Engineering Applications of Artificial Intelligence 177 (2026) 114979

Contents lists available at ScienceDirect

Engineering Applications of Artificial Intelligence

journal homepage: www.elsevier.com/locate/engappai

Research paper

Granular-ball based robust representation learning for social recommendation

Xiaofei Zhu [a],*, Shiyan Wu [a], Li Liu [b], Shuyin Xia [c], Yi Wang [d], Guoyin Wang [e]

[a] *School of Computer Science and Engineering, Chongqing University of Technology, Chongqing 400054, China*
[b] *School of Big Data and Software Engineering, Chongqing University, Chongqing 401331, China*
[c] *School of Computer Science and Technology, Chongqing University of Posts and Telecommunications, Chongqing 400065, China*
[d] *Chongqing Ant Consumer Finance Co., Ltd., Ant Group, Chongqing 400064, China*
[e] *National Center for Applied Mathematics in Chongqing, Chongqing Normal University, Chongqing 401331, China*

ARTICLE INFO

Keywords:
Recommender systems
Social recommendation
Granular-ball computing
Knowledge distillation
Diffusion model

ABSTRACT

Social recommendation systems seek to leverage social relationships to mitigate data sparsity and cold-start issues by augmenting user–item interactions. However, existing methods encounter two critical limitations: (1) They predominantly model user–item interactions at a fine-grained granular level of user/item nodes, neglecting the potential coarse-grained collaborative patterns; and (2) They usually suppress noisy edges in social graphs from a single granular perspective, failing to adjust the denoising granularity according to the actual strength of relationships between users. To address these challenges, we propose GBRSR, a novel Granular-ball based Robust Representation Learning framework. Inspired by the "Global-first" cognitive principle, Granular-ball Computing (GBC), which represents data as granular-ball units with geometric significance, has garnered significant attention due to its outstanding performance in many fields. We leverage GBC theory for representation distillation, transferring coarse-grained knowledge to enhance fine-grained node-level representations. In addition, we employs a granular-ball based structure denoising strategy to prune noisy user relationships, while simultaneously alleviating noise in user representations through a diffusion process. Extensive experiments on three real-world benchmark datasets validate the superiority of GBRSR in recommendation accuracy and robustness, particularly under noisy and sparse conditions.

1. Introduction

Recommendation systems seek to help users cope with information overload by providing personalized recommendations based on their historical interactions (Sharma et al., 2024; Wu et al., 2023). However, previous research faces significant limitations when handling cold-start users with sparse interaction history. To address this, social recommendation has emerged as a promising solution which leverages social relationships (e.g., friendships, trust networks) as auxiliary signals to supplement behavioral data (S. Li et al., 2024; Xiong et al., 2025). By jointly modeling user–item interactions and social influence dynamics, these methods enhance recommendation robustness, particularly in data-scarce scenarios (Chen et al., 2024; Huang et al., 2025).

Early social recommendation approaches (Wang et al., 2019; Song et al., 2019; Shao and Liu, 2021) usually aggregate features from neighbors in both interaction and social graphs through heuristic fusion mechanisms, aiming to enrich user representations and alleviate data sparsity limitations. For example, GraphRec (Fan et al., 2019) combines user–item interactions and social relationships, jointly incorporating interactions and opinions to enhance user representations. SMIN (Long et al., 2021) integrates self-supervised learning and metagraph-based heterogeneous graph neural networks to capture multi-hop user–item relationships, strengthening user representations (Afoudi et al., 2023). While these methods have demonstrated relatively remarkable performance and can handle the limitation of data sparsity, they may suffer from noisy relationships within social connections, which may compromise the robustness of user representations. Recently, some methods propose to utilize supervisory signals (Long et al., 2021; Yu et al., 2021; Du et al., 2022; Wu et al., 2022; Xiong et al., 2024) and diffusion models (Wu et al., 2019; Z. Li et al., 2024) to effectively alleviate the noise issue (Jiang and Zuo, 2025). For instance, DSL (Wang et al., 2023) employs a denoised self-augmented learning framework to minimize the impact of noisy social connections. RecDiff (Z. Li et al., 2024) employ diffusion models to remove noisy social influences and enhance recommendation accuracy.

---

* Corresponding author.
E-mail address: zxf@cqut.edu.cn (X. Zhu).

https://doi.org/10.1016/j.engappai.2026.114979
Received 24 October 2025; Received in revised form 3 April 2026; Accepted 26 April 2026

Although above-mentioned approaches have garnered significant attention and achieved considerable progress, they still encounter two key challenges: (i) In the interaction graph, prior approaches mainly concentrate on learning representations at the most granular level of user/item points. To be specific, these methods usually capture user–item fine-grained interaction relationships from a local perspective. However, they neglect modeling potential coarse-grained interactions between users and items from a broader perspective, resulting in limitations in comprehensive modeling capabilities for collaborative signals. (ii) In the social graph, there may be a number of irrelevant relationships, posing challenges for models to capture high-quality user preferences. Traditional methods tend to adopt fixed denoising strategies that only remove irrelevant users at a single granularity and cannot adjust the denoising granularity based on the actual relationship strength between users. This limits their ability to handle noise in social networks.

Inspired by the “Global-first” cognitive mechanism of Granular-ball Computing (Xia et al., 2025a,b), we attempt to learn representations from multiple granular levels by exploring the stability of coarse-grained granular balls to guide the refinement of fine-grained node representations. Furthermore, as the granular-ball partition process is capable of learning high-quality granular balls from the user–item interaction graph via removing irrelevant edges, we argue that it can serve as an effective way to alleviate the noise issue. In this paper, we propose a Granular-ball based Robust Representation Learning (GBRSR) model for social recommendation. To be specific, GBRSR develops a cross-granularity representation distillation module to learn informative user/item representations and incorporate a diffusion-enhanced granular-ball denoising module for adaptive noise elimination. (i) In cross-granularity representation distillation, we partition user/item points into multiple adaptive-size granular-balls based on user–item interaction relationships in order to capture user preference patterns across different granularities. Subsequently, a distillation process between granular-ball and user/item points is utilized to feed the information from user/item balls to individual user/item points for modeling multi-granularity signals. (ii) In diffusion-enhanced granular-ball denoising, we address the noise issue by simultaneously apply a granular-ball based structural denoising process and a diffusion based representation denoising process. To be specific, GBRSR adaptively removes noisy edges from the social graph using Granular-ball Computing technique, ensuring the retention of relevant user groups. Then, the denoised representation generated by the structural denoising process is further fed into a diffusion denoising process to address the semantic representation noise. Unlike traditional methods that learns user/item representation at the most granular level of individual points, we encapsulate and represent user/item points using granular-balls, thereby enhancing the robustness of representation learning. Finally, we integrate the representations from both modules and utilize them to predict items of potential interest to users. The major contributions of this paper are summarized as follows:

- Inspired by the “Global-first” cognitive principle, we propose to incorporate granular-ball computing technique for representation distillation by transferring coarse-grained knowledge to fine-grained node-level representations.
- We develop a diffusion-enhanced granular-ball denoising module that synergistically prunes noisy user relationships via a granular-ball based structure denoising strategy and alleviates noise within user representations based on a diffusion process.
- We conduct extensive experimentation on three real-world datasets, and validate the superiority of GBRSR in recommendation accuracy and robustness, particularly under noisy and sparse conditions.

The remainder of this paper is organized as follows. Section 2 reviews related work in Social Recommendation and Granular-ball Computing. Section 3 introduces the preliminaries and details the proposed GBRSR framework. Section 4 presents the experimental settings, results, and comprehensive analysis. Finally, Section 5 concludes the paper and discusses future research directions.

