Extended Abstract Track

Which Hyperbolic Model? An Analytical and Empirical Comparison

Editors: List of editors' names

Abstract

Hyperbolic geometry has become central to machine learning for hierarchical and graph-structured data. Several isometric models of hyperbolic space — including the Poincaré disk, Lorentz (hyperboloid), Klein, Half-space, and Hemisphere — are theoretically equivalent but behave differently in in machine learning tasks. This work addresses the question: Which hyperbolic model should one use? We contribute (i) an analytical and visual comparison of the five classical models, (ii) an experimental framework for recommender systems with hyperbolic matrix factorization (HMF) and hyperbolic Bayesian personalized ranking (HBPR), and (iii) preliminary insights on optimization stability across models. The present extended abstract emphasizes contributions (i) and (ii); full experimental results will appear in a longer version.

Keywords: List of keywords

1. Introduction

Hyperbolic geometry provides an expressive embedding space for data with latent hierarchical or tree-like structure. Successful applications include NLP, graphs neural network Nickel and Kiela (2017); Ganea et al. (2018); Chami et al. (2019). Multiple models exist Cannon et al. (1997)— Poincaré, Lorentz, Klein, Half-space, Hemisphere — each an isometric representation of $\mathbb{H}^n$. While theoretically equivalent, they differ in optimization properties, interpretability, and visualization. The open question remains: *which model should one use in practice?* Prior surveys Bachmann et al. (2022); Gao et al. (2024) have introduced hyperbolic methods, but systematic model comparisons are limited. Our contribution includes:

- (i) a detailed analytical and visual comparison of the five models,
- (ii) an empirical framework in recommender systems, and
- (iii) preliminary observations on optimization stability.

This extended abstract focuses on (i) and (ii), leaving full experiments for future work.

2. Five Hyperbolic Models

While mathematically equivalent, the various isometric models of hyperbolic space—differing in coordinate representations and metric properties—exhibit distinct practical advantages and drawbacks. We examine five principal models, comparing them along three key dimensions: their mathematical formalisms (Table 1), the computational complexity of their distance functions (Table 2), and their qualitative trade-offs in applied machine learning contexts such as optimization stability and visualization (Table 3).

Model	Domain	Distance
Poincaré disk	$\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < 1\}$	$d_{\mathbb{D}}(\mathbf{x}, \mathbf{y}) = \operatorname{arcosh}\left(1 + 2\frac{\|\mathbf{x}-\mathbf{y}\|^2}{(1-\|\mathbf{x}\|^2)(1-\|\mathbf{y}\|^2)}\right)$
Lorentz	$\{\mathbf{x} \in \mathbb{R}^{n+1} : \langle \mathbf{x}, \mathbf{x} \rangle_{\mathcal{L}} = -1,\, x_0 > 0\}$	$d_{\mathbb{L}}(\mathbf{x}, \mathbf{y}) = \operatorname{arcosh}(-\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}})$
Klein	$\{\mathbf{x} \in \mathbb{R}^n : \|\mathbf{x}\| < 1\}$	$d_{\mathbb{K}}(\mathbf{x}, \mathbf{y}) = \operatorname{arcosh}\left(\frac{1-\langle \mathbf{x}, \mathbf{y} \rangle}{\sqrt{(1-\|\mathbf{x}\|^2)(1-\|\mathbf{y}\|^2)}}\right)$
Half-space	$\{(\mathbf{x}, x_n) \in \mathbb{R}^{n-1} \times \mathbb{R}_{>0}\}$	$d_{\mathbb{H}}(\mathbf{x}, \mathbf{y}) = \operatorname{arcosh}\left(1 + \frac{\|\mathbf{x}-\mathbf{y}\|^2}{2x_n y_n}\right)$
Hemisphere	$\{\mathbf{x} \in \mathbb{R}^{n+1} : \|\mathbf{x}\| = 1,\, x_0 > 0\}$	$d_{\mathbb{S}}(\mathbf{x}, \mathbf{y}) = \operatorname{arcosh}(\langle \mathbf{x}, \mathbf{y} \rangle_{\mathcal{L}})$

Table 1: Five isometric models of hyperbolic space with their domains and distance functions.

Model	Distance	Computational Cost
Lorentz (Hyperboloid)	$d_{\mathbb{L}}(x, y)$	Lowest
Hemisphere	$d_{\mathbb{S}}(x, y)$	Very Low
Half-Space (Poincaré)	$d_{\mathbb{H}}(u, v)$	Medium
Poincaré Disk	$d_{\mathbb{D}}(x, y)$	High
Klein	$d_{\mathbb{K}}(x, y)$	Highest

Table 2: Computational cost ranking of hyperbolic distance metrics.

