

HS-GCN: Hamming Spatial Graph Convolutional Networks for Recommendation

Han Liu, Yinwei Wei, *Member, IEEE*, Jianhua Yin, *Member, IEEE*,
and Liqiang Nie, *Senior Member, IEEE*

Abstract—An efficient solution to the large-scale recommender system is to represent users and items as binary hash codes in the Hamming space. Towards this end, existing methods tend to code users by modeling their Hamming similarities with the items they historically interact with, which are termed as the first-order similarities in this work. Despite their efficiency, these methods suffer from the suboptimal representative capacity, since they forgo the correlation established by connecting multiple first-order similarities, i.e., the relation among the indirect instances, which could be defined as the high-order similarity. To tackle this drawback, we propose to model both the first- and the high-order similarities in the Hamming space through the user-item bipartite graph. Therefore, we develop a novel learning to hash framework, namely Hamming Spatial Graph Convolutional Networks (HS-GCN), which explicitly models the Hamming similarity and embeds it into the codes of users and items. Extensive experiments on three public benchmark datasets demonstrate that our proposed model significantly outperforms several state-of-the-art hashing models, and obtains performance comparable with the real-valued recommendation models.

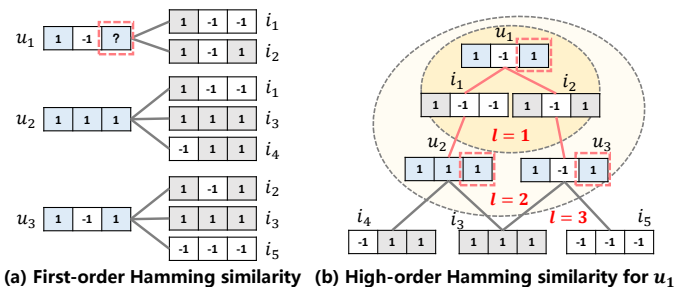
Index Terms—Hashing, Efficient Recommendation, High-order Similarity, Graph Convolutional Network, Hamming Space.

1 INTRODUCTION

THE recommender system is developed to locate the interested items from the overwhelming information according to users' preferences. Hence, how to measure the similarities between users and items is at the core of the personalized recommendation. Towards this end, existing studies [1], [2], [3] tend to follow a two-stage pipeline: representing the users and items with vectors [4], and then predicting their interactions by measuring the similarities between vectors. Despite the remarkable performance, these methods still face an inevitable problem that the computation grows exponentially with increasing users and items [5]. Theoretically, for recommending top- k preferred items for each user, the time complexity is $\mathcal{O}(NMK + NM\log k)$ when there are N users and M items, represented by K -dimensional embeddings in the latent space.

Diving into these methods, we can easily find that the problem mainly comes from the user-item similarity computations [6]. However, it is virtually impossible to design a new algorithm that not only can compute the similarity between two vectors but is more efficient than the inner-product. Therefore, some efforts have been dedicated to learning a new kind of representation—hash code—for the user and item, so as to alleviate the complexity [7], [8]. In particular, benefiting from such a code consisting of ± 1 bits, the measurement can be accelerated via XOR bit operation [9]. Hence, the time complexity of recommendation is significantly reduced and even constant time search is made possible by exploiting lookup tables [6]. Since the

- Han Liu, Jianhua Yin, and Liqiang Nie are with School of Computer Science and Technology, Shandong University, Qingdao 266200, China. E-mail: hanliu.sdu@gmail.com, jhyin@sdu.edu.cn, nieliqiang@gmail.com.
- Yinwei Wei is with School of Computing, National University of Singapore, Singapore. E-mail: weiyinwei@hotmail.com.
- Liqiang Nie is the corresponding author.



(a) First-order Hamming similarity (b) High-order Hamming similarity for u_1

Fig. 1. Illustration of the high-order and the first-order Hamming similarity in learning to hash.

computation efficiency of the user-item similarity is supercharged, the large-scale recommendation could be conducted efficiently [10], [11], especially on mobile application where the computational resource is very limited [12], [13]. Nevertheless, since users and items are approximately represented as the binary vectors, the performance of hashing-based recommendation models tends to be suboptimal. To deal with this drawback, several methods are developed to enhance the representation ability of the hash codes [14], [15], [16]. For instance, HashNet [17] utilizes the deep neural networks to learn the user and item embeddings and then map them into the Hamming space to calculate the Hamming distance between two vectors. More recently, inspired by the success of graph convolutional networks in representation learning, HashGNN [18] has been proposed to encode the local structural information into each node in the user-item interaction graph, and binarize their enhanced representations via a hash layer.

Although these methods enrich the user and item representations, we argue that they learn the hash codes merely with the similarities of interacted user-item pairs, but ignore the similarities hidden in the indirect interactions. To distinguish these similarities with the similarity computing

in Euclidean space [2], we term them as the first-order and high-order Hamming similarities, respectively. To illustrate our argument, we connect the direct interactions to model the indirect interactions, which builds a user-item bipartite graph shown in Figure 1(b), and compare it with the direct interactions shown in Figure 1(a) on hash code learning. In particular, the users (*i.e.*, u_1 , u_2 , and u_3) and items (*i.e.*, i_1 , i_2 , and i_3) are represented by binary vectors, and the presence of an edge signifies that the user interacts with the item. To learn the hash codes for the users, both the deep learning based and GNN-based models code their interacted items and minimize their first-order Hamming similarities (*i.e.*, $\langle u_1, i_1 \rangle$ and $\langle u_1, i_2 \rangle$), as shown in Figure 1(a). However, for u_1 , we argue that its code is hard to determine merely by the first-order Hamming similarity, because it lies right in the middle of two items. On the contrary, by incorporating the high-order Hamming similarity, as illustrated in Figure 1(b), the hidden cues can be captured by measuring its high-order Hamming similarities, such as $\langle u_1, u_2 \rangle$ and $\langle u_1, u_3 \rangle$, largely facilitating the hash code learning.

Therefore, we propose to code the users and items by explicitly modeling the high-order Hamming similarity on the user-item bipartite graph. It is worth noticing that this is different with HashGNN that learns the nodes' representations in Euclidean space and then transforms them into hash codes. We argue that the supervision signal for the hash code optimization merely comes from the first-order interactions in the history, which may aggravate the information loss of high-order Hamming similarity during the transformation.

However, conducting the graph convolutional operations in the Hamming space to capture the Hamming similarity is non-trivial. Following the terms in graph convolutional networks, we attribute the challenges into two aspects:

- Different from the operations on the continuous vectors, each element in hash codes is restricted to the binary values. Therefore, for each node, *how to aggregate the information from its neighbors in the Hamming space is the first challenge we face.*
- In order to represent the nodes in Hamming space, it is essential to preserve their Hamming similarities with the codes of their neighbors. Hence, *how to explicitly encode the Hamming similarities into each node is another technical challenge in this work.*

In order to address the outlined challenges, we develop a novel hashing-based recommendation model, named Hamming Spatial Graph Convolutional Networks (**HS-GCN**), consisting of the initial, propagation, and prediction layers. Specifically, in the initial layer, we recognize the user-item interactions as a bipartite graph and initialize the nodes with a hash code generation method proposed in [17], which guarantees the feasibility of end-to-end optimization. Upon the constructed graph, we devise a code propagation layer to implement the graph convolutional operations in Hamming space. More specifically, the layer could be divided into two components — hash code aggregation and hash code encoding. The former aggregates information from first-order similar neighbors by counting the bit-wise signs (*i.e.*, finding out the dominated bit value in each dimension)

of their corresponding hash codes. And the latter injects the Hamming similarities into the node hash codes by refining their bits consistent with the bit-wise signs of their neighbors. With the stacked propagation layers, we iteratively embed the learned hash codes into the nodes. And then, the interactions of user-item pairs could be predicted by scoring their affinities at the prediction layer. To evaluate our proposed model, we conduct extensive experiments on three publicly accessible datasets and compare the performance with several state-of-the-art baselines. In addition to outperforming the hash-based models (e.g., HashNet and HashGNN), our proposed method achieves the results comparable with the real-valued based models, such as PinSage and GraphSAGE.

Overall, the main contributions of our work are summarized in three-folds:

- To the best of our knowledge, this is the first attempt to explicitly model the high-order Hamming similarity between users and items in the recommender system, which enhances the representation ability of hash codes and accordingly optimizes the user-item interaction prediction.
- We present a GCN-based hashing recommendation model, named Hamming Spatial Graph Convolutional Network (HS-GCN), which explicitly captures the first- and high-order Hamming similarities. In particular, we develop the novel graph convolutions operation—hash code aggregation and hash code encoding—in the Hamming space.
- Extensive experimental results on three real-world datasets have demonstrated that the proposed model yields a substantial performance improvement compared with several state-of-the-art baselines. As a side contribution, we have released the codes to facilitate other researchers: <https://github.com/hanliu95/HS-GCN>.

2 RELATED WORK

In this section, we briefly review graph-based methods for recommendation and learning to hash for recommendation.

2.1 Graph-based Methods for Recommendation

Machine learning on graphs is an important task with the advantage of incorporating structural information. As one of the primary applications, representation learning on the user-item graph structure has been widely-studied in recommendation scenarios by utilizing information propagation. In this line of research, ItemRank [19] and BiRank [20] make early efforts for label propagation. In particular, these methods directly propagate the user preference scores (*i.e.*, labels) on the graph, scoring connected items with similar labels for a user. However, these methods are essentially neighbor-based and insufficient to encode structural information of a graph.

Recently, GNNs have received focused attention [21], [22], since GNNs have special advantage on modeling the graph structure, especially information propagation, to guide the representation learning. However, early GNN-based methods suffer from expensive computation costs as the graph convolution on the spectral domain. Subsequently, GCNs exploit a graph convolution operation on the

spatial domain, which aggregates the embeddings of neighbors to refine the embedding of the target node [3], [23]. Motivated by the efficiency of graph convolution, GC-MC [24], PinSage [25], and NGCF [2] employ GCNs to capture the propagation of latent Collaborative Filtering (CF) signals in the user-item interaction graph for recommendation. However, later studies argue the excessive complexity of GCNs. For example, LightGCN [26] develops a linear model by removing all redundant parameters. Experimental results demonstrate that the simplified design outperforms original GCNs by a large margin. This shows that the nonlinearity of GCNs is unnecessary for CF recommendation.

As aforementioned, we summarized two key points: 1) graph-based methods shed light on modeling the relational information propagation like the high-order similarity in the Hamming space, based on a user-item graph; 2) graph convolution has better performance after being simplified, verifying that the nonlinearity is not necessary in graph-based recommendation. Our work is highly inspired by these analyses, since the high-order Hamming similarity can be effectively modeled by the graph-based mechanism, and all the operations are simply linear in the Hamming space. Moreover, Heterogeneous GNN is proposed to be adapted upon structurally complex graphs to utilize richer information within them. Based on this technology, some methods effectively learn representations by aggregating different types of neighboring information in heterogeneous information network, such as HERec [27], MCRec [28], MEIRec [29], and ie-HGCN [30]. On contrary, our method works with simple interaction information which can be easily collected.

