperbolic ones or relying on extrinsic parameterizations. To address it, we propose the *Intrinsic Lorentz Neural Network* (ILNN), a fully intrinsic hyperbolic architecture that conducts all computations within the Lorentz model. At its core, the network introduces a novel *point-to-hyperplane* fully connected layer (FC), replacing traditional Euclidean affine logits with closed-form hyperbolic distances from features to learned Lorentz hyperplanes, thereby ensuring that the resulting geometric decision functions respect the inherent curvature. Around this fundamental layer, we design intrinsic modules: GyroLBN, a Lorentz batch normalization that couples gyro-centering with gyro-scaling, consistently outperforming both LBN and GyroBN while reducing training time. We additionally proposed a gyro-additive bias for the FC output, a Lorentz patch-concatenation operator that aligns the expected log-radius across feature blocks via a digamma-based scale, and a Lorentz dropout layer. Extensive experiments conducted on CIFAR-10/100 and two genomic benchmarks (TEB and GUE) illustrate that ILNN achieves state-of-the-art performance and computational cost among hyperbolic models and consistently surpasses strong Euclidean baselines.

1 INTRODUCTION

Hierarchical structures exist across a wide range of machine learning applications, including computer vision (Khrulkov et al., 2020; Ghadimi Atigh et al., 2022; Bdeir et al., 2024; Pal et al., 2025), natural language processing (Ganea et al., 2018; Tifrea et al., 2019; Yang et al., 2024a), knowledge-graph reasoning (Nickel & Kiela, 2017; Balazevic et al., 2019; Welz et al., 2025), graph learning (Chami et al., 2019; Yang et al., 2022; Li et al., 2024), and genomics tasks (Khan et al., 2025; Zhou & Sharpee, 2021; Yang et al., 2025). While Euclidean neural networks often incur high distortion or require excessive dimensionality when embedding hierarchical or scale-free structures (Nickel & Kiela, 2018), hyperbolic geometry, characterized by constant negative curvature, offers exponential representational capacity and enables more compact embeddings (Ganea et al., 2018; Shimizu et al., 2020). Consequently, *hyperbolic neural networks* (HNNs) have demonstrated notable success in applications with hierarchical structure (Ganea et al., 2018; Chami et al., 2019; Shimizu et al., 2020; Gulcehre et al., 2019; Li et al., 2024; Fan et al., 2024; Leng et al., 2024; Bdeir et al., 2024; Pal et al., 2025; Chen et al., 2025a; Nguyen et al., 2025; He et al., 2025).

Early hyperbolic neural networks predominantly operated in the Poincaré model because its gyrovector formalism makes many neural primitives straightforward to define (Ganea et al., 2018; Shimizu et al., 2020). However, the unit-ball constraint and boundary saturation render the Poincaré ball more susceptible to numerical instabilities than the Lorentz model, motivating a recent shift toward Lorentz neural networks (LNNs) that exhibit improved optimization stability (Bdeir et al., 2024). Despite this trend leading to the emergence of excellent work (He et al., 2024; Fan et al., 2024; Liang et al., 2024), recent advanced LNNs (Bdeir et al., 2024; Khan et al., 2025) architectures remain only partially intrinsic: they mix Euclidean affine transformations with manifold operations or rely on extrinsic parameterizations that compromise geometric consistency. For example, the Lorentz fully connected layer (LFC) (Bdeir et al., 2024) that applies Euclidean mappings to Lorentz vector, then extract the space part from mapping output and accordingly calculate the time part to form the

054 output Lorentz vector. Normalization further illustrates this tension. Lorentz batch normalization
 055 (LBN) (Bdeir et al., 2024) recenter features but either ignore gyro-variance, while GyroBN (Chen
 056 et al., 2025a) offers gyrogroup-based control yet can still be computationally heavy due to depending
 057 on Fréchet-type statistics. These compromises limit the representational power and efficiency of
 058 Lorentz neural networks.

059 An effective way to mitigate partially intrinsic designs is the *point-to-hyperplane* formulation, which
 060 has been verified on hyperbolic Poincaré (Shimizu et al., 2020) and matrix manifolds (Nguyen
 061 et al., 2025; 2024). Inspired by this, we introduce ILNN, a fully intrinsic hyperbolic network in
 062 which every operation, parameter, and update is defined inside the Lorentz model. At its core
 063 is a *Point-to-hyperplane Lorentz Fully Connected* (PLFC) layer that replaces Euclidean affine
 064 transformation with closed-form Lorentzian distances to learned hyperplanes, yielding curvature-
 065 aware, margin-interpretable decision functions. Surrounding PLFC, we develop intrinsic components:
 066 *GyroLBN*, a batch normalization that couples gyro-centering with variance-controlled gyro-scaling
 067 and outperforms LBN and GyroBN while reducing wall-clock time. Furthermore, we introduced a
 068 log-radius algin Lorentz patch-concatenation to build the stable CNN module, a gyro-additive bias,
 069 and a Lorentz dropout layer to further improve its performance. By designing intrinsic geometric
 070 operations, ILNN maintains mathematical consistency and geometric interpretability throughout the
 071 network.

072 The main contributions of our work can be summarized as follows:

- 073 1. We propose an *intrinsic hyperbolic neural network* that eliminates reliance on extrinsic Euclidean
 074 operations, thus fully harnessing hyperbolic geometry.
- 075 2. We introduce a novel *point-to-hyperplane Lorentz fully connected layer*(PLFC), which replaces
 076 traditional affine transformations with intrinsic hyperbolic distances, significantly enhancing
 077 representational fidelity.
- 078 3. We introduce a GyroLBN, a Lorentz batch normalization that combines gyro-centering with
 079 gyro-scaling, consistently outperforming both LBN and GyroBN while reducing training time.
- 080 4. Extensive experiments demonstrate the effectiveness of ILNN, achieving state-of-the-art per-
 081 formance on CIFAR-10/100 datasets and genomic benchmarks (TEB and GUE), consistently
 082 outperforming Euclidean and existing hyperbolic counterparts.

083 2 RELATED WORK

084 **Hyperbolic embeddings.** A large body of work demonstrates that negatively curved representations
 085 can efficiently encode hierarchical and scale-free structure across modalities. Foundational results
 086 on hyperbolic embeddings (eg, Poincaré embeddings) show strong benefits for symbolic data with
 087 latent trees (Nickel & Kiela, 2017), and early hyperbolic neural architectures extend core layers and
 088 classifiers to the Poincaré model for language tasks (Ganea et al., 2018). On graphs, hyperbolic
 089 GNNs capture hierarchical neighborhoods and outperform Euclidean counterparts on link prediction
 090 and node classification (Chami et al., 2019); knowledge graphs likewise benefit from multi-relational
 091 hyperbolic embeddings (Balazevic et al., 2019). In vision, hyperbolic image embeddings improve
 092 retrieval and classification under class hierarchies (Khrulkov et al., 2020), and fully hyperbolic
 093 CNNs (HCNN) generalize convolution, normalization, and MLR directly in the Lorentz model,
 094 yielding strong encoder-side gains (Bdeir et al., 2024). Hyperbolic modeling has also been adapted to
 095 genomics, where hyperbolic genome embeddings leverage evolutionary signal to surpass Euclidean
 096 baselines across diverse benchmarks (Khan et al., 2025). In this paper, we target these two application
 097 fronts, *vision* and *genomics*.

098 **Hyperbolic neural networks.** Designing *fully hyperbolic* neural networks means replacing every
 099 Euclidean building block with operations that are intrinsic to a negatively curved manifold, without
 100 shuttling features back and forth between Euclidean and hyperbolic spaces. Early work on the
 101 Poincaré ball showed how to lift common layers and primitives (linear/FC maps, activations, concate-
 102 nation, MLR) using gyrovector-space calculus and Möbius operations (Ganea et al., 2018; Shimizu
 103 et al., 2020), while concurrent efforts formulated end-to-end hyperbolic models beyond simple
 104 projection heads (Chen et al., 2022). Due to the Poincaré ball’s greater susceptibility to numerical
 105 instability, subsequent work tried to generalize network components to the Lorentz model (Gulcehre
 106 et al., 2019; Chami et al., 2019; Fan et al., 2023). Beyond classification heads, attention mecha-
 107 nisms

nisms (Gulcehre et al., 2019), graph convolutional layers (Chami et al., 2019), and fully hyperbolic generative models (Qu & Zou, 2022) were also developed in this setting. A major advance came with HCNN, which introduced Lorentz-native formulations of convolution, batch normalization, and multinomial logistic regression, bringing vision encoders fully into the hyperbolic domain (Chen et al., 2022; Bdeir et al., 2024; Yang et al., 2024b). But it still remains partially intrinsic, mixing Euclidean operations with hyperbolic ones. Our work follows this *intrinsic-first* principle and departs in key aspects: a point-to-hyperplane fully connected layer in the Lorentz model with closed-form h -distances for logits and GyroLBN, a Lorentz batch normalization that couples gyro-centering with gyro-scaling.

Normalization in HNNs. Normalization layers are central to stable and efficient training, yet extending them to curved spaces remains challenging. A general Riemannian batch normalization (RBN) based on the Fréchet mean was first made practical by differentiable solvers, enabling manifold-aware centering and variance control; however, its iterative nature can be slow in practice (Lou et al., 2020). Recently, Bdeir et al. (2024) introduced *Lorentz Batch Normalization* (LBN), which uses the closed-form Lorentzian centroid (Law et al., 2019) for re-centering and a principled tangent-space rescaling. Concurrently, *Gyrogroup Batch Normalization* (GyroBN) generalizes RBN to manifolds with gyro-structure, performing gyro-centering and variance-controlled gyro-scaling that better preserve distributional shape under gyro-operations (Chen et al., 2025a). While LBN provides efficient Lorentzian re-centering and global rescaling, it does not explicitly align batch statistics under gyro-operations; GyroBN performs such alignment but is computationally heavier because it relies on Fréchet means. Motivated by these trade-offs, we propose *GyroLBN*, which aligns batch statistics under gyro-operations within the Lorentz model while retaining the speed and practicality of closed-form re-centering.

3 BACKGROUND

There are five isometric hyperbolic models that one can work with (Cannon et al., 1997). We adopt the Lorentz (hyperboloid) model due to numerical stability (Mishne et al., 2023).

Lorentz model. The Lorentz model embeds hyperbolic space as the upper sheet of a two-sheeted hyperboloid in $(n+1)$ -dimensional Minkowski space. It is defined as

$$\mathbb{L}_K^n = \left\{ x \in \mathbb{R}^{n+1} \mid \langle x, x \rangle_L = \frac{1}{K}, x_0 > 0 \right\}.$$

where $K < 0$ denotes the constant sectional curvature, and $\langle x, y \rangle_L := -x_0 y_0 + \sum_{i=1}^n x_i y_i$ is the Lorentz inner product. Following Ratcliffe (2006), we write $x = (x_0, x_s)$ with *time* coordinate x_0 and *space* coordinates $x_s \in \mathbb{R}^n$. The origin is $o = (K^{-1/2}, 0, \dots, 0)$. The closed-form squared distance is given by $d_{\mathcal{L}}^2(x, y) = \|x - y\|_{\mathcal{L}}^2 = \frac{2}{K} - 2\langle x, y \rangle_{\mathcal{L}}$. Other Riemannian operators are presented in Appendix B, such as exponential and logarithmic maps, and parallel transport.

Gyrovector space. It forms the algebraic foundation for the hyperbolic space, as the vector space for the Euclidean space (Ungar, 2022). A gyrovector space has gyroaddition and scalar gyromultiplication, corresponding to the vector addition and scalar multiplication in Euclidean space. Recently, Chen et al. (2025b) proposed the gyroaddition and gyromultiplication over the Lorentz model:

$$x \oplus y = \text{Exp}_x(\text{PT}_{\bar{0} \rightarrow x}(\text{Log}_{\bar{0}}(y))), \quad \forall x, y \in \mathbb{L}_K^n, \quad (1)$$

$$t \odot x = \text{Exp}_{\bar{0}}(t \text{Log}_{\bar{0}}(x)), \quad \forall t \in \mathbb{R}, \forall x \in \mathbb{L}_K^n. \quad (2)$$

Particularly, the inverse is $\ominus x := (-1) \odot x = [x_t, -x_s]$, which satisfies $x \oplus (\ominus x) = (\ominus x) \oplus x = \bar{0}$. Eqs. 1 and 2 admit closed-form expressions, which are more efficient.

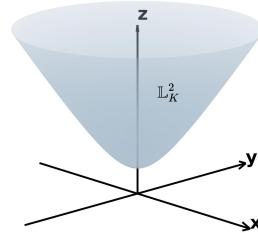


Figure 1: 2-dimensional Lorentz model in the 3-dimensional Minkowski space.

162 4 INTRINSIC LORENTZ NEURAL NETWORK

164 We introduce the Intrinsic Lorentz Neural Network (ILNN), whose components are entirely defined
 165 by the Lorentz geometry. Specifically, we propose (i) a point-to-hyperplane fully connected (FC)
 166 layer, (ii) GyroLBN, a Lorentz batch normalization layer, and (iii) several other intrinsic modules,
 167 including log-radius concatenation, Lorentz activations, and Lorentz dropout.

169 4.1 POINT-TO-HYPERPLANE LORENTZ FULLY-CONNECTED LAYER

171 FC layers perform an affine transformation defined by $\mathbf{y} = \mathbf{A}\mathbf{x} - \mathbf{b}$, which can be element-wisely
 172 expressed as $y_k = \mathbf{a}_k \mathbf{x} - b_k$, where $\mathbf{x}, \mathbf{a}_k \in \mathbb{R}^n$ and $b_k \in \mathbb{R}$ (Shimizu et al., 2020). Geometrically,
 173 this transformation can be interpreted as mapping the input vector \mathbf{x} to an output score y_k , representing
 174 either the coordinate value or the signed distance relative to a hyperplane passing through the origin
 175 and orthogonal to the k -th coordinate axis in the output space \mathbb{R}^m . Motivated by this geometric
 176 interpretation, in this section, we first derive the Lorentz multinomial logistic regression (Lorentz
 177 MLR) to obtain the signed distance, and subsequently present the formulation of the point-to-
 178 hyperplane Lorentz fully connected (PLFC) layer.

