

Modeling Scale-free Graphs with Hyperbolic Geometry for Knowledge-aware Recommendation

Yankai Chen¹, Menglin Yang¹, Yingxue Zhang²,
Mengchen Zhao², Ziqiao Meng¹, Jianye Hao², Irwin King¹

¹The Chinese University of Hong Kong, ²Huawei Noah's Ark Lab

{ykchen,mlyang,zqmeng,king}@cse.cuhk.edu.hk,{yingxue.zhang,zhaomengchen,haojianye}@huawei.com

ABSTRACT

Aiming to alleviate data sparsity and cold-start problems of traditional recommender systems, incorporating knowledge graphs (KGs) to supplement auxiliary information has recently gained considerable attention. Via unifying the KG with *user-item interactions* into a tripartite graph, recent works explore the graph topologies to learn the low-dimensional representations of users and items with rich semantics. These real-world tripartite graphs are usually *scale-free*, however, the intrinsic hierarchical graph structures of which are underemphasized in existing works, consequently, leading to suboptimal recommendation performance. To address this issue and provide more accurate recommendation, we propose a knowledge-aware recommendation method with *Lorentz model* of the hyperbolic geometry, namely *Lorentzian Knowledge-enhanced Graph convolutional networks for Recommendation (LKGR)*. LKGR facilitates better modeling of scale-free tripartite graphs after the data unification. Specifically, we employ different information propagation strategies in the hyperbolic space to explicitly encode heterogeneous information from historical interactions and KGs. Additionally, our proposed knowledge-aware attention mechanism enables the model to automatically measure the information contribution, producing the coherent information aggregation in the hyperbolic space. Extensive experiments on three real-world benchmarks demonstrate that LKGR outperforms state-of-the-art methods by 3.6-15.3% of Recall@20 on Top-K recommendation.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender system, Knowledge graph, Hyperbolic geometric

ACM Reference Format:

Yankai Chen¹, Menglin Yang¹, Yingxue Zhang², Mengchen Zhao², Ziqiao Meng¹, Jianye Hao², Irwin King¹. 2022. Modeling Scale-free Graphs with Hyperbolic Geometry for Knowledge-aware Recommendation. In *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining (WSDM '22)*, February 21–25, 2022, Tempe, AZ, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3488560.3498419>

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 ACM 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](https://permissions.acm.org).

WSDM '22, February 21–25, 2022, Tempe, AZ, USA

© 2022 Association for Computing Machinery.

ACM ISBN 978-1-4503-9132-0/22/02...\$15.00

<https://doi.org/10.1145/3488560.3498419>




Figure 1: (a) Degree distribution of two real benchmarks. (b) Distance comparison in Euclidean and hyperbolic spaces.

1 INTRODUCTION

To alleviate the data sparsity and cold-start problems in traditional recommender systems [12, 16, 26, 32, 34, 44], incorporating *knowledge graphs* (KGs) into the recommender systems as side information has attracted growing attention in recent years [35–39]. A KG is a heterogeneous graph, where nodes represent *entities* (i.e., products or items, as well as related attributes and properties) and edges represent mutual *relations* between entities. Instead of relying on user-item historical records only, recommender systems (RS) extracting rich relational information in KGs can well compensate for the sparsity. Recently, several works [35–39] develop *graph convolutional networks* (GCNs) in recommendation, thanks to their capability of modeling complex data dependency into the graph format and summarizing the semantic information behind the topology.

Major motivation. These GCN-based models for knowledge-aware recommendation [35–39] usually unify the historical *user-item interactions* with KGs into the tripartite graphs, as shown in Figure 2(a). Then they learn the latent representations of users and items to estimate their matching probabilities in the Euclidean space. However, after the data unification, these tripartite graphs present the *scale-free* (hierarchical) graph characteristic, which is ignored by existing works [35–39]. We analyze the real-world benchmarks that are widely studied in these works and show the degree distributions of two large datasets in Figure 1(a). These two representative benchmarks are used to recommend movies¹ and restaurants² (details are in Section 4.1). We observe that these graphs approximate the power-law distribution. According to Bourgain's theorem [23], Euclidean space is however unable to obtain comparably low distortion for tree-like (power-law distributed) data [31]. Consequently, traditional graph embedding in the Euclidean space may not effectively capture the intrinsic hierarchical structures of these scale-free

¹MovieLens-20M: <https://grouplens.org/datasets/movielens/>

²Meituan-Dianping: <https://www.dianping.com/>

graphs, which leads to the high distortion of node embeddings and ultimately suppresses the recommendation performance [2, 17, 28].

To address this problem, we model these scale-free graphs with *Lorentz model* of the hyperbolic geometry for knowledge-aware recommendation. We propose an end-to-end model, namely *Lorentzian Knowledge-enhanced Graph convolutional networks for Recommendation (LKGR)*. Generally, LKGR learns better representations for users, items, and KG entities for recommendation. Concretely:

- LKGR projects node representations onto the Lorentzian manifold, i.e., one specific Riemannian manifold. It lives in the hyperbolic space, i.e., a continuous tree space with exponential volume growth property, producing less distortion for scale-free data with intrinsic hierarchical structures [2, 17, 28, 41]. As shown in Figure 1(b), in the hyperbolic space, graph nodes that are closer to the graph center show a smaller distance, while nodes near the graph boundary present a larger distance. This exponentially-evolved distance measurement actually fits well with the tree-like graph structure, where nodes with small degrees can be viewed as leaves on the boundary, and nodes with large degrees are like roots located in the central positions. On the contrary, Euclidean space embeddings are not position-sensitive and thus may not capture the latent information of the graph hierarchy.
- At information propagation stage, unlike previous work [36, 37] mainly focusing on mining KGs, LKGR summarizes the *interactive signals* from the user-item interactions and extracts *informative knowledge* from KGs simultaneously. Then we detach these two heterogeneous information from each other to update the node embeddings accordingly. This is main because of the heterogeneity of graph nodes under the recommendation scenarios, which is also different from most previous works that conduct undifferentiated convolutional operations to all nodes. These interactive signals are vital that directly reveal the user preferences and item targeting customers; along with the information extracted from KGs, LKGR further enriches the embeddings of users and items with diverse information components.
- We propose Lorentzian Knowledge-aware Attention Mechanism by considering the local graph structures on the Lorentzian manifold. LKGR with our proposed attention mechanism can well weigh the relative importance of neighboring information and selectively propagate information on the associated manifold. In addition, LKGR is equipped with high-order information propagation techniques in the hyperbolic space, which enables it to be extensible for different recommendation benchmarks.

LKGR learns the hierarchical structures of scale-free graphs with the hyperbolic geometry, and coherently summarizes interactive signals and knowledge associations into the low-dimensional embeddings. We extensively evaluate LKGR in three real-world scenarios of book, movie, and restaurant recommendation, compared to recent state-of-the-art methods. Experimental results not only prove the effectiveness of all proposed model components, but also demonstrate the superiority of LKGR over compelling baselines: we achieve the improved performance by 3.61-15.32% of Recall metric for Top-20 recommendation.

Organization. We define the problem in Section 2 and present the detailed methodology of LKGR model in Section 3. In Section 4,

we then report the experimental results. Finally, we review the related works in Section 5 and conclude the paper in Section 6.