2. Related work

2.1. Social recommendation

Social recommendation has evolved significantly over the past decade. Initially, research focused on integrating social networks as static regularization terms into matrix factorization frameworks. With the advent of deep learning, Graph Neural Networks (GNNs) become the dominant paradigm due to their ability to model high-order connectivity. Pioneering works like GraphRec (Fan et al., 2019) and DiffNet (Wu et al., 2019) introduce attention mechanisms and layer-wise diffusion to aggregate social neighbors’ preferences, significantly alleviating the cold-start problem. Subsequent developments, such as SAMN (Chen et al., 2019) and MHCN (Yu et al., 2021), refine this process by employing multi-channel hypergraphs and aspect-level attention to capture complex social motifs.

However, a growing consensus in the community is that social networks contain unavoidable noise (e.g., casual acquaintances) which can degrade recommendation performance. DSL (Wang et al., 2023) employs cross-view contrastive learning to align social and interaction views, thereby implicitly filtering noise via self-supervision. GDMSR (Quan et al., 2023) proposes curriculum learning to progressively mask unreliable edges based on confidence scores. Most recently, generative diffusion models like RecDiff (Z. Li et al., 2024) have been applied to reconstruct clean user embeddings from noisy inputs. Despite these advances, most existing denoising methods still operate on the *fine-grained* node level. They lack a mechanism to model the macroscopic structural stability of user communities, which limits their robustness in highly sparse environments.

2.2. Granular-ball computing

Granular-ball Computing (GBC) originates from Granular Computing (GrC), a paradigm for human-inspired information processing (Pedrycz, 2001). Traditional GrC methods mainly focus on constructing information granules using fuzzy sets, rough sets, or quotients spaces to solve complex problems at different levels of abstraction. To address the efficiency bottlenecks of traditional irregular granules, Wang (2017) formally introduces the concept of “Granular-ball Computing” based on the cognitive mechanism of “Global first”. This approach innovatively utilizes hyperspheres (balls) to adaptively cover the data space, ensuring that the computation cost scales with the number of balls rather than the massive number of individual samples.

Since its inception, Granular-ball Computing has witnessed rapid development across various domains. In supervised learning, GBC (Xia et al., 2019; Xia et al., 2023) and its variants demonstrate that coarse-grained ball centers could effectively replace fine-grained samples for efficient classification. In unsupervised learning, subsequent works extend the theory to clustering (Cheng et al., 2024; Xie et al., 2024; Zhang et al., 2025; Xie et al., 2025) and outlier detection, proving that the geometric regularity of granular balls offers superior robustness against label noise and outliers compared to point-based methods.

Recently, integrating Granular-ball Computing with deep representation learning has become a burgeoning research direction. Researchers have begun to explore mapping discrete granular balls into continuous vector spaces to handle high-dimensional data (Su et al., 2025). In particular, we draw significant inspiration from GBGC (Xia et al., 2025b), which pioneeringly introduces Granular-ball Computing into the field of graph coarsening. This work demonstrates that constructing adaptive granular structures can efficiently preserve global topological integrity while filtering out local noise. Complementarily, recent studies on robust granular computing (Yang et al., 2025) have introduced the principle of justifiable granularity to explicitly model data uncertainty, further validating the efficacy of adaptive granular boundaries in mitigating decision risks. Motivated by these granular paradigms, we leverage the stability of coarse-grained granular balls to guide the refinement of fine-grained node representations in social recommendation, specifically addressing the challenges of noisy interactions and structural uncertainty.

3. Methodology

3.1. Preliminaries

Problem Definition. Let $\mathcal{U} = (u_1, \dots, u_M)$ represent the set of users, and $\mathcal{I} = (v_1, \dots, v_N)$ represent the set of items. The user–item interaction matrix $\mathcal{R} \in \mathbb{R}^{M\times N}$ and social relation matrix $A_s \in \mathbb{R}^{M\times M}$ form the interaction graph G_r and social graph G_s respectively. The goal is to learn effective user and item embeddings that accurately predict the likelihood of user u_i interacting with item v_j, while addressing varying scales of relevant users/item and multi-granularity noise in social connections. In social recommendation, the goal is to predict the next item by leveraging the interaction graph user–item G_r and social graph G_s.

Granular-Ball on Graph. To address the limitations of fine-grained node modeling, we introduce the concept of the granular-ball as the fundamental unit for coarse-grained representation. According to the Granular-ball Computing theory proposed by Xia et al. (2023), a granular-ball GB is defined as a hyperspherical granular structure $GB = \{x_i \mid \|x_i - \mathbf{c}\| \le r\}$, where x represents the set of data points contained in the ball, with c and r specifying the center and the radius, respectively. Physically, in the context of our graph-based task, a granular-ball GB corresponds to a cohesive subgraph containing a set of nodes with high structural homogeneity. Geometrically, we formalize it as a hyperspherical region in the embedding space, characterized by a tuple $(\mathbf{c}, r)$, where $\mathbf{c}$ denotes the semantic center (a representative embedding of the subgraph) and r represents the coverage radius. For any node u belonging to GB, its representation $\mathbf{e}_u$ is constrained within the boundary, i.e., $\|\mathbf{e}_u - \mathbf{c}\|_2 \le r$. This formulation allows the model to capture the dominant structural skeleton of social or interaction graphs while filtering out local noise through the coarse-grained encapsulation.

3.2. Overview

To address the challenges of inconsistent interaction granularity and multi-granularity noise in social networks, we propose the **Granular-ball based Robust Representation Learning (GBRSR)**, a robust framework for social recommendation. As shown in Fig. 1, GBRSR consists of two key modules: **Cross-granularity Representation Distillation (CRD)** and **Diffusion-Enhanced Granular-ball Denoising (DEGD)**. In CRD, we perform granular-ball clustering on similarity-based user and item graphs to form adaptive coarse-grained groups, which are linked to construct a user–item Ball Graph. A Ball-based GCN is trained on this graph to capture collaborative signals at the ball level, and a distillation mechanism is utilized to transfer the coarse-grained information to enhance the fine-grained user/item representations. In DEGD, the social graph is decomposed into compact granular-ball subgraphs, where weak or noisy edges are gradually pruned to suppress structural noise. A diffusion process is then applied to further refine user embeddings and mitigate representational noise. By jointly modeling multi-granular preferences and applying dual-channel denoising, GBRSR effectively improves recommendation robustness in sparse and noisy social environments.

3.3. Cross-granularity representation distillation

3.3.1. Backbone interaction modeling

Given the effectiveness of lightweight graph convolutional architectures in collaborative filtering, we adopt a simplified GCN architecture as the backbone to capture user–item collaborative signals. Specifically, we denote the user–item bipartite graph as $\mathcal{G}_r$, and define the interaction matrix as $\mathcal{E} \in \mathbb{R}^{I\times J}$, where I and J denote the number of users and items, respectively. The corresponding normalized symmetric adjacency matrix is constructed as:

$$L_r = \mathbf{D}_r^{-\frac{1}{2}}\mathbf{A}_r\mathbf{D}_r^{-\frac{1}{2}}, \quad \mathbf{A}_r = \begin{bmatrix} \mathbf{0} & \mathcal{E} \\ \mathcal{E}^\top & \mathbf{0} \end{bmatrix}, \tag{1}$$

where $\mathbf{D}_r$ is the diagonal degree matrix of $\mathbf{A}_r$. The representation of users and items is propagated across L layers. At the lth layer, the embedding update is defined as:

$$\mathcal{E}_r^{(l)} = (\mathcal{L}_r + \mathbf{I}) \cdot \mathcal{E}_r^{(l-1)}, \quad \bar{\mathcal{E}}_r = \sum_{l=0}^{L} \mathcal{E}_r^{(l)}, \tag{2}$$

We emphasize that only the user–item interaction graph $\mathcal{G}_r$ is involved in backbone embedding learning. Although the user–user social graph is introduced later for structural denoising and auxiliary optimization, it is not involved in representation propagation or item prediction in this stage.