179 **Lorentz MLR.** In Euclidean space, the *multinomial logistic regression* (MLR) for class $c \in$
 180 $\{1, \dots, C\}$ can be formulated as the logits of the Euclidean MLR classifier using the distance from
 181 instances to hyperplanes describing the class region (Lebanon & Lafferty, 2004), which can be written
 182 as

$$183 \quad p(y=c \mid \mathbf{x}) \propto \exp(v_c(\mathbf{x})), \quad v_c(\mathbf{x}) = \text{sign}(\langle \mathbf{a}_c, \mathbf{x} - \mathbf{p}_c \rangle) \|\mathbf{a}_c\| d(\mathbf{x}, H_{\mathbf{a}_c, \mathbf{p}_c}), \quad \mathbf{a}_c \in \mathbb{R}^n, \quad (3)$$

185 where $H_{\mathbf{a}_c, \mathbf{p}_c} = \{\mathbf{x} \in \mathbb{R} : \langle \mathbf{a}_c, \mathbf{x} - \mathbf{p}_c \rangle\}$ is the Euclidean hyperplane of class c and $d(\cdot, \cdot)$ denotes
 186 Euclidean distance. Once we have the Lorentz hyperplane and closed-form point-hyperplane distance,
 187 we can extend Euclidean MLR to the Lorentz model.

188 In Lorentz model \mathbb{L}_K^n , following Bdeir et al. (2024), for $\mathbf{p} \in \mathbb{L}_K^n$ and $\mathbf{w} \in \mathcal{T}_{\mathbf{p}}\mathbb{L}_K^n$, the hyperplane
 189 passing through \mathbf{p} and perpendicular to \mathbf{w} is given by $H_{\mathbf{w}, \mathbf{p}} = \{\mathbf{x} \in \mathbb{L}_K^n \mid \langle \mathbf{w}, \mathbf{x} \rangle_{\mathcal{L}} = 0\}$, where \mathbf{w}
 190 should satisfy the condition $\langle \mathbf{w}, \mathbf{w} \rangle > 0$. To eliminate this condition, $\mathbf{w} \in \mathcal{T}_{\mathbf{p}}\mathbb{L}_K^n$ is parameterized by
 191 a vector $\bar{\mathbf{z}} \in \mathcal{T}_{\mathbf{0}}\mathbb{L}_K^n = [0, az/\|\mathbf{z}\|]$, where $a \in \mathbb{R}$ and $\mathbf{z} \in \mathbb{R}^n$. As $\mathbf{w} \in \mathcal{T}_{\mathbf{p}}\mathbb{L}_K^n$, $\bar{\mathbf{z}}$ is parallel transport
 192 to \mathbf{p} , the Lorentz hyperplane defined by

$$193 \quad \tilde{H}_{\mathbf{z}, a} = \{\mathbf{x} \in \mathbb{L}_K^n \mid \cosh(\sqrt{-K}a) \langle \mathbf{z}, \mathbf{x}_s \rangle - \sinh(\sqrt{-K}a) \|\mathbf{z}\| x_t = 0\}, \quad (4)$$

195 where a and \mathbf{z} represent the distance and orientation to the origin, respectively. Due to the $\mathbf{z} \in \mathcal{T}_{\mathbf{0}}\mathbb{L}_K^n$
 196 and the sign-preserving property of parallel transported, this construction naturally satisfies $\langle \mathbf{w}, \mathbf{w} \rangle >$
 197 0. Then for $\mathbf{x} \in \mathbb{L}_K^n$, the distance to a Lorentz hyperplane is given by

$$199 \quad d_{\mathcal{L}}(\mathbf{x}, \tilde{H}_{\mathbf{z}, a}) = \frac{1}{\sqrt{-K}} \left| \sinh^{-1} \left(\sqrt{-K} \frac{\cosh(\sqrt{-K}a) \langle \mathbf{z}, \mathbf{x}_s \rangle - \sinh(\sqrt{-K}a) \|\mathbf{z}\|_2 x_t}{\sqrt{\|\cosh(\sqrt{-K}a) \mathbf{z}\|_2^2 - (\sinh(\sqrt{-K}a) \|\mathbf{z}\|_2)^2}} \right) \right|, \quad (5)$$

202 Substituting Eq. 4 and Eq. 5 into Eq. 3, we obtain that the Lorentz MLR's output logit corresponding to class c
 203 is given by

$$204 \quad v_{\mathbf{z}_c, a_c}(\mathbf{x}) = \frac{1}{\sqrt{-K}} \text{sign}(\alpha_c) \beta_c \left| \sinh^{-1} \left(\sqrt{-K} \frac{\alpha_c}{\beta_c} \right) \right|, \quad (6)$$

$$206 \quad \alpha_c = \cosh(\sqrt{-K}a) \langle \mathbf{z}, \mathbf{x}_s \rangle - \sinh(\sqrt{-K}a),$$

$$208 \quad \beta_c = \sqrt{\|\cosh(\sqrt{-K}a) \mathbf{z}\|_2^2 - (\sinh(\sqrt{-K}a) \|\mathbf{z}\|_2)^2}.$$

211 **Lorentz fully-connected layer (PLFC).** Shimizu et al. (2020) interpreted the Euclidean FC layer as
 212 an operation that transforms the input \mathbf{x} via $v_k(\mathbf{x})$, treating the output \mathbf{y}_k as the signed distance from the
 213 hyperplane $H_{\mathbf{e}^{(k)}, 0}$ passing through the origin and orthogonal to the k -th axis of the output space \mathbb{R}^m , which
 214 can be written as

$$215 \quad d^{\pm}(\mathbf{y}_k, H_{\mathbf{e}^{(k)}, 0}) = v_k(\mathbf{x}), \quad k = 1, \dots, m. \quad (7)$$

216 where $d^\pm(\cdot, \cdot)$ denote the signed distance in Euclidean space. We now collect the above ingredients into an
 217 *intrinsic* FC layer that maps an input $\mathbf{x} \in \mathbb{L}_K^n$ to an output $\mathbf{y} = (y_0, y_s) \in \mathbb{L}_K^m$. Let $\{(\mathbf{z}_k, a_k)\}_{k=1}^m$ be learnable
 218 hyperplane parameters and define $v_k(\mathbf{x}) = v_{\mathbf{z}_k, a_k}(\mathbf{x})$ as in Eq. 6. Matching each $v_k(\mathbf{x})$ to the signed h -distance
 219 from \mathbf{y} to the k -th *coordinate hyperplane* $\tilde{H}_{e^{(k)}, 0} = \{\mathbf{x} = (x_0, x_1, \dots, x_m)^T \in \mathbb{L}_K^m \mid \langle \mathbf{e}^{(k)}, \mathbf{x} \rangle = \mathbf{x}_k = 0, k = 1, 2, \dots, m\}$ fixes the spatial coordinates, while the time coordinate follows from the hyperboloid
 220 constraint.
 221

222 **Theorem 1** (PLFC layer). *Let $\mathbf{x} \in \mathbb{L}_K^n$, $Z = \{\mathbf{z}_k\}_{k=1}^m \subset \mathbb{R}^n$ and $a = \{a_k\}_{k=1}^m \subset \mathbb{R}$. The point-to-hyperplane*
 223 *Lorentz fully connected layer $\text{PLFC}_K: \mathbb{L}_K^n \rightarrow \mathbb{L}_K^m$ is*

$$224 \quad v_k(\mathbf{x}) = v_{\mathbf{z}_k, a_k}(\mathbf{x}), \quad k = 1, \dots, m, \quad (8a)$$

$$225 \quad y_{s,k} = \frac{1}{\sqrt{-K}} \sinh(\sqrt{-K} v_k(\mathbf{x})), \quad (8b)$$

$$226 \quad 227 \quad 228 \quad \mathbf{y} = \left[\sqrt{(-K)^{-1} + \|\mathbf{y}_s\|_2^2}, \mathbf{y}_s \right], \quad (8c)$$

229 where $\mathbf{y}_s = (y_{s,1}, \dots, y_{s,m})^\top$. In the flat-space limit $K \rightarrow 0$, equation 8 reduces to the Euclidean affine map
 230 $\mathbf{y} = A\mathbf{x} + b$ with $A_k = \mathbf{z}_k^\top$ and bias $b = a$.
 231

232 It can be shown that the signed distance from \mathbf{y} to each Lorentz hyperplane passing through the origin and
 233 orthogonal to the k -th coordinate axis is given by $v_k(\mathbf{x})$, as detailed in the Appendix E.1, thereby fulfilling the
 234 properties described above.

235 **Gyro-bias.** A learnable offset $\mathbf{b} \in \mathbb{L}_K^m$ can be added intrinsically via the gyroaddition, $\mathbf{y} \leftarrow \mathbf{y} \oplus \mathbf{b}$, yielding
 236 the final PLFC output.
 237

238 **Discussion.** The previous Lorentz fully connected layer (LFC) used in HCNN was formed by

$$239 \quad \mathbf{y} = [\sqrt{|\phi(W\mathbf{x}, \mathbf{v})|^2 - 1/K}, \phi(W\mathbf{x}, \mathbf{v})],$$

240 where $\mathbf{x} \in \mathcal{L}_K^n$ and $W\mathbf{x}$ is a standard matrix–vector product in \mathbb{R}^n . This $W\mathbf{x}$ is defined using the ambient linear
 241 structure of the Minkowski space, not any operation on the Lorentz manifold so itself only partially intrinsic.
 242 The PLFC depends *only* on Lorentzian operations, avoids tangent-space linearisation, and enjoys closed-form
 243 gradients, making it an efficient and curvature-consistent replacement for Euclidean FC layers in hyperbolic
 244 networks. More detailed discussions are in Appendix D and a theoretical advantage of PLFC over LFC is shown
 245 in Appendix E.2.
 246

4.2 GYROGROUP LORENTZ BATCH NORMALIZATION (GYROLBN)

248 Batch Normalization (BN) facilitates training by normalizing batch statistics. Recently, Bdeir et al. (2024)
 249 proposed LBN, which uses a Lorentzian centroid to efficiently compute the batch mean. However, their approach
 250 fails to normalize sample statistics. Besides, Chen et al. (2025a;b) extended BN into manifolds based on
 251 gyrogroup structures, referred to GyroBN (Chen et al., 2025a). Although it can normalize sample mean/variance
 252 on Lorentz spaces, the involved Fréchet mean could be inefficient. Therefore, we propose *GyroLBN* to combine
 253 the gyrogroup normalization with Lorentzian centroid & radius statistics, retaining GyroBN’s effectiveness
 254 while eliminating computationally expensive Fréchet mean.
 255

256 **Batch centering and dispersion.** Generally, the Fréchet mean must be solved iteratively, which significantly
 257 delays the training speed. Therefore, following Law et al. (2019); Bdeir et al. (2024), with $\mathbf{x}_i \in \mathbb{L}_K^n$,
 258 $\nu_i \geq 0$, $\sum_{i=1}^m \nu_i > 0$ and a batch of features $B = \{\mathbf{x}_i \in \mathbb{L}_K^n\}_{i=1}^m$, is given by, we define the batch mean by
 259 Lorentzian centroid which can be calculated efficiently in closed form (Law et al., 2019):
 260

$$261 \quad \mu_B = \frac{\sum_{i=1}^m \nu_i \mathbf{x}_i}{\sqrt{-K} \left| \left| \sum_{i=1}^m \nu_i \mathbf{x}_i \right| \right|_{\mathcal{L}}}, \quad (9)$$

262 which solves $\min_{\mu_B \in \mathbb{L}_K^n} \sum_{i=1}^m \nu_i d_{\mathcal{L}}^2(\mathbf{x}_i, \mu_B)$. Moreover, the mean is not weighted, which gives $\nu_i = \frac{1}{m}$.
 263

264 As for dispersion, we adopt the Fréchet variance $\sigma^2 \in \mathbb{R}^+$, defined as the expected squared Lorentzian distance
 265 between a point \mathbf{x}_i and the mean μ_B , and given by $\sigma_B^2 = \frac{1}{m} \sum_{i=1}^m d_{\mathcal{L}}^2(\mathbf{x}_i, \mu)$ (Kobler et al., 2022).
 266

267 **Normalization map.** Denote learned scale $\gamma \in \mathbb{R}_{>0}$ (per-channel) and learned bias $\beta \in L_K^d$; For each
 268 sample x , the GyroLBN output is
 269

$$270 \quad \forall i \leq N, \quad \tilde{\mathbf{x}}_i \leftarrow \widehat{\beta \oplus} \left(\overbrace{\frac{\gamma}{\sqrt{\sigma_B + \epsilon}} \odot}^{\text{Scaling}} \left(\overbrace{\ominus \mu_B \oplus \mathbf{x}_i}^{\text{Centering}} \right) \right), \quad \varepsilon > 0. \quad (10)$$

270 where \oplus denotes left gyroaddition Eq. 1, \ominus denotes gyro inverse, and \odot denotes gyro scalar product Eq. 2.
 271 Unless otherwise stated, during inference, we maintain per-channel running statistics by updating the Lorentzian
 272 centroid (mean) and dispersion (variance) with momentum and, at test time, substitute these for batch stats.
 273

274 **Discussion.** GyroLBN unifies the *gyrogroup* normalization paradigm of GyroBN with the *efficient* Lorentz
 275 statistics of LBN, differing from GyroBN only in the choice of statistics by replacing Fréchet-based batch
 276 estimates with closed-form Lorentzian centroid and variance while retaining the same gyrogroup centering and
 277 scaling scheme. Compared to Fréchet variance used in generic Riemannian BN (Lou et al., 2020; Chen et al.,
 278 2022), the mean-radius statistic avoids iterative solvers and proved numerically stable in our vision/genomics
 279 settings. It therefore (i) remains fully intrinsic, (ii) avoids iterative Fréchet solvers (important for large batches
 280 and 2D convolutions), and (iii) integrates naturally with gyro-additive residual/bias layers used in our encoder.
 281 Compared to LBN and Fréchet-based GyroBN (Lou et al., 2020), GyroLBN consistently reduced wall-clock
 282 time while improving accuracy in our settings, as shown in Table 4.
 283

