

Geometric Disentangled Collaborative Filtering

Yiding Zhang*
zyd@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Beijing, China

Xiao Wang
xiaowang@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Beijing, China

Hao Sun
hasun@microsoft.com
Microsoft
Beijing, China

Chaozhuo Li*
cli@microsoft.com
Microsoft Research Asia
Beijing, China

Chuan Shi†
shichuan@bupt.edu.cn
Beijing University of Posts and
Telecommunications
Beijing, China

Liangjie Zhang
liazha@microsoft.com
Microsoft
Beijing, China

Qi Zhang
qizhang@microsoft.com
Microsoft
Beijing, China

Xing Xie
xing.xie@microsoft.com
Microsoft Research Asia
Beijing, China

Yuming Liu
yumliu@microsoft.com
Microsoft
Beijing, China

Weiwei Deng
dedeng@microsoft.com
Microsoft
Beijing, China

ABSTRACT

Learning informative representations of users and items from the historical interactions is crucial to collaborative filtering (CF). Existing CF approaches usually model interactions solely within the Euclidean space. However, the sophisticated user-item interactions inherently present highly non-Euclidean anatomy with various types of geometric patterns (i.e., tree-likeness and cyclic structures). The Euclidean-based models may be inadequate to fully uncover the intent factors beneath such hybrid-geometry interactions. To remedy this deficiency, in this paper, we study the novel problem of Geometric Disentangled Collaborative Filtering (GDGF), which aims to reveal and disentangle the latent intent factors across multiple geometric spaces. A novel generative GDGF model is proposed to learn geometric disentangled representations by inferring the high-level concepts associated with user intentions and various geometries. Empirically, our proposal is extensively evaluated over five real-world datasets, and the experimental results demonstrate the superiority of GDGF.

*Both authors contributed equally to this research.

†Chuan Shi is a corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '22, July 11–15, 2022, Madrid, Spain.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8732-3/22/07...\$15.00

<https://doi.org/10.1145/3477495.3531982>

CCS CONCEPTS

• Information systems → Collaborative filtering.

KEYWORDS

Collaborative Filtering, Disentangled Representation Learning, Non-Euclidean Geometry

ACM Reference Format:

Yiding Zhang, Chaozhuo Li, Xing Xie, Xiao Wang, Chuan Shi, Yuming Liu, Hao Sun, Liangjie Zhang, Weiwei Deng, and Qi Zhang. 2022. Geometric Disentangled Collaborative Filtering. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '22)*, July 11–15, 2022, Madrid, Spain. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3477495.3531982>

1 INTRODUCTION

The rapid development of information technology has facilitated an explosion of information, leading to the challenge of information overload [1]. Recommender systems mitigate the information overload by suggesting a small set of items for users to meet their personalized interests [26, 27, 47]. Basically, recommender systems aim at modeling a user's preferences based on her historical interactions (e.g., ratings and clicks) with different items, known as collaborative filtering (CF) [44], and further recommending the user some items that she might be interested in [1].

Most existing CF models follow the paradigm to first learn a set of user/item representations and then build an interaction function to make recommendations based on the learned embeddings [44]. Learning representations that precisely reflect users' preferences based chiefly on user historical behaviors, has been a central point of interest. Matrix Factorization (MF) [24] embeds users and




Figure 1: An illustration of the user-item interactions and the non-Euclidean structures.

items as distributed vectors via matrix decomposition. Deep neural networks (DNNs) are further introduced to capture the latent preferences beneath the highly non-linear interactions [15, 47]. Recently, graph neural networks (GNNs) have demonstrated great potential in boosting recommendation performance. High-order user-item interactions are explicitly captured by the stacked GNN layers to learn expressive embeddings [14, 48].

The representation learning of most CF models is defined in the Euclidean space due to the computational simplicity [43]. However, recent researches have demonstrated that a myriad of data (e.g., data with tree-likeness or cyclic structures) exhibits the highly non-Euclidean latent anatomy [3, 25, 40, 50]. Compared to the conventional Euclidean geometry, non-Euclidean geometries (i.e., hyperbolic or spherical geometry) are more suitable for modeling data with non-Euclidean characteristics. Coincidentally, such non-Euclidean characteristics also inherently exist in the CF scenario. As shown in Figure 1(b), the high-order interactions of user D_1 can be naturally extended to a tree-likeness structure based on the interactions in Figure 1(a). This is reasonable as the receptive field tends to be exponentially larger in the higher orders [42]. Euclidean geometry is insufficient to obtain relatively low distortion for embedding such tree structures even using an unbounded number of dimensions [28], but this task would be surprisingly easy for hyperbolic spaces with only 2 dimensions [41]. In addition, users sharing similar preferences may interact with similar items, leading to the cyclic structures in Figure 1(c) (e.g., $D_4 \rightarrow \mathcal{I}_2 \rightarrow D_5 \rightarrow \mathcal{I}_3 \rightarrow D_4$). Such cyclic structures indicate the behavioral similarities between users, which are essentially the collaborative signals leveraged by CF methods [44]. Representing the cyclic structures into the Euclidean space with limited expressivity may result in inferior representations, while spherical geometry is more powerful in modeling data with cyclic structures [3, 9].

Recently, several works focus on learning user/item representations within a single type of non-Euclidean space (e.g., hyperbolic space) [36, 45, 46, 52]. However, the underlying structures of user-item interactions are more complicated than the pure trees or spheres. As shown in Figure 1(b), the tree-likeness structure and the cyclic structure (e.g., $D_1 \rightarrow \mathcal{I}_1 \rightarrow D_3 \rightarrow \mathcal{I}_2 \rightarrow D_1$) are nested together, leading to the hybrid geometric characteristics. From this point of view, the single geometry-based hypothesis is inadequate to capture latent intent factors beneath such hybrid structures. To reveal the latent reasons why a user interacted with an item, disentangled representation learning is introduced to learn factorized

embeddings to disentangle the latent intent factors [31, 32, 49, 54]. Although existing disentangled CF models learn representations solely within the Euclidean geometry, they motivate us to disentangle the sophisticated characteristics with different geometries.

In this paper, we propose to perform CF over hybrid geometries via disentangled learning, namely Geometric Disentangled Collaborative Filtering (GDCF). Based on the basic assumption that user-item interactions are latently generated from highly sophisticated intent factors, we further presume that these intent factors should be associated with various geometries. Different intent factors are responsible for generating user-item interactions belonging to different types of geometries. However, how to learn a desirable GDCF model is still obscure. Previous disentangled CF models only focus on uncovering the intent factors, while GDCF further needs to capture the latent correlations between the intent factors and different geometries. In addition, the closeness measurements in different geometries are distinct. For example, distances in spherical spaces are finite while the distances can be infinite in Euclidean and hyperbolic spaces. Since the representations of users and items are defined within different geometries, it is essential to measure the closeness across different geometries in a principled manner.

To address the mentioned challenges, we make the first attempt to learn geometric disentangled representations and propose a novel generative GDCF model on the basis of variational autoencoders (VAEs). Specifically, the geometric disentangled representations are learned by identifying the high-level geometric concepts associated with user intentions. The concepts represent the intent factors associated with different geometries, which contribute to capturing the task-relevant correlations between latent factors and geometries. To measure the user-item similarities across multiple geometries, we further propose to project the disentangled representations from different spaces into a shared latent space, and thus a universal measurement can be leveraged to calculate their closeness. Our proposal is thoroughly evaluated on five publicly available datasets, and the experimental results demonstrate its superiority.

We summarize our main contributions as follows:

- To the best of our knowledge, we are the first to study the novel problem of geometric disentangled collaborative filtering, which is capable of capturing fine-grained geometric characteristics beneath the hybrid user-item interactions.
- We propose a novel GDCF model to learn geometric disentangled representations by identifying the high-level geometry-aware concepts associated with user intentions.
- Extensively, we evaluate GDCF on five real-world datasets. Experimental results demonstrate its superiority.

2 PROBLEM DEFINITION

In this section, we will formally define the studied problem. A recommendation dataset can be formulated as $\mathbf{D} = \{\mathbf{U} \cdot \mathbf{I} \cdot \mathbf{E}\}$, in which $\mathbf{U} = \{D_1 \cdot D_2 \cdot \dots \cdot D_n\}$ denotes a set of n users, $\mathbf{I} = \{\mathcal{I}_1 \cdot \mathcal{I}_2 \cdot \dots \cdot \mathcal{I}_\# \}$ is the set of $\#$ items and $\mathbf{E} \in \{0 \cdot 1\}^{n \times \#}$ consists of the historical interactions between users and items. For convenience, we use $\mathcal{I}_D = \{\mathcal{I} : \mathbf{E}_{D \cdot \mathcal{I}} = 1\}$ to represent items interacted with user D . Given a candidate pair $(D \cdot \mathcal{I})$ consisting of a target user D and a potential item \mathcal{I} , we aim to learn a preference score $f_{D \cdot \mathcal{I}} \in \{0 \cdot 1\}$ to indicate how likely this item should be recommended to the target user.

Different from existing single geometry based approaches, we aim to learn the representations for users $\{z_D\}_{D=1}^{\#}$ and items $\{h_\beta\}_{\beta=1}^n$ by disentangling user embeddings with different geometries to comprehensively understand the sophisticated interactions.