2 PROBLEM AND MOTIVATION

2.1 Problem Formulation

A KG is formally defined as a set of the knowledge triplets: $\{(e_1, r, e_2) | e_1, e_2 \in \mathcal{E}, r \in \mathcal{R}\}$, denoting that relation r connects entity e_1 and e_2 . \mathcal{E} and \mathcal{R} represent the set of entities and relations. The KG is used to provide external knowledge for items, e.g., (*The Godfather*, *ActedBy*, *Marlon Brando*). User-item interactions can be similarly represented as: $\{(u, r^*, i) | u \in \mathcal{U}, i \in \mathcal{I}\}$. \mathcal{U} and \mathcal{I} denote the set of users and items, respectively, and r^* generalizes all interaction types, e.g., *browse*, *click*, or *purchase*, as one relation between u and i . The user-item interaction matrix $Y \in \mathcal{R}^{|\mathcal{U}| \times |\mathcal{I}|}$ is defined according to user-item interaction, where $y_{u,i} = 1$ indicates there is an observed interaction, otherwise $y_{u,i} = 0$. Moreover, each item can be matched with an entity in the KG to elucidate alignments between items and entities [38, 39]. Unifying user behaviors and item knowledge into the *Unified Knowledge Graph (UKG)*, which essentially is a tripartite graph and can be defined as $\mathcal{G} = \{(h, r, t) | h, t \in \mathcal{E}', r \in \mathcal{R}'\}$ where $\mathcal{E}' = \mathcal{E} \cup \mathcal{U}$ and $\mathcal{R}' = \mathcal{R} \cup \{r^*\}$. For example, the UKG shown in Figure 2(a) unifies various relations as well as users, items, and entities in a graph for movie recommendation.

Notations. In this paper, we use the bold lowercase characters, bold uppercase characters and calligraphy characters to denote the vectors, matrices and sets, respectively. Non-bold characters are used to denote scalars or graph nodes.

Task description. Given the UKG, the recommendation task studied in this paper is to train a RS model estimating the probability $\hat{y}_{u,i}$ that user u may adopt item i .

3 PROPOSED LKGR APPROACH

3.1 Hyperbolic Geometry Preliminaries

Hyperbolic geometry is a non-Euclidean geometry with a constant negative curvature measuring how a geometric object deviates from a flat plane [30]. In this paper, we adopt the **Lorentz model** that well describes hyperbolic geometry, for its unique simplicity and numerical stability [1, 18, 24, 27].

Lorentzian manifold and tangent space. Let $\langle \cdot, \cdot \rangle_{\mathbb{H}}: \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ represent the Lorentzian inner product in the hyperbolic space:

$$\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{H}} = -x_0 y_0 + x_1 y_1 + \dots + x_{d-1} y_{d-1}. \quad (1)$$

For ease of presentation, the d -dimensional Lorentzian manifold is denoted by $\mathbb{H}^{d,c}$ with the negative curvature $-1/c$, and the Euclidean tangent space with d dimensions centered at vector $\mathbf{x} \in \mathbb{H}^{d,c}$ is denoted by $\mathbb{T}_{\mathbf{x}}^{d,c}$. $\mathbb{T}_{\mathbf{x}}^{d,c}$ is a local, first-order approximation of the Lorentzian manifold at \mathbf{x} , which is useful to perform graph convolutional operations in the hyperbolic space [1].

Exponential and logarithmic mappings. Hyperbolic and tangent space can be bridged by exponential and logarithmic mappings. Given $\mathbf{x}, \mathbf{y} \in \mathbb{H}^{d,c}$ ($\mathbf{x} \neq \mathbf{y}$), $\mathbf{z} \in \mathbb{T}_{\mathbf{x}}^{d,c}$ ($\mathbf{z} \neq \mathbf{0}$), the exponential mapping $\Pi_{\mathbf{x}}^{exp,c}(\mathbf{z}): \mathbb{T}_{\mathbf{x}}^{d,c} \rightarrow \mathbb{H}^{d,c}$, maps \mathbf{z} to the hyperbolic space; the reverse logarithmic mapping projects vectors back to the tangent space

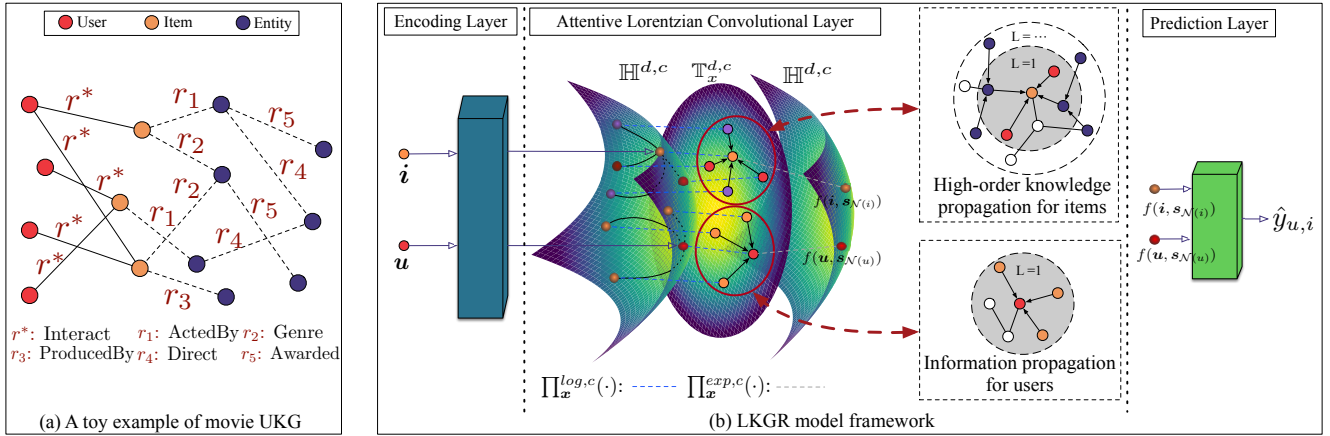


Figure 2: (a) The tripartite graph modeling users, items and KG entities. (b) Illustration of the proposed LKGR model.

centered at \mathbf{x} , where $\|z\|_{\mathbb{H}} = \sqrt{\langle z, z \rangle_{\mathbb{H}}}$ is the norm of z :

$$\begin{aligned} \prod_{\mathbf{x}}^{exp,c}(z) &= \cosh\left(\frac{\|z\|_{\mathbb{H}}}{\sqrt{c}}\right)\mathbf{x} + \sqrt{c} \sinh\left(\frac{\|z\|_{\mathbb{H}}}{\sqrt{c}}\right) \frac{z}{\|z\|_{\mathbb{H}}}, \\ \prod_{\mathbf{x}}^{log,c}(\mathbf{y}) &= \sqrt{c} \operatorname{arccosh}\left(-\frac{\langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{H}}}{c}\right) \cdot \frac{\mathbf{y} + \frac{1}{c} \langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{H}} \mathbf{x}}{\|\mathbf{y} + \frac{1}{c} \langle \mathbf{x}, \mathbf{y} \rangle_{\mathbb{H}} \mathbf{x}\|_{\mathbb{H}}}. \end{aligned} \quad (2)$$

3.2 LKGR Model Overview

In the following content, we introduce the proposed Lorentzian Knowledge-enhanced Graph Convolutional Networks (LKGR) in detail. Figure 2(b) illustrates the framework of the LKGR model. Generally, it consists of three components:

- **Encoding Layer.** In practice, input embeddings may live in the Euclidean space or hyperbolic space. If they are in the Euclidean space, Encoding Layer first projects them to the hyperbolic space to make sure they are on the Lorentzian manifold.
- **Attentive Lorentzian Convolutional Layer.** Aiming to accurately profile the latent user-item representations, the Lorentzian Convolutional Layer respectively updates their embeddings based on their sampled ego-networks, e.g., Figures 2(b) (unsampled nodes are colored white). In the hyperbolic space, users and items propagate *interactive signals* back and forth, which actually simulates the collaborative filtering effect. To automatically learn the relative importance of information from the KG, we design *Lorentzian Knowledge-aware Attention Mechanism*, which enables the selective and biased information aggregation in the hyperbolic space. Furthermore, by stacking multiple attentive Lorentzian Convolutional Layers, LKGR can explicitly explore the high-order KG subgraphs to further extract distant knowledge.
- **Prediction Layer.** To achieve the efficient estimation in the matching stage, our Prediction Layer directly collects the learned representations of users and items and outputs the matching score, by conducting the inner product in the hyperbolic space.

