

Elsevier required licence: © 2021

This manuscript version is made available under the
CC-BY-NC-ND 4.0 license

<http://creativecommons.org/licenses/by-nc-nd/4.0/>

The definitive publisher version is available online at

<https://doi.org/10.1016/j.ins.2021.07.084>

Leveraging Implicit Relations for Recommender Systems

Anchen Li^{a,b}, Bo Yang^{a,b,*}, Huan Huo^c and Farookh Khadeer Hussain^c

^aCollege of Computer Science and Technology, Jilin University, China

^bKey Laboratory of Symbolic Computation and Knowledge Engineering of Ministry of Education, China

^cUniversity of Technology Sydney, Australia

ARTICLE INFO

Keywords:
Recommender systems
Collaborative filtering
Implicit relations
Deep learning

ABSTRACT

Collaborative filtering (CF) is one of the dominant techniques used in recommender systems. Most CF-based methods treat every user (or item) as an isolated existence, without explicitly modeling potential mutual relations among users (or items), which are latent in user-item interactions. In this paper, we design a novel strategy to mine user-user and item-item implicit relations and propose a natural way of utilizing the implicit relations for recommendation. Specifically, our method contains two major phases: neighbor construction and a recommendation framework. The first phase constructs an implicit neighbor set for each user and item according to historical user-item interaction. In the second phase, based on the constructed neighbor sets, we propose a deep framework to generate recommendations. We conduct extensive experiments with four datasets on movie, business, book, and restaurant recommendations and compare our methods with seven baselines, e.g., feature-based, neighborhood-based, and graph-based models. The experiment results demonstrate that our method achieves superior performance in rating prediction and top- k recommendation.

1. Introduction

In the era of information explosion, recommender systems play an indispensable role in identifying user preferences by recommending products or services. Collaborative filtering (CF) is one of the state-of-art techniques in recommender systems [1, 2, 3, 4, 5, 6, 7]. In a typical CF scenario with user-item interaction history, matrix factorization (MF), which embeds users and items in a shared latent space and models the user preference to an item as the inner product between the corresponding user and item embeddings, has become one of the most popular approaches [8]. However, due to the complex interaction between users and items, the shallow representations in the MF-based methods lack the expressiveness to model features for users and items [9, 10].

Recent years have witnessed the great success of deep neural network techniques in many research areas such as computer vision and natural language processing. Some recently proposed recommendation approaches utilize deep neural networks to capture the complex relationships between user-item interactions, which enhance the performance of

the previous shallow models [9, 11]. Though successful, most deep recommendation models treat every user (or item) as an isolated existence and have tended not to focus on potential user-user or item-item relations. Such potential re-

lations are latent in user-item interactions and could provide valuable information to infer user or item features [10, 12]. Although some existing works [13, 14] utilize graph neural networks (GNNs) on the user-item bipartite graph to capture high-order relations among users (or items), a more explicit and straightforward way is to directly construct user-user and item-item relations. Empirical evidence is from the recent work MMCF [10], which explicitly utilizes the *co-occurrence* relation (i.e., users who have interacted with the same items or items with which the same users have interacted⁴²) to define the neighbors for users and items. For instance, Figure 1(a) shows a simple user-item interaction in the movie domain, where each user rates movies on a 5-point integer scale to express their preference for movies. Take the user co-occurrence relation for example: for user d , the co-occurrence relation defines

user a , user b , and user c as her neighbors (as shown in Figure 1(b)) because user d and these users have interacted with common items.

Although MMCF has shown promising results, we argue that such a co-occurrence relation is macro-level and coarse-grained. For instance, for user d , the co-occurrence relation defines user e as one of her neighbors, but user d and user e have different preferences for movie d . By observing the user-item interaction in Figure 1(a), although user d and user a are not in a co-occurrence relation, they both share common preferences with user b and user c . Such a high-order

51 transitive relation is also a very significant signal for revealing user preferences and item properties, while it is ignored
 52 by most existing works. Therefore, we believe that (i) the co-occurrence relation contains useful information,
 53 but not all co-occurrence relations help; (ii) using the co-occurrence relation directly without filtering may introduce
 54 some irrelevant information or even noise, which will mislead the learning of user and item feature representations;
 55 (iii) the co-occurrence relation ignores high-order transitive relations for users and items.

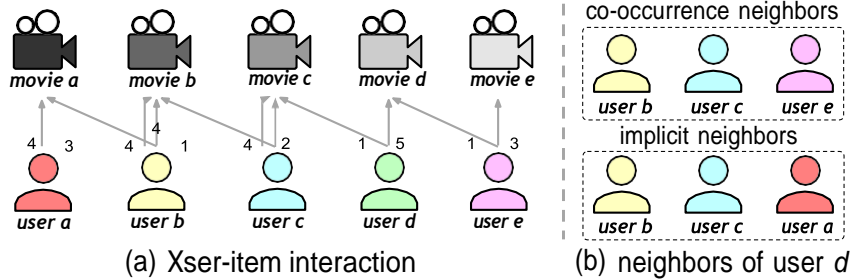


Figure 1: Illustration of the problem. (a): A simple user-item interaction scenario in a 5-star system. (b): Comparison between the co-occurrence neighbors and implicit neighbors.

56 To overcome the aforementioned limitation, we propose our method IRec which leverages user-user and item-item
 57 **I**mplicit relations for **R**ecommendation. IRec contains two major phases: neighbor construction and a recommendation
 58 framework. In the first phase, we construct an implicit neighbor set for each user and each item. More specifically,
 59 we first utilize the user-item interaction information to construct a user relational graph and an item relational graph.
 60 We then map each graph to a latent continuous space to find the implicit neighbors for users and items. In this way,
 61 the
 62 constructed neighborhoods not only filter out some irrelevant (or noisy) co-occurrence relations, they also may contain
 63 high-order transitive relations. In the second phase, we design a deep framework based on graph neural networks
 64 (GNNs) which utilizes the constructed user and item neighbor sets for recommendation. The key component of the frame-
 65 work is that we devise an aggregator on the neighbor sets to update the feature representations of users and items.
 66 Empirically, we apply IRec to four real-world scenarios of the movie, business, book, and restaurant recommenda-
 67 tions. The experiment results show that IRec outperforms the state-of-the-art approaches in both rating prediction and
 top- k recommendation. In summary, our main contributions in this paper are listed as follows:

- 68 • We provide a novel approach to find implicit neighbors for users and items.
- 69 • We propose an end-to-end framework that integrates implicit neighbors into recommendations.
- 70 • The experimental results on four real-world datasets show the effectiveness of IRec.

71 The remainder of this paper is organized as follows. Section 2 reviews the work related to our methods. In Section 3,
 72 we present the problem formulation and introduce the proposed method IRec. In Section 4, we describe the
 experiments conducted

73 on four real-world datasets and present the experiment results, followed by a conclusion and suggestions for future work
 in Section 5.

74 2. Related Work

75 In this section, we provide a brief overview of four areas that are highly relevant to our work.

76 2.1. Collaborative Filtering

77 Collaborative filtering (CF) can generally be grouped into three categories: neighborhood-based model, latent factor
 78 model, and hybrid model [15, 16]. Neighborhood-based methods identify neighborhoods of similar
 79 users or items based on the user-item interaction history [12, 16]. For example, ItemKNN utilizes collaborative item-
 80 item similarities (e.g. cosine similarity) to generate recommendations [17]. The latent factor model, such as matrix
 81 factorization [8], projects users and items into low-dimensional feature vector spaces. The interactions between users

82 and items are modelled as the inner product of their latent vectors. With the development of deep learning, some latent
83 factor models utilize deep neural networks as representation learning tools to capture complex user-item interactions
84 [9, 11]. As for the hybrid model, it merges the latent factor model and the neighborhood-based model. SVD++ is a

well-known and commonly used hybrid model that leverages users' explicit feedback and implicit feedback to predict user preferences [16]. Recently, a line of work leverages co-occurrence relations to define the neighbors for users and items and integrates deep components into the hybrid model [10, 18]. Since the co-occurrence relation is coarse-grained and lacks high-order semantics, these methods are insufficient to generate better recommendations. Different from the aforementioned work, our proposed method IRec is a unified hybrid model using implicit relations to define user and item neighbors, which accounts for both co-occurrence relations and high-order transitive relations.

2.2. Graph Representation

Graph representation learning is a significant method to learn latent, low-dimensional representations of vertices in the graph, while preserving both graph structure and node content. In general, there are two types of graph representation learning methods: unsupervised methods and semi-supervised methods [19]. Unsupervised graph representation approaches focus on preserving the graph topology structure [20, 21]. For instance, LINE designs the objective function that preserves first-order and second-order proximity for learning node representations [20]. DeepWalk utilizes local information obtained from random walks to learn node representations [21]. As for the semi-supervised method, it utilizes some labeled vertices for representation learning [19]. Graph convolutional networks (GCN) [22] and graph attention networks (GAT) [23] are two powerful semi-supervised methods for solving the classification problem in the graph. GCN learns node representation by aggregating the features of its neighbor nodes. GAT further enhances the performance of GCN by using the attention mechanism. In this work, we first use unsupervised representation approaches to construct the implicit neighbor sets. We then borrow the recent advances of GCN and design a framework that utilizes constructed neighbor sets for recommendation.

Some recent studies like GCMC [14], NGCF [13] also adopt GCN ideas for recommendation, and they are designed for the user-item bipartite graph. Different from the above literature, we provide a new perspective for recommendation with the assistance of constructed user-user and item-item implicit neighborhoods.

2.3. Deep Learning

Deep learning is an emerging field of machine learning and is receiving a huge amount of attention at the moment [24]. Since deep learning technology can solve complex tasks while providing start-of-the-art results, it has achieved great success in many research areas [25], such as computer vision [26, 27], speech recognition [28, 29], and natural language processing [30, 31]. Due to the effectiveness of deep components, an increasing number of researchers are interested in integrating deep models into recommender systems [9, 10, 11, 32, 33, 34, 35, 36, 37, 38, 39, 40]. For instance, NeuCF is proposed to model the user-item interactions with a multi-layer perceptron [9]. DKEN is a deep end-to-end framework that uses deep neural networks and knowledge graph embedding for knowledge-enhanced recommendation [35]. RM-DRL utilizes convolutional neural networks and recurrent neural networks to produce user and item semantic feature vectors, respectively [36]. R-ConvMF integrates convolutional neural networks into probabilistic matrix factorization for document-based recommendations [39]. LUAR designs a neural attention mechanism to find important auxiliary reviews to address the sparsity problem in review-based recommendation [40]. In this work, we propose a deep recommendation framework in our method IRec. The framework is based on graph neural networks. The key component of the framework is that we devise an aggregator on the neighbor sets to update the feature representations of users and items. With the help of deep learning, we can obtain sufficient representation power for building a successful recommender system.