3 PRELIMINARY

In this section, we will introduce some preliminary knowledge from the perspective of manifolds, including the basic notations and operators in three geometries (i.e., hyperbolic, Euclidean and spherical geometries). Three types of manifolds \mathcal{M} are defined with the constant sectional curvature \wedge [5]. The general realizations of these manifolds are the hyperboloid H^\wedge , the Euclidean space E , as well as the hypersphere S^\wedge :

$$\mathcal{M}^\wedge = \begin{cases} H^\wedge: \{\mathbf{x} \in \mathbb{R}^{+1}: \langle \mathbf{x}, \mathbf{x} \rangle^\wedge = 1/\wedge \cdot G_0 > 0\} & \text{for } \wedge < 0 \\ E^\wedge: \mathbb{R}^3 & \text{for } \wedge = 0 \\ S^\wedge: \{\mathbf{x} \in \mathbb{R}^{+1}: \langle \mathbf{x}, \mathbf{x} \rangle^\wedge = 1/\wedge\} & \text{for } \wedge > 0 \end{cases} \quad (1)$$

where $\langle \cdot, \cdot \rangle^\wedge$ denotes the curvature-aware scalar product, namely $\langle \mathbf{x}, \cdot \rangle^\wedge = \int_{\beta=0}^{\wedge} G_{\beta-\beta}$ for $\wedge > 0$ and $\langle \mathbf{x}, \cdot \rangle^\wedge = \langle \mathbf{x}, \cdot \rangle_{\mathcal{L}} = -G_0 + \int_{\beta=1}^{\wedge} G_{\beta-\beta}$ for $\wedge < 0$. For $\mathcal{M}^\wedge \in \{H^\wedge, S^\wedge\}$, a critical notion is **tangent space** [16], which is useful when the operators are not explicitly defined in the manifolds. The tangent space of \mathcal{M}^\wedge at $\mathbf{x} \in \mathcal{M}^\wedge$ is defined as a n -dimensional vector space approximating \mathcal{M}^\wedge around \mathbf{x} :

$$\mathcal{T}_\mathbf{x} \mathcal{M}^\wedge := \{\mathbf{v} \in \mathbb{R}^{+1}: \langle \mathbf{v}, \mathbf{x} \rangle^\wedge = 0\} \quad (2)$$

The mapping between manifold $\mathcal{M}^\wedge \in \{H^\wedge, S^\wedge\}$ and its tangent space $\mathcal{T}_\mathbf{x} \mathcal{M}^\wedge$ can be achieved by **exponential map** and **logarithmic map** [16]. The exponential map projects points from a subset of the tangent space $\mathcal{T}_\mathbf{x} \mathcal{M}^\wedge$ to the original manifold \mathcal{M}^\wedge , while the logarithmic map projects points in the opposite direction. For points $\mathbf{x}, \cdot \in \mathcal{M}^\wedge$, $\mathbf{v} \in \mathcal{T}_\mathbf{x} \mathcal{M}^\wedge$, such that $\mathbf{x} < \cdot$, $\mathbf{v} < \mathbf{0}$ and $\mathbf{0} = (1/\wedge) \cdot \mathbf{0} \dots \cdot \mathbf{0}$, the exponential map $\exp_\mathbf{x}^\wedge(\cdot)$ and logarithmic map $\log_\mathbf{x}^\wedge(\cdot)$ are formally defined as follows:

$$\begin{aligned} \exp_\mathbf{x}^\wedge(\mathbf{v}) &= \cos^\wedge\left(\frac{\rho_{\wedge} \|\mathbf{v}\|^\wedge}{\rho_{\wedge}}\right) \mathbf{x} + \sin^\wedge\left(\frac{\rho_{\wedge} \|\mathbf{v}\|^\wedge}{\rho_{\wedge}}\right) \frac{\mathbf{v}}{\|\mathbf{v}\|^\wedge} \\ \log_\mathbf{x}^\wedge(\cdot) &= \frac{\cos^\wedge^{-1}(\wedge \langle \mathbf{x}, \cdot \rangle^\wedge)}{\sin^\wedge \cos^\wedge^{-1}(\wedge \langle \mathbf{x}, \cdot \rangle^\wedge)} (\cdot - \wedge \langle \mathbf{x}, \cdot \rangle^\wedge \mathbf{x}) \end{aligned} \quad (3)$$

where $\|\mathbf{v}\|^\wedge = \frac{\rho_{\wedge}}{\rho_{\wedge}} \langle \mathbf{v}, \mathbf{v} \rangle^\wedge$ denotes the norm of tangent vector, and the curvature-aware trigonometric functions are:

$$\sin^\wedge = \begin{cases} \sin & \text{if } \wedge > 0 \\ \sinh & \text{if } \wedge < 0 \end{cases}, \quad \cos^\wedge = \begin{cases} \cos & \text{if } \wedge > 0 \\ \cosh & \text{if } \wedge < 0 \end{cases}$$

To connect vectors in tangent spaces, we use parallel transport $PT_{\mathbf{x} \rightarrow \cdot}: \mathcal{T}_\mathbf{x} \mathcal{M}^\wedge \rightarrow \mathcal{T}_\cdot \mathcal{M}^\wedge$ [16]:

$$PT_{\mathbf{x} \rightarrow \cdot}(\cdot) = \mathbf{v} - \frac{\wedge \langle \cdot, \mathbf{v} \rangle^\wedge}{1 + \wedge \langle \mathbf{x}, \cdot \rangle^\wedge} (\mathbf{x} + \cdot) \quad (4)$$

which ensures the transported vectors stay parallel to the connection. In addition, the geodesic distance is defined as:

$$3^\wedge(\mathbf{x}, \cdot) = \frac{1}{\rho_{\wedge}} \cos^\wedge^{-1}(\wedge \langle \mathbf{x}, \cdot \rangle^\wedge) \quad (5)$$

There are also several widely-used realizations of these spaces, i.e., Poincaré ball \mathcal{P}^\wedge for hyperbolic spaces, projected hypersphere

D^\wedge for spherical spaces. Please refer to [43] for more details. Moreover, for Euclidean space E^\wedge , let $\mathbf{x}, \cdot \in E^\wedge$ and $\mathbf{v} \in \mathcal{T}_\mathbf{x} E^\wedge$, the exponential map is defined as $\exp_\mathbf{x}(\mathbf{v}) = \mathbf{x} + \mathbf{v}$, and the logarithmic map is $\log_\mathbf{x}(\cdot) = \cdot - \mathbf{x}$. The parallel transport is defined as $PT_{\mathbf{x} \rightarrow \cdot}(\mathbf{v}) = \mathbf{v}$.

4 METHODOLOGY

In this section, we will introduce the details of the GDCF model. First, we will expound the basic model paradigm. Then, the implementations of several major components will be presented. Finally, we will discuss several key points within our proposal.

4.1 Mathematical Paradigm

In this subsection, we will introduce the basic mathematical paradigm of GDCF. Users may have diverse interests, and these interests can be attributed to various latent concepts. Therefore, we aim to learn a factorized representation for user D composed of high-level concepts. The factorized representation of user D is denoted as $\mathbf{z}_D = [z_D^{(1)}, z_D^{(2)}, \dots, z_D^{(\cdot)}]$, where $z_D^{(\cdot)} \in \mathcal{M}^\wedge$ is a 3-dimensional concept representation living in the manifold \mathcal{M}^\wedge . $z_D^{(\cdot)}$ depicts the preference of user D regarding to the \cdot th concept. Different from the diverse interests of users, an item usually represents a specific object, which is fixed and explicit [32]. Thus, items are represented by the 3-dimensional embeddings $\{h_\beta\}_{\beta=1}^n \in E^3$ in Euclidean spaces instead of the factorized vectors.

In order to model the correlations between items and latent concepts, we design the concept assignment relations for items: $\mathbf{C} = \{c_\beta\}_{\beta=1}^n$, where $c_\beta = [2_{\beta,1}, 2_{\beta,2}, \dots, 2_{\beta,\cdot}]$ and $2_{\beta,\cdot}$ indicates whether item β belongs to concept \cdot . The concept assignment of item β is described by a distribution $\mathcal{P}(2_{\beta,\cdot})$, which denotes the probability of item β belonging to concept \cdot and $\sum_{\cdot=1}^{\#} \mathcal{P}(2_{\beta,\cdot}) = 1$. The assignments among different concepts ensure that the user-item interactions can be reconstructed based on different manifolds.

Following previous works [17, 23, 31], we adopt the variational auto-encoders (VAEs) as the basic paradigms and propose a generative GDCF model. Assuming that the interactions are generated from the true world simulator using the factorized user embedding \mathbf{z}_D along with the concept assignment distribution $\mathcal{P}(\mathbf{C})$. An appropriate objective function should maximize the marginal likelihood of the user's observed data f_D in expectation over the distribution of user representation \mathbf{z}_D and item concept assignment \mathbf{C} :

$$\max_{\mathbf{z}_D, \mathbf{C}} E_{\mathcal{P}(\mathbf{C})} \mathcal{P}(f_D | \mathbf{z}_D, \mathbf{C}) \mathcal{P}(\mathbf{z}_D) \mathcal{P}(\mathbf{C}) \quad (6)$$

Given the historical interactions of user D : $f_D = \{\beta: E_{D\beta} = 1\}$, the inferred posterior configurations of \mathbf{z}_D and \mathbf{C} can be naturally described by a probability distribution $\mathcal{Q}(\mathbf{z}_D | f_D, \mathbf{C})$. Moreover, we are also interested in disentangling a user's preferences at a more granular level regarding various aspects of an item. This can be achieved by introducing a constraint over $\mathcal{Q}(\mathbf{z}_D | f_D, \mathbf{C})$ and matching $\mathcal{Q}(\mathbf{z}_D | f_D, \mathbf{C})$ to a unit Gaussian prior $\mathcal{P}(\mathbf{z}_D)$. Therefore, the constrained optimization problem can be formalized as:

$$\begin{aligned} \max_{\mathbf{z}_D, \mathbf{C}} E_{f_D \sim \mathcal{D}} E_{\mathcal{P}(\mathbf{C})} E_{\mathcal{Q}(\mathbf{z}_D | f_D, \mathbf{C})} \log \mathcal{P}(f_D | \mathbf{z}_D, \mathbf{C}) \\ \text{subject to } \int \mathcal{Q}(\mathbf{z}_D | f_D, \mathbf{C}) \|\mathcal{P}(\mathbf{z}_D) - \mathcal{Q}(\mathbf{z}_D | f_D, \mathbf{C})\| < n \end{aligned} \quad (7)$$

