

A Contrastive Multi-Decomposition Approach with Social-Guided Signals for Robust Recommender Systems

Mohammad Aliabadi

University of Kurdistan

Alireza Abdollahpouri

`abdollahpouri@uok.ac.ir`

University of Kurdistan

Parham Moradi

Victoria University

Research Article

Keywords: Recommender Systems, Graph Neural Networks, Contrastive Learning, Low-Rank Graph Decomposition, Social Recommendation, Robust Recommendation, Multi-Resolution Representation Learning

Posted Date: August 12th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10454737/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at International Journal of Data Science and Analytics on October 22nd, 2026. See the published version at <https://doi.org/10.1007/s41060-026-01320-5>.

A Contrastive Multi-Decomposition Approach with Social-Guided Signals for Robust Recommender Systems

Mohammad Aliabadi¹, Alireza Abdollahpouri^{1*}, Parham Moradi²

¹Department of Computer Engineering, University of Kurdistan,
Sanandaj, Iran.

²School of Engineering, Victoria University, Melbourne, Australia.

*Corresponding author(s). E-mail(s): abdollahpouri@uok.ac.ir;

Abstract

Graph neural networks (GNNs) have achieved promising results in recommender systems, yet their performance often degrades in the presence of noisy implicit feedback and sparse interactions. Existing approaches commonly represent the user-item graph as a single structure, which limits their ability to distinguish stable collaborative patterns from unreliable local connections. This paper presents a Contrastive Multi-Decomposition Approach with Social-Guided Signals (COMPASS), which combines multi-resolution low-rank decomposition with social-guided contrastive learning. Instead of relying on stochastic graph augmentations, COMPASS constructs multiple views from rank-truncated singular value decomposition, where lower-rank components capture global collaborative information and higher-rank components preserve finer interaction details. These views are integrated through a Gumbel-Softmax gating mechanism that learns the relative importance of different approximation levels. Social relations are incorporated as a soft contrastive regularizer to encourage socially consistent representations. In addition, an adaptive gradient-balanced weighting strategy is adopted to improve the stability of multi-objective optimization. Experiments conducted on Yelp, LastFM, and Douban show that COMPASS consistently achieves competitive improvements over strong baselines, with gains of up to 5.76% in NDCG@20. Additional noise injection experiments indicate that the proposed framework maintains stable performance under corrupted interactions, suggesting favorable robustness to noisy feedback.

Keywords: Recommender Systems, Graph Neural Networks, Contrastive Learning, Low-Rank Graph Decomposition, Social Recommendation, Robust Recommendation, Multi-Resolution Representation Learning

1 Introduction

Graph Neural Networks (GNNs) have become an effective paradigm for collaborative filtering by modeling high-order relationships between users and items Wang et al. (2019); He et al. (2020). Despite their success, many GNN-based recommenders remain sensitive to the quality of the interaction graph. Real-world feedback is often noisy and sparse, and incidental interactions may introduce unreliable edges that affect message propagation Wu et al. (2021); Yu et al. (2022). Since GNNs mainly emphasize low-frequency information, distinguishing noisy signals from genuine preferences becomes increasingly difficult when the interaction data are sparse.

Several approaches have attempted to alleviate this issue through contrastive learning Wu et al. (2021); Yu et al. (2022) or by incorporating auxiliary social information Wu et al. (2019); Fan et al. (2019). However, most existing methods process structural information at a single level of granularity. For example, random edge perturbations used in frameworks such as SGL may remove informative connections together with noisy ones. Likewise, existing SVD-based approaches typically rely on a single fixed-rank approximation, which may not capture the diverse structural requirements of different users. Users with limited interaction histories tend to rely more on global collaborative patterns, while users with richer histories may benefit from finer structural details that provide additional preference information.

Motivated by this observation, we propose Contrastive Multi-Decomposition Approach with Social-Guided Signals (COMPASS), a framework designed to capture user preferences across multiple structural resolutions. Rather than relying on stochastic graph perturbations, COMPASS constructs multiple views from rank-truncated singular value decompositions of the interaction matrix. Lower-rank approximations capture dominant collaborative patterns, while higher-rank components retain finer interaction details. This design is supported by the classical low-rank approximation theorem, which states that rank- q truncation provides the optimal approximation under the Frobenius norm. A Gumbel-Softmax gating mechanism then learns how to combine these complementary views according to user characteristics. In addition, social relations are incorporated as a soft contrastive signal to encourage consistency between social homophily and collaborative preferences. To improve training stability, adaptive gradient-balanced loss weighting is adopted to coordinate the recommendation objective and auxiliary contrastive learning.

The main contributions of this work are summarized as follows:

- **Multi-resolution graph decomposition.** A multi-view framework based on hierarchical SVD truncation is developed to represent interactions at different approximation levels. The resulting views capture both dominant collaborative patterns and finer residual structures, allowing the model to preserve community-level information

while maintaining personalized preferences.

- **Data-driven view fusion.** A Gumbel-Softmax gating mechanism is introduced to learn different rank preferences for sparse and active users. The fusion process adaptively balances global and local information while introducing only $3|\mathcal{Q}|$ additional parameters.
- **Adaptive optimization.** A gradient-balanced learning strategy with bounded adaptive loss weights is employed to improve the stability of joint optimization between the recommendation task and auxiliary contrastive objectives.
- **Robustness evaluation.** Extensive experiments, including noise injection and sparsity analysis, are conducted to assess the behavior of the proposed framework under challenging conditions. The results indicate consistent improvements over competitive baselines and show stable performance in the presence of corrupted interactions.

2 Related Work

Graph neural networks have become a central paradigm for collaborative filtering because they effectively model high-order relationships in user-item graphs Wang et al. (2019); He et al. (2020); Wu et al. (2022). Beyond early graph convolutional models such as GCMC and NGCF Wang et al. (2019); Berg et al. (2017), recent studies have explored richer structures and auxiliary information, including node-personalized graph representations Zhou et al. (2024), multimodal mutual learning Li et al. (2024), causal collaborative filtering Wang et al. (2024), higher-order hypergraph architectures Sakong et al.

(2024), and multi-aspect graph learning Zhang et al. (2023). In parallel, contrastive learning has emerged as an effective strategy for improving representation quality and robustness. Representative approaches include SGL Wu et al. (2021), SimGCL Yu et al. (2022), XSimGCL Yu et al. (2023), neighborhood-enriched contrastive learning Lin et al. (2022), negative sampling strategies Miao et al. (2022), denoised graph contrastive learning Xu et al. (2024), semantic prototype learning Wen et al. (2025), multimodal contrastive knowledge graph learning Wei et al. (2024), and multi-level graph contrastive learning for cold-start recommendation Sun et al. (2025). Multi-view learning has further been investigated through heterogeneous graph structures and semantic views Wang et al. (2020); Yu et al. (2021); Won et al. (2023); Yang et al. (2025); Tao et al. (2025); He et al. (2025). Despite these advances, most methods rely on predefined views or stochastic perturbations and typically construct representations at a single structural granularity.

Social information has long been employed to alleviate sparsity and improve recommendation quality Ma et al. (2011); Tang et al. (2013). Recent graph-based social recommenders, including GraphRec, DANSER, and DiffNet++, integrate interaction and social graphs through attention and diffusion mechanisms Fan et al. (2019); Wu et al. (2019, 2020). More recent studies have recognized that social links themselves may be noisy or weakly aligned with preference similarity. Existing solutions address this issue through denoised self-augmentation Wang et al. (2023), homophily-aware reconstruction Jiang et al. (2024), bottleneck regularization Yang et al. (2024), preference-guided

denoising Quan et al. (2023), diffusion-based social modeling Li et al. (2024), and contrastive multi-view frameworks Xiong et al. (2025); Meng et al. (2025). Although these approaches improve robustness, social information is commonly coupled with fixed propagation architectures and is rarely used as soft semantic guidance.

Recent studies have also explored deterministic and spectral perspectives for recommendation. LightGCL replaces stochastic graph augmentation with a global SVD-derived collaborative view Cai et al. (2023), while frequency-domain methods exploit spectral representations to enhance robustness and capture latent collaborative patterns Shirali et al. (2024). Related developments in spectral graph analysis have shown that graph signals naturally contain information at different frequencies Von Luxburg (2007); Hammond et al. (2011). However, existing spectral recommendation models mainly rely on a single global representation and do not explicitly characterize complementary structures across multiple approximation levels.