4.3 OTHER LORENTZ MODULES

284 **Log-radius concatenation.** When concatenating N Lorentz patches, each with $(1 + d)$ coordinates (one
 285 time and d spatial), naively stacking the Nd spatial components biases the resulting feature norm toward higher
 286 dimensions: the expected radius grows with Nd , skewing subsequent layers. We introduce a *log-radius*-preserving
 287 concatenation that makes the expected *log* spatial radius invariant to the number of concatenated blocks. Let
 288 $v \in \mathbb{R}^{n_i}$ denote the spatial part of a block and assume its radius factorizes as $\|v\| = \sigma\sqrt{T}$ with $T \sim \chi_{n_i}^2$. Then

$$\mathbb{E}[\log \|v\|] = \log \sigma + \frac{1}{2}(\psi\left(\frac{n_i}{2}\right) + \log 2),$$

290 where ψ is the digamma function. To keep $\mathbb{E}[\log \|v\|]$ constant across dimensions, we scale each block's spatial
 291 part by

$$s(n, n_i) = \exp\left(\frac{1}{2}\left[\psi\left(\frac{n}{2}\right) - \psi\left(\frac{n_i}{2}\right)\right]\right),$$

295 with $n = Nd$ the total post-concat spatial dimension and $n_i = d$ the per-block spatial dimension. Concretely,
 296 given per-window tensors (t_1, \dots, t_N) (times) and (u_1, \dots, u_N) with $u_i \in \mathbb{R}^d$ (spaces), we form the scaled
 297 space $\tilde{u}_i = s u_i$ and recompute a single time coordinate that keeps the output on the hyperboloid:

$$t' = \sqrt{-\frac{1}{K} + s^2 \sum_{i=1}^N (t_i^2 + \frac{1}{K})},$$

301 where $k > 0$ sets the origin time via $t_0 = \sqrt{k}$ (in practice, with $K = -1, k = 1$). The final concatenated vector
 302 is $[t', \tilde{u}_1, \dots, \tilde{u}_N] \in \mathbb{R}^{1+Nd}$. This *log-radius* alignment (i) is parameter-free, (ii) is robust to heavy-tailed radii
 303 because it matches *geometric* means, (iii) preserves the Lorentz constraint by design, and (iv) avoids domination
 304 by any single wide block, improving stability as kernel size or channel count grows.
 305

306 **Lorentz convolutional layer.** We use channel-last feature map representations throughout HCNNs, and
 307 add the Lorentz model's time component as an additional channel dimension, following Bdeir et al. (2024). A
 308 hyperbolic feature map can be define as an ordered set of n -dimensional hyperbolic vectors, where every spatial
 309 position contains a vector that can be combined with its neighbors

310 The convolutional layer can be formulated as a matrix multiplication between a linearized kernel and the
 311 concatenation of values within its receptive field (Shimizu et al., 2020). Then, we extend this definition by
 312 replacing the Euclidean FC and concatenation with our PLFC and *log-radius*-preserving concatenation.

313 Let $\mathbf{x} = \{\mathbf{x}_{h,w} \in \mathbb{L}_K^n\}_{h,w=1}^{H,W}$ be an input hyperbolic feature map and let $\tilde{H} \times \tilde{W}$ denotes the kernel size with
 314 stride δ . For each spatial location (h, w) we gather the patch $\{\mathbf{x}_{h+\delta\tilde{h}, w+\delta\tilde{w}}\}_{\tilde{h}, \tilde{w}=1}^{\tilde{H}, \tilde{W}} \subset \mathbb{L}_K^n$, pad with the origin
 315 $(\sqrt{1/(-K)}, 0, \dots, 0)$ and concatenate these $\tilde{H}\tilde{W}$ vectors by the *log-radius* scheme introduced above. We
 316 defined the Lorentz convolutional layer as

$$\mathbf{y}_{h,w} = \text{PLFC}(\text{LogCat}(\{\mathbf{x}_{h'+\delta\tilde{h}, w'+\delta\tilde{w}} \in \mathbb{L}_K^n\}_{\tilde{h}, \tilde{w}=1}^{\tilde{H}, \tilde{W}})), \quad (11)$$

317 where h' and w' denote the starting position, and LogCat denotes our *log-radius*-preserving concatenation.
 318

319 **Lorentz dropout.** To regularize features without leaving the manifold, we adopt a dropout operator that
 320 acts on Lorentz coordinates and then *reprojects* to the hyperboloid. Concretely, during training, we apply an
 321 elementwise Bernoulli mask to the current representation $x \in \mathbb{L}_K^n$ (probability p of zeroing each entry), yielding
 322 \tilde{x} in ambient \mathbb{R}^{n+1} . Because naive masking can violate the hyperboloid constraint and the positivity of the time

component, we immediately map back via a projection $\text{proj}_{\mathcal{L}}(\tilde{x})$ that restores $\langle x, x \rangle_L = -1/K$ and $x_0 > 0$. At evaluation, the operator is the identity. In practice, we observe that the “mask & project” scheme outperforms the variant “log-exp” scheme that applies a log map to $T_0 \mathbb{L}_K^n$, performs Euclidean dropout, and then returns via the exponential map, because the nonlinearity of \exp_0 couples all coordinates so that masking any tangent component perturbs the entire point after mapping back to the manifold. It is parameter-free, numerically stable, and compatible with subsequent intrinsic layers (e.g., GyroLBN and PLFC), since its output again lies on \mathbb{L}_K^n .

Lorentz activation. Following Bdeir et al. (2024), we define activations *directly* in the Lorentz model by acting on the spatial coordinates only and then recomputing the time coordinate from the hyperboloid constraint. For example, given $x = (x_0, x_s) \in \mathbb{L}_K^n$ and an elementwise nonlinearity activation function ReLU, *Lorentz ReLU* can be defined as

$$\text{L-ReLU}(x) = \left[\sqrt{\frac{1}{-K} + \|\text{ReLU}(x_s)\|_2^2}, \text{ReLU}(x_s) \right].$$

5 EXPERIMENT

We evaluate ILNN on image classification dataset CIFAR-10/100 (Krizhevsky et al., 2009) and genomics (TEB (Khan et al., 2025), GUE (Zhou et al., 2023)) and compare against Euclidean and multiple hyperbolic baselines under matched training recipes. For fairness, each Euclidean backbone is translated to the hyperbolic setting via *one-to-one* module replacement, keeping depth/width, parameter count, and schedule as close as possible. All models are implemented in PyTorch (Paszke et al., 2019) with 32-bit precision. We fix curvature to $K = -1$ and train with Riemannian optimizers from Geoopt (Kochurov et al., 2020): RiemannianSGD for CIFAR and RiemannianAdam for genomics. Unless otherwise stated, results are averaged over five random seeds and reported as mean \pm std; classification uses accuracy on CIFAR-10/100, and MCC on genomics. All experiments are conducted on NVIDIA A100 80GB GPUs. Additional configuration details (batch size, learning rate) appear in Appendix F.

5.1 IMAGE CLASSIFICATION

Experimental setup. Following the evaluation setup of HCNN (Bdeir et al., 2024) and symmetric space (Nguyen et al., 2025), we benchmark ILNN on CIFAR-10 and CIFAR-100. For each method, we instantiate a ResNet-18 backbone (He et al., 2015) in the corresponding geometry: The original ResNet-18 with BN and linear classifier (Euclidean); Euclidean encoder + hyperbolic head (Poincaré or Lorentz); all layers fully to the target manifold (HCNN (Bdeir et al., 2024) and our ILNN). The relative distortion from input manifold to classifier (Bdeir et al., 2024) is comparable across runs ($\delta_{\text{rel}} = 0.26$ for CIFAR-10; 0.23 for CIFAR-100).

Main results. Table 1 summarizes test accuracy as mean \pm sd over five runs. ILNN attains the highest average accuracy on both datasets, 95.36% on CIFAR-10 and 78.41% on CIFAR-100, exceeding the Euclidean ResNet-18 baseline by +0.22 and +0.69 percentage points (pp), respectively. Relative to the strongest prior hyperbolic competitor (HCNN-Lorentz), ILNN also improves by +0.22 pp on CIFAR-10 and +0.34 pp on CIFAR-100. Although ILNN exhibits a slightly larger standard deviation than

HCNN-Lorentz, on CIFAR-10, even the conservative lower bound of ILNN (95.23%) exceeds the upper bound of HCNN (95.22%); on CIFAR-100, the lower bound of ILNN is specifically 78.24%, equal to the upper bound of HCNN. These gains demonstrate that our intrinsic approach preserves the geometry of the Lorentz model without distortion: by combining point-to-hyperplane FC and GyroBN within fully negative-curvature space, ILNN leverages the native manifold structure more effectively than others.

Visualization. We visualize embeddings by mapping network outputs from the Lorentz model to the Poincaré ball and, in parallel, applying the logarithmic map at the origin to view them in the tangent plane; colors indicate output label prediction. In Figure 2 (CIFAR-10), the ILNN embeddings form ten compact clusters, whose areas

Table 1: Classification accuracy (%) of ResNet-18 models. We estimate the mean and standard deviation from five runs. The best performance is highlighted in bold (higher is better).

	CIFAR-10 ($\delta_{\text{rel}} = 0.26$)	CIFAR-100 ($\delta_{\text{rel}} = 0.23$)
Euclidean (He et al., 2015)	95.14 ± 0.12	77.72 ± 0.15
Hybrid Poincaré (Guo et al., 2022)	95.04 ± 0.13	77.19 ± 0.50
Poincaré ResNet (Van Spengler et al., 2023)	94.51 ± 0.15	76.60 ± 0.32
Euclidean-Poincaré-H (Fan et al., 2023)	81.72 ± 7.84	44.35 ± 2.93
Euclidean-Poincaré-G (Ganea et al., 2018)	95.14 ± 0.11	77.78 ± 0.09
Euclidean-Poincaré-B (Nguyen et al., 2025)	95.23 ± 0.08	77.78 ± 0.15
Hybrid Lorentz (Bdeir et al., 2024)	94.98 ± 0.12	78.03 ± 0.21
HCNN Lorentz (Bdeir et al., 2024)	95.14 ± 0.08	78.07 ± 0.17
ILNN (ours)	95.36 ± 0.13	78.41 ± 0.23

378 are visibly smaller than the corresponding clusters produced by HCNN. Only marginal colour bleeding occurs
 379 at the boundaries, indicating that the decision margins learned by the point-to-hyperplane FC induce margins
 380 aligned with the data geometry. The effect becomes even more pronounced on the harder CIFAR-100 task, as
 381 shown in Figure 3. Despite packing 100 classes into the same 2D manifold, ILNN still yields dense colour
 382 "islands" with clear gaps in between, whereas HCNN exhibits overlapping clouds and several mixed-colour
 383 zones. These visual trends are consistent with the quantitative improvements in Table 1.

384 5.2 GENOMIC CLASSIFICATION

385 **Experimental setup.** We evaluate on the most challenging subsets of the **TEB** and **GUE** genomics benchmarks,
 386 **specifically those on which the published HCNN (Khan et al., 2025) does not surpass a Euclidean**
 387 **convolutional baseline.** For every dataset, we instantiate four models that share the same convolutional stem
 388 and classifier width: In our comparisons, the Euclidean CNN uses standard convolutions, Euclidean BN, and
 389 a linear classifier; HCNN-S employs Lorentz modules with a single global curvature K across the network;
 390 HCNN-M uses Lorentz modules with layer-wise curvatures K ; and ILNN is fully intrinsic, featuring GyroBN
 391 and a point-to-hyperplane FC with gyro-bias, while fixing a global curvature $K = -1$. All hyperbolic models
 392 are optimized with RiemannianAdam (lr 10^{-3}); the Euclidean baseline uses Adam with an identical schedule.
 393 We measure performance by the Matthews correlation coefficient (MCC).

394 **Main results.** Table 2 shows that ILNN achieves the best score on every task. On the two TEB pseudogene
 395 sets, it improves over the Euclidean baseline by +9.6 and +13.0 pp, and still exceeds the stronger HCNN-S
 396 by +2.0 and 8.8 pp, respectively, demonstrating clear gains where previous hyperbolic models were already
 397 competitive. The advantage is even more striking on GUE. For Covid-variant classification, HCNN collapses
 398 (MCC 36.7/14.8) much lower than the Euclidean model, whereas ILNN matches and slightly surpasses the
 399 Euclidean score (64.8 vs. 63.6). On promoter-related tasks, ILNN consistently raises the bar: Tata core-promoter
 400 detection jumps from 79.9 (best prior) to 83.9, and the most difficult "all" split goes from 67.2 to 70.7. Across
 401 the board, ILNN tightens the worst-case gap between hyperbolic and Euclidean models and converts several
 402 cases of HCNN under-performance into clear wins. These results confirm that fully intrinsic design choices,
 403 point-to-hyperplane FC, GyroLBN, and other proposed components, translate into tangible accuracy gains on
 404 real-world genomic data.

405 Table 2: Model performance (MCC) on all real-world genomics datasets, averaged over five random
 406 seeds (mean \pm standard deviation). The two highest-scoring models are in bold. * denotes that the
 407 result was reproduced under the same setting.

409 Benchmark	410 Task	411 Dataset	412 Model			
			413 Euclidean 414 CNN	415 Hyperbolic 416 HCNN-S	417 Hyperbolic 418 HCNN-M	419 Hyperbolic 420 ILNN
421 TEB	422 Pseudogenes	423 processed	424 60.66 ± 0.82	425 68.30 ± 0.93	426 65.41 ± 5.54	427 70.26 ± 0.32
		428 unprocessed	429 51.94 ± 2.69	430 56.13 ± 0.56	431 58.36 ± 1.80	432 64.90 ± 0.74
	433 Covid Variant Classification	434 Covid	435 $63.62 \pm 1.34^*$	436 36.71 ± 9.69	437 14.81 ± 0.46	438 64.76 ± 0.54
		439 tata	440 78.26 ± 2.85	441 79.54 ± 1.61	442 79.87 ± 2.50	443 83.90 ± 0.53
		444 notata	445 66.60 ± 1.07	446 66.52 ± 0.28	447 65.95 ± 0.51	448 72.59 ± 0.69
	449 GUE	450 all	451 66.47 ± 0.74	452 65.26 ± 1.11	453 67.16 ± 0.55	454 70.89 ± 0.43
		455 Core Promoter Detection	456 tata	457 78.58 ± 3.39	458 79.74 ± 2.66	459 78.77 ± 0.78
			460 notata	461 90.81 ± 0.51	462 89.86 ± 0.76	463 90.28 ± 0.37
			464 all	465 88.00 ± 0.39	466 87.60 ± 0.51	467 87.93 ± 0.76
	468 Promoter Detection	469 tata	470 78.58 ± 3.39	471 79.74 ± 2.66	472 78.77 ± 0.78	473 83.26 ± 1.90
		474 notata	475 90.81 ± 0.51	476 89.86 ± 0.76	477 90.28 ± 0.37	478 92.48 ± 0.35
		479 all	480 88.00 ± 0.39	481 87.60 ± 0.51	482 87.93 ± 0.76	483 91.34 ± 0.38

421 5.3 GRAPHS

422 **Experimental setup.** To further evaluate the effectiveness of our method, we extend it to three widely-
 423 used graph datasets (AIRPORT, CORA and PUBMED), which exhibit intricate topological and hierarchical
 424 relationships, making them an ideal testbed for evaluating the effectiveness of hyperbolic networks (e.g., HGNN,
 425 HGCN, HGAT, HAN, HNN++ and Hypformer). We choose the Hypformer as baseline and replace the Linear
 426 layer in Hypformer.