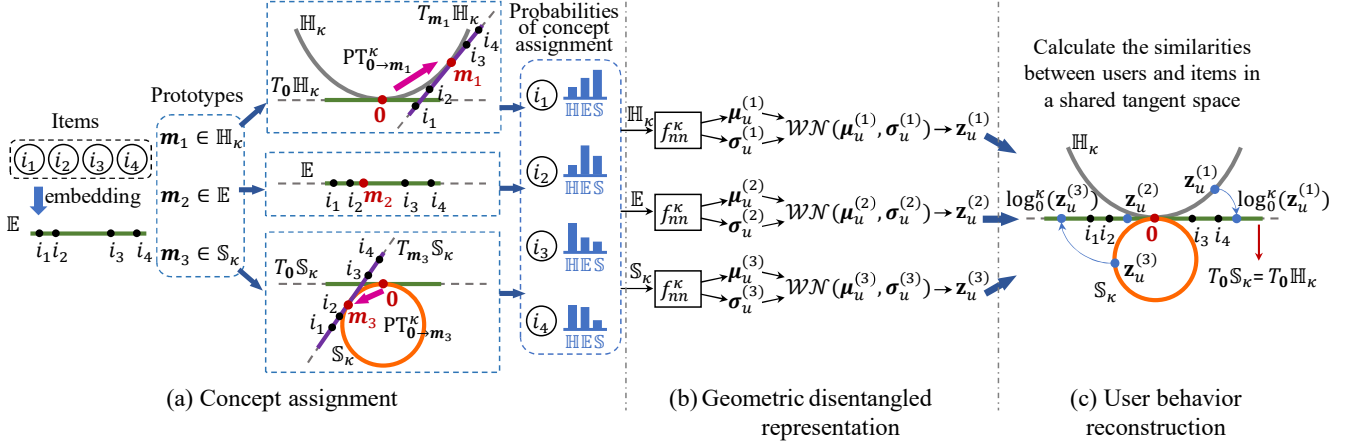


Figure 2: The framework of GDCF.

where $n \geq 0$ indicates the strength of the applied constraint. We rewrite Eq. (7) as a Lagrangian under the KKT conditions [19, 21]:

$$\mathcal{F}(\cdot; \mathbf{f}_D; \mathbf{Z}_D) = E_{\gamma(\cdot)} E_{\mathcal{M}(\cdot)} \log \gamma(\mathbf{f}_D | \mathbf{Z}_D; \mathbf{C}) - V \|\mathcal{M}(\mathbf{Z}_D | \mathbf{f}_D; \mathbf{C})\|_{\gamma(\cdot)} - n \cdot \quad (8)$$

where the KKT multiplier $V \geq 0$ is the regularization coefficient. Since $V \cdot n \geq 0$, according to the complementary slackness KKT condition, Eq. (8) can be re-written as [17]:

$$\mathcal{F}(\cdot; \mathbf{f}_D; \mathbf{Z}_D) \geq \mathcal{L}(\cdot; \mathbf{f}_D; \mathbf{Z}_D; V) = E_{\gamma(\cdot)} E_{\mathcal{M}(\cdot)} \log \gamma(\mathbf{f}_D | \mathbf{Z}_D; \mathbf{C}) - V \|\mathcal{M}(\mathbf{Z}_D | \mathbf{f}_D; \mathbf{C})\|_{\gamma(\cdot)} \quad (9)$$

which is the final optimization objective of our model.

4.2 Implementation

In this subsection, we will introduce the details of three major components in the implementation of the proposed GDCF model as shown in Figure 2.

4.2.1 Concept assignment. Concept assignment calculates the probabilities of an item assigning to different concepts. We propose a prototype-based concept assignment strategy to prevent over-parameterization and low sampling efficiency. Specifically, prototypes $\{\mathbf{m}_\beta\}_{\beta=1}^n$ are proposed to describe the anchors of the concepts across manifolds. Since items and the prototype \mathbf{m}_β might come from different geometries, we need to map them into a shared space for comparison. One straightforward approach is to assume all the items are represented in the shared tangent space at 0, i.e., $T_0 \mathcal{M}$, and then transform them into \mathcal{M} via $\exp_0^\wedge(\cdot)$. However, the transformed item representations would be identical and indistinguishable if the prototypes share the same latent space, i.e., $\mathcal{M}_\beta = \mathcal{M}_{\beta+1}$ [31]. Thus, the prototypes of items should be explicitly incorporated in the transformation process. As shown in Figure 2 (a), we assume that item representations $\{\mathbf{h}_\beta\}_{\beta=1}^n$ come from $T_0 \mathcal{M}_\beta$, and transform them into $T_0 \mathcal{M}$ via parallel transport to obtain

the prototype-aware item representations:

$$\mathbf{h}'_\beta = \text{PT}_{0 \rightarrow \mathbf{m}_\beta}^\wedge(\mathbf{h}_\beta) \quad (10)$$

Then, we define the similarity measurement between \mathbf{h}'_β and \mathbf{m}_β . A common strategy is to transform \mathbf{h}'_β into manifold \mathcal{M}_β via $\exp_{\mathbf{m}_\beta}^\wedge(\cdot)$. Similarities between \mathbf{h}'_β and \mathbf{m}_β can be measured by their geodesic distance on \mathcal{M}_β . We theoretically prove this approach is equivalent to a simpler operation, i.e., computing the norm of \mathbf{h}'_β on $T_0 \mathcal{M}_\beta$:

THEOREM 4.1. For $\mathbf{x}, \mathbf{y} \in \mathcal{M}$ and $\mathcal{M} \in \{\mathbb{E}, \mathbb{S}^\wedge, \mathbb{D}^\wedge, \mathbb{H}^\wedge, \mathbb{P}^\wedge\}$, the geodesic distance between \mathbf{x} and \mathbf{y} , is equal to the norm of $\log_{\mathbf{x}}^\wedge(\mathbf{y})$ in the tangent space $T_0 \mathcal{M}_\beta$, i.e., $\mathfrak{z}(\mathbf{x}, \mathbf{y}) = \|\log_{\mathbf{x}}^\wedge(\mathbf{y})\|$.

Proof of Theorem 4.1: For Euclidean space \mathbb{E} , it is easy to prove that $\|\log_{\mathbf{x}}^\wedge(\cdot)\| = \|\cdot - \mathbf{x}\| = \mathfrak{z}_{\mathbb{E}}(\mathbf{x}, \cdot)$. For the hypersphere model, let $\mathbf{v} = \log_{\mathbf{x}}^\wedge(\cdot)$, we have $\cdot = \exp_{\mathbf{x}}^\wedge(\mathbf{v})$. The distance between $\mathbf{x}, \cdot \in \mathbb{H}^\wedge$ can be rewritten as:

$$\begin{aligned} \mathfrak{z}_{\mathbb{H}^\wedge}(\mathbf{x}, \cdot) &= \mathfrak{z}_{\mathbb{H}^\wedge}(\mathbf{x}, \exp_{\mathbf{x}}^\wedge(\mathbf{v})) \\ &= \frac{1}{\sqrt{-\lambda}} \cosh^{-1} \left(\langle \mathbf{x}, \cosh(\sqrt{-\lambda} \|\mathbf{v}\|_{\mathcal{L}}) \mathbf{x} \right. \\ &\quad \left. + \sinh(\sqrt{-\lambda} \|\mathbf{v}\|_{\mathcal{L}}) \frac{\mathbf{v}}{\sqrt{-\lambda} \|\mathbf{v}\|_{\mathcal{L}}} \rangle_{\mathcal{L}} \right) \\ &= \frac{1}{\sqrt{-\lambda}} \cosh^{-1} \left(\frac{1}{\lambda} \cosh(\sqrt{-\lambda} \|\mathbf{v}\|_{\mathcal{L}}) \right) \\ &= \|\mathbf{v}\|_{\mathcal{L}} = \|\log_{\mathbf{x}}^\wedge(\cdot)\|_{\mathcal{L}} \end{aligned} \quad (11)$$

This theorem also holds for the hypersphere model \mathbb{S}^\wedge following above proposition. For the Poincaré ball model \mathbb{P}^\wedge , the distance between $\mathbf{x}, \cdot \in \mathbb{P}^\wedge$ can be rewritten as:

$$\begin{aligned} \mathfrak{z}_{\mathbb{P}^\wedge}(\mathbf{x}, \cdot) &= \mathfrak{z}_{\mathbb{P}^\wedge}(\mathbf{x}, \exp_{\mathbf{x}}^\wedge(\mathbf{v})) \\ &= \frac{2}{\sqrt{-\lambda}} \tanh^{-1} \left(\frac{\sqrt{-\lambda} \tanh \left(\frac{\sqrt{-\lambda} \|\mathbf{v}\|}{2} \right) \frac{1}{\sqrt{-\lambda}} \right) \\ &= \|\mathbf{v}\| = \|\log_{\mathbf{x}}^\wedge(\cdot)\|_{\mathcal{L}} \end{aligned} \quad (12)$$

The formulas about \mathbb{P}^\wedge are referred from [43]. The proof of the projected hypersphere model \mathbb{D}^\wedge is similar to the above process.

Based on Theorem 4.1, we can measure the similarities by the norm of \mathbf{h}'_θ as:

$$b_{\theta,:} = -\|\mathbf{h}'_\theta\|^\wedge \quad (13)$$

The probability of item θ belonging to concept c can be achieved after a softmax function:

$$p(z_{\theta,:}) = \frac{\exp(b_{\theta,:}/g)}{\sum_{c=1} \exp(b_{\theta,c}/g)} \quad (14)$$

where g is a hyper-parameter to control the scale of values, and we set $g \in (0,1)$ to obtain a more skewed distribution.

