Bipartite Graph Convolutional Hashing for Effective and Efficient Top-N Search in Hamming Space

Yankai Chen[1], Yixiang Fang[2], Yifei Zhang[1], Irwin King[1]

{ykchen, yfzhang, king}@cse.cuhk.edu.hk   fangyixiang@cuhk.edu.cn

[1]The Chinese University of Hong Kong   [2]The Chinese University of Hong Kong, Shenzhen

Abstract

Searching on bipartite graphs is basal and versatile to many real-world Web applications, e.g., online recommendation, database retrieval, and query-document searching. Given a query node, the conventional approaches rely on the similarity matching with the vectorized node embeddings in the continuous Euclidean space. To efficiently manage intensive similarity computation, developing hashing techniques for graph-structured data has recently become an emerging research direction. Despite the retrieval efficiency in Hamming space, prior work is however confronted with *catastrophic performance decay*. In this work, we investigate the problem of hashing with Graph Convolutional Network on bipartite graphs for effective Top-N search. We propose an end-to-end *Bipartite Graph Convolutional Hashing* approach, namely BGCH, which consists of three novel and effective modules: (1) *adaptive graph convolutional hashing*, (2) *latent feature dispersion*, and (3) *Fourier serialized gradient estimation*. Specifically, the former two modules achieve the substantial retention of the structural information against the inevitable information loss in hash encoding; the last module develops Fourier Series decomposition to the hashing function in the frequency domain mainly for more accurate gradient estimation. The extensive experiments on six real-world datasets not only show the performance superiority over the competing hashing-based counterparts, but also demonstrate the effectiveness of all proposed model components contained therein.

CCS Concepts

• **Computing methodologies** → **Learning latent representations**; • **Information systems** → **Information retrieval**.

Keywords

Representation Learning; Learning to Hash; Graph Convolutional Network; Bipartite Graph; Hamming Space Search

(a) Bipartite graph modeling

(b) Effectiveness (NDCG@1000) versus matching computation latency (log-scale)
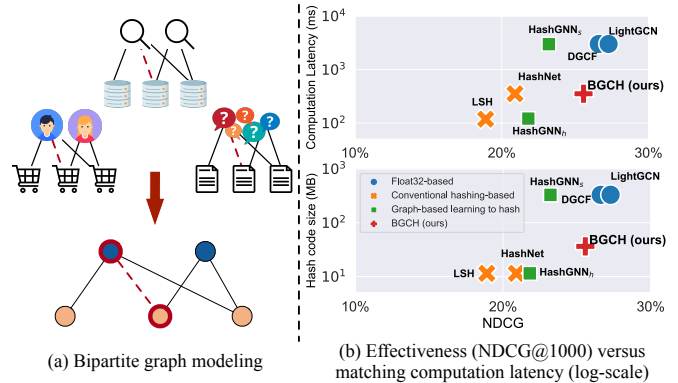
Figure 1: Illustration of bipartite graph modeling and overall model performance visualization on Dianping dataset.

1 Introduction

Bipartite graphs are ubiquitous in the real world for the ease of modeling various Web applications, e.g., as shown in Figure 1(a), user-product recommendation [41, 63] and online query-document matching [66]. A fundamental task, *Top-N search*, is to filter out N best-matched graph nodes for a query node, e.g., recommending Top-N attractive products to a target user in the user-product graph. With the development of the recent machine learning research, learning vectorized representations (*a.k.a.* embeddings) nowadays has become one of the standard procedures for similarity matching [10, 16]. Among existing techniques, graph-based neural methods, i.e., *Graph Convolutional Networks* (GCNs), have recently present remarkable model performance [17, 19]. Due to the ability to capture high-order connection information, GCN models can thus produce semantic enrichment to the node embeddings. Based on the learned embeddings, similarity estimation is then exhaustively proceeded in the continuous Euclidean space.

Apart from embedding informativeness, *computation latency* and *embedding memory overhead* are two important criteria for realistic application deployment. With the explosive data growth, *learning to hash* [25, 51] recently provides an alternative option to graph-based models for optimizing the model scalability. Generally, it learns to convert the vectorized list of continuous values into the finite binarized hash codes. In lieu of using *full-precision*[1] embeddings, the learned hash codes have the promising potential to achieve, not only the space reduction, but also the computation acceleration for Top-N object matching and retrieval in the Hamming space.

Despite the promising advantages of bridging GCNs and learning to hash, simply stacking these two techniques is trivial and thus falls short of performance satisfaction with several inadequacies:

---

[1]The term "full-precision" generally refers to single-precision and double-precision. And we use float32 by default throughout this work for illustration.

- **Coarse-grained similarity measurement.** Compared to continuous embeddings, hash codes with the same vector dimension are naturally *less expressive* with finite encoding permutation in Hamming space (e.g., 2^d if the dimension is d). Consequently, this leads to a coarse-grained estimation of the pairwise node similarity, thus drawing a conspicuous performance decay with inaccurate Top-N matching.
- **Feature erosion issue.** Recent work [38, 44, 45, 49] usually adopts $\text{sign}(\cdot)$ function for $O(1)$ complexity encoding. However, hashing via $\text{sign}(\cdot)$ will inevitably smooth the embedding feature informativeness, via converting each digit of continuous embeddings into the hamming space, no matter what specific value it used to be. Thus the latent features in these learned hash embeddings become less informative, and existing models lack certain mechanisms to hedge the feature erosion in hashing.
- **Intractable model optimization.** Since $\text{sign}(\cdot)$ is not differentiable at 0 and its derivatives are 0 anywhere else, previous models usually use *visually similar* but not necessarily *theoretically relevant* functions, e.g., $\tanh(\cdot)$, for gradient estimation. This may lead to inconsistent optimization directions in model training. Moreover, because of the embedding discreteness, the associated loss landscape[2] (Figure 2(a)) are steep and bumping [1], which further increases the difficulty in optimization.

In this paper, we study the problem of learning to hash with Graph Convolutional Network (GCN) on bipartite graphs for effective Top-N Hamming space search. We propose a model namely *Bipartite Graph Convolutional Hashing* (BGCH), with three effective modules: (1) *adaptive graph convolutional hashing*, (2) *latent feature dispersion*, and (3) *Fourier serialized gradient estimation*. While the former two modules significantly enrich the informativeness and expressivity to the learned hash codes, the last one provides an accordant and tractable optimization flow in forward and backward propagation of model optimization. Concretely:

- **Adaptive graph convolutional hashing.** Our first module designs a *topology-aware* convolutional hashing that employs the layer-wise hash encoding (from low- to high-order sub-structures of bipartite graphs) to consecutively binarize the node features with different semantics. To boost the expressivity, the convolutional hashing is equipped an effective approximation technique for *embedding rescaling*, which does not undermine the efficiency of Hamming distance computation. Intuitively, these two designs make the learned hash codes more informative and expressive for preserving fine-grained similarity in the Hamming space.
- **Latent feature dispersion**. Our second module, i.e., feature dispersion, aims to hedge the inevitable information loss from the numerical binarization. In conventional continuous embeddings, major features however condense in a small region of embedding structures. Since these vectorized latent features tend to be inevitably smoothed by the hashing discreteness, it is natural to preserve information as much as possible by spreading out those decisive features to the majority of embedding dimensions, instead of the very few of them. To achieve this, our proposed module aims to explicitly disperse informative latent node features, which can be further diffused to each convolutional layer when exploring the bipartite graph.

- **Fourier serialized gradient estimation.** Furthermore, to provide accurate gradient estimation, BGCH proposes to decompose $\text{sign}(\cdot)$ function with Fourier Series in the frequency domain. Compared to existing gradient estimators [11, 15, 39, 44, 59], this estimator better follows the *main direction* of factual gradients to enable an accordant and tractable model optimization in forward and backward propagation. With the limited number of decomposition terms, BGCH can well provide more accurate gradient estimation to $\text{sign}(\cdot)$ within the acceptable training cost.

Based on the learned hash codes, BGCH maintains moderate resource consumption whilst providing substantial performance improvement in Top-N Hamming space retrieval. The quality-cost trade-off is summarized in Figure 1(b), which compares BGCH against a list of representative counterparts (*including float32-based and hashing-based*) on a real-world bipartite graph with over 10 million observed edges (experimental details are reported in § 5.1). As the lower-right corner of Figure 1(b) indicates the ideal optimal performance, BGCH can deliver over 8× computation acceleration and space reduction relative to existing full-precision models, while being more effective than each hash-based method (§ 5.2 and § 5.3). To summarize, our main contributions are organized as follows:

- We study the problem of learning to hash with Graph Convolutional Network on bipartite graphs. We propose a novel approach BGCH with three effective modules for effective and efficient Top-N search in Hamming space (§ 4).
- We conduct extensive experiments on six real-world datasets to evaluate the retrieval quality. In-depth analyses are also provided towards the necessity of all proposed model components from both technical and empirical perspectives (§ 5).
- We theoretically prove the model effectiveness and provide complexity analyses in terms of time and space costs (Appendix C).

2 Related Work

Graph convolution network (GCN). Early work studies the graph convolutions mainly on the *spectral domain*, such as Laplacian eigen-decomposition [3] and Chebyshev polynomials [12]. One major issue is that these models are usually computationally expensive. To tackle this problem, *spatial-based* GCN models are proposed to re-define the graph convolution operations by aggregating the embeddings of neighbors to refine and update the target node's embedding. Due to its scalability to large graphs, spatial-based GCN models are widely used in various applications [17, 19, 52]. For example, to capture high-order structural information, NGCF [52] and LightGCN [19] learn the collaborative filtering signals on bipartite interaction graphs for recommendation. Despite the effectiveness in embedding latent features for graph nodes, they usually suffer from inference inefficiency due to the high computational cost of similarity calculation between continuous embeddings [49]. To address this issue, *learning to hash* provides the feasibility.

Learning to hash. Learning to hash models are promising to achieve computation acceleration and storage reduction for general information retrieval and processing tasks [5, 13, 19, 22, 23, 36, 40]. More than reducing conflicts [33], similarity-preserving hashing maps high-dimensional dense vectors to a low-dimensional Hamming space for efficiently processing downstream tasks. A representative model is Locality Sensitive Hashing (LSH) [14] that uses

---

[2]Details about the visualization construction are attached in Appendix A.

random projections as the hash functions. Recent work focuses on integrating the deep neural network architectures for model improvement [51]. They inspire a series of follow-up work for various tasks, such as fast retrieval of images [4, 38, 44], documents [9, 35], categorical information [27], e-commerce products [6, 65].