Overall, existing methods still face three limitations. First, most approaches construct a single collaborative representation or rely on manually designed views, limiting their ability to capture structural information at different resolutions. Second, social information is typically incorporated through dedicated propagation mechanisms rather than being utilized as soft guidance for representation learning. Third, robustness and self-supervised objectives are often sensitive to manually selected augmentations and balancing strategies. These observations motivate the development of COMPASS, which combines multi-resolution decomposition, adaptive view fusion, and socially guided contrastive learning within a unified framework.

3 Methodology

This section presents COMPASS, a multi-view framework designed to improve robustness under noisy and sparse interaction data. Instead of relying on stochastic graph perturbations, COMPASS constructs complementary representations through multi-resolution low-rank decomposition and integrates collaborative and social signals within a unified learning framework. The proposed framework consists of four main components:

- **Multi-Resolution View Generation:** Multiple rank-specific SVD approximations of the interaction matrix are used to construct complementary views. Lower-rank components emphasize dominant collaborative patterns, whereas higher-rank components retain finer interaction details, resulting in a hierarchy of deterministic representations across different approximation levels.
- **Data-Driven View Fusion:** A Gumbel-Softmax gating mechanism learns separate rank distributions for sparse and active users. This design allows the model to place more emphasis on global collaborative information for users with limited interactions while incorporating finer structural details for users with richer histories, using only $3|\mathcal{Q}|$ additional scalar parameters.
- **Social-Guided Contrastive Regularization:** Social relations are incorporated as a soft contrastive signal that encourages consistency between collaborative and social representations without imposing rigid constraints on the learning process.
- **Adaptive Gradient-Balanced Optimization:** To improve the stability of multi-objective learning, COMPASS employs a bounded learnable weighting

scheme together with a dual-optimizer strategy to coordinate the BPR and InfoNCE objectives during training.

During inference, only the computationally efficient collaborative branch is retained, which keeps the model practical for large-scale recommendation scenarios. The overall architecture and the interactions among its components are illustrated in Fig. 1.

3.1 Preliminaries

This subsection introduces the problem setting and the notation used throughout the paper. We first define the collaborative and social graphs, followed by the recommendation objective and the embedding representations used in the proposed framework. Table 1 summarizes the main symbols and variables.

3.1.1 Graph Structures and Interaction Matrices

Let $\mathcal{U} = \{u_1, u_2, \dots, u_m\}$ denote the set of users, where $m = |\mathcal{U}|$, and let $\mathcal{I} = \{i_1, i_2, \dots, i_n\}$ denote the set of items, where $n = |\mathcal{I}|$. User-item interactions are modeled as a bipartite graph $\mathcal{G}_c = (\mathcal{U}, \mathcal{I}, \mathcal{E}_c)$, referred to as the collaborative graph, where $\mathcal{E}_c \subseteq \mathcal{U} \times \mathcal{I}$ denotes the set of observed interactions.

The interactions are represented by an adjacency matrix $\mathbf{A} \in \{0, 1\}^{m \times n}$, where $\mathbf{A}_{ij} = 1$ indicates that user u_i has interacted with item i_j , and $\mathbf{A}_{ij} = 0$ otherwise.

We also consider a social graph $\mathcal{G}_s = (\mathcal{U}, \mathcal{E}_s)$, where $\mathcal{E}_s \subseteq \mathcal{U} \times \mathcal{U}$ represents social relations among users. Its adjacency matrix is denoted by $\mathbf{S} \in \{0, 1\}^{m \times m}$, where $\mathbf{S}_{ij} = 1$ if users u_i and u_j are socially connected, and $\mathbf{S}_{ij} = 0$ otherwise.

3.1.2 Recommendation Task

The goal is to learn a recommendation function that maps user-item pairs to real-valued preference scores for ranking. Formally, the recommendation model is defined by:

$$f : \mathcal{U} \times \mathcal{I} \rightarrow \mathbb{R}, \quad (1)$$

where, $f(u_i, i_j)$ represents the predicted preference score of user u_i for item i_j .

The scoring function is parameterized by the learnable parameters Θ and depends on both the collaborative graph $\mathcal{G}_c$ and the social graph $\mathcal{G}_s$. Accordingly, the predicted score for a user-item pair is given by:

$$\hat{y}_{ij} = f(u_i, i_j \mid \mathcal{G}_c, \mathcal{G}_s, \Theta), \quad (2)$$

where, $\hat{y}_{ij}$ denotes the estimated preference of user u_i for item i_j . During training, the model learns to assign higher scores to observed interactions than to unobserved ones. At inference time, candidate items are ranked according to $\hat{y}_{ij}$ to generate personalized top- N recommendations.

3.1.3 Input and Embedding Dimensions

The collaborative adjacency matrix $\mathbf{A}$ has dimensions $m \times n$, whereas the social adjacency matrix $\mathbf{S}$ has dimensions $m \times m$. User and item representations are initialized by the embedding matrices $\mathbf{E}_u^{(0)} \in \mathbb{R}^{m \times d}$ and $\mathbf{E}_i^{(0)} \in \mathbb{R}^{n \times d}$, respectively, where d denotes the embedding dimension. Subsequent representations maintain compatible dimensions, which enables consistent operations across different components of the framework.

Fig. 1 presents the collaborative view, the adaptive multi-resolution decomposition view, the social view, and the joint

Table 1. Summary of the main symbols used throughout the paper.

Notation	Description
Sets and graphs	
$\mathcal{U}, \mathcal{I}$	Sets of users and items.
m, n	Number of users and items.
$\mathcal{G}_c, \mathcal{G}_s$	Collaborative and social graphs.
$\mathcal{E}_c, \mathcal{E}_s$	Edge sets of the collaborative and social graphs.
$\mathcal{O}$	Set of training triplets (u, i, j) used in BPR.
$\mathcal{Q}$	Set of rank values used for SVD.
Scalars and parameters	
d	Embedding dimension.
L	Number of GNN layers.
l	Layer index, where $l \in \{0, \dots, L\}$.
q	A rank value from $\mathcal{Q}$.
λ	Coefficient of ℓ_2 regularization.
τ_c	Temperature parameter of InfoNCE.
τ_g	Temperature parameter of Gumbel-Softmax.
p_s	Social edge dropout probability.
B	Mini-batch size.
$\sigma(\cdot)$	Sigmoid activation function.
Θ	Set of learnable model parameters.
Matrices and embeddings	
$\mathbf{A} \in \{0, 1\}^{m \times n}$	User-item interaction matrix.
$\mathbf{S} \in \{0, 1\}^{m \times m}$	User-user social adjacency matrix.
$\hat{\mathbf{A}}, \hat{\mathbf{S}}$	Symmetrically normalized adjacency matrices.
$\mathbf{E}_u, \mathbf{E}_i$	Final user and item embeddings from the collaborative view.
$\mathbf{e}_u, \mathbf{e}_i$	Embedding vectors of user u and item i .
$\mathbf{G}_u, \mathbf{G}_i$	Final user and item embeddings from the multi-resolution view.
$\mathbf{S}_u$	Final user embeddings from the social view.
$\mathcal{U}_{\text{sparse}}, \mathcal{U}_{\text{active}}$	User groups partitioned according to interaction frequency.
$\boldsymbol{\pi}^{\hat{\mathbf{S}}}, \boldsymbol{\pi}^{\mathbf{A}} \in \mathbb{R}^{ \mathcal{Q} }$	Gumbel-Softmax logits for sparse and active user groups.
$\boldsymbol{\pi}^{\mathbf{I}} \in \mathbb{R}^{ \mathcal{Q} }$	Gumbel-Softmax logits for item-side fusion.
y_q^u	Fusion weight associated with user group u and rank q .
$\hat{\mathbf{A}}_q \in \mathbb{R}^{m \times n}$	Rank- q approximation of $\hat{\mathbf{A}}$.
$\bar{y}_q$	Average fusion weight for rank q used in item embeddings.
$\mathbf{E}^{(l)}, \mathbf{G}^{(l,q)}, \mathbf{S}^{(l)}$	Layer-wise embeddings for the collaborative, multi-resolution, and social views.
Loss functions	
$\mathcal{L}_{BPR}$	Bayesian Personalized Ranking loss.
$\mathcal{L}_{InfoNCE}$	Contrastive InfoNCE loss.
$\mathcal{L}_{total}$	Overall training objective.

optimization strategy. During inference, which keeps the model computationally only the collaborative branch is retained, efficient.