2.4. Feedback Information

The recommender systems collect user feedback information through the feedback techniques, and then utilize the feedback information to generate recommendations [41]. User feedback information can be roughly divided into two categories: explicit feedback (e.g., ratings) and implicit feedback (e.g., clicking and browsing history). There are several differences between the two types of feedback information [42]: (i) explicit feedback can capture both positive and negative user preferences, while implicit feedback can only be positive; (ii) compared with implicit feedback, explicit feedback can more accurately and unequivocally reflect users' interest in items; and (iii) explicit feedback is scarce and difficult to collect whereas implicit feedback is abundant and far outweighs the quantity of explicit data. To cater for different types of user feedback, researchers have designed corresponding explicit feedback recommendation methods [43, 44, 45] and implicit feedback recommendation methods [41, 46]. Some studies also use both explicit and implicit data for personalized recommendations [11, 47]. In this paper, we focus on explicit feedback recommender systems with users' rating information. We first utilize user explicit feedback to construct a neighbor set for each user and item. We then develop a deep framework that utilizes constructed neighbor sets for recommendation.

3. Methodology

In this section, we first introduce the notations and formulate the problems. We then describe two phases of IRec: neighbor construction and the recommendation framework.

3.1. Notations and Problem Formulation

In a typical recommendation scenario, we suppose there are M users $V' = \{u_1, u_2, \dots, u_M\}$ and N items $V = \{v_1, v_2, \dots, v_N\}$. We define $\mathbf{Y} \in \mathbb{R}^{M \times N}$ as the user-item historical interaction matrix. If user u_a has rated item v_i , i -th element y_{ai} in \mathbf{Y} is the rating score from u_a to v_i , otherwise we employ $y_{ai} = 0$ to represent the unknown rating.

Given the above information (V', V, \mathbf{Y}) , the first phase of IRec outputs user relational data $\mathcal{N} = \{\mathcal{N}(1), \dots, \mathcal{N}(M)\}$ and item relational data $\mathcal{N}_v = \{\mathcal{N}_v(1), \dots, \mathcal{N}_v(N)\}$. \mathcal{N}_u and \mathcal{N}_v contain implicit neighbors for users u and items v respectively. The details of building up relational data \mathcal{N}_u and \mathcal{N}_v are discussed in subsection 3.2. In the IRec's second phase, given the user-item interaction matrix \mathbf{Y} , user neighbor data \mathcal{N}_u and item neighbor data \mathcal{N}_v , the recommendation framework aims to learn a prediction function $y_{ai} = F(u_a, v_i, \mathbf{Y}, \mathcal{N}_u, \mathcal{N}_v)$, where y_{ai} is the predicted rating from user u_a to item v_i , an item with which they have never engaged before, and F is the framework parameters of function F . The details of this phase are discussed in subsection 3.3.

3.2. Neighbor Construction

In this subsection, we describe the neighbor construction strategy. This strategy contains three steps: (i) construct the user relational graph and item relational graph; (ii) map the user and item relational graphs to latent spaces respectively; and (iii) find implicit neighbors for users and items from their latent spaces. Next, we detail each step.

3.2.1. Step 1: Construction of relational graphs

The user-item interaction data can be represented as a bipartite graph structure. We first transform the bipartite graph structure to construct a user relational graph and item relational graph to identify user-user relationships and item-item relationships. In addition, to reflect the strength of the relationship in a fine-grained way, we construct the relational graphs as weighted graphs. To this end, we utilize users' opinions on items to construct two weighted relational graphs $\mathcal{G}_u = (V', \mathcal{E}_u)$ and $\mathcal{G}_v = (V, \mathcal{E}_v)$ for users and items respectively. In the user relational graph \mathcal{G}_u , the edge $e_{ab}^u \in \mathcal{E}_u$ connects two users if they have engaged with at least one common item before. In addition, e_{ab}^u is associated with a weight $w_{ab}^u > 0$ to indicate the relational strength between two users. Similarly, for the item relational graph \mathcal{G}_v , the edge $e_{ij}^v \in \mathcal{E}_v$ connects two items if they at least have been engaged by one common user and its weight $w_{ij}^v > 0$ indicates the relational strength between two items.

The opinions of users on items, such as ratings or reviews, play a crucial role in reflecting user preferences and item attributes. Here, we employ users' ratings to items as the user opinions. We utilize the difference in the user opinions to define the weight of the edges in the relational graphs. Specifically, for the edge e_{ab}^u between user u_a and user u_b in \mathcal{G}_u , the edge weight w_{ab}^u is defined as follows:

$$w_{ab}^u = Y_{max} - \frac{1}{C} \sum_{v_i \in C_{ab}} y_{ai} - y_{bi}, \quad (1)$$

where Y_{max} is the max score in all ratings (e.g., 5 in a 5-star system) and C_{ab} is the subset of V containing the items that u_a and u_b both rated before. Similarly, the weight w_{ij}^v of the edge e_{ij}^v in \mathcal{G}_v which connects item v_i and item v_j is defined as follows:

$$w_{ij}^v = Y_{max} - \frac{1}{\sum_{u \in D_{ij}} y_{ai} - y_{aj}}, \quad (2)$$

165 where D_{ij} is the subset of V' containing the users who has rated both v_i and v_j before.

166 Figure 2(a1) and (a2) are toy examples of building relational graphs for users and items in the simple user-item
 167 interaction scenario (in Figure 1 (a)). The advantage is that such relational graphs not only reflect the co-occurrence
 168 relations (one-hop neighbors) but also infer the high-order transitive relations (multi-hop neighbors).

169 For particular recommendation scenarios, one can employ other metrics to reflect users' opinions and design user-
 170 user and item-item opinion relations based on the metrics.

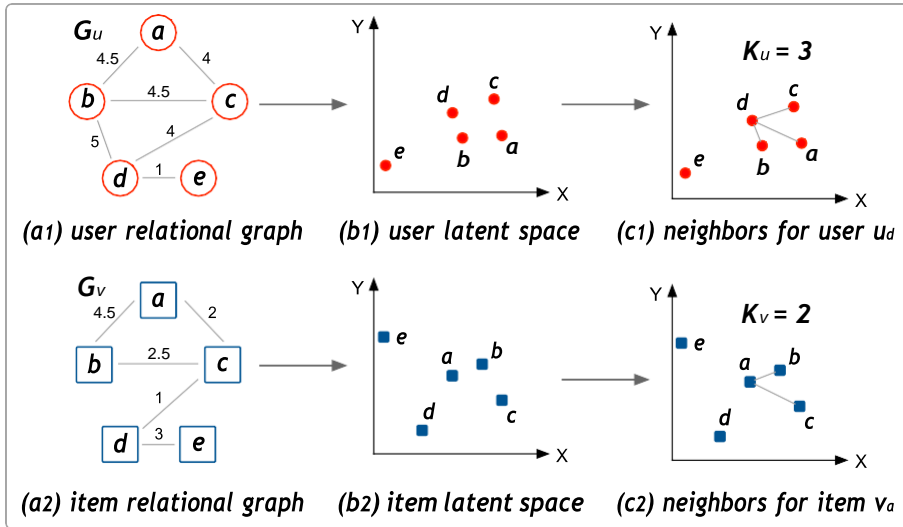


Figure 2: An illustration of the neighbor construction. (a) \rightarrow (b): Mapping user and item relational graphs to latent spaces respectively. (b) \rightarrow (c): Finding implicit neighbors for users and items, respectively.

3.2.2. Step 2: Relational graph mapping

After the relational graph construction, we utilize the node embedding method to map each relational graph to a latent continuous space. Specifically, for the user relational graph, we use a function $J_u: u \rightarrow z^u$ to map a user node $u \in V'$ from $\{i_u\}$ to a low-dimensional vector $z^u \in \mathbb{R}^{l_u}$ in a latent continuous space, where l_u is the dimension number of the vector for users. Similarly, for the item relational graph, we utilize another function $J_v: v \rightarrow z^v$ to map an item node $v \in V$ from $\{i_v\}$ to a low-dimensional vector $z^v \in \mathbb{R}^{l_v}$, where l_v is the vector dimension for items. Note that z^u and z^v can also be considered as positions for user u and item v in their latent spaces. After the mapping, both the structures and properties of the relational graphs are preserved and presented as the geometry in the corresponding latent space. Also, for a target node, nodes with important high-order transitive relations will appear near the target node, while nodes with irrelevant co-occurrence relations will appear far away from the target node.

Recent research reveals that a common embedding method that only preserves the connection patterns of a graph can be effective [48]. In this paper, we employ LINE [20], which can preserve both the local and global network structures, as our embedding method to map the user and item relational graphs to their corresponding latent continuous spaces. Note that one can employ or redesign other embedding methods to create suitable latent spaces, such as struc2vec [49], DeepWalk [21], and Poincare [50].

Figure 2(b1) and (b2) are examples of the latent space after mapping when the space dimensions $l_u = l_v = 2$. Although there is a co-occurrence relation between user d and user e , their distance in the latent space may be far away due to the small weight of their edge in the user relational graph.

3.2.3. Step 3: Construction of relational data

Based on the latent spaces, users' and items' relational data \mathcal{N} and \mathcal{N} can be constructed. Specifically, user u_a 's relational data $\mathcal{N}_{u_a}(a)$ is a user set which contains K_u (a pre-defined hyper-parameter) nearest neighbors in the user latent space based on the particular distance metric in the space. The construction of item neighbors is similar to the user. For instance, item v_i 's neighbors are defined as $\mathcal{N}_v(i)$ which contains top- K nearest neighbors items in the latent space to v_i . In this way, constructed relational data not only contains nodes with important high-order transitive relations, but also filters some nodes with irrelevant co-occurrence relations. Compared with only accounting for co-occurrence relations, our method reveals an in-depth understanding of potential user-user and item-item relations.

Figure 2(c1) and (c2) show examples of constructed relational data for user u_d and item v_a . The neighbor set for user u_d is $\mathcal{N}_{u_d}(d) = \{u_a, u_b, u_c\}$ when $K_u = 3$, and the neighbor set for item v_a is $\mathcal{N}_{v_a}(a) = \{v_b, v_c\}$ when $K_v = 2$.

3.2.4. Complexity analysis

Since we utilize LINE as our embedding method, the overall time complexity of relational graph mapping is $O(l_u \cdot S_u \cdot \mathcal{E}_u + l_v \cdot S_v \cdot \mathcal{E}_v)$, where S_u and S_v is the number of negative samples for users and items [20].

To obtain the nearest neighbors, the time complexity for a user is $O(M \cdot l_u)$. Similarly, for an item, the time complexity is $O(N \cdot l_v)$. In practice, we can utilize some acceleration computation methods proposed by previous works [51, 52, 53] to speed

up the process of obtaining implicit neighbors. Note that the relational data \mathcal{N} and \mathcal{N} can be computed offline in advance, so we can prepare \mathcal{N}_u and \mathcal{N}_v before generating recommendations. In this work, the construction of the relational data \mathcal{N}_u and \mathcal{N}_v is constrained to utilizing the user-item interaction records in the training split.