2.2 Learning to Hash for Recommendation

Another relevant research line is learning to hash for recommendation, which proceeds along two directions: unsupervised hashing and supervised hashing [31]. The former [32], [33] learns hash functions that encode objects to binary codes by training from unlabeled data, while the latter [34], [35], [36], [37] aims to learn more discriminative hash codes by exploring labeled signals, such as the feedback between users and items in recommendation scenarios. In general, early learning to hash for recommendation methods are essentially two-stage approaches [38]. However, the Hamming similarity between learned hash codes might not correspond to the original relevance between a user and an item, since there are quantization deviations when thresholding real values to binary bits during the binarization step. DCF [6] tackles the challenging discrete optimization problem and learns user and item hash codes directly. Therefore, the learned hash codes are able to model the intrinsic user-item relevance.

A prevailing trend is to leverage deep learning for recommendation. For instance, NFM [39] successfully introduces deep learning to enhance representation learning and matching function modeling in recommendation. In the light of this, deep learning to hash is subsequently developed and yields promising recommendation performance. Early efforts of deep learning to hash also adopt a two-stage strategy: the first stage employs deep networks to learn continuous representations and the second one uses the $\text{sign}(\mathbf{x})$

function to binarize the learned representations into binary hash codes. This category of methods, such as CNNH [40], DNNH [41], and DHN [15], also suffer from the quantization deviation. To alleviate it, HashNet [17] is proposed to devise a one-stage learning to hash technique to decrease the quantization deviation of binarization. It approximates the $\text{sign}(\mathbf{x})$ function with function $\tanh(\beta\mathbf{x})$, where β is a scaled parameter that increases during training. The infinite approximation makes the deviation negligible, and thus contributes to better recommendation performance.

In a sense, the aforementioned hashing techniques only consider the first-order similarity between hash codes. Therefore, we introduced graph-based techniques in hash learning to capture the high-order Hamming similarity. It is worth mentioning that several recent efforts have incorporated GNN insights into hashing, such as DGCN-BinCF [42], GCNH [43], and HashGNN [18]. Particularly, HashGNN [18] sets up a GNN in the continuous space, followed by a straight through estimator [44] for generating hash codes in the Hamming space. However, the hashing step leads to information loss, which impedes the capturing of intrinsic high-order Hamming similarity. To bridge this gap, we directly constructed a GCN in the Hamming space to model the high-order similarity.

Different from the diffusion models [45], [46], [47] which combine different similarity measures by multiple similarity graphs [48], [49], [50], [51], our method extends the Hamming similarity from first-order to high-order via a single graph, and the high-order similarity is exploited to learn hash codes with better representative capacity. Also different from hypergraph convolution networks that exploit multi-modal data in hypergraph for representation learning [52], our method is simply based on the single-modal interaction information in user-item graph, which can be obtained more easily.

3 PRELIMINARIES

We first give the definition of Hamming similarity, a similarity measurement for the hash codes in the Hamming space, highlighting that it is related to the number of the same bits. We then formulate the problem to be solved in our work.

3.1 Hamming Similarity

In the Hamming space, the user-item similarity is equivalent to the similarity between their corresponding hash codes, called Hamming similarity. Given user u and item i , we denote their hash codes as $\mathbf{h}_u \in \{\pm 1\}^K$ and $\mathbf{h}_i \in \{\pm 1\}^K$ respectively, where K represents the length of the codes. Accordingly, the Hamming similarity between them is defined as:

$$\text{sim}(u, i) = \frac{1}{K} \sum_{k=1}^K \mathbb{I}(h_{uk} = h_{ik}), \quad (1)$$

where h_{uk} and h_{ik} are the k -th bits of \mathbf{h}_u and \mathbf{h}_i , respectively, and $\mathbb{I}(\cdot)$ denotes the indicator function that returns 1 if the statement is true and 0 otherwise. Furthermore, it could be proven that $\text{sim}(u, i)$ is proportional to the number of same bits in \mathbf{h}_u and \mathbf{h}_i . As such, we rewrite the formulation [6] as:

$$\text{sim}(u, i) = \frac{1}{2} + \frac{1}{2K} \mathbf{h}_u^\top \mathbf{h}_i \propto \mathbf{h}_u^\top \mathbf{h}_i, \quad (2)$$

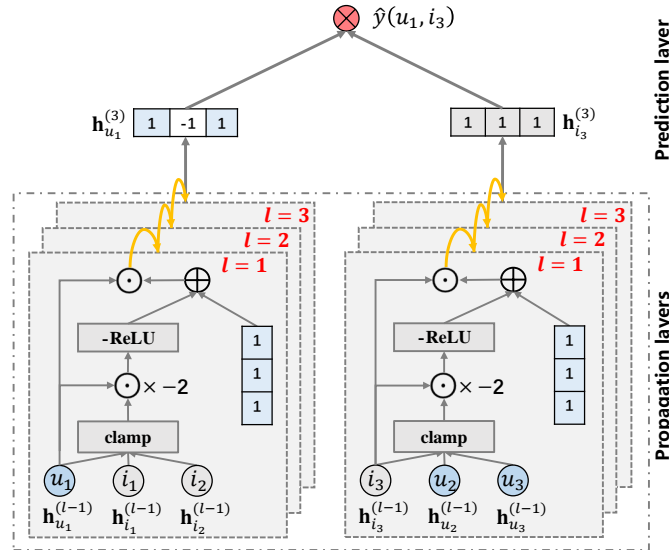


Fig. 2. Illustration of our proposed HS-GCN model (the arrow lines present the flow of information). The hash codes of user u_1 (left) and item i_3 (right) are refined with multiple hash code propagation layers.

where the derivation process is omitted. Based on this formulation, we directly employ the inner product of \mathbf{h}_u and \mathbf{h}_i to measure the similarity between u and i .

3.2 Problem Formulation

In this work, we aim to tackle the problem of mapping nodes in a user-item graph to hash codes for recommendation. We formulate the problem to be addressed in this paper as the following:

- **Input:** the user-item bipartite graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ that records historical user-item interactions in a graph structure. Whereinto, the node set consists of the user set \mathcal{U} of N users and the item set \mathcal{I} of M items, formally $\mathcal{V} = \mathcal{U} \cup \mathcal{I}$. And, the edge set \mathcal{E} contains the observed user-item interactions in the training phase.
- **Output:** the hash codes of all users and items, defined as $\{\mathbf{h}_u^* | u \in \mathcal{U}\} \cup \{\mathbf{h}_i^* | i \in \mathcal{I}\}$. Since $\mathbf{h}_u^* \top \mathbf{h}_i^*$ denotes the similarity score between user u and item i , and on top of that we can generate a ranking list of items for a given user and hence solve the practical problem of recommendation. To accurately measure the similarity score, we focus on devising a GCN for hash code propagation, which is able to learn node hash codes with high-order Hamming similarities latent in the user-item bipartite graph.

4 METHODOLOGY

In this section, we detail our proposed HS-GCN model, as illustrated in Figure 2. Specifically, the model consists of three components: 1) the initial layer that yields the trainable hash codes for the nodes in the user-item bipartite graph; 2) the propagation layer that explicitly model the Hamming similarities of node pairs and inject them into their codes for enhancement of the representation ability; and 3) the prediction layer that scores the similarity between the user and item by conducting the inner product of their hash codes. Moreover, we provide the optimization of HS-GCN in the form of matrix and compare HS-GCN with a graph-based learning to hash method HashGNN [18].

4.1 Initial Layer

The initial layer serves to generate the initial hash codes $\mathbf{h}_u \in \{\pm 1\}^K$ ($\mathbf{h}_i \in \{\pm 1\}^K$) for user u (item i), where K is the size of hash codes. Obviously, due to its discrete value, the hash code cannot be trained by the gradient descent. Hence, we alternatively employ a proxy to represent the users and items during the training phase. To achieve this goal, we follow the existing work [17] to reformulate the hash codes as,

$$\mathbf{h}_u = \lim_{\beta \rightarrow \infty} \tanh(\beta \mathbf{e}_u), \mathbf{h}_i = \lim_{\beta \rightarrow \infty} \tanh(\beta \mathbf{e}_i), \quad (3)$$

where $\mathbf{e}_u \in \mathbb{R}^K$ and $\mathbf{e}_i \in \mathbb{R}^K$ denote the parameter vector of user u and item i , respectively. By scaling the value of β , the outputs can approximate the desired codes. With the help of such a proxy, it is capable of addressing the ill-posed gradient issue and optimizing the representations during the training process. When the model reaches the convergence, we could infer the interactions between users and items via the learned hash codes, instead of their approximations.

In what follows, we take the approximate codes of users and items as the initial representations of the nodes in the user-item graph and detail our devised graph convolutional operations in Hamming space. For notational convenience, we denote them as $\mathbf{h}_u^{(0)}$ and $\mathbf{h}_i^{(0)}$, identifying the nodes' codes at the 0-th layer.

4.2 Propagation Layer

To perform the graph convolutional operations in Hamming space, we devise a novel propagation layer, which disentangles the operations into the hash code aggregation and hash code encoding. With the stacked propagation layers, we could explicitly model the first- and high-order Hamming similarities hidden in the user-item graph and optimize the nodes' representations by injecting the learned Hamming similarities into the codes.

In this section, without particular clarification, we elaborate our proposed model on the user nodes and conduct the same operations on item nodes.

4.2.1 Hash Code Aggregation

Regarding the graph convolutional operations, it is essential to capture the local structure information for each centric-node. To the end, we are inspired by the prior studies [2], [3] and develop a new aggregation to explicitly model the local structure information in Hamming space, which reveals the first-order Hamming similarity. Formally, the aggregation at $(l+1)$ -th propagation layer can be defined as,

$$\mathbf{m}_u^{(l)} = f(\mathbf{h}_u^{(l)}, \{\mathbf{h}_i^{(l)} | i \in \mathcal{N}_u\}), \quad (4)$$

where $\mathbf{h}_u^{(l)} \in \{\pm 1\}^K$ and $\mathbf{h}_i^{(l)} \in \{\pm 1\}^K$ recursively denote the hash codes of user u and item i at the l -th layer propagation, respectively. And, \mathcal{N}_u denotes the set of item neighbors of user u , which u directly interacted with. Using the function $f(\cdot)$, we could capture the informative signal $\mathbf{m}_u^{(l)}$ distilled by incorporating the node with its neighbors.