427 **Main results.** The quantitative results are summarized in Table 3. Overall, hyperbolic models consistently
 428 outperform their Euclidean counterparts on all three benchmarks, confirming that negatively curved representa-
 429 tions are well-suited for graphs with rich hierarchical structure. Among existing methods, HAN, HNN++, and
 430 SGFormer constitute the strongest baselines, while Hypformer further improves their performance. Building on

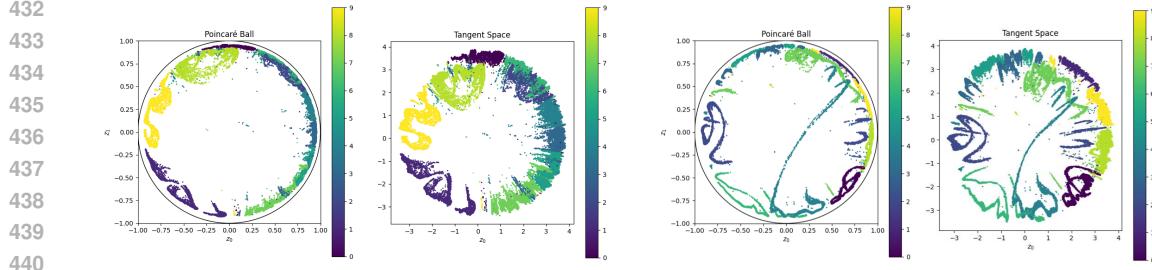


Figure 2: Embedding visualization of CIFAR-10 dataset in Poincaré and Tangent Space. Colors represent labels. HCNN (94.98, left) and ILNN (95.48, right).

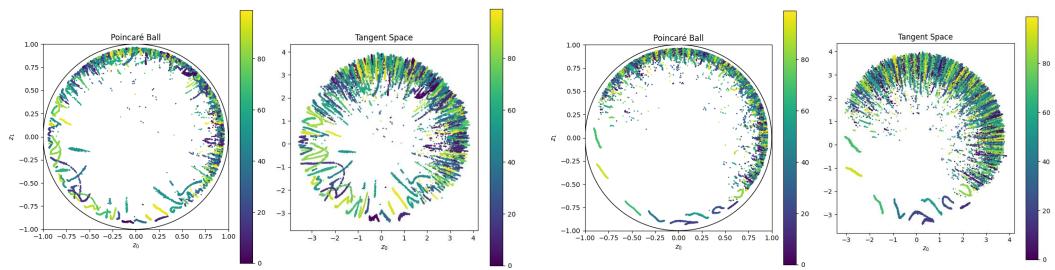


Figure 3: Embedding visualization of CIFAR-100 dataset in Poincaré and Tangent Space. Colors represent labels. HCNN (77.67, left) and ILNN (78.64, right).

this backbone, Hypformer+PLFC achieves the best results on all datasets, reaching $96.03\% \pm 0.34$ on AIRPORT, $85.68\% \pm 0.19$ on CORA, and $82.52\% \pm 0.33$ on PUBMED. Compared with the original Hypformer, this corresponds to absolute gains of $+1.03$, $+0.68$, and $+1.22$ percentage points, respectively. The standard deviations remain small, and the confidence intervals of Hypformer+PLFC are largely separated from those of competing methods on AIRPORT and PUBMED, indicating that the observed gains are statistically stable. These results demonstrate that replacing the Euclidean linear classifier in Hypformer with our hyperbolic point-to-hyperplane PLFC layer provides a simple yet effective way to enhance graph representation learning on graph benchmarks.

5.4 ABLATION STUDY

To better understand the contributions of each architectural component, we conduct an ablation study on our two main innovations: the point-to-hyperplane fully connected head (PLFC) and the Gyrogroup Lorentz Batch Normalization (GyroLBN). We compare them against their Lorentz counterparts (LFC, LBN) as well as GyroBN, measuring both classification accuracy across CIFAR-10 and genomics tasks, and training efficiency to disentangle the effects of the classifier design and the normalization strategy, and to assess whether their combination leads to complementary improvements.

Table 3: Testing results (Accuracy) on AIRPORT, CORA and PUBMED. The best results are in bold, respectively.

Models	AIRPORT	CORA	PUBMED
GCN (Kipf & Welling, 2017)	81.4 ± 0.6	81.3 ± 0.3	78.1 ± 0.2
GAT (Veličković et al., 2017)	81.5 ± 0.3	83.0 ± 0.7	79.0 ± 0.3
SGC (Wu et al., 2019)	82.1 ± 0.5	80.1 ± 0.2	78.7 ± 0.1
HGNN (Liu et al., 2019)	84.7 ± 1.0	77.1 ± 0.8	78.3 ± 1.2
HGCN (Chami et al., 2019)	89.3 ± 1.2	76.5 ± 0.6	78.0 ± 1.0
HGAT (Chami et al., 2019)	89.6 ± 1.0	77.4 ± 0.7	78.3 ± 1.4
GraphFormer (Ying et al., 2021)	88.1 ± 1.2	60.0 ± 0.5	73.3 ± 0.7
GraphTrans (Wu et al., 2021)	94.3 ± 0.6	77.6 ± 0.8	77.5 ± 0.7
GraphGPS (Rampášek et al., 2022)	94.5 ± 0.9	73.0 ± 1.4	72.8 ± 1.4
FPS-T (Cho et al., 2023)	96.0 ± 0.6	82.3 ± 0.7	78.5 ± 0.6
HAN (Gulcehre et al., 2019)	92.9 ± 0.6	83.1 ± 0.5	79.0 ± 0.6
HNN++ (Shimizu et al., 2020)	92.3 ± 0.3	82.8 ± 0.6	79.9 ± 0.4
F-HNN (Chen et al., 2022)	93.0 ± 0.7	81.0 ± 0.7	77.5 ± 0.8
NodeFormer (Wu et al., 2022)	80.2 ± 0.6	82.2 ± 0.9	79.9 ± 1.0
SGFormer (Wu et al., 2023)	92.9 ± 0.5	83.2 ± 0.9	80.0 ± 0.8
Hypformer (Yang et al., 2024b)	95.0 ± 0.5	85.0 ± 0.3	81.3 ± 0.3
Hypformer+PLFC	96.0 ± 0.3	85.7 ± 0.2	82.5 ± 0.3

PLFC vs. LFC. Holding the normalizer fixed, replacing the Lorentz fully connected head with the point-to-hyperplane fully connected head consistently improves accuracy. With GyroLBN, PLFC outperforms LFC on CIFAR-10 (95.36 vs. 95.19,

486
487
488
Table 4: Ablation on PLFC and GyroLBN. Fit time denotes the training time of the CIFAR-10 dataset
487
488
per epoch.

489 490 491 Benchmark	492 493 494 495 Task	496 497 498 Dataset	499 500 Model			
			LFC 501 LBN	503 LFC 504 GyroLBN	505 PLFC 506 GyroBN	507 PLFC 508 GyroLBN
CIFAR-10			95.14 \pm 0.08	95.19 \pm 0.15	95.28 \pm 0.17	95.36\pm0.13
GUE	Core Promoter Detection	tata	78.26 \pm 2.85	81.33 \pm 3.19	80.89 \pm 3.11	83.90\pm0.53
		notata	66.60 \pm 1.07	71.92 \pm 0.52	72.22 \pm 0.82	72.59\pm0.69
		all	66.47 \pm 0.74	69.74 \pm 1.3	70.14 \pm 0.45	70.89\pm0.43
	Promoter Detection	tata	78.58 \pm 3.39	80.46 \pm 0.99	81.16 \pm 1.99	83.26\pm1.90
		notata	90.81 \pm 0.51	91.88 \pm 1.01	91.67 \pm 0.49	92.48\pm0.35
		all	88.00 \pm 0.39	90.28 \pm 1.04	91.02 \pm 0.56	91.34\pm0.38
Fit Time(s)			142	125	314	169

+0.17) and on all six genomics subsets, with gains of +2.57 on *tata* core-promoter detection (83.90 vs. 81.33), +0.67 on *notata* (72.59 vs. 71.92), and +0.97 on the *all* split (70.89 vs. 69.74); for promoter detection, the improvements are +2.80 (*tata*: 83.26 vs. 80.46), +0.60 (*notata*: 92.48 vs. 91.88), and +1.06 (*all*: 91.34 vs. 90.28). Although the PLFC variant with GyroBN also performs strongly, exceeding the LFC baseline by more than +3 on all genomics subsets and matching or surpassing it on CIFAR-10 (95.28 vs. 95.14). These results indicate that decision functions based on point-to-hyperplane distance provide a more effective inductive bias than affine logits in both vision and genomics.

**509
510
511
512
513
514
515
516
GyroLBN vs. LBN, GyroBN.** Under the same FC, GyroLBN improves over LBN across all tasks. With the LFC head, CIFAR-10 increases from 95.14 to 95.19; on genomics, the gains range from +1.07 to +5.32 (e.g., core-promoter *notata*: 71.92 vs. 66.60). With the PLFC head, GyroLBN is superior to GyroBN on every dataset, including *notata* in core-promoter detection, where it reaches 72.59 versus 72.22. The improvements span +0.32 to +3.01 on genomics and +0.08 on CIFAR-10. Training time measurements show that GyroLBN attains these gains with favorable efficiency, running faster than PLFC with GyroBN (169 s vs. 314 s) and even faster than LBN for the comparison of LFC between GyroLBN and LBN (125 s vs. 142 s). Overall, GyroLBN offers a better accuracy and efficiency trade-off than both LBN and GyroBN in our settings.

517
518
519
Furthermore, to ensure a fair comparison with GyroBN, we further vary the number of Fréchet mean iterations
520
used by the normalizers (1, 2, 5, 10, and a fixed-point solve denoted by ∞ ; see Appendix G). Across all iteration
521
cases, *PLFC+GyroLBN* remains the best-performing configuration on CIFAR-10 and on every genomics split,
522
exhibiting the effectiveness of our GyroLBN again.

523 524 6 CONCLUSION

525
526
527
528
529
This study presented the Intrinsic Lorentz Neural Network, an architecture whose computations remain entirely
530
531
532
within the Lorentz model of hyperbolic space. We introduced a point-to-hyperplane fully connected layer
533
534
535
that converts signed hyperbolic distances into logits, together with GyroLBN and log-radius concatenation for
536
537
538
numerically stable normalization and feature aggregation. Integrated into a coherent network, these components
539
540
yield superior performance on CIFAR-10, CIFAR-100, and two challenging genomics benchmarks, surpassing
541
542
543
Euclidean, hybrid, and prior hyperbolic baselines while preserving competitive training cost. The results
544
545
546
underscore the value of keeping every layer intrinsic to the manifold and provide practical building blocks for
547
548
549
future work on representation learning in negatively curved geometries.

550 551 552 REPRODUCIBILITY STATEMENT

553
554
555
All theoretical results are established under explicit conditions. Experimental details are given in Appendix F.
556
557
The code will be released upon acceptance.

558 559 560 ETHICS STATEMENT

561
562
563
This work only uses publicly available datasets and does not involve human subjects or sensitive information.
564
565
We identify no specific ethical concerns.

540 REFERENCES
541

542 Ivana Balazevic, Carl Allen, and Timothy Hospedales. Multi-relational poincaré graph embeddings. *Advances
543 in neural information processing systems*, 32, 2019.

544 Ahmad Bdeir, Kristian Schwegel, and Niels Landwehr. Fully hyperbolic convolutional neural networks for
545 computer vision. In *International Conference on Learning Representations (ICLR)*, 2024.

546 James W. Cannon, William J. Floyd, Richard Kenyon, and Walter R. Parry. Hyperbolic geometry. In Silvio
547 Levy (ed.), *Flavors of Geometry*, volume 31 of *MSRI Publications*. Mathematical Sciences Research Institute,
548 1997.

549 Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks.
550 *Advances in neural information processing systems*, 32, 2019.

551 Weize Chen, Xu Han, Yankai Lin, Hexu Zhao, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. Fully
552 hyperbolic neural networks. In *Proceedings of the 60th Annual Meeting of the Association for Computational
553 Linguistics (Volume 1: Long Papers)*, pp. 5672–5686, 2022.

554 Ziheng Chen, Yue Song, Xiao-Jun Wu, and Nicu Sebe. Gyrogroup batch normalization. In *International Confer-
555 ence on Learning Representations (ICLR)*, 2025a. URL https://proceedings iclr cc/paper_files/paper/2025/file/fb19caf91f4966e426026685028cb73a-Paper-Conference.pdf.

556 Ziheng Chen, Xiao-Jun Wu, and Nicu Sebe. Riemannian batch normalization: A gyro approach. *arXiv preprint
557 arXiv:2509.07115*, 2025b.

558 Sungjun Cho, Seunghyuk Cho, Sungwoo Park, Hankook Lee, Honglak Lee, and Moontae Lee. Curve your
559 attention: Mixed-curvature transformers for graph representation learning. *arXiv preprint arXiv:2309.04082*,
560 2023.

561 Xiran Fan, Chun-Hao Yang, and Baba C. Vemuri. Horospherical decision boundaries for large margin classifica-
562 tion in hyperbolic space. *arXiv preprint arXiv:2302.06807*, 2023.

563 Xiran Fan, Minghua Xu, Huiyuan Chen, Yuzhong Chen, Mahashweta Das, and Hao Yang. Enhancing hy-
564 perbolic knowledge graph embeddings via lorentz transformations. In *Findings of the Association for
565 Computational Linguistics: ACL 2024*, pp. 4575–4589, 2024. URL <https://aclanthology.org/2024.findings-acl.272.pdf>.