Fig. 1. Overview of the COMPASS architecture.

3.2 Collaborative Graph Learning

The collaborative view provides the primary preference signal from the user-item interaction graph $\mathcal{G}_c$. This component adopts the propagation scheme of LightGCN He et al. (2020).

To alleviate degree-related bias, the interaction matrix is symmetrically normalized. The normalized adjacency matrix is given by Eq. (3):

$$\hat{\mathbf{A}} = \mathbf{D}_u^{-1/2} \mathbf{A} \mathbf{D}_i^{-1/2} \quad (3)$$

where, $\mathbf{A} \in \{0,1\}^{m \times n}$ denotes the interaction matrix, and $\mathbf{D}_u \in \mathbb{R}^{m \times m}$ and $\mathbf{D}_i \in \mathbb{R}^{n \times n}$ are diagonal degree matrices Kipf and Welling (2016).

Starting from the trainable embeddings $\mathbf{E}_u^{(0)} \in \mathbb{R}^{m \times d}$ and $\mathbf{E}_i^{(0)} \in \mathbb{R}^{n \times d}$,

layer-wise propagation iteratively updates user and item representations. For each layer $l \in \{1, \dots, L\}$, the propagation process is expressed by Eq. (4):

$$\mathbf{E}_u^{(l)} = \hat{\mathbf{A}} \mathbf{E}_i^{(l-1)}, \quad \mathbf{E}_i^{(l)} = \hat{\mathbf{A}}^\top \mathbf{E}_u^{(l-1)}. \quad (4)$$

Representations obtained at different depths are then aggregated to capture information from multiple propagation layers. The final embeddings are computed according to Eq. (5):

$$\mathbf{E}_u = \sum_{l=0}^L \alpha_l \mathbf{E}_u^{(l)} \quad \mathbf{E}_i = \sum_{l=0}^L \alpha_l \mathbf{E}_i^{(l)} \quad (5)$$

Following LightGCN He et al. (2020), uniform weights with $\alpha_l = 1/(L+1)$ are

adopted in all experiments. The resulting embeddings are used in the recommendation objective and are subsequently involved in cross-view contrastive learning.

3.3 Multi-Resolution Low-Rank Decomposition

Although the collaborative view captures the main interaction patterns, it represents the graph at a single structural resolution and treats all connections uniformly. To characterize user-item relationships at different levels of granularity, COMPASS incorporates a multi-resolution low-rank decomposition module. In contrast to stochastic graph perturbations such as edge dropout, this component adopts a deterministic formulation based on classical matrix approximation theory [Eckart and Young \(1936\)](#); [Mirsky \(1960\)](#).

For a given rank q , the truncated singular value decomposition yields a rank- q approximation $\hat{\mathbf{A}}_q = \mathbf{U}_q \mathbf{\Sigma}_q \mathbf{V}_q^\top$, which preserves the most informative components of the normalized interaction matrix. As the rank increases, additional structural information is gradually incorporated into the approximation. Lower-rank components emphasize dominant collaborative patterns, whereas higher-rank components retain finer interaction details that may capture more localized preferences. Consequently, multiple truncation levels provide complementary views of the interaction graph and offer a deterministic alternative to stochastic graph augmentation.

3.3.1 Multi-Rank SVD Approximation Hierarchy

Consider the singular value decomposition of the normalized interaction matrix $\hat{\mathbf{A}}$

defined in Eq. (3). The decomposition can be expressed as:

$$\hat{\mathbf{A}} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^\top \quad (6)$$

where, $\mathbf{U} \in \mathbb{R}^{m \times m}$ and $\mathbf{V} \in \mathbb{R}^{n \times n}$ contain the left and right singular vectors, respectively, and $\mathbf{\Sigma}$ is the diagonal matrix of singular values.

For a given rank q , retaining the leading singular components yields the truncated approximation

$$\hat{\mathbf{A}} \approx \hat{\mathbf{A}}_q = \mathbf{U}_q \mathbf{\Sigma}_q \mathbf{V}_q^\top \quad (7)$$

which provides the optimal rank- q approximation under the Frobenius norm [Eckart and Young \(1936\)](#); [Mirsky \(1960\)](#). Here, $\mathbf{U}_q \in \mathbb{R}^{m \times q}$, $\mathbf{\Sigma}_q \in \mathbb{R}^{q \times q}$, and $\mathbf{V}_q \in \mathbb{R}^{n \times q}$. Lower-rank approximations emphasize dominant collaborative patterns, whereas higher-rank components retain finer interaction details. Consequently, computing truncated decompositions for multiple values of $q \in \mathcal{Q}$ produces complementary views that characterize the interaction graph at different structural resolutions. The choice of $\mathcal{Q}$ is discussed in Section 5.4.

For each rank q , message propagation is performed through the corresponding low-rank approximation. The layer-wise updates are given by:

$$\begin{aligned} \mathbf{G}_u^{(l,q)} &= (\mathbf{U}_q \mathbf{\Sigma}_q) \left(\mathbf{V}_q^\top \mathbf{E}_i^{(l-1)} \right), \\ \mathbf{G}_i^{(l,q)} &= (\mathbf{V}_q \mathbf{\Sigma}_q) \left(\mathbf{U}_q^\top \mathbf{E}_u^{(l-1)} \right). \end{aligned} \quad (8)$$

where, $\mathbf{G}_u^{(l,q)} \in \mathbb{R}^{m \times d}$ and $\mathbf{G}_i^{(l,q)} \in \mathbb{R}^{n \times d}$ denote the user and item representations obtained from the rank- q view at layer l .

Representations from different propagation layers are then aggregated to obtain the final embeddings associated with rank q :

$$\begin{aligned}\mathbf{G}_u^{(q)} &= \frac{1}{L} \sum_{l=1}^L \mathbf{G}_u^{(l,q)}, \\ \mathbf{G}_i^{(q)} &= \frac{1}{L} \sum_{l=1}^L \mathbf{G}_i^{(l,q)}.\end{aligned}\tag{9}$$

which provide a collection of multi-resolution representations for subsequent view fusion.

3.3.2 Activity-Aware Data-Driven Fusion

The rank-specific embeddings are combined through a learnable fusion mechanism. Instead of using uniform averaging or manually selected weights, COMPASS learns separate rank distributions for sparse and active users.

Users are partitioned according to their interaction counts relative to the median number of training interactions $\tilde{n}$. Users satisfying $|\mathcal{N}_u| \geq \tilde{n}$ are assigned to the active group $\mathcal{G}_A$, while the remaining users belong to the sparse group $\mathcal{G}_S$.

Each group is associated with a learnable logit vector, denoted by $\boldsymbol{\pi}^S \in \mathbb{R}^{|\mathcal{Q}|}$ and $\boldsymbol{\pi}^A \in \mathbb{R}^{|\mathcal{Q}|}$. For a user u , the corresponding logits are selected according to:

$$\boldsymbol{\pi}^u = \begin{cases} \boldsymbol{\pi}^S, & \text{if } u \in \mathcal{G}_S, \\ \boldsymbol{\pi}^A, & \text{if } u \in \mathcal{G}_A. \end{cases}\tag{10}$$

where, π_q^u denotes the logit associated with rank q for the activity group of user u . The fusion weights are then computed using the Gumbel-Softmax distribution:

$$y_q^u = \frac{\exp((\log \pi_q^u + g_q)/\tau_g)}{\sum_{k \in \mathcal{Q}} \exp((\log \pi_k^u + g_k)/\tau_g)}\tag{11}$$

where, g_q are i.i.d. Gumbel(0, 1) samples and τ_g is the temperature parameter. During training, τ_g is gradually annealed

and clipped to $\tau_{\min} = 0.5$ Jang et al. (2016).

Since the logits are shared within each activity group, users in $\mathcal{G}_S$ and $\mathcal{G}_A$ share their corresponding fusion distributions. Owing to the softmax normalization, the resulting weights satisfy $\sum_{q \in \mathcal{Q}} y_q^S = 1$ and $\sum_{q \in \mathcal{Q}} y_q^A = 1$.