However, the implementation of this function is not straightforward in Hamming space. In this work, we implement it in two steps. In particular, considering the fact that

the transformation of the weighted matrix is inapplicable in Hamming space, we first adopt the non-weighted sum to integrate the bit-wise signals of the hash codes. Accordingly, we rewrite the aggregation in Hamming space:

$$\mathbf{m}_u^{(l)} = f(\mathbf{h}_u^{(l)} + \sum_{i \in \mathcal{N}_u} \mathbf{h}_i^{(l)}). \quad (5)$$

Then, since the hash codes constrain binary values (*i.e.*, +1 and -1), we capture the bit-wise dominate value in the hash codes of centric node and its neighbors, which is different with the graph operations in Euclidean space, like mean-and-max-pooling. We attribute it to the binomial distribution of the hash code. As such, we aim at designing a bit-wise function $f(\cdot)$ for aggregated codes and use each bit in $\mathbf{m}_u^{(l)}$ as the sign to reflect the dominate value at the corresponding bit of codes of u associated with its neighbors. Thereby, jointly considering the continuity of function, we implement $f(\cdot)$ with the interval limited clamp function as follows:

$$\mathbf{m}_u^{(l)} = \text{clamp}(\mathbf{x}) = \begin{cases} 1, & \text{if } \mathbf{x} > 1 \\ \mathbf{x}, & \text{if } -1 \leq \mathbf{x} \leq 1, \\ -1, & \text{if } \mathbf{x} < -1 \end{cases}, \quad (6)$$

where $\mathbf{x} = (\mathbf{h}_u^{(l)} + \sum_{i \in \mathcal{N}_u} \mathbf{h}_i^{(l)})$. Owing to the characteristic of the clamp function, we can guarantee $\mathbf{m}_u^{(l)} \in \{-1, 0, 1\}^K$. For each node, the function is expected to reflect the distribution of its neighbors' hash codes. It is able to determine the value of centric node, making it similar with its neighbors. Taking one bit in hash code as an example, the function will output -1 on this bit when the majority of corresponding bits in neighbors' hash codes are -1 , and vice versa. In a nutshell, we propose an aggregation operation in Hamming space, which integrates the codes from the centric node associated with its neighbors and distills the signal to reflect the bit-wise value distributions of hash codes.

4.2.2 Hash Code Encoding

Obtaining the aggregated codes from the local structure, we then encode the information into the nodes to preserve their first-order Hamming similarities towards their connected nodes. As for each node, it is of importance to minimize the total of bits which are different from the corresponding bits of its neighbors' codes. Therefore, we first compare the code of user u (*i.e.*, $\mathbf{h}_u^{(l)}$) with the obtained signs (*i.e.*, $\mathbf{m}_u^{(l)}$) to capture the number of different bits, as,

$$\mathbf{d}_u^{(l)} = \mathbf{h}_u^{(l)} \odot \mathbf{m}_u^{(l)}, \quad (7)$$

where \odot denotes the element-wise product, and $\mathbf{d}_u^{(l)} \in \{-1, 0, 1\}^K$ is a vector consisting of the bit-wise comparison result. More specifically, the value -1 in the output indicates the corresponding bit differs from most of neighbors'; whereas, the value 1 means the node has the same bit value with the majority of nodes around it. Besides, when the value equals to 0 , it implies that the corresponding bit is caught in the middle.

Then, we could refine the codes of user u according to each value in the obtained vector. Considering the element 0 existing in the vector, this operation is hardly implemented by conducting the element-wise product of two vectors, which probably outputs zero in addition to the binary

value (*i.e.*, +1 and -1). Therefore, we design a new function to avoid the negative influence of the value 0 . It is formally defined as,

$$\mathbf{c}_u^{(l)} = -\text{ReLU}(-2 \times \mathbf{d}_u^{(l)}) + \text{ones}, \quad (8)$$

where $\text{ones} = \{1\}^K$ is the vector that adjusts the outputs following the binary constraint. Based on this transform function, the values are tuned within $\{-1, 1\}$. Note that we convert the values 0 to 1 rather than -1 , making the -1 values in $\mathbf{c}_u^{(l)}$ exactly indicate the unquestioned bit differences between the node and its neighbors. This will improve the reliability of the following bit-wise refinement.

With the tuned vector, we could refine the hash code $\mathbf{h}_u^{(l)}$ by maximizing the similarity with its neighbors. Since $\mathbf{c}_u^{(l)}$ indicates the bits which are distinct with the majority of its neighbors, we utilize $\mathbf{c}_u^{(l)}$ to refine the bits of $\mathbf{h}_u^{(l)}$ with the element-wise product, formally,

$$\mathbf{h}_u^{(l+1)} = \mathbf{c}_u^{(l)} \odot \mathbf{h}_u^{(l)}, \quad (9)$$

where $\mathbf{h}_u^{(l+1)}$ denotes the hash code of user u learned at the $(l+1)$ -th propagation layer. Analogously, we can obtain the hash code $\mathbf{h}_i^{(l+1)}$ of item i by propagating hash codes from itself and its connected users. In summary, the designed hash code propagation layer contributes to exploit the first-order similarity information to relate user and item hash codes in the Hamming space.

4.2.3 Hash Code Propagation Rule in the Matrix Form

To offer a holistic view of the hash code propagation architecture and facilitate the batch implementation, we provide the matrix form of the layer-wise propagation rule as:

$$\mathbf{H}^{(l+1)} = \left(-\text{ReLU}(-2 \times \text{clamp}((\mathbf{A} + \mathbf{I})\mathbf{H}^{(l)}) \odot \mathbf{H}^{(l)}) + \text{Ones} \right) \odot \mathbf{H}^{(l)}, \quad (10)$$

where $\mathbf{H}^{(l)} \in \{\pm 1\}^{(N+M) \times K}$ are the hash codes for users and items obtained after l steps of propagation. $\mathbf{H}^{(0)}$ is the matrix of the initial hash codes after the initial layer, where $\mathbf{h}_u^{(0)} = \mathbf{h}_u$ and $\mathbf{h}_i^{(0)} = \mathbf{h}_i$. \mathbf{I} denotes the identity matrix, and $\text{Ones} \in \{1\}^{(N+M) \times K}$ is a matrix of all ones. \mathbf{A} denotes the adjacency matrix for the user-item graph, which is formulated as:

$$\mathbf{A} = \begin{bmatrix} \mathbf{0} & \mathbf{R} \\ \mathbf{R}^\top & \mathbf{0} \end{bmatrix}, \quad (11)$$

where $\mathbf{R} \in \{0, 1\}^{N \times M}$ is the user-item interaction matrix, and $\mathbf{0}$ is a all zero matrix. By implementing the matrix form propagation rule, we can simultaneously update the hash codes for all users and items in a rather efficient way. Moreover, it assists us to discard the node sampling procedure, which is commonly used to make graph convolution networks applicable on large-scale graph.

4.3 Prediction Layer

We stack multiple propagation layers to explore the high-order similarities between hash codes, which are vital for upgrading the hash codes to finally estimate the relevance between the user-item pair. In particular, we define L as the

depth of the final propagation layer, and the hash code of user u is recursively defined as:

$$\begin{cases} \mathbf{h}_u^{(1)} = \mathbf{c}_u^{(0)} \odot \mathbf{h}_u, \\ \mathbf{h}_u^{(2)} = \mathbf{c}_u^{(1)} \odot \mathbf{h}_u^{(1)} = \mathbf{c}_u^{(1)} \odot \mathbf{c}_u^{(0)} \odot \mathbf{h}_u, \\ \dots\dots \\ \mathbf{h}_u^{(L)} = \mathbf{c}_u^{(L-1)} \odot \mathbf{h}_u^{(L-1)} = \mathbf{c}_u^{(L-1)} \odot \dots \odot \mathbf{c}_u^{(0)} \odot \mathbf{h}_u. \end{cases} \quad (12)$$

It shows that our model obtains the final hash codes by element-wise multiplying the initial ones with a series of refinement vectors. These refinement vectors then modulate the bit signals of user u 's hash code layer by layer. As thus, the high-order Hamming similarity is injected into the hash code learning process. Note that the length of hash code is identically K in each layer according to Eqn.(12). Analogously, we can obtain the hash code for item i at the L -th layer.

After performing L propagation layers, we obtain multiple hash codes for user u , termed $\{\mathbf{h}_u^{(1)}, \dots, \mathbf{h}_u^{(L)}\}$. It is worth noting that we choose not to concatenate them as the final representation for a user, which is widely used in the graph-based recommendation models [2], [26]. The reason is that the length of concatenation codes probably harms the computational efficiency and memory space. Alternately, we take the hash codes at the last layer (i.e., $\mathbf{h}_u^{(L)}$ and $\mathbf{h}_i^{(L)}$) as the representations of user u and item i to predict their interaction, formally,

$$\mathbf{h}_u^* = \mathbf{h}_u^{(L)}, \quad \mathbf{h}_i^* = \mathbf{h}_i^{(L)}. \quad (13)$$

Finally, we apply the inner product to estimate the similarity between hash codes of the target user and item, which can be equivalently treated as the user's preference towards the item:

$$\hat{y}_{ui} = \mathbf{h}_u^{*\top} \mathbf{h}_i^*. \quad (14)$$

4.4 Optimization

In this work, we optimize the proposed model based on the users' implicit feedback, such as observations and purchases [53], [54]. Compared to explicit ratings, the implicit feedback is easier to collect but more challenging for exploring the user preference, due to its scarcity of negative feedback. To optimize the model by the maximization likelihood estimation (MLE), we formulate the likelihood function $P(r_{ui}|\mathbf{h}_u^*, \mathbf{h}_i^*)$ as the probability of the interaction between user u and item i equals r_{ui} given their trainable codes \mathbf{h}_u^* and \mathbf{h}_i^* . As such, its optimal objective can be instantiated as the pairwise logistic function as follows:

$$P(r_{ui}|\mathbf{h}_u^*, \mathbf{h}_i^*) = \begin{cases} \sigma(\hat{y}_{ui}), & r_{ui} = 1, \\ 1 - \sigma(\hat{y}_{ui}), & r_{ui} = 0, \end{cases} \quad (15)$$

where $\sigma(\cdot)$ is the standard sigmoid function. And, $\hat{y}_{ui} = \mathbf{h}_u^{*\top} \mathbf{h}_i^*$ is the estimated Hamming similarity between \mathbf{h}_u^* and \mathbf{h}_i^* . In this case, the more similar the hash codes are, the larger the conditional probability $P(1|\mathbf{h}_u^*, \mathbf{h}_i^*)$ will be, and vice versa. Therefore, Eqn.(15) is a reasonable extension of the logistic regression classifier for the pairwise classifica-

tion problem based on hash codes. Further, we can optimize the parameter by minimizing the following loss function:

$$\mathcal{L}_{cross} = - \sum_{r_{ui} \in \mathbf{R}} r_{ui} \log(\sigma(\hat{y}_{ui})) + (1 - r_{ui}) \log(1 - \sigma(\hat{y}_{ui})), \quad (16)$$

where \mathbf{R} is the aforementioned user-item interaction matrix that records the ground-truth implicit feedback. With this loss, the optimal object turns to reconstruct the observed interactions in the training phase.