To leverage hashing techniques with GCNs, the state-of-the-art work HashGNN [49] investigates learning to hash for online matching and recommendation. Specifically, HashGNN consectively combines the GraphSage [17] as the embedding encoder and learning to hash method to get the corresponding binary encodings afterwards. Its hash encoding process only proceeds at the end of multi-layer graph convolutions, i.e., using the aggregated output of Graph-Sage for representation binarization. However, this fails to capture intermediate semantics from nodes' different layers of receptive fields [30]. The other issue of HashGNN is using *Straight-Through Estimator (STE)* [2] to assume all gradients of sign($\cdot$) as 1 in back-propagation. However, the integral of 1 is a certain linear function other than the sign($\cdot$), whereas this may lead to inconsistent optimization directions in the model training. To address these issues, our model BGCH is proposed with effectiveness justification in § 5.

3 Preliminaries and Problem Formulation

Graph Convolution Network (GCN). The general idea of GCN is to learn node embeddings by *iteratively propagating and aggregating* latent features of node neighbors via the graph topology [19, 30, 56]:

$$V_x^{(l)} = AGG\left(V_x^{(l-1)}, \{V_z^{(l-1)} : z \in \mathcal{N}(x)\}\right), \tag{1}$$

where $V_x^{(l)} \in \mathbb{R}^d$ denotes node x's embedding after l-th iteration of graph convolutions, indexed in the embedding matrix V. $\mathcal{N}(x)$ is the set of x's neighbors. Function $AGG(\cdot, \cdot)$ is the information aggregation function, with several implementations in previous work [17, 30, 50, 57], mainly aiming to transform the center node feature and the neighbor features. In this work, we adopt the graph convolution paradigm from the state-of-the-art model LightGCN [19].

Bipartite Graph and Adjacency Matrix. The bipartite graph is denoted as $\mathcal{G} = \{\mathcal{V}_1, \mathcal{V}_2, \mathcal{E}\}$, where $\mathcal{V}_1$ and $\mathcal{V}_2$ are two *disjoint* node sets and $\mathcal{E}$ is the set of edges between nodes in $\mathcal{V}_1$ and $\mathcal{V}_2$. We can use $Y \in \mathbb{R}^{|\mathcal{V}_1| \times |\mathcal{V}_2|}$ to indicate the edge transactions, where 1-valued entries, i.e., $Y_{x,y} = 1$, indicate there is an observed edge between nodes $x \in \mathcal{V}_1$ and $y \in \mathcal{V}_2$, otherwise $Y_{x,y} = 0$. Then the adjacency matrix A of the whole graph can be defined as:

$$A = \begin{bmatrix} 0 & Y \\ Y^T & 0 \end{bmatrix}. \tag{2}$$

Problem Formulation. Give a bipartite graph $\mathcal{G} = \{\mathcal{V}_1, \mathcal{V}_2, \mathcal{E}\}$ and its adjacency matrix A, we devote to learn a hashing function:

$$F(A|\Theta) \rightarrow Q, \tag{3}$$

where Θ is the set of all learnable parameters. Given two nodes in the bipartite graph, e.g., $x \in \mathcal{V}_1$ and $y \in \mathcal{V}_2$, their hash codes are Q_x and Q_y. Then the probability of edge existence $\widehat{Y}_{x,y}$ between nodes $x \in \mathcal{V}_1$ and $y \in \mathcal{V}_2$ can be effectively and efficiently measured by the hash codes Q_x and Q_y, i.e., $\widehat{Y}_{x,y} = f(Q_x, Q_y)$ where f is a score function. Intuitively, the larger value $\widehat{Y}_{x,y}$ is, the more likely x and y are matched, i.e., an edge between x and y exists. Explanations of key notations used in this paper are attached in Appendix B.

4 BGCH: Methodology

4.1 Overview

We formally introduce our BGCH model. Notice that since the effect of feature dispersion module propagates along with convolutional hashing, we then introduce these modules in the following order: (1) *latent feature dispersion* (§ 4.2) aims to disperse the embedded features into wider embedding structures to hedge the inevitable information loss in hashing; (2) *adaptive graph convolutional hashing* (§ 4.3) provides an effective encoding approach to significantly improve the hashed feature expressivity whilst maintaining the matching efficiency in the hamming space; (3) *Fourier serialized gradient estimation* (§ 4.4) introduces the Fourier Series decomposition for sign($\cdot$) in the frequency domain to provide more accurate gradient approximation. Based on the learned hash codes, BGCH develops efficient online matching with the Hamming distance measurement (§ 4.5). Our model illustration is attached in Figure 2(b).

4.2 Latent Feature Dispersion

To tackle the feature erosion issue, we seek to disperse the embedded features as one effective strategy to hedge the inevitable information loss caused by numerical binarization. From the perspective of singular value decomposition (SVD), singular values and corresponding singular vectors reconstruct the original matrix; normally, large singular values can be interpreted to associate with *major feature structures* of the matrix [54]. Since we want to avoid condensing and gathering informative features in (relatively small) embedding sub-structures, it is natural to bridge the target by working on these singular values. Hence, based on this intuition, we aim to *normalize singular values for equalizing their respective contributions in constituting latent features*. To achieve this, Power Normalization [31, 68] is one of the solutions that tackle related problems such as feature imbalance [32]. Inspired by the recent approximation attempt [61], we now introduce a lightweight feature dispersion technique in graph convolution as follows.

Concretely, let I denote the identity matrix, we start from generating a *standard normal random vector* $p^{(0)} \sim \mathcal{N}(0, I)$ where $p^{(0)} \in \mathbb{R}^c$. Based on the embedding matrix to conduct feature dispersion, e.g., let $V = V^{(0)}$, we compute the desired **dispersing vector** $p^{(k)}$ by iteratively performing $p^{(k)} = V^\top V p^{(k-1)}$. The iteration for generating dispersing vectors is independent of the graph convolution iterations[3]. We have the projection matrix P of $p^{(K)}$ via:

$$P = \frac{p^{(K)} p^{(K)\top}}{||p^{(K)}||_2^2}. \tag{4}$$

Then we have the feature-dispersed representation matrix with the hyper-parameter $\epsilon \in (0, 1)$ as follows:

$$\widetilde{V} = V(I - \epsilon P). \tag{5}$$

Consequently, integrating the dispersed matrix $\widetilde{V}$, we have the **feature-dispersed** graph convolution as:

$$\widetilde{V}^{(l+1)} = (D^{-\frac{1}{2}} A D^{-\frac{1}{2}}) \widetilde{V}^{(l)}, \text{ where } \widetilde{V}^{(0)} = V^{(0)}(I - \epsilon P). \tag{6}$$

Note that we explicitly conduct this feature dispersion operation one time only at the initial step, i.e., $\widetilde{V}^{(0)}$, and, more importantly, such feature dispersion can be diffused via the multi-layer graph

---

[3]In our work, we set $K \leq L$ mainly to enable the associated complexity of dispersing vector generation is upper bounded by the graph convolution complexity.
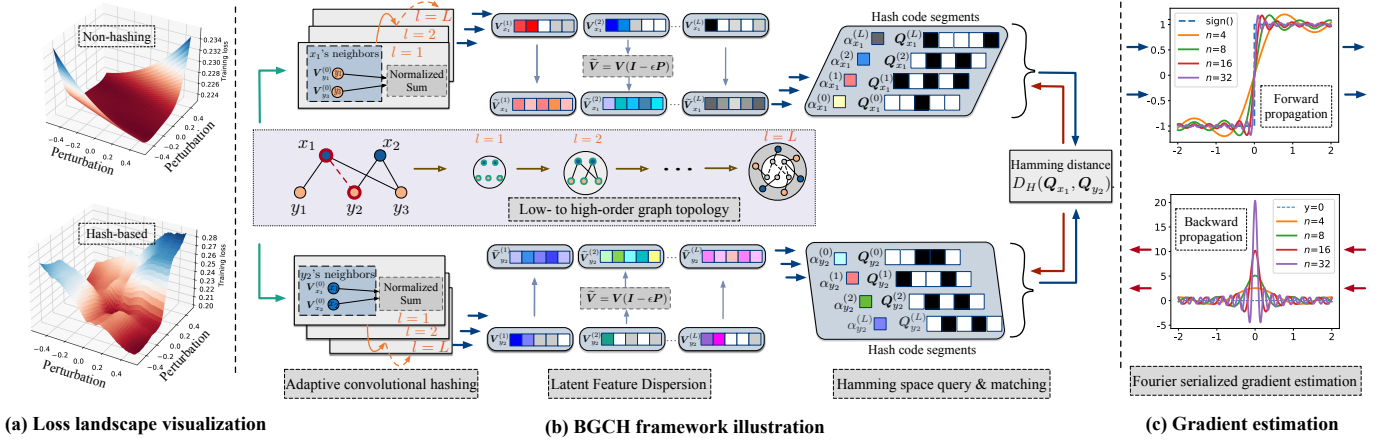
Figure 2: (a) Visualized loss landscape comparison; (b) BGCH model framework (best view in color); (c) Fourier Serialized gradient estimation in forward and bachward propagation.

convolutions from 0 to L. Compare to the unprocessed embedding counterpart, e.g., $V^{(l)}$, embedding matrix $\widetilde{V}^{(l)}$ at each layer presents a dispersed feature structure with a *more balanced distribution of singular values in expection*. We formally explain this as follows:

Theorem 1 (Feature Dispersion). *Let $V^{(l)} = U_1 \Sigma U_2^{\mathsf{T}}$, where U_1 and U_2 are unitary matrices and descending singular value matrix $\Sigma = \text{diag}(\sigma_1, \sigma_2, \cdots, \sigma_c)$. Then $\mathbb{E}(\widetilde{V}^{(l)}) = U_1 \Sigma \Sigma_\mu U_2^{\mathsf{T}}$ where $\Sigma_\mu = \text{diag}(\mu_1, \mu_2, \cdots, \mu_c)_{0 < \mu_{1 \cdots c} < 1}$ is in ascending order.*

Intuitively, given the same orthonormal bases, compared to $V^{(l)}$, it is harder in expectation to reconstruct $\widetilde{v}^{(l)}$ with informative features being dispersed out in larger matrix sub-structures. This eventually provides the functionality to hedge the information loss in numerical binarization. We attach the theorem proof in Appendix C and evaluate the module effectiveness later in § 5.4.