4.2.2 Geometric disentangled representation. Users may have diverse interests, leading to the disentangled representations w.r.t. different concepts. The prior distribution $p_\lambda(\mathbf{z}_D)$ is set to a unit Gaussian distribution. The encoder $\mathcal{E}_\lambda(\mathbf{z}_D|\mathbf{f}_D \cdot \mathbf{C})$ learns the representations for user D based on the high-level concept distributions. Assume $\mathcal{E}_\lambda(\mathbf{z}_D|\mathbf{f}_D \cdot \mathbf{C}) = \prod_{c=1}^m \mathcal{E}_\lambda(\mathbf{z}_D^{(c)}|\mathbf{f}_D \cdot \mathbf{C})$, and each $\mathcal{E}_\lambda(\mathbf{z}_D^{(c)}|\mathbf{f}_D \cdot \mathbf{C})$ is defined as a multivariate normal distribution with a diagonal co-variance matrix [17]. Let $\mathbf{t}_D = [s_{D,1} \cdot s_{D,2} \cdots \cdot s_{D,m}] \odot [p(z_{1,:}) \cdot p(z_{2,:}) \cdots \cdot p(z_{m,:})] \in \mathbb{E}^m$ be the probabilities of the interacted items for user D belonging to c^{th} concept, where \odot is the point-wise product. The mean vector $\mathbf{a}_D^{(c)} \in \mathcal{M}_c^3$ and the standard deviation $\mathbf{b}_D^{(c)} \in \mathcal{T}_D \mathcal{M}_c^3$ are calculated by a neural network $\mathcal{F}_{\text{nn}}^\wedge : \mathcal{M}_c^m \rightarrow \mathcal{M}_c^{23}$ with input \mathbf{t}_D :

$$\begin{aligned} (\mathbf{a}_D^{(c)} \cdot \mathbf{b}_D^{(c)}) &= \mathcal{F}_{\text{nn}}^\wedge \exp_0^\wedge(\mathbf{t}_D) \\ \mathbf{b}_D^{(c)} &= \mathbf{a}_D^{(c)} \cdot \mathbf{2}_D^{(c)} = \mathbf{f}_0 \cdot 4^{(-\frac{1}{2} \log_0^\wedge(\mathbf{b}_D^{(c)}))} \end{aligned} \quad (15)$$

Here we implement the $\mathcal{F}_{\text{nn}}^\wedge$ as non-Euclidean neural networks [3, 10, 53]. Specifically, for the projected manifolds, i.e., \mathcal{D}^\wedge and \mathcal{P}^\wedge , we follow the neural structures proposed in [3, 10], which leverage the \log_0^\wedge and \exp_0^\wedge to transform the m -dimensional representations between tangent spaces and the manifolds. The Euclidean operations are conducted over all the m elements for the m -dimensional features. However, a m -dimensional feature vector in \mathcal{H}^\wedge or \mathcal{S}^\wedge has $m+1$ elements. If the Euclidean operations \mathcal{F}_{nn} are directly applied on all the $m+1$ dimensions, the transformed features would be out of the tangent spaces [53]. To address this challenge, we propose to conduct the Euclidean operations on the last m elements. The non-Euclidean operations defined in \mathcal{P}^\wedge or \mathcal{D}^\wedge are equivalent to those defined in \mathcal{H}^\wedge or \mathcal{S}^\wedge , respectively. Please refer to Appendix A.1 for the proposed operations and proofs.

As shown in Figure 2 (b), based on the obtained $\mathbf{a}_D^{(c)}$ and $\mathbf{2}_D^{(c)}$, we leverage the wrapped normal distribution [37, 43] to build the Gaussian distribution on Riemannian manifolds to generate user representations. Specifically, to sample instances from the distribution $\mathcal{WN}(\mathbf{a}_D^{(c)} \cdot \mathbf{2}_D^{(c)})$, we need to first sample instances from the tangent space of the manifold $\mathcal{T}_D \mathcal{M}_c$, i.e., $\mathbf{v} \sim \mathcal{N}(\mathbf{0}, \mathbf{2}_D) \in \mathcal{T}_D \mathcal{M}_c$. The geometric disentangled representation $\mathbf{z}_D^{(c)}$ can be obtained via parallel transporting \mathbf{v}_D from $\mathcal{T}_D \mathcal{M}_c$ to $\mathcal{T}_{\mathbf{0}} \mathcal{M}_c$, and then projecting it to the manifold $\mathcal{T}_{-D} \mathcal{M}_c$ by exponential map:

$$\mathbf{z}_D^{(c)} = \exp_{-\mathbf{0}}^\wedge \left(\text{PT}_{\mathbf{0} \rightarrow -\mathbf{0}}^\wedge(\mathbf{v}_D) \right) \quad (16)$$

4.2.3 User behavior reconstruction. Here we aim to predict which items are most likely to be clicked by the target user. Since the item representations live in the Euclidean space, we can view this space as a tangent space and transform the item representations $\{\mathbf{h}_\theta\}_{\theta=1}^m$ to \mathcal{M}_c , and then compute their similarities based on the distances. However, the distances in different spaces may have different value scales. For example, the distances in Euclidean or hyperbolic spaces could be theoretically infinite, while for spherical spaces, the distances are finite [16]. To tackle this challenge, we propose to project the user representations $\mathbf{z}_D^{(c)}$ to the tangent space $\mathcal{T}_D \mathcal{M}_c$. As shown in Figure 2 (c), the tangent spaces of these manifolds are isomorphic to Euclidean spaces. Therefore, we can leverage Euclidean measurements to calculate the similarities of different geometries. Here we only consider the angle between the representations to avoid the influence of the norm of the representations in different spaces. Therefore, the Euclidean inner product can be employed to measure the similarities between the normalized c^{th} concept of user D and item θ in the shared tangent space $\mathcal{T}_D \mathcal{M}_c$:

$$B(\mathbf{z}_D^{(c)} \cdot \mathbf{h}_\theta) = \left\langle \frac{\log_0^\wedge(\mathbf{z}_D^{(c)}) \cdot \mathbf{h}_\theta}{\|\log_0^\wedge(\mathbf{z}_D^{(c)})\| \|\mathbf{h}_\theta\|} \right\rangle \quad (17)$$

User behaviors can be further reconstructed via the categorical distribution over the m items: $p_\lambda(s_{D,\theta}|\mathbf{z}_D \cdot \mathbf{C}) \propto \sum_{c=1}^m B(\mathbf{z}_D^{(c)} \cdot \mathbf{h}_\theta)$.

4.3 Optimization

The trainable parameters λ of GDCF include concept prototypes $\{\mathbf{m}_c \in \mathcal{M}_c^3\}_{c=1}^m$, item representations $\{\mathbf{h}_\theta \in \mathbb{E}^3\}_{\theta=1}^m$, and the parameters in the neural network $\mathcal{F}_{\text{nn}}^\wedge$. We optimize λ to maximize the training objective in Eq. (9) via Adam [22] and RiemannianAdam [4] for Euclidean and non-Euclidean parameters, respectively. Moreover, since a 3-dimensional embedding in manifold $\mathcal{M}_c \in \{\mathcal{H}^\wedge \cdot \mathcal{S}^\wedge\}$ has $3+1$ elements, we concatenate or remove an element "0" at the first dimension of features in the tangent spaces to ensure the vector operations can be conducted in a principal manner.

4.4 Discussions

In this section, we will discuss two key points, i.e., similarity measurements and probability distributions on Riemannian manifolds, as well as the relations to some related methods including disentangled VAEs and mix-curvature representation learning.

Similarity measurement. We leverage two types of metrics to measure the similarities: the geodesic distance for concept assignment in Eq. (13), and the Euclidean distance on the tangent spaces for user behavior reconstruction in Eq. (17). These two metrics are employed for different purposes. Concept assignment aims to infer the most likely manifold for item θ while the geodesic distance thrives on measuring the manifold-specific similarity. Thus, the metric used in the concept assignment is set to geodesic distance. For the behavior reconstruction, we need to measure the similarities between item representations and disentangled user representations on all the manifolds in a universal manner. The disentangled user representations are projected to a shared tangent space, where the Euclidean distance can be utilized as the closeness metric.

Probability distribution on Riemannian manifolds. Several approaches have been proposed to generalize the Normal distribution

Table 1: Statistics of five datasets.

Dataset	MovieLens-100k	MovieLens-1M	MovieLens-20M	AliShop	LastFM
#users	603	6,038	136,677	10,668	1,862
#items	5,697	3,605	20,108	20,591	14,795
#interactions	47,922	836,452	9,990,030	767,493	89,805
#held-out users	50	500	10,000	4,000	200

to Riemannian manifolds, including restricted normal, Riemannian normal, and wrapped normal [9, 13, 37]. Restricted normal distributions restrict a point of the ambient space sampled from a Gaussian distribution to the manifolds. Von Mises-Fisher distribution [9] is a kind of restricted normal, but it has only a single scalar covariance parameter and cannot parameterize covariance for different dimensions. Riemannian normal distributions [13, 34, 35] are based on geodesic distance in the manifolds. These distributions aim to maximize the entropy distributions and resemble the Gaussian distribution’s properties the closest. However, the resemble process is quite computationally expensive, leading to low efficiency. Wrapped normal distributions [37, 43] first sample features from a Gaussian distribution in the tangent space and the sampled features are further transformed to Riemannian manifolds via parallel transport and exponential map. We select wrapped normal distributions in our model due to its computational efficiency. Note that we do not follow the wrapped normal distributions when calculating the probability density function. This is because the density functions are expected to be directly computed on the tangent space instead of the manifolds.

Connections to disentangled VAEs. VAEs have been widely exploited to learn the disentangled representations [17, 18, 20, 31]. Existing disentangled VAEs usually leverage the cosine function to measure the similarities, which ignore the norms of representations. Thus, these disentangled VAEs can be viewed as representation learning on the hypersphere [31]. The key difference between these methods and our proposal is that GDCF learns the representations on different Riemannian manifolds. Multiple Riemannian manifolds endow our proposal with powerful modeling capacity. The major difference between MacridVAE and GDCF is that MacridVAE only focuses on a single manifold (i.e., hypersphere) while GDCF is capable of disentangling representations for different types of manifolds. In addition, MacridVAE aims at learning the one-hot concept assignment relations \mathbf{C} , while GDCF does not require the \mathbf{C} to be one-hot vectors, leading to a general solution with less restrictions.

Connections to mix-curvature representation learning. Mix-curvature methods [3, 11, 43] learn the representations by balancing the expressive capacities of different kinds of manifolds. Cross-manifold representations learned by existing mix-curvature methods are defined in the product spaces, i.e., $\mathcal{M}_1 \times \mathcal{M}_2 \times \dots \times \mathcal{M}_n$. GDCF is fundamentally different from these methods because the representations of GDCF on different manifolds are independent from each other, i.e., $\mathcal{M}_1 \cdot \mathcal{M}_2 \cdot \dots \cdot \mathcal{M}_n$, which can preserve the uniqueness of different geometries.

Complexity analysis. The time complexity of GDCF is $\mathcal{O}(3\# \cdot \mathcal{Z})$, where \mathcal{Z} is the number of concepts, \mathcal{Z} is the dimension of user/item

representations, $\#$ and \mathcal{Z} are the number of users and items, respectively. The time complexity is on par with the SOTA disentangled models like MacridVAE [31].

5 EXPERIMENTS

5.1 Experimental Setup

Datasets. We conduct the experiments on five real-world datasets, including three MovieLens datasets with different numbers of interactions (MovieLens-100k, MovieLens-1M and MovieLens-20M) [12], AliShop [31] and LastFM [6]. The training/validation/testing partitions are following previous works [27, 31]. The statistics of these datasets are shown in Table 1.