In general, Eqn.(16) is effective in learning useful hash codes for users and items, and guarantees that the interacted user-item pairs have similar hash codes. However, the ranking among candidate items tends to be more important than their absolute scores, since recommender systems suggest items to users mainly according to their rankings. For implicit feedback data, the ranking relation can be obtained by sampling from the interaction matrix \mathbf{R} [1]. Specifically, we assume that the observed interactions, which are more reflective of a user's preference, should be assigned with higher estimated values than the unobserved ones. Inspired by this, we introduce a ranking order reinforced loss function as:

$$\mathcal{L}_{rank} = \sum_{(u,i,j) \in \mathcal{D}} \max(0, -\sigma(\hat{y}_{ui}) + \sigma(\hat{y}_{uj}) + \alpha), \quad (17)$$

where $\mathcal{D} = \{(u, i, j) | (u, i) \in \mathcal{R}^+, (u, j) \in \mathcal{R}^-\}$ denotes the triplet training data, \mathcal{R}^+ indicates the observed interactions, and \mathcal{R}^- is the unobserved interactions; α is the margin parameter that helps to control the observed and unobserved interactions. By minimizing \mathcal{L}_{rank} , the interacted items will gather around the user u more closely than items that are not interacted with u in the Hamming space.

By combining the aforementioned factors, the overall loss function to be minimized is formulated as:

$$\begin{aligned} \mathcal{L} &= \mathcal{L}_{cross} + \lambda_1 \mathcal{L}_{rank} + \lambda_2 \|\mathbf{E}\|_2^2 \\ &= - \sum_{r_{ui} \in \mathbf{R}} r_{ui} \log(\sigma(\hat{y}_{ui})) + (1 - r_{ui}) \log(1 - \sigma(\hat{y}_{ui})) \\ &\quad + \lambda_1 \sum_{(u,i,j) \in \mathcal{D}} \max(0, -\sigma(\hat{y}_{ui}) + \sigma(\hat{y}_{uj}) + \alpha) + \lambda_2 \|\mathbf{E}\|_2^2, \end{aligned} \quad (18)$$

where λ_1 is the trade-off parameter to balance the proportion between the entropy and the ranking loss. Note that both the pair-wise ranking loss and the point-wise reconstruction loss are necessary for the model optimization; $\mathbf{E} \in \mathbb{R}^{(N+M) \times K}$ denotes the real-valued parameter matrix in the initial layer as trainable model parameters, and λ_2 controls the L_2 regularization strength to prevent overfitting. In addition, we adopt mini-batch Adam [55] to optimize our model and update the model parameters. Particularly, for a batch of randomly sampled triplets $(u, i, j) \in \mathcal{D}$, we obtain their final hash codes \mathbf{h}_u^* , \mathbf{h}_i^* , and \mathbf{h}_j^* , and then update model parameters by using the gradients of the loss function.

4.5 Comparison with HashGNN

We compare our hash code propagation layer with HashGNN [18], which adopts the embedding propagation in Euclidean space to learn hash codes. HashGNN designs

the graph convolution operation not on the hash code $\mathbf{h}_u^{(l)}$ (or $\mathbf{h}_i^{(l)}$) but on its real-valued parameter approximation $\mathbf{e}_u^{(l)}$ (or $\mathbf{e}_i^{(l)}$). There obviously exists a quantization deviation between them, *i.e.*, $\mathbf{e}_u^{(l)} = \mathbf{h}_u^{(l)} + \epsilon_u^{(l)}$. Hence, the propagation layer of HashGNN can be formulated as:

$$\mathbf{e}_u^{(l+1)} = \tanh(\mathbf{W}^{(l)} \cdot \text{mean}(\mathbf{h}_u^{(l)} + \epsilon_u^{(l)} + \sum_{i \in \mathcal{N}_u} (\mathbf{h}_i^{(l)} + \epsilon_i^{(l)}))), \quad (19)$$

where the quantization deviations are accumulated with the aggregation. The nonlinear tanh function and weight matrix $\mathbf{W}^{(l)} \in \mathbb{R}^{K \times K}$ cause a new deviation between $\mathbf{e}_u^{(l+1)}$ and $\mathbf{h}_u^{(l+1)}$, which will make the final hash codes unable to obtain the high-order Hamming similarity. Compared with HashGNN, our proposed HS-GCN model implements the propagation of hash codes without introducing quantization deviations, effectively capturing the high-order Hamming similarity between hash codes.

5 EXPERIMENTAL SETUP

In this section, we first present the evaluation datasets, and then introduce our experimental settings, followed by elaboration of baseline methods.

5.1 Datasets

To justify the effectiveness of our proposed model, we conducted experiments on six publicly accessible datasets: MovieLens¹, Yelp², Amazon³, Gowalla, Pinterest, and Netflix⁴. Table 1 summarizes the detailed statistics of the evaluated datasets.

MovieLens: This is a widely used movie rating dataset, and we adopted the 1M version in our experiments. Similar to [56], we transformed the rating scores into implicit feedback, so that the label of each user-item pair is either 1 or 0, indicating whether the user rated the movie. The other datasets are processed in a similar way.

Yelp: Yelp is a famous online review platform for business, such as restaurants, bars, and spas. We selected the dataset from the latest version and used a 20-core setting to provide a denser dataset.

Amazon: Amazon-review is a popular dataset for commodity recommendation [57]. We selected Amazon-book from the collection. Similarly, we used the 20-core setting to ensure that each user and item have 20 interactions at least.

Gowalla: This is the check-in dataset [58] obtained from Gowalla, where users share their locations by checking-in. To ensure the quality of the dataset, we used the 10-core setting, *i.e.*, retaining users and items with at least ten interactions similar to [2].

Pinterest: This implicit feedback dataset is constructed by [59] for evaluating image recommendation. We adopted its processed version shared by [60].

Netflix: This is the large-scale movie rating dataset used in the Netflix challenge. We applied the 20-core setting to obtain a denser dataset.

TABLE 1
Statistics of the datasets.

Dataset	#Users	#Items	#Interactions	Density
MovieLens	6,040	3,706	1,000,209	4.47%
Yelp	40,500	58,755	2,024,283	0.09%
Amazon	36,783	77,379	2,402,416	0.08%
Gowalla	29,585	40,981	1,027,370	0.09%
Pinterest	55,187	9,916	1,500,809	0.27%
Netflix	429,584	17,764	99,884,887	1.31%

For each dataset, we randomly split its implicit feedback into two parts: 70% for training and the rest 30% for testing. Moreover, 10% interactions in the training set are randomly selected as the validation set for hyper-parameter tuning. For each observed user-item interaction in the training set, we treated it as a positive instance, and then adopted the negative sampling strategy to pair it with five negative items that the user did not interact with in the training set.

5.2 Experimental Settings

Evaluation Metric. For each user in the testing set, we regarded all items with no interaction with her/him as negative ones. Then the recommendation method predicts the user's preference scores (*e.g.*, Hamming similarity) over all the items, except the ones used in the training set. To measure the performance of top- K recommendation and preference ranking, we adopted HR@ K (Hit Ratio) [60] and NDCG@ K (Normalized Discounted Cumulative Gain) as the evaluation metrics. Whereinto, NDCG@ K are formally defined as:

$$\text{NDCG@}K = \frac{\text{DCG}_K}{\text{IDCG}_K}, \text{ and } \text{DCG}_K = \sum_{i=1}^K \frac{2^{r_{ui}} - 1}{\log_2(1 + i)},$$

where IDCG is the ideal DCG, and r_{ui} denotes the interaction status of the i -th recommended item. By default, we set $K = 50$ and 100. We reported the average metrics for all users in the testing set.

Model Implementation. We implemented our HS-GCN model via the development tool Pytorch⁵ and Pytorch Geometric⁶. We set the depth of HS-GCN L as two to model the second-order Hamming similarity. The hash code size is set to be 64. We optimized our model by Adam optimizer with a batch size of 3,000. Besides, we utilized the popular approach of Xavier [61] to initialize all the parameters in our model. And we applied the grid search for tuning the hyper-parameters based on the results from the validation set: the learning rate is tuned amongst $\{1\text{e-}4, 3\text{e-}4, 1\text{e-}3, 3\text{e-}3\}$, and finally set to be $3\text{e-}4$; the trade-off coefficient λ_1 is tuned from 0.1 to 1 with step size of 0.1, and finally set to be 0.1; the coefficient λ_2 of L_2 normalization is searched within $\{1\text{e-}8, 1\text{e-}7, \dots, 1\}$, and $1\text{e-}7$ is the optimal value. The margin parameter α is fixed to 0.2. Moreover, early stopping is adopted if HR@ K on the validation data does not increase for 10 successive epochs.

1. <https://grouplens.org/datasets/movielens/>.

2. <https://www.yelp.com/dataset>.

3. <http://jmcauley.ucsd.edu/data/amazon/>.

4. <http://www.netflixprize.com/>.

5. <https://pytorch.org/>.

6. <https://pytorch-geometric.readthedocs.io/en/latest/>.

5.3 Baseline Comparison

To demonstrate the effectiveness, we compared the performance of our proposed method with a series of state-of-the-art hash learning based models (*i.e.*, DHN, Hash_ste, HashNet, and HashGNN). Beyond these methods, we also introduce several real-valued recommendation models (*i.e.*, MF, BiNE, PinSage, and GraphSAGE) to justify that HS-GCN could achieve comparable results.

- **DCF** [6]: It is the first method to directly optimize hash codes based on the rating matrix. In our experiments, we adapted DCF for addressing the interaction matrix.
- **DFM** [5]: This is a discretely parameterized factorization machine for rating prediction. We adapted it to be trained on implicit data.
- **DHN** [15]: This is a popular deep hashing method for similar image recommendation. It learns high-quality hash codes by controlling the quantization deviation. We adapted it for user-item recommendation by replacing the original AlexNet [62] framework with Graph Convolution Networks (GCNs).
- **Hash_ste**: This is an effective end-to-end hash learning method based on straight through estimator [44], which enables the optimization of discrete variables by directly copying approximated gradients of them.
- **HashNet** [17]: This is a state-of-the-art deep hashing method, which devises a continuous and differentiable function to approximate the sign function. Similarly, we used GCN to replace the AlexNet in HashNet.
- **GCNH** [43]: It is a graph-based hashing method, which introduces a novel and efficient asymmetric graph convolution network, followed by a binarization step to learn similarity-preserving hash codes. We adapted it for recommendation via the user and item ID information as its input feature.
- **DGCN-BinCF** [42]: It is a graph-based hashing method. This method relaxes the binary constraint and makes continuous optimization possible to distill the ranking information from GCN into hash codes.
- **HashGNN** [18]: It is a graph-based hashing model, where GNN is utilized towards high quality hash codes.
- **MF** [63]: Matrix Factorization is the most common embedding model for recommendation, which only exploits the user-item direct interactions for recommendation.
- **BiNE** [64]: This embedding model adopts biased random walks for representation learning based on the bipartite graph. It applies an optimization framework for both explicit ratings and implicit feedback.
- **PinSage** [25]: PinSage is a graph embedding model. It combines efficient random walks and graph convolutions on the item-item graph, which incorporates both the graph structure and the node feature information. Hereon, we applied it on the user-item interaction graph.
- **GraphSage** [65]: As a famous graph embedding model, it is a general inductive framework that learns embeddings by sampling and aggregating features from a node's local neighborhood.