3.3. Recommendation Framework

In this subsection, we present the recommendation framework of the IRec, as illustrated in Figure 3. By taking a user u_a , an item v_i and their neighbors $\mathcal{N}_u(a)$ and $\mathcal{N}_v(i)$ as inputs, the framework outputs the predicted rating \hat{y}_{ai} from u_a to v_i . The recommendation framework consists of three parts: the embedding layer, the aggregation layer, and the prediction layer. Details of each part are described in the following.

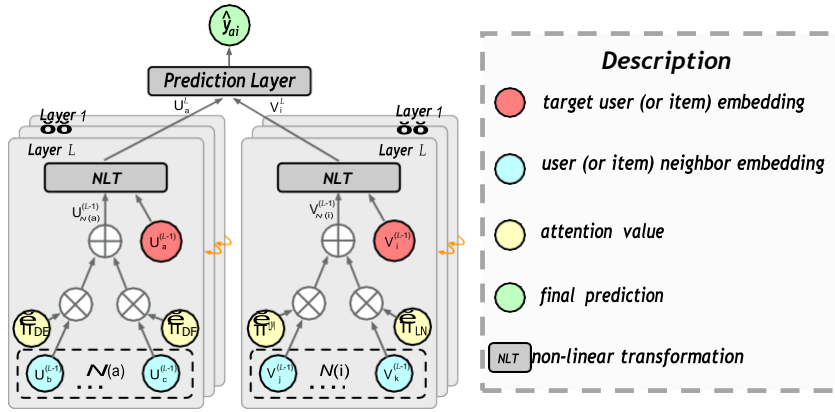


Figure 3: IRec's recommendation framework's architecture.

3.3.1. Embedding layer

The embedding layer transforms the primitive features of users and items (e.g., ID, user gender, item category etc.) into low-dimensional dense vectors called *embeddings*. Similar to mainstream embedding based recommender models [9, 10], we use one-hot vectors $\mathbf{u}^{id} \in \mathbb{R}^{M \times 1}$ and $\mathbf{v}^{id} \in \mathbb{R}^{N \times 1}$ to encode the ID features of user u_a and item v_i ,

respectively. We can obtain u 's embedding \mathbf{u} and v 's embedding \mathbf{v} , as follows:

$$\mathbf{u}_a = \mathbf{U}^T \cdot \mathbf{u}_a^{id}, \quad (3)$$

$$\mathbf{v}_i = \mathbf{V}^T \cdot \mathbf{v}_i^{id}, \quad (4)$$

where $\mathbf{U} \in \mathbb{R}^{M \times d}$ and $\mathbf{V} \in \mathbb{R}^{N \times d}$ are the embedding matrices for the user features and item features, respectively.

Here

d is the dimension size of the embeddings.

3.3.2. Aggregation layer

Aggregation is a key component in the framework because the user and item representations are bound up with the implicit neighbors by aggregation. By taking related inputs u_a , v_i , $\mathcal{N}_u(a)$ and $\mathcal{N}_v(i)$, we design an aggregator to update u 's and v 's feature representations as follows:

$$\mathbf{u}_a^* = \text{Agg}_u(u_a, \mathcal{N}_u(a)) = \text{agg}_u^{(H)}(\mathbf{u}_a, \text{agg}_u^{(L)}(\mathcal{N}_u(a))), \quad (5)$$

$$\mathbf{v}_i^* = \text{Agg}_v(v_i, \mathcal{N}_v(i)) = \text{agg}_v^{(H)}(\mathbf{v}_i, \text{agg}_v^{(L)}(\mathcal{N}_v(i))), \quad (6)$$

216 where \mathbf{u}_a and \mathbf{v}_i are the representations for user u_a and item v_i from their embedding tables \mathbf{U} and \mathbf{V} . Agg is a
 217 compound aggregation function used to update user and item representations. $\text{agg}_v^{(L)}$ is the low-level aggregation

function which maps the user (or item) neighbor set into a single embedding vector, and $\text{agg}^{(H)}$ is the high-level aggregation function which integrates the target user (or item) representation and the neighbor representation into a new representation of the target user (or item).

Specifically, we present the implementation of the function Agg_u . We illustrate the process for users and the same process works for items. We first compute the score n_{ab} between user u_a and one neighbor $u_b \in \mathcal{N}_u(a)$ via the attention mechanism as:

$$n_{ab} = (\mathbf{u}_a \circ \mathbf{u}_b)^T \tanh(\mathbf{w}_{u(L)} \cdot [\mathbf{u}_b \ \mathbf{v}_i] + \mathbf{b}_{u(L)}), \quad (7)$$

where \circ and \cdot mean concatenation operation and element-wise product between two vectors. \mathbf{w} and \mathbf{b} are parameters of the attention mechanism in the user low-level aggregation function. We also employ \tanh as the nonlinear activation function. In general, n_{ab} characterizes the importance of one neighbor u_b for the target user u_a . The intuition is as follows: the first term calculates the compatibility between user u_a and her neighbor u_b , and the second term computes the opinions of neighbor u_b on the target item v_i . Here, we simply employ the inner product on the two terms, however one can design a more sophisticated attention mechanism.

We then implement $\text{agg}_u^{(L)}$ to characterize user u 's implicit neighbors by the linear combination:

$$\mathbf{u}_{\mathcal{N}_u(a)} = \text{agg}_u^{(L)}(\mathcal{N}_u(a)) = \sum_{u_b \in \mathcal{N}_u(a)} n_{ab} \cdot \mathbf{u}_b, \quad (8)$$

where n_{ab} denotes the normalized attention coefficient:

$$n_{ab} = \frac{\exp(n_{ab})}{\sum_{u_b \in \mathcal{N}_u(a)} \exp(n_{ab})}. \quad (9)$$

For the high-level function $\text{agg}_u^{(H)}$, it aggregates the user representation \mathbf{u}_a and its neighbor representations $\mathbf{u}_{\mathcal{N}_u(a)}$ as the new representation of user u . We implement $\text{agg}_u^{(H)}$ by performing the summation operation between two representation vectors before employing nonlinear transformation:

$$\mathbf{u}_a^* = \text{agg}_u^{(H)}(\mathbf{u}_a, \mathbf{u}_{\mathcal{N}_u(a)}) = \alpha(\mathbf{w}_{u(H)} \cdot (\mathbf{u}_a + \mathbf{u}_{\mathcal{N}_u(a)}) + \mathbf{b}_{u(H)}), \quad (10)$$

where $\mathbf{w}_{u(H)}$, $\mathbf{b}_{u(H)}$ are parameters in the user high-level aggregation function and α is the nonlinear activation function. We try various kinds of operations, such as concatenation, and find the summation operation always shows the best performance.

Through a single aggregation layer, user (or item) representation is dependent on itself as well as the direct neighbors. We can further stack more layers to obtain high-order information from the multi-hop neighbors of users (or items). More formally, in the l -th layer, for user u_a and item v_i , their representations are defined as:

$$\mathbf{u}_a^l = \text{agg}_u^{(H)}(\mathbf{u}_a^{(l-1)}, \mathbf{u}_{\mathcal{N}_u(a)}^{(l-1)}), \quad (11)$$

$$\mathbf{v}_i^l = \text{agg}_v^{(H)}(\mathbf{v}_i^{(l-1)}, \mathbf{v}_{\mathcal{N}_v(i)}^{(l-1)}), \quad (12)$$

where $\mathbf{u}_{\mathcal{N}_u(a)}^{(l-1)}$ and $\mathbf{v}_{\mathcal{N}_v(i)}^{(l-1)}$ are defined as:

$$\mathbf{u}_{\mathcal{N}_u(a)}^{(l-1)} = \sum_{u_b \in \mathcal{N}_u(a)} n_{ab} \cdot \mathbf{u}_b^{(l-1)}, \quad (13)$$

$$\mathbf{v}_{\mathcal{N}_v(i)}^{(l-1)} = \sum_{v_j \in \mathcal{N}_v(i)} n_{ij} \cdot \mathbf{v}_j^{(l-1)}. \quad (14)$$

3.3.3. Prediction layer

After the L aggregation layer, we feed user representation \mathbf{u}^L and item representation \mathbf{v}^L into a function p

$\mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ for rating prediction.

a

$$\hat{y}_{ai} = p(\mathbf{u}_a^L, \mathbf{v}_i^L).$$

(15)

Here we implement the prediction function p as the MLP component [9], which can model complicated interactions between users and items. Specifically, the MLP component is implemented with two hidden layers (tower structure: $2d - d - 1$) as:

$$y_{ai} = \mathbf{w}_p^2 \cdot a \cdot \mathbf{w}_p^1 \cdot [\mathbf{u}_a^L \mathbf{v}_i^L] + \mathbf{b}_p^1 + \mathbf{b}_p^2 \quad (16)$$

where \mathbf{w}^* and \mathbf{b}^* denote the weight matrix and bias parameters in the MLP. One can utilize various prediction functions, such as inner product, to generate recommendations.

3.3.4. Learning algorithm

To estimate the parameters of the recommendation framework, we have the following objective function:

$$\min \mathcal{L} = \mathcal{L}_{\text{Rec}} + \lambda \|\Theta\|^2, \quad (17)$$

where \mathcal{L}_{Rec} measures the loss in the recommendation framework. Our paper is centered on the rating prediction for recommendation which is a regression problem. For the regression, we formulate \mathcal{L}_{Rec} as the squared loss:

$$\mathcal{L}_{\text{Rec}} = \frac{1}{|\mathcal{G}|} \sum_{(a,i) \in \mathcal{G}} (y_{ai} - \hat{y}_{ai})^2, \quad (18)$$

where \mathcal{G} denotes the observed ratings in \mathbf{Y} .

The second term in Equation (17) is the L2 regularization term to control model complexity and to avoid over-fitting. $\Theta = \{\mathbf{U}, \mathbf{V}, \mathbf{w}^a, \mathbf{w}^b, \mathbf{w}^{l(L)}, \mathbf{w}^{l(H)}, \mathbf{w}^{b(L)}, \mathbf{w}^{b(H)}, \lambda\} \in \{1, \dots, L\}$ is the parameter set in the framework.

The training process of the recommendation framework is summarized in Algorithm 1.