3.3 Encoding Layer

If input embeddings are in the Euclidean space, before inputting to the following layers, we first need to explicitly encode the Euclidean input onto the Lorentzian manifold. Let $\mathbf{x}_{\mathbb{E}} \in \mathbb{R}^{d-1}$ and

$\mathbf{x}_{\mathbb{H}} \in \mathbb{H}^d$ denote Euclidean inputs and the transformed hyperbolic feature embedding, respectively. $\mathbf{x}_{\mathbb{H}}$ can be encoded as follows:

$$\mathbf{x}_{\mathbb{H}} = \prod_{\mathbf{o}}^{exp,c}([0, \mathbf{x}_{\mathbb{E}}]) = \left[\sqrt{c} \cosh\left(\frac{\|\mathbf{x}_{\mathbb{E}}\|_2}{\sqrt{c}}\right), \sqrt{c} \sinh\left(\frac{\|\mathbf{x}_{\mathbb{E}}\|_2}{\sqrt{c}}\right) \frac{\mathbf{x}_{\mathbb{E}}}{\|\mathbf{x}_{\mathbb{E}}\|_2} \right], \quad (3)$$

where $(0, \mathbf{x}_{\mathbb{E}})$ is a d -dimensional vector in (Euclidean) tangent space and vector $\mathbf{o} = \{\sqrt{c}, 0, \dots, 0\} \in \mathbb{H}^{d,c}$ denotes the origin in $\mathbb{H}^{d,c}$. \mathbf{o} is used as a reference vector to perform tangent space operations. In the context of LKGR, we set the curvature $-1/c$ as a trainable variable, which dynamically measures how hierarchical the embedding space is [1]. Unless otherwise specified, we use \mathbf{x} to denote $\mathbf{x}_{\mathbb{H}}$ in the following sections.

Encoding Layer enables our model to be compatible with upstream Euclidean inputs. To evaluate the holistic performance, e.g., recommendation accuracy, training time, of all proposed LKGR modules, in this paper, we initialize the node embeddings in the Euclidean space. Experimental details can be found in Section 4.3.

3.4 Attentive Lorentzian Convolutional Layer

We describe a single layer on Lorentzian Knowledge-aware Attention, information propagation, and neighbor aggregation.

3.4.1 Lorentzian Knowledge-aware Attention. Our proposed attention mechanism considers the local structures of knowledge triplets on the Lorentzian manifold. We first compute the attentive weight $\pi(h, r, t)$ of an edge between entity h and entity t connected by relation r :

$$\pi(h, r, t) = \prod_{\mathbf{o}}^{log,c}(\mathbf{h})^T \mathbf{W}_r \prod_{\mathbf{o}}^{log,c}(\mathbf{t}), \quad (4)$$

which is further normalized, denoted by $\hat{\pi}(h, r, t)$, across all edges connected with h by adopting the softmax function:

$$\hat{\pi}(h, r, t) = \frac{\exp(\pi(h, r, t))}{\sum_{t' \in \mathcal{N}(h)} \exp(\pi(h, r, t'))}. \quad (5)$$

Our attention mechanism depends on the node embedding \mathbf{h} , \mathbf{t} and weight matrix \mathbf{W}_r , enabling LKGR to automatically measure different contributions of knowledge-based neighbors. Based on these learned weights, neighboring information can be selectively propagated and aggregated.

3.4.2 Lorentzian Information Propagation. To propagate the neighbor information on the Lorentzian manifold, we compute the Lorentzian linear combination of neighborhoods for users and items, respectively. As shown in Figure 2(a), since users only interact with items, the embedding of user u 's ego-network (i.e., $\mathcal{N}(u)$) representing u 's historical interactive information, is defined in Equation 6. Here $\hat{\pi}(u, r^*, i)$ is the normalized weight of edge (u, r^*, i) .

$$\mathbf{s}_{\mathcal{N}(u)} = \prod_{\mathbf{u}}^{exp,c} \left(\sum_{i \in \mathcal{N}(u)} \hat{\pi}(u, r^*, i) \prod_{\mathbf{u}}^{log,c}(\mathbf{i}) \right). \quad (6)$$

Similarly, items connect to users and KG entities; therefore, item i collectively receives the interactive signals from the user neighbors (i.e., $\mathcal{N}^{UI}(i)$), and knowledge from the KG side (i.e., $\mathcal{N}^{KG}(i)$). Let $u \in \mathcal{N}^{UI}(i)$ and $e \in \mathcal{N}^{KG}(i)$ denote a user and an entity connecting with item i by relation r^* and r . LKGR collectively summarizes the interactive information and knowledge backgrounds for item i as:

$$\mathbf{s}_{\mathcal{N}(i)} = \prod_{\mathbf{i}}^{exp,c} \left(\sum_{u \in \mathcal{N}^{UI}(i)} \hat{\pi}(i, r^*, u) \prod_{\mathbf{i}}^{log,c}(\mathbf{u}) + \sum_{e \in \mathcal{N}^{KG}(i)} \hat{\pi}(i, r, e) \prod_{\mathbf{i}}^{log,c}(\mathbf{e}) \right). \quad (7)$$

We implement a fixed-size random neighbor sampling instead of using full node neighbors, which is particularly useful for web-scale recommender systems [10, 42].

3.4.3 Lorentzian Neighbor Aggregation. After obtaining the propagated neighbor information (i.e., $\mathbf{s}_{\mathcal{N}(u)}$ and $\mathbf{s}_{\mathcal{N}(i)}$), the next step is to perform Lorentzian neighbor aggregation and update the embeddings for users and items. Generally, for node embedding \mathbf{x} and its neighbor representation $\mathbf{s}_{\mathcal{N}(x)}$, we use function $f(\mathbf{x}, \mathbf{s}_{\mathcal{N}(x)})$: $\mathbb{H}^d \times \mathbb{H}^d \rightarrow \mathbb{H}^d$ to update the representation for node x , i.e., $\mathbf{x} = f(\mathbf{x}, \mathbf{s}_{\mathcal{N}(x)})$. In this paper, we utilize three types of Lorentzian aggregators to implement $f(\cdot)$ as:

- **Sum Aggregator** [15] takes the summation of two inputs and conducts a nonlinear transformation, followed by a nonlinear activation on the Lorentzian manifold:

$$f_{sum} = \sigma^{\otimes c} \left(\mathbf{A} \odot^c (\mathbf{x} \oplus^c \mathbf{s}_{\mathcal{N}(x)}) \oplus^c \mathbf{b} \right), \quad (8)$$

where \mathbf{A} and \mathbf{b} are the trainable weights and bias defined in the associated tangent space. All Lorentzian basic operations e.g., \odot^c , will be introduced later.