3.3.2. Adaptive granular-ball generation

To reduce graph redundancy while preserving structural coherence, we propose a graph coarsening mechanism driven by a density-constrained splitting strategy. Prior to executing this granular partition, we first establish the foundational homogeneous graphs to capture latent semantic relations. Following methodology of Xu et al. (2025), we construct the user graph $\mathcal{G}_u$ and item graph $\mathcal{G}_i$ using the aggregated user and item representations derived from $\bar{\mathcal{E}}_r$. Specifically, we extract the user embeddings $\bar{\mathcal{E}}_u$ and item embeddings $\bar{\mathcal{E}}_i$ from $\bar{\mathcal{E}}_r$, which encode high-order collaborative signals and serve as semantic foundations for the subsequent coarsening process. For each pair of nodes (i, j), the similarity score is computed as:

$$S_{ij} = \frac{\bar{\mathbf{e}}_i^\top \bar{\mathbf{e}}_j}{\|\bar{\mathbf{e}}_i\|\|\bar{\mathbf{e}}_j\|}, \quad \tilde{\mathbf{S}} = \tilde{\mathbf{D}}^{-1/2}\mathbf{S}\tilde{\mathbf{D}}^{-1/2}, \tag{3}$$

where $\bar{\mathbf{e}}_i$ and $\bar{\mathbf{e}}_j$ denote the embedding vectors corresponding to nodes i and j in $\bar{\mathcal{E}}_r$, and $\tilde{\mathbf{S}}$ is the normalized similarity matrix, ensuring scale-invariant message propagation in subsequent GNN layers. The constructed graph topology serves as the structural substrate for the subsequent adaptive granular partition.

Quality-Guided Adaptive Splitting. The core of our method is a hierarchical granular-ball generation process governed by a specific granularity density function. Let G_B be a candidate granular ball inducing a subgraph $\mathcal{G}_{G_B} = (\mathcal{V}_{G_B}, \mathcal{E}_{G_B})$. We define its structural quality score $Q(G_B)$ as:

$$Q(G_B) = \frac{|\mathcal{E}_{G_B}|}{|\mathcal{V}_{G_B}|^2}, \tag{4}$$

where $|\mathcal{E}_{G_B}|$ and $|\mathcal{V}_{G_B}|$ denote the number of edges and nodes within a ball, respectively. This function evaluates the structural tightness of any candidate subgraph, ensuring that the splitting process preserves local connectivity.

The generation process adopts a top-down approach, as summarized in Algorithm 1. For a given parent ball G_B, we employ a modularity-based clustering operator to tentatively partition it into a set of child sub-balls $\{G_{S_1}, G_{S_2}, \dots, G_{S_m}\}$. Specifically, we adopt the Louvain (Blondel et al., 2024) algorithm to perform this task, as it directly optimizes graph modularity to uncover natural community structures without requiring a predefined number of clusters. This ensures that the generated sub-balls inherently reflect the topological connectivity of the graph. To determine when to stop partitioning users or items into smaller balls,

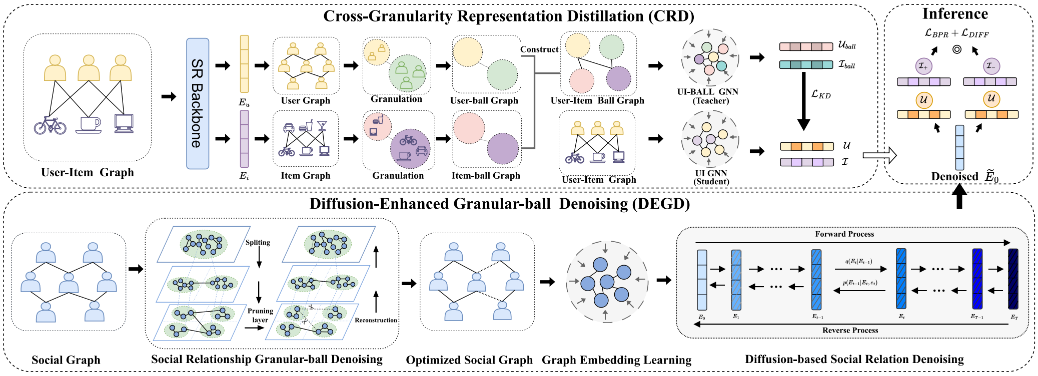

Fig. 1. An overview of the GBRSR framework.

we propose three stopping criteria, including quality score, maximum recursion depth, and the number of edges within a ball. Formally, the splitting terminates if:

$$\exists G_{S_i} \in \{G_{S_1}, \dots, G_{S_m}\}, \quad \text{s.t.} \quad Q(G_{S_i}) < Q(G_B). \tag{5}$$

This constraint acts as a self-adaptive regularizer, preventing over-segmentation in sparse regions while allowing fine-grained granulation in dense interaction areas.

Granular-ball Representation and Initialization. Upon convergence of the adaptive splitting, we map the topological granular balls into the semantic embedding space to initialize the coarse graph features. For each generated ball G_{B_k}, we compute its semantic center $\mathbf{c}_k$ and dispersion radius r_k:

$$\mathbf{c}_k = \frac{1}{|G_{B_k}|} \sum_{i \in G_{B_k}} \mathbf{e}_i, \quad r_k = \frac{1}{|G_{B_k}|} \sum_{i \in G_{B_k}} \|\mathbf{e}_i - \mathbf{c}_k\|_2, \tag{6}$$

where $\mathbf{e}_i$ represents the initial embedding of node i. Crucially, the center vector $\mathbf{c}_k$ serves as the initialized node feature for the ball in the coarsened graph, effectively aggregating the semantic information of its constituent nodes. The radius r_k quantifies the internal semantic consistency.

User–Item Ball Graph Construction. Finally, we construct the User–Item Ball Graph to facilitate high-order message propagation. An edge is established between a user ball B_i^u and an item ball B_j^v if their cumulative interaction strength exceeds a significance threshold τ:

$$w_{ij} = \mathbb{I}\left[\sum_{u \in B_i^u} \sum_{v \in B_j^v} \mathcal{A}_{uv} \geq \tau\right]. \tag{7}$$

The resulting graph provides a noise-resilient and structurally concise representation of the original user–item interactions.

3.3.3. *Cross-granularity distillation*

To effectively integrate multi-granularity structural signals, we propose a cross-granularity distillation framework to enhance representation robustness and efficiency. This framework leverages coarse-grained representations derived from the user-ball graph $\mathcal{G}_u^{\text{ball}}$ and item-ball graph $\mathcal{G}_i^{\text{ball}}$, learned via a teacher Graph Convolutional Network (GCN). Specifically, the teacher GCN aggregates node embeddings within each granular-ball to obtain coarse-grained representations that capture high-order and high-level structural semantics.

For any user u, we associate it with its corresponding user-ball $G^u_{B_k}$, which contains structurally similar users. The teacher model computes the coarse-grained representation of this user group by averaging their node embeddings:

$$u_{\text{ball}} = \frac{1}{|G^u_{B_k}|} \sum_{u' \in G^u_{B_k}} e_{u'}, \tag{8}$$

where $e_{u'}$ denotes the fine-grained embedding of user u', and $|G^u_{B_k}|$ is the number of users in the group. Similarly, for each item v, we compute its coarse-grained embedding v_{ball} by aggregating the embeddings of all items in its corresponding granular-ball $G^v_{B_l}$.

To bridge the semantic gap between coarse-grained and fine-grained levels, we train a student GCN to align its node-level representations with the high-level signals from the teacher. This alignment is achieved through a cross-granularity distillation loss, defined as:

$$\mathcal{L}_{DIST} = \sum_u \|u - u_{\text{ball}}\|^2 + \sum_v \|v - v_{\text{ball}}\|^2, \tag{9}$$

where u and v denote the student model's fine-grained embeddings of users and items, while u_{ball} and v_{ball} are the corresponding coarse-grained supervision signals, respectively.