Items are shared across all users and therefore do not require activity-specific grouping. A separate learnable logit vector $\boldsymbol{\pi}^I \in \mathbb{R}^{|\mathcal{Q}|}$ is introduced for item-side fusion. The corresponding weights are computed using the same Gumbel-Softmax formulation:

$$\bar{y}_q = \frac{\exp((\pi_q^I + g_q)/\tau_g)}{\sum_{q' \in \mathcal{Q}} \exp((\pi_{q'}^I + g_{q'})/\tau_g)}, \quad q \in \mathcal{Q}.\tag{12}$$

Using these weights, the final multi-resolution embeddings are obtained as:

$$\mathbf{G}_u = \sum_{q \in \mathcal{Q}} y_q^{g(u)} \mathbf{G}_u^{(q)}, \quad \mathbf{G}_i = \sum_{q \in \mathcal{Q}} \bar{y}_q \mathbf{G}_i^{(q)},\tag{13}$$

where, $g(u) \in \{S, A\}$ denotes the activity group of user u . Since the fusion weights are normalized, both $\mathbf{G}_u$ and $\mathbf{G}_i$ are convex combinations of the rank-specific representations. The fusion module introduces only $3|\mathcal{Q}|$ additional scalar parameters.

This design allows different activity groups to place different levels of emphasis on the available structural resolutions. Users with limited interaction histories tend to rely more on dominant collaborative patterns, whereas users with richer histories may benefit from finer structural details. An analysis of the learned fusion weights is provided in Section 5.4.2.

3.4 Social Trust Integration

To enrich user representations beyond direct item interactions, COMPASS incorporates a social view that exploits social links between users. This auxiliary information helps alleviate cold-start effects and provides preference signals that may not be fully observable from interaction data alone [Tang et al. \(2013\)](#). Reciprocal links are considered because mutual connections often correspond to stronger and more reliable social relations [Ma et al. \(2011\)](#); [Yang et al. \(2012\)](#).

A symmetric social adjacency matrix $\mathbf{S} \in \{0, 1\}^{m \times m}$ is constructed from $\mathcal{G}_s$. To reduce the influence of high-degree users, symmetric normalization is applied:

$$\hat{\mathbf{S}} = \mathbf{D}_s^{-1/2} \mathbf{S} \mathbf{D}_s^{-1/2} \quad (14)$$

where, $\mathbf{D}_s$ denotes the diagonal degree matrix [Fan et al. \(2019\)](#).

To maintain a shared representation space across different views, social user embeddings are initialized using the collaborative embeddings:

$$\mathbf{S}_u^{(0)} = \mathbf{E}_u^{(0)} \quad (15)$$

At each layer l , social embeddings are updated through neighborhood aggregation, with edge dropout applied to the normalized social graph:

$$\mathbf{S}_u^{(l)} = \text{EdgeDrop}(\hat{\mathbf{S}}, p_s) \mathbf{S}_u^{(l-1)} \quad (16)$$

Here, $\text{EdgeDrop}(\hat{\mathbf{S}}, p_s)$ randomly removes individual nonzero entries of $\hat{\mathbf{S}}$ with probability p_s during training while preserving the node set. The final social embeddings $\mathbf{S}_u$ are obtained by aggregating layer representations in a manner analogous to Eq. (5).

The collaborative embeddings $\mathbf{e}_u$ and social embeddings $\mathbf{s}_u$ are derived from different graph structures. Their alignment

is motivated by the principle of social homophily, which suggests that socially connected users often exhibit similar characteristics and preferences [Fan et al. \(2019\)](#). In this setting, the collaborative representation reflects preferences inferred from item interactions, whereas the social representation summarizes information propagated from reciprocal neighbors.

The two views are aligned through a contrastive objective, where social information acts as a soft regularizer rather than a hard constraint. As described in Section 3.6, the contribution of the social contrastive term is controlled by a bounded learnable weight, allowing the model to reduce the influence of noisy or weakly aligned neighborhoods when necessary.

3.5 Multi-Objective Loss Formulation

Training combines a recommendation objective with auxiliary contrastive learning. The recommendation component is based on Bayesian Personalized Ranking (BPR), where each triplet (u, i, j) consists of a user u , an observed item i , and an unobserved item j . The BPR loss is given by Eq. (17):

$$\mathcal{L}_{BPR} = - \sum_{(u, i, j) \in \mathcal{O}} \ln \sigma(\mathbf{e}_u^\top \mathbf{e}_i - \mathbf{e}_u^\top \mathbf{e}_j) \quad (17)$$

where, $\mathcal{O}$ denotes the set of training triplets, $\mathbf{e}_u$ is the final collaborative embedding of user u , and $\mathbf{e}_i$ and $\mathbf{e}_j$ denote the final collaborative embeddings of items i and j , respectively.

To encourage consistency across different views, contrastive learning is applied between the collaborative and multi-resolution representations, as well as

between the collaborative and social representations. Specifically, user embeddings are aligned between the collaborative and multi-resolution views and between the collaborative and social views, while item embeddings are aligned between the collaborative and multi-resolution views. The corresponding objectives are defined by Eqs. (18)-(20):

$$\mathcal{L}_{u,mg} = - \sum_{u \in \mathcal{U}} \log \frac{\exp(s(\mathbf{e}_u, \mathbf{g}_u)/\tau_c)}{\sum_{v \in \mathcal{U}} \exp(s(\mathbf{e}_u, \mathbf{g}_v)/\tau_c)} \quad (18)$$

$$\mathcal{L}_{i,mg} = - \sum_{i \in \mathcal{I}} \log \frac{\exp(s(\mathbf{e}_i, \mathbf{g}_i)/\tau_c)}{\sum_{j \in \mathcal{I}} \exp(s(\mathbf{e}_i, \mathbf{g}_j)/\tau_c)} \quad (19)$$

$$\mathcal{L}_{u,s} = - \sum_{u \in \mathcal{U}} \log \frac{\exp(s(\mathbf{e}_u, \mathbf{s}_u)/\tau_c)}{\sum_{v \in \mathcal{U}} \exp(s(\mathbf{e}_u, \mathbf{s}_v)/\tau_c)} \quad (20)$$

where, $s(\cdot, \cdot)$ denotes cosine similarity and τ_c is the temperature parameter. The overall contrastive objective is obtained by summing the three components:

$$\mathcal{L}_{InfoNCE} = \mathcal{L}_{u,mg} + \mathcal{L}_{i,mg} + \mathcal{L}_{u,s} \quad (21)$$

3.6 Gradient-Balanced Joint Optimization

The recommendation and contrastive objectives introduced above contribute differently to parameter updates because their gradient magnitudes may differ considerably. As a result, directly combining these objectives with fixed coefficients may lead to imbalanced optimization, where one objective contributes disproportionately during training.

To illustrate this behavior, Fig. 2 presents the gradient magnitude distributions observed during training. Fig. 2(a)-(c) show the distributions associated with the BPR loss, the user contrastive loss, and the item contrastive loss, respectively. The results indicate that the recommendation objective generally produces smaller gradients than the contrastive objectives. In our experiments, $\|\nabla_{\Theta} \mathcal{L}_{BPR}\|$ is typically on the order of 10^{-5} , whereas $\|\nabla_{\Theta} \mathcal{L}_{InfoNCE}\|$ is commonly on the order of 10^{-4} . Fig. 2(d) further presents a batch-level decomposition of the total gradient vector and illustrates the relative contributions of the individual objectives. These observations suggest that fixed loss weights may place relatively less emphasis on the recommendation objective. This motivates the use of adaptive weighting to coordinate the contributions of the recommendation and contrastive losses.

3.6.1 Gradient-Balancing Weights

To account for the gradient-scale differences illustrated in Fig. 2, bounded learnable coefficients are introduced for the recommendation and contrastive objectives. Two scalar parameters, α_{BPR} and $\alpha_{InfoNCE}$, are mapped to the corresponding loss weights through sigmoid functions:

$$w_{BPR} = w_{\min}^b + (w_{\max}^b - w_{\min}^b) \sigma(\alpha_{BPR}) \quad (22)$$

$$w_{InfoNCE} = w_{\min}^c + (w_{\max}^c - w_{\min}^c) \sigma(\alpha_{InfoNCE}) \quad (23)$$

The bounds constrain the weights to a finite interval while allowing their relative contributions to adapt during training. In our implementation, $w_{\min}^b = 0.7$, $w_{\max}^b = 1.0$, $w_{\min}^c = 10^{-6}$, and $w_{\max}^c = 10^{-2}$.