Baseline Implementation. For fair comparison, all baseline methods are implemented in Pytorch and Pytorch Geometric, except the publicly available implementation of

DCF⁷, DFM⁸, GCNH⁹, and BiNE¹⁰. Without specification, the default size of embedding or hash code is 64. For all graph-based baselines, we adopted a two-layer GCN initialized from Xavier, and used the Adam optimizer with a well-chosen mini-batch size for model optimization. The learning rate is tuned amongst {1e-4, 3e-4, 1e-3, 3e-3}. For HashGNN, we set the architecture of the graph layer, and the hyper-parameters, following the original paper [18]. For HashNet, we initialized the scale parameter β for $\tanh(\beta x)$ as 1, and exponentially increased it after constant epochs as suggested in [17]. For DGCN-BinCF, PinSage, and GraphSage, we implemented them according to the default architectures in the original papers [25], [42], [64], [65], and tuned their hyper-parameters based upon the model performance on validation set.

6 EXPERIMENTAL RESULTS

In order to validate the effectiveness of our proposed method, we conducted overall experiments to compare our proposed HS-GCN model with baseline methods in performance, efficiency, and sparsity issue. Moreover, in order to evaluate the effectiveness of components in HS-GCN, we performed ablation experiments on the key components including: propagation layers, ranking loss, and dropout technique.

6.1 Overall Experiments

We started by comparing the performance of all the methods, and then analyzed the efficiency of these methods. Finally, we explored the effectiveness of modeling high-order similarity under the sparse settings.

6.1.1 Performance Comparison

The results of our method and baselines over the experimented datasets are presented in Table 2. Besides, we reported the improvements and statistical significance test in Table 2, which are calculated between our proposed method and the strongest hashing baseline (highlighted with underline). Observing the results shown in Table 2 from top to bottom, we obtained the following key findings:

- DCF and DFM underperform on the implicit feedback datasets. This reflects the limitation of hashing methods that are designed on the basis of explicit feedback. DHN obtains poor performance on all the datasets. This indicates that the two-stage hashing method is insufficient to capture first-order similarities between users and items in the Hamming space, which is caused by the quantization deviation. Hash_ste consistently exceeds DHN across all cases, demonstrating the advantage of directly learning to hash strategy. However, Hash_ste fails to completely solve the quantization deviation between the straight through estimator and the sign function.
- Compared with DHN and Hash_ste, HashNet effectively eliminates the quantization deviation, and the better performance verifies that HashNet captures the first-order

7. <https://github.com/hanwangzhang/Discrete-Collaborative-Filtering>.

8. <https://github.com/hanliu95/DFM>.

9. <https://github.com/zxJohnFly/GCN>.

10. <https://github.com/clhctcjj/BiNE>.

TABLE 2

Overall performance comparison on six datasets. % Improv. and p-value denote the relative improvements (%) and t-test results of HS-GCN compared with HashGNN, respectively.

Methods	MovieLens		Yelp		Amazon		Gowalla		Pinterest		Netflix	
	HR@50	HR@100	HR@50	HR@100	HR@50	HR@100	HR@50	HR@100	HR@50	HR@100	HR@50	HR@100
DCF	0.0469	0.0910	0.0143	0.0281	0.0145	0.0288	0.1376	0.1471	0.0723	0.0840	0.0667	0.1143
DFM	0.0549	0.1068	0.0163	0.0310	0.0171	0.0336	0.1753	0.1820	0.0849	0.0931	0.0819	0.1330
DHN	0.0782	0.1393	0.0228	0.0430	0.0238	0.0448	0.2281	0.2415	0.1200	0.1328	0.1084	0.1813
Hash_ste	0.0955	0.1684	0.0349	0.0620	0.0327	0.0557	0.3268	0.3606	0.1125	0.2032	0.1211	0.2007
HashNet	0.1359	0.2232	0.0364	0.0651	0.0353	0.0599	0.3370	0.3830	0.1539	0.2730	0.1463	0.2337
GCNH	0.1552	0.2488	0.0389	0.0714	0.0394	0.0636	0.3979	0.4209	0.1847	0.3129	0.1515	0.2402
DGCN-BinCF	0.1554	0.2546	0.0358	0.0666	0.0373	0.0623	0.3420	0.3853	0.1834	0.3112	0.1744	0.2883
<u>HashGNN</u>	<u>0.1598</u>	<u>0.2557</u>	<u>0.0447</u>	<u>0.0774</u>	<u>0.0363</u>	<u>0.0611</u>	<u>0.3932</u>	<u>0.4329</u>	<u>0.1947</u>	<u>0.3214</u>	<u>0.2058</u>	<u>0.3185</u>
HS-GCN	0.2052	0.3169	0.0497	0.0883	0.0523	0.0885	0.4084	0.4480	0.2066	0.3322	0.2192	0.3432
MF	0.1332	0.2042	0.0319	0.0513	0.0349	0.0566	0.3360	0.3675	0.1400	0.2472	0.1981	0.3042
BiNE	0.1308	0.1868	0.0345	0.0546	0.0335	0.0560	0.3603	0.3812	0.1526	0.2582	0.2097	0.3346
PinSage	0.1587	0.2501	0.0503	0.0812	0.0402	0.0656	0.3791	0.4131	0.1837	0.3008	0.2279	0.3585
GraphSage	0.2132	0.3326	0.0581	0.0942	0.0408	0.0672	0.3894	0.4305	0.1919	0.3162	0.2334	0.3616
% Improv.	28.41%	23.93%	11.19%	14.08%	44.08%	44.84%	3.87%	3.49%	6.11%	3.36%	6.51%	7.76%
p-value	2.77e-4	7.01e-4	5.62e-3	1.95e-4	1.04e-3	1.27e-4	1.81e-4	1.39e-4	1.37e-3	5.08e-4	1.35e-4	2.04e-4

Methods	MovieLens		Yelp		Amazon		Gowalla		Pinterest		Netflix	
	N@50	N@100	N@50	N@100	N@50	N@100	N@50	N@100	N@50	N@100	N@50	N@100
DCF	0.0641	0.0785	0.0097	0.0137	0.0113	0.0173	0.2066	0.2599	0.0756	0.1109	0.0862	0.1014
DFM	0.0767	0.0937	0.0105	0.0165	0.0131	0.0187	0.2616	0.3041	0.0912	0.1242	0.1019	0.1294
DHN	0.1038	0.1251	0.0149	0.0226	0.0176	0.0264	0.3440	0.4080	0.1212	0.1710	0.1373	0.1660
Hash_ste	0.1305	0.1608	0.0225	0.0322	0.0263	0.0354	0.4692	0.5114	0.1248	0.1748	0.1515	0.1817
HashNet	0.1940	0.2233	0.0232	0.0335	0.0281	0.0378	0.4825	0.5245	0.1441	0.1877	0.1710	0.1900
GCNH	0.1443	0.2283	0.0294	0.0450	0.0319	0.0408	0.4957	0.5490	0.1616	0.2105	0.1723	0.2049
DGCN-BinCF	0.2298	0.2688	0.0268	0.0345	0.0302	0.0390	0.4646	0.5234	0.1436	0.2072	0.2230	0.2745
<u>HashGNN</u>	<u>0.2335</u>	<u>0.2630</u>	<u>0.0308</u>	<u>0.0430</u>	<u>0.0273</u>	<u>0.0374</u>	<u>0.5155</u>	<u>0.5512</u>	<u>0.1693</u>	<u>0.2198</u>	<u>0.3056</u>	<u>0.3502</u>
HS-GCN	0.3081	0.3376	0.0336	0.0483	0.0436	0.0585	0.5351	0.5696	0.1786	0.2325	0.3883	0.4501
MF	0.2023	0.2216	0.0240	0.0312	0.0303	0.0392	0.4977	0.5221	0.1029	0.1514	0.2469	0.2883
BiNE	0.1370	0.1765	0.0231	0.0320	0.0278	0.0370	0.5104	0.5346	0.1169	0.1858	0.2584	0.2999
PinSage	0.2316	0.2517	0.0380	0.0496	0.0337	0.0438	0.5136	0.5447	0.1536	0.2063	0.2923	0.3261
GraphSage	0.3195	0.3490	0.0434	0.0570	0.0349	0.0457	0.5368	0.5581	0.1643	0.2124	0.3703	0.4114
% Improv.	31.95%	28.37%	9.09%	12.33%	59.71%	56.42%	3.80%	3.34%	5.49%	5.78%	27.1%	28.5%
p-value	3.28e-3	2.87e-3	1.41e-2	8.47e-4	3.51e-4	1.44e-4	2.21e-3	4.09e-4	4.11e-3	1.69e-3	4.15e-3	3.72e-3

Hamming similarity more accurately. However, none of these methods explicitly models the high-order Hamming similarity in the hash code learning process, which leads to suboptimal results.

- GNN-based hashing methods significantly outperform the hashing methods based on deep learning. It makes sense since GNN improves the quality of user and item continuous representations prior to the binarization step, which indirectly enriches the final hash codes. Within GNN-based methods, HashGNN is superior since it is a real sense of end-to-end hash learning method, due to its ability of solving the back-propagation issue of sign function by gradient copy. In contrast, DGCN-BinCF and GCNH both use sign function as an extra step for binarization, inevitably resulting in quantization loss.
- It is clear that HS-GCN yields the best performance among hashing methods. In particular, the average improvements of HS-GCN over the strongest hashing baseline HashGNN *w.r.t.* HR@K are 26.17%, 12.64%, 44.46%,

3.68%, 4.74%, and 7.14%, and *w.r.t.* NDCG@K are 30.16%, 10.71%, 58.07%, 3.57%, 5.64%, and 27.8% in MovieLens, Yelp, Amazon, Gowalla, Pinterest, and Netflix, respectively. The reason is that HS-GCN is capable of capturing the high-order similarity of hash codes by directly constructing GCN in the Hamming space, while HashGNN only uses the first-order similarity to guide the hash learning. Additionally, we conducted the cross validation t-test, and p-value < 0.05 indicates that the improvements of HS-GCN over the strongest hashing baseline are statistically significant.