- **Concatenate Aggregator** [10] concatenates two vectors, followed by a nonlinear transformation and activation as:

$$f_{concat} = \sigma^{\otimes c} \left(\mathbf{A} \odot^c (\mathbf{x} \otimes^c \mathbf{s}_{\mathcal{N}(x)}) \oplus^c \mathbf{b} \right), \quad (9)$$

where \otimes^c denotes the operation:

$$\mathbf{a} \otimes^c \mathbf{b} = \prod_{\mathbf{o}}^{exp,c} \left(\prod_{\mathbf{o}}^{log,c}(\mathbf{a}) \parallel \prod_{\mathbf{o}}^{log,c}(\mathbf{b}) \right). \quad (10)$$

- **Neighbor Aggregator** [33] directly updates the output representation with the input $\mathbf{s}_{\mathcal{N}(x)}$:

$$f_{neighbor} = \sigma^{\otimes c} \left(\mathbf{A} \odot^c \mathbf{s}_{\mathcal{N}(x)} \oplus^c \mathbf{b} \right). \quad (11)$$

Lorentzian linear transformation and activation. Given \mathbf{A} and \mathbf{b} , Lorentzian linear transformation of the hyperbolic geometry can be well extended from the Euclidean geometry as [1, 7]:

$$\begin{aligned} \mathbf{A} \odot^c \mathbf{x} &= \prod_{\mathbf{o}}^{exp,c} \left(\mathbf{A} \prod_{\mathbf{o}}^{log,c}(\mathbf{x}) \right), \\ \mathbf{x} \oplus^c \mathbf{b} &= \prod_{\mathbf{x}}^{exp,c} \left(\mathbf{b} - \gamma \left(\prod_{\mathbf{o}}^{log,c}(\mathbf{x}) + \prod_{\mathbf{x}}^{log,c}(\mathbf{o}) \right) \right), \end{aligned} \quad (12)$$

where $\gamma = \frac{\langle \prod_{\mathbf{o}}^{log,c}(\mathbf{x}), \mathbf{b} \rangle_{\mathbb{H}}}{c \operatorname{arccosh} \left(-\frac{(\mathbf{o}, \mathbf{x})_{\mathbb{H}}}{c} \right)^2}$. The hyperbolic activation σ is defined:

$$\sigma^{\otimes c} = \prod_{\mathbf{o}}^{exp,c} \left(\sigma \left(\prod_{\mathbf{o}}^{log,c}(\mathbf{x}) \right) \right). \quad (13)$$

Algorithm 1: LKGR algorithm

Input: UKG \mathcal{G} ; trainable parameters $\Theta: c, \{\mathbf{u}\}_{u \in \mathcal{U}}, \{\mathbf{i}\}_{i \in \mathcal{I}}, \{\mathbf{e}\}_{e \in \mathcal{E}}, \{\mathbf{W}_r\}_{r \in \mathcal{R}}, \{\mathbf{A}_j, \mathbf{b}_j\}_{j=0}$; hyper-parameters: $d, L, \eta, \lambda, f(\cdot)$.
Output: Prediction function $\mathcal{F}(u, i | \Theta, \mathcal{G})$

```

1 while LKGR not converge do
2   for  $(u, i) \in \mathcal{G}$  that  $y_{u,i} = 1$  do
3      $\mathcal{N}(u), \mathcal{N}^{UI}(i) \leftarrow$  get sampled user-item neighbors for  $u, i$ ;
4      $\mathbf{s}_{\mathcal{N}(u)}, \mathbf{s}_{\mathcal{N}^{UI}(i)} \leftarrow$  propagate interactive information;
5      $\mathbf{u} \leftarrow f(\mathbf{u}, \mathbf{s}_{\mathcal{N}(u)})$ ;
6      $\mathcal{N}^{KG}(i) \leftarrow$  get sampled  $L$ -hops of KG neighbors for  $i$ ;
7     for  $l = L, \dots, 1$  do
8       for  $e \in (l-1)$ -hop neighbor of  $i$  in  $\mathcal{N}^{KG}(i)$  do
9          $\mathcal{N}(e) \leftarrow$   $e$ 's entity neighbors in  $\mathcal{N}^{KG}(i)^{(l)}$ ;
10         $\mathbf{s}_{\mathcal{N}(e)} \leftarrow$  propagate KG information;
11        if  $l = 1$  then
12           $\mathbf{s}_{\mathcal{N}(e)} \leftarrow$  propagate interactive information
13          and KG backgrounds to  $e$ ;
14           $\mathbf{e} \leftarrow f(\mathbf{e}, \mathbf{s}_{\mathcal{N}(e)})$ ;
15         $\hat{y}_{u,i} \leftarrow$  compute estimated matching score;
16         $\mathcal{L} \leftarrow$  compute loss and optimize LKGR model;
17 return  $\mathcal{F}$ .
```

3.4.4 High-order Knowledge Extraction. To further explore the high-order information in KGs and propagate distant knowledge to items, as shown in Figure 2(b), we can explore the multi-hop subgraphs and stack more propagation layers in LKGR accordingly. Concretely, we first conduct l -hop neighbor sampling for item i to reach its high-order subgraph $\mathcal{N}^{KG}(i)$ in KG, where we use $\mathcal{N}^{KG}(i)^{(l)}$ to represent i 's l -hop neighbors. Then we propagate knowledge from the l -hop subgraph and iteratively aggregate to the centric node i . For example, entity e is the l -hop neighbor of item i in KG, i.e., $e \in \mathcal{N}^{KG}(i)^{(l)}$. Then in the l -hop of propagation, we formulate e ' neighbor representation by exploring e 's adjacent subgraph $\mathcal{N}^{KG}(i)^{(l+1)}$ as follows:

$$\mathbf{s}_{\mathcal{N}(e), e \in \mathcal{N}^{KG}(i)^{(l)}} = \prod_{\mathbf{e}}^{exp,c} \left(\sum_{e' \in \mathcal{N}^{KG}(i)^{(l+1)}} \hat{\pi}(e, r', e') \prod_{\mathbf{e}}^{log,c}(\mathbf{e}') \right), \quad (14)$$

where \mathbf{e}' , the embedding of entity $e' \in \mathcal{N}^{KG}(i)^{(l+1)}$, and coefficient $\hat{\pi}(e, r', e')$ are updated based on the previous step computation.

Specifically, high-order propagation relies on the neighbor sampling to generate a multi-hop sub-graph where edges live in the consecutive hops. As illustrated by the pseudocodes in Algorithm 1, the l -hop KG information can be iteratively propagated from $l = L$ to $l = 1$ via message passing along these edges (lines 7-13). At the 1-hop subgraph, the condensed KG information and user neighbor information collectively enriches i 's representation (lines 11-12). Please notice that 0-hop neighbor of item i in $\mathcal{N}^{KG}(i)$ is i itself (line 8), so that if $l = 1, e = i$ and $\mathbf{s}_{\mathcal{N}(e)} = \mathbf{s}_{\mathcal{N}(i)}$ (lines 11-13).

Time complexity analysis. Let Y denote the number of user-item interactions. α is the average time cost of numerical computation between Euclidean and hyperbolic spaces. The training time cost per epoch (iteration) is $O(\alpha \cdot Y \cdot (|\mathcal{N}(u)| + |\mathcal{N}^{UI}(i)| + |\mathcal{N}^{KG}(i)|^L))$. In this paper, for all benchmarks, the sampling size is no more than 8. Although the theoretical time complexity is exponential to L , in our work, $L \leq 2$. This is because stacking too many propagation hops incurs performance detriment, the main cause of which lies in the well-known *over-smoothing* [19, 21] problem. As we will show in Section 4.5, compared to recent state-of-the-art models stacking limited hops ($L \leq 2$), LKGR is comparably efficient in practice.

3.5 Prediction Layer and Model Optimization

3.5.1 Prediction Layer. In traditional embedding-based matching models, inner product and L2 distance are widely adopted, mainly because this simple but effective interaction modeling enables accelerated computation on the online matching stage. Therefore, we use the learned hyperbolic representations of users and items from the previous layer and take the inner product in the hyperbolic space to estimate their matching score:

$$\hat{y}_{u,i} = g(\mathbf{u}, \mathbf{i}) = (\prod_o^{log,c}(\mathbf{u}))^T \prod_o^{log,c}(\mathbf{i}). \quad (15)$$

During the evaluation stage, items with top scores $\hat{y}_{u,i}$ are selected as recommended items to a given user u .