4.3 Adaptive Graph Convolutional Hashing

One feasible solution for increasing expressivity and smoothing loss landscapes is to include the *relaxation strategy*. Hence, apart from the topology-aware embedding binarization with $\text{sign}(\cdot)$:

$$Q_x^{(l)} = \text{sign}(\widetilde{V}_x^{(l)}), \tag{7}$$

our model BGCH additionally computes a layer-wise positive rescaling factor for each node, e.g., $\alpha_x^{(l)} \in \mathbb{R}^+$, such that $\widetilde{V}_x^{(l)} \approx \alpha_x^{(l)} Q_x^{(l)}$. In this work, we introduce a simple but effective approach to directly calculate the rescaling factors as follows:

$$\alpha_x^{(l)} = \frac{1}{d} ||\widetilde{V}_x^{(l)}||_1. \tag{8}$$

Instead of setting these factors as learnable, such deterministic computation substantially prunes the parameter search space while attaining the adaptive approximation functionality for different graph nodes. We demonstrate this in § 5.4 of experiments.

After L iterations of feature propagation and hashing, we obtain the table of **adaptive hash codes** $Q = \{\alpha, Q\}$, where $\alpha \in \mathbb{R}^{(|\mathcal{V}_1|+|\mathcal{V}_2|) \times (L+1)}$ and $Q \in \mathbb{R}^{(|\mathcal{V}_1|+|\mathcal{V}_2|) \times d}$. For each node x, its corresponding hash codes are indexed and assembled:

$$\boldsymbol{\alpha}_x = \alpha_x^{(0)} || \alpha_x^{(1)} || \cdots || \alpha_x^{(L)}, \text{ and } Q_x = Q_x^{(0)} || Q_x^{(1)} || \cdots || Q_x^{(L)}. \tag{9}$$

Intuitively, the hash code table Q represents the bipartite structural information that is propagated back and forth at different iteration steps l, i.e., from 0 to the maximum step L. It not only tracks the intermediate knowledge hashed for all graph nodes, but also maintains the value approximation to their original continuous embeddings, e.g., $\widetilde{v}_x^{(l)}$. In addition, with the slightly more space cost (complexity analysis in Appendix C, such detached hash encoding approach still supports the bitwise operations (§ 4.5) for accelerating inference and matching.

4.4 Fourier Serialized Gradient Estimation

To provide the accordant gradient estimation for hash function $\text{sign}(\cdot)$, we approximate it by introducing its Fourier Series decomposition in the frequency domain. Specifically, $\text{sign}(\cdot)$ can be viewed as a special case of the periodical Square Wave Function $t(x)$ within the length $2H$, i.e., $\text{sign}(\phi) = t(\phi), |\phi| < H$. Since $t(x)$ can be decomposed in Fourier Series, we shall have:

$$\text{sign}(\phi) = \frac{4}{\pi} \sum_{i=1,3,5,\cdots}^{+\infty} \frac{1}{i} \sin(\frac{\pi i \phi}{H}), \text{ where } |\phi| < H. \tag{10}$$

Fourier Series decomposition of $\text{sign}(\cdot)$ with infinite terms is a lossless transformation [46]. Thus, as shown in Figure 2(c), we can set the finite expanding term n to obtain its approximation version as follows:

$$\text{sign}(\phi) \doteq \frac{4}{\pi} \sum_{i=1,3,5,\cdots}^{n} \frac{1}{i} \sin(\frac{\pi i \phi}{H}). \tag{11}$$

The corresponding derivatives can be derived accordingly as:

$$\frac{\partial \text{sign}(\phi)}{\partial \phi} \doteq \frac{4}{H} \sum_{i=1,3,5,\cdots}^{n} \cos(\frac{\pi i \phi}{H}). \tag{12}$$

Different from other gradient estimators such as tanh-alike [15, 44] and SignSwish [11], approximating $\text{sign}(\cdot)$ function with its Fourier Series will not corrupt the main direction of factual gradients in model optimization [58]. This is beneficial to bridge a coordinated transformation from the continuous values to its corresponding binarization for node representations, which significantly retains the discriminability of binarized representations and

produces better retrieval accuracy accordingly. We present this performance comparison in § 5.6 of experiments. To summarize, as shown in Equation (13), to learn and optimize the binarized embeddings for graph nodes, we apply the strict sign($\cdot$) function for forward propagation and estimate the gradients $\frac{\partial \text{sign}(\phi)}{\partial \phi}$ for backward propagation.

$$
\begin{cases}
Q^{(l)} = \text{sign}(\phi), & \text{Forward propagation.} \\
\frac{\partial Q^{(l)}}{\partial \phi} \doteq \frac{4}{H} \sum_{i=1,3,5,\cdots}^{n} \cos(\frac{\pi i \phi}{H}). & \text{Backward propagation.}
\end{cases}
\tag{13}
$$

4.5 Score Prediction and Model Optimization

4.5.1 ***Matching score prediction.*** Given two nodes $x \in \mathcal{V}_1$ and $y \in \mathcal{V}_2$, one natural manner to implement the score function is *inner-product*, mainly for its simplicity as:

$$
\widehat{Y}_{x,y} = (\alpha_x Q_x)^\mathsf{T} \cdot (\alpha_y Q_y).
\tag{14}
$$

However, the inner product in Equation (14) is still conducted in the (continuous) Euclidean space with *full-precision arithmetics*. To bridge the connection between the inner product and Hamming distance measurement, we introduce Theorem 2 as follows:

Theorem 2 (Hamming Distance Matching). *Given two hash codes, we have* $(\alpha_x Q_x)^\mathsf{T} \cdot (\alpha_y Q_y) = \alpha_x \alpha_y (d - 2D_H(Q_x, Q_y))$.

$D_H(\cdot, \cdot)$ denotes the Hamming distance between two inputs. Based on Theorem 2, we transform the score computation to the Hamming distance matching. By doing so, we can reduce most number of the floating-point operations (#FLOPs) in the original score computation formulation (Equation (14)) to efficient hamming distance matching. This can develop substantial computation acceleration that is analyzed in Appendix C.

4.5.2 ***Multi-loss Objective Function.*** Our objective function consists of two components, i.e., graph reconstruction loss $\mathcal{L}_{rec}$ and BPR loss $\mathcal{L}_{bpr}$. Generally, these two loss functions harness the regularization effect to each other. The intuition of such design is:

- $\mathcal{L}_{rec}$ reconstructs the observed bipartite graph topology;
- $\mathcal{L}_{bpr}$ ranks the matching scores computed from the hash codes.

Concretely, we implement $\mathcal{L}_{rec}$ with Cross-entropy loss:

$$
\mathcal{L}_{rec} = \sum_{x \in \mathcal{V}_1} \Big(\sum_{y \in \mathcal{N}(x)} \ln \sigma\big((V_x^{(0)})^\mathsf{T} \cdot V_y^{(0)}\big) + \sum_{y' \notin \mathcal{N}(x)} \ln \big(1 - \sigma((V_x^{(0)})^\mathsf{T} \cdot V_{y'}^{(0)})\big) \Big),
\tag{15}
$$

where σ is the activation function, e.g., Sigmoid. $\mathcal{L}_{rec}$ bases on the initial continuous embeddings before the graph convolution, e.g., $V_x^{(0)}$, providing the most fundamental information for topology reconstruction. As for $\mathcal{L}_{bpr}$, we employ *Bayesian Personalized Ranking* (BPR) loss as:

$$
\mathcal{L}_{bpr} = - \sum_{x \in \mathcal{V}_1} \sum_{\substack{y \in \mathcal{N}(x) \\ y' \notin \mathcal{N}(x)}} \ln \sigma(\widehat{Y}_{x,y} - \widehat{Y}_{x,y'}).
\tag{16}
$$

$\mathcal{L}_{bpr}$ encourages the predicted score of an observed edge to be higher than its unobserved counterparts [19]. Let Θ denote the set of trainable embeddings regularized by the parameter λ_2 to avoid over-fitting. our final objective function is finally defined as:

$$
\mathcal{L} = \mathcal{L}_{rec} + \lambda_1 \mathcal{L}_{bpr} + \lambda_2 ||\Theta||_2^2.
\tag{17}
$$

Table 1: The statistics of datasets.

	MovieLens	Gowalla	Pinterest	Yelp2018	AMZ-Book	Dianping
$|\mathcal{V}_1|$	6,040	29,858	55,186	31,668	52,643	332,295
$|\mathcal{V}_2|$	3,952	40,981	9,916	38,048	91,599	1,362
$|\mathcal{E}|$	1,000,209	1,027,370	1,463,556	1,561,406	2,984,108	10,000,014
Density	0.04190	0.00084	0.00267	0.00130	0.00062	0.02210

So far, we have introduced all technical parts of BGCH and attached the pseudocodes in Appendix B. We present all the theorem proofs and complexity analyses in Appendix C.

5 Experimental Evaluation

We evaluate BGCH to answer the following research questions:

- **RQ1.** How does BGCH perform compared to state-of-the-art hashing-based models in the Top-N Hamming space retrieval?
- **RQ2.** what is the performance gap between BGCH and the full-precision models in terms of long-list retrieval quality?
- **RQ3.** What are the benefits of proposed components in BGCH?
- **RQ4.** what is the practical BGCH resource consumption?
- **RQ5.** How does the Fourier Series decomposition perform *w.r.t.* retrieval accuracy and training efficiency?

5.1 Experiment Setup

Datasets and evaluation metrics. We include six real-world bipartite graphs in Table 1 that are widely evaluated [7, 8, 19, 52, 60, 64]. We adopt evaluation protocols Recall@N and NDCG@N to measure the Top-N Hamming space ranking capability. Dataset details and evaluation procedure are explained in Appendix D.

Baselines. We include the following representative hashing-based models for (1) general object retrieval (LSH [14]), (2) image search (HashNet [4]), and (3) Top-N candidate generation for recommendation (Hash_Gumbel [24, 42], CIGAR [28] and HashGNN [49]). We also include several state-of-the-art full-precision[4] recommender models, i.e., NeurCF [20], NGCF [52], DGCF [53], LightGCN [19], for the long-list ranking quality comparison. Model introductions are referred in Appendix D. Early hashing methods, e.g., SH [55], RMMH [26], LCH [62], are excluded mainly because the above competing models [4, 28] have already validated the performance superiority over them.