By aligning representations across distinct granularities via this teacher–student distillation mechanism, the student model integrates fine-grained interaction patterns and high-level structural semantics, enhancing its capability to learn robust representations, especially in the sparse interaction data scenario.

Algorithm 1 Adaptive Granular-ball Generation

```
Input: Graph $\mathcal{G}$, Pruning Probability $\mathcal{P}_{ic}$ (Optional), Max Depth $D_{max}$
Output: Granular Balls $GB_{final}$ with Features $\mathbf{c}_k, r_k$
Initialize $GB_{final} \leftarrow \emptyset$
$\mathcal{P}_{init} \leftarrow$ Louvain($\mathcal{G}$) {Global coarse initialization}
For each initial ball $G_B \in \mathcal{P}_{init}$ do
    $GB_{final} \leftarrow GB_{final} \cup$ AdaptiveSplit($G_B$, 1)
End For
Construct edges between balls in $\mathcal{GB}_{final}$ with pruning probability $\mathcal{P}_{ic}$
For each ball $G_{B_k} \in GB_{final}$ do
    Compute center $\mathbf{c}_k$ and radius $r_k$ via Eq. (6)
End For
Return $GB_{final}$

Function AdaptiveSplit($GB_{curr}$, $depth$)
$Q_{parent} \leftarrow$ Quality($GB_{curr}$)
$\{S_1, \dots, S_m\} \leftarrow$ Louvain($GB_{curr}$) {Tentative split}
// Adaptive Stop: If quality drops or depth limit reached
If $\min_i(Q(S_i)) < Q_{parent}$ or $depth > D_{max}$ or no edges in $GB_{curr}$ then
    Return $\{GB_{curr}\}$ {Stop splitting and Return atomic ball}
End If
$Result \leftarrow \emptyset$
For $S_i$ in $\{S_1, \dots, S_m\}$ do
    $Result \leftarrow Result \cup$ AdaptiveSplit($S_i$, $depth + 1$)
End For
Return $Result$
End Function
```

3.4. Diffusion-enhanced granular-ball denoising

3.4.1. Social relationship granular-ball denoising

To suppress structural noise in social graphs, we perform topological denoising via granular-ball partition. Specifically, the user social graph $\mathcal{G}_u$ is partitioned into a set of structurally compact and semantically coherent granular balls, denoted as $\mathcal{GB} = \{G_{B_1}, G_{B_2}, \dots, G_{B_K}\}$. Here, each granular ball G_{B_i} (where $i \in \{1, \dots, K\}$) encapsulates a group of users exhibiting highly consistent local connectivity patterns and behavioral similarities. Given that our quality-guided generation process explicitly optimizes for structural tightness, we argue that the resulting intra-ball edges represent reliable, high-confidence dependencies driven by local homogeneity. Consequently, our denoising strategy prunes inter-ball connections to eliminate accidental noise. At the same time, it preserves the integrity of intra-ball structures.

To formalize this, we define an edge removal probability $P_{\text{ic}}(u, v)$ for each edge (u, v) as:

$$P_{\text{ic}}(u, v) = \begin{cases} 0, & \text{if } u, v \in G_{B_k}, \\ \alpha, & \text{if } u \in G_{B_i},\ v \in G_{B_j},\ i \neq j, \end{cases} \tag{10}$$

where $\alpha \in [0, 1]$ is a hyperparameter that controls the pruning strength for inter-ball connections. With this formulation, we construct a denoised adjacency matrix $\tilde{A}$ by pruning potentially noisy edges:

$$\tilde{A}_{uv} = A_{uv} \cdot (1 - P_{\text{ic}}(u, v)). \tag{11}$$

This approach preserves important intra-cluster connections while suppressing noisy inter-cluster links, thereby improving the structural quality of the social graph.

However, structural filtering alone may not fully eliminate semantic noise within learned embeddings due to inconsistent neighborhood semantics or embedding drift. To further eliminate latent semantic noise, we introduce a diffusion-based representation denoising module in the subsequent stage.

3.4.2. Diffusion-based social relation denoising

To further eliminate residual noise in user representations after structural denoising, we adopt a latent diffusion strategy for social relation denoising. Following the design of RecDiff (Z. Li et al., 2024), we implement a simplified diffusion model that injects Gaussian noise into edge embeddings and learns to reverse this process to recover denoised representations.

Let E_0 denote the initial edge embedding derived from the structurally denoised user graph. The forward diffusion process progressively corrupts E_0 over T steps as:

$$q(E_t \mid E_{t-1}) = \mathcal{N}(E_t; \sqrt{1 - \beta_t} E_{t-1}, \beta_t I), \tag{12}$$

where $\beta_t \in (0, 1)$ controls the noise level at each step t. After T steps, we obtain the noisy embedding E_T. The reverse-time generative process is defined as:

$$p_\theta(E_{t-1} \mid E_t, t) = \mathcal{N}(E_{t-1}; \mu_\theta(E_t, t), \sigma_t^2 I), \tag{13}$$

where μ_θ and σ_t^2 are predicted by a neural network conditioned on E_t and the timestep t. The denoised embedding $\tilde{E}_0$ is obtained via iterative sampling from E_T back to E_0.

To train the denoising model, we minimize the reconstruction loss between the predicted embedding and the ground truth. This diffusion-based denoising loss is defined as:

$$\mathcal{L}_{\text{DIFF}} = \mathbb{E}_{q(E_t|E_0)} \left[\left\| \hat{E}_\theta(E_1, 1) - E_0 \right\|_2^2 \right], \tag{14}$$

where $\hat{E}_\theta(E_1, 1)$ is the predicted clean embedding and E_0 is the ground-truth input at time step 0.

Table 1
Statistics of experimental datasets.

Dataset	Ciao	Yelp	Epinions
# of Users	1925	99,262	14,680
# of Items	15,053	105,142	233,261
# of User–Item Interactions	23,223	672,513	447,312
# of Social Interactions	65,084	1,298,522	632,144

Compared with traditional structure-based denoising strategies, this latent diffusion approach operates in a lower-dimensional semantic space and avoids explicit manipulation of graph topology. More importantly, it enhances the semantic fidelity of user embeddings by filtering out residual noise and mitigating heterophily effects.

The final denoised representation $\tilde{E}_0$ is incorporated into the recommendation model as a purified social signal. Together with the structural filtering stage, we ensure both topological and representational robustness, thereby improving recommendation accuracy and overall robustness.

3.5. Joint optimization with recommendation and denoising objectives

To ensure robust and accurate recommendation under noisy and sparse social environments, we formulate a unified optimization framework that combines three objectives: (i) contrastive preference learning, (ii) diffusion-based representation denoising, and (iii) cross-granularity knowledge distillation.

For preference learning, we adopt the widely used Bayesian Personalized Ranking (BPR) loss ($\mathcal{L}_{\text{BPR}}$), which optimizes pairwise rankings between positive and negative items based on user implicit feedback. On top of this foundation, we introduce two auxiliary objectives to enhance robustness.

The overall objective is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{BPR}} + \mathcal{L}_{\text{DIFF}} + \lambda \mathcal{L}_{\text{DIST}}. \tag{15}$$

where λ balances the distillation term. This joint training scheme enables GBRSR to integrate multi-scale semantics and effectively denoise both structural and representational signals, leading to improved accuracy and robustness.