566 Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. *Advances in neural
567 information processing systems*, 31, 2018.

568 Mina Ghadimi Atigh, Julian Schoep, Erman Acar, Nanne van Noord, and Pascal Mettes. Hyperbolic image
569 segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
(CVPR)*, pp. 4429–4439, 2022.

570 Caglar Gulcehre, Misha Denil, Mateusz Malinowski, Ali Razavi, Razvan Pascanu, Karl Moritz Hermann, Peter
571 Battaglia, Victor Bapst, David Raposo, Adam Santoro, et al. Hyperbolic attention networks. In *ICLR*, 2019.

572 Y. Guo, X. Wang, Y. Chen, and S. X. Yu. Clipped hyperbolic classifiers are super-hyperbolic classifiers. In
573 *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1–10, Los Alamitos,
574 CA, USA, jun 2022. IEEE Computer Society. doi: 10.1109/CVPR52688.2022.00010. URL <https://doi.ieee.org/10.1109/CVPR52688.2022.00010>.

575 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*,
576 abs/1512.03385, 2015. URL <http://arxiv.org/abs/1512.03385>.

577 Neil He, Menglin Yang, and Rex Ying. Lorentzian residual neural networks. *arXiv preprint arXiv:2412.14695*,
578 2024. URL <https://arxiv.org/abs/2412.14695>.

579 Neil He, Rishabh Anand, Hiren Madhu, Ali Maatouk, Smita Krishnaswamy, Leandros Tassiulas, Menglin Yang,
580 and Rex Ying. Helm: Hyperbolic large language models via mixture-of-curvature experts. *arXiv preprint
581 arXiv:2505.24722*, 2025.

582 Raiyan R Khan, Philippe Chlenski, and Itsik Pe'er. Hyperbolic genome embeddings. *arXiv preprint
583 arXiv:2507.21648*, 2025.

584 Andrei Khrulkov, Leyla Mirvakhanova, Evgeniya Ustinova, Ivan Oseledets, and Victor Lempitsky. Hyperbolic
585 image embeddings. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition
(CVPR)*, pp. 6418–6428, 2020.

594 Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *ICLR*,
 595 2017.

596 Reinmar J. Kobler, Jun-ichiro Hirayama, and Motoaki Kawanabe. Controlling the fréchet variance improves
 597 batch normalization on the symmetric positive definite manifold. In *ICASSP 2022 - 2022 IEEE International*
 598 *Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 3863–3867, 2022. doi: 10.1109/
 599 ICASSP43922.2022.9746629.

600 Max Kochurov, Rasul Karimov, and Serge Kozlukov. Geoopt: Riemannian optimization in pytorch. *arXiv*
 601 *preprint arXiv:2005.02819*, 2020. URL <https://arxiv.org/abs/2005.02819>.

602 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.

603 Marc T. Law, Renjie Liao, Jake Snell, and Richard S. Zemel. Lorentzian distance learning for hyperbolic
 604 representations. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of
 605 *Proceedings of Machine Learning Research*, pp. 3672–3681. PMLR, 2019. URL <https://proceedings.mlr.press/v97/law19a.html>.

606 Guy Lebanon and John Lafferty. Hyperplane margin classifiers on the multinomial manifold. In *Proceedings of*
 607 *the twenty-first international conference on Machine learning*, pp. 66, 2004.

608 Zhiying Leng, Tolga Birdal, Xiaohui Liang, and Federico Tombari. Hypersdfusion: Bridging hierarchical
 609 structures in language and geometry for enhanced 3d text2shape generation. In *Proceedings of the IEEE/CVF*
 610 *Conference on Computer Vision and Pattern Recognition*, 2024. URL https://openaccess.thecvf.com/content/CVPR2024/papers/Leng_HyperSDFusion_Bridging_Hierarchical_Structures_in_Language_and_Geometry_for_Enhanced_CVPR_2024_paper.pdf.

611 Yancong Li, Xiaoming Zhang, Ying Cui, and Shuai Ma. Hyperbolic graph neural network for temporal
 612 knowledge graph completion. In *Proceedings of the 2024 Joint International Conference on Computational*
 613 *Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pp. 8474–8486, Torino, Italia, 2024.
 614 ELRA and ICCL. URL <https://aclanthology.org/2024.lrec-main.743/>.

615 Qiuyu Liang, Weihua Wang, Feilong Bao, and Guanglai Gao. Fully hyperbolic rotation for knowledge graph
 616 embedding. *arXiv preprint arXiv:2411.03622*, 2024. URL <https://arxiv.org/abs/2411.03622>.

617 Qi Liu, Maximilian Nickel, and Douwe Kiela. Hyperbolic graph neural networks. In *NeurIPS*, pp. 8230–8241,
 618 2019.

619 Aaron Lou, Nina Miolane, Tien-Minh Le, Xavier Pennec, Susan Holmes, and Quoc Gu. Differentiating
 620 through the fréchet mean. In *Proceedings of the 37th International Conference on Machine Learning (ICML)*,
 621 volume 119 of *Proceedings of Machine Learning Research*, pp. 6393–6403, 2020. URL <https://proceedings.mlr.press/v119/lou20a.html>.

622 Gal Mishne, Zhengchao Wan, Yusu Wang, and Sheng Yang. The numerical stability of hyperbolic representation
 623 learning. In *International Conference on Machine Learning*, pp. 24925–24949. PMLR, 2023.

624 Xuan Son Nguyen and Shuo Yang. Building neural networks on matrix manifolds: A gyrovector space approach.
 625 In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of*
 626 *Machine Learning Research*. PMLR, 2023.

627 Xuan Son Nguyen, Boris Muzellec, and Agnieszka Słowik. A gyrovector space approach for symmetric positive
 628 semi-definite matrices. In *Proceedings of the 17th European Conference on Computer Vision (ECCV)*.
 629 Springer, 2022.

630 Xuan Son Nguyen, Shuo Yang, and Aymeric Histace. Matrix manifold neural networks++. *arXiv preprint*
 631 *arXiv:2405.19206*, 2024.

632 Xuan Son Nguyen, Shuo Yang, and Aymeric Histace. Neural networks on symmetric spaces of noncompact
 633 type. In *The Thirteenth International Conference on Learning Representations*, 2025.

634 Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In
 635 *Advances in Neural Information Processing Systems*, volume 30, 2017.

636 Maximilian Nickel and Douwe Kiela. Learning continuous hierarchies in the lorentz model of hyperbolic geom-
 637 etry. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of*
 638 *Machine Learning Research*, pp. 3776–3785. PMLR, 2018.

639 Avik Pal, Max van Spengler, Guido Maria D’Amely di Melendugno, Alessandro Flaborea, Fabio Galasso, and
 640 Pascal Mettes. Compositional entailment learning for hyperbolic vision-language models. In *The Thirteenth*
 641 *International Conference on Learning Representations*, 2025.

648 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen,
 649 Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zachary
 650 DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and
 651 Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *Advances in
 652 Neural Information Processing Systems (NeurIPS)*, 2019. URL [https://proceedings.neurips.
 653 cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html](https://proceedings.neurips.cc/paper/2019/hash/bdbca288fee7f92f2bfa9f7012727740-Abstract.html).

654 Eric Qu and Dongmian Zou. Lorentzian fully hyperbolic generative adversarial network. *arXiv preprint
 655 arXiv:2201.12825*, 2022.

656 Ladislav Rampášek, Michael Galkin, Vijay Prakash Dwivedi, Anh Tuan Luu, Guy Wolf, and Dominique Beaini.
 657 Recipe for a general, powerful, scalable graph transformer. In *NeurIPS*, volume 35, pp. 14501–14515, 2022.

658 John G. Ratcliffe. *Foundations of Hyperbolic Manifolds*, volume 149 of *Graduate Texts in Mathematics*. Springer,
 659 New York, 2 edition, 2006.

660 Ryohei Shimizu, Yusuke Mukuta, and Tatsuya Harada. Hyperbolic neural networks++. *arXiv preprint
 661 arXiv:2006.08210*, 2020.

662 Ondrej Skopek, Octavian-Eugen Ganea, and Gary Bécigneul. Mixed-curvature variational autoencoders. In
 663 *International Conference on Learning Representations*, 2020.

664 Alexandru Tifrea, Gary Bécigneul, and Octavian-Eugen Ganea. Poincaré glove: Hyperbolic word embeddings.
 665 *arXiv preprint arXiv:1810.06546*, 2019. URL <https://arxiv.org/abs/1810.06546>.

666 Abraham A. Ungar. *Analytic Hyperbolic Geometry and Albert Einstein’s Special Theory of Relativity*. World
 667 Scientific, 2008.

668 Abraham Albert Ungar. *Analytic Hyperbolic Geometry in N Dimensions: An Introduction*. CRC Press, Boca
 669 Raton, 2014.

670 Abraham Albert Ungar. *Analytic Hyperbolic Geometry and Albert Einstein’s Special Theory of Relativity
 671 (Second Edition)*. World Scientific, 2022.

672 Max Van Spengler, Erwin Berkhou, and Pascal Mettes. Poincare resnet. In *Proceedings of the IEEE/CVF
 673 International Conference on Computer Vision*, pp. 5419–5428, 2023.

674 Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph
 675 attention networks. *arXiv preprint arXiv:1710.10903*, 2017.

676 Simon Welz, Lucie Flek, and Akbar Karimi. Multi-hop reasoning for question answering with hyperbolic
 677 representations. *arXiv preprint arXiv:2507.03612*, 2025.

678 Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. Simplifying graph
 679 convolutional networks. In *ICML*, pp. 6861–6871. PMLR, 2019.

680 Qitian Wu, Wentao Zhao, Zenan Li, David P Wipf, and Junchi Yan. Nodeformer: A scalable graph structure
 681 learning transformer for node classification. In *NeurIPS*, volume 35, pp. 27387–27401, 2022.

682 Qitian Wu, Wentao Zhao, Chenxiao Yang, Hengrui Zhang, Fan Nie, Haitian Jiang, Yatao Bian, and Junchi Yan.
 683 Simplifying and empowering transformers for large-graph representations. In *NeurIPS*, volume 36, 2023.

684 Zhanghao Wu, Paras Jain, Matthew Wright, Azalia Mirhoseini, Joseph E Gonzalez, and Ion Stoica. Representing
 685 long-range context for graph neural networks with global attention. In *NeurIPS*, volume 34, pp. 13266–13279,
 686 2021.

687 Jiaqi Yang, Wenting Chen, Xiaohan Xing, Sean He, Xiaoling Luo, Xinheng Lyu, Linlin Shen, and Guoping
 688 Qiu. Hysurvpred: Multimodal hyperbolic embedding with angle-aware hierarchical contrastive learning and
 689 uncertainty constraints for survival prediction. *arXiv preprint arXiv:2503.13862*, 2025.

690 Menglin Yang, Min Zhou, Zhihao Li, Jiahong Liu, Lujia Pan, Hui Xiong, and Irwin King. Hyperbolic graph
 691 neural networks: A review of methods and applications. *arXiv preprint arXiv:2202.13852*, 2022.

692 Menglin Yang, Aosong Feng, Bo Xiong, Jihong Liu, Irwin King, and Rex Ying. Hyperbolic fine-tuning for large
 693 language models. *arXiv preprint arXiv:2410.04010*, 2024a.

694 Menglin Yang, Harshit Verma, Delvin Ce Zhang, Jiahong Liu, Irwin King, and Rex Ying. Hypformer: Exploring
 695 efficient transformer fully in hyperbolic space. In *Proceedings of the 30th ACM SIGKDD Conference on
 696 Knowledge Discovery and Data Mining*, pp. 3770–3781, 2024b.

702 Chengxuan Ying, Tianle Cai, Shengjie Luo, Shuxin Zheng, Guolin Ke, Di He, Yanming Shen, and Tie-Yan Liu.
703 Do transformers really perform badly for graph representation? In *NeurIPS*, 2021.
704

705 Yuansheng Zhou and Tatyana O. Sharpee. Hyperbolic geometry of gene expression. *iScience*, 24(3):102225,
706 2021. doi: 10.1016/j.isci.2021.102225.

707 Zhihan Zhou, Yanrong Ji, Weijian Li, Pratik Dutta, Ramana Davuluri, and Han Liu. Dnabert-2: Efficient
708 foundation model and benchmark for multi-species genome. *arXiv preprint arXiv:2306.15006*, 2023.
709
710
711
712
713
714
715
716
717
718
719
720
721
722
723
724
725
726
727
728
729
730
731
732
733
734
735
736
737
738
739
740
741
742
743
744
745
746
747
748
749
750
751
752
753
754
755

756	APPENDIX	
757		
758		
759	A Use of large language models	16
760		
761	B Operations in the Lorentz model	16
762		
763	C Gyrovector structures on the Lorentz model	17
764	C.1 Gyrogroups	17
765	C.2 From gyrogroups to gyrovector spaces	17
766	C.3 Gyrostructures induced by Riemannian geometry	18
767	C.4 Constant curvature model and Möbius operations	18
768	C.5 Induced gyrovector operations on the Lorentz model	18
769		
770		
771		
772	D Discussion about Intrinsic	19
773		
774	E Proofs	19
775	E.1 Proof of Theorem 1	19
776	E.2 Proof of Theorem 2	20
777		
778	F Implementation Details	22
779	F.1 Datasets	22
780	F.2 Settings	23
781		
782		
783	G More Ablation Studies	23
784		
785		
786		
787		
788		
789		
790		
791		
792		
793		
794		
795		
796		
797		
798		
799		
800		
801		
802		
803		
804		
805		
806		
807		
808		
809		

810 A USE OF LARGE LANGUAGE MODELS
811812 We use large language models to aid or polish writing.
813814 B OPERATIONS IN THE LORENTZ MODEL
815816 **Setup and notation** Fix a negative curvature $K < 0$. Let $\langle \cdot, \cdot \rangle_{\mathcal{L}}$ denote the Minkowski bilinear form with
817 signature $(-, +, \dots, +)$ on \mathbb{R}^{n+1} , and write $\|\cdot\|_{\mathcal{L}} = \sqrt{\langle \cdot, \cdot \rangle_{\mathcal{L}}}$ for the induced (Riemannian) norm on tangent
818 vectors. The n -dimensional hyperbolic space in the Lorentz (hyperboloid) model is
819

820
$$\mathbb{L}_K^n := \left\{ \mathbf{x} \in \mathbb{R}^{n+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = \frac{1}{K}, x_0 > 0 \right\},$$