5.2 Top-N Hamming Space Query (RQ1)

To evaluate **fine-to-coarse** Top-N ranking capability, we set N=1000. We first report the results of Recall@20$_{1000}$ and NDCG@20$_{1000}$[5] in Top-1000 search in Table 2 and then plot the holistic Recall and NDCG metric curves of {20, 50, 100, 200, 500, 1000} of Top-1000 in Figure 3. We set convolution iteration number as 2 and embedding dimension as 256 for BGCH and baselines for fair comparison.

- **The results demonstrate the superiority of BGCH model over prior hashing-based models.** (1) As shown in Table 2, the state-of-the-art model, i.e., HashGNN, works better than traditional hashing-based baselines, e.g., LSH, HashNet, CIGAR. This indicates that, compared to graph-based models, a direct adaptation of conventional (i.e., non-graph-based) hashing methods may

---

[4]They are denoted by FT32 as we implement them with float32 in the experiments.
[5]We then use simple notation Recall@20, NDCG@20 if there is no ambiguity caused.

Table 2: Results of Recall@20 and NDCG@20 in Top-1000 retrieval: (1) "R" and "N" denote the Recall and NGCG; (2) the bold indicate BGCH and the underline represents the best-performing models; (3) Mark * denotes scenarios where Wilcoxon signed-rank tests indicate statistically significant improvements over the second-best models over 95% confidence level.

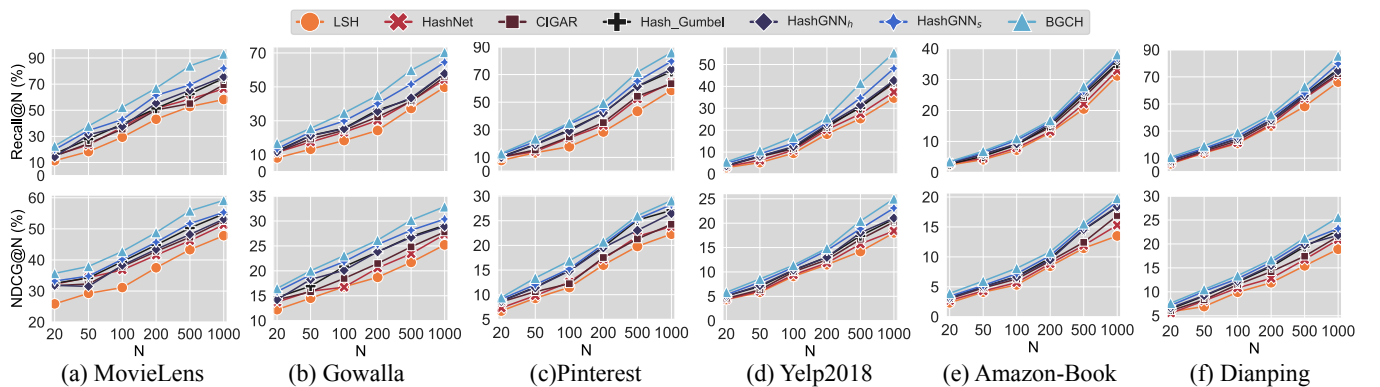| Dataset | MovieLens (%) | | Gowalla (%) | | Pinterest (%) | | Yelp2018 (%) | | AMZ-Book (%) | | Dianping (%) | |
Metric	$R@20_{1000}$	$N@20_{1000}$	$R@20_{1000}$	$N@20_{1000}$	$R@20_{1000}$	$N@20_{1000}$	$R@20_{1000}$	$N@20_{1000}$	$R@20_{1000}$	$N@20_{1000}$	$R@20_{1000}$	$N@20_{1000}$
LSH	11.38	25.87	8.14	12.23	7.88	6.71	2.91	4.35	2.41	2.34	5.85	5.84
HashNet	15.43	32.23	11.38	13.74	10.27	7.33	3.37	4.41	2.86	2.71	6.24	5.59
CIGAR	14.84	31.73	11.57	14.21	10.34	8.53	3.65	4.57	3.05	3.03	6.91	6.03
Hash_Gumbel	16.62	32.48	12.26	14.68	10.53	8.74	3.85	5.12	2.69	3.24	8.29	6.43
HashGNN$_h$	14.21	31.83	11.63	14.21	10.15	8.67	3.77	5.04	3.09	3.15	8.34	6.68
HashGNN$_s$	19.87	33.21	13.45	14.87	12.38	9.11	4.86	5.34	3.34	3.45	9.57	7.13
BGCH	**22.86***	**36.26***	**16.73***	**16.48***	**12.78***	**9.42***	**5.51***	**5.84***	**3.48***	**3.92***	**10.66***	**7.63***
% Gain	15.05%	9.18%	24.39%	10.83%	3.23%	3.40%	13.37%	9.36%	4.19%	13.62%	11.39%	7.01%

Figure 3: Top-N retrieval quality with N in {20, 50, 100, 200, 500, 1000} (best view in color).

be hard to achieve comparable performance, mainly because of the effectiveness of *graph convolutional* architecture in capturing latent information within the bipartite graph topology for hash encoding preparation. (2) Owing to our proposed model components, e.g., *adaptive graph convolutional hashing*, BGCH consistently outperforms HashGNN over all datasets, by 3.23%~24.39%, and 3.40%~ 13.62% *w.r.t.* Recall@20 and NDCG@20, respectively. (3) Furthermore, we conduct the Wilcoxon signed-rank tests at BGCH. The results verify that all BGCH improvements over the second-best model are statistically significant over 95% confidence level. (4) To explain these, our proposed topology-aware graph convolutional hashing approach effectively enriches the graph node embeddings. Our proposed feature dispersion further alleviates the feature erosion issue caused by numerical binarization. Last but not least, our proposed Fourier serialized gradient estimation is also vital to provide accurate gradients for model optimization. We conduct the ablation study later in § 5.4.

- **By varying N from 20 to 1000, BGCH consistently shows competitive performance compared to the baselines.** While Recall@N indicates the fraction of relevant objects in Top-N retrieval, NDCG@N measures the ranking capability for relative orders. As shown in Figure 3: (1) Compared to the approximated version of HashGNN, i.e., HashGNN$_s$, BGCH generally obtains stable and significant improvements of both Recall and NDCG metrics over all six benchmarks with N from 20 to 1000. (2)

Table 3: NDCG@1000 results of Float32-based models.

	Movie	Gowalla	Pinterest	Yelp2018	AMZ-Book	Dianping
NeurCF	58.76	32.07	28.79	24.69	19.83	25.54
NGCF	60.28	32.13	29.78	25.23	20.37	25.76
DGCF	62.41	34.97	31.47	26.28	21.74	26.87
LightGCN	62.88	35.26	31.32	26.55	21.92	27.28
BGCH	**59.16**	**32.87**	**29.09**	**25.01**	**19.79**	**25.57**
% capacity	94.08%	93.22%	92.44%	94.20%	90.28%	93.73%

Apart from the higher retrieval quality, another advantage of BGCH over HashGNN$_s$ is that it still supports bitwise operations, i.e., hamming distance matching, for inference acceleration. This is because, to improve the prediction accuracy, HashGNN$_s$ adopts a Bernoulli random variable to provide the probability of replacing the certain digits in the hash codes with the original continuous values, which thus disables the bitwise computation. As we present in § 5.5, BGCH achieves over 8× inference acceleration over HashGNN$_s$, which is particularly promising for query-based online matching and retrieval applications.

5.3 Comparing to FT32-based Models (RQ2)

In this section, we also compare BGCH with several full-precision (FT32-based) models to evaluate the long-list search quality. As we can observe from Table 3, we have the following analyses. (1) We notice that our model BGCH generally performs competitively with early full-precision models, e.g., NeurCF and NGCF, over all datasets. As for the state-of-the-art model LightGCN, our model

Table 4: Ablation study.

Variant	MovieLens		Gowalla		Pinterest		Yelp2018		AMZ-Book		Dianping	
	R@20	N@20	R@20	N@20	R@20	N@20	R@20	N@20	R@20	N@20	R@20	N@20
w/o FD	22.82 (-0.17%)	35.87 (-1.08%)	15.92 (-4.84%)	15.79 (-4.19%)	12.25 (-4.15%)	9.07 (-3.72%)	5.16 (-6.35%)	5.49 (-2.49%)	3.26 (-6.32%)	3.57 (-8.93%)	10.46 (-1.88%)	7.50 (-1.70%)
w/o AH-TA	19.54(-14.52%)	29.17(-19.55%)	13.49(-19.37%)	12.38(-24.88%)	12.24 (-4.23%)	8.86 (-5.94%)	4.77(-13.43%)	5.18(-11.30%)	2.49(-28.45%)	2.86(-27.04%)	9.83 (-7.79%)	6.87 (-9.96%)
w/o AH-RF	16.73(-26.82%)	26.97(-25.62%)	11.24(-32.82%)	11.29(-31.49%)	10.18(-20.34%)	7.33(-22.19%)	3.76(-31.76%)	4.30(-26.37%)	3.27 (-6.03%)	3.64 (-7.14%)	8.33(-21.86%)	6.93 (-9.17%)
w/in LF	21.06 (-7.87%)	34.59 (-4.61%)	15.48 (-7.47%)	15.38 (-6.67%)	11.94 (-6.57%)	8.89 (-5.63%)	4.86(-11.80%)	5.17(-11.47%)	3.14 (-9.77%)	3.62 (-7.65%)	9.40(-11.82%)	7.27 (-4.72%)
w/o $\mathcal{L}_{bpr}$	21.42 (-6.30%)	34.83 (-3.94%)	15.87 (-5.14%)	15.66 (-4.98%)	12.33 (-3.52%)	9.17 (-2.65%)	5.31 (-3.63%)	5.61 (-3.94%)	3.35 (-3.74%)	3.77 (-3.83%)	10.21 (-4.22%)	7.38 (-3.28%)
w/o $\mathcal{L}_{rec}$	17.01(-25.59%)	27.16(-25.10%)	12.27(-26.66%)	12.63(-23.36%)	10.81(-15.41%)	7.86(-16.56%)	3.93(-28.68%)	4.37(-25.17%)	3.19 (-8.33%)	3.73 (-4.85%)	8.82(-17.26%)	7.26 (-4.85%)
BGCH	**22.86**	**36.26**	**16.73**	**16.48**	**12.78**	**9.42**	**5.51**	**5.84**	**3.48**	**3.92**	**10.66**	**7.63**

can generally achieve over 90% of the Top-1000 ranking capability. (2) The performance of BGCH demonstrates its effectiveness in guaranteeing the long-list Top-N retrieval quality. This is useful for some industrial applications, e.g., recommender systems, which usually consist of two major stages: *candidate generation* and *re-ranking*. Thus, obviously, the good quality of candidate generation directly reduces the complexity of next-stage re-ranking, as the search space is substantially pruned. (3) Considering the *efficiency in Hamming space retrieval* and the *reduced space cost* of those learned hash codes, we believe that BGCH can provide the optional alternative to these full-precision models, especially in scenarios with limited computation resources.