Algorithm 1 Training algorithm for the recommendation framework

Input: Interaction matrix \mathbf{Y} ; user neighbor set \mathcal{N}_u ; item neighbor set \mathcal{N}_v ; balancing factors λ ; learning rate η

Output: Prediction function $F(u, v, \mathbf{Y}, \mathcal{N}_u, \mathcal{N}_v)$

- 1: Initialize all parameters in Θ
 - 2: **repeat**
 - 3: Sample a minibatch of user-item interaction data from \mathbf{Y}
 - 4: Calculate $\mathcal{L} \leftarrow \mathcal{L} + \mathcal{L}_{\text{Rec}} \cdot \frac{1}{|\mathcal{G}|}$
 - 5: **for** each parameter $\theta \in \Theta$ **do**
 - 6: Calculate $\frac{\partial \mathcal{L}}{\partial \theta}$ on the minibatch by backpropagation
 - 7: Update θ by gradient descent with learning rate η
 - 8: **end for**
 - 9: **until** \mathcal{L} converges or is sufficiently small
 - 10: **return** $F(u, v, \mathbf{Y}, \mathcal{N}_u, \mathcal{N}_v)$
-

3.3.5. Complexity analysis

In this subsection, we discuss the time complexity and space complexity for the recommendation framework.

Time complexity. The time cost of the recommendation framework mainly comes from the aggregation layer. For users, the matrix multiplication in the aggregation layer has computational complexity $O(M \cdot K_u \cdot L \cdot d^2)$, where M is the number of users, K_u is the number of neighbors for each user, L is the total layers in the aggregation layer and d denotes the embedding size. Similarly, the time consumption for items in the aggregation layer is $O(N \cdot K_v \cdot L \cdot d^2)$, where N is the number of items and K_v is the number of neighbors for each item. In general, the overall training complexity is $O(M \cdot K_u \cdot L \cdot d^2 + N \cdot K_v \cdot L \cdot d^2)$. In fact, as shown in our experiment section, the framework reaches the best performance when $L = 1$. Also, the number of neighbors for users and items are limited with $K_u = K_v \ll \min\{M, N\}$. Therefore, the total time complexity of the recommendation framework in IRec is acceptable.

Space complexity. As shown in Equation (17), the model parameters Θ comprise two parts: embedding

249 tables for users and items $\Theta_1 = \{\mathbf{U}, \mathbf{V}\}$, and weight parameters in neural components $\Theta_2 = \{\mathbf{w}^1, \mathbf{w}^2, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l, \mathbf{w}^l\}$,
250 $\mathbf{w}_{v(H)}^l, \forall l \in \{1, \dots, L\}$. Parameter set Θ_1 is identical to that of the classical embedding-based models, such

as MF [8], BPR [54]. Parameter set O_2 is lighter than O_1 and can be neglected because (i) the parameters in O_2 are shared by all users and items; (ii) the dimension of each parameter in O_2 is far less than the number of users and items. Therefore, the space complexity of the framework is the same as the classical embedding models.

4. Experiments

In this section, we evaluate our method IRec in four real-world scenarios: movie, business, book, and restaurant recommendations. We first introduce the experiment settings, then present the experiment results. We also analyze the choice of hyper-parameters, the training efficiency, and some case studies in this section.

4.1. Experiment Setup

In this subsection, we introduce the datasets, baselines, evaluation protocols, and the choice of hyper-parameters.

4.1.1. Datasets.

Four datasets DVD1 movie dataset, Yelp2 business dataset, Douban³ book dataset, and Dianping⁴ restaurant dataset are used in our experiments. Each dataset contains users' ratings (ranging from 1 to 5) on the items. The statistics of the datasets⁵ are summarized in Table 1.

Table 1

Basic statistics for the four datasets: Movie (DVD), Business (Yelp), Book (Douban), and Restaurant (Dianping).

dataset	# users	# items	# interactions	density
Movie	2,433	12,838	32,893	0.105%
Business	10,580	13,870	171,102	0.117%
Book	11,777	20,697	190,590	0.078%
Restaurant	10,549	17,707	188,813	0.101%

4.1.2. Baselines.

To verify the performance of our proposed method IRec, we compared it with the following state-of-art recommendation methods. The characteristics of the comparison methods are listed as follows:

- **SVD++** is a well-known baseline, which is a hybrid model combining the latent factor model and the neighborhood model [16].
- **NFM** is a feature-based factorization model, which improves FM [55] by using the MLP component to capture the high-order feature interaction [37]. Here we concatenate user ID embeddings and item ID embeddings as input for NFM.
- **GCMC** is a graph-based recommendation framework which adopts a graph auto-encoder in a user-item bipartite graph to learn user and item embeddings for rating prediction [14].
- **NGCF** is a state-of-the-art graph-based recommender system which utilizes multiple propagation layers to learn user and item representations by propagating embeddings on the user-item bipartite graph [13]. For the rating prediction task, we replace the inner product with a two-layer MLP component in the prediction layer to enhance its performance.
- **CUNE** is a semantic social recommendation method which identifies semantic social friends from the collaborative user network and models these semantic relations as regularization terms to constrain the matrix factorization model [58]. In particular, they do not consider the neighbor information of items.

¹DVD: <https://www.librec.net/datasets.html>

²Yelp: <http://www.yelp.com/>

³Douban: <https://www.douban.com/>

⁴Dianping: <https://www.dianping.com/>

⁵Datasets are published at: <https://www.dropbox.com/s/dzr0uk6zj155w0z/data.zip?dl=0>

⁶We have tried other deep factorization models DeepFM [56] and Wide&Deep [57], and find that NFM is slightly better than them. Therefore, Anchen Li, Bo Yang, Huan Huo, and Farookh Khadeer Hussain: *Preprint submitted to Elsevier*

- **CMN** is a state-of-the-art memory-based model which designs the memory slots of similar users to learn user embeddings [12]. Note that it only focuses on the user's neighbors without accounting for the information about similar items.
- **MMCF** is another state-of-the-art memory-based model which models user-user and item-item co-occurrence contexts by memory networks [10]. It is the work which is most related to us. Different from our methods, it only focuses on co-occurrence relations and ignores high-order transitive relations among users and items.

4.1.3. Evaluation Protocols.

Three classes of metrics are adopted to evaluate the recommendation quality:

(i) For the rating prediction task, we utilize mean absolute error (MAE) and root mean square error (RMSE) as the performance metrics, which are widely adopted in many related works [14, 16, 37, 59]. Smaller values of MAE and RMSE indicate a better recommendation and are defined as follows:

$$\text{MAE} = \frac{1}{|D_{\text{test}}|} \sum_{(a,i) \in D_{\text{test}}} |y_{ai} - \hat{y}_{ai}|, \quad (19)$$

$$\text{RMSE} = \sqrt{\frac{1}{|D_{\text{test}}|} \sum_{(a,i) \in D_{\text{test}}} (y_{ai} - \hat{y}_{ai})^2}, \quad (20)$$

where D denotes the test set of the rating records.

(ii) We further adopt three rank-based metrics precision, recall and F1-score to evaluate different methods. Following the approaches in [60], precision, recall and F1-score are defined as follows:

$$\text{Precision} = \frac{1}{M} \sum_{i=1}^M \frac{|\text{Fav}(i) \cap \text{Rec}(i)|}{|\text{Rec}(i)|}, \quad (21)$$

$$\text{Recall} = \frac{1}{M} \sum_{i=1}^M \frac{|\text{Fav}(i) \cap \text{Rec}(i)|}{|\text{Fav}(i)|}, \quad (22)$$

$$\text{F1-score} = \frac{2 \times \text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}, \quad (23)$$

where $\text{Fav}(i) = \{j \in Q \mid y_{ij} \geq 4\}$ and $\text{Rec}(i)$ is the item set that user i has interacted with in the test set. Since the users' ratings on the items range from 1 to 5, a rating of 4 or 5 usually indicate that the users like the items. Therefore, we define $\text{Fav}(i)$ as the favorite item set of user i . $\text{Rec}(i) = \{j \in Q \mid \hat{y}_{ij} \geq 4\}$ is the set of items which will be recommended to user i and $b(\cdot)$ is the rounding function, which rounds the predicted rating \hat{y}_{ij} to an integer rating. The bigger the precision and recall values, the better the ranking.

(iii) In addition, we also use the three top-k based ranking metrics in our experiment: Precision@k (Pre@k for short), Recall@k (Rec@k for short), F1@k (for short). Following the approaches in [60], three metrics are defined as follows:

$$\text{Pre@k} = \frac{1}{M} \sum_{i=1}^M \text{Pre@k}(i) = \frac{1}{M} \sum_{i=1}^M \frac{|\text{Fav}(i) \cap \text{Rec}(i)|}{k}, \quad (24)$$

$$\text{Rec@k} = \frac{1}{M} \sum_{i=1}^M \text{Rec@k}(i) = \frac{1}{M} \sum_{i=1}^M \frac{|\text{Fav}(i) \cap \text{Rec}(i)|}{|\text{Fav}(i)|}, \quad (25)$$

$$, \quad (26)$$

295

$$F1@k = \frac{2 \times \text{Pre}@k \times \text{Rec}@k}{\text{Pre}@k + \text{Rec}@k}$$

296

where (i) is the top k item set in the ranking list determined by the trained models for user i .

297

298

4.1.4. Parameter Settings.

299

For the neighbor construction phase, we define $l_u = l_v = 8$, $S_u = S_v = 5$ and utilize Euclidean distance to calculate the distance in the latent spaces. We implemented the recommendation framework of IRec with Tensorflow which is a Python library for deep learning. For each dataset, we randomly split it into training, validation, and test sets following the

6 : 2 : 2 ratio. We repeated each experiment 5 times and reported the average performance. The framework parameters are first initialized by the Xavier initializer [61], and then updated them by conducting mini-batch Adam [62]. For the selection of the activation function, we utilized LeakyReLU [63] by default.

For our framework, there are six key hyper-parameters that need to be tuned, including the dimension of embeddings d , layer size L , number of neighbors K , batch size b , balancing factor -1 and learning rate Y . The hyper-parameters were tuned on the validation set using a grid search which is widely used in many deep models [13, 19, 36]. Specifically, we

divide the hyper-parameters into two sets, model hyper-parameters $\{d, L, K\}$ and training hyper-parameters $\{b, -1, Y\}$, and then apply the grid search to these sets separately. One can utilize other methods proposed by previous works [64, 65, 66] for hyper-parameter optimization. Table 2 shows our hyper-parameter settings. The key hyper-parameter settings for the baselines are defined as follows. For NFM, we utilize a one-layer MLP component according to the original paper [37]. Regarding NGCF, we tune the depth of layer L between $\{1, 2, 3, 4\}$, and find NGCF performs best with $L = 2$ for the movie and book datasets, and $L = 3$ for the business and restaurant datasets. For CMN and MDCF, the memory hop H is tuned between $\{1, 2, 3, 4\}$, and we find $H=2$ reaches the best performance. The settings for the other hyper-parameters for all the baselines are reached by either empirical study or following the original papers.

Deep models have a strong representation ability but they usually suffer from the over-fitting problems. To prevent over-fitting, we adopt L2 regularization (as mentioned in subsection 3.3.4) and the early stopping strategy [67] (i.e., premature stopping if RMSE on the validation set does not increase for 3 successive epochs). Figure 4 shows the training and validation error of each epoch of IRec. From the figure, we can see that 20 epochs are sufficient for our method to train and converge. If the model continues learning, then a situation of over-fitting will occur (i.e., the validation loss has begun to increase). We also tried the dropout technique [68], and found that introducing dropout masks slightly decreases the performance. A possible reason for this is that we do not introduce too many weight parameters with large dimensions in neural components, thus the dropout technique may not be as helpful to our model as it is to other deep and large models. Therefore, we do not introduce dropout mechanisms.