4. Experimental results and analysis

In this section, we conduct extensive experiments to evaluate the effectiveness of our proposed granular-ball based robust representation learning framework. We aim to answer the following research questions:

- **RQ1.** How does our model compare to strong baselines on standard benchmarks in Recall@K and NDCG@K, and can it better balance accuracy and robustness?
- **RQ2.** What are the effects of different designed modules in our GBRSR framework on recommendation performance?
- **RQ3.** How do key hyperparameters, such as inter-cluster probabilities in granular-ball construction, affect performance and stability?
- **RQ4.** Is GBRSR sufficiently robust to handle noisy in social recommendation?
- **RQ5.** What is the computational complexity of our method in comparison to existing baseline approaches?

4.1. Experimental settings

4.1.1. Datasets

We conduct experiments on three publicly available datasets: Yelp, Ciao, and Epinions. The details of these datasets are summarized in Table 1.

- **Yelp:** Originating from the Yelp review platform, this dataset contains abundant user–item interactions in the form of ratings and reviews on local businesses such as restaurants and shops. It additionally provides social links among users, which can be exploited in social recommendation scenarios.
- **Ciao:** The Ciao dataset is derived from a product review community, where users not only rate and comment on items but also build explicit trust relationships. Such information enables the exploration of trust-aware recommendation methods.
- **Epinions:** As one of the earliest and most widely used social recommendation benchmarks, the Epinions dataset records users' product reviews and ratings while maintaining a large-scale trust network. Its relatively complex social structure makes it suitable for testing recommendation models under diverse network conditions.

4.1.2. Evaluation metrics

To evaluate the performance of each model, we ranked all candidate items for each user and adopted two widely used metrics:

- **Recall:** Recall at rank K measures the proportion of users for whom the relevant item appears in the top-K ranked items. It evaluates the model's ability to retrieve relevant items from the candidate pool.
- **Normalized Discounted Cumulative Gain (NDCG):** NDCG considers both the position and relevance of recommended items. It assigns higher scores to ground-truth items ranked higher in the top-K list, reflecting the ranking quality.

4.1.3. Baselines

To comprehensively evaluate our model, we compare it with a variety of representative social recommendation baselines across four categories.

TrustMF (Yang et al., 2017) extends traditional matrix factorization by jointly modeling users' latent preferences and trust relationships. It learns separate embeddings for trustors and trustees, thereby incorporating social trust into rating prediction.

NGCF (Wang et al., 2019) is a graph-based collaborative filtering framework that captures high-order connectivity in user–item graphs through message passing and embedding propagation.

GraphRec (Fan et al., 2019) is designed for social recommendation, integrating user–item interactions, user–user social relations, and item attributes into a unified GNN framework.

DGRec (Song et al., 2019) addresses dynamic social recommendation by applying graph attention networks to capture the temporal evolution of user preferences and social influence in session-based interactions.

SAMN (Chen et al., 2019) introduces a social attentional memory network that considers both aspect-level preferences and friend-level influences, enabling personalized recommendation via adaptive attention to different friends.

KCGN (Huang et al., 2021) couples the social graph with an external knowledge graph using a knowledge-aware GNN, thus enriching user and item embeddings with structured semantic and relational information.

SMIN (Long et al., 2021) develops a self-supervised metagraph informax framework that maximizes mutual information between users and their higher-order metagraph-based neighbors, enhancing representation learning in social recommendation.

MHCN (Yu et al., 2021) constructs multi-channel hypergraphs from users' social connections, user–item interactions, group affiliations, and applies self-supervised hypergraph convolution to capture complementary social signals.

DiffNet (Wu et al., 2019) models the social influence diffusion process in a recursive manner, where multi-step propagation effectively smooths noisy user representations. By iteratively refining embeddings through layer-wise diffusion, DiffNet reduces the impact of unreliable social links and enhances the robustness of user preference modeling.

GDMSR (Quan et al., 2023) proposes a denoising-based self-supervised learning approach for social recommendation, which reduces noise in graph-based representations and improves robustness by preference-guided denoising.

DSL (Wang et al., 2023) designs a denoised self-augmented learning framework that disentangles user preference and social influence, and employs self-supervised contrastive tasks to refine social recommendation representations.

RecDiff (Z. Li et al., 2024) introduces a diffusion-based generative framework for social recommendation, where the progressive denoising process reconstructs clean user–item representations from noisy social signals, thereby improving robustness and enhancing recommendation accuracy.

4.1.4. Implementation details

For fair comparison, we adopt experimental settings aligned with prior works (Z. Li et al., 2024). Specifically, our model is optimized with Adam using learning rates from {1e−4, 5e−4, 1e−3}. We set the embedding size to 64, use 3 LightGCN layers, and train the model with a batch size of 2048 and early stopping (patience = 10). User (item) graphs are first constructed via k-Nearest Neighbors (kNN) and then refined with recursive density-based granular-ball clustering. The diffusion step is set to 50. To ensure statistical reliability, we conduct five independent runs for main experiments and three runs for ablation studies, reporting the average results as final outcomes. All experiments are conducted on an NVIDIA RTX 4090 GPU with PyTorch 2.4.0.

4.2. Overall comparison (RQ1)

We compare the overall recommendation performance of GBRSR with baselines. Based on the results reported in Table 2, we summarize our main findings as follows.

- Graph-based social recommenders such as GraphRec generally outperform traditional collaborative filtering methods like TrustMF. This is because these methods can enhance representation learning and better capture complex semantic relations in user–item interactions by modeling high-order connectivity structures, thereby improving recommendation accuracy.
- Compared with graph-based methods, self-supervised approaches (e.g., SMIN, MHCN, DSL) often achieve stronger results. This is due to that they leverage self-supervision signals from unlabeled data to alleviate sparsity in social structures and refine user and item representations.
- Among all baselines, denoising-based methods (e.g., RecDiff, GDMSR) deliver relatively the best performance, underscoring the importance of mitigating noise in social relations. In particular, RecDiff employs a diffusion-based generative mechanism that enhances robustness against noisy edges and yields high recommendation accuracy.
- Comparing with all state-of-the-art methods, our proposed GBRSR achieves the best performance on all three datasets. For instance, on the dataset Ciao, GBRSR outperforms the best performing baseline RecDiff by 10.28% on R@10 and 12.65% on N@20. The superior performance of GBRSR is due to that it capability to better capture user preferences at different levels of granularity while reducing the influence of noisy social signals.

4.3. Ablation study (RQ2)

To further investigate the effectiveness of each key component in GBRSR, we perform ablation studies by removing its two major modules. To be specific, we consider two variants:

Table 2
Overall recommendation performance comparison between GBRSR and state-of-the-art baselines across three datasets. Results demonstrate that GBRSR consistently achieves superior or highly competitive accuracy and ranking quality in most scenarios. The best results are bolded, the second best are underlined.

Datasets	Ciao				Yelp				Epinions			
	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20
TrustMF	0.0341	0.0539	0.0289	0.0343	0.0224	0.0371	0.0149	0.0193	0.0165	0.0265	0.0163	0.0195
NGCF	0.0366	0.0559	0.0301	0.0363	0.0276	0.0450	0.0177	0.0230	0.0217	0.0353	0.0206	0.0243
GraphRec	0.0322	0.0540	0.0266	0.0335	0.0233	0.0419	0.0144	0.0201	0.0207	0.0334	0.0206	0.0246
DGRec	0.0296	0.0517	0.0254	0.0319	0.0245	0.0410	0.0158	0.0209	0.0197	0.0326	0.0194	0.0236
SAMN	0.0345	0.0604	0.0289	0.0384	0.0289	0.0403	0.0195	0.0208	0.0193	0.0329	0.0181	0.0226
KCGN	0.0360	0.0602	0.0263	0.0350	0.0284	0.0460	0.0182	0.0234	0.0221	0.0370	0.0219	0.0264
SMIN	0.0326	0.0588	0.0275	0.0354	0.0316	0.0485	0.0198	0.0251	0.0203	0.0333	0.0186	0.0228
MHCN	0.0343	0.0621	0.0286	0.0378	0.0342	0.0567	0.0225	0.0292	0.0272	0.0438	0.0272	0.0321
DiffNet	0.0328	0.0528	0.0271	0.0328	0.0381	0.0557	0.0247	0.0292	0.0238	0.0384	0.0227	0.0273
GDMSR	0.0340	0.0560	0.0276	0.0355	0.0369	0.0513	0.0196	0.0246	0.0226	0.0368	0.0206	0.0241
DSL	0.0412	0.0606	0.0329	0.0389	0.0315	0.0504	0.0203	0.0259	0.0229	0.0365	0.0226	0.0267
RecDiff	0.0457	0.0712	0.0361	0.0419	**0.0391**	0.0597	**0.0249**	0.0308	0.0282	**0.0460**	0.0275	0.0336
GBRSR (Ours)	**0.0504***	**0.0747***	**0.0391***	**0.0472***	0.0377	**0.0617***	**0.0249***	**0.0320***	**0.0289***	0.0459	**0.0293***	**0.0339***
Improv.	10.28%	4.91%	8.31%	12.65%	−3.58%	3.35%	0.00%	3.90%	2.48%	−0.22%	6.55%	0.89%