3.5.2 Model Optimization. Let \mathcal{S}_u^+ denote the positive interacted item set of user u , i.e., $\hat{y}_{u,i} = 1$, and \mathcal{S}_u^- represent the corresponding sampling set of negative items, i.e., $\hat{y}_{u,i} = 0$. To effectively optimize LKGR for training, in this paper, we set $|\mathcal{S}_u^+| = |\mathcal{S}_u^-|$. In each iteration of model training, \mathcal{S}_u^+ and \mathcal{S}_u^- are updated accordingly. Finally, the loss function of LKGR is defined as follows:

$$\mathcal{L} = \sum_{u \in \mathcal{U}} \left(\sum_{i \in \mathcal{S}_u^+} \mathcal{J}(y_{u,i}, \hat{y}_{u,i}) - \sum_{i \in \mathcal{S}_u^-} \mathcal{J}(y_{u,i}, \hat{y}_{u,i}) \right) + \lambda \|\Theta\|_2^2. \quad (16)$$

where $\mathcal{J}(\cdot)$ denotes the cross-entropy loss, Θ is the set of trainable model parameters and embeddings, and $\|\Theta\|_2^2$ is the L2-regularizer parameterized by λ .

4 EXPERIMENTS

We evaluate LKGR model under the three real-world scenarios, with the aim of answering the following research questions:

- **RQ1.** How does LKGR perform compared to state-of-the-art KG-enhanced recommendation models?
- **RQ2.** How is the time efficiency of LKGR in model training?
- **RQ3.** How do proposed model components and different hyperparameter settings affect LKGR?

4.1 Datasets

To evaluate the effectiveness of LKGR, we utilize the following three benchmarks for movie, book, and restaurant recommendations. In terms of the diversity in domain, size, and sparsity, all these three datasets are frequently evaluated in recent works [35–37, 39].

Specifically, Book-Crossing³ (**Book**) is a dataset of book ratings (ranging from 0 to 10) extracted from the Book-Crossing community. MovieLens-20M (**Movie**) is a widely adopted benchmark for movie

Table 1: Three datasets used in this paper.

	Book	Movie	Restaurant
# users	17,860	138,159	2,298,698
# items	14,967	16,954	1,362
# interactions	139,746	13,501,622	23,416,418
# entities	77,903	102,569	28,115
# relations	25	32	7
# KG triples	151,500	499,474	160,519

recommendation. It contains approximately 20 million explicit ratings (ranging from 1 to 5) on the MovieLens website. Dianping-Food (**Restaurant**) is a commercial dataset from Dianping.com that consists of over 10 million diverse interactions, e.g., clicking, saving, and purchasing, between about 2 million users and 1 thousand restaurants. The first two datasets are publicly accessible and the last one is contributed by Meituan-dianping Inc. [36]. Dataset statistics are summarized in Table 1.

4.2 Baselines

To demonstrate the effectiveness of the proposed model, we compare LKGR with three types of state-of-the-art recommendation methods: CF-based methods (BPRMF, NFM), regularization-based methods (CKE and KGAT), and propagation-based methods (RippleNet, KGCN, KGNN-LS, CKAN), as follows.

- **BPRMF** [29] is a representative CF-based method that performs matrix factorization with implicit feedback, optimized by the Bayesian Personalized ranking optimization criterion.
- **NFM** [11] is a state-of-the-art neural factorization machine model for recommendation.
- **CKE** [43] is a classical regularization-based method. CKE learns semantic embeddings using TransR [22] with structural, textual and visual information to subsume matrix factorization under a unified Bayesian framework.
- **KGAT** [38] is another representative regularization-based model that collectively refines user and item embeddings via an attentive embedding propagation layer. We use the pre-trained embeddings of users and items from BPRMF to initialize the model.
- **RippleNet** [35] is a recent state-of-the-art propagation-based model. Aiming at enriching the users’ representations, RippleNet uses a memory-like network to propagate users’ preferences towards items by following paths in KGs.
- **KGCN** [37] is another state-of-the-art propagation-based method which extends spatial GCN approaches to the KG domain. By aggregating high-order neighbor information, both structure information and semantic information of the KG can be learned to capture users’ potential long-distance interests.
- **KGNN-LS** [36] is a state-of-the-art propagation-based method that applies graph neural network architecture to KGs with label smoothness regularization for recommendation.
- **CKAN** [39] is the latest state-of-the-art propagation-based method which employs a heterogeneous propagation strategy to encode diverse information in KGs for recommendation.

³Book-Crossing: <http://www2.informatik.uni-freiburg.de/~chiegler/BX/>

Table 2: Average results of Top@20 recommendation task. Underline indicates the second-best model performance. Bold denotes the empirical improvements against second-best models, and * denotes scenarios where a Wilcoxon signed-rank test indicates a statistically significant improvement under 95% confidence level between our model and second-best models.

Model	Book-Crossing		MovieLens-20M		Dianping-Food	
	Recall@20(%)	NDCG@20(%)	Recall@20(%)	NDCG@20(%)	Recall@20(%)	NDCG@20(%)
BPRMF	4.67 (8.7e-3)	2.80 (4.3e-3)	20.48 (1.6e-2)	15.77 (9.1e-3)	19.90 (3.0e-2)	10.79 (2.0e-2)
NFM	3.93 (2.2e-2)	2.17 (1.5e-2)	19.79 (3.3e-2)	14.28 (1.1e-2)	23.85 (3.9e-2)	<u>12.48</u> (3.0e-2)
CKE	4.38 (9.6e-3)	2.24 (4.2e-2)	21.52 (1.2e-2)	15.73 (1.3e-2)	22.24 (3.1e-2)	12.09 (1.5e-2)
RippleNet	7.12 (2.1e-2)	5.09 (1.7e-2)	13.74 (2.6e-2)	9.77 (1.7e-2)	21.20 (4.1e-2)	10.99 (2.0e-2)
KGNN-LS	<u>8.51</u> (2.9e-2)	<u>6.06</u> (1.7e-2)	20.20 (1.0e-2)	15.49 (1.3e-2)	15.52 (4.9e-2)	7.92 (2.9e-2)
KGCN	7.85 (2.9e-2)	5.93 (2.3e-2)	19.24 (3.2e-2)	13.87 (1.6e-2)	19.03 (3.0e-2)	9.34 (1.5e-2)
KGAT	5.34 (6.1e-3)	3.01 (7.9e-3)	<u>21.80</u> (7.7e-3)	<u>16.81</u> (1.1e-2)	15.57 (2.4e-2)	7.67 (1.7e-2)
CKAN	6.19 (1.1e-2)	3.47 (5.3e-3)	17.48 (1.7e-2)	12.48 (1.4e-2)	<u>24.10</u> (3.9e-2)	13.33 (2.0e-2)
LKGR	9.48* (2.3e-2)	6.50* (1.6e-2)	25.14* (4.1e-2)	18.34* (4.7e-2)	24.97* (4.4e-2)	10.45 (1.9e-2)
% Improv.	11.40%	7.26%	15.32%	9.10%	3.61%	N/A




Figure 3: Average results of Recall@K and NDCG@K in Top-K Recommendation.