- HS-GCN significantly outperforms the embedding based baselines MF and BiNE. Moreover, compared with the graph embedding models GraphSage and PinSage, our model yields comparable performance. GraphSage performs better than most hash-based methods, since it utilizes GCN to learn real-valued embeddings for users and items, which have better representative ability than binary codes. It is worth mentioning that our performance on

Amazon is superior than the best graph embedding model GraphSage. This also verifies the importance of capturing the high-order similarity for hash learning.

6.1.2 Efficiency Comparison

In this section, we analyzed the efficiency of HS-GCN compared with baselines. In particular, we focused on both the training and the testing efficiency of these models, and recorded their running time. For fair comparison, all the models are trained on Ubuntu 16.04.5 with NVIDIA TITAN Xp, 12GB frame buffer and Python3.7 until the convergence, and tested on Windows 10 with Intel Core i7, 16GB RAM and Python3.7. During the training process, the hyper-parameters (*e.g.*, the training epoch, the batch size, and the learning rate) of all models follow the optimal settings. Though we adopted linear scan as the testing technique to recommend items for all methods, it is worth mentioning some techniques like multi-level indexing [66] can further accelerate the testing of hashing methods. Table 3 summarizes the results on different-scale Amazon dataset, where “1/8Amazon” is constructed by randomly selecting one in eight interactions from the complete dataset, and the rest can be done in the same manner. There are similar efficiency comparisons on other datasets, and thus we omitted them for space saving. Regarding the experiments, we found that:

- Our model costs much less training time than other methods. Compared to MF, HS-GCN and other graph-based models are more efficient in training since they quickly converge in experiments. This is reasonable since indirectly connected users and items are involved through the graph structure when optimizing the interaction pairs in mini-batch. Compared to other graph-based models, HS-GCN performs more efficiently during the training process, since it has the most concise model size that is the same as MF. Specifically, HS-GCN introduces no additional parameters except for the initial real-valued matrix E , while other models introduce the trainable weight matrix as transfer parameters in each propagation layer.
- Similar with most hashing models, HS-GCN achieves significant speedup compared with continuous recommendation models regarding the testing time. By jointly analyzing Table 2 and Table 3, we can see that HS-GCN not only can achieve competitive performance compared to state-of-the-art continuous models, but also is time-efficient. This verifies that HS-GCN is an operable solution for large-scale Web services to substantially reduce the computation cost of their recommendation systems [67].
- Across the different-scale datasets, the training efficiency ratio of HS-GCN is stable around 2.4 times based on HashGNN, and the testing efficiency ratio is stable around 4.9 times based on GraphSage. From Table 3, we can observe that both training and testing efficiency ratios of our method show upward trends with the scale of the dataset gradually increasing.

6.1.3 Performance Comparison *w.r.t.* Interaction Sparsity Levels

The sparsity issue usually limits the expressiveness of recommender systems, since inactive users lack sufficient

TABLE 3

Training/Testing time on different-scale Amazon dataset with 64 bits. The results are reported in hours and seconds, respectively. Training/Testing efficiency ratio is computed between HS-GCN and HashGNN/GraphSage, which is the best graph-based hashing/embedding baseline.

Methods	1/8Amazon		1/4Amazon		1/2Amazon		Amazon	
	Train	Test	Train	Test	Train	Test	Train	Test
HS-GCN	1.3	3.0	2.4	6.1	4.6	13.3	9.1	24.8
DHN	1.4	3.0	2.7	6.1	5.4	13.2	11.2	24.9
Hash_ste	1.5	3.0	3.0	6.0	6.1	13.3	12.6	24.7
HashNet	2.1	3.0	4.3	6.2	9.0	13.1	18.3	24.8
GCNH	1.9	3.1	3.6	6.1	7.2	13.3	14.8	25.0
DGCN-BinCF	1.6	3.1	3.2	6.1	6.2	13.3	12.6	24.8
HashGNN	2.7	3.0	5.6	6.1	11.1	13.2	23.3	24.9
MF	2.3	14.1	4.3	30.5	9.1	66.2	18.9	127.3
BiNE	1.4	14.2	2.6	30.2	5.4	66.2	10.9	127.0
PinSage	1.6	14.1	3.2	30.4	6.5	66.0	13.1	127.2
GraphSage	1.5	14.1	2.9	30.1	6.0	65.9	12.1	127.0
Efficiency Ratio	2.1	4.7	2.3	4.9	2.4	5.0	2.6	5.1

interactions to generate high-quality representations. This especially impedes the learning to hash methods only based on the first-order Hamming similarity. In this section, we attempted to answer whether exploring the high-order Hamming similarity is useful to alleviate this issue.

Towards this end, we first divided users into different sparsity-level groups. In particular, according to the interaction number per user in the training set, we divided the testing set into four groups, each of which has the same interaction sum. Taking Yelp dataset as an example, the interaction numbers per user are less than 23, 42, 89, and 1,713, respectively. Figure 3 shows the experiment results *w.r.t.* NDCG@50 on different user groups in MovieLens, Yelp, and Amazon. There is a similar performance trend *w.r.t.* HR@50, and thus we omitted this part for space saving. Regarding the results, we found that:

- HS-GCN consistently outperforms all other hashing baselines on all user groups. This demonstrates that exploiting the high-order Hamming similarity can improve the hash code learning for both active and inactive users.
- After analyzing Figures 3, we observed that HS-GCN achieves larger improvements in the first two groups (*e.g.*, 9.31% and 6.88% over the best baseline separately for < 23 and < 42 in Yelp) than that of the others (*e.g.*, 0.52% for <1,713 Yelp group). It verifies that capturing the high-order Hamming similarity is especially beneficial to the inactive users, since their hash codes are learned from more sufficient similar information besides first-order similarities. Hence, exploiting the high-order similarity is promising to solve the sparsity issue in hashing-based recommender systems.

6.2 Ablation Experiments

In this section, we studied the effect of key components in our proposed model, including propagation layers, ranking loss, and dropout. As the hash code propagation layer plays an important role in our model, we started by investigating the influence of layer numbers on the performance. Then

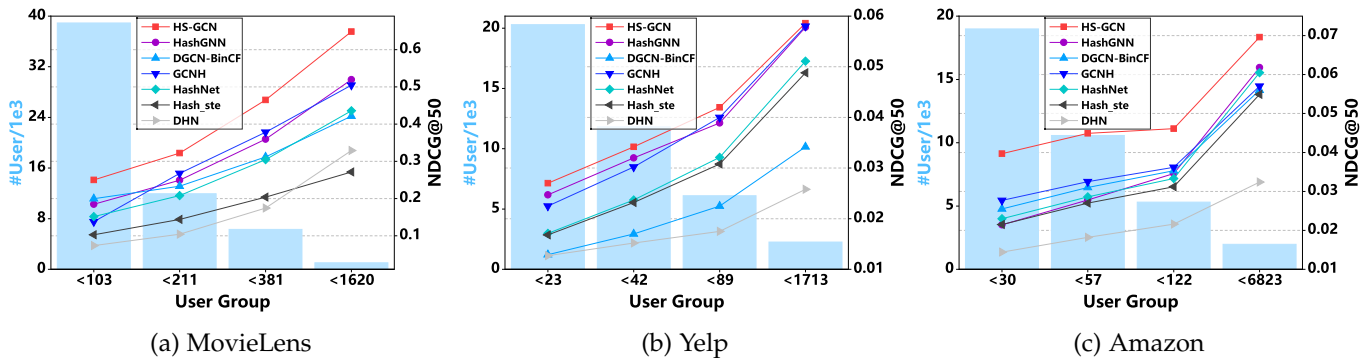


Fig. 3. Performance comparison over the sparsity distribution of user groups on three different datasets. Wherein, the background histograms indicate the number of users involved in each group, and the lines demonstrate the performance *w.r.t.* NDCG@50.

TABLE 4
Effect of hash codes propagation layer numbers (L).

Methods	MovieLens		Yelp		Amazon	
	HR@50	N@50	HR@50	N@50	HR@50	N@50
HS-GCN-1	0.1822	0.2774	0.0452	0.0316	0.0407	0.0340
HS-GCN-2	0.2052	0.3081	0.0497	0.0336	0.0523	0.0436
HS-GCN-3	0.1840	0.2798	0.0482	0.0329	0.0500	0.0421

we analyzed the ranking reinforced loss in the model optimization. Moreover, we jointly analyzed the affects of node dropout and bit dropout ratios.

6.2.1 Effect of Layer Numbers

To investigate the optimal number of hash code propagation layers, we varied the model depth. Particularly, we searched the layer numbers within $\{1, 2, 3\}$. Table 4 summarizes the experimental results, where HS-GCN-2 denotes the model with two propagation layers, and similar notations for the others. Jointly analyzing Table 2 and Table 4, we have the following observations:

- Obviously, HS-GCN-2 and HS-GCN-3 consistently outperform HS-GCN-1 across all datasets. The reason is that HS-GCN-1 only considers the first-order Hamming similarity, while HS-GCN-2 and HS-GCN-3 both utilize the high-order Hamming similarity.
- HS-GCN-2 achieves the best performance. When further stacking the propagation layer, we found that HS-GCN-3 leads to overfitting on all the datasets. The optimal layer number of our model is consistent with the prior HashGNN [18]. Therefore, the second-order similarity is sufficient for hash learning.
- When varying the number of propagation layers, HS-GCN is consistently superior to the learning to hash baselines across three datasets. This verifies that directly capturing the high-order Hamming similarity can facilitate the quality of hash codes.

6.2.2 Effect of the Ranking Loss

For capturing the relative item ranking in hash learning, we introduced the ranking reinforced loss in model optimization. During the experiments, we found that the ranking loss can contribute to both the initial layer and the prediction layer, where the ranking losses are computed by the initial state $\{\mathbf{h}_u, \mathbf{h}_i, \mathbf{h}_j\}$ and the final hash codes $\{\mathbf{h}_u^*, \mathbf{h}_i^*, \mathbf{h}_j^*\}$, respectively. To study the influence of these two types of

TABLE 5
Performance comparison of the models with and without different losses.

Methods	MovieLens		Yelp		Amazon	
	HR@50	N@50	HR@50	N@50	HR@50	N@50
HS-GCN	0.2052	0.3081	0.0497	0.0336	0.0523	0.0436
W/o Initial Rank	0.1773	0.2747	0.0468	0.0313	0.0509	0.0428
W/o Final Rank	0.1769	0.2566	0.0418	0.0281	0.0460	0.0385
W/o CR-Entropy	0.1564	0.2426	0.0156	0.0100	0.0289	0.0239

ranking losses, we conducted ablation study to validate how each of them contributes to the overall performance of our model.