Baselines. We select the following SOTA CF models as baselines:

- NGCF [48] is a graph-based CF model to incorporate the high-order connectivity of user-item interactions.
- LightGCN [14] is a SOTA CF recommendation model based on graph convolution network.
- DGCF [49] is a disentangled CF model to learn representations for different latent user intentions via graph convolutional network.
- MacridVAE [31] is also a disentangled model, which leverages VAE to disentangle macro and micro components for user behaviors.
- HyperML [46] is a hyperbolic CF model, which explores metric learning in hyperbolic spaces for recommendation.
- H-VAE [36] is a hyperbolic VAE based CF approach, which employs a hyperbolic VAE to solve the collaborative filtering problem.
- HGCF [45] is a hyperbolic CF method, which is devised to capture the high-order correlations between users and items by using hyperbolic GNN.
- $\hat{\Lambda}$ -GCN [3] is a mix-curvature GNN. We modify its loss to the hinge loss to solve the CF problem.

Hyper-parameter. We perform the hyper-parameter search on a small validation set. For the proposed GDCF, we set the dimension of embedding as 100 and tune the temperature $g \in [0.0^1]$, $V \in [0.100]$ in Eq. (9), the number of factors $\in \{1.2.3. \dots .20\}$, the standard deviation of the prior $f_0 \in [0.075.0.5]$ and the number of hidden layers $\in \{0.1.2.3\}$ by grid search. We set $\hat{\Lambda} \in \{-1.0.1\}$ for hyperbolic, Euclidean and spherical geometry, respectively. Moreover, we also tune the following parameters for all methods: learning rate $\in [10^{-6}.10^{-2}]$, dropout probability $\in [0.1]$, l_2 regularization strength $\in [10^{-8}.10^{-2}]$. We tune these hyper-parameters automatically via Optuna [2].

Infrastructure. We implement our model with Pytorch, and conduct the experiments with: CPU: Intel(R) Xeon(R) Platinum 8168 CPU @ 2.70GHz GPU: 24 \times NVIDIA Tesla V100.

5.2 Main Results

We evaluate the performance of our approach on the five popular datasets. Here we implement GDCF with three kinds of manifolds i.e., H^\wedge , E , and S^\wedge . To evaluate the effectiveness on top-K recommendation task, we adopt two popular evaluation metrics: Recall@20 and NDCG@20 [14, 48, 49]. We repeat this process 5 times and report the average scores.

Table 2: The recommendation performance comparison. Best results are in bold and second best results are underlined.

Datasets	MovieLens-100k		MovieLens-1M		MovieLens-20M		Alishop		LastFM	
Metric	NDCG@20	Recall@20	NDCG@20	Recall@20	NDCG@20	Recall@20	NDCG@20	Recall@20	NDCG@20	Recall@20
NGCF	0.2437	0.2997	0.2956	0.2852	0.2979	0.3214	0.1058	0.1175	0.2681	0.2924
LightGCN	<u>0.2596</u>	0.3112	0.3115	0.3016	0.3172	0.3498	0.1573	0.1512	<u>0.3116</u>	<u>0.3258</u>
HyperML	0.2160	0.2549	0.2719	0.2839	0.2631	0.2809	0.1023	0.1061	0.2546	0.2852
H-VAE	0.2157	0.2706	0.2729	0.3049	0.2405	0.2737	0.1087	0.1101	0.2667	0.2795
HGCF	0.2218	0.2701	0.2844	0.2955	0.3093	0.3536	0.1115	0.1135	0.2706	0.2985
\wedge -GCN	0.1912	0.2408	0.2443	0.2351	0.2155	0.2419	0.0846	0.0816	0.2489	0.2582
DGCF	0.2572	<u>0.3148</u>	<u>0.3252</u>	0.3128	0.3193	0.3626	0.1556	0.1547	0.3062	0.3242
MacridVAE	0.2310	0.3095	<u>0.3241</u>	<u>0.3605</u>	<u>0.3205</u>	<u>0.3965</u>	<u>0.1831</u>	<u>0.1862</u>	0.3072	0.3213
GDCF	0.2652	0.3261	0.3314	0.3637	0.3317	0.4052	0.1874	0.1886	0.3173	0.3296

Experimental results are summarized in Table 2, and we can observe that the proposed GDCF achieves SOTA performance on the five datasets across all metrics, demonstrating the effectiveness of the geometric disentangled representation learning. GDCF outperforms the hyperbolic CF methods (e.g., HyperML, H-VAE and HGCF) by a large margin, which verifies that a single hyperbolic space may be insufficient to capture the hybrid geometric characteristics. The disentanglement-based methods (DGCF and MacridVAE) outperform other baselines in most cases, which proves the benefits of factorized representations under the CF scenario. By incorporating the geometric concepts into the disentangled representation learning, GDCF surpasses DGCF and MacridVAE over all the datasets, confirming the superior modeling capacity of our proposal.

5.3 Ablation Study

Since three types of geometries are incorporated in the GDCF model, we further conduct an ablation study to investigate the effectiveness of different geometry combinations. Here we leverage hyperboloid model H^\wedge and Poincaré ball P^\wedge to describe hyperbolic geometry, hypersphere model S^\wedge and projected hypersphere D^\wedge to describe spherical geometry, respectively. We evaluate the performance of GDCF with three types of geometry combinations, and the subscripts are omitted for simplicity: the single geometry $H \cdot P \cdot E \cdot S \cdot D$; the combinations of two geometries $H \& E \cdot P \& E \cdot E \& S \cdot E \& D \cdot H \& S \cdot P \& D$; and the combinations of three geometries $H \& E \& S$, $P \& E \& D$. The numbers of components for different geometries are set to approximately equal.

The results are reported in Table 3. The best results of all the methods are marked in bold, while the best results within each type of geometry combinations are marked with underlines. One can see that GDCF implemented with three kinds of geometries achieve the best performance in most cases, which demonstrates that all the three geometries contribute to learning expressive and meaningful representations. For the single geometry-based variations, non-Euclidean methods (H, P, S and D) generally outperform Euclidean-based model (E). It demonstrates that the non-Euclidean geometric structures are ubiquitous under the recommendation scenario and our proposal can effectively encode such patterns into representations. Moreover, the simplified version of GDCF with only two geometries (H&E and H&S) already achieve promising

performance on MovieLens datasets, revealing the effectiveness of representation learning in multiple geometries. Ablation models implemented with H or S consistently perform better than the counterparts with P and D, which is reasonable as P and D suffer from numerical instability in the optimization process [7, 29, 39].

5.4 Universal Curvature Analysis

Previous ablation studies raise an interesting question: how to find the most effective combination of these geometries. The simplest approach is to enumerate all the possible combinations and select the candidate with best performance. This enumeration method would be optimal but impractical due to the inefficiency and low scalability [43]. Here we propose an approximate method to tackle this challenge. Specifically, we randomly initialize the curvatures of the components within $[-1, 1]$ and view the curvatures \wedge of all components as trainable parameters in the training process. Moreover, the sign of curvature is also not constrained. Thus, components are able to change their geometries from hyperbolic (spherical) to spherical (hyperbolic) space to achieve the optimal solutions. Note that $\|0\|^\wedge = \|(1/\sqrt{|\wedge|} \cdot 0 \dots 0)\|^\wedge \xrightarrow{\wedge \rightarrow 0} \pm\infty$, for the hyperboloid H^\wedge and hypersphere S^\wedge [43]. Therefore, the universal curvature variation of GDCF leverages the Poincaré ball P and the projected hypersphere D for hyperbolic and spherical geometry, respectively. GDCF with the universal curvature is denoted as U . We are interested in comparing U with $P \& D$, since both of them leverage Poincaré ball P and projected hypersphere D .

Table 4 exhibits the results. The combinations with the best performance are marked in bold. One can see that the variations with universal curvature U outperform $P \& D$ in most cases, which indicates U can automatically select more proper geometries. However, U does not converge to some specific curvatures [43], leading to the low generalization. Another limitation of U is that the Euclidean component cannot be incorporated as it cannot ensure \wedge will converge to 0 in the training process.

5.5 Parameter Sensitivity Analysis

In this subsection, we will investigate the effect of two key parameters (the number of concepts and the dimension of embedding \mathcal{J}) on the MovieLens-100k and LastFM datasets. Figure 3(a) and 3(b) show the results of GDCF under varying . With the increase of the number of concepts , the performance of GDCF first improves

Table 3: Ablation study of the geometry combinations. Best results of all the methods are in bold, and the best results within each type of geometry combinations are underlined.

Datasets Metric	MovieLens-100k		MovieLens-1M		MovieLens-100M		Alishop		LastFM	
	NDCG@20	Recall@20	NDCG@20	Recall@20	NDCG@20	Recall@20	NDCG@20	Recall@20	NDCG@20	Recall@20
H	0.2340	0.3127	0.3236	0.3539	0.3223	0.3858	<u>0.1870</u>	<u>0.1833</u>	<u>0.3106</u>	<u>0.3275</u>
P	0.2353	0.3063	0.3245	0.3517	0.3216	0.3827	0.1862	0.1824	0.309	0.3225
E	0.2252	0.2965	0.3253	0.3548	0.3207	0.3851	0.1805	0.1797	0.3023	0.3172
S	<u>0.2361</u>	<u>0.3141</u>	<u>0.3266</u>	<u>0.3554</u>	<u>0.3239</u>	<u>0.3892</u>	0.1728	0.1733	0.3076	0.3212
D	0.2357	0.3046	0.3196	0.3501	0.3216	0.3864	0.1734	0.1751	0.3018	0.3176
H&E	<u>0.2620</u>	0.3135	0.3263	0.3597	0.3263	0.3991	0.1885	0.1833	0.3147	<u>0.3288</u>
P&E	0.2542	0.3118	0.3257	0.3540	0.3248	0.3954	0.1875	0.1826	0.3105	0.3238
E&S	0.2517	0.3126	0.3261	0.3595	0.3250	0.3973	0.1798	0.1770	<u>0.3149</u>	0.3281
E&D	0.2523	<u>0.3231</u>	0.3254	0.3538	0.3244	0.3922	0.1757	0.1785	0.3067	0.3192
H&S	0.2582	0.3155	<u>0.3271</u>	<u>0.3599</u>	<u>0.3295</u>	<u>0.4006</u>	0.1841	0.1804	0.3141	0.3239
P&D	0.2526	0.3164	0.3241	0.3562	0.3266	0.3934	0.1818	0.1794	0.3099	0.3227
P&E&D	0.2549	0.3196	0.3256	0.3572	0.3284	0.3986	0.1811	0.1802	0.3124	0.3276
H&E&S	<u>0.2652</u>	<u>0.3261</u>	<u>0.3314</u>	<u>0.3637</u>	<u>0.3317</u>	<u>0.4052</u>	<u>0.1874</u>	0.1886	<u>0.3173</u>	<u>0.3296</u>

Table 4: Results of the universal curvature analysis.