821

822 and we use $\bar{\mathbf{0}} := (\frac{1}{\sqrt{-K}}, \mathbf{0})$ as the pole (“origin”). When no confusion can arise, $\|\cdot\|$ denotes the Euclidean
823 norm in a tangent space.
824825 **Distance** For $\mathbf{x}, \mathbf{y} \in \mathbb{L}_K^n$, the geodesic distance inherited from Minkowski space is
826

827
$$D_{\mathcal{L}}(\mathbf{x}, \mathbf{y}) = \frac{1}{\sqrt{-K}} \cosh^{-1}(K \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}).$$

828

829 A useful identity for computations is the “Lorentzian chord” expression (squared distance)
830

831
$$d_{\mathcal{L}}^2(\mathbf{x}, \mathbf{y}) = \|\mathbf{x} - \mathbf{y}\|_{\mathcal{L}}^2 = \frac{2}{K} - 2 \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}.$$

832

833 Specializing to the pole $\bar{\mathbf{0}}$,
834

835
$$D_{\mathcal{L}}(\mathbf{x}, \bar{\mathbf{0}}) = \|\log_{\bar{\mathbf{0}}}^K(\mathbf{x})\|, \quad d_{\mathcal{L}}^2(\mathbf{x}, \bar{\mathbf{0}}) = \frac{2}{K} - 2 \langle \mathbf{x}, \bar{\mathbf{0}} \rangle_{\mathcal{L}} = \frac{2}{K} + \frac{2x_0}{\sqrt{-K}},$$

836

837 and, equivalently,
838

839
$$D_{\mathcal{L}}(\mathbf{x}, \bar{\mathbf{0}}) = \frac{1}{\sqrt{-K}} \cosh^{-1}(\sqrt{-K} x_0).$$

840

841 **Tangent space** The tangent space at $\mathbf{x} \in \mathbb{L}_K^n$ is the Minkowski-orthogonal complement of \mathbf{x} ,
842

843
$$\mathcal{T}_{\mathbf{x}} \mathbb{L}_K^n = \{ \mathbf{v} \in \mathbb{R}^{n+1} : \langle \mathbf{v}, \mathbf{x} \rangle_{\mathcal{L}} = 0 \}.$$

844

845 Restricted to $\mathcal{T}_{\mathbf{x}} \mathbb{L}_K^n$, the metric is positive definite and coincides with the Riemannian metric of \mathbb{L}_K^n .
846847 **Exponential and logarithmic maps** For $\mathbf{z} \in \mathcal{T}_{\mathbf{x}} \mathbb{L}_K^n$,

848
$$\exp_{\mathbf{x}}^K(\mathbf{z}) = \cosh(\alpha) \mathbf{x} + \sinh(\alpha) \frac{\mathbf{z}}{\alpha}, \quad \alpha = \sqrt{-K} \|\mathbf{z}\|_{\mathcal{L}}.$$

849

850 The inverse map $\log_{\mathbf{x}}^K : \mathbb{L}_K^n \rightarrow \mathcal{T}_{\mathbf{x}} \mathbb{L}_K^n$ sends $\mathbf{y} \in \mathbb{L}_K^n$ to
851

852
$$\log_{\mathbf{x}}^K(\mathbf{y}) = \frac{\cosh^{-1}(\beta)}{\sqrt{\beta^2 - 1}} (\mathbf{y} - \beta \mathbf{x}), \quad \beta = K \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}.$$

853

854 At the pole $\bar{\mathbf{0}}$ these simplify to
855

856
$$\exp_{\bar{\mathbf{0}}}^K(\mathbf{z}) = \frac{1}{\sqrt{-K}} \left(\cosh(\sqrt{-K} \|\mathbf{z}\|), \sinh(\sqrt{-K} \|\mathbf{z}\|) \frac{\mathbf{z}}{\|\mathbf{z}\|} \right),$$

857
$$\log_{\bar{\mathbf{0}}}^K(\mathbf{y}) = \begin{cases} \mathbf{0}, & \mathbf{y} = \bar{\mathbf{0}}, \\ \frac{\cosh^{-1}(\beta)}{\sqrt{\beta^2 - 1}} (\mathbf{y} - \beta \bar{\mathbf{0}}), & \text{otherwise,} \end{cases} \quad \beta := K \langle \bar{\mathbf{0}}, \mathbf{y} \rangle_{\mathcal{L}}.$$

858

859 with the convention $\mathbf{z}/\|\mathbf{z}\| = \mathbf{0}$ when $\mathbf{z} = \mathbf{0}$.
860861 **Parallel transport** Transporting $\mathbf{v} \in \mathcal{T}_{\mathbf{x}} \mathbb{L}_K^n$ along the geodesic from \mathbf{x} to \mathbf{y} yields
862

863
$$\text{PT}_{\mathbf{x} \rightarrow \mathbf{y}}^K(\mathbf{v}) = \mathbf{v} - \frac{\langle \log_{\mathbf{x}}^K(\mathbf{y}), \mathbf{v} \rangle_{\mathcal{L}}}{d_{\mathcal{L}}(\mathbf{x}, \mathbf{y})} \left(\log_{\mathbf{x}}^K(\mathbf{y}) + \log_{\mathbf{y}}^K(\mathbf{x}) \right) = \mathbf{v} + \frac{\langle \mathbf{y}, \mathbf{v} \rangle_{\mathcal{L}}}{\frac{1}{-K} - \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}} (\mathbf{x} + \mathbf{y}).$$

864 **Lorentzian centroid and average pooling** Given points $\mathbf{x}_1, \dots, \mathbf{x}_m \in \mathbb{L}_K^n$ with nonnegative
 865 weights ν_i (not all zero), the weighted Fréchet mean with respect to the squared Lorentzian distance,
 866 $\min_{\mu \in \mathbb{L}_K^n} \sum_{i=1}^m \nu_i d_{\mathcal{L}}^2(\mathbf{x}_i, \mu)$, is obtained in closed form by
 867

$$\mu = \frac{\sum_{i=1}^m \nu_i \mathbf{x}_i}{\sqrt{-K} \left\| \sum_{i=1}^m \nu_i \mathbf{x}_i \right\|_{\mathcal{L}}}.$$

870 In neural architectures, an “average pooling” over a hyperbolic receptive field can be implemented by taking this
 871 Lorentzian centroid of the features in the field.
 872

873 **Lorentz transformations** A matrix $\mathbf{A} \in \mathbb{R}^{(n+1) \times (n+1)}$ is a Lorentz transformation if it preserves the
 874 Minkowski product: $\langle \mathbf{Ax}, \mathbf{Ay} \rangle_{\mathcal{L}} = \langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}}$ for all \mathbf{x}, \mathbf{y} . Such matrices form the Lorentz group $\mathbf{O}(1, n)$
 875 (equivalently, $\mathbf{A}^T \eta \mathbf{A} = \eta$ for the Minkowski metric η). Restricting to transformations that map the upper sheet
 876 to itself yields the time-orientation-preserving subgroup

$$\mathbf{O}^+(1, n) = \{ \mathbf{A} \in \mathbf{O}(1, n) : (\mathbf{Ax})_0 > 0 \text{ for all } \mathbf{x} \in \mathbb{L}_K^n \},$$

877 which acts by isometries on \mathbb{L}_K^n .
 878

880 Every $\mathbf{A} \in \mathbf{O}^+(1, n)$ admits a polar decomposition into a Lorentz rotation and a Lorentz boost, $\mathbf{A} = \mathbf{RB}$. The
 881 rotation fixes the time axis and rotates spatial coordinates:

$$\mathbf{R} = \begin{bmatrix} 1 & \mathbf{0}^\top \\ \mathbf{0} & \tilde{\mathbf{R}} \end{bmatrix}, \quad \tilde{\mathbf{R}} \in \mathbf{SO}(n).$$

885 A boost with velocity vector $\mathbf{v} \in \mathbb{R}^n$ ($\|\mathbf{v}\| < 1$) has the block form

$$\mathbf{B} = \begin{bmatrix} \gamma & -\gamma \mathbf{v}^\top \\ -\gamma \mathbf{v} & \mathbf{I}_n + \frac{\gamma^2}{1+\gamma} \mathbf{v} \mathbf{v}^\top \end{bmatrix}, \quad \gamma = \frac{1}{\sqrt{1 - \|\mathbf{v}\|^2}}.$$

890 (Equivalently, the spatial block can be written $\mathbf{I}_n + \frac{\gamma-1}{\|\mathbf{v}\|^2} \mathbf{v} \mathbf{v}^\top$ when $\mathbf{v} \neq \mathbf{0}$.)
 891

C GYROVECTOR STRUCTURES ON THE LORENTZ MODEL

C.1 GYROGROUPS

896 We recall the algebraic notion of a gyrogroup, which extends the concept of a group to settings where associativity
 897 is relaxed and corrected by gyrations (Ungar, 2008; 2014).

898 **Definition 1** (Gyrogroup). *Let G be a nonempty set endowed with a binary operation $\oplus: G \times G \rightarrow G$. The
 899 pair (G, \oplus) is a gyrogroup if, for all $a, b, c \in G$, the following axioms hold:*

900 (G1) *There exists an element $e \in G$ such that $e \oplus a = a$ (left identity).*

901 (G2) *For each $a \in G$ there exists $\ominus a \in G$ such that $\ominus a \oplus a = e$ (left inverse).*

902 (G3) *There exists a map $\text{gyr}[a, b]: G \rightarrow G$ (the gyration generated by a, b) such that*

$$a \oplus (b \oplus c) = (a \oplus b) \oplus \text{gyr}[a, b](c) \quad (\text{left gyroassociative law}).$$

904 (G4) *$\text{gyr}[a, b] = \text{gyr}[a \oplus b, b]$ (left reduction law).*

905 *If in addition*

$$a \oplus b = \text{gyr}[a, b](b \oplus a), \quad \forall a, b \in G,$$

907 *then (G, \oplus) is called gyrocommutative.*

909 Gyrogroups generalize groups: when all gyrations are the identity map, (G, \oplus) reduces to a usual group. The
 910 lack of strict associativity is compensated by the gyration operators, which encode curvature induced nonlinearity.
 911

C.2 FROM GYROGROUPS TO GYROVECTOR SPACES

914 To model both addition and scalar multiplication in curved geometries, gyrogroups can be enriched to gyrovector
 915 spaces (Ungar, 2008; Nguyen et al., 2022).

916 **Definition 2** (Gyrovector space). *Let (G, \oplus) be a gyrocommutative gyrogroup and let $\odot: \mathbb{R} \times G \rightarrow G$ be a
 917 map called scalar multiplication. The triple (G, \oplus, \odot) is a gyrovector space if, for all $a, b, c \in G$ and $s, t \in \mathbb{R}$,*

918 (VI) *$1 \odot a = a$, $0 \odot a = e$, $t \odot e = e$, and $(-1) \odot a = \ominus a$.*

918 (V2) $(s+t) \odot a = s \odot a \oplus t \odot a$.
 919 (V3) $(st) \odot a = s \odot (t \odot a)$.
 920 (V4) $\text{gyr}[a, b](t \odot c) = t \odot \text{gyr}[a, b](c)$.
 921 (V5) $\text{gyr}[s \odot a, t \odot a]$ is the identity map on G .

922 These axioms mirror the familiar properties of vector spaces, with gyrations accounting for the deviation from
 923 linearity. In particular, (V2) and (V3) play the role of distributivity and associativity for scalar multiplication,
 924 while (V4)–(V5) guarantee a consistent interaction between gyrations and scaling.

926 C.3 GYROSTRUCTURES INDUCED BY RIEMANNIAN GEOMETRY

928 The above algebraic objects can be constructed on a Riemannian manifold from the exponential and logarithmic
 929 maps at a distinguished origin. Following Nguyen et al. (2022); Nguyen & Yang (2023); Chen et al. (2025a),
 930 let (\mathcal{M}, g) be a complete Riemannian manifold with identity element $E \in \mathcal{M}$. Denote by Exp_x and Log_x the
 931 Riemannian exponential and logarithmic maps at $x \in \mathcal{M}$, and by $\text{PT}_{x \rightarrow y}$ the parallel transport from $T_x \mathcal{M}$ to
 $T_y \mathcal{M}$. For $P, Q, R \in \mathcal{M}$ and $t \in \mathbb{R}$, define

$$932 P \oplus Q = \text{Exp}_P(\text{PT}_{E \rightarrow P}(\text{Log}_E Q)), \quad (12)$$

$$933 t \odot P = \text{Exp}_E(t \text{Log}_E P), \quad (13)$$

$$935 \ominus P = (-1) \odot P = \text{Exp}_E(-\text{Log}_E P), \quad (14)$$

$$936 \text{gyr}[P, Q]R = (\ominus(P \oplus Q)) \oplus (P \oplus (Q \oplus R)). \quad (15)$$

937 The induced gyro inner product, norm, and distance are

$$938 \langle P, Q \rangle_{\text{gr}} = \langle \text{Log}_E P, \text{Log}_E Q \rangle_{T_E \mathcal{M}}, \quad (16)$$

$$939 \|P\|_{\text{gr}} = \sqrt{\langle P, P \rangle_{\text{gr}}}, \quad (17)$$

$$940 d_{\text{gr}}(P, Q) = \|\ominus P \oplus Q\|_{\text{gr}}. \quad (18)$$

942 Under mild regularity assumptions, $(\mathcal{M}, \oplus, \odot)$ forms a gyrovector space and the gyrodistance d_{gr} coincides
 943 with the geodesic distance on a wide class of manifolds, including constant curvature spaces (Nguyen et al.,
 944 2022; Chen et al., 2025a). In Euclidean space, these constructions reduce exactly to standard vector addition,
 945 scalar multiplication, and the Euclidean metric.

946 C.4 CONSTANT CURVATURE MODEL AND MÖBIUS OPERATIONS

948 For hyperbolic geometry it is convenient to introduce the constant curvature model

$$949 950 \mathcal{M}_K = \begin{cases} \mathbb{P}_K^n, & K < 0, \\ \mathbb{R}^n, & K = 0, \end{cases}$$

951 where \mathbb{P}_K^n is the Poincaré ball of curvature $K < 0$ and radius $1/\sqrt{-K}$. On \mathbb{P}_K^n the gyrostructures in Eqs. 12
 952 and 13 admit closed form expressions known as Möbius addition and Möbius scalar multiplication (Ungar, 2008;
 953 Ganea et al., 2018; Skopek et al., 2020).