Fig. 2. Gradient magnitude analysis. (a) Gradient distribution of the BPR loss. (b) Gradient distribution of the user contrastive loss. (c) Gradient distribution of the item contrastive loss. (d) Batch-level decomposition of the total gradient vector.

These ranges avoid degenerate solutions in which a loss weight approaches zero while maintaining flexibility in the optimization process.

To improve stability, a dual-optimizer strategy is adopted. The primary optimizer updates the model parameters Θ at every mini-batch, whereas a secondary optimizer updates α_{BPR} and $\alpha_{InfoNCE}$ at a lower frequency. This separation moderates rapid fluctuations in the weights and improves convergence. The overall training procedure is summarized in Algorithm 1.

The two-timescale update scheme is consistent with optimization strategies commonly employed in differentiable architecture search and meta-learning Liu et al. (2019); Finn et al. (2017). The bounded parameterization is further inspired by learnable multi-task weighting approaches Kendall et al. (2018). Guided by the gradient magnitudes shown in Fig. 2, the lower bound $w_{min}^c = 10^{-6}$ keeps the contrastive objective active throughout training, whereas the upper bound $w_{max}^c = 10^{-2}$ limits excessive influence from the auxiliary objective.

The overall training objective is given by Eq. (24)

$$\begin{aligned} \mathcal{L}_{total} = & w_{BPR} \mathcal{L}_{BPR} \\ & + w_{InfoNCE} \mathcal{L}_{InfoNCE} \\ & + \lambda \|\Theta\|_2^2. \end{aligned} \quad (24)$$

where, λ denotes the ℓ_2 regularization coefficient.

3.7 Complexity Analysis

Computational efficiency is an important consideration for practical recommendation systems. This subsection analyzes the computational cost of COMPASS and summarizes the design choices that keep the additional overhead moderate He et al. (2020); Wu et al. (2022).

The overall complexity consists of a one-time pre-computation stage and a per-epoch training stage. The main pre-computation step involves constructing the SVD-based views. Using randomized SVD reduces the cost for each rank q to approximately $\mathcal{O}(q^2(m+n))$, resulting in a total one-time complexity of $\mathcal{O}(|\mathcal{Q}|q_{max}^2(m+n))$ Halko et al. (2011). User activity groups are determined through a single $\mathcal{O}(m)$ pass over interaction counts. Both operations are performed once and reused throughout training.

Training procedure of COMPASS.

Input: User-item graph $\mathcal{G}_c$, social graph $\mathcal{G}_s$, hyperparameters $\lambda, d, L, \tau_g, \tau_c$, rank set $\mathcal{Q}$, adaptive-weight update frequency N_{update} .

Output: Trained model parameters Θ .

```

1: Initialize user and item embeddings  $\mathbf{E}_u^{(0)} \in \mathbb{R}^{m \times d}$  and  $\mathbf{E}_i^{(0)} \in \mathbb{R}^{n \times d}$ .
2: Initialize Gumbel-Softmax logits  $\pi^S, \pi^A$ , and  $\pi^I$  for user-group and item-side fusion.
3: Initialize adaptive loss-weight parameters  $\alpha \leftarrow \{\alpha_{\text{BPR}}, \alpha_{\text{InfoNCE}}\}$ .
4: Define model parameters  $\Theta \leftarrow \{\mathbf{E}_u^{(0)}, \mathbf{E}_i^{(0)}, \pi^S, \pi^A, \pi^I\}$ .
5: Initialize the model optimizer  $\text{Optimizer}_{\Theta}$  and the adaptive-weight optimizer  $\text{Optimizer}_{\alpha}$ .
6: Pre-compute  $\hat{\mathbf{A}}$  using Eq. (3) and  $\hat{\mathbf{S}}$  using Eq. (14).
7: Pre-compute SVD factors  $\{\mathbf{U}_q, \Sigma_q, \mathbf{V}_q\}_{q \in \mathcal{Q}}$  for  $\hat{\mathbf{A}}$ .
8: Partition users into  $\mathcal{G}_{\text{sparse}}$  and  $\mathcal{G}_{\text{active}}$  according to interaction counts.
9: for epoch = 1 to max_epochs do
10:   for each mini-batch  $\mathcal{O}_b$  do
11:     Compute collaborative embeddings  $\mathbf{E}_u, \mathbf{E}_i$  using Eq. (5).
12:     Compute social embeddings  $\mathbf{S}_u$  using Eq. (16).
13:     for each  $q \in \mathcal{Q}$  do
14:       Compute rank-specific embeddings  $\mathbf{G}_{u,q}$  and  $\mathbf{G}_{i,q}$  using Eq. (8).
15:     end for
16:     Compute activity-aware fusion weights using Eqs. (10) and (11).
17:     Compute item-side fusion weights using Eq. (12).
18:     Obtain fused multi-resolution embeddings  $\mathbf{G}_u$  and  $\mathbf{G}_i$  using Eq. (13).
19:     Compute  $\mathcal{L}_{\text{BPR}}$  on  $\mathcal{O}_b$  using Eq. (17).
20:     Compute  $\mathcal{L}_{\text{InfoNCE}}$  using Eqs. (18)-(21).
21:     Compute  $w_{\text{BPR}}$  and  $w_{\text{InfoNCE}}$  from  $\alpha$ .
22:     Compute the total loss  $\mathcal{L}_{\text{total}}$  using Eq. (24).
23:     Update model parameters  $\Theta$  using  $\text{Optimizer}_{\Theta}$ .
24:     if epoch mod  $N_{\text{update}} = 0$  then
25:       Update adaptive loss-weight parameters  $\alpha$  using  $\text{Optimizer}_{\alpha}$ .
26:     end if
27:   end for
28: end for
29: return  $\Theta$ 

```

During training, the collaborative branch propagates information over the user-item graph with complexity $\mathcal{O}(|\mathcal{E}_c|d)$, while the social branch incurs a cost of $\mathcal{O}(|\mathcal{E}_s|d)$ He et al. (2020); Fan et al. (2019). The multi-resolution branch reuses the pre-computed SVD factors and performs matrix multiplications with complexity $\mathcal{O}(|\mathcal{Q}|q_{\max}(m+n)d)$ Cai et al. (2023).

The activity-aware fusion mechanism introduces only $3|\mathcal{Q}|$ additional scalar parameters, corresponding to π^S, π^A , and π^I . For $|\mathcal{Q}| = 3$, this amounts to

nine scalar parameters and a single group lookup per user, resulting in negligible computational overhead.

Finally, computing the BPR and InfoNCE objectives with mini-batches of size B requires approximately $\mathcal{O}(B^2d)$ operations when in-batch negatives are employed. Therefore, the overall per-epoch complexity is

$$\mathcal{O}(|\mathcal{E}_c|d + |\mathcal{E}_s|d + |\mathcal{Q}|q_{\max}(m+n)d + B^2d). \quad (25)$$

Table 2. One-time SVD pre-computation time using randomized SVD [Halko et al. \(2011\)](#).

Dataset	Users	Items	Interaction Count	Pre-comp. Time (s)
LastFM	1,892	17,632	92,834	0.22
Yelp	16,239	14,284	198,397	0.39
Douban	13,024	22,347	792,062	0.71

Table 2 reports the wall-clock time required for the one-time SVD pre-computation. Even for the largest dataset, the pre-computation stage completes in less than one second, indicating that its contribution is small relative to the repeated training cost.

This analysis shows that COMPASS preserves the same linear dependence on the number of graph edges as LightGCN and SimGCL [He et al. \(2020\)](#); [Yu et al. \(2022\)](#), with an additional term associated with the multi-resolution branch. Since the SVD factors are computed once and both $|Q|$ and $q_{\max}$ remain small constants in practice, the additional cost remains moderate.

During inference, only the collaborative branch is retained. Consequently, the inference complexity reduces to $\mathcal{O}(|\mathcal{E}_c|d)$, which is identical to that of LightGCN [He et al. \(2020\)](#). Empirical running-time comparisons are provided in Section 5 through Fig. 3, complementing the asymptotic analysis with measured execution times under a fixed hardware setting.

4 Experiments

This section evaluates COMPASS and examines the contribution of its main components. The experiments are designed to address the following research questions:

RQ1. Performance comparison:

How does COMPASS compare with

strong baselines across different evaluation metrics?