Dataset	Metric	P&D	U	Best
MovieLens-100k	NDCG@20	0.2526	0.2564	H&E&S
	Recall@20	0.3164	0.3218	H&E&S
MovieLens-1M	NDCG@20	0.3241	0.3261	H&R&S
	Recall@20	0.3562	0.3557	H&E&S
MovieLens-20M	NDCG@20	0.3266	0.3301	H&E&S
	Recall@20	0.3934	0.4049	H&E&S
AliShop	NDCG@20	0.1818	0.1821	H&E
	Recall@20	0.1794	0.1814	H&E&S
LastFM	NDCG@20	0.3099	0.3134	H&E&S
	Recall@20	0.3227	0.3305	U

and then drops on both datasets across all metrics. This phenomenon suggests that a suitable number of concepts is needed for GDCF, while redundant concepts may introduce noise and lead to the performance decay. Also, figure 3(c) and 3(d) present the results of GDCF under varying β . With the increase of the dimension of embeddings β , model performance first improves and then keeps steady. This is reasonable as representations of proper dimensions are already capable of fully modeling the interactions, and too large dimensions may increase model complexity and effect the training efficiency.

5.6 Case Study

In this subsection, we will investigate whether GDCF can capture the correlations between the latent concepts and the geometric structures. To this end, we provide two cases (items with ID i961 and i4837) in MovieLens-100k dataset. Figure 4 displays the 2-hop ego-graphs centered with different items. One can see that the ego-graph of i961 is closer to the tree-likeness structures with less cycles compared to i4837. After model training, the concept assignment probabilities calculated by GDCF on different geometries are shown

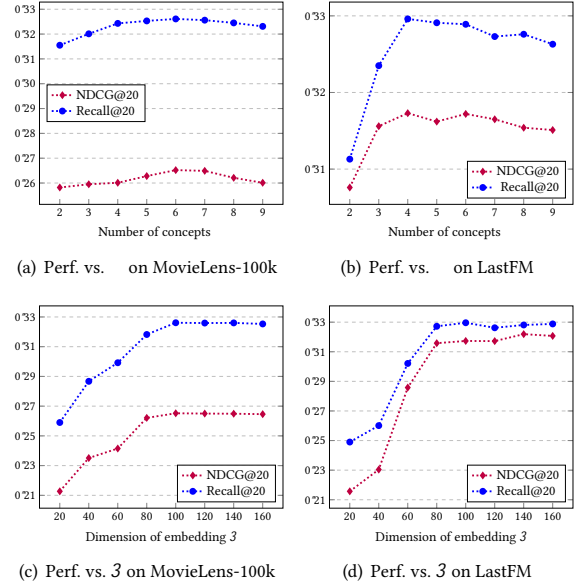


Figure 3: Results of the parameter sensitivity analysis.

in Figure 4(c) and Figure 4(d). These probabilities reveal that the hyperbolic model H is more important in modeling tree-structures (i961) and spherical model S is crucial in modeling cyclical structures (i4837). This phenomenon demonstrates that our proposal is capable of effectively capturing the intrinsic correlations between the latent concepts and the geometric structures.

6 RELATED WORK

In this section, we will briefly summarize two related techniques: the disentangled representation learning and the non-Euclidean representation learning, along with their applications on the recommendation scenario.

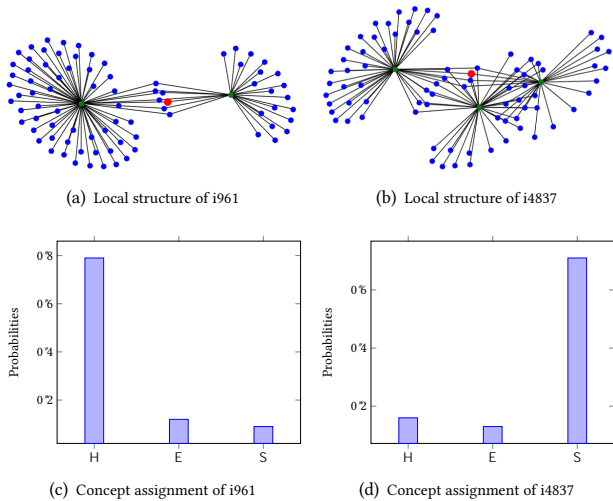


Figure 4: Local topological structures and concept assignments of two cases in MovieLens-100k. The center nodes, 1-hop nodes, 2-hop nodes are marked by red, green, blue circles, respectively.

6.1 Disentangled Representation Learning

Disentangled representation learning aims at learning factorized representations that reveal and disentangle the latent intent factors beneath the input dataset, which has been widely used in a myriad of applications. Higgins et al. [17] propose beta-VAE to learn interpretable factorised latent representations from raw image data in an unsupervised manner. Ma et al. [30] introduce the disentangled graph convolutional network to learn disentangled node representations. Yang et al. [51] propose a factorizable graph convolutional network for explicitly disentangling the intertwined relations encoded in a graph. Recently, disentangled representation learning is introduced into the recommendation scenario to boost performance. MacridVAE [31] is designed to infer the high-level concepts associated with user intentions. Wang et al. [49] further devise a disentangled graph collaborative filtering model to encode the high-order interactions into the factorized representations. For the sequential recommendation, Ma et al. [33] introduce a sequence-to-sequence training strategy based on latent self-supervision and disentanglement. By combining the curriculum learning with the disentangled representation learning, Curriculum Disentangled Recommendation model [8] is proposed to efficiently learn disentangled representations from complex and noisy multi-feedbacks. Despite the promising performance, existing disentangled approaches learn representations solely within the Euclidean geometry and ignore the hybrid geometric characteristics of the user-item interactions.

6.2 Non-Euclidean Representation Learning

A myriad of data exhibits the highly non-Euclidean latent anatomy [3, 25, 40, 50], and non-Euclidean geometries (i.e., hyperbolic or spherical geometry) are more suitable for modeling data with non-Euclidean characteristics compared to Euclidean geometry. Nickel

et al. [38] first leverage the hyperbolic space to learn hierarchical representations of symbolic data. HGCN [7] is an inductive hyperbolic GCN designed for hierarchical and scale-free graphs. Mixed-curvature Variational Autoencoder [43] exploits both the hyperbolic and spherical geometries to better modeling data characteristics. In the recommendation scenario, existing works usually focus on modeling interactions with hyperbolic models [46]. Feng et al. [46] propose to learn the representations of check-in activities in a hyperbolic space. Mirvakhabova et al. [36] propose an autoencoder based on hyperbolic geometry for solving collaborative filtering problem. HGCF [45] is devised to capture the high-order correlations between users and items by integrating hyperbolic geometry into graph neural networks. Most existing works focus on learning representations in a single non-Euclidean space, which may be incapable of modeling the hybrid geometric interactions as discussed in the introduction.

7 CONCLUSION

In this paper, we study the novel problem of geometric disentangled collaborative filtering. Different from existing single geometry-based CF models, we propose a novel GDCF model to incorporate multiple types of geometries to disentangle the sophisticated hybrid patterns of user-item interactions. Specifically, GDCF learns factorized representations that uncover and disentangle the intent factors hidden in the historical interactions, which is capable of learning expressive user/item representations. Extensive experiments are conducted over five publicly available datasets, and the experimental results demonstrate the superiority of GDCF.

ACKNOWLEDGMENTS

This work is supported in part by the National Natural Science Foundation of China (No. U20B2045, 62192784, 62172052, 62002029, 61772082).

REFERENCES

- [1] Gediminas Adomavicius and Alexander Tuzhilin. 2005. Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE transactions on knowledge and data engineering* 17, 6 (2005), 734–749.
- [2] Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. 2019. Optuna: A Next-generation Hyperparameter Optimization Framework. In *SIGKDD*.
- [3] Gregor Bachmann, Gary Bécigneul, and Octavian-Eugen Ganea. 2020. Constant Curvature Graph Convolutional Networks. In *ICML*.
- [4] Gary Bécigneul and Octavian-Eugen Ganea. 2019. Riemannian Adaptive Optimization Methods. In *ICLR*.
- [5] Marcel Berger. 2012. *A panoramic view of Riemannian geometry*. Springer Science & Business Media.
- [6] Iván Cantador, Peter Brusilovsky, and Tsvi Kuflik. 2011. 2nd Workshop on Information Heterogeneity and Fusion in Recommender Systems (HetRec 2011). In *RecSys*. ACM.
- [7] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. 2019. Hyperbolic graph convolutional neural networks. In *NeurIPS*. 4869–4880.
- [8] Hong Chen, Yudong Chen, Xin Wang, Ruobing Xie, Rui Wang, Feng Xia, and Wenwu Zhu. 2021. Curriculum Disentangled Recommendation with Noisy Multi-feedback. *NeurIPS* 34 (2021).
- [9] Tim R Davidson, Luca Falorsi, Nicola De Cao, Thomas Kipf, and Jakub M Tomczak. 2018. Hyperspherical variational auto-encoders. *arXiv preprint arXiv:1804.00891* (2018).
- [10] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic neural networks. In *NeurIPS*. 5350–5360.
- [11] Albert Gu, Frederic Sala, Beliz Gunel, and Christopher Ré. 2018. Learning mixed-curvature representations in product spaces. In *ICLR*.