5.4 Ablation Study (RQ3)

We evaluate the necessity of model components with Top-20 search metrics and report the results in Table 4.

Effect of Feature Dispersion. We first analyze the effect of our proposed feature dispersion approach for hedging the feature erosion in hash encoding. We introduce the model variant, denoted by *w/o FD*, to directly disable it by setting η as 0. As shown in Table 4, the performance gap between *w/o FD* and BGCH well demonstrates the effectiveness of dispersing the latent features before embedding binarization for hashing over these six datasets. Moreover, let the density summarized in Table 1 be computed by $\frac{|\mathcal{V}_1|\times|\mathcal{V}_2|}{|\mathcal{E}|}$. In sparse datasets, i.e., Gowalla (0.00084), Pinterest (0.00267), Yelp2018 (0.00130), and AMZ-Book (0.00062), the performance decay between BGCH and *w/o FD* is much larger than on the other two datasets, i.e., MovieLens (0.04190) and Dianping (0.02210). This is because sparse datasets are more sensitive to hashing as they may not have insufficient training edges to abridge the gap against their unhashed version. Another promising approach to tackle data sparsity issue is *data augmentation* [67] and we leave it for future work.

Effect of Adaptive Graph Convolutional Hashing. Then we study this model component by setting two variants, where: (1) *w/o AH-TA* only disables the *topology-awareness of hashing* and sets it as the final encoder after all graph convolutions (just like the conventional manner [4, 49]); (2) *w/o AH-RF* removes the *rescaling factors*. From Table 4 results, we have the following observations:

(1) The variant *w/o AH-TA* consistently underperforms BGCH. This demonstrates that simply using the rear output embeddings from the GCN framework may not sufficiently model the unique latent node features for hashing, especially for the rich structural information within different graph depths. While in BGCH, by capturing the intermediate information for representation

enrichment, the topology-aware hashing can effectively alleviate the limited expressivity of discrete hash codes.

(2) Apart from the topology-aware hashing, another key point for contributing to the performance improvement is the *rescaling factor* that we introduced in Equation (8). After removing it from BGCH, variant *w/o AH-RF* presents huge performance decay. Although these factors are directly calculated and may not be theoretically optimal, they reflect the numerical uniqueness of embeddings for later hash encoding, which substantially improves BGCH's prediction capability. We study the *determinacy* design of factor computation in the following section.

Design of Learnable Rescaling. We include another variant namely *w/in LF* to indicate the model version using *learnable rescaling factors*. As shown in Table 4, the design of learnable rescaling factors in *w/in LF* does not achieve good performance as expected. One explanation is that, our proposed model currently does not post a strong mathematical constraint to the learnable factors (α_x), e.g., $\alpha_x^{(l)} = \mathrm{argmin}(\widetilde{V}_x^{(l)}, \alpha_x^{(l)}Q_x^{(l)})$, mainly because of its additional training complexity; and purely relying on the stochastic optimization, e.g., stochastic gradient descent (SGD), may hardly reach the optimum. Considering the additional search space introduced from this regularization design, we argue that our deterministic rescaling method is simple yet effective in practice.

Effect of Multi-loss in Optimization. Lastly, to study the effect of BPR loss $\mathcal{L}_{bpr}$ and graph reconstruction loss $\mathcal{L}_{rec}$, we set two variants, termed by *w/o $\mathcal{L}_{bpr}$* and *w/o $\mathcal{L}_{rec}$*, to optimize BGCH separately. As shown in Table 4, with all other model components, partially using each one of $\mathcal{L}_{bpr}$ and $\mathcal{L}_{rec}$ can not achieve the expected performance. This confirms the effectiveness of our proposed multi-loss design: while $\mathcal{L}_{bpr}$ learns to assign higher prediction values to observed edges, i.e., $Y_{x,y} = 1$, than the unobserved node pair counterparts, $\mathcal{L}_{rec}$ transfers the graph reconstruction problem to a classification task by using the original embeddings in training. By collectively optimizing these two loss functions, our model BGCH can learn precise intermediate embeddings from $\mathcal{L}_{rec}$, and generate targeted hash codes with high-quality relative order information regularized by $\mathcal{L}_{bpr}$ accordingly.

5.5 Resource Consumption Analysis (RQ4)

Due to the various value ranges over all six datasets, we compactly report the value ratios of BGCH over the state-of-the-art hashing-based model HashGNN$_s$ in Figure 4.

Model Training Time Cost. As indicated by the metric "*T-Time*" in Figure 4, we notice that training HashGNN$_s$ is more time-consuming than our proposed model. The main reason is

Figure 4: Resource consumption ratios.

(a) Top-20 retrieval metrics (b) Training time costs

Figure 5: Fourier Series decomposition term n in BGCH.

that HashGNN adopts the early GCN framework [17] as the model backbone, while our model follows the latest framework [19] to remove operations, e.g., self-connection, feature transformation, and nonlinear activation. In addition, on the two largest datasets AMZ-Book and Dianping, the training cost ratio further increases to around 5~6 times. This is because we have to decrease the batch size of $HashGNN_s$ for tractable training process.

Online Inference Time Cost. We randomly generate 1,000 queries and evaluate the computation time cost. To present a fair comparison, we disable all parallel arithmetic techniques (e.g., MKL, BLAS) by using the open-source toolkit[6]. Indicated by "*I-Time*" in Figure 4, our model with Hamming distance matching generally achieves over $8\times$ computation acceleration over $HashGNN_s$ on all datasets. This is because, as we have explained in § 5.2, $HashGNN_s$ randomly replaces the hash codes with their original continuous embeddings for relaxation and adopts floating-point arithmetics to pursue performance improvement while sacrificing the computation acceleration from the bitwise operations.

Hash Codes Memory Footprint. Binarized embeddings can largely reduce memory space consumption. Compared to the state-of-the-art hashing-based model $HashGNN_s$, our BGCH further achieves about $9\times$ of memory space reduction for the hash codes. As we have just explained, since $HashGNN_s$ interprets hash codes with random real-value digits, it thus requires additional cost to distinguish binary digits from full-precision ones. On the contrary, BGCH separates the storage of binarized encoding parts and corresponding rescaling factors, thus providing the advantage for space overhead optimization.

5.6 Study of Fourier Gradient Estimation (RQ5)

We take our largest dataset Dianping for illustration and the analysis can be generally popularized to the other datasets.

Effect of Decomposition Term n. We vary the decomposition term n from 1 to 16. As shown in Figure 5, we have two observations: (1) Different decomposition terms will surely affect the final retrieval quality, as theoretically, the larger n increases, the more accurate gradients can be estimated. However, in practice, too large values of n may introduce the overfitting risk, which implies that keeping a moderate n, e.g., n=4 in Figure 5(a), can already maximize the model performance. (2) By varying n from 1 to 16, the training time per iteration of BGCH slowly increases. This generally coincides with our complexity analysis in Appendix C, in which the majority of training cost lies in our feature dispersion and graph convolutional hashing, as $O(\frac{2cs(K+L)|\mathcal{E}|^2}{B}) \gg O(snd|\mathcal{E}|)$.
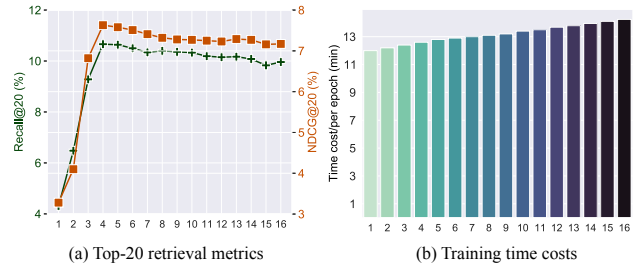
---

[6]https://www.lfd.uci.edu/~gohlke/pythonlibs/

Table 5: Gradient estimator comparison on Recall@20.

	Movie	Gowalla	Pinterest	Yelp2018	AMZ-Book	Dianping
STE	20.93 (-8.44%)	14.85 (-11.24%)	12.35 (-3.36%)	5.24 (-4.90%)	3.12(-10.34%)	10.34 (-3.00%)
Tanh	21.75 (-4.86%)	15.06 (-9.98%)	12.36 (-3.29%)	5.43 (-1.45%)	3.21 (-7.76%)	10.41 (-2.34%)
SignSwish	22.13 (-3.19%)	15.62 (-6.63%)	12.44 (-2.66%)	5.50 (-0.18%)	3.34 (-4.02%)	10.43 (-2.16%)
Sigmoid	22.08 (-3.41%)	15.21 (-9.09%)	12.52 (-2.03%)	5.53 (+0.03%)	3.18 (-8.62%)	10.38 (-2.63%)
PBE	21.68 (-5.16%)	15.05(-10.04%)	12.32 (-3.60%)	5.35 (-2.90%)	3.13(-10.06%)	10.47 (-1.78%)
BGCH	**22.86**	**16.73**	**12.78**	**5.51**	**3.48**	**10.66**

Comparison with Other Gradient Estimators. We include several recent gradient estimators, i.e., *Tanh-like* [15, 44], *Sign-Swish* [11], *Sigmoid* [59], and *projected-based estimator* [39] (denoted as PBE). (1) The results summarized in Table 5 well demonstrate the superiority of our proposed Fourier Series decomposition to sign(·) function in gradient estimation. As we have briefly explained, most existing estimators employ the *visually similar* function approximation to sign(·); compared to STE, they generally provide better gradient estimation. (2) However, for those bipartite graphs with heavy sparsity, e.g., Gowalla (0.00084) and AMZ-Book (0.00062), graph-based models may hardly collect enough structural information for effective hash codes training. Based on the limited training samples, these *theoretically irrelevant* estimators may not effectively rectify the optimization deviation, and thus present a recognizable performance gap against our proposed Fourier serialized estimator.