Table 5 records the results of the ablation study on three datasets, where HS-GCN is our intact model with two types of ranking losses, which are computed upon triplets of the initial state $\{\mathbf{h}_u, \mathbf{h}_i, \mathbf{h}_j\}$ and the final hash codes $\{\mathbf{h}_u^*, \mathbf{h}_i^*, \mathbf{h}_j^*\}$, respectively. W/o Initial Rank denotes the ablation variant of HS-GCN via removing the ranking loss on the triplets of initial hash codes, while W/o Final Rank is the variant via removing the ranking loss on the final hash codes. From the table, we can observe that HS-GCN performs worse when ignoring the ranking loss for prediction. This demonstrates the capacity of ranking loss for more effective hashing. Moreover, we found that the ranking loss can improve the model performance more significantly when acting on the initial layer, since it injects the order information into the initial state. This also demonstrates that initialization enhances the quality of final hash codes, which is consistent with the findings of prior efforts [6], [11].

Following the operations in HashGNN [18], we introduce the cross-entropy loss to supervise the interactions between users and items they consumed before. We suggest that the cross-entropy loss maximization provides the signal to preserve the bipartite graph structure, which is the base of graph neural networks. Since this is not originally presented by us, we did not claim this design in our paper. To evaluate the effectiveness of each design, we do ablation study on the several datasets. In particular, we removed the binary cross-entropy loss from the final loss function, denoted as W/o CR-Entropy in Table 5, and observed that the variant is suboptimal compared with the intact one. The possible reason is that the binary cross-entropy loss is capable of supervising the hash codes to reconstruct the interactions between users and items, and thus optimizes their hash coding.

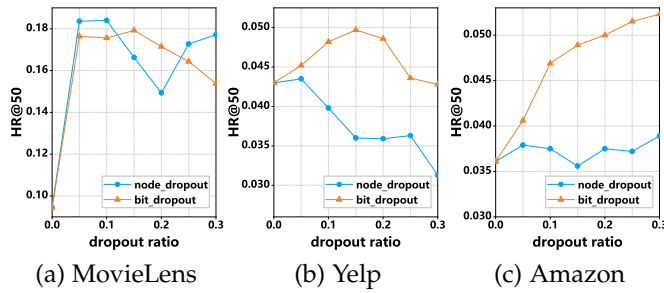


Fig. 4. Effect of node dropout and bit dropout ratios.

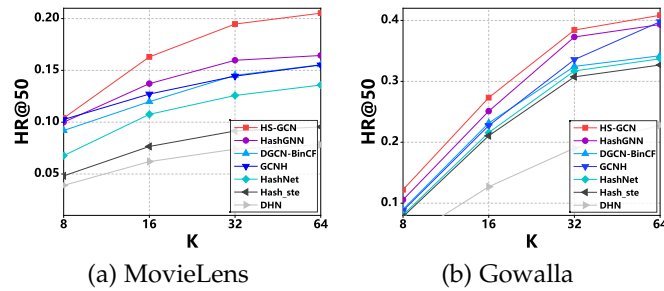


Fig. 5. Effect of the length of hash codes.

6.2.3 Effect of Dropout

Although GCNs have strong representation ability, they usually suffer from the overfitting problem. Dropout is an effective solution to prevent models from overfitting [68]. Following the prior GCN-based researches and considering our model architecture, we attempted to employ two dropout techniques: node dropout and bit dropout to improve the performance of HS-GCN. Node dropout randomly discards particular neighbor nodes during the aggregation of hash codes. For the l -th propagation layer, we randomly dropped $p_1 \times (N + M)$ nodes of the adjacency matrix, where p_1 is the dropout ratio. We also conducted bit dropout that randomly drops a few bits of the final user and item hash codes before the prediction, with a probability p_2 .

Note that dropout is only used during the training process, and should be disabled in the testing process. Figure 4 plots the effect of node dropout ratio p_1 and bit dropout ratio p_2 under HR@50 evaluation metric on different datasets. Regarding the two dropout strategies, bit dropout delivers much better performance especially on sparse datasets. Node dropout even has a negative effect on Yelp dataset. One reason might be that dropping neighbor nodes damages the hash code aggregation operation in the propagation process of HS-GCN, especially when the neighbor nodes are sparse. Bit dropout is more effective, which means that it can be an effective strategy to address the overfitting of learning to hash on GCNs.

6.2.4 Effect of the Length of Hash Codes

In order to explore how the length of hash code affects HS-GCN and other baselines, we changed the length of their hash codes within $\{8, 16, 32, 64\}$, and repeated the experiments. Figure 5 shows the performance of these models with different length K on MovieLens and Gowalla datasets. It can be observed that the hash code length has positive influence on hashing models, while the influence weakens

as length K increases. Similar trends are observed on other datasets.

7 CONCLUSION AND FUTURE WORK

In this work, we explicitly incorporate the high-order Hamming similarity into the hash code learning. Particularly, we devise a novel framework HS-GCN, which yields binary hash codes by propagating the similarity information on the user-item graph structure. HS-GCN is the first proposed graph convolutional network in the Hamming space, where the propagation layers are built upon Hamming space operations to directly capture the high-order Hamming similarity. Extensive experiments on three real-world datasets demonstrate the rationality and effectiveness of capturing the high-order similarity for learning to hash.

In future, we plan to strengthen HS-GCN by incorporating multi-form knowledge [69], like attributes in tables [70], and celebrities in matrices [71], [72]. Moreover, we are interested in building recommender systems for micro videos [73]. Considering the micro videos constantly emerge in large numbers, they thus require more efficient and accurate recommendation. Another emerging research direction is to explore the interpretable hashing-based recommendation [74].

ACKNOWLEDGMENTS

This work is supported by the National Natural Science Foundation of China, No.:U1936203

REFERENCES

- [1] S. Rendle, C. Freudenthaler, Z. Gantner, and L. Schmidt-Thieme, "Bpr: Bayesian personalized ranking from implicit feedback," in *UAI*, 2012, pp. 452–461.
- [2] X. Wang, X. He, M. Wang, F. Feng, and T.-S. Chua, "Neural graph collaborative filtering," in *SIGIR*, 2019, pp. 165–174.
- [3] Y. Wei, X. Wang, L. Nie, X. He, R. Hong, and T.-S. Chua, "Mmgcn: Multi-modal graph convolution network for personalized recommendation of micro-video," in *MM*, 2019, pp. 1437–1445.
- [4] Z. Li, J. Tang, and T. Mei, "Deep collaborative embedding for social image understanding," *TPAMI*, vol. 41, no. 9, pp. 2070–2083, 2018.
- [5] H. Liu, X. He, F. Feng, L. Nie, R. Liu, and H. Zhang, "Discrete factorization machines for fast feature-based recommendation," in *IJCAI*, 2018, pp. 3449–3455.
- [6] H. Zhang, F. Shen, W. Liu, X. He, H. Luan, and T.-S. Chua, "Discrete collaborative filtering," in *SIGIR*, 2016, pp. 325–334.
- [7] Y. Zhang, H. Wang, D. Lian, I. W. Tsang, H. Yin, and G. Yang, "Discrete ranking-based matrix factorization with self-paced learning," in *KDD*, 2018, pp. 2758–2767.
- [8] J. Wang, S. Kumar, and S.-F. Chang, "Semi-supervised hashing for large-scale search," *TPAMI*, vol. 34, no. 12, pp. 2393–2406, 2012.
- [9] Y. Weiss, A. Torralba, and R. Fergus, "Spectral hashing," in *NIPS*, 2009, pp. 1753–1760.
- [10] W.-C. Kang and J. McAuley, "Candidate generation with binary codes for large-scale top-n recommendation," in *CIKM*, 2019, pp. 1523–1532.
- [11] D. Lian, R. Liu, Y. Ge, K. Zheng, X. Xie, and L. Cao, "Discrete content-aware matrix factorization," in *KDD*, 2017, pp. 325–334.
- [12] T. Chen, H. Yin, Y. Zheng, Z. Huang, Y. Wang, and M. Wang, "Learning elastic embeddings for customizing on-device recommenders," in *KDD*, 2021, pp. 138–147.
- [13] Q. Wang, H. Yin, T. Chen, Z. Huang, H. Wang, Y. Zhao, and N. Q. Viet Hung, "Next point-of-interest recommendation on resource-constrained mobile devices," in *WWW*, 2020, pp. 906–916.
- [14] Y. Zhang, H. Yin, Z. Huang, X. Du, G. Yang, and D. Lian, "Discrete deep learning for fast content-aware recommendation," in *WSDM*, 2018, pp. 717–726.