RQ2. Ablation study: What are the contributions of the multi-resolution view, the social view, and the adaptive optimization strategy to the overall performance?

RQ3. Hyperparameter analysis:

How do the hyperparameters associated with the proposed components influence performance and training behavior?

RQ4. Robustness analysis: How robust is COMPASS under challenging conditions, including varying levels of sparsity and injected interaction noise?

The experimental setup is first described, including the datasets, baselines, evaluation protocol, and implementation details. The results corresponding to each research question are then presented and discussed.

4.1 Experimental Setup

This subsection introduces the experimental setting and implementation details, including the datasets, baseline methods, evaluation metrics, and hyperparameter settings. For consistency, all recommendation and contrastive objectives are evaluated under the same protocol.

4.1.1 Evaluation Metrics

The quality of the recommendation lists is evaluated at cutoff values $K \in 20, 40$ using Recall@K, NDCG@K, and HR@K. Recall@K measures the proportion of relevant items retrieved, NDCG@K accounts for the ranking positions of relevant items, and HR@K indicates whether at least one relevant item appears among the top- K recommendations. Unless stated otherwise, all reported results correspond to the average performance over five independent runs with different random seeds to reduce the influence of stochastic optimization.

To assess statistical significance, a two-sided paired t -test is performed for each dataset, metric, and cutoff value. For every configuration, each model is trained using five random seeds, yielding five values of Recall@K, NDCG@K, and HR@K. The paired t -test is then conducted between COMPASS and the best-performing baseline using these per-seed results, and the corresponding p -values are reported in Tables 4 and 5.

4.1.2 Datasets

COMPASS is evaluated on three widely used real-world datasets that represent different recommendation scenarios and exhibit varying levels of sparsity. Table 3 summarizes their statistics.

- **Yelp**: A business recommendation dataset in which user interactions with local businesses are accompanied by social links among users.
- **LastFM**: A music recommendation dataset containing listening histories together with social connections between users.
- **Douban**: A book recommendation dataset derived from an online review community, where user interactions are complemented by social relations.

Following common practice, each dataset is randomly split into 70% training, 20% validation, and 10% test interactions.

Table 3. Statistics of the evaluation datasets.

Dataset	Users	Items	Interactions	Social Links
Yelp	16,239	14,284	198,397	158,590
LastFM	1,892	17,632	92,834	25,434
Douban	13,024	22,347	792,062	169,150

4.1.3 Baselines

COMPASS is compared with ten representative and competitive baselines, including graph collaborative filtering models, contrastive learning approaches, and recent social recommendation methods designed to address noisy interactions and unreliable social relations.

- **LightGCN** He et al. (2020): A simplified graph collaborative filtering model based on linear neighborhood propagation.
- **SGL** Wu et al. (2021): A self-supervised recommendation framework that generates augmented graph views through stochastic perturbations.
- **SimGCL** Yu et al. (2022): A contrastive model that replaces graph augmentations with embedding perturbations.
- **NCL** Lin et al. (2022): A neighborhood-enriched contrastive framework that exploits structural and semantic neighbors.
- **LightGCL** Cai et al. (2023): A deterministic contrastive model that constructs global collaborative views using SVD.
- **xSimGCL** Yu et al. (2023): A lightweight extension of SimGCL that employs cross-layer contrastive learning.
- **MHCN** Yu et al. (2021): A social recommendation model based on multi-channel hypergraph convolution.
- **SHaRe** Jiang et al. (2024): A preference-aware framework that refines noisy social relations through graph rewiring.
- **DSL** Wang et al. (2023): A dual-view social recommendation model with denoising and self-supervised learning.
- **RecDiff** Li et al. (2024): A diffusion-based approach for mitigating social noise.

- **RGCML** [Xiong et al. \(2025\)](#): A robust social recommendation framework that combines denoised social relations with global and local contrastive learning.

For a fair comparison, all baseline models are independently tuned on each dataset using the same validation protocol. Hyperparameters, including embedding dimension, learning rate, and the number of layers, are selected from a common search space whenever applicable, and the best-performing configuration of each baseline is used for final evaluation.

4.1.4 Parameter Settings

All experiments are conducted on a single NVIDIA Tesla T4 GPU. Hyperparameters are selected according to validation performance to ensure fair and reproducible comparisons. For each dataset, all methods, including the baselines, are independently tuned via grid search. The search space includes embedding dimension, learning rate, number of layers, and method-specific parameters. The configuration achieving the best validation performance is retained for final evaluation. Consequently, all reported baseline results correspond to their best-performing configurations rather than fixed default settings.

To reduce the influence of random initialization and stochastic optimization, each experiment is repeated independently with at least five different random seeds, and the reported results correspond to the average performance across runs. The observed variations are consistently small, indicating stable training behavior. In addition, statistical significance analysis based on p-values is performed to verify that the observed improvements are unlikely to be caused by random fluctuations.

Unless stated otherwise, COMPASS uses a common configuration across all datasets. The embedding dimension is set to $d = 256$, selected from 16, 32, 64, 128, 256, 512 according to validation performance. The same search space is applied to all methods, and the best value for each model is retained. A two-layer GNN is employed, with the number of layers selected from 1, 2, 3, 4. Training uses the Adam optimizer with learning rate 10^{-3} , chosen from 5×10^{-4} , 10^{-3} , 5×10^{-3} , 10^{-2} , and a batch size of 4096. Model parameters are initialized using Xavier initialization [Glorot and Bengio \(2010\)](#).

Settings associated with the proposed components are summarized below:

- **SVD ranks:** The multi-resolution branch employs $Q = 3, 6, 9$ to capture interaction patterns at different structural resolutions, ranging from dominant collaborative structure to finer interaction details.
- **Contrastive temperature:** The InfoNCE temperature is fixed at $\tau_c = 0.5$, selected from 0.1, 0.2, $\dots$, 0.9.
- **Social edge dropout:** The edge dropout probability is set to $p_s = 0.1$, selected from 0.0, 0.1, 0.2, 0.3.
- **Gumbel-Softmax temperature:** The fusion temperature is initialized at $\tau_g = 1.0$, multiplied by 0.97 after each epoch, and clipped at $\tau_{\min} = 0.5$ to avoid premature convergence to a hard argmax [Jang et al. \(2016\)](#).

The loss weights are learned rather than manually tuned. Specifically, w_{BPR} and w_{InfoNCE} are updated using a separate Adam optimizer operating at a lower frequency. Unless stated otherwise, an update interval of five epochs is adopted based on preliminary experiments.

4.2 Overall Performance Comparison (RQ1)

To answer RQ1, COMPASS is compared with representative graph-based, contrastive learning, and social recommendation methods, including the recent RGCML model. Tables 4 and 5 summarize the results for $K = 20$ and $K = 40$, respectively. Across all datasets and evaluation metrics, COMPASS consistently provides the strongest performance. In addition, the reported p -values obtained from two-sided paired t -tests indicate that the improvements over the best-performing baseline are unlikely to be explained by random variation.

For $K = 20$, COMPASS consistently improves upon RGCML across all datasets. On LastFM, the gains are moderate, with relative improvements of 2.14% in Recall@20, 1.15% in HR@20, and 3.08% in NDCG@20. Similar improvements are observed on Yelp, where Recall@20, HR@20, and NDCG@20 increase by 1.99%, 1.94%, and 2.75%, respectively. The largest gains appear on Douban, with improvements of 6.32% in Recall@20, 4.69% in HR@20, and 5.76% in NDCG@20. These results suggest that the proposed multi-resolution and socially guided representations are particularly beneficial when interaction signals are relatively sparse and additional structural information becomes more valuable.

A similar trend is observed for $K = 40$. COMPASS maintains consistent improvements over RGCML and preserves its advantage when longer recommendation lists are considered. On LastFM, Recall@40, HR@40, and NDCG@40 improve by 2.03%, 2.99%, and 2.68%, respectively. On Yelp, the corresponding gains are 2.01%, 1.64%, and 1.83%. Performance improvements remain most pronounced on Douban, reaching




Fig. 3. Per-epoch training time comparison across methods. SVD pre-computation is a negligible one-time cost performed before training, excluded from per-epoch measurements and reported in Table 2.

5.61% for Recall@40, 4.56% for HR@40, and 5.24% for NDCG@40. The consistent improvements across both cutoff values indicate that the learned representations remain effective not only for the highest-ranked recommendations but also for deeper ranking positions.