* Indicates that the improvement is statistically significant ($p < 0.01$).

Table 3
Ablation results on three datasets, illustrating the necessity of the CRD and DEGD modules for achieving robust social recommendation.

Method	Ciao		Yelp		Epinions	
	R@20	N@20	R@20	N@20	R@20	N@20
w/o CRD	0.0728	0.0423	0.0597	0.0311	0.0443	0.0329
w/o DEGD	0.0602	0.0359	0.0589	0.0298	0.0440	0.0315
GBRSR	**0.0747**	**0.0472**	**0.0617**	**0.0320**	**0.0459**	**0.0339**

- **GBRSR w/o CRD**, which replaces the cross-granularity representation distillation (CRD) module by only applying GNN over user–item graph to learn user/item-representations. It does not enhance the fine-grained user/item representations with the guidance from the coarse-grained information.
- **GBRSR w/o DEGD**, which replaces DEGD by solely remaining the graph embedding learning on the original social graph. This variant ignores both the topological denoising via granular-ball partition and social relation denoising via diffusion strategy in the diffusion-enhanced granular-ball denoising module (DEGD).

From the results shown in Table 3, we can observe that eliminating either CRD or DEGD lead to a significant drop in performance across all datasets, which validates the important roles of each component in our proposed model GBRSR. To be specific, removing CRD results in a performance drop of 2.54% and 10.38% in terms of R@20 and N@20 on the dataset Ciao, respectively. We can also observe a similar performance drop on the other two datasets. This indicates that the CRD module can effectively learn user/item representations by capturing user preference patterns across different granularities.

Similarly, without DEGD, the model becomes more vulnerable to noisy social relations, resulting in degraded recommendation performance. For instance, removing DEGD degrades the performance of the model by 19.41% and 23.94% in terms of R@20 and N@20 on the dataset Ciao, respectively. The result demonstrates the effectiveness of introducing the two denoising sub-modules, i.e., the granular-ball based structural denoising and the diffusion based representation denoising.

To verify the rationality of using the L_2 distance in Eq. (9), we investigate the impact of different distance metrics on knowledge transfer, including cosine similarity, KL divergence and L_2 distance. As illustrated in Fig. 2, the L_2 distance consistently yields better performance than both cosine similarity and KL divergence. We attribute this to the geometric nature of granular balls: minimizing the L_2 distance strictly constrains both the magnitude and direction of the node embeddings towards the semantic center. In contrast, metrics like cosine similarity and KL divergence fail to strictly constrain nodes within the Euclidean hyperspherical boundary, violating the geometric compactness assumption of the granular-ball structure. Thus, L_2 distance is both empirically and theoretically the optimal choice for maintaining precise spatial positioning.

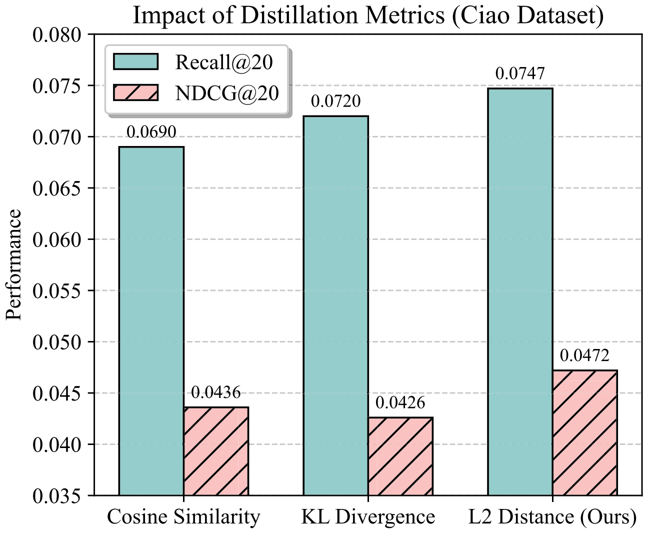

Fig. 2. Results based on different distance metrics.

4.4. Hyperparameter sensitivity analysis (RQ3)

4.4.1. Impact of inter-cluster probability P_{ic}

P_{ic} is the edge removal probability utilized for the topological denoising via granular-ball partition. To study the influence of P_{ic} to our model performance, we tune it $\in \{0.01, 0.05, 0.1, 0.3\}$ on the Ciao and Yelp datasets. The results are presented in Fig. 3. We observe that performance improves when P_{ic} increases from 0.01 to 0.05, and reaches a peak when $P_{\text{ic}} = 0.1$. It starts to decline when P_{ic} exceeds 0.1. This is because it may fail to effectively eliminate noisy edges when the threshold P_{ic} is too small. Conversely, an overly large threshold, while removing noise, may aggressively prune useful edges, resulting in the loss of structural information and further harms model performance.

4.4.2. Impact of distillation loss weight λ

The loss weight λ is utilized to control the influence of the distillation term in the overall objective, which reflects the relative importance of coarse-grained guidance in the distillation process. To investigate the

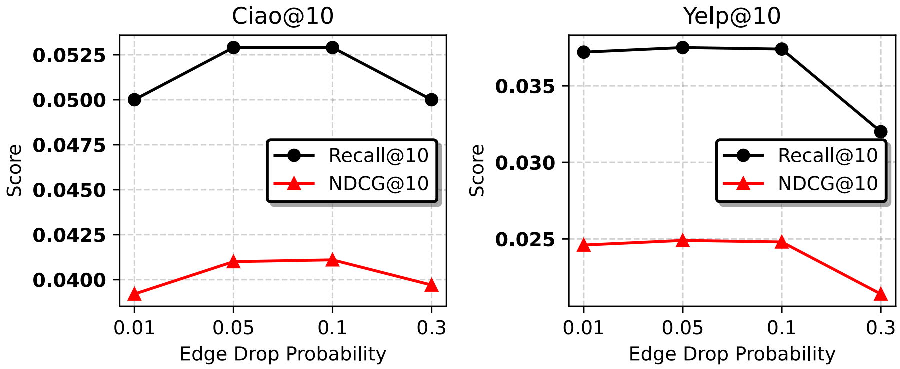

Fig. 3. The effect of P_{ic} on model performance.