Table 2

Hyper-parameter settings for the four datasets: Movie (DVD), Business (Yelp), Book (Douban), and Restaurant

(Dianping).dataset	hyper-parameter settings					
Movie	$d = 8$	$L = 1$	$K = 6$	$b = 256$	$-1 = 10^{-4}$	$Y = 10^{-2}$
Business	$d = 16$	$L = 1$	$K = 6$	$b = 1024$	$-1 = 10^{-4}$	$Y = 2 \times 10^{-2}$
Book	$d = 16$	$L = 1$	$K = 5$	$b = 1024$	$-1 = 2 \times 10^{-4}$	$Y = 10^{-2}$
Restaurant	$d = 16$	$L = 1$	$K = 8$	$b = 1024$	$-1 = 10^{-4}$	$Y = 2 \times 10^{-2}$

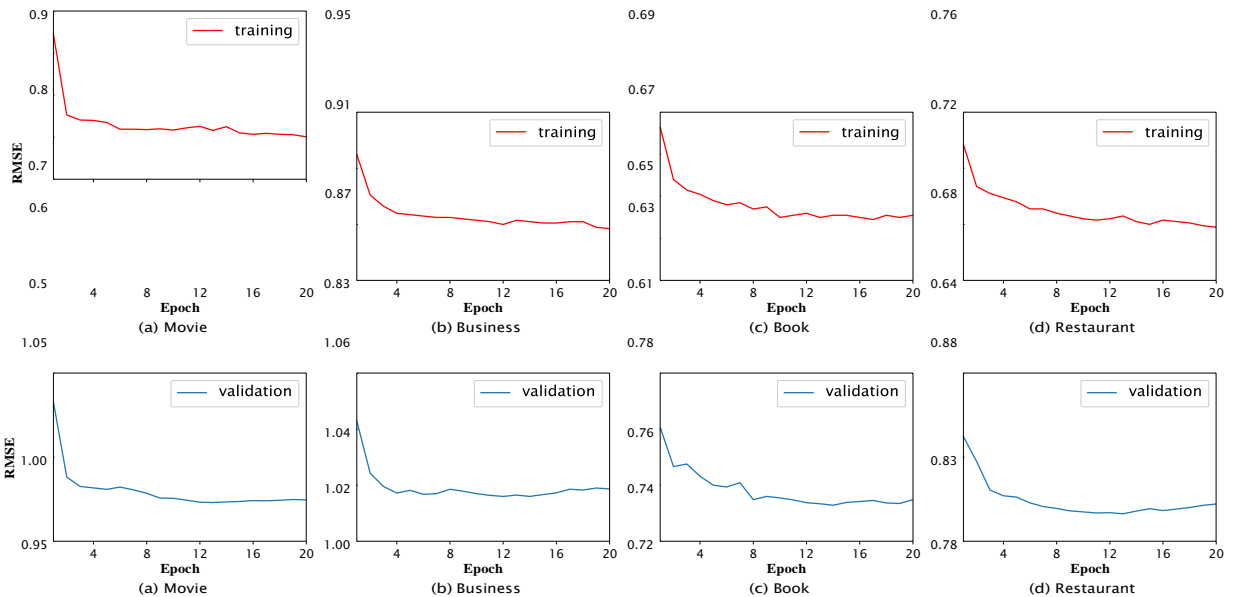


Figure 4: Training and validation error of each epoch of IRec on the four datasets.

4.2. Empirical Study

323 We conduct an empirical study to investigate the correlation between users (or items) and their co-occurrence
 324 neighbors. To formulate this issue, we utilize the difference in the explicit feedback (i.e. ratings) as the index. Taking
 325 users for example, if a co-occurrence user pair had engaged a common item before and the difference in their ratings
 326 on the item is greater than or equal to 2 (in a 5-star system), we believe this co-occurrence user pair encodes irrelevant
 327 (or noisy) information. To this end, we make statistics ???DO YOU MEAN we calculate the statistics??? on the four
 328 datasets used in this paper. The results are presented in Figure 5. We observed that both the user co-occurrence relation
 329 and item co-occurrence relation encode irrelevant (or noisy) information in the four datasets. In particular, more than 1/5
 330 of the user co-occurrence relations and 1/4 of the item co-occurrence relations in the business dataset exist in such cases.
 331 The above findings empirically demonstrate that not all co-occurrence relations contain useful information, so it is
 332 important to filter out irrelevant information to construct meaningful implicit neighbors for users and items.
 333




Figure 5: Empirical study on the four datasets. (a) Percentage of co-occurrence user pairs encoding irrelevant (or noisy) information. (b) Percentage of co-occurrence item pairs encoding irrelevant (or noisy) information.

4.3. Performance Comparison

334 Table 3 and Figures 6, 7, 8 show the performance of all the compared methods on the four datasets. From the
 335 results, we make the following main observations:

- 336 (i) SVD++ achieves poor performance on the four datasets, which indicates that shallow representation is insufficient
 337 to capture complex user-item interactions. NFM consistently outperforms SVD++, which suggests the significance of
 338 non-linear feature interactions between user and item embeddings in recommender systems. However, both SVD++
 339 and NFM ignore user-user and item-item relations.
- 340 (ii) Both SVD++ and CUNE are shallow models based on matrix factorization, while CUNE achieves better
 341 performance than SVD++; meanwhile, for the deep recommendation models, CMN and MMCF generally achieve
 342 better performance than NFM in most cases. These results suggest that considering potential user-user and item-item
 343 neighbors can enhance the recommendation performance. In addition, MMCF consistently outperforms CUNE and
 344 CMN. This makes sense since CUNE and CMN only account for user neighbor information, while MMCF considers
 345 co-occurrence information for both users and items.
- 346 (iii) For GNN-based models, NGCF achieves better performance than GCMC in most cases. The reason might be
 347 because GCMC only incorporates first-order neighbors for users and items in the bipartite graph, while NGCF models
 348 the high-order information.
- 349 (iv) Our method IRec consistently yields the best performance on the four datasets, which demonstrates the
 350 effective-ness of IRec on rating prediction and top- k recommendation. We also conduct one-sample t-tests and $p <$
 351 0.05 indicates that the improvements of IRec over the best baseline are statistically significant.

4.4. Data Sparsity and Cold Start Issues

352 As mentioned in many studies in the literature [43, 60, 69], data sparsity and cold start are two challenges faced
 353 by most rec-ommenders. In this subsection, we in

investigate the ability of our model to handle these two issues.

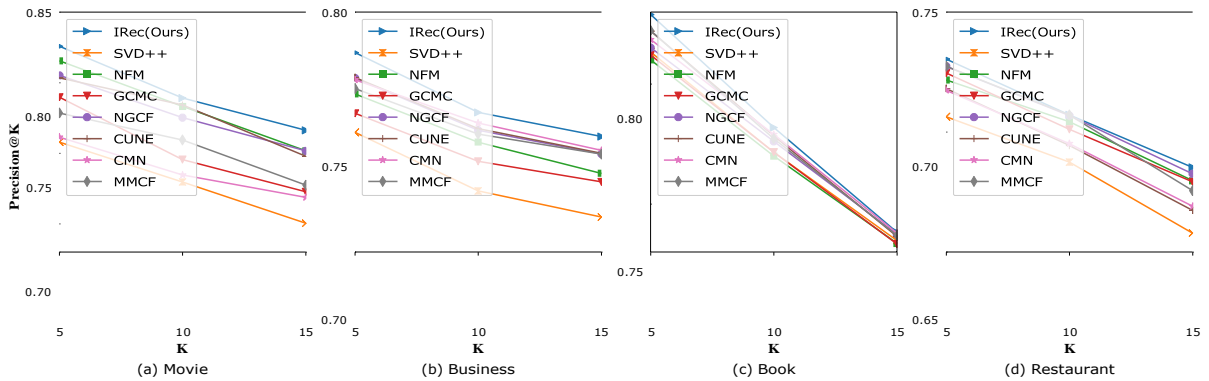
c

h
a
l
l
e
n
g
e
s
f
a
c
e
d
b
y
m
o
s
t
r
e
c
o
m
m
e
n
d
e
r
s
.
I
n
t
h
i
s
s
u
b
s
e
c
t
i
o
n
,
w
e
i
n

Table 3

Recommendation performance of seven competing methods and our method IRec on the four datasets. The proposed method IRec achieves the best performance on all metrics, as shown in boldface. * denotes the statistical significance for $p < 0.05$ compared to the best baseline.

Movie	SVD++	NFM	GCMC	NGCF	CUNE	CMN	MMCF	IRec
MAE	0.8183	0.8178	0.7790	0.7893	0.8181	0.8097	0.7804	0.7594*
RMSE	1.0503	1.0391	1.0177	1.0120	1.0406	1.0192	1.0068	0.9876*
Precision	0.8715	0.8908	0.8940	0.9027	0.8918	0.8733	0.8971	0.9166*
Recall	0.7386	0.8285	0.8197	0.7990	0.7985	0.8270	0.8640	0.8757*
F1-score	0.7996	0.8585	0.8552	0.8477	0.8426	0.8495	0.8802	0.8957*
Business	SVD++	NFM	GCMC	NGCF	CUNE	CMN	MMCF	IRec
MAE	0.8318	0.8045	0.8167	0.7995	0.8035	0.8122	0.8032	0.7889*
RMSE	1.0469	1.0357	1.0437	1.0266	1.0351	1.0348	1.0344	1.0164*
Precision	0.8187	0.8196	0.8507	0.8509	0.8514	0.8302	0.8349	0.8569*
Recall	0.6330	0.6538	0.6697	0.6769	0.6490	0.7022	0.7264	0.7361*
F1-score	0.7140	0.7274	0.7494	0.7540	0.7365	0.7609	0.7769	0.7919*
Book	SVD++	NFM	GCMC	NGCF	CUNE	CMN	MMCF	IRec
MAE	0.6079	0.5909	0.6040	0.5958	0.5960	0.5947	0.5926	0.5881*
RMSE	0.7648	0.7436	0.7549	0.7523	0.7442	0.7419	0.7387	0.7331*
Precision	0.8674	0.8820	0.8760	0.8835	0.8812	0.8571	0.8799	0.8856*
Recall	0.7122	0.7902	0.8168	0.7853	0.7639	0.7911	0.8137	0.8494*
F1-score	0.7822	0.8336	0.8454	0.8315	0.8184	0.8228	0.8455	0.8671*
Restaurant	SVD++	NFM	GCMC	NGCF	CUNE	CMN	MMCF	IRec
MAE	0.6640	0.6459	0.6480	0.6263	0.6365	0.6414	0.6238	0.6207*
RMSE	0.8682	0.8488	0.8385	0.8239	0.8356	0.8473	0.8198	0.8077*
Precision	0.8491	0.8460	0.8669	0.8746	0.8495	0.8524	0.8654	0.8750*
Recall	0.6907	0.7058	0.7156	0.7460	0.7404	0.7084	0.7252	0.7845*
F1-score	0.7618	0.7696	0.7840	0.8052	0.7912	0.7738	0.7891	0.8273*