Several trends can be observed from Tables 4 and 5. First, contrastive learning methods generally outperform conventional graph collaborative filtering models, highlighting the importance of self-supervised regularization under sparse implicit feedback. Second, recent social denoising approaches such as SHaRe, DSL, RecDiff, and RGCML further improve recommendation quality by exploiting auxiliary social information. Finally, COMPASS consistently achieves additional gains over these models. This behavior suggests that combining complementary structural resolutions with activity-aware fusion and socially guided contrastive regularization provides more robust preference representations than relying on a single graph view or a fixed social propagation mechanism.

Table 4. Performance comparison of recommendation models across three datasets ($k = 20$).

Method	LastFM			Douban			Yelp		
	R@20	HR@20	NDCG@20	R@20	HR@20	NDCG@20	R@20	HR@20	NDCG@20
LightGCN	0.2352	0.8460	0.2280	0.1345	0.4523	0.1089	0.1425	0.2845	0.0891
SGL	0.2417	0.8523	0.2345	0.1389	0.4598	0.1123	0.1466	0.2893	0.0922
SimGCL	0.2486	0.8588	0.2412	0.1434	0.4652	0.1156	0.1495	0.2956	0.0947
NCL	0.2498	0.8609	0.2425	0.1468	0.4701	0.1178	0.1508	0.2901	0.0956
LightGCL	0.2553	0.8680	0.2482	0.1502	0.4768	0.1201	0.1542	0.2974	0.0978
xSimGCL	0.2585	0.8745	0.2508	0.1523	0.4826	0.1234	0.1555	0.3005	0.0986
MHCN	0.2600	0.8750	0.2520	0.1556	0.4859	0.1256	0.1562	0.3009	0.0994
SHaRe	0.2611	0.8760	0.2545	0.1578	0.4892	0.1278	0.1569	0.3017	0.1000
DSL	0.2634	0.8824	0.2557	0.1634	0.4945	0.1312	0.1587	0.3015	0.1005
RecDiff	0.2641	0.8813	0.2560	0.1656	0.4968	0.1334	0.1593	0.3020	0.1002
RGCML	<u>0.2668</u>	<u>0.8856</u>	<u>0.2598</u>	<u>0.1740</u>	<u>0.5012</u>	<u>0.1390</u>	<u>0.1608</u>	<u>0.3045</u>	<u>0.1018</u>
COMPASS	0.2725	0.8958	0.2678	0.1850	0.5247	0.1470	0.1640	0.3104	0.1046
<i>p-value</i>	2.3e-3	4.1e-3	1.8e-3	3.5e-3	1.2e-3	2.8e-3	1.5e-3	4.2e-3	3.1e-3
Average Impr %	+2.14%	+1.15%	+3.08%	+6.32%	+4.69%	+5.76%	+1.99%	+1.94%	+2.75%

Table 5. Performance comparison of recommendation models across three datasets ($k = 40$).

Method	LastFM			Douban			Yelp		
	R@40	HR@40	NDCG@40	R@40	HR@40	NDCG@40	R@40	HR@40	NDCG@40
LightGCN	0.3156	0.8867	0.2645	0.2613	0.6589	0.1423	0.1987	0.3412	0.1056
SGL	0.3234	0.8924	0.2712	0.2673	0.6694	0.1467	0.2021	0.3456	0.1083
SimGCL	0.3289	0.8959	0.2767	0.2774	0.6742	0.1500	0.2090	0.3489	0.1123
NCL	0.3312	0.8991	0.2789	0.2823	0.6798	0.1534	0.2102	0.3498	0.1136
LightGCL	0.3378	0.9056	0.2845	0.2867	0.6867	0.1578	0.2156	0.3523	0.1152
xSimGCL	0.3423	0.9116	0.2878	0.2934	0.6945	0.1623	0.2178	0.3567	0.1168
MHCN	0.3450	0.9150	0.2885	0.2956	0.6989	0.1645	0.2190	0.3570	0.1172
SHaRe	0.3470	0.9170	0.2884	0.3001	0.7014	0.1689	0.2200	0.3590	0.1180
DSL	0.3512	0.9234	0.2923	0.3045	0.7103	0.1734	0.2198	0.3589	0.1178
RecDiff	0.3517	0.9202	0.2922	<u>0.3089</u>	0.7089	0.1767	0.2210	0.3600	0.1185
RGCML	<u>0.3562</u>	<u>0.9289</u>	<u>0.2978</u>	0.3030	<u>0.7184</u>	<u>0.1834</u>	<u>0.2234</u>	<u>0.3648</u>	<u>0.1205</u>
COMPASS	0.3634	0.9567	0.3058	0.3200	0.7512	0.1930	0.2279	0.3708	0.1227
<i>p-value</i>	1.7e-3	2.5e-4	1.1e-3	1.1e-3	4.3e-4	2.9e-3	2.1e-3	3.7e-3	2.9e-3
Average Impr %	+2.03%	+2.99%	+2.68%	+5.61%	+4.56%	+5.24%	+2.01%	+1.64%	+1.83%

Besides recommendation accuracy, computational efficiency is also important in practical systems. Fig. 3 compares the average per-epoch training time of representative social recommendation models. Although COMPASS introduces

additional multi-resolution components, its training cost remains comparable to recent methods such as RGCML and RecDiff. Combined with the complexity analysis in Section 4 and the negligible one-time SVD preprocessing cost reported

in Table 2, these results indicate that the accuracy gains are achieved without a substantial increase in computational overhead.

4.3 Ablation Study (RQ2)

To investigate the contribution of individual components, the complete COMPASS model is compared with three ablated variants: (i) $\text{COMPASS}_{w/o-s}$, which removes the social view, (ii) $\text{COMPASS}_{w/o-m}$, which removes the multi-resolution decomposition branch, and (iii) $\text{COMPASS}_{w/o-a}$, which replaces adaptive gradient-balanced optimization with fixed equal loss weights. The results are summarized in Table 6.

Performance decreases are observed for all ablated variants across datasets and evaluation metrics, indicating that the three components provide complementary benefits. Their relative contributions, however, vary across datasets.

Among the three components, the multi-resolution branch produces the largest and most consistent impact. Removing this component results in the greatest performance degradation on all datasets. On LastFM, Recall@20 decreases from 0.2725 to 0.2571, corresponding to an approximate relative drop of 5.7%, while on Yelp the reduction reaches about 6.0%. The effect is even more pronounced on Douban, where Recall@20 decreases by approximately 10.4%. These observations suggest that modeling collaborative structure at multiple resolutions provides information that cannot be captured by a single graph representation.

The contribution of the social view depends more strongly on the characteristics of the dataset. On LastFM and Yelp, removing the social branch leads to moderate performance degradation. In

contrast, the impact becomes substantially larger on Douban, where Recall@20 decreases from 0.1850 to 0.1701. This behavior indicates that social information provides additional preference signals that complement collaborative interactions, particularly in settings where social relations are more informative. The simultaneous performance reductions observed for both $w/o-s$ and $w/o-m$ further suggest that the two branches capture different aspects of user preference rather than redundant information.

Replacing adaptive gradient-balanced optimization with fixed loss weights yields smaller but consistently negative effects across all datasets. Relative reductions in Recall@20 are approximately 2.5%, 4.0%, and 2.6% on LastFM, Douban, and Yelp, respectively. Although the magnitude of these decreases is smaller than those caused by removing the social or multi-resolution branches, the trend remains consistent across datasets with different characteristics. This observation indicates that adaptive weighting mainly serves as a stabilizing mechanism that improves the coordination between recommendation and contrastive objectives, rather than acting as the primary source of performance gains.

Overall, the ablation results indicate that the largest improvements originate from the multi-resolution decomposition module, while the social view and adaptive optimization provide additional and complementary benefits. The combination of all three components yields the strongest performance, suggesting that the gains achieved by COMPASS arise from their interaction rather than from a single dominant component.

Table 6. Ablation study of COMPASS components across three datasets ($k = 20$).