6 Conclusion

We study the graph convolutional hashing over bipartite graphs for efficient Hamming space search, by proposing BGCH with three effectual modules. Extensive experiments demonstrate the model superiority over conventional counterparts and validate the effectiveness of all proposed modules. As for future work, we plan to investigate modeling with the *semi-supervised* graph setting [47, 48] mainly for its commonality in practice. Moreover, another promising direction is to upgrade BGCH for *inductive learning* [17], such that it can make adaptive matching and prediction for evolving graphs with structural updates.

References

[1] Haoli Bai, Wei Zhang, Lu Hou, Lifeng Shang, Jing Jin, Xin Jiang, Qun Liu, Michael Lyu, and Irwin King. 2020. Pushing the limit of bert quantization. arXiv (2020).

[2] Yoshua Bengio, Nicholas Léonard, and Aaron Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv (2013).

[3] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann LeCun. 2013. Spectral networks and locally connected networks on graphs. arXiv (2013).

[4] Zhangjie Cao, Mingsheng Long, Jianmin Wang, and Philip S Yu. 2017. Hashnet: Deep learning to hash by continuation. In ICCV. 5608–5617.

[5] Minyu Chen, Guoqiang Li, Chen Ma, Jingyang Li, and Hongfei Fu. 2022. Answering Coding Questions via Dense Retrieval on GitHub Repositories. In COLING.

[6] Yankai Chen, Huifeng Guo, Yingxue Zhang, Chen Ma, Ruiming Tang, Jingjie Li, and Irwin King. 2022. Learning binarized graph representations with multi-faceted quantization reinforcement for top-k recommendation. In SIGKDD.

[7] Yankai Chen, Menglin Yang, Yingxue Zhang, Mengchen Zhao, Ziqiao Meng, Jianye Hao, and Irwin King. 2022. Modeling Scale-free Graphs with Hyperbolic Geometry for Knowledge-aware Recommendation. In WSDM. 94–102.

[8] Yankai Chen, Yaming Yang, Yujing Wang, Jing Bai, Xiangchen Song, and Irwin King. 2022. Attentive Knowledge-aware Graph Convolutional Networks with Collaborative Guidance for Personalized Recommendation. In ICDE.

[9] Yankai Chen, Yifei Zhang, Huifeng Guo, Ruiming Tang, and Irwin King. 2022. An Effective Post-training Embedding Binarization Approach for Fast Online Top-K Passage Matching. In AACL-IJCNLP. 102–108.

[10] Zhiyong Cheng, Ying Ding, Lei Zhu, and Mohan Kankanhalli. 2018. Aspect-aware latent factor model: Rating prediction with ratings and reviews. In WWW.

[11] Sajad Darabi, Mouloud Belbahri, Matthieu Courbariaux, and Vahid Partovi Nia. 2018. Bnn+: Improved binary network training. (2018).

[12] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. CNNs on graphs with fast localized spectral filtering. NeurIPS 29 (2016).

[13] Yifan Gao, Chien-Sheng Wu, Jingjing Li, Shafiq Joty, Steven CH Hoi, Caiming Xiong, Irwin King, and Michael Lyu. 2020. Discern: Discourse-Aware Entailment Reasoning Network for Conversational Machine Reading. In EMNLP. 2439–2449.

[14] Aristides Gionis, Piotr Indyk, Rajeev Motwani, et al. 1999. Similarity search in high dimensions via hashing. In PVLDB, Vol. 99. 518–529.

[15] Ruihao Gong, Xianglong Liu, Shenghu Jiang, Tianxiang Li, Peng Hu, Jiazhen Lin, Fengwei Yu, and Junjie Yan. 2019. Differentiable soft quantization: Bridging full-precision and low-bit neural networks. In ICCV. 4852–4861.

[16] Aditya Grover and Jure Leskovec. 2016. node2vec: Scalable feature learning for networks. In SIGKDD. 855–864.

[17] William L Hamilton, Rex Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In NeurIPS. 1025–1035.

[18] Ruining He and Julian McAuley. 2016. Modeling the visual evolution of fashion trends with one-class collaborative filtering. In WWW. 507–517.

[19] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In SIGIR. 639–648.

[20] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural collaborative filtering. In WWW. 173–182.

[21] Xiangnan He, Hanwang Zhang, Min-Yen Kan, and Tat-Seng Chua. 2016. Fast matrix factorization for online recommendation with implicit feedback. In ICCV.

[22] Xuming Hu, Fukun Ma, Chenyao Liu, Chenwei Zhang, Lijie Wen, and Philip S Yu. 2021. Semi-supervised Relation Extraction via Incremental Meta Self-Training. In EMNLP.

[23] Xuming Hu, Lijie Wen, Yusong Xu, Chenwei Zhang, and Philip S. Yu. 2020. Self-ORE: Self-supervised Relational Feature Learning for Open Relation Extraction. In EMNLP. 3673–3682.

[24] Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. (2017).

[25] Herve Jegou, Matthijs Douze, and Cordelia Schmid. 2010. Product quantization for nearest neighbor search. TPAMI 33, 1 (2010), 117–128.

[26] Alexis Joly and Olivier Buisson. 2011. Random maximum margin hashing. In CVPR. IEEE, 873–880.

[27] Wang-Cheng Kang, Derek Zhiyuan Cheng, Tiansheng Yao, Xinyang Yi, Ting Chen, Lichan Hong, and Ed H Chi. 2021. Learning to Embed Categorical Features without Embedding Tables for Recommendation. In SIGKDD. 840–850.

[28] Wang-Cheng Kang and Julian McAuley. 2019. Candidate generation with binary codes for large-scale top-n recommendation. In CIKM. 1523–1532.

[29] Diederik Kingma and Jimmy Ba. 2015. Method for stochastic optimization. ICLR.

[30] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. (2017).

[31] Piotr Koniusz, Fei Yan, Philippe-Henri Gosselin, and Krystian Mikolajczyk. 2016. Higher-order occurrence pooling for bags-of-words: Visual concept detection. TPAMI 39, 2 (2016), 313–326.

[32] Piotr Koniusz, Hongguang Zhang, and Fatih Porikli. 2018. A deeper look at power normalizations. In CVPR. 5774–5783.

[33] Tim Kraska, Alex Beutel, Ed H Chi, Jeffrey Dean, and Neoklis Polyzotis. 2018. The case for learned index structures. In SIGMOD. 489–504.

[34] Guohao Li, Matthias Muller, Ali Thabet, and Bernard Ghanem. 2019. Deepgcns: Can gcns go as deep as cnns?. In ICCV. 9267–9276.

[35] Hao Li, Wei Liu, and Heng Ji. 2014. Two-Stage Hashing for Fast Document Retrieval.. In ACL. 495–500.

[36] Jingjing Li, Zichao Li, Lili Mou, Xin Jiang, Michael Lyu, and Irwin King. 2020. Unsupervised text generation by learning from search. NeurIPS.

[37] Dawen Liang, Laurent Charlin, James McInerney, and David M Blei. 2016. Modeling user exposure in recommendation. In WWW. 951–961.

[38] Xiaofan Lin, Cong Zhao, and Wei Pan. 2017. Towards accurate binary convolutional neural network. (2017).

[39] Chunlei Liu, Wenrui Ding, Xin Xia, Yuan Hu, Baochang Zhang, Jianzhuang Liu, Bohan Zhuang, and Guodong Guo. 2019. Rectified Binary Convolutional Networks for Enhancing the Performance of 1-bit DCNNs. arXiv (2019).

[40] Chen Ma, Peng Kang, and Xue Liu. 2019. Hierarchical gating networks for sequential recommendation. In SIGKDD.

[41] Chen Ma, Liheng Ma, Yingxue Zhang, Ruiming Tang, Xue Liu, and Mark Coates. 2020. Probabilistic metric learning with adaptive margin for top-k recommendation. In SIGKDD.

[42] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. 2017. The concrete distribution: A continuous relaxation of discrete random variables. (2017).

[43] Yury Nahshan, Brian Chmiel, Chaim Baskin, Evgenii Zheltonozhskii, Ron Banner, Alex M Bronstein, and Avi Mendelson. 2021. Loss aware post-training quantization. Machine Learning (2021), 1–18.

[44] Haotong Qin, Ruihao Gong, Xianglong Liu, Mingzhu Shen, Ziran Wei, Fengwei Yu, and Jingkuan Song. 2020. Forward and backward information retention for accurate BNNs. In CVPR. 2250–2259.

[45] Mohammad Rastegari, Vicente Ordonez, Joseph Redmon, and Ali Farhadi. 2016. Imagenet classification using binary CNNs. In ECCV. Springer, 525–542.

[46] BRYAN RUST. 2013. CONVERGENCE OF FOURIER SERIES. (2013).

[47] Zixing Song, Ziqiao Meng, Yifei Zhang, and Irwin King. 2021. Semi-supervised Multi-label Learning for Graph-structured Data. In CIKM.

[48] Zixing Song, Yifei Zhang, and Irwin King. 2022. Towards an Optimal Asymmetric Graph Structure for Robust Semi-supervised Node Classification. In SIGKDD.

[49] Qiaoyu Tan, Ninghao Liu, Xing Zhao, Hongxia Yang, Jingren Zhou, and Xia Hu. 2020. Learning to hash with GNNs for RecSys. In WWW. 1988–1998.

[50] Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. In ICLR.

[51] Jingdong Wang, Ting Zhang, Nicu Sebe, Heng Tao Shen, et al. 2017. A survey on learning to hash. TPAMI 40, 4 (2017), 769–790.

[52] Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. 2019. Neural graph collaborative filtering. In SIGIR. 165–174.

[53] Xiang Wang, Hongye Jin, An Zhang, Xiangnan He, Tong Xu, and Tat-Seng Chua. 2020. Disentangled graph collaborative filtering. In ICLR. 1001–1010.

[54] Xing Wei, Yue Zhang, Yihong Gong, Jiawei Zhang, and Nanning Zheng. 2018. Grassmann pooling as compact homogeneous bilinear pooling for fine-grained visual classification. In ECCV. 355–370.