4.3 Experiment Setup

To evaluate LKGR, we randomly divide each dataset 5 times into training, evaluation, and test sets with the ratio of 6:2:2. In the evaluation of Top-K recommendation task, we use the model learned from the training set to rank K items for each user in the test set with the highest predicted scores $\hat{y}_{i,j}$. We choose two widely-used evaluation protocols, i.e., Recall@K and NDCG@K. We optimize all models with Adam optimizer [14] and adopt the default Xavier initializer [9] to initialize the model parameters.

We implement the LKGR model under Python 3.7 and Pytorch 1.4.0. The experiments are run on a Linux machine with 1 V100 GPU, 4 Intel Xeon CPU (Skylake, IBRS), 16 GB of RAM. For all the baselines, we follow the official hyper-parameter settings from original papers or as default in corresponding codes. For methods lacking recommended settings, we apply a grid search for hyper-parameters. The learning rate, denoted by η , is tuned within $\{10^{-3}, 5 \times 10^{-2}, 10^{-2}, 5 \times 10^{-1}\}$ and the coefficient of L2 normalization is tuned among $\{10^{-6}, 10^{-5}, 10^{-4}, 10^{-3}\}$. The embedding size is searched in $\{8, 16, 32, 64, 128\}$.

4.4 Performance Comparison (RQ1)

4.4.1 Top-K recommendation. We evaluate LKGR on Top-K recommendation task by varying K in $\{1, 5, 10, 20, 50, 100\}$. Table 2 reports Top-20 recommendation results for detailed comparison and analysis and Figure 3 contains the complete performance curves of LKGR and baselines. From these results, we can observe that:

- *The results of Top-20 recommendation demonstrate the improvements of LKGR are statistically stable and significant.* We report the details of Top-20 recommendation of all models in Table 2. For example, we can observe that LKGR significantly outperforms all baselines by 11.40% (9.48→8.51), 15.32% (25.14→21.80) and 3.61% (24.97→24.10) of Recall metric for Top-20 recommendation on three datasets. As for the NDCG metric, LKGR achieves 7.26% (6.50→6.06) and 9.10% (18.34→16.81) improvement of NDCG@20 on Book and Movie datasets but does not perform the best on Restaurant. This shows that LKGR can make good recalling of Top-20 candidates from the item corpus but may not well capture the ranking within for this specific dataset, leaving the space of further optimization. The reported standard deviations of LKGR


are also comparable with all baselines, showing the stability of our proposed model. Furthermore, we conduct *Wilcoxon signed-rank tests* to verify that the achieved performance improvements are statistically significant over the second-best recommendation model under the 95% confidence level.

- As K increases, *LKGR* consistently performs well on three benchmarks. Results in Figure 3 demonstrate the effectiveness of *LKGR* on improving the performance of the ranking recommendation task, i.e., Top- K recommendation. This shows that the hyperbolic geometry does well capture the hierarchical properties of graphs, meaning that it can preserve the order of users’ preferences towards the items, which benefits the ranking recommendation task. Furthermore, our discrete information propagation and aggregation strategies for users and items on the Lorentzian manifold are effective. We will conduct a more comprehensive ablation study in the later section, analyzing the contribution of all *LKGR* module components to the recommendation performance.

4.5 Time Efficiency Comparison (RQ2)

Model	BK	MV	RT
BPRMF	4.58	109.21	134.97
NFM	26.50	92.52	313.87
CKE	10.15	82.83	98.61
RippleNet	12.46	1,159.32	2,343.93
KGNN-LS	1.62	36.51	81.17
KGCN	4.47	15.27	58.80
KGAT	61.88	333.86	2,254.83
CKAN	3.36	439.23	529.21
LKGR	43.26	250.88	337.10

(a) Training time cost per epoch (s).



(b) Efficiency v.s. accuracy

Figure 4: Model comparison on efficiency and accuracy.

In this section, we study how time efficient our *LKGR* and baselines are. All methods are run on the same aforementioned running environment, and we use the default hyper-parameters that are reported in papers or official codes. We use BK, MV, and RT to denote Book, Movie, Restaurant, and report the results in Figure 4.

We can observe that: (1) as shown in Figure 4(a), compared to CF-based methods, i.e., BPRMF and NFM, *LKGR* requires more time for model training as it needs to explore the KG for information propagation; (2) compared to the remaining KG-based recommendation models, *LKGR* requires the analogous time cost of model training, which may dispel concerns of large computation overhead in the hyperbolic space. (3) Figure 4(b) visualizes the overall performances in terms of efficiency and accuracy among all models. As the upper-right corner of the figure means the ideal optimal performance, *LKGR* makes an excellent trade-off w.r.t efficiency and accuracy, especially on Movie and Restaurant datasets.

4.6 Ablation Study of LKGR (RQ3.A)

In this section, we first conduct a comprehensive ablation study to evaluate the effectiveness of all model components. We use $R@20$ and $N@20$ to denote $\text{recall}@20$ and $\text{NGCG}@20$.

Table 3: Ablation study on Top- K recommendation (%).

Dataset	w/o IS	w/o KE	w/o HG	w/o LKA	LKGR
BK- $R@20$	4.75 (-49.89%)	5.79 (-38.92%)	8.11 (-14.45%)	8.35 (-11.92%)	9.48
BK- $N@20$	3.81 (-41.38%)	4.22 (-35.08%)	6.21 (-4.46%)	6.43 (-1.08%)	6.50
MV- $R@20$	12.45 (-50.48%)	11.73 (-53.34%)	22.19 (-11.73%)	23.94 (-4.77%)	25.14
MV- $N@20$	8.76 (-52.23%)	7.92 (-56.82%)	15.47 (-15.65%)	17.03 (-7.14%)	18.34
RT- $R@20$	11.22 (-55.07%)	9.86 (-60.51%)	21.47 (-14.02%)	23.17 (-7.21%)	24.97
RT- $N@20$	6.54 (-37.42%)	4.74 (-54.64%)	10.41 (-0.38%)	11.03 (5.55%)	10.45

Impact of interactive signal propagation. To investigate the impact of information propagation and aggregation of interactive signal in *LKGR*, we set a variant model, denoted by $LKGR_{w/o IS}$, which disables the information passing between users and items in our proposed Equations (6) and (7), so that items receive information only from KG entities. Based on the results reported in Table 3, without information propagation between users and items, the performance of Top- K recommendation drops dramatically, which demonstrates that learning historical interactive information for users and items is essential for the performance improvement.

Impact of knowledge extraction. Similarly, we set another variant $LKGR_{w/o KE}$ to study the influence of knowledge extraction, by cutting the KG information propagation to items (formulated in Equations (7) and (14)). As shown in Table 3, purely relying on the user-item interaction modeling is not enough to boost performance; thus integrating KGs in recommendation is also very important.

Impact of hyperbolic geometry. We verify the impact of the hyperbolic geometry by replacing all hyperbolic operations into Euclidean space and retaining the computation logic of *LKGR*. We denote the variant as $LKGR_{w/o HG}$. As shown in Table 3, even with adequate information propagation from user-item interactions and KGs, directly modeling these two data sources in Euclidean space brings about large performance decay on Top- K recommendation task. For example, by removing the hyperbolic geometry, $LKGR_{w/o HG}$ produces worse listwise ranking results on all three benchmarks: 14.45% (8.11→9.48), 11.73% (22.19→25.14), and 14.02% (21.47→24.97) of Recall for Top-20 recommendation, respectively. To conclude, with all other essential components, embedding graph nodes on the Lorentzian manifold is substantial to *LKGR*, and modeling the scale-free graphs in the hyperbolic space can learn better representation for recommendation.