- [12] F Maxwell Harper and Joseph A Konstan. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)* 5, 4 (2015), 1–19.
- [13] Søren Hauberg. 2018. Directional statistics with the spherical normal distribution. In *FUSION*. 704–711.
- [14] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [15] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In *Proceedings of the 26th international conference on world wide web*. 173–182.
- [16] Sigurdur Helgason. 1979. *Differential geometry, Lie groups, and symmetric spaces*. Vol. 80. Academic press.
- [17] Irina Higgins, Loic Matthey, Arka Pal, Christopher Burgess, Xavier Glorot, Matthew Botvinick, Shakir Mohamed, and Alexander Lerchner. 2017. beta-vaes: Learning basic visual concepts with a constrained variational framework. In *ICLR*. 1024–1034.
- [18] Jun-Ting Hsieh, Bingbin Liu, De-An Huang, Li Fei-Fei, and Juan Carlos Niebles. 2018. Learning to decompose and disentangle representations for video prediction. In *NeurIPS*. 515–524.
- [19] Kuhn HW and T AW. 1951. Nonlinear programming. In *Proceedings of 2nd Berkeley Symposium, Berkeley: University of California Press*. 481–492.
- [20] Vineet John, Lili Mou, Hareesh Bahuleyan, and Olga Vechtomova. 2019. Disentangled Representation Learning for Non-Parallel Text Style Transfer. In *ACL*. 424–434.
- [21] William Karush. 1939. *Minima of Functions of Several Variables with Inequalities as Side Conditions*. Master’s thesis. Department of Mathematics, University of Chicago, Chicago, IL, USA.
- [22] Diederik P Kingma and Jimmy Ba. 2015. Adam: A method for stochastic optimization. *ICLR* (2015).
- [23] Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114* (2013).
- [24] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.
- [25] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. 2010. Hyperbolic geometry of complex networks. *Physical Review E* 82, 3 (2010), 036106.
- [26] Xiaopeng Li and James She. 2017. Collaborative variational autoencoder for recommender systems. In *SIGKDD*. 305–314.
- [27] Dawen Liang, Rahul G Krishnan, Matthew D Hoffman, and Tony Jebara. 2018. Variational autoencoders for collaborative filtering. In *WWW*. 689–698.
- [28] Nathan Linial, Eran London, and Yuri Rabinovich. 1995. The geometry of graphs and some of its algorithmic applications. *Combinatorica* 15, 2 (1995), 215–245.
- [29] Qi Liu, Maximilian Nickel, and Douwe Kiela. 2019. Hyperbolic graph neural networks. In *NeurIPS*. 8228–8239.
- [30] Jianxin Ma, Peng Cui, Kun Kuang, Xin Wang, and Wenwu Zhu. 2019. Disentangled graph convolutional networks. In *ICML*. PMLR, 4212–4221.
- [31] Jianxin Ma, Chang Zhou, Peng Cui, Hongxia Yang, and Wenwu Zhu. 2019. Learning disentangled representations for recommendation. In *NeurIPS*. 5711–5722.
- [32] Jianxin Ma, Chang Zhou, Hongxia Yang, Peng Cui, Xin Wang, and Wenwu Zhu. 2020. Disentangled Self-Supervision in Sequential Recommenders. In *SIGKDD*. 483–491.
- [33] Jianxin Ma, Chang Zhou, Hongxia Yang, Peng Cui, Xin Wang, and Wenwu Zhu. 2020. Disentangled self-supervision in sequential recommenders. In *SIGKDD*. 483–491.
- [34] Emile Mathieu, Charline Le Lan, Chris J Maddison, Ryota Tomioka, and Yee Whye Teh. 2019. Continuous Hierarchical Representations with Poincaré Variational Auto-Encoders. *arXiv preprint arXiv:1901.06033* (2019).
- [35] Emile Mathieu and Maximilian Nickel. 2020. Riemannian Continuous Normalizing Flows. In *NeurIPS*. 2503–2515.
- [36] Leyla Mirvakhabova, Evgeny Frolov, Valentin Khulkov, Ivan Oseledets, and Alexander Tuzhilin. 2020. Performance of hyperbolic geometry models on top-N recommendation tasks. In *RecSys*. 527–532.
- [37] Yoshihiro Nagano, Shoichiro Yamaguchi, Yasuhiro Fujita, and Masanori Koyama. 2019. A wrapped normal distribution on hyperbolic space for gradient-based learning. In *ICML*. 4693–4702.
- [38] Maximilian Nickel and Douwe Kiela. 2017. Poincaré embeddings for learning hierarchical representations. In *NeurIPS*. 6338–6347.
- [39] Maximilian Nickel and Douwe Kiela. 2018. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. In *ICML*. 3779–3788.
- [40] Wei Peng, Tuomas Varanka, Abdelrahman Mostafa, Henglin Shi, and Guoying Zhao. 2021. Hyperbolic deep neural networks: A survey. *arXiv preprint arXiv:2101.04562* (2021).
- [41] Frederic Sala, Chris De Sa, Albert Gu, and Christopher Ré. 2018. Representation tradeoffs for hyperbolic embeddings. In *ICML*. 4457–4466.
- [42] Aravind Sankar, Yozen Liu, Jun Yu, and Neil Shah. 2021. Graph Neural Networks for Friend Ranking in Large-scale Social Platforms. In *WWW*. 2535–2546.
- [43] Ondrej Skopec, Octavian-Eugen Ganea, and Gary Bécigneul. 2019. Mixed-curvature Variational Autoencoders. In *ICLR*.
- [44] Xiaoyuan Su and Taghi M Khoshgoftaar. 2009. A survey of collaborative filtering techniques. *Advances in artificial intelligence* 2009 (2009).
- [45] Jianing Sun, Zhaoyue Cheng, Saba Zuberi, Felipe Pérez, and Maksims Volkovs. 2021. HGCF: Hyperbolic Graph Convolution Networks for Collaborative Filtering. In *WWW*. 593–601.
- [46] Lucas Vinh Tran, Yi Tay, Shuai Zhang, Gao Cong, and Xiaoli Li. 2020. HyperML: A Boosting Metric Learning Approach in Hyperbolic Space for Recommender Systems. In *WSDM*. 609–617.
- [47] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative deep learning for recommender systems. In *SIGKDD*. 1235–1244.
- [48] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In *SIGIR*. 165–174.
- [49] Xiang Wang, Hongye Jin, An Zhang, Xiangnan He, Tong Xu, and Tat-Seng Chua. 2020. Disentangled Graph Collaborative Filtering. In *SIGIR*. 1001–1010.
- [50] Richard C Wilson, Edwin R Hancock, Elżbieta Pekalska, and Robert PW Duin. 2014. Spherical and hyperbolic embeddings of data. *IEEE transactions on pattern analysis and machine intelligence* 36, 11 (2014), 2255–2269.
- [51] Yiding Yang, Zunlei Feng, Mingli Song, and Xinchao Wang. 2020. Factorizable Graph Convolutional Networks. *NeurIPS* 33 (2020).
- [52] Sixiao Wang, Hongxu Chen, Xiao Ming, Lizhen Cui, Hongzhi Yin, and Guandong Xu. 2021. Where are we in embedding spaces? A Comprehensive Analysis on Network Embedding Approaches for Recommender Systems. *arXiv preprint arXiv:2105.08908* (2021).
- [53] Yiding Zhang, Xiao Wang, Chuan Shi, Nian Liu, and Guojie Song. 2021. Lorentzian Graph Convolutional Networks. In *WWW*. 1249–1261.
- [54] Yu Zheng, Chen Gao, Xiang Li, Xiangnan He, Depeng Jin, and Yong Li. 2020. Disentangling User Interest and Conformity for Recommendation with Causal Embedding. *arXiv preprint arXiv:2006.11011* (2020).

A SUPPLEMENTARY MATERIAL

A.1 Neural Networks in the hypersphere model

To ensure the transformed features satisfy the spherical geometry, we define hypersphere matrix-vector multiplication to perform the hypersphere matrix-vector multiplication:

DEFINITION A.1. (Hypersphere matrix-vector multiplication) *If $M : \mathbb{R}^e \rightarrow \mathbb{R}^c$ is a linear map with matrix representation, given two points $\mathbf{x} = (G_0 \cdot \dots \cdot G_e) \in S_{\mathbb{R}^e}$, $\mathbf{v} = (E_0 \cdot \dots \cdot E_e) \in \mathcal{T}_0 S_{\mathbb{R}^e}$, we have:*

$$M^{\otimes \wedge}(\mathbf{x}) = \exp_0^{\wedge}(\hat{M}(\log_0^{\wedge}(\mathbf{x}))) \cdot \hat{M}(\mathbf{v}) = (0 \cdot M(E_1 \cdot \dots \cdot E_e)) \quad (18)$$

Let M be a $\langle \times \rangle$ matrix, M' be a $\langle \times \rangle$ matrix, $\mathbf{x} \in S_{\mathbb{R}^e}$, $M^{\otimes \wedge} \mathbf{x} := M^{\otimes \wedge}(\mathbf{x})$, we have matrix associativity as: $(M' M)^{\otimes \wedge} \mathbf{x} = M' \otimes^{\wedge} (M^{\otimes \wedge} \mathbf{x})$.

The hypersphere matrix-vector multiplication also satisfies following theorem:

THEOREM A.1. *Given a point in spherical space, which is represented by $\mathbf{x}_{S_{\mathbb{R}^e}}^- \in S_{\mathbb{R}^e}$ using hypersphere model or $\mathbf{x}_{D_{\mathbb{R}^e}}^- \in D_{\mathbb{R}^e}$ using projected hypersphere ball model [3], respectively. Let M be a $\langle \times \rangle$ matrix, hypersphere matrix-vector multiplication $M^{\otimes \wedge} \mathbf{x}_{S_{\mathbb{R}^e}}^-$ used in hyperboloid model is equivalent to Möbius matrix-vector multiplication $M^{\otimes \wedge} \mathbf{x}_{D_{\mathbb{R}^e}}^-$ used in projected hypersphere model.*

Proof: Let $\mathbf{x}_{S_{\mathbb{R}^e}}^- \in S_{\mathbb{R}^e}$, $\mathbf{v} = (E_0 \cdot E_1 \cdot \dots \cdot E_e) \in \mathcal{T}_0 S_{\mathbb{R}^e}$, and M be a $\langle \times \rangle$ matrix, hypersphere matrix-vector multiplication is shown as follows:

$$M^{\otimes \wedge} \mathbf{x}_{S_{\mathbb{R}^e}}^- := M^{\otimes \wedge}(\mathbf{x}_{S_{\mathbb{R}^e}}^-) = \exp_0^{\wedge} \hat{M} \log_0^{\wedge}(\mathbf{x}_{S_{\mathbb{R}^e}}^-) = \sim_{S_{\mathbb{R}^e}}^{-} \cdot \hat{M}(\mathbf{v}) = 0 \cdot M(E_1 \cdot \dots \cdot E_e) \quad (19)$$