[55] Yair Weiss and Antonio Torralba. 2008. Spectral hashing. NeurIPS (2008).

[56] Felix Wu, Amauri Souza, Tianyi Zhang, Christopher Fifty, Tao Yu, and Kilian Weinberger. 2019. Simplifying graph convolutional networks. In ICML.

[57] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. 2019. How powerful are graph neural networks? (2019).

[58] Yixing Xu, Kai Han, Chang Xu, Yehui Tang, Chunjing Xu, and Yunhe Wang. 2021. Learning frequency domain approximation for BNNs. arXiv (2021).

[59] Jiwei Yang, Xu Shen, Jun Xing, Xinmei Tian, Houqiang Li, Bing Deng, Jianqiang Huang, and Xian-Sheng Hua. 2019. Quantization Networks. In CVPR. 7308–7316.

[60] Menglin Yang, Min Zhou, Jiahong Liu, Defu Lian, and Irwin King. 2022. HRCF: Enhancing collaborative filtering via hyperbolic geometric regularization. In WWW. 2462–2471.

[61] Tan Yu, Yunfeng Cai, and Ping Li. 2020. Toward faster and simpler matrix normalization via rank-1 update. In ECCV. Springer, 203–219.

[62] Dell Zhang, Jun Wang, Deng Cai, and Jinsong Lu. 2010. Laplacian co-hashing of terms and documents. In ECIR. Springer, 577–580.

[63] Jiani Zhang, Xingjian Shi, Shenglin Zhao, and Irwin King. 2019. Stacked and reconstructed graph convolutional networks for recommender systems. IJCAI.

[64] Xinni Zhang, Yankai Chen, Cuiyun Gao, Qing Liao, Shenglin Zhao, and Irwin King. 2022. Knowledge-aware Neural Networks with Personalized Feature Referencing for Cold-start Recommendation. arXiv (2022).

[65] Yan Zhang, Defu Lian, and Guowu Yang. 2017. Discrete personalized ranking for fast collaborative filtering from implicit feedback. In AAAI, Vol. 31.

[66] Yifei Zhang and Hao Zhu. 2019. Doc2hash: Learning discrete latent variables for documents retrieval. In NAACL. 2235–2240.

[67] Yifei Zhang, Hao Zhu, Zixing Song, Piotr Koniusz, and Irwin King. 2022. COSTA: Covariance-Preserving Feature Augmentation for Graph Contrastive Learning. In SIGKDD.

[68] Yifei Zhang, Hao Zhu, Zixing Song, Piotr Koniusz, and Irwin King. 2022. Spectral Feature Augmentation for Graph Contrastive Learning and Beyond. In arXiv.

Table 6: Notations and meanings.

Notation	Meaning
$\mathcal{G}, \mathcal{V}_1, \mathcal{V}_2, \mathcal{E}$	Bipartite graph with sets of nodes and edges.
c, d	Convolution dimension and hash code dimension.
$V_x^{(l)}$	Node x's embedding at iteration l.
Y	Edge transactions where $Y_{x,y} = 1$ indicates the interaction existence between nodes x and y, and otherwise $Y_{x,y} = 0$.
A, D	Adjacency matrix and associated diagonal degree matrix.
$p^{(k)}, P$	Dispersing vector at iteration k and the projection matrix.
$\widetilde{V}$	Feature-dispersed embedding matrix.
$\widehat{Y}$	Estimated matching scores.
$Q_x^{(l)}$	Hash code segment of node x at iteration l.
$\alpha^{(l)}$	x's rescaling factor computed at the l-th convolution.
Q_x	Target hash codes of node x.
L, K	Numbers of convolutional hashing and dispersion generation.
$\mathcal{L}_{rec}, \mathcal{L}_{bpr}, \mathcal{L}$	Two loss terms of final objective function $\mathcal{L}$.
$\eta, H, n, \lambda_1, \lambda_2$	hyper-parameters.

Algorithm 1: BGCH algorithm.

1 **while** <u>model not converge</u> **do**
2   **for** <u>$k = 0, \cdots, K - 1$</u> **do**
3      $p^{(k+1)} \leftarrow (V^{(0)})^\top V^{(0)} p^{(k)}$;
4    $P \leftarrow$ obtain the projection matrix ;   ▷ Eq. (4)
5    $\widetilde{V}^{(0)} \leftarrow$ obtain the feature-dispersed embeddings ;   ▷ Eq. (5)
6   **for** <u>$l = 0, \cdots, L - 1$</u> **do**
7      $\widetilde{V}^{(l+1)} \leftarrow (D^{-\frac{1}{2}} A D^{-\frac{1}{2}}) \widetilde{V}^{(l)}$;   ▷ Eq. (6)
8      $Q^{(l+1)} \leftarrow \text{sign}(\widetilde{V}^{(l+1)})$;   ▷ Eq. (7)
9      $\alpha \leftarrow$ calculate the rescaling factors ;   ▷ Eq. (8)
10   **for** <u>$x \in \mathcal{V}_1, y \in \mathcal{N}(x)$</u> **do**
11      $\widehat{Y}_{x,y} \leftarrow \alpha_x \alpha_y (d - 2D_H(Q_x, Q_y))$;   ▷ Eq. (14)&Thm. 2
12    $\mathcal{L} \leftarrow$ compute loss and optimize the model ; ▷ Eq's. (15-17)
13 **Function** Gradient_estimator($\mathcal{L}$):
14    $\frac{\partial \mathcal{L}}{\partial V} \leftarrow \frac{\partial \mathcal{L}}{\partial Q} \cdot \frac{4}{H} \sum_{i=1,3,5,\cdots}^{n} \cos(\frac{\pi i V}{H})$;   ▷ Eq. (12)

A Loss Landscape Visualization

We simulate the optimization trajectories of learnable embeddings and visually compare the loss landscapes of non-hashing and hashing versions in Figure 2(a). Specifically, we manually assign perturbations [1, 43] to the embeddings on MovieLens dataset as: $V_x^{(l)} = V_x^{(l)} \pm p \cdot |\overline{V_x^{(l)}}| \cdot \mathbf{1}^{(l)}$. where $|\overline{V_x^{(l)}}|$ represents the absolute mean of entries in $V_x^{(l)}$ and perturbation magnitudes p are from $\{0.01, \cdots, 0.50\}$. $\mathbf{1}$ is an all-one vector. For pairs of perturbed node embeddings, we plot their loss distribution accordingly. As we can observe, the non-hashing version produces a flat loss surface, showing the local convexity. On the contrary, the hashing counterpart has a bumping and complex loss landscape.

B Notation Table and BGCH Pseudo-codes

We use bold uppercase and calligraphy characters for matrices and sets. The non-bolded denote graph nodes or scalars. Key notations and Pseudocodes are explained in Table 6 and Algorithm 1.

C Theoretical Proofs and Analyses

Theorem 1 (Feature Dispersion). *Let $V^{(l)} = U_1 \Sigma U_2^\top$, where U_1 and U_2 are unitary matrices and descending singular value matrix $\Sigma = \text{diag}(\sigma_1, \sigma_2, \cdots, \sigma_c)$. Then $\mathbb{E}(\widetilde{V}^{(l)}) = U_1 \Sigma \Sigma_\mu U_2^\top$ where $\Sigma_\mu = \text{diag}(\mu_1, \mu_2, \cdots, \mu_c)_{0 < \mu_1 \cdots c < 1}$ is in ascending order.*

PROOF. We focus on comparing $(\widetilde{V}^{(0)}, V^{(0)})$, which can be easily popularized to any convolution layer, i.e., $(\widetilde{V}^{(l)}, V^{(l)})$. Conducting SVD decomposition on $V^{(0)}$, we have $V^{(0)} = U_1 \Sigma U_2^\top$, where U_1 and U_2 are unitary matrices of singular vectors. Then following $p^{(k)} = V^{(0)\top} V^{(0)} p^{(k-1)}$, we shall have $p^{(k)} = (V^{(0)\top} V^{(0)})^k p^{(0)}$. Replacing $V^{(0)}$ with its SVD decomposition, we get the following equation:

$$p^{(k)} = (U_2 \Sigma^{2k} U_2^\top) p^{(0)}. \tag{18}$$

Then we transform the projection matrix in Equation (4) as follows:

$$P = \frac{p^{(k)} p^{(k)\top}}{p^{(k)\top} p^{(k)}} = \frac{(U_2 \Sigma^{2k} U_2^\top) p^{(0)} p^{(0)\top} (U_2 \Sigma^{2k} U_2^\top)}{p^{(0)\top} (U_2 \Sigma^{2k} U_2^\top)(U_2 \Sigma^{2k} U_2^\top) p^{(0)}}$$
$$= U_2 \Sigma^{2k} \frac{U_2^\top p^{(0)} p^{(0)\top} U_2}{p^{(0)\top} U_2 \Sigma^{4k} U_2^\top p^{(0)}} \Sigma^{2k} U_2^\top. \tag{19}$$

Let $t = U_2^\top p^{(0)}$, we can further simplify the above equation to:

$$P = U_2 \Sigma^{2k} \frac{t t^\top}{t^\top \Sigma^{4k} t} \Sigma^{2k} U_2^\top, \text{ where scalar } t^\top \Sigma^{4k} t = \sum_{j=1}^{c} t_j^2 \sigma_j^{4k}. \tag{20}$$

Recalling that $\widetilde{V}^{(0)} = V^{(0)}(I - \epsilon P)$, $\mathbb{E}(\widetilde{V}^{(0)}) = V^{(0)} - \epsilon \cdot \mathbb{E}(V^{(0)} P)$. Then we focus on the term $\mathbb{E}(V^{(0)} P)$ as follows:

$$\mathbb{E}(V^{(0)} P) = \frac{1}{t^\top \Sigma^{4k} t} U_1 \Sigma^{2k+1} \cdot \mathbb{E}(t t^\top) \cdot \Sigma^{2k} U_2^\top \tag{21}$$