**Figure 6:** The results of Precision@K on the four datasets.

363

356 364

357

358

359

360

361

362

4.4.1. *R*The data sparsity problem is a great challenge for most recommender systems. To investigate the effect of data sparsity, we **bin** the test users into four groups with different sparsity levels based on the number of observed ratings in the training data, and **keep each group** including a similar number of interactions. For example, [10,29) in the Movie dataset means each user in this group has at least 10 rating records and less than 29 rating records. Figure 9 shows the RMSE results for the different user groups with different models on the four datasets. From the results, we observe that our IRec outperforms the other methods in most cases. It is worth mentioning that IRec consistently outperforms all baselines in the first group on the four datasets, which verifies that our method IRec can maintain a good performance when data are extremely sparse.

n
d
a
t
a
s
p
a
r
s
e
s
c
e
n
a
r
i
o
s

T

h
e
d
a
t
a
s
p
a
r
s
i
t
y
p
r
o
b
l
e
m
i
s
a
g
r

Leveraging Implicit Relations for Recommender Systems

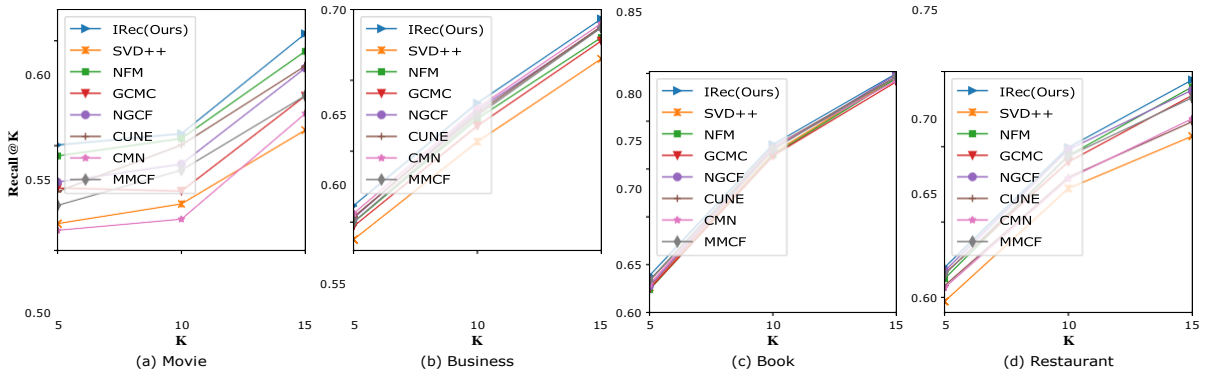


Figure 7: The results of Recall@K on the four datasets.




Figure 8: The results of F1@K on the four datasets.

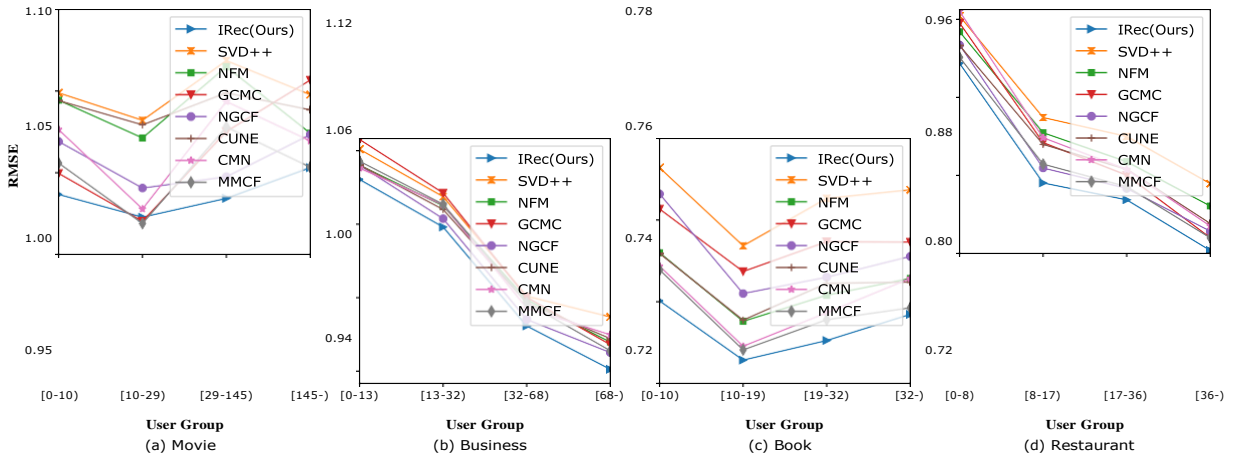


Figure 9: Performance comparison over the sparsity distribution of user groups on the four datasets.

370
365 371
366 372
367 373
368 374
369

4.4.2. Results in cold-start scenarios

We consider two cold-start scenarios, namely the cold-start user problem and cold-start item problem. We treat those who have rated x or fewer ratings as cold-start users and those that have been rated less than x as cold-start items. Following other work [60, 70], we set $x = 5$. Tables 4 and 5 illustrate the RMSE results of our method IRec and the other baselines in two cold-start scenarios on four datasets. In the tables, * denotes the statistical significance for $p < 0.05$, compared to the best baseline. We can see that our method IRec is beneficial to relatively inactive users and items in four recommendation scenarios.

4.5. Parameter Sensitivity

We explore the impact of three hyper-parameters: embedding size d , neighbor size k , and number of layers L in the aggregation layer. The results on the Movie and Business datasets are plotted in Figure 10. We make the following observations: (i) a proper embedding size d is needed. If d is too small, the model lacks expressiveness, while a too large d increases the complexity of the recommendation framework and may overfit the datasets. (ii) The performance

Table 4

RMSE results on testing cold start users on four datasets. The proposed method IRec achieves best performances on all metrics which are in boldface. * denotes the statistical significance for $p < 0.05$, compared to the best baseline.

Method	SVD++	NFM	GCMC	NGCF	CUNE	CMN	MMCF	IRec
Movie	1.0617	1.0525	1.0465	1.0336	1.0516	1.0424	1.0272	0.9980*
Business	1.0880	1.0987	1.0851	1.0743	1.0696	1.0778	1.0643	1.0399*
Book	0.7794	0.7713	0.7690	0.7621	0.7752	0.7700	0.7591	0.7409*
Restaurant	0.9824	0.9798	0.9792	0.9533	0.9605	0.9572	0.9443	0.9325*

Table 5

RMSE results on testing cold start items on four datasets. The proposed method IRec achieves best performances on all metrics which are in boldface. * denotes the statistical significance for $p < 0.05$, compared to the best baseline.

Method	SVD++	NFM	GCMC	NGCF	CUNE	CMN	MMCF	IRec
Movie	1.0861	1.0705	1.0446	1.0308	1.0574	1.0333	1.0205	0.9959*
Business	1.1676	1.1826	1.1539	1.1344	1.1612	1.1563	1.1440	1.1213*
Book	0.7981	0.7873	0.7880	0.7942	0.7910	0.7943	0.7708	0.7610*
Restaurant	0.8880	0.8958	0.9084	0.8876	0.8997	0.8929	0.8858	0.8671*

377 is improved with an increase in neighbor size K at the beginning because more neighbors provide more information.
 378 However, the performance takes a downward trend when K is larger than 6, since too many neighbors may introduce
 379 noise which reduces the accuracy of the prediction. (iii) In relation to the number of layers L in the aggregation layer,
 380 we find that when $L=1$, it is good enough because a larger L will bring massive noise when generating high-quality
 381 user and item representations and may lead to over-fitting. Similar results can be found in many other studies [19, 71, 72].

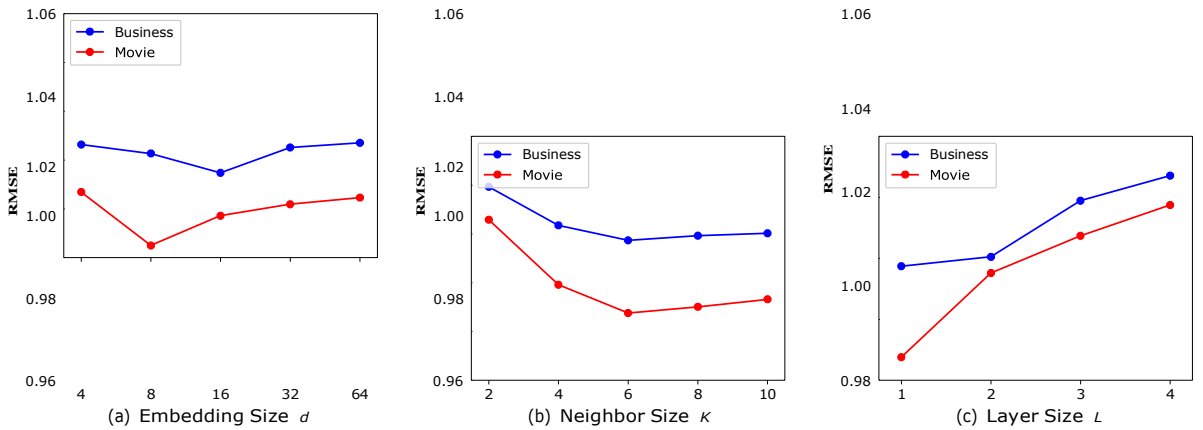


Figure 10: Parameter sensitivity of IRec's recommendation framework on the Movie and Business datasets w.r.t. (a) embedding size d , (b) neighbor size K , and (c) number of layers L in the aggregation layer.

389
 382
 390
 383
 391
 384
 392
 385
 386
 387
 388

In this section we discuss the experiments to explore the training efficiency of our IRec and two related methods CMN and MMCF, which explicitly account for the user-user and item-item neighbor information for recommendation. We first evaluate the training time of one iteration in the same environment (1.8 GHz Intel Core i5 and 8GB of RAM memory). Three methods are executed with 20 iterations and we report the average runtime. Table 6 shows the computation time for the four datasets. We observe that IRec is 1.5-2 times faster than CMN and 4-5 times faster than MMCF in one iteration, which confirms that our IRec has better training efficiency.

We then compare the number of trainable parameters for CMN, MMCF, and IRec. Table 7 summarizes the number of parameters of each method on embedding size 16 on the four datasets. We observe that IRec needs the least trainable parameters compared with CMN and MMCF. Specifically, MMCF requires more than double the number of parameters compared with our method IRec. This demonstrates IRec is a light yet effective model for recommendation.