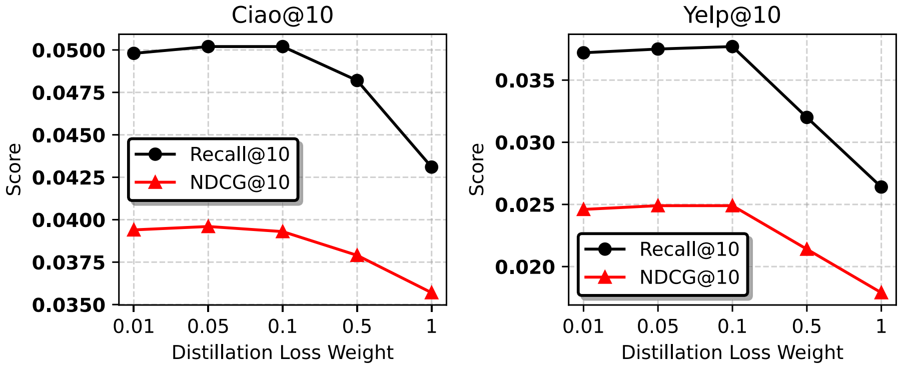

Fig. 4. The effect of λ on model performance.

influence of λ, we test $\lambda \in \{0.01, 0.05, 0.1, 0.2\}$ on Ciao and Yelp. The result is reported in Fig. 4. We can observe that with the increase of the distillation loss weight λ, the performance raises slightly first and then start to drop dramatically. The reason behind is that the student network cannot fully exploit the teacher information when λ is small, while the teacher signal overwhelms the task objective if λ becomes too large. For the sake of effectively balancing the distillation term, we set $\lambda = 0.1$ in our experiments.

4.5. Model robustness to noise (RQ4)

We investigate the robustness of our proposed model under noisy interaction scenarios. Specifically, we inject synthetic user–item interactions at four noise levels (5%, 10%, 15%, and 20%) into the training set of the **Ciao** dataset, while keeping the test set unchanged. The experimental outcomes are illustrated in Fig. 5. First, we find that when the introduced noise level is below a specific threshold, both RecDiff and our model exhibit progressively enhanced robustness with increasing noise levels. For example, the performance of both comparing methods increase slightly when the noise level under 15%. If we further increasing the noise level, we can observe a performance decline of both methods. In addition, our model consistently surpasses RecDiff across all noise levels, underscoring the stronger generalization capability of our proposed approach.

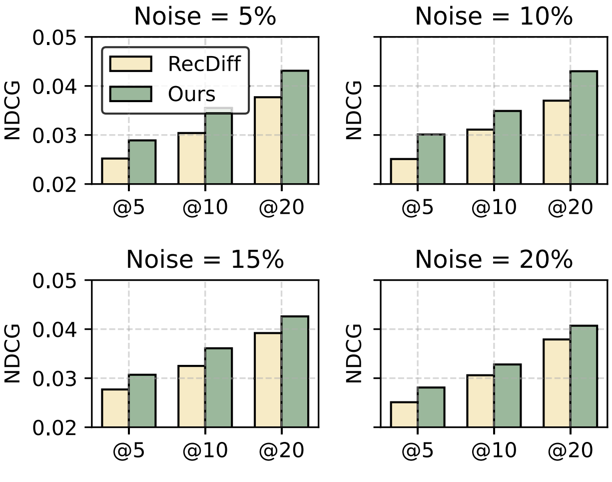

Fig. 5. Model robustness study to noise on dataset ciao.

4.6. Complexity analysis (RQ5)

To evaluate the computational efficiency and resource requirements of our framework, we compare the training memory consumption and execution time of GBRSR with several state-of-the-art methods. As reported in Table 4, GBRSR exhibits a moderate memory consumption. While there is a slight increase in memory usage compared to point-level models (e.g., SMIN, RecDiff) due to the construction of granular-ball structures, our model maintains a competitive balance between resource overhead and recommendation performance.

To further evaluate the efficiency of our training framework, we conduct additional experiments to investigate the temporal cost of different models.

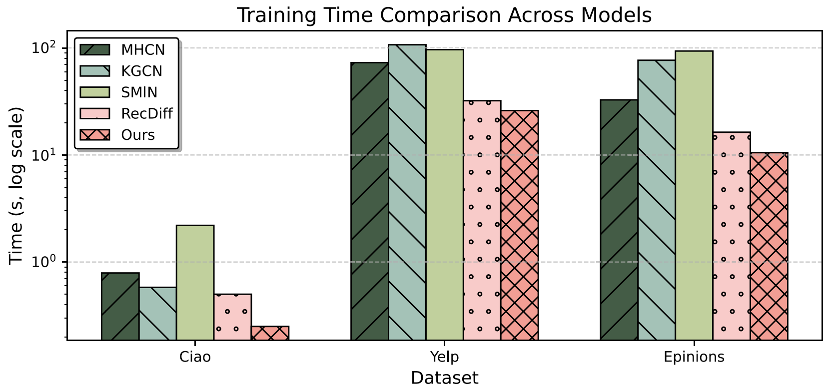

Fig. 6. Per-epoch training time comparison for evaluating computational efficiency.

Table 4
Evaluation of training memory consumption (MB) for GBRSR and competing methods. The results indicate that GBRSR maintains a competitive balance between memory efficiency and recommendation effectiveness across three datasets.

Model	Ciao (MB)	Yelp (MB)	Epinions (MB)
MHCN	1195	11,230	11,800
KCGN	1137	10,842	11,450
SMIN	1112	10,610	11,070
RecDiff	1096	10,456	11,408
GBRSR (Ours)	**1276**	**13,102**	**12,406**

Fig. 6 reports the per-epoch training and testing time across datasets on a logarithmic scale. GBRSR achieves consistently lower training and inference time compared with all baselines, demonstrating its potential scalability to large-scale recommendation scenarios. The main reason is that while existing social recommenders, such as MHCN and SMIN, leverage self-supervised augmentation, the additional mutual information maximization often increases computational burden. In contrast, GBRSR employs coarse-grained graph pruning and a diffusion-enhanced denoising mechanism, which does not incur extensive computational costs. This simple yet effective design not only improves recommendation performance but also demonstrates superior training efficiency.

5. Conclusion

In this paper, we propose GBRSR, a novel social recommendation framework designed to address two core challenges: adaptively capturing user–item relations at varying granularities, and mitigating multi-scale noise in social networks. GBRSR introduces a cross-granularity representation distillation mechanism that partitions user–item interactions into adaptive granular-balls, enabling multi-level preference modeling beyond fixed-hop constraints. In addition, we develop a diffusion-enhanced granular-ball denoising module that combines structural and latent-space denoising to suppress noise across different levels of semantic abstraction. Extensive experiments on three real-world datasets demonstrate that GBRSR consistently outperforms state-of-the-art baselines.

In the future work, we will focus on leveraging the powerful semantic reasoning capabilities of LLMs to guide the process of granular-ball generation, aiming to achieve a better alignment between high-level structural patterns and fine-grained semantic knowledge in social recommendation. Additionally, exploring more dynamic and self-evolving mechanisms tailored for temporally evolving social network is also a valuable direction for future study.

CRediT authorship contribution statement

Xiaofei Zhu: Writing – review & editing, Supervision, Project administration, Funding acquisition, Conceptualization, Methodology. **Shiyan Wu:** Writing – original draft, Software, Methodology, Investigation, Conceptualization. **Li Liu:** Writing – review & editing, Supervision. **Shuyin Xia:** Writing – review & editing, Supervision, Conceptualization, Methodology. **Yi Wang:** Supervision, Methodology, Writing – review & editing. **Guoyin Wang:** Supervision, Methodology, Conceptualization, Writing – review & editing.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This work was supported by the Science and Technology Innovation Key R&D Program of Chongqing (CSTB2024TIAD-STX0027), the National Natural Science Foundation of China (62472059), the Chongqing Talent Plan Project, China (CSTC2024YCJH-BGZXM0022), the Open Research Fund of Key Laboratory of Cyberspace Big Data Intelligent Security (Chongqing University of Posts and Telecommunications), Ministry of Education (CBDIS202403).

Data availability

Data will be made available on request.

References

Afoudi, Y., Lazaar, M., Hmaidi, S., 2023. An enhanced recommender system based on heterogeneous graph link prediction. Eng. Appl. Artif. Intell. 124, 106553.