Since $p^{(0)} \sim \mathcal{N}(0, I)$ and U_2 is a unitary matrix, thus $t \sim \mathcal{N}(0, I)$. This indices that each element of t, e.g., $t_j \in t$, is *i.i.d.* random variable. Thus, $\mathbb{E}(t_j \cdot t_k) = 0$ for $j \neq k$ and $\mathbb{E}(t t^\top)$ is a diagonal matrix, i.e., $\mathbb{E}(t t^\top) = \text{diag}(t_1^2, t_2^2, \cdots, t_c^2)$. We then have:

$$\mathbb{E}(V^{(0)} P) = U_1 \cdot \text{diag}\left(\frac{\sigma_1 t_1^2 \sigma_1^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}, \cdots, \frac{\sigma_c t_c^2 \sigma_c^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}\right) \cdot U_2^\top. \tag{22}$$

Therefore,

$$\mathbb{E}(\widetilde{V}^{(0)}) = U_1 \cdot \text{diag}\left(\sigma_1 - \epsilon \frac{\sigma_1 t_1^2 \sigma_1^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}, \cdots, \sigma_c - \epsilon \frac{\sigma_c t_c^2 \sigma_c^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}\right) \cdot U_2^\top. \tag{23}$$

Let $\mu_k = 1 - \epsilon \frac{t_k^2 \sigma_k^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}$, with $\epsilon \in (0, 1)$, obviously, $0 < \mu_k < 1$. Furthermore, $\forall k_1 \geq k_2$, we have:

$$\mu_{k_1} - \mu_{k_2} = \epsilon \mathbb{E}\left(\frac{t_{k_1}^2 \sigma_{k_1}^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}} - \frac{t_{k_2}^2 \sigma_{k_2}^{4k}}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}\right) \geq \epsilon \sigma_{k_1}^{4k} \cdot \mathbb{E}\left(\frac{t_{k_1}^2 - t_{k_2}^2}{\sum_{j=1}^{c} t_j^2 \sigma_j^{4k}}\right) = 0, \tag{24}$$

as $\sigma_{k_2}^{4k} \geq \sigma_{k_1}^{4k}$, and t_{k_1} and t_{k_2} are *i.i.d.* random variables with same normal distribution. Equation (24) proves that μ_k is *monotone non-decreasing* in expectation, which completes the proof. □

Theorem 2 (Hamming Distance Matching). *Given two hash codes, we have $(\alpha_x Q_u)^\top \cdot (\alpha_y Q_y) = \alpha_x \alpha_y (d - 2D_H(Q_x, Q_y))$.*

PROOF.
$$Q_x^\top \cdot Q_y = |\{i|(Q_x)_i = (Q_y)_i = 1\}| + |\{i|(Q_x)_i = (Q_y)_i = -1\}|$$
$$- |\{i|(Q_x)_i \neq (Q_y)_i\}|$$
$$= d - 2 \cdot |\{i|(Q_x)_i \neq (Q_y)_i\}| = d - 2D_H(Q_x, Q_y), \tag{25}$$

which completes the proof. □

Training time complexity. As shown in Table 7, $|\mathcal{E}|$, B, s, and n are the edge number, batch size, numbers of train iterations and

Table 7: Traing time complexity.

Graph Norm.	Feat. Disp.	Conv. & Hash.	Loss Comp.	Grad. Est.
$O(2|\mathcal{E}|)$	$O(\frac{2csK|\mathcal{E}|^2}{B})$	$O(\frac{2csL|\mathcal{E}|^2}{B})$	$O(2s(c+d)|\mathcal{E}|)$	$O(snd|\mathcal{E}|)$

Table 8: Complexity of space and computation.

	Space cost	#FLOP	#BOP
float32-based	$32|\mathcal{V}_1 \cup \mathcal{V}_2|d$	$O(|\mathcal{V}_1| \cdot |\mathcal{V}_2|d)$	-
BGCH	$|\mathcal{V}_1 \cup \mathcal{V}_2|(d + 32(L+1))$	$O(4|\mathcal{V}_1| \cdot |\mathcal{V}_2|)$	$O(|\mathcal{V}_1| \cdot |\mathcal{V}_2|d)$

Fourier Series decomposition terms. (1) The time complexity of the graph normalization, i.e., $D^{-\frac{1}{2}}AD^{-\frac{1}{2}}$, is $O(2|\mathcal{E}|)$. (2) Before the graph convolution, we first conduct the feature dispersion only for the initial node embeddings, e.g., $V_x^{(0)}$, which takes $O(\frac{2csK|\mathcal{E}|^2}{B})$ complexity. In our experiment, hyper-parameter $K \leq 3$. (3) In graph convolution, the time complexity is $O(\frac{2csL|\mathcal{E}|^2}{B})$, where $L \leq 4$ is a common setting [17, 19, 30, 52] to avoid *over-smoothing* [34]. (4) As for the loss computation, BGCH takes $O(2sc|\mathcal{E}|)$ to compute $\mathcal{L}_{rec}$ and $O(2sd|\mathcal{E}|)$ for $\mathcal{L}_{bpr}$. (5) Lastly, BGCH takes $O(snd|\mathcal{E}|)$ to estimate the gradients for the d-dimension hash codes. Thus, thee total complexity in total is quadratic to the graph edge numbers, i.e., $|\mathcal{E}|$, which is common in GCN frameworks.

Hash codes space cost. As shown in Table 8, the total space cost of hash codes is $O(|\mathcal{V}_1 \cup \mathcal{V}_2|(d + 32(L+1)))$ bits, supposing that we use float32 for those rescaling factors in $L + 1$ iterations. Compared to the continuous embedding size, i.e., $32|\mathcal{V}_1 \cup \mathcal{V}_2|d$ bits, the theoretical size reduction ratio thus is:

$$ratio = \frac{32|\mathcal{V}_1 \cup \mathcal{V}_2|d}{|\mathcal{V}_1 \cup \mathcal{V}_2|(d + 32(L+1))} = \frac{32d}{d + 32(L+1)}. \quad (26)$$

As we just explained, stacking too many iteration layers will incurring performance detriment [34]. Hence, if $L \leq 4$ and $d = 1024$, BGCH can achieve considerable size compression.

Online matching. The original score formulation in Equation (14) contains d floating-point operations (#FLOPs). As shown in Table 8, using Hamming distance matching can convert the most of floating-point arithmetics to binary operations (#BOPs), with slightly more #FLOPs for scalar computations, i.e., $4 \ll d$.

D Experiment Setup Details

Datasets. We evaluate our model on the following six six datasets:

(1) **MovieLens**[7] is a widely adopted benchmark between *users* and *movies*. Similar to the setting in [21, 49], if the user x has rated item y, we set the edge $Y_{x,y} = 1$, otherwise $Y_{x,y} = 0$.

(2) **Gowalla**[8] [19, 49, 52, 53] is the dataset [37] between *customers* and *their check-in locations* collected from Gowalla.

(3) **Pinterest**[9] is an open dataset for image recommendation between *users* and *images*. Edges represent the pins over images initiated by users.

(4) **Yelp2018**[10] is from Yelp Challenge 2018 Edition, bipartitely modeling between *users* and *local businesses*.

(5) **AMZ-Book**[11] is the bipartite graph between *readers* and *books*, organized from the book collection of Amazon-review [18].

(6) **Dianping**[12] is a commercial dataset between *users* and *local businesses* recording their diverse interactions, e.g., clicking, saving, and purchasing.

Evaluation metrics. To evaluate the model performance of Hamming space retrieval over bipartite graphs, we directly deploy our model BGCH in the basic user-item recommendation scenarios. Specifically, given a query node, we apply the hash codes to match Top-N answers for the query with the closest Hamming distances, and thus adopt two widely-used evaluation protocols Recall@N and NDCG@N to measure the ranking capability.

Implementations. We implement our models under Python 3.6 and PyTorch 1.14.0 on a Linux machine with 1 Nvidia GeForce RTX 3090 GPU, 4 Intel Core i7-8700 CPUs, 32 GB of RAM with 3.20GHz. For all the baselines, we follow the official hyper-parameter settings. We apply a grid search if lacking recommended model settings. The dimension is searched in {32, 64, 128, 256, 512}. The learning rate η is tuned within $\{10^{-3}, 5 \times 10^{-3}, 10^{-2}, 5 \times 10^{-2}\}$ and the coefficient λ is tuned among $\{10^{-5}, 10^{-4}, 10^{-3}\}$. We initialize and optimize all models with default normal initializer and Adam optimizer [29].

Baselines. All baselines studied in this paper are introduced as:

(1) **LSH** [14] is a classical hashing method. LSH is proposed to approximate the similarity search for massive high-dimensional data and we introduce it for Top-N object search by following the adaptation in [49].

(2) **HashNet** [4] is a representative deep hashing method that is originally proposed for multimedia retrieval tasks. Similar to [49], we adapt it for graph data by modifying it with the general graph convolutional network.

(3) **CIGAR** [28] is a state-of-the-art neural-network-based framework for fast Top-N candidate generation in recommendation. CIGAR can be further followed by a full-precision re-ranking algorithm. And we only use its hashing part for fair comparison.

(4) **Hash_Gumbel** is a variance of BGCH with Gumbel-softmax for hash encoding and gradient estimation [24, 42]. Specifically, we first expand each embedding bit to a size-two one-hot encoding. Then it utilizes the Gumbel-softmax trick to replace $sign(\cdot)$ as relaxation for binary hash code generation.

(5) **HashGNN** [49] is the state-of-the-art learning to hash based method with GCN framework. We use HashGNN$_h$ to denote the vanilla version with *hard encoding* proposed in [49], where each element of user-item embeddings is strictly binarized. We use HashGNN$_s$ to denote its proposed approximated version.

(6) **NeurCF** [20] is one representative deep neural network model for collaborative filtering in recommendation.

(7) **NGCF** [52] is one of the representative graph-based recommender models with collaborative filtering methodology.

(8) **DGCF** [53] is a state-of-the-art graph-based model that learns disentangled user intents for better recommendation.

(9) **LightGCN** [19] is another latest state-of-the-art GCN-based recommender model that has been widely evaluated.

---

[7] https://grouplens.org/datasets/movielens/1m/
[8] https://github.com/gusye1234/LightGCN-PyTorch/tree/master/data/gowalla
[9] https://sites.google.com/site/xueatalphabeta/dataset-1/pinterest_iccv
[10] https://github.com/gusye1234/LightGCN-PyTorch/tree/master/data/yelp2018

---

[11] https://github.com/gusye1234/LightGCN-PyTorch/tree/master/data/amazon-book
[12] https://www.dianping.com/