Table 6

Training time of one iteration of CMN, MMCF, and IRec on the four datasets. 's' denotes 'second'.

Method	Movie	Business	Book	Restaurant
CMN	0.87s	7.59s	8.83s	7.59s
MMCF	2.58s	17.09s	19.98s	19.29s
IRec	0.59s	3.62s	5.06s	4.48s

Table 7

Number of parameters of CMN, MMCF, and IRec on embedding size 16 on the four datasets. 'k' denotes '10³'.

Method	Movie	Business	Book	Restaurant
CMN	300.7k	577.9k	725.4k	638.3k
MMCF	492.3k	786.0k	1042.8k	907.8k
IRec	245.9k	392.7k	521.1k	453.6k

4.7. Co-occurrence Relation vs Implicit Relation

In this subsection, we compare the co-occurrence relation and the implicit relation. To this end, we compare IRec with its variant CRec. CRec utilizes the recommendation framework of IRec and leverages the co-occurrence relation instead of the implicit relation for recommendation. Table 8 shows the performance of IRec and CRec on the four datasets. From the results, we find that the implicit relation can further improve the recommendation performance compared with the co-occurrence relation.

In the above experiments, we have validated the effectiveness of IRec which leverages the implicit relation for recommendation. Next, we investigate whether the co-occurrence relation and constructed implicit relation overlap with each other. For this purpose, we make statistics??we calculate the statistics on the Movie dataset from the perspective of users and show an overlapping relationship between the co-occurrence user pairs and the implicit user pairs.

Table 8

The results of MAE between IRec and its variant CRec on the four datasets.

Method	Movie	Business	Book	Restaurant
IRec	0.7594	0.7889	0.5881	0.6207
CRec	0.7784	0.7925	0.5971	0.6310




Figure 11: Overlapping relationship between the co-occurrence user pairs and implicit user pairs.

From Figure 11, we make the following findings: (i) only around 30% of implicit user pairs are co-occurrence relations. That is to say, if we only consider co-occurrence user pairs, a large portion of potentially relevant

potentially relevant users is missed; (ii) most co-occurrence user pairs have less relevance, which demonstrates the necessity of filtering the co-occurrence relation. Based on the above observations, we conclude that only accounting for the co-occurrence relations may result in a loss of useful information and it introduces some irrelevant information, while our method IRec leverages embedding methods to generate meaningful implicit relations for users and items.

4.8. Attention Analysis

Benefiting from the attention mechanism, we can visualize the attention weights placed on the neighbors for users and items, which reflects how the model learns. In this subsection, we analyze the attention mechanism from the perspective of users to show the learning process of IRec's recommendation framework and we obtain similar observations for items. We randomly selected one user u_{72} from the Restaurant dataset, and three relevant items v_{513} , v_{186} , v_{1917} (from the test set). Figure 12 shows the attention weights of the user u_{72} 's neighbors for the three user-item pairs. For convenience, we label the neighbor IDs starting from 1, which may not necessarily reflect the true ID from the dataset. From the heatmap, we make the following findings: (i) Not all neighbors make the same contribution when generating recommendations. For instance, for the user-item pair (u_{72}, v_{513}) , the attention weights of user u_{72} 's neighbor # 2 and # 5 are relatively high. The reason for this may be that neighbor # 2 and # 5 have rated item v_{513} in the training set. Therefore, neighbor # 2 and # 5 will provide more useful information when making recommendations.

(ii) For different items, the attention distributions of the neighbors are different, which reflects the attention mechanism that can adaptively measure the influence strength of the neighbors.

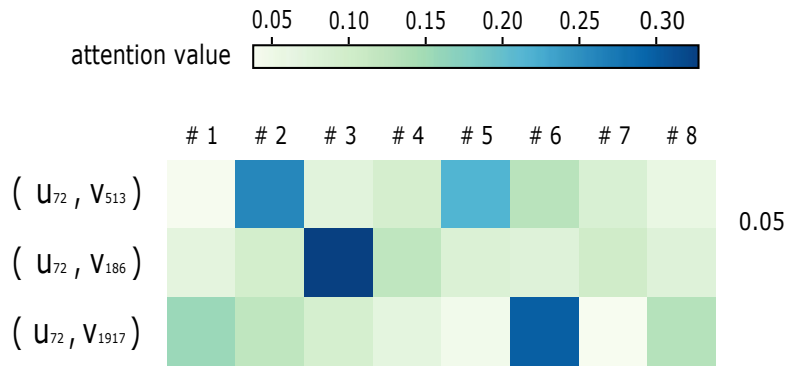


Figure 12: Attention heatmap for the neighbors of three user-item pairs from the restaurant dataset.

5. Conclusion and Future Work

In this work, we propose a method ca

In this work, we proposed a method called IRec to leverage implicit neighbors for better recommendations. IRec includes (i) a neighbor construction method that utilizes the user-item interaction information to construct implicit neighbor sets for each user and item; and (ii) a novel framework that integrates constructed neighbor sets into the recommendation task. We conducted extensive experiments on four real-world datasets. The experiment results demonstrate the superiority of IRec over several state-of-the-art methods in rating prediction and top- k recommendation. The results also show that our method is beneficial for relatively inactive and cold-start users.

For future work, we will (i) integrate side information into IRec such as knowledge graphs and social networks to further enhance the recommendation; (ii) employ more embedding methods in the neighbor construction phase to dig out user-user and item-item relationships; and (iii) try to generate recommendation explanations to comprehend user behaviors and item attributes.

Acknowledgements

h
i
s
w
o
r
k
w
a
s
s
u
p
p
o
r
t
e
d
b
y
t
h
e
N
a
t
i
o
n
a
l
N
a
t
u
r
a
l
S
c
i
e
n
c
e
F
o
u
n
d
a
t
i
o
n of China under Grant No. 61876069; Jilin Province Key Scientific and Technological Research and Development Project under Grant Nos. 20180201067GX and 20180201044GX; and Jilin Province Natural Science Foundation under Grant No. 20200201036JC.

CRediT authorship contribution statement

Anchen Li: Conceptualization, Methodology, Data curation, Formal analysis, Writing - original draft. **Bo Yang:** Methodology, Writing - review & editing, Funding acquisition, Supervision. **Huan Huo:** Writing - review & editing, Supervision. **Farookh Khadeer Hussain:** Writing - review & editing, Supervision.

References

- [1] B. Sarwar, G. Karypis, J. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in: Proceedings of the 10th international conference on World Wide Web, 2001, pp. 285–295.
- [2] X. He, H. Zhang, M.-Y. Kan, T.-S. Chua, Fast matrix factorization for online recommendation with implicit feedback, in: Proceedings of the 39th international ACM SIGIR conference on Research and Development in Information Retrieval, 2016, pp. 549–558.
- [3] B. Sarwar, G. Karypis, J. Konstan, J. Riedl, Item-based collaborative filtering recommendation algorithms, in: Proceedings of the 10th international conference on World Wide Web, 2001, pp. 285–295.
- [4] Y.-C. Lee, T. Kim, J. Choi, X. He, S.-W. Kim, M-bpr: A novel approach to improving bpr for recommendation with multi-type pair-wise preferences, *Information Sciences* 547 (2021) 255–270.
- [5] F. S. de Aguiar Neto, A. F. da Costa, M. G. Manzano, R. J. Campello, Pre-processing approaches for collaborative filtering based on hierarchical clustering, *Information Sciences* 534 (2020) 172–191.
- [6] G. R. Lima, C. E. Mello, A. Lyra, G. Zimbrao, Applying landmarks to enhance memory-based collaborative filtering, *Information Sciences* 513 (2020) 412–428.
- [7] X.-Y. Huang, B. Liang, W. Li, Online collaborative filtering with local and global consistency, *Information Sciences* 506 (2020) 366–382.
- [8] Y. Koren, R. Bell, C. Volinsky, Matrix factorization techniques for recommender systems, *Computer* 42 (8) (2009) 30–37.
- [9] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, T.-S. Chua, Neural collaborative filtering, in: Proceedings of the 26th international conference on World Wide Web, 2017, pp. 173–182.
- [10] X. Jiang, B. Hu, Y. Fang, C. Shi, Multiplex memory network for collaborative filtering, in: Proceedings of the 20th SIAM International Conference on Data Mining, SIAM, 2020, pp. 91–99.
- [11] H.-J. Xue, X. Dai, J. Zhang, S. Huang, J. Chen, Deep matrix factorization models for recommender systems., in: Proceedings of the 26th International Joint Conference on Artificial Intelligence, Vol. 17, Melbourne, Australia, 2017, pp. 3203–3209.
- [12] T. Ebesu, B. Shen, Y. Fang, Collaborative memory network for recommendation systems, in: Proceedings of the 41st international ACM SIGIR conference on Research and Development in Information Retrieval, 2018, pp. 515–524.
- [13] X. Wang, X. He, M. Wang, F. Feng, T.-S. Chua, Neural graph collaborative filtering, in: Proceedings of the 42nd international ACM SIGIR conference on Research and Development in Information Retrieval, 2019, pp. 165–174.
- [14] R. v. d. Berg, T. N. Kipf, M. Welling, Graph convolutional matrix completion, *arXiv preprint arXiv:1706.02263*.
- [15] F. Ricci, L. Rokach, B. Shapira, Introduction to recommender systems handbook, in: *Recommender systems handbook*, Springer, 2011, pp. 1–35.
- [16] Y. Koren, Factorization meets the neighborhood: a multifaceted collaborative filtering model, in: Proceedings of the 14th ACM SIGKDD international conference on knowledge discovery and data mining, 2008, pp. 426–434.
- [17] J. Wang, A. P. De Vries, M. J. Reinders, Unifying user-based and item-based collaborative filtering approaches by similarity fusion, in: Proceedings of the 29th international ACM SIGIR conference on Research and Development in Information Retrieval, 2006, pp. 501–508.
- [18] J. Sun, Y. Zhang, C. Ma, M. Coates, H. Guo, R. Tang, X. He, Multi-graph convolution collaborative filtering, in: Proceedings of the 19th IEEE International Conference on Data Mining, IEEE, 2019, pp. 1306–1311.
- [19] Y. Qu, T. Bai, W. Zhang, J. Nie, J. Tang, An end-to-end neighborhood-based interaction model for knowledge-enhanced recommendation, in: Proceedings of the 1st International Workshop on Deep Learning Practice for High-Dimensional Sparse Data, 2019, pp. 1–9.
- [20] J. Tang, M. Qu, M. Wang, M. Zhang, J. Yan, Q. Mei, Line: Large-scale information network embedding, in: Proceedings of the 24th international conference on World Wide Web, 2015, pp. 1067–1077.
- [21] B. Perozzi, R. Al-Rfou, S. Skiena, Deepwalk: Online learning of social representations, in: Proceedings of the 20th ACM SIGKDD international conference on knowledge discovery and data mining, 2014, pp. 701–710.
- [22] T. N. Kipf, M. Welling, Semi-supervised classification with graph convolutional networks, *arXiv preprint arXiv:1609.02907*.
- [23] P. Velickovic, G. Cucurull, A. Casanova, A. Romero, P. Lio, Y. Bengio, Graph attention networks, *arXiv preprint arXiv:1710.10903*.
- [24] S. Zhang, L. Yao, A. Sun, Y. Tay, Deep learning based recommender system: A survey and new perspectives, *ACM Computing Surveys* 52 (1) (2019) 1–38.
- [25] I. Goodfellow, Y. Bengio, A. Courville, Y. Bengio, *Deep learning*, Vol. 1, MIT press Cambridge, 2016.
- [26] K. He, X. Zhang, S. Ren, J. Sun, Deep residual learning for image recognition, in: Proceedings of the 2016 IEEE Conference on Computer Vision and Pattern Recognition, 2016, pp. 770–778.
- [27] A. Krizhevsky, I. Sutskever, G. E. Hinton, Imagenet classification with deep convolutional neural networks, in: Proceedings of the 26th Annual Conference on Neural Information Processing Systems, 2012, pp. 1097–1105.
- [28] D. Amodei, S. Ananthanarayanan, R. Anubhai, J. Bai, E. Battenberg, C. Case, J. Casper, B. Catanzaro, Q. Cheng, G. Chen, et al., Deep speech 2: End-to-end speech recognition in english and mandarin, in: Proceedings of the 33rd International Conference on Machine Learning, PMLR, 2016, pp. 173–182.
- [29] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath, et al., Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups, *IEEE Signal processing magazine* 29 (6) (2012) 82–97.