Blondel, V., Guillaume, J.-L., Lambiotte, R., 2024. Fast unfolding of communities in large networks: 15 years later. J. Stat. Mech. Theory Exp. 2024 (10), 10R001.

Chen, X., Lei, P.I., Sheng, Y., Liu, Y., Gong, Z., 2024. Social influence learning for recommendation systems. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. pp. 312–322.

Chen, C., Zhang, M., Liu, Y., Ma, S., 2019. Social attentional memory network: Modeling aspect- and friend-level differences in recommendation. In: Proceedings of the 20th ACM International Conference on Web Search and Data Mining. pp. 177–185.

Cheng, D., Li, Y., Xia, S., Wang, G., Huang, J., Zhang, S., 2024. A fast granular-ball-based density peaks clustering algorithm for large-scale data. IEEE Trans. Neural Netw. Learn. Syst. 35 (12), 17202–17215.

Du, J., Ye, Z., Yao, L., Guo, B., Yu, Z., 2022. Socially-aware dual contrastive learning for cold-start recommendation. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 1927–1932.

Fan, W., Ma, Y., Li, Q., He, Y., Zhao, Y.E., Tang, J., Yin, D., 2019. Graph neural networks for social recommendation. In: Proceedings of the ACM on Web Conference 2019. pp. 417–426.

Huang, Y., Liang, K., Huang, Y., Zeng, X., Chen, K., Zhou, B., 2025. Social recommendation via graph-level counterfactual augmentation. In: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence. pp. 334–342.

Huang, C., Xu, H., Xu, Y., Dai, P., Xia, L., Lu, M., Bo, L., Xing, H., Lai, X., Ye, Y., 2021. Knowledge-aware coupled graph neural network for social recommendation. In: Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence. pp. 4115–4122.

Jiang, J., Zuo, Y., 2025. GREEN: graph reasoning enhanced encoder network for social intention-aware forecast of vessel navigating trajectory. Eng. Appl. Artif. Intell. 161, 112092.

Li, S., Gan, M., Xu, J., 2024. Socialcu: integrating commonalities and uniqueness of users and items for social recommendation. Proc. ACM Web Conf. 2024 27 (6), 71.

Li, Z., Xia, L., Huang, C., 2024. RecDiff: Diffusion model for social recommendation. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management. pp. 1346–1355.

Long, X., Huang, C., Xu, Y., Xu, H., Dai, P., Xia, L., Bo, L., 2021. Social recommendation with self-supervised metagraph informax network. In: Proceedings of the 30th ACM International Conference on Information and Knowledge Management. pp. 1160–1169.

Pedrycz, W., 2001. Granular computing: an introduction. In: Proceedings Joint 9th IFSA World Congress and 20th NAFIPS International Conference, vol. 3, pp. 1349–1354.

Quan, Y., Ding, J., Gao, C., Yi, L., Jin, D., Li, Y., 2023. Robust preference-guided denoising for graph based social recommendation. In: Proceedings of the ACM Web Conference 2023. pp. 1097–1108.

Shao, Y., Liu, C., 2021. H2rec: Homogeneous and heterogeneous network embedding fusion for social recommendation. Int. J. Comput. Intell. Syst. 14 (1), 1303–1314.

Sharma, K., Lee, Y., Nambi, S., Salian, A., Shah, S., Kim, S., Kumar, S., 2024. A survey of graph neural networks for social recommender systems. ACM Comput. Surv. 56 (10), 265.

Song, W., Xiao, Z., Wang, Y., Charlin, L., Zhang, M., Tang, J., 2019. Session-based social recommendation via dynamic graph attention networks. In: Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining. pp. 555–563.

Su, P., Huang, S., Ma, W., Xiong, D., Lv, J., 2025. Multi-view granular-ball contrastive clustering. In: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence. pp. 20637–20645.

Wang, G., 2017. Data-driven granular cognitive computing. In: Proceedings of the International Joint Conference 2017. pp. 13–24.

Wang, X., He, X., Wang, M., Feng, F., Chua, T., 2019. Neural graph collaborative filtering. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 165–174.

Wang, T., Xia, L., Huang, C., 2023. Denoised self-augmented learning for social recommendation. In: Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence. pp. 2324–2331.

Wu, J., Fan, W., Chen, J., Liu, S., Li, Q., Tang, K., 2022. Disentangled contrastive learning for social recommendation. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. pp. 4570–4574.

Wu, L., Sun, P., Fu, Y., Hong, R., Wang, X., Wang, M., 2019. A neural influence diffusion model for social recommendation. In: Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 235–244.

Wu, S., Sun, F., Zhang, W., Xie, X., Cui, B., 2023. Graph neural networks in recommender systems: A survey. ACM Comput. Surv. 55 (5), 97:1–97:37.

Xia, S., Liu, Y., Ding, X., Wang, G., Yu, H., Luo, Y., 2019. Granular ball computing classifiers for efficient, scalable and robust learning. Inf. Sci. 483, 136–152.

Xia, S., Ma, X., Liu, Z., Liu, C., Zhao, S., Wang, G., 2025a. Graph coarsening via supervised granular-ball for scalable graph neural network training. In: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence. pp. 12872–12880.

Xia, S., Wang, G., Gao, X., Lian, X., 2023. Granular-ball computing: an efficient, robust, and interpretable adaptive multi-granularity representation and computation method. ArXiv E-Prints. arXiv:2304.11171.

Xia, S., Wang, G., Xu, G., Zhao, S., Wang, G., 2025b. GBGC: efficient and adaptive graph coarsening via granular-ball computing. In: Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence. pp. 3489–3497.

Xie, J., Cheng, Y., Xia, S., Hua, C., Wang, G., Gao, X., 2025. AW-GBGAE: an adaptive weighted graph autoencoder based on granular-balls for general data clustering. IEEE Trans. Pattern Anal. Mach. Intell. 47 (12), 11298–11314.

Xie, J., Dai, M., Xia, S., Zhang, J., Wang, G., Gao, X., 2024. An efficient fuzzy stream clustering method based on granular-ball structure. In: Proceedings of the 40th IEEE International Conference on Data Engineering. pp. 901–913.

Xiong, F., Sun, H., Luo, G.X., Pan, S., Qiu, M., Wang, L., 2024. Graph attention network with high-order neighbor information propagation for social recommendation. In: Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence. pp. 2478–2486.

Xiong, F., Zhang, T., Pan, S., Luo, G., Wang, L., 2025. Robust graph based social recommendation through contrastive multi-view learning. In: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence. pp. 12890–12898.

Xu, J., Chen, Z., Yang, S., Li, J., Wang, H., Ngai, E.C.H., 2025. MENTOR: multi-level self-supervised learning for multimodal recommendation. In: Proceedings of the Thirty-Ninth AAAI Conference on Artificial Intelligence. pp. 12908–12917.

Yang, B., Lei, Y., Liu, J., Li, W., 2017. Social collaborative filtering by trust. IEEE Trans. Pattern Anal. Mach. Intell. 39 (8), 1633–1647.

Yang, J., Xiaodiao, L., Wang, G., Pedrycz, W., Xia, S., Zhang, Q., Wu, D., 2025. A robust three-way classifier with shadowed granular balls based on justifiable granularity. IEEE Trans. Neural Netw. Learn. Syst. 36 (9), 16534–16548.

Yu, J., Yin, H., Li, J., Wang, Q., Hung, N.Q.V., Zhang, X., 2021. Self-supervised multi-channel hypergraph convolutional network for social recommendation. In: Proceedings of the ACM on Web Conference 2021. pp. 413–424.

Zhang, X., Xu, L., Jia, W., 2025. Adaptive granular-ball based density peak clustering. Neurocomputing 656, 131458.