Discussion. The three tables provide a complete framework for selecting a hyperbolic model. Table 1 establishes their fundamental mathematical equivalence, showing that the choice is, in theory, arbitrary. However, this theoretical symmetry is broken in practice. Table 2 provides a quantitative criterion, ranking models by computational cost and identifying the Lorentz model as the most efficient. Table 3 complements this with qualitative practical trade-offs: the Lorentz model's computational advantage is paired with stable optimization, while the visually intuitive Poincaré disk suffers from unstable gradients. The other models offer specialized niches—the Klein model is midpoint-friendly, the Half-space offers interpretable scaling, and the Hemisphere exhibits symmetry—but their use cases are often limited by higher computational cost or other drawbacks. Consequently, the model selection process is guided not by geometry, but by the specific computational and representational demands of the application, with the Lorentz model often being the default choice for large-scale optimization.

3. Analytical and Visual Comparison

To build intuition about the five hyperbolic models, we visualize them in low dimensions: $n = 1$ embedded in $\mathbb{R}^2$ and $n = 2$ embedded in $\mathbb{R}^3$. Although isometric, their embeddings and geodesic structures appear quite different (Figures 1 and 2).

Each model offers a unique perspective on hyperbolic geodesics. The most visually intuitive are perhaps those of the Poincaré disk (circular arcs meeting the boundary at right angles) and the Klein model (straight lines), which are Euclideanly simple but metrically hyperbolic. In contrast, the Lorentz model's geodesics are determined by planar intersec-

Model	Optimization	Visualization
Poincaré	− unstable gradients	+ intuitive disk
Lorentz	+ stable optimization	− less interpretable
Klein	± midpoint-friendly	− distorted geodesics
Half-space	+ interpretable scaling	− boundary issues
Hemisphere	+ symmetry	± less common

Table 3: Qualitative pros and cons of hyperbolic models for machine learning tasks.

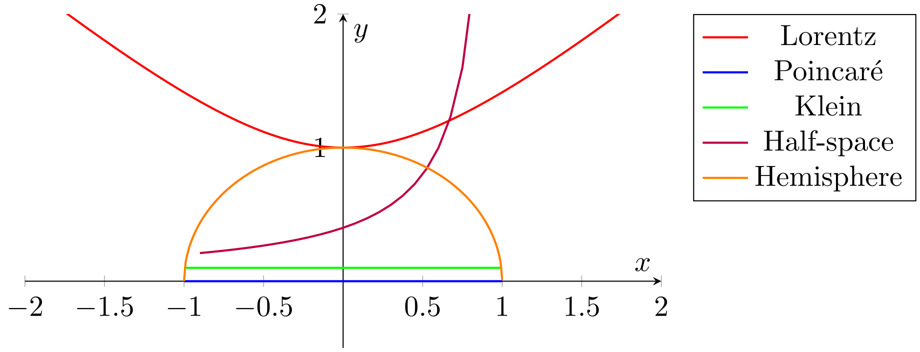

Figure 1: One-dimensional case ($n = 1$) of the five hyperbolic models in $\mathbb{R}^2$.

tions with the hyperboloid, a structure that lends itself well to stable computation. The Half-space model portrays geodesics as vertical rays or semicircles, while the Hemisphere model represents them as great-circle arcs, highlighting its connection to spherical geometry.

4. Experiments in Recommender Systems

4.1. Setup

To evaluate the practical impact of our analysis, we conduct experiments on recommender systems by integrating hyperbolic embeddings into two classical models: Matrix Factorization (MF) Koren et al. (2009) and Bayesian Personalized Ranking (BPR) Rendle et al. (2009).

Datasets. We focus on the MovieLens 100K dataset as a benchmark, and plan to extend to larger datasets such as LastFM and Amazon reviews. Table 4 summarizes the statistics.

Dataset	Users	Items	Interactions
MovieLens 100K	943	1,682	100,000
LastFM (planned)	∼2K	∼17K	∼90K
Amazon Music (planned)	∼20K	∼10K	∼200K

Table 4: Datasets for experiments (completed and planned).