Let $\mathbf{x}_{D^{\bar{\lambda}}}^- \in D^{\bar{\lambda}}$, Möbius matrix-vector multiplication has the formulation as [3]:

$$\begin{aligned} \mathbf{M} \otimes^{\wedge} \mathbf{x}_{D^{\bar{\lambda}}}^- &:= \mathbf{M}^{\otimes^{\wedge}}(\mathbf{x}_{D^{\bar{\lambda}}}^-) \\ &= (1/\sqrt{\kappa}) \tan \frac{\|\mathbf{M}\mathbf{x}_{D^{\bar{\lambda}}}^-\|}{\|\mathbf{x}_{D^{\bar{\lambda}}}^-\|} \tan^{-1}(\sqrt{\kappa}\|\mathbf{x}_{D^{\bar{\lambda}}}^-\|) \frac{\mathbf{M}\mathbf{x}_{D^{\bar{\lambda}}}^-}{\|\mathbf{M}\mathbf{x}_{D^{\bar{\lambda}}}^-\|} = \sim_{D^{\bar{\lambda}}}^{\leq} \end{aligned} \quad (20)$$

For $\%_{S^{\bar{\lambda}} \rightarrow D^{\bar{\lambda}}}(\mathbf{x}_{S^{\bar{\lambda}}}^-) = \mathbf{x}_{D^{\bar{\lambda}}}^-$ and a shared $\times =$ matrix \mathbf{M} , we aim to prove $\%_{S^{\bar{\lambda}} \rightarrow D^{\bar{\lambda}}}(\sim_{S^{\bar{\lambda}}}^{\leq}) = \sim_{D^{\bar{\lambda}}}^{\leq}$

For $\mathbf{x}_{S^{\bar{\lambda}}}^- = (G_{0S^{\bar{\lambda}}} \cdot G_{1S^{\bar{\lambda}}} \cdot \dots \cdot G_{=S^{\bar{\lambda}}}) \in S^{\bar{\lambda}}$, let $\hat{\mathbf{x}}_{S^{\bar{\lambda}}} = (G_{1S^{\bar{\lambda}}} \cdot \dots \cdot G_{=S^{\bar{\lambda}}})$, and the logarithmic map of $\mathbf{x}_{S^{\bar{\lambda}}}^-$ at $\mathbf{0} = (1/\sqrt{\kappa} \cdot \mathbf{0} \cdot \dots \cdot \mathbf{0}) \in S^{\bar{\lambda}}$ is shown as follows:

$$\log_0^{\wedge}(\mathbf{x}_{S^{\bar{\lambda}}}^-) = \frac{\cos^{-1}(\sqrt{\kappa}G_{0S^{\bar{\lambda}}})}{1 - \kappa G_{0S^{\bar{\lambda}}}^2} (\mathbf{0} \cdot \hat{\mathbf{x}}_{S^{\bar{\lambda}}}) = \frac{\cos^{-1}(\sqrt{\kappa}G_{0S^{\bar{\lambda}}})}{\sqrt{\kappa}\|\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|} (\mathbf{0} \cdot \hat{\mathbf{x}}_{S^{\bar{\lambda}}}) \quad (21)$$

Let $\otimes = \frac{\cos^{-1}(\sqrt{\kappa}G_{0S^{\bar{\lambda}}})}{\sqrt{\kappa}\|\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|}$, $\log_0^{\wedge}(\mathbf{x}_{S^{\bar{\lambda}}}^-) = \otimes(\mathbf{0} \cdot \hat{\mathbf{x}}_{S^{\bar{\lambda}}})$, so we have:

$$\hat{\mathbf{M}} \log_0^{\wedge}(\mathbf{x}_{S^{\bar{\lambda}}}^-) = (\mathbf{0} \cdot \otimes \mathbf{M} \hat{\mathbf{x}}_{S^{\bar{\lambda}}}) = \mathbf{m} \quad (22)$$

The hypersphere matrix-vector multiplication is given as following:

$$\mathbf{M} \otimes^{\wedge} \mathbf{x}_{S^{\bar{\lambda}}}^- = \frac{1}{\sqrt{\kappa}} \cos(\sqrt{\kappa}\|\mathbf{m}\|) \cdot \frac{\sin(\sqrt{\kappa}\|\mathbf{m}\|)}{\sqrt{\kappa}\|\mathbf{m}\|} \mathbf{M} \hat{\mathbf{x}}_{S^{\bar{\lambda}}} = \sim_{S^{\bar{\lambda}}}^{\leq} \quad (23)$$

Then we map $\sim_{S^{\bar{\lambda}}}^{\leq}$ to the projected hypersphere model:

$$\%_{S^{\bar{\lambda}} \rightarrow D^{\bar{\lambda}}}(\sim_{S^{\bar{\lambda}}}^{\leq}) = \frac{1}{\sqrt{\kappa}} \tan \frac{\|\mathbf{M}\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\| \cos^{-1}(\sqrt{\kappa}G_{0S^{\bar{\lambda}}})}{2\|\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|} \frac{\mathbf{M}\hat{\mathbf{x}}_{S^{\bar{\lambda}}}}{\|\mathbf{M}\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|} \quad (24)$$

Note that $\|\mathbf{m}\| = \frac{\rho}{\langle \mathbf{m} \cdot \mathbf{m} \rangle} = \frac{\rho}{\|\otimes \mathbf{M} \hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|}$, and $G_{0S^{\bar{\lambda}}} = \frac{\rho}{1 + \kappa \|\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|^2}$

Moreover, the point $\mathbf{x}_{D^{\bar{\lambda}}}^- = (G_{1D^{\bar{\lambda}}} \cdot \dots \cdot G_{=D^{\bar{\lambda}}}) \in D^{\bar{\lambda}}$ can be mapped into the hypersphere model as following:

$$\begin{aligned} \%_{D^{\bar{\lambda}} \rightarrow S^{\bar{\lambda}}}(\mathbf{x}_{D^{\bar{\lambda}}}^-) &= \frac{(1/\sqrt{\kappa} - \sqrt{\kappa}\|\mathbf{x}_{D^{\bar{\lambda}}}^-\|^2 \cdot 2G_{1D^{\bar{\lambda}}} \cdot \dots \cdot 2G_{=D^{\bar{\lambda}}})}{1 + \kappa \|\mathbf{x}_{D^{\bar{\lambda}}}^-\|^2} \\ &= (G_{0S^{\bar{\lambda}}} \cdot G_{1S^{\bar{\lambda}}} \cdot \dots \cdot G_{=S^{\bar{\lambda}}}) \end{aligned} \quad (25)$$

Thus, the squared norm of $\hat{\mathbf{x}}_{S^{\bar{\lambda}}}$ is given as:

$$\|\hat{\mathbf{x}}_{S^{\bar{\lambda}}}\|^2 = \frac{\Theta}{\beta=1} \frac{2}{1 + \kappa \|\mathbf{x}_{D^{\bar{\lambda}}}^-\|^2} (G_{D^{\bar{\lambda}}})^2 = \frac{2\|\mathbf{x}_{D^{\bar{\lambda}}}^-\|^2}{1 + \kappa \|\mathbf{x}_{D^{\bar{\lambda}}}^-\|^2} \quad (26)$$

By combining Eq. (24) and Eq. (26), we have:

$$\begin{aligned} \%_{S^{\bar{\lambda}} \rightarrow D^{\bar{\lambda}}}(\sim_{S^{\bar{\lambda}}}^{\leq}) \\ &= (1/\sqrt{\kappa}) \tan \frac{\|\mathbf{M}\mathbf{x}_{D^{\bar{\lambda}}}^-\|}{\|\mathbf{x}_{D^{\bar{\lambda}}}^-\|} \tan^{-1}(\sqrt{\kappa}\|\mathbf{x}_{D^{\bar{\lambda}}}^-\|) \frac{\mathbf{M}\mathbf{x}_{D^{\bar{\lambda}}}^-}{\|\mathbf{M}\mathbf{x}_{D^{\bar{\lambda}}}^-\|} = \sim_{D^{\bar{\lambda}}}^{\leq} \end{aligned} \quad (27)$$

Therefore, hypersphere matrix-vector multiplication is equivalence to Möbius matrix-vector multiplication.

Also, to apply **non-linear transformation** on the hypersphere model, the hypersphere pointwise non-linear activation can be derived as:

DEFINITION A.2. Hypersphere pointwise non-linear activation If $f: \mathbb{R}^{\bar{\lambda}} \rightarrow \mathbb{R}^{\bar{\lambda}}$ is a pointwise non-linearity map, given two points $\mathbf{x} = (G_0 \cdot \dots \cdot G_{\bar{\lambda}}) \in S^{\bar{\lambda}}$ and $\mathbf{v} = (E_0 \cdot \dots \cdot E_{\bar{\lambda}}) \in \mathcal{T}_0 S^{\bar{\lambda}}$, the hypersphere version $f_{\otimes^{\wedge}}$ is:

$$f_{\otimes^{\wedge}}(\mathbf{x}) = \exp_0^{\wedge}(f_{\otimes^{\wedge}}(\log_0^{\wedge}(\mathbf{x}))) \cdot f_{\otimes^{\wedge}}(\mathbf{v}) = (\mathbf{0} \cdot f(E_1) \cdot \dots \cdot f(E_{\bar{\lambda}})) \quad (28)$$

The hypersphere pointwise non-linear activation has the following property:

THEOREM A.2. Given a point in hyperbolic space, it is modeled by $\mathbf{x}^{\bar{V}} \in H^{\bar{V}}$ using hyperboloid model and $\mathbf{x}^{\bar{U}} \in D^{\bar{U}}$ using Poincaré ball model, respectively. Lorentzian pointwise non-linearity $f^{\otimes^{\wedge V}}(\mathbf{x}^{\bar{V}})$ in the hyperboloid model is equivalent to Möbius pointwise non-linearity $f^{\otimes^{\wedge U}}(\mathbf{x}^{\bar{U}})$ in the Poincaré ball model [10], when $f(\cdot)$ indicates some specific non-linear activation, e.g., Relu, leaklyRelu.

The proof of Theorem A.2 is similar to Theorem A.1, and we omit it due to the page length limit.