Impact of Lorentzian Knowledge-aware Attention mechanism. We also examine the impact of our proposed attention mechanism. For comparison, we use a variant that set $\pi(h, r, t) = 1$ to make equal contribution for computing $s_{N(\cdot)}$, termed by $LKGR_{w/o LKA}$. As shown in Table 3, disabling our attention mechanism leads to the decreased $\text{Recall}@20$ by 11.92% (8.35→9.48), 4.77% (23.94→25.14), and 7.21% (23.17→24.97). This proves that Lorentzian Knowledge-aware Attention mechanism is also effective, as it adaptively measures the importance weights for different nodes on the Lorentzian manifold and enables more coherent embedding aggregation.

4.7 Effect of Key Hyper-parameters (RQ3.B)

Due to the space limit, we only report the effect of some key hyper-parameters on the model performance.

Effect of Lorentzian aggregators. To explore the influence of aggregating neighbor information on the Lorentzian manifold, we conduct experiments over different selections of Lorentzian aggregators $f(\cdot)$. Based on the Recall and NDCG metrics, we observe that in Table 4 *sum aggregator*, i.e., $f_{sum}(\cdot)$, is superior on datasets of Book and Movie. For the Restaurant dataset, $f_{concat}(\cdot)$ is much more suitable. This is because sum and concat aggregators are capable of retaining external neighbor information as well as the internal information when comparing to $f_{neighbor}(\cdot)$.

Table 4: Top-K recommendation (%) w.r.t different f .

f	f_{sum}		f_{concat}		$f_{neighbor}$	
	R@20	N@20	R@20	N@20	R@20	N@20
BK	9.48	6.50	8.19	6.06	8.10	5.84
MV	25.14	18.34	21.42	16.33	20.87	15.23
RT	22.55	9.83	24.97	10.45	21.48	9.29

Effect of KG propagation depth L . We verify how the propagation depth affects the performance by varying L from 0 to 3. Depth 0 means no local information aggregation from KG. As shown in Table 5, we notice that LKGR achieves the best performance when L is 1, 2, and 1 for Book, Movie, and Restaurant. In addition to the over-smoothing issue that we explain in Section 3.4.4, another possible reason is: when long-distance propagation introduces distant knowledge, it may also bring about irrelevant information, especially when the dataset is large and dense. Thus, preserving an appropriate depth in the high-order information propagation enables maximized performance over different recommendation benchmarks adaptively.

Table 5: Top-K recommendation (%) w.r.t different L .

L	$L = 0$		$L = 1$		$L = 2$		$L = 3$	
	R@20	N@20	R@20	N@20	R@20	N@20	R@20	N@20
BK	5.79	4.22	9.48	6.50	8.16	5.73	7.43	5.28
MV	22.19	15.47	20.34	15.06	25.14	18.34	21.48	15.17
RT	9.86	4.74	24.97	10.45	22.12	9.47	21.67	9.02

5 RELATED WORKS

With the rapid development of information networks, studying the ubiquitous graph data has aroused various interests in both industry and research communities [3–5, 40, 45]. Among these research topics, KG-enhanced recommender systems receive much attention recently as they can alleviate data sparsity and cold-start problems for better recommendation. These methods can be generally categorized into three branches: (1) path-based, (2) regularization-based, and (3) propagation-based.

(1) In **path-based methods**, high level meta-paths are extracted from KGs and then input into the predictive model. Such meta-paths are usually manually selected or generated, which require intensive input of domain knowledge and labor resources [6, 13]. Furthermore, it is difficult to optimize the path retrieval for large and complex graphs, while the selected paths have a great impact

on the final performance. Thus we exclude path-based methods for model comparison. (2) In **regularization-based methods**, additional loss terms are devised to capture the KG structure and regularize the model training. For example, KGAT [38] merges the two tasks of recommendation and KG completion to jointly train the model. HyperKnow [25] further embeds the joint training in Poincaré Ball with an adaptive regularization mechanism to balance two loss terms. However, one deficiency is that these methods usually rely on the traditional knowledge graph embedding methods to complete the KG training, high-order semantic information in the KG and user-item interactions is not explicitly modeled, which may result in suboptimal representation learning for users and items. (3) **Propagation-based methods**, performing iterative message passing with the guidance of the topological structure in the KGs [35–37, 39], has attracted much attention recently. With the auxiliary information passed along links in the KGs, the representations of users and items can be refined to provide more accurate recommendation service. However, all these methods assume the learning process in the low-dimensional Euclidean space; nevertheless, because of the intrinsic hierarchical structure of KGs, whether the Euclidean space is appropriate for all kinds of scenarios is still an open question. Conceptually, LKGR is inspired by the graph-based message passing mechanism and performs graph convolutional operations in the hyperbolic space.

6 CONCLUSION AND FUTURE WORKS

In this paper, we propose a KG-enhanced recommendation model, namely LKGR, which learns the embeddings of users and items as well as the KG entities in the hyperbolic space. We propose a knowledge-aware attention mechanism on the Lorentzian manifold to discriminate the contribution of graph node informativeness, which is followed by multi-layer aggregation for high-order information propagation. The experimental results over three real-world datasets not only validate the performance improvements of LKGR over recent state-of-the-art solutions, but also demonstrate the effectiveness of all proposed model components.

As for the future works, there are two potential directions: (1) since all existing works unify the user-item interactions and the KG into a static graph, while in practice, users and items usually contain temporal and dynamic interactions. How to simultaneously learn the graph structure and temporal information in a unified framework is a good direction to work on. (2) It is worth exploring other application scenarios with hyperbolic geometry for performance improvement, e.g., information retrieval [46] and language processing [8, 20], such that hyperbolic modeling can well learn the embedding of interrelated data with less distortion.

ACKNOWLEDGMENTS

The work described in this paper was partially supported by the Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 2410021, Research Impact Fund, No. R5034-18) and Research Grants Council of the Hong Kong Special Administrative Region, China (CUHK 3133238, Research Sustainability of Major RGC Funding Schemes (RSFS)).