954 Let $x, y \in \mathbb{P}_K^n$ and set $c = -K > 0$. The Möbius addition is

$$955 956 x \oplus_K y = \frac{(1 + 2c\langle x, y \rangle + c\|y\|^2)x + (1 - c\|x\|^2)y}{1 + 2c\langle x, y \rangle + c^2\|x\|^2\|y\|^2}, \quad (19)$$

958 and the Möbius scalar multiplication is

$$959 960 t \otimes_K x = \begin{cases} \frac{1}{\sqrt{c}} \tanh(t \text{artanh}(\sqrt{c}\|x\|)) \frac{x}{\|x\|}, & x \neq 0, \\ 961 0, & x = 0. \end{cases} \quad (20)$$

962 Equipped with \oplus_K and \otimes_K , the ball \mathbb{P}_K^n is a real gyrovector space whose gyrodistance coincides with the
 963 hyperbolic geodesic distance.

965 C.5 INDUCED GYROVECTOR OPERATIONS ON THE LORENTZ MODEL

966 The Lorentz model \mathbb{L}_K^n is isometric to the Poincaré ball \mathbb{P}_K^n via a standard mapping. Let $K < 0$ and denote
 967 $r = 1/\sqrt{-K}$. For $x = (x_0, x_s) \in \mathbb{L}_K^n$ with $x_0 > 0$ and $-x_0^2 + \|x_s\|^2 = 1/K$, define

$$968 969 \psi: \mathbb{L}_K^n \rightarrow \mathbb{P}_K^n, \quad \psi(x) = \frac{x_s}{x_0 + r}, \quad (21)$$

$$970 971 \psi^{-1}: \mathbb{P}_K^n \rightarrow \mathbb{L}_K^n, \quad \psi^{-1}(u) = \left(\frac{1 + c\|u\|^2}{1 - c\|u\|^2} r, \frac{2u}{1 - c\|u\|^2} \right), \quad (22)$$

972 where $c = -K > 0$. The map ψ is a Riemannian isometry between (\mathbb{L}_K^n, g_K) and (\mathbb{P}_K^n, g_K) (Ratcliffe, 2006;
 973 Skopek et al., 2020).

974 We transfer the gyrovector structure from the ball to the hyperboloid by conjugation with ψ . For $x, y \in \mathbb{L}_K^n$ and
 975 $t \in \mathbb{R}$, define
 976

$$x \oplus_K^{\mathcal{L}} y = \psi^{-1}(\psi(x) \oplus_K \psi(y)), \quad (23)$$

$$t \otimes_K^{\mathcal{L}} x = \psi^{-1}(t \otimes_K \psi(x)). \quad (24)$$

977 Substituting the explicit expressions Eqs. 19 and 20 and the coordinate maps Eqs. 21 and 22 yields closed form
 978 formulas for $x \oplus_K^{\mathcal{L}} y$ and $t \otimes_K^{\mathcal{L}} x$ in Lorentz coordinates. By construction, $(\mathbb{L}_K^n, \oplus_K^{\mathcal{L}}, \otimes_K^{\mathcal{L}})$ is a gyrovector
 979 space that is isomorphic to $(\mathbb{P}_K^n, \oplus_K, \otimes_K)$. In particular, the gyrodistance induced by Eq. 23 agrees with the
 980 hyperbolic geodesic distance on \mathbb{L}_K^n , and the gyroaddition and gyro scalar multiplication act as hyperbolic
 981 analogues of Euclidean vector addition and Euclidean scaling.

982 D DISCUSSION ABOUT INTRINSIC

983 In our paper, “intrinsic Lorentz” refers to using only operations that are well-defined on the Lorentz model itself,
 984 rather than on its ambient Minkowski space. We define a layer $(F : \mathbb{L}_K^n \rightarrow \mathbb{L}_K^m)$ intrinsic if

- 985 1. its input, output, and all intermediate states lie on some \mathbb{L}_K^d (they always satisfy $\langle z, z \rangle_L = -1/K$, $z_0 > 0$).
- 986 2. it is expressed entirely in terms of the Lorentzian geometry: $\langle \cdot, \cdot \rangle_L$, the induced distance d_L , and operators
 987 derived from them (exp/log maps, parallel transport, gyroaddition/gyroscaling, Lorentzian centroids, etc.),
 988 without ever using arbitrary Euclidean linear maps on Lorentz vectors in the ambient space.

989 Under this definition, the previous Lorentz fully connected layer (LFC) used in HCNN is not intrinsic. Its update
 990 has the form

$$991 \mathbf{y} = [\sqrt{|\phi(W\mathbf{x}, \mathbf{v})|^2 - 1/K}, \phi(W\mathbf{x}, \mathbf{v})],$$

992 where $\mathbf{x} \in \mathcal{L}_K^n$, $W\mathbf{x}$ is a standard matrix–vector product in \mathbb{R}^n and the operation

$$993 \phi(W\mathbf{x}, \mathbf{v}) = \lambda \sigma(\mathbf{v}^T \mathbf{x} + b') \frac{\mathbf{W}\psi(\mathbf{x}) + \mathbf{b}}{\|\mathbf{W}\psi(\mathbf{x}) + \mathbf{b}\|}. \quad (25)$$

994 This $W\mathbf{x}$ is defined using the ambient linear structure of the Minkowski space, not any operation on the
 995 Lorentz manifold itself. Because the core transformation is an ambient Minkowski multiplication rather than a
 996 Lorentzian/geodesic operation, this layer is only partially intrinsic.

997 By contrast, our PLFC is constructed entirely from Lorentz-geometry primitives that have closed-form definitions
 998 on \mathcal{L}_K^n . Each logit is the signed Lorentzian distance from the input point to a learned Lorentz hyperplane, and
 999 the output \mathbf{y} is then recovered in closed form as the unique point on \mathcal{L}_K^m whose signed distances to a set of
 1000 coordinate hyperplanes equal these logits. All steps (hyperplane parameterization, point-to-hyperplane distance,
 1001 reconstruction of \mathbf{y}) are expressed only via the Lorentzian inner product and distance; no ambient Euclidean
 1002 affine map $W\mathbf{x} + b$ is ever applied. In this sense, PLFC is an intrinsic Lorentz FC layer.

1003 E PROOFS

1004 E.1 PROOF OF THEOREM 1

1005 *Proof.* We work in the hyperboloid model $\mathbb{L}_K^m = \{\mathbf{x} = [x_t, \mathbf{x}_s] \in \mathbb{R}^{1+m} : \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = 1/K, x_t > 0\}$ with
 1006 $K < 0$ and Minkowski bilinear form $\langle [x_t, \mathbf{x}_s], [y_t, \mathbf{y}_s] \rangle_{\mathcal{L}} = -x_t y_t + \langle \mathbf{x}_s, \mathbf{y}_s \rangle$. Let $\bar{\mathbf{0}} = [(-K)^{-1/2}, \mathbf{0}]$ be the
 1007 basepoint. Denote by $\mathbf{e}^{(k)} \in \mathbb{R}^m$ the k -th Euclidean basis vector in the spatial block and set $\mathbf{e}^{(k)} = [0, \mathbf{e}^{(k)}]$.

1008 **Lorentz coordinate hyperplanes.** The k -th spatial axis is the geodesic through $\bar{\mathbf{0}}$ in the direction $\mathbf{e}^{(k)}$.
 1009 The Lorentz coordinate hyperplane through $\bar{\mathbf{0}}$, orthogonal to this axis, is given by

1010 **Definition 3.** (Lorentz hyperplane containing $\bar{\mathbf{0}}$ and orthogonal to the k -th axis)

$$1011 \tilde{H}_{\mathbf{e}^{(k)}, 0}^K = \{\mathbf{x} = [x_t, \mathbf{x}_s]^\top \in \mathbb{L}_K^m \mid \langle \mathbf{e}^{(k)}, \mathbf{x} \rangle_{\mathcal{L}} = \langle \mathbf{e}^{(k)}, \mathbf{x}_s \rangle = x_{s,k} = 0\}. \quad (26)$$

1012 This is the special case of the Lorentz hyperplane $\tilde{H}_{\mathbf{z}, a}$ in Eq. 4 with $a = 0$ and $\mathbf{z} = \mathbf{e}^{(k)}$ (hence $\|\mathbf{z}\|_2 = 1$),
 1013 for which $\cosh(\sqrt{-K}a) = 1$ and $\sinh(\sqrt{-K}a) = 0$, giving $\langle \mathbf{z}, \mathbf{x}_s \rangle = x_{s,k} = 0$.

1014 With Definition 3, the preparation for constructing \mathbf{y} in Eq. 8 is complete.

1026 **Derivation of \mathbf{y} .** Let $\mathbf{x} \in \mathbb{L}_K^n$ and $\mathbf{y} = [y_t, \mathbf{y}_s]^\top \in \mathbb{L}_K^m$ be the input and output of the PLFC layer,
 1027 respectively. As in Eq. 8a, for $k = 1, \dots, m$ we define the scores $v_k(\mathbf{x}) = v_{\mathbf{z}_k, a_k}(\mathbf{x})$ via the Lorentz MLR
 1028 logits in Eq. 6.

1029 To endow \mathbf{y} with the desired property—that the *signed* hyperbolic distance from \mathbf{y} to the k -th coordinate
 1030 hyperplane equals $v_k(\mathbf{x})$ —we impose the simultaneous system
 1031

$$1032 d_{\mathcal{L}}^{\pm}(\mathbf{y}, \bar{H}_{\mathbf{e}^{(k)}, 0}^K) = v_k(\mathbf{x}), \quad k = 1, \dots, m. \quad (27)$$

1033 Using Eqs. 4 and 5 with $a = 0$ and $\mathbf{z} = \mathbf{e}^{(k)}$ (so that the denominator equals 1), the unsigned
 1034 point-to-hyperplane distance specializes to
 1035

$$1036 d_{\mathcal{L}}(\mathbf{y}, \bar{H}_{\mathbf{e}^{(k)}, 0}^K) = \frac{1}{\sqrt{-K}} |\sinh^{-1}(\sqrt{-K} y_{s,k})|. \quad (28)$$

1037 Orienting $\bar{H}_{\mathbf{e}^{(k)}, 0}^K$ by the unit normal $\mathbf{e}^{(k)}$ and recalling the sign convention used in Eq. 6, the *signed* distance is
 1038 therefore
 1039

$$1040 d_{\mathcal{L}}^{\pm}(\mathbf{y}, \bar{H}_{\mathbf{e}^{(k)}, 0}^K) = \frac{1}{\sqrt{-K}} \sinh^{-1}(\sqrt{-K} y_{s,k}). \quad (28)$$

1041 Substituting Eq. 28 into Eq. 27 yields, for each k ,

$$1042 \frac{1}{\sqrt{-K}} \sinh^{-1}(\sqrt{-K} y_{s,k}) = v_k(\mathbf{x}), \quad (29)$$

1043 and hence
 1044

$$1045 y_{s,k} = \frac{1}{\sqrt{-K}} \sinh(\sqrt{-K} v_k(\mathbf{x})), \quad k = 1, \dots, m, \quad (30)$$

1046 which is exactly Eq. 8b. Collecting these coordinates gives $\mathbf{y}_s = (y_{s,1}, \dots, y_{s,m})^\top$.

1047 Finally, the time coordinate y_t is fixed by the hyperboloid constraint $\langle \mathbf{y}, \mathbf{y} \rangle_{\mathcal{L}} = 1/K$, i.e.,
 1048

$$1049 -y_t^2 + \|\mathbf{y}_s\|_2^2 = \frac{1}{K} \implies y_t = \sqrt{(-K)^{-1} + \|\mathbf{y}_s\|_2^2}, \quad (31)$$

1050 with the positive root chosen to remain on the top sheet, which is Eq. 8c. Thus Eq. 8 follows.
 1051

1052 **Confirmation of the existence of \mathbf{y} .** For any real scores $v_k(\mathbf{x})$, Eq. 30 yields real $y_{s,k}$ because \sinh is
 1053 entire. Since $-K > 0$, we have $(-K)^{-1} + \|\mathbf{y}_s\|_2^2 > 0$, so y_t in Eq. 31 is real and strictly positive. Therefore
 1054 $\mathbf{y} = [y_t, \mathbf{y}_s] \in \mathbb{L}_K^m$ always exists and lies on the correct sheet. Moreover, Eq. 29 guarantees that the signed
 1055 distance from \mathbf{y} to each coordinate hyperplane $\bar{H}_{\mathbf{e}^{(k)}, 0}^K$ is exactly $v_k(\mathbf{x})$, as required.
 1056

1057 **Flat-space limit.** As $K \rightarrow 0^-$, we have $\sinh(\sqrt{-K} v) = \sqrt{-K} v + O(K^{3/2})$, hence $y_{s,k} \rightarrow v_k(\mathbf{x})$ from
 1058 Eq. 30. Using Eq. 6 and the expansion $\cosh(\sqrt{-K} a) = 1 + O(K)$, $\sinh(\sqrt{-K} a) = \sqrt{-K} a + O(K^{3/2})$,
 1059 one obtains $v_k(\mathbf{x}) \rightarrow \langle \mathbf{z}_k, \mathbf{x}_s \rangle - a_k$, so the spatial part reduces to the Euclidean affine map with row vectors
 1060 $A_k = \mathbf{z}_k^\top$ and bias $b_k = a_k$, as stated below Theorem 1. \square
 1061

E.2 PROOF OF THEOREM 2

1063 **Theorem 2** (Margin preservation and contraction of PLFC and LFC). *Fix a curvature $K < 0$. For any input x ,
 1064 let the penultimate layer produce*

$$1065 u(x) = (u_1(x), \dots, u_m(x)) \in \mathbb{R}^m,$$

1066 and let c denote the true class. Define the pre-logit margin

$$1067 \Delta(x) := u_c(x) - \max_{j \neq c} u_j(x).$$

1068 Consider two Lorentz output-layer designs that use signed geodesic distances from the output point $y(x)$ to the
 1069 coordinate hyperplanes as logits:

1070 • PLFC head (intrinsic). The spatial coordinates are

$$1071 y_{s,k}^{\text{PLFC}}(x) = \frac{1}{\sqrt{-K}} \sinh(\sqrt{-K} u_k(x)),$$

1072 and the signed Lorentzian distance to the k -th coordinate hyperplane is

$$1073 d_k^{\text{PLFC}}(x) = \frac{1}{\sqrt{-K}} \text{asinh}(\sqrt{-K} y_{s,k}^{\text{PLFC}}(x)).$$

1080 • LFC head (extrinsic / linear). The spatial coordinates are taken directly as
 1081

$$1082 y_{s,k}^{\text{LFC}}(x) = u_k(x),$$

1083 and the signed Lorentzian distance to the same hyperplane is
 1084

$$1085 d_k^{\text{LFC}}(x) = \frac{1}{\sqrt{-K}} \operatorname{asinh}(\sqrt{-K} u_k(x)).$$

1086 Define the distance-based margins
 1087

$$1088 \Delta^{\text{PLFC}}(x) := d_c^{\text{PLFC}}(x) - \max_{j \neq c} d_j^{\text{PLFC}}(x), \quad \Delta^{\text{LFC}}(x) := d_c^{\text{LFC}}(x) - \max_{j \neq c} d_j^{\text{LFC}}(x).$$

1090 Then, for every sample x :
 1091

1092 1. **(Margin preservation of PLFC)**

$$1093 \Delta^{\text{PLFC}}(x) = \Delta(x).$$

1094 2. **(Margin contraction of LFC)** The LFC head preserves the sign of the margin and contracts its magnitude:

$$1095 \operatorname{sign}(\Delta^{\text{LFC}}(x)) = \operatorname{sign}(\Delta(x)), \quad |\Delta^{\text{LFC}}(x)| \leq |\Delta(x)|,$$

1097 with strict inequality $|\Delta^{\text{LFC}}(x)| < |\Delta(x)|$ whenever $\Delta(x) \neq 0$.
 1098

1099 *Proof.* Define
 1100

$$h(t) := \frac{1}{\sqrt{-K}} \operatorname{asinh}(\sqrt{-K} t), \quad K < 0.$$

1102 By the point-to-hyperplane distance formula in the Lorentz model, for any Lorentz point $y = (y_0, y_s)$ the
 1103 signed distance to the k -th coordinate hyperplane is exactly $h(y_{s,k})$.
 1104

(1) *PLFC preserves the margin.* For the PLFC head we have
 1105

$$1106 y_{s,k}^{\text{PLFC}}(x) = \frac{1}{\sqrt{-K}} \sinh(\sqrt{-K} u_k(x)),$$

1108 and thus
 1109

$$1110 d_k^{\text{PLFC}}(x) = h(y_{s,k}^{\text{PLFC}}(x)) = \frac{1}{\sqrt{-K}} \operatorname{asinh}\left(\sqrt{-K} \cdot \frac{1}{\sqrt{-K}} \sinh(\sqrt{-K} u_k(x))\right) = \frac{1}{\sqrt{-K}} \operatorname{asinh}(\sinh(\sqrt{-K} u_k(x))).$$

1111 Since $\sinh : \mathbb{R} \rightarrow \mathbb{R}$ is a bijection with inverse asinh , we have $\operatorname{asinh}(\sinh z) = z$ for all $z \in \mathbb{R}$, hence
 1112

$$1113 d_k^{\text{PLFC}}(x) = u_k(x).$$

1114 Consequently,
 1115

$$1116 \Delta^{\text{PLFC}}(x) = d_c^{\text{PLFC}}(x) - \max_{j \neq c} d_j^{\text{PLFC}}(x) = u_c(x) - \max_{j \neq c} u_j(x) = \Delta(x),$$

1118 which proves (1).
 1119

(2) *LFC contracts the margin.* For the LFC head we have $y_{s,k}^{\text{LFC}}(x) = u_k(x)$, hence
 1120

$$1121 d_k^{\text{LFC}}(x) = h(u_k(x)).$$

1123 We first record two basic properties of h . Differentiating,
 1124

$$h'(t) = \frac{d}{dt} \left[\frac{1}{\sqrt{-K}} \operatorname{asinh}(\sqrt{-K} t) \right] = \frac{1}{\sqrt{-K}} \cdot \frac{\sqrt{-K}}{\sqrt{1 + (-K)t^2}} = \frac{1}{\sqrt{1 + (-K)t^2}},$$

1126 so $h'(t) > 0$ for all t and $h'(t) \leq 1$ with $h'(t) = 1$ if and only if $t = 0$. Therefore, h is strictly increasing and
 1127 1-Lipschitz.
 1128

1129 Let j^* be any index of a maximizer of the competing pre-logits:
 1130

$$1131 j^* \in \arg \max_{j \neq c} u_j(x).$$

1132 Since h is strictly increasing, the same index maximizes the distance-based logits:
 1133

$$\max_{j \neq c} d_j^{\text{LFC}}(x) = \max_{j \neq c} h(u_j(x)) = h(u_{j^*}(x)).$$

1134
1135
1136

Hence

$$\Delta(x) = u_c(x) - u_{j^*}(x), \quad \Delta^{\text{LFC}}(x) = h(u_c(x)) - h(u_{j^*}(x)).$$

1137 If $\Delta(x) = 0$, then $u_c(x) = u_{j^*}(x)$ and consequently $\Delta^{\text{LFC}}(x) = 0$, so the conclusion holds trivially. Suppose
1138 now $\Delta(x) \neq 0$ and, without loss of generality, set1139
1140

$$a := u_c(x), \quad b := u_{j^*}(x), \quad a \neq b.$$

1141 Assume $a > b$; the case $a < b$ is analogous by symmetry. Using the fundamental theorem of calculus,1142
1143
1144

$$h(a) - h(b) = \int_b^a h'(t) dt.$$

1145 Because $h'(t) > 0$ for all t , we have $h(a) - h(b) > 0$, so $\text{sign}(h(a) - h(b)) = \text{sign}(a - b)$. Moreover, since
1146 $h'(t) \leq 1$ everywhere and $h'(t) < 1$ for all $t \neq 0$, the integrand is strictly less than 1 on a subset of $[b, a]$ with
1147 positive measure whenever $a \neq b$. Thus1148
1149

$$0 < h(a) - h(b) = \int_b^a h'(t) dt < \int_b^a 1 dt = a - b.$$

1150

Taking absolute values yields

1151
1152

$$|h(a) - h(b)| < |a - b|.$$

1153

Substituting back $a = u_c(x)$ and $b = u_{j^*}(x)$, we obtain1154
1155

$$\text{sign}(\Delta^{\text{LFC}}(x)) = \text{sign}(\Delta(x)), \quad |\Delta^{\text{LFC}}(x)| = |h(a) - h(b)| < |a - b| = |\Delta(x)|.$$

1156

Combining with the $\Delta(x) = 0$ case, this gives1157
1158

$$|\Delta^{\text{LFC}}(x)| \leq |\Delta(x)|,$$

1159

with strict inequality whenever $\Delta(x) \neq 0$, which proves (2). \square 1160
1161

F IMPLEMENTATION DETAILS

1162

F.1 DATASETS

1163
1164
1165
1166
1167All these datasets exhibit hierarchical class relations and high hyperbolicity (low δ_{rel}) as shown in Table 5, making the use of hyperbolic models well-motivated.1168
1169
1170
1171
1172
1173
1174
1175

Image Datasets For image classification, we adopt the standard benchmarks CIFAR-10 and CIFAR-100 (Krizhevsky et al., 2009). CIFAR-10 and CIFAR-100 each contain 60,000 32×32 color images drawn from 10 and 100 classes, respectively. Following the PyTorch setup and (Bdeir et al., 2024), we use 50,000 images for training and 10,000 for testing. CIFAR-10 and CIFAR-100 are standard proxies for visual object recognition whose categories naturally admit semantic hierarchies (e.g., animal \rightarrow mammal \rightarrow dog \rightarrow specific breed). In CIFAR-100, this is made explicit by grouping the 100 fine-grained classes into 20 coarse superclasses in the original dataset design. This hierarchical structure has been extensively verified in prior hyperbolic vision work (Bdeir et al. (2024); Nguyen et al. (2025) on CIFAR-10/100).

1176
1177
1178
1179
1180
1181
1182
1183
1184
1185
1186
1187

Gene Datasets For genomic sequence classification, we evaluate on the Transposable Elements Benchmark (TEB) (Khan et al., 2025) and the Genome Understanding Evaluation (GUE) (Zhou et al., 2023) suite. TEB is a curated, multi-species collection of binary classification tasks spanning seven transposable-element families across retrotransposons, DNA transposons, and pseudogenes; we follow the authors’ released preprocessing and data partitions. GUE aggregates 28 datasets covering seven biologically meaningful tasks—including transcription-factor binding, epigenetic mark prediction, promoter and splice-site detection—with sequences ranging roughly from 70 to 1000 base pairs originating from yeast, mouse, human, and viral genomes. Unless otherwise noted, we adopt the official train/validation/test splits and report the Matthews correlation coefficient (MCC) as our primary metric. The TEB and GUE suites are constructed directly from natural genomic sequences and inherit the biological hierarchies of their domains. TEB tasks span multiple transposable-element families across retrotransposons, DNA transposons, and pseudogenes, which themselves sit in a multi-level taxonomic hierarchy (orders \rightarrow superfamilies \rightarrow families). GUE aggregates datasets for transcription-factor binding, promoter and core-promoter detection, splice-site prediction, and COVID-variant classification across several species. These tasks are all manifestations of hierarchical regulatory structure (e.g., motifs \rightarrow modules \rightarrow promoters \rightarrow gene expression).

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198
1199
1200
1201
1202
1203
1204
1205
1206
1207
1208
1209
1210
1211
1212
1213
1214
1215
1216
1217
1218
1219
1220
1221
1222
1223
1224
1225
1226
1227
1228
1229
1230
1231
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241
Table 5: Hyperbolicity values of the datasets used in our experiments (δ_{rel}).

Benchmark	Task	Dataset	δ_{rel}
CIFAR	Image classification	CIFAR-10	0.26
		CIFAR-100	0.23
TEB	Pseudogenes	processed	0.19
	Pseudogenes	unprocessed	0.16
GUE	Covid variant classification	covid	0.42
		all	0.23
GUE	Core Promoter detection	notata	0.21
		tata	0.14
GUE	Promoter detection	all	0.26
		notata	0.26
		tata	0.14

F.2 SETTINGS

Table 6 summarizes the hyperparameters used to train the model. We additionally note a dataset-specific choice regarding normalization statistics. On *CIFAR-10/100* we enable running statistics: we maintain per-channel running statistics by updating the Lorentzian centroid (mean) and dispersion (variance) with a momentum term and, at test time, substitute these running estimates for the batch statistics. In contrast, on the *TEB* and *GUE* genomic suites, we disable running statistics entirely, because enabling them consistently led to a collapse of MCC after a few dozen epochs. For these genomic datasets, we therefore compute statistics on-the-fly from each evaluation batch (i.e., no moving averages are used at test time). Compared with natural images, the genomic tasks exhibit stronger distributional non-stationarity (heterogeneous sequence lengths and tasks) and higher batch-to-batch variability. Under these conditions, momentum-based running estimates accumulate bias and lag behind the true data distribution; in a Lorentzian normalization layer, a biased centroid and underestimated dispersion can over- or under-normalize timelike features, shrinking margins and destabilizing optimization.

Table 6: Summary of hyperparameters used in different datasets.

Hyperparameter	CIFAR-10&100	TEB	GUE
Epochs	200	150	150
Batch size	128	256	256
Learning rate (LR)	1e-1	8e-4	9e-4
Drop LR epochs	60, 120, 160	100,130	100,130
Drop LR gamma	0.2	0.1	0.1
Weight decay	5e-4	6e-3	5e-3
Optimizer	(Riemannian)SGD	(Riemannian)Adam	(Riemannian)Adam
Floating point precision	32 bit	32 bit	32 bit
GPU type	RTX A100	RTX A100	RTX A100
Num. GPUs	1	1	1
Hyperbolic curvature K	-1	-1	-1
Dropout rate	0.05	0.05	0.05

G MORE ABLATION STUDIES

Table 7 above expands the ablation to the number of Fréchet mean iterations used by each normalizer while keeping all other components fixed. Across all GUE datasets and both tasks, *PLFC+GyroLBN* achieves the best accuracy under every iteration budget, and the relative ordering of methods is unchanged as the budget increases. Moving from 1 to 2 and 5 iterations yields modest but consistent gains, whereas 10 steps and the fixed-point solution (cells with gray background, denoted by ∞) offer only marginal improvements at a higher computational cost. These trends indicate that the advantage of GyroLBN stems from the normalization rule itself rather than from merely computing a tighter Fréchet mean. In practice, allocating a small budget of two to

1242 five iterations recovers nearly all of the attainable accuracy while preserving efficiency, making *PLFC+GyroLBN*
 1243 the most effective and economical choice in our setting.
 1244

1245 Table 7: Ablation study on the number of Fréchet mean iterations. The symbol ∞ indicates that
 1246 iterations are performed until convergence, which is the setting used in our main ablation experiments.
 1247

1248 1249 1250	1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295	1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295	1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295	1251 1252 1253 1254 1255 1256 1257 1258 1259 1260 1261 1262 1263 1264 1265 1266 1267 1268 1269 1270 1271 1272 1273 1274 1275 1276 1277 1278 1279 1280 1281 1282 1283 1284 1285 1286 1287 1288 1289 1290 1291 1292 1293 1294 1295	Model			
				LFC LBN	LFC GyroLBN	PLFC GyroBN	PLFC GyroLBN	
GUE	Core Promoter Detection	tata	1			80.09 \pm 1.90		
			2			81.38 \pm 2.74		
			5	78.26 \pm 2.85	81.33 \pm 3.19	79.47 \pm 2.21	83.90\pm0.53	
			10			81.27 \pm 1.75		
			∞			80.89 \pm 3.11		
	Promoter Detection	notata	1			71.99 \pm 0.68		
			2			71.06 \pm 0.49		
			5	66.60 \pm 1.07	71.92 \pm 0.52	71.65 \pm 0.79	72.59\pm0.69	
			10			71.12 \pm 1.75		
			∞			72.22 \pm 1.44		
	all	all	1			70.47 \pm 0.85		
			2			70.42 \pm 0.49		
			5	66.47 \pm 0.74	69.74 \pm 1.3	70.75 \pm 0.45	70.89\pm0.43	
			10			69.81 \pm 0.67		
			∞			70.14 \pm 0.45		