- [15] H. Zhu, M. Long, J. Wang, and Y. Cao, "Deep hashing network for efficient similarity retrieval," in *AAAI*, 2016, pp. 2415–2421.
- [16] H. Liu, R. Wang, S. Shan, and X. Chen, "Deep supervised hashing for fast image retrieval," in *CVPR*, 2016, pp. 2064–2072.
- [17] Z. Cao, M. Long, J. Wang, and P. S. Yu, "Hashnet: Deep learning to hash by continuation," in *ICCV*, 2017, pp. 5608–5617.
- [18] Q. Tan, N. Liu, X. Zhao, H. Yang, J. Zhou, and X. Hu, "Learning to hash with graph neural networks for recommender systems," in *WWW*, 2020, pp. 1988–1998.
- [19] M. Gori, A. Pucci, V. Roma, and I. Siena, "Itemrank: A random-walk based scoring algorithm for recommender engines," in *IJCAI*, 2007, pp. 2766–2771.
- [20] X. He, M. Gao, M.-Y. Kan, and D. Wang, "BiranK: Towards ranking on bipartite graphs," *TKDE*, vol. 29, no. 1, pp. 57–71, 2016.
- [21] Y. Wei, X. Wang, X. He, L. Nie, Y. Rui, and T.-S. Chua, "Hierarchical user intent graph network for multimedia recommendation," *TMM*, 2021.
- [22] Y. Wei, X. Wang, L. Nie, X. He, and T.-S. Chua, "Graph-refined convolutional network for multimedia recommendation with implicit feedback," in *MM*, 2020, pp. 3541–3549.
- [23] F. Liu, Z. Cheng, L. Zhu, Z. Gao, and L. Nie, "Interest-aware message-passing GCN for recommendation," in *WWW*, 2021, pp. 1296–1305.
- [24] Y. Wu, H. Liu, and Y. Yang, "Graph convolutional matrix completion for bipartite edge prediction," in *IC3K*, 2018, pp. 49–58.
- [25] R. Ying, R. He, K. Chen, P. Eksombatchai, W. L. Hamilton, and J. Leskovec, "Graph convolutional neural networks for web-scale recommender systems," in *KDD*, 2018, pp. 974–983.
- [26] X. He, K. Deng, X. Wang, Y. Li, Y. Zhang, and M. Wang, "LightGCN: simplifying and powering graph convolution network for recommendation," in *SIGIR*, 2020, pp. 639–648.
- [27] C. Shi, B. Hu, W. X. Zhao, and S. Y. Philip, "Heterogeneous information network embedding for recommendation," *TKDE*, vol. 31, no. 2, pp. 357–370, 2018.
- [28] B. Hu, C. Shi, W. X. Zhao, and P. S. Yu, "Leveraging meta-path based context for top-n recommendation with a neural co-attention model," in *KDD*, 2018, pp. 1531–1540.
- [29] S. Fan, J. Zhu, X. Han, C. Shi, L. Hu, B. Ma, and Y. Li, "Metapath-guided heterogeneous graph neural network for intent recommendation," in *KDD*, 2019, pp. 2478–2486.
- [30] Y. Yang, Z. Guan, J. Li, W. Zhao, J. Cui, and Q. Wang, "Interpretable and efficient heterogeneous graph convolutional network," *TKDE*, 2021.
- [31] Z. Li, J. Tang, L. Zhang, and J. Yang, "Weakly-supervised semantic guided hashing for social image retrieval." *IJCV*, vol. 128, no. 8/9, pp. 2265–2278, 2020.
- [32] H. Jegou, M. Douze, and C. Schmid, "Product quantization for nearest neighbor search," *TPAMI*, vol. 33, no. 1, pp. 117–128, 2010.
- [33] W. Liu, J. Wang, S. Kumar, and S.-F. Chang, "Hashing with graphs," in *ICML*, 2011, pp. 1–8.
- [34] B. Kulis and T. Darrell, "Learning to hash with binary reconstructive embeddings," in *NIPS*, 2009, pp. 1042–1050.
- [35] W. Liu, J. Wang, R. Ji, Y.-G. Jiang, and S.-F. Chang, "Supervised hashing with kernels," in *CVPR*, 2012, pp. 2074–2081.
- [36] M. Norouzi and D. J. Fleet, "Minimal loss hashing for compact binary codes," in *ICML*, 2011, pp. 353–360.
- [37] F. Shen, C. Shen, W. Liu, and H. Tao Shen, "Supervised discrete hashing," in *CVPR*, 2015, pp. 37–45.
- [38] Y. Gong, S. Lazebnik, A. Gordo, and F. Perronnin, "Iterative quantization: A procrustean approach to learning binary codes for large-scale image retrieval," *TPAMI*, vol. 35, no. 12, pp. 2916–2929, 2012.
- [39] X. He and T.-S. Chua, "Neural factorization machines for sparse predictive analytics," in *SIGIR*, 2017, pp. 355–364.
- [40] R. Xia, Y. Pan, H. Lai, C. Liu, and S. Yan, "Supervised hashing for image retrieval via image representation learning," in *AAAI*, 2014, pp. 2156–2162.
- [41] H. Lai, Y. Pan, Y. Liu, and S. Yan, "Simultaneous feature learning and hash coding with deep neural networks," in *CVPR*, 2015, pp. 3270–3278.
- [42] H. Wang, D. Lian, and Y. Ge, "Binarized collaborative filtering with distilling graph convolutional networks," *arXiv preprint arXiv:1906.01829*, 2019.
- [43] X. Zhou, F. Shen, L. Liu, W. Liu, L. Nie, Y. Yang, and H. T. Shen, "Graph convolutional network hashing," *TC*, vol. 50, no. 4, pp. 1460–1472, 2018.
- [44] Y. Bengio, N. Léonard, and A. Courville, "Estimating or propagating gradients through stochastic neurons for conditional computation," *arXiv preprint arXiv:1308.3432*, 2013.
- [45] Y. Zhou, X. Bai, W. Liu, and L. J. Latecki, "Fusion with diffusion for robust visual tracking," in *NIPS*, 2012, pp. 2978–2986.
- [46] F. Yang, R. Hinami, Y. Matsui, S. Ly, and S. Satoh, "Efficient image retrieval via decoupling diffusion into online and offline processing," in *AAAI*, 2019, pp. 9087–9094.
- [47] B. Wang, A. M. Mezlini, F. Demir, M. Fiume, Z. Tu, M. Brudno, B. Haibe-Kains, and A. Goldenberg, "Similarity network fusion for aggregating data types on a genomic scale," *Nature Methods*, vol. 11, no. 3, pp. 333–337, 2014.
- [48] H. Liu, R. Ji, Y. Wu, F. Huang, and B. Zhang, "Cross-modality binary code learning via fusion similarity hashing," in *CVPR*, 2017, pp. 7380–7388.
- [49] S. Bai, Z. Zhou, J. Wang, X. Bai, L. Jan Latecki, and Q. Tian, "Ensemble diffusion for retrieval," in *ICCV*, 2017, pp. 774–783.
- [50] Y. Wang, Y. Wang, J. Yang, and Z. Lin, "Dissecting the diffusion process in linear graph convolutional networks," *arXiv preprint arXiv:2102.10739*, 2021.
- [51] B. Jiang, D. Lin, J. Tang, and B. Luo, "Data representation and learning with graph diffusion-embedding networks," in *CVPR*, 2019, pp. 10 414–10 423.
- [52] Y. Feng, H. You, Z. Zhang, R. Ji, and Y. Gao, "Hypergraph neural networks," in *AAAI*, 2019, pp. 3558–3565.
- [53] Y. Wei, X. Wang, Q. Li, L. Nie, Y. Li, X. Li, and T.-S. Chua, "Contrastive learning for cold-start recommendation," in *MM*, 2021, pp. 5382–5390.
- [54] X. Yu, T. Gan, Y. Wei, Z. Cheng, and L. Nie, "Personalized item recommendation for second-hand trading platform," in *MM*, 2020, pp. 3478–3486.
- [55] Z. Zhang, "Improved adam optimizer for deep neural networks," in *IWQoS*, 2018, pp. 1–2.
- [56] X. He, H. Zhang, M.-Y. Kan, and T.-S. Chua, "Fast matrix factorization for online recommendation with implicit feedback," in *SIGIR*, 2016, pp. 549–558.
- [57] X. Wang, X. He, L. Nie, and T.-S. Chua, "Item silk road: Recommending items from information domains to social users," in *SIGIR*, 2017, pp. 185–194.
- [58] D. Liang, L. Charlin, J. McInerney, and D. M. Blei, "Modeling user exposure in recommendation," in *WWW*, 2016, pp. 951–961.
- [59] X. Geng, H. Zhang, J. Bian, and T.-S. Chua, "Learning image and user features for recommendation in social networks," in *ICCV*, 2015, pp. 4274–4282.
- [60] X. He, L. Liao, H. Zhang, L. Nie, X. Hu, and T.-S. Chua, "Neural collaborative filtering," in *WWW*, 2017, pp. 173–182.
- [61] X. Glorot and Y. Bengio, "Understanding the difficulty of training deep feedforward neural networks," in *AISTATS*, 2010, pp. 249–256.
- [62] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Communications of the ACM*, vol. 60, no. 6, pp. 84–90, 2017.
- [63] A. Mnih and R. R. Salakhutdinov, "Probabilistic matrix factorization," in *NIPS*, 2008, pp. 1257–1264.
- [64] M. Gao, L. Chen, X. He, and A. Zhou, "Bine: Bipartite network embedding," in *SIGIR*, 2018, pp. 715–724.
- [65] W. Hamilton, Z. Ying, and J. Leskovec, "Inductive representation learning on large graphs," in *NIPS*, 2017, pp. 1024–1034.
- [66] M. Norouzi, A. Punjani, and D. J. Fleet, "Fast search in hamming space with multi-index hashing," in *CVPR*, 2012, pp. 3108–3115.
- [67] J. Liu, Z. Li, J. Tang, Y. Jiang, and H. Lu, "Personalized geo-specific tag recommendation for photos on social websites," *TMM*, vol. 16, no. 3, pp. 588–600, 2014.
- [68] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, "Dropout: a simple way to prevent neural networks from overfitting," *JMLR*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [69] F. Liu, Z. Cheng, C. Sun, Y. Wang, L. Nie, and M. S. Kankanhalli, "User diverse preference modeling by multimodal attentive metric learning," in *MM*, 2019, pp. 1526–1534.
- [70] X. Wang, X. He, Y. Cao, M. Liu, and T.-S. Chua, "Kgat: Knowledge graph attention network for recommendation," in *KDD*, 2019, pp. 950–958.
- [71] X. Song, F. Feng, X. Han, X. Yang, W. Liu, and L. Nie, "Neural compatibility modeling with attentive knowledge distillation," in *SIGIR*, 2018, pp. 5–14.

- [72] C. Sun, X. Song, F. Feng, W. X. Zhao, H. Zhang, and L. Nie, "Supervised hierarchical cross-modal hashing," in *SIGIR*, 2019, pp. 725–734.
- [73] Y. Wei, X. Wang, W. Guan, L. Nie, Z. Lin, and B. Chen, "Neural multimodal cooperative learning toward micro-video understanding," *TIP*, vol. 29, pp. 1–14, 2019.
- [74] Y. Zhang, Y. Zhang, M. Zhang, and C. Shah, "Ears 2019: the 2nd international workshop on explainable recommendation and search," in *SIGIR*, 2019, pp. 1438–1440.



Han Liu received his B.E. and MA. from Shandong University (SDU). Currently, he is working towards the Ph.D. degree in the School of Computer Science and Technology, SDU. His main research interests include information retrieval and recommendation. He has published papers in referred conference proceedings, such as IJCAI'2018 and WSDM'2020.



Yinwei Wei received his MS degree from Tianjin University and Ph.D. degree from Shandong University, respectively. Currently, he is a research fellow with NEXt++, National University of Singapore. His research interests include multimedia computing and recommendation. Several works have been published in top forums, such as MM and TIP. Dr. Wei has served as the PC member for several conferences, such as MM, AAAI, and IJCAI, and the reviewer for TMM, TKDE, and TIP.



Jianhua Yin received the Ph.D. degree in computer science and technology from Tsinghua University, Beijing, China, in 2017. He is currently an Associate Professor with the School of Computer Science and Technology, Shandong University, Jinan, China. He has published several papers in the top venues, such as ACM TOIS, IEEE TKDE, ACM MM, ACM SIGKDD, ACM SIGIR, and IEEE ICDE. His research interests mainly include data mining and machine learning applications.



Liqiang Nie received the B.Eng. degree from Xi'an Jiaotong University and the Ph.D. degree from the National University of Singapore (NUS). After Ph.D., he continued his research in NUS as a Research Fellow for more than three years. He is currently a professor with the School of Computer Science and Technology, Shandong University. Meanwhile, he is the Adjunct Dean with the Shandong AI Institute. His research interests lie primarily in multimedia computing and information retrieval. He has published around

100 papers in the top conferences or journals, with more than 10,000 Google Scholar citations. He was granted several awards like SIGIR and ACM MM best paper nomination award, ACM SIGMM Rising Star, DAMO Academy Young Fellow, AI 2000 the most influential scholar in artificial intelligence and MIT TR35 China. He is an AE of IEEE TKDE, TMM, ACM ToMM and Information Science. He serves as a PC Chair of ICIMCS 2017, PCM 2018, ChinaMM 2020 and an Area Chair of ACM MM 2018-2021.