REFERENCES

- [1] Ines Chami, Zhitao Ying, Christopher Ré, and Jure Leskovec. 2019. Hyperbolic graph convolutional neural networks. In *Neural Information Processing Systems (NeurIPS)*. 4868–4879.
- [2] Wei Chen, Wenjie Fang, Guangda Hu, and Michael W Mahoney. 2013. On the hyperbolicity of small-world and treelike random graphs. *Internet Mathematics* 9, 4 (2013), 434–491.
- [3] Yankai Chen, Jie Zhang, Yixiang Fang, Xin Cao, and Irwin King. 2020. Efficient community search over large directed graphs: An augmented index-based approach. In *International Joint Conference on Artificial Intelligence (IJCAI)*. 3544–3550.
- [4] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *Neural Information Processing Systems (NeurIPS)*.
- [5] Yixiang Fang, Reynold Cheng, Yankai Chen, Siqiang Luo, and Jiafeng Hu. 2017. Effective and efficient attributed community search. *The VLDB Journal* 26, 6 (2017), 803–828.
- [6] Xinyu Fu, Jiani Zhang, Ziqiao Meng, and Irwin King. 2020. MAGNN: Metapath Aggregated Graph Neural Network for Heterogeneous Graph Embedding. In *World Wide Web (WWW)*. 2331–2341.
- [7] Octavian Ganea, Gary Bécigneul, and Thomas Hofmann. 2018. Hyperbolic neural networks. In *Neural Information Processing Systems (NeurIPS)*. 5345–5355.
- [8] Yifan Gao, Chien-Sheng Wu, Jingjing Li, Shafiq Joty, Steven CH Hoi, Caiming Xiong, Irwin King, and Michael R Lyu. 2020. Discern: Discourse-aware entailment reasoning network for conversational machine reading. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- [9] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*. 249–256.
- [10] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Neural Information Processing Systems (NeurIPS)*. 1024–1034.
- [11] Xiangnan He and Tat-Seng Chua. 2017. Neural factorization machines for sparse predictive analytics. In *International Conference on Research and Development in Information Retrieval (SIGIR)*. 355–364.
- [12] Jonathan L Herlocker, Joseph A Konstan, Loren G Terveen, and John T Riedl. 2004. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems (TOIS)* 22, 1 (2004), 5–53.
- [13] Binbin Hu, Chuan Shi, Wayne Xin Zhao, and Philip S Yu. 2018. Leveraging metapath based context for top-n recommendation with a neural co-attention model. In *International Conference on Knowledge Discovery & Data Mining (SIGKDD)*. 1531–1540.
- [14] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *The International Conference on Learning Representations (ICLR)* (2014).
- [15] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. *The International Conference on Learning Representations (ICLR)* (2017).
- [16] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009), 30–37.
- [17] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguńá. 2010. Hyperbolic geometry of complex networks. *Physical Review E* 82, 3 (2010), 036106.
- [18] Marc Law, Renjie Liao, Jake Snell, and Richard Zemel. 2019. Lorentzian distance learning for hyperbolic representations. In *The International Conference on Machine Learning (ICML)*. PMLR, 3672–3681.
- [19] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. 2019. Deepgcns: Can gcns go as deep as cnns?. In *International Conference on Computer Vision (ICCV)*. 9267–9276.
- [20] Jingjing Li, Zichao Li, Lili Mou, Xin Jiang, Michael R Lyu, and Irwin King. 2020. Unsupervised text generation by learning from search. *Neural Information Processing Systems (NeurIPS)* (2020).
- [21] Qimai Li, Zhichao Han, and Xiao-Ming Wu. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *AAAI Conference on Artificial Intelligence (AAAI)*, Vol. 32.
- [22] Hailun Lin, Yong Liu, Weiping Wang, Yinliang Yue, and Zheng Lin. 2017. Learning entity and relation embeddings for knowledge resolution. *Procedia Computer Science* 108 (2017), 345–354.
- [23] Nathan Linial, Eran London, and Yuri Rabinovich. 1995. The geometry of graphs and some of its algorithmic applications. *Combinatorica* 15, 2 (1995), 215–245.
- [24] Qi Liu, Maximilian Nickel, and Douwe Kiela. 2019. Hyperbolic graph neural networks. *Neural Information Processing Systems (NeurIPS)* (2019).
- [25] Chen Ma, Liheng Ma, Yingxue Zhang, Haolun Wu, Xue Liu, and Mark Coates. 2021. Knowledge-Enhanced Top-K Recommendation in Poincaré Ball. *Association for the Advancement of Artificial Intelligence (AAAI)* (2021).
- [26] Hao Ma, Irwin King, and Michael R Lyu. 2007. Effective missing data prediction for collaborative filtering. In *International Conference on Research and Development in Information Retrieval (SIGIR)*. 39–46.
- [27] Maximilian Nickel and Douwe Kiela. 2018. Learning continuous hierarchies in the lorentz model of hyperbolic geometry. *The International Conference on Machine Learning (ICML)* (2018).
- [28] Erzsébet Ravasz and Albert-László Barabási. 2003. Hierarchical organization in complex networks. *Physical review E* 67, 2 (2003), 026112.
- [29] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2012. BPR: Bayesian personalized ranking from implicit feedback. *The Conference on Uncertainty in Artificial Intelligence (UAI)* (2012).
- [30] Joel W Robbin and Dietmar A Salamon. 2011. Introduction to differential geometry. *ETH, Lecture Notes, preliminary version* (2011), 18.
- [31] Frederic Sala, Chris De Sa, Albert Gu, and Christopher Ré. 2018. Representation tradeoffs for hyperbolic embeddings. In *International Conference on Machine Learning (ICML)*. PMLR, 4460–4469.
- [32] Andrew I Schein, Alexandrin Popescul, Lyle H Ungar, and David M Pennock. 2002. Methods and metrics for cold-start recommendations. In *International Conference on Research and Development in Information Retrieval (SIGIR)*. 253–260.
- [33] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. 2018. Graph attention networks. *The International Conference on Learning Representations (ICLR)* (2018).
- [34] Maksims Volkovs, Guangwei Yu, and Tomi Poutanen. 2017. Dropoutnet: Addressing cold start in recommender systems. In *Neural Information Processing Systems (NeurIPS)*. 4957–4966.
- [35] Hongwei Wang, Fuzheng Zhang, Jialin Wang, Miao Zhao, Wenjie Li, Xing Xie, and Minyi Guo. 2018. Ripplenet: Propagating user preferences on the knowledge graph for recommender systems. In *The Conference on Information and Knowledge Management (CIKM)*. 417–426.
- [36] Hongwei Wang, Fuzheng Zhang, Mengdi Zhang, Jure Leskovec, Miao Zhao, Wenjie Li, and Zhongyuan Wang. 2019. Knowledge-aware graph neural networks with label smoothness regularization for recommender systems. In *International Conference on Knowledge Discovery & Data Mining (SIGKDD)*. 968–977.
- [37] Hongwei Wang, Miao Zhao, Xing Xie, Wenjie Li, and Minyi Guo. 2019. Knowledge graph convolutional networks for recommender systems. In *World Wide Web (WWW)*. 3307–3313.
- [38] Xiang Wang, Xiangnan He, Yixin Cao, Meng Liu, and Tat-Seng Chua. 2019. Kgat: Knowledge graph attention network for recommendation. In *International Conference on Knowledge Discovery & Data Mining (SIGKDD)*. 950–958.
- [39] Ze Wang, Guangyan Lin, Huobin Tan, Qinghong Chen, and Xiyang Liu. 2020. CKAN: Collaborative Knowledge-aware Attentive Network for Recommender Systems. In *International Conference on Research and Development in Information Retrieval (SIGIR)*. 219–228.
- [40] Menglin Yang, Ziqiao Meng, and Irwin King. 2020. FeatureNorm: L2 Feature Normalization for Dynamic Graph Embedding. In *2020 IEEE International Conference on Data Mining (ICDM)*. IEEE, 731–740.
- [41] Menglin Yang, Min Zhou, Marcus Kalander, Zengfeng Huang, and Irwin King. 2021. Discrete-time Temporal Network Embedding via Implicit Hierarchical Learning in Hyperbolic Space. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery & Data Mining (SIGKDD)*. 1975–1985.
- [42] Rex Ying, Ruining He, Kaifeng Chen, Pong Eksombatchai, William L Hamilton, and Jure Leskovec. 2018. Graph convolutional neural networks for web-scale recommender systems. In *International Conference on Knowledge Discovery & Data Mining (SIGKDD)*. 974–983.
- [43] Fuzheng Zhang, Nicholas Jing Yuan, Defu Lian, Xing Xie, and Wei-Ying Ma. 2016. Collaborative knowledge base embedding for recommender systems. In *International Conference on Knowledge Discovery & Data Mining (SIGKDD)*. 353–362.
- [44] Jiani Zhang, Xingjian Shi, Shenglin Zhao, and Irwin King. 2019. Star-gcn: Stacked and reconstructed graph convolutional networks for recommender systems. *International Joint Conference on Artificial Intelligence (IJCAI)* (2019).
- [45] Yingxue Zhang, Soumyasundar Pal, Mark Coates, and Deniz Üstebay. 2019. Bayesian graph convolutional neural networks for semi-supervised classification. In *Association for the Advancement of Artificial Intelligence (AAAI)*.
- [46] Yifei Zhang and Hao Zhu. 2019. Doc2hash: Learning discrete latent variables for documents retrieval. In *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT)*. 2235–2240.