Leveraging Implicit Relations for Recommender Systems

- [30] K. Cho, B. Van Merriënboer, C. Gulcehre, D. Bahdanau, F. Bougares, H. Schwenk, Y. Bengio, Learning phrase representations using rnn encoder-decoder for statistical machine translation, arXiv preprint arXiv:1406.1078.
- [31] T. Mikolov, M. Karafiát, L. Burget, J. Cernocký, S. Khudanpur, Recurrent neural network based language model, in: Proceedings of the 11th conference of the international speech communication association, 2010, pp. 1045–1048.
- [32] J. Bobadilla, R. Lara-Cabrera, Á. González-Prieto, F. Ortega, Deepfair: deep learning for improving fairness in recommender systems, arXiv preprint arXiv:2006.05255.
- [33] J. Bobadilla, F. Ortega, A. Gutiérrez, S. Alonso, Classification-based deep neural network architecture for collaborative filtering recommender systems, *International Journal of Interactive Multimedia & Artificial Intelligence* 6 (1) (2020) 68–77.
- [34] M. Gao, J. Zhang, J. Yu, J. Li, J. Wen, Q. Xiong, Recommender systems based on generative adversarial networks: A problem-driven perspective, *Information Sciences* 546 (2020) 1166–1185.
- [35] X. Guo, W. Lin, Y. Li, Z. Liu, L. Yang, S. Zhao, Z. Zhu, Dken: Deep knowledge-enhanced network for recommender systems, *Information Sciences* 540 (2020) 263–277.
- [36] J. Ni, Z. Huang, J. Cheng, S. Gao, An effective recommendation model based on deep representation learning, *Information Sciences* 542 (2021) 324–342.
- [37] X. He, T.-S. Chua, Neural factorization machines for sparse predictive analytics, in: Proceedings of the 40th international ACM SIGIR conference on Research and Development in Information Retrieval, 2017, pp. 355–364.
- [38] J. Han, L. Zheng, H. Huang, Y. Xu, S. Y. Philip, W. Zuo, Deep latent factor model with hierarchical similarity measure for recommender systems, *Information Sciences* 503 (2019) 521–532.
- [39] D. Kim, C. Park, J. Oh, H. Yu, Deep hybrid recommender systems via exploiting document context and statistics of items, *Information Sciences* 417 (2017) 72–87.
- [40] D. Hyun, C. Park, J. Cho, H. Yu, Learning to utilize auxiliary reviews for recommendation, *Information Sciences* 545 (2020) 595–607.
- [41] E. R. Núñez-Valdéz, J. M. C. Lovelle, O. S. Martínez, V. García-Díaz, P. O. De Pablos, C. E. M. Marín, Implicit feedback techniques on recommender systems applied to electronic books, *Computers in Human Behavior* 28 (4) (2012) 1186–1193.
- [42] G. Jawaheer, M. Szomszor, P. Kostkova, Comparison of implicit and explicit feedback from an online music recommendation service, in: Proceedings of the 1st international workshop on information heterogeneity and fusion in recommender systems, 2010, pp. 47–51.
- [43] H.-C. Wang, H.-T. Jhou, Y.-S. Tsai, Adapting topic map and social influence to the personalized hybrid recommender system, *Information Sciences*.
- [44] L. Sheugh, S. H. Alizadeh, A novel 2d-graph clustering method based on trust and similarity measures to enhance accuracy and coverage in recommender systems, *Information Sciences* 432 (2018) 210–230.
- [45] Z. Zhang, H. Lin, K. Liu, D. Wu, G. Zhang, J. Lu, A hybrid fuzzy-based personalized recommender system for telecom products/services, *Information Sciences* 235 (2013) 117–129.
- [46] E. R. Núñez-Valdez, D. Quintana, R. G. Crespo, P. Isasi, E. Herrera-Viedma, A recommender system based on implicit feedback for selective dissemination of ebooks, *Information Sciences* 467 (2018) 87–98.
- [47] G. Li, Q. Chen, Exploiting explicit and implicit feedback for personalized ranking, *Mathematical Problems in Engineering* 2016.
- [48] H. Pei, B. Wei, K. C.-C. Chang, Y. Lei, B. Yang, Geom-gen: Geometric graph convolutional networks, arXiv preprint arXiv:2002.05287.
- [49] M. Nickel, D. Kiena, Poincare embeddings for learning hierarchical representations, arXiv preprint arXiv:1705.08039.
- [50] L. F. Ribeiro, P. H. Saverese, D. R. Figueiredo, struc2vec: Learning node representations from structural identity, in: Proceedings of the 23rd ACM SIGKDD international conference on knowledge discovery and data mining, 2017, pp. 385–394.
- [51] J. L. Bentley, Multidimensional binary search trees used for associative searching, *Communications of the ACM* 18 (9) (1975) 509–517.
- [52] S. M. Omohundro, Five balltree construction algorithms, *International Computer Science Institute Berkeley*, 1989.
- [53] T. Liu, A. W. Moore, A. Gray, C. Cardie, New algorithms for efficient high-dimensional nonparametric classification, *Journal of Machine Learning Research* 7 (6) (2006) 1135–1158.
- [54] S. Rendle, C. Freudenthaler, Z. Gantner, L. Schmidt-Thieme, Bpr: Bayesian personalized ranking from implicit feedback, arXiv preprint arXiv:1205.2618.
- [55] S. Rendle, Factorization machines, in: Proceedings of the 10th IEEE International Conference on Data Mining, IEEE, 2010, pp. 995–1000.
- [56] H. Guo, R. Tang, Y. Ye, Z. Li, X. He, Deepfm: a factorization-machine based neural network for ctr prediction, arXiv preprint arXiv:1703.04247.
- [57] H.-T. Cheng, L. Koc, J. Harmsen, T. Shaked, T. Chandra, H. Aradhya, G. Anderson, G. Corrado, W. Chai, M. Ispir, et al., Wide & deep learning for recommender systems, in: Proceedings of the 1st workshop on deep learning for recommender systems, 2016, pp. 7–10.
- [58] C. Zhang, L. Yu, Y. Wang, C. Shah, X. Zhang, Collaborative user network embedding for social recommender systems, in: Proceedings of the 17th SIAM international conference on data mining, SIAM, 2017, pp. 381–389.
- [59] C. Feng, J. Liang, P. Song, Z. Wang, A fusion collaborative filtering method for sparse data in recommender systems, *Information Sciences* 521 (2020) 365–379.
- [60] B. Yang, Y. Lei, J. Liu, W. Li, Social collaborative filtering by trust, *IEEE transactions on pattern analysis and machine intelligence* 39 (8) (2016) 1633–1647.
- [61] X. Glorot, Y. Bengio, Understanding the difficulty of training deep feedforward neural networks, in: Proceedings of the 13th international conference on artificial intelligence and statistics, JMLR Workshop and Conference Proceedings, 2010, pp. 249–256.
- [62] D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, arXiv preprint arXiv:1412.6980.
- [63] A. L. Maas, A. Y. Hannun, A. Y. Ng, Rectifier nonlinearities improve neural network acoustic models, in: Proceedings of the 30th International Conference on Machine Learning, Citeseer, 2013, p. 3.
- [64] J. Snoek, H. Larochelle, R. P. Adams, Practical bayesian optimization of machine learning algorithms, arXiv preprint arXiv:1206.2944.
- [65] J. Bergstra, R. Bardenet, Y. Bengio, B. Kégl, Algorithms for hyper-parameter optimization, in: Proceedings of the 25th Annual Conference on Neural Information Processing Systems, Neural Information Processing Systems Foundation, 2011, pp. 2546–2554.

Leveraging Implicit Relations for Recommender Systems

- 559 [66] J. Bergstra, Y. Bengio, Random search for hyper-parameter optimization, *Journal of machine learning research* 13 (2) (2012) 281–305.
- 560 [67] L. Prechelt, Early stopping-but when?, in: *Neural Networks: Tricks of the trade*, Springer, 1998, pp. 55–69.
- 561 [68] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, R. Salakhutdinov, Dropout: a simple way to prevent neural networks from overfitting, *The journal of machine learning research* 15 (1) (2014) 1929–1958.
- 562 [69] J. Herce-Zelaya, C. Porcel, J. Bernabé-Moreno, A. Tejada-Lorente, E. Herrera-Viedma, New technique to alleviate the cold start problem in recommender systems using information from social media and random decision forests, *Information Sciences* 536 (2020) 156–170.
- 563 [70] M. Jamali, M. Ester, A matrix factorization technique with trust propagation for recommendation in social networks, in: *Proceedings of the 4th ACM conference on Recommender systems*, 2010, pp. 135–142.
- 564 [71] H. Wang, F. Zhang, J. Wang, M. Zhao, W. Li, X. Xie, M. Guo, Exploring high-order user preference on the knowledge graph for recommender systems, *ACM Transactions on Information Systems* 37 (3) (2019) 1–26.
- 565 [72] H. Wang, M. Zhao, X. Xie, W. Li, M. Guo, Knowledge graph convolutional networks for recommender systems, in: *Proceedings of the 28th international conference on World Wide Web*, 2019, pp. 3307–3313.
- 566
- 567
- 568
- 569
- 570