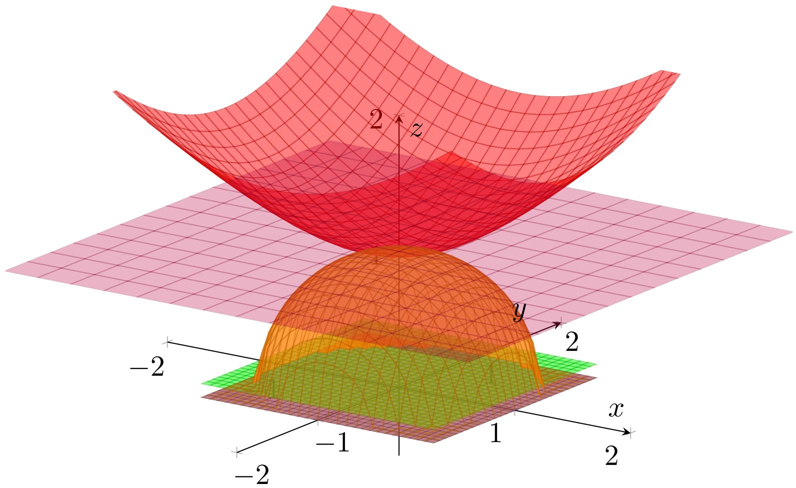

Figure 2: Two-dimensional case ($n = 2$) of the five hyperbolic models in $\mathbb{R}^3$: Lorentz (red), Poincaré (blue), Klein (green), Hemisphere (orange), Half-space (purple).

Models. We benchmark standard Euclidean Matrix Factorization (MF) Koren et al. (2009) and Bayesian Personalized Ranking (BPR) Rendle et al. (2009) against their hyperbolic variants (HMF, HBPR). Each hyperbolic variant is implemented across the five models, with embeddings constrained to the appropriate manifold and predictions made via hyperbolic distance.

Metrics. Performance is evaluated using rating prediction metrics (RMSE, MAE) and ranking metrics (HR@K, NDCG@K):

$$\text{RMSE} = \sqrt{\frac{1}{|\mathcal{T}|}\sum_{u,i}(\hat{r}_{ui} - r_{ui})^2}, \quad \text{MAE} = \frac{1}{|\mathcal{T}|}\sum_{u,i}|\hat{r}_{ui} - r_{ui}|$$

$$\text{HR@K} = \frac{1}{|\mathcal{U}|}\sum_{u}\mathbb{1}(\text{rank}_u \leq K), \quad \text{NDCG@K} = \frac{1}{|\mathcal{U}|}\sum_{u}\frac{\mathbb{1}(\text{rank}_u \leq K)}{\log_2(\text{rank}_u + 1)}$$

where $\mathbb{1}$ is the indicator function, $\mathcal{T}$ is the test set, and $\mathcal{U}$ is the user set.

4.2. Status and Ongoing Work

A full suite of experimental results will be presented in the final version of this work. Our ongoing efforts are focused on completing: (i) a comprehensive comparison across all five hyperbolic models, (ii) an extensive evaluation on multiple datasets, and (iii) detailed ablation studies on the effects of embedding dimensionality and optimization schemes.

5. Conclusion

This study provides a unified analytical and visual comparison of five hyperbolic models and an experimental framework for recommender systems. Theoretical isometries do not imply empirical interchangeability, making the choice of the model a crucial design decision. Future work will extend experiments and provide practical guidelines.

6. Citations and Bibliography

References

Gregor Bachmann, Gary Bécigneul, and Octavian Ganea. Hyperbolic deep neural networks: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. doi: 10.1109/TPAMI.2022.3147916.

James W. Cannon, William J. Floyd, Richard Kenyon, and Walter R. Parry. *Hyperbolic Geometry*. MSRI Publications, 1997. URL http://library.msri.org/books/Book31/files/cannon.pdf.

Ines Chami, Rex Ying, Christopher Ré, and Jure Leskovec. Hyperbolic graph convolutional neural networks. *Advances in Neural Information Processing Systems (NeurIPS)*, 32, 2019.

Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. Hyperbolic neural networks. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 5345–5355, 2018.

X. Gao et al. Hyperbolic brain representations. *arXiv preprint arXiv:2409.12990*, 2024.

Yehuda Koren, Robert Bell, and Chris Volinsky. Matrix factorization techniques for recommender systems. In *Computer*, volume 42, pages 30–37. IEEE, 2009.

Maximilian Nickel and Douwe Kiela. Poincaré embeddings for learning hierarchical representations. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 6338–6347, 2017.

Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the 25th Conference on Uncertainty in Artificial Intelligence (UAI)*, pages 452–461, 2009.