Dataset	Metric	COMPASS _{w/o-s}	COMPASS _{w/o-m}	COMPASS _{w/o-a}	COMPASS
Yelp	Recall@20	0.1565	0.1542	0.1598	0.1640
	NDCG@20	0.0989	0.0971	0.1012	0.1046
	HR@20	0.2986	0.2945	0.3034	0.3104
LastFM	Recall@20	0.2615	0.2571	0.2658	0.2725
	NDCG@20	0.2565	0.2518	0.2612	0.2678
	HR@20	0.8753	0.8672	0.8834	0.8958
Douban	Recall@20	0.1701	0.1658	0.1776	0.1850
	NDCG@20	0.1349	0.1312	0.1411	0.1470
	HR@20	0.4982	0.4901	0.5116	0.5247

Note: w/o-s = without social view; w/o-m = without multi-resolution view; w/o-a = without adaptive gradient-balanced optimization.

4.4 Hyperparameter Sensitivity Analysis (RQ3)

To evaluate the robustness of COMPASS with respect to general hyperparameters, four settings are examined while keeping the remaining configurations fixed, namely the embedding dimension d , propagation depth L , batch size, and contrastive temperature τ_c . Fig. 4 reports Recall@20 and NDCG@20 on the three datasets.

Overall, similar trends are observed across datasets, suggesting that the model is not overly sensitive to moderate variations in these parameters. Increasing the embedding dimension improves performance up to $d = 256$, after which a slight decline is observed. This behavior indicates that larger embedding spaces provide additional representational capacity, while excessively large dimensions offer limited benefit and may reduce generalization.

Finally, the contrastive temperature τ_c reaches its best performance near 0.5 across all datasets. Smaller temperatures produce sharper similarity distributions, whereas larger values reduce the discriminative effect of contrastive learning. Taken

together, these results indicate that COMPASS maintains stable performance over a relatively broad range of hyperparameter values and does not require highly specialized tuning.

The propagation depth exhibits a comparable pattern. Moving from one to two layers consistently improves performance, whereas deeper architectures gradually degrade recommendation quality. This tendency is consistent with the over-smoothing phenomenon commonly observed in graph neural networks, indicating that shallow propagation is sufficient for capturing informative neighborhood structure.

Batch size also demonstrates a clear optimum around 4096. Smaller batches lead to slightly lower performance, while excessively large batches reduce recommendation quality, suggesting that moderate batch sizes provide a favorable trade-off between optimization stability and generalization.

4.4.1 Multi-Resolution Rank Configuration

The influence of the rank set $\mathcal{Q}$ is examined from two perspectives, namely the




Fig. 4. Hyperparameter sensitivity analysis across three datasets: (a) LastFM dataset, (b) Douban dataset, (c) Yelp dataset, (●) Recall@20, (■) NDCG@20.

number of decomposition levels and the choice of rank values.

Table 7 shows that incorporating multiple decomposition levels consistently improves recommendation accuracy compared with a single-rank configuration. Performance increases when moving from one to three ranks, whereas introducing a fourth rank yields only marginal gains and occasionally produces slight degradation. This behavior suggests that a small number of complementary approximation levels is sufficient to capture useful structural diversity. Consequently, $|\mathcal{Q}| = 3$ is adopted throughout the experiments.

Given $|\mathcal{Q}| = 3$, several candidate rank combinations with different ranges and spacings were evaluated. Configurations with very closely spaced ranks provide limited diversity among the generated views, while uniformly high-rank settings tend to introduce redundant fine-grained information. Among the examined alternatives, $\mathcal{Q} = 3, 6, 9$ consistently achieved the best overall performance and was therefore selected as the default configuration. These results indicate that combining low, intermediate, and relatively higher-rank approximations provides complementary collaborative information without introducing excessive redundancy.

4.4.2 Analysis of Learned Fusion Weights

Beyond the choice of the rank set itself, it is also informative to examine how the learned Gumbel-Softmax fusion mechanism allocates importance to different decomposition levels. Users are divided into sparse and active groups according to the median number of training interactions, and Fig. 5 shows the average learned fusion weights over five independent runs.

A consistent pattern emerges across datasets. Sparse users assign a larger proportion of their weight to the low-rank component ($q = 3$), whereas active users place relatively greater emphasis on the higher-rank component ($q = 9$). The intermediate component ($q = 6$) receives comparable weights across groups and acts as a transition between coarse and fine structural representations.

These observations are consistent with the intuition underlying the proposed activity-aware fusion mechanism. Users with limited interaction histories benefit primarily from low-rank views that emphasize broad collaborative patterns and provide stronger regularization. In contrast, users with richer histories can exploit higher-rank components that preserve finer structural details and capture more specific preferences.

The learned weight distributions therefore indicate that the model adapts the contribution of different decomposition levels according to user activity, rather than relying on a fixed rank combination for all users. This behavior complements the rank sensitivity analysis in Section 5.4.1 and suggests that the multi-resolution representations provide useful structural diversity across different data regimes.

4.4.3 Sensitivity of Adaptive Optimization

The dual-optimizer strategy introduces an update interval T , which determines how frequently the loss weights are adjusted relative to the model parameters. Table 8 summarizes the results obtained with different values of T .

Updating the loss weights at every iteration ($T = 1$) leads to slightly inferior performance, whereas excessively large intervals reduce the ability of the adaptive

Table 7. Effect of rank set cardinality $|\mathcal{Q}|$ on recommendation performance. Each row reports the best result found across all evaluated configurations at that cardinality level.

$ \mathcal{Q} $	LastFM		Douban		Yelp	
	R@20	N@20	R@20	N@20	R@20	N@20
1	0.2604	0.2551	0.1715	0.1366	0.1556	0.0978
2	0.2663	0.2621	0.1798	0.1432	0.1608	0.1021
3	0.2727	0.2678	0.1850	0.1470	0.1640	0.1046
4	0.2714	0.2659	0.1831	0.1451	0.1629	0.1038




Fig. 5. Activity-aware fusion weights per SVD rank for sparse and active users on (a) LastFM and (b) Yelp. Sparse users favor lower-rank global components ($q = 3$), while active users assign greater weight to higher-rank residual components ($q = 9$).

weights to respond to changing optimization dynamics. The best performance is obtained at $T = 5$, with only minor differences observed between $T = 5$ and $T = 10$. These results indicate that the proposed optimization strategy is relatively insensitive to moderate variations in the update interval. Therefore, $T = 5$ is adopted as the default setting, providing a favorable balance between stability and adaptability.

4.5 Robustness Analysis (RQ4)

The practical value of a recommender system depends not only on average ranking accuracy but also on its ability to remain reliable under challenging conditions. To examine this aspect, robustness

Table 8. Sensitivity to the loss-weight update interval T on LastFM. $T=5$ is the default.

Update interval T	R@20	N@20
$T = 1$	0.2664	0.2612
$T = 5$	0.2727	0.2678
$T = 10$	0.2714	0.2663
$T = 20$	0.2689	0.2638

is evaluated from three complementary perspectives. First, performance is analyzed across different levels of user activity to assess behavior under sparse interaction histories. Second, the diversity of learned preference scores is investigated to understand whether the model can distinguish varying degrees of user interest. Finally, interaction noise is artificially injected into the training graph to evaluate resilience to corrupted feedback. Together, these analyses provide a more comprehensive assessment beyond conventional offline evaluation.

4.5.1 Robustness under User Sparsity

To investigate performance under different levels of interaction sparsity, users are partitioned according to their activity level, and Recall@20 is evaluated separately for each group. Fig. 6 reports the results.

Across all datasets, recommendation quality improves with increasing user




Fig. 6. Recall@20 across user activity groups.

activity, which is expected because richer interaction histories provide stronger preference signals. More importantly, COMPASS consistently achieves the best performance in every activity group, with the largest margins observed for low-activity users. On LastFM, the advantage is particularly evident for users with fewer than 20 interactions, while on Yelp and Douban the gains are most noticeable for users with fewer than five interactions.

These results suggest that the proposed framework remains effective when interaction histories are extremely limited. The combination of multi-resolution decomposition and social guidance provides complementary information when collaborative signals alone are insufficient. As user activity increases, the performance gap gradually becomes smaller because abundant interaction data already provides stronger supervision. Nevertheless, COMPASS maintains a consistent advantage across all activity groups, indicating robustness to different degrees of sparsity.

4.5.2 Preference Differentiation Analysis

The superior performance observed for sparse users suggests that COMPASS may

capture preference intensity more effectively. To provide qualitative insight, preference scores are computed as $p(u, i) = \sigma(\mathbf{z}_u^\top \mathbf{z}_i)$ and visualized for a randomly selected user and several interacted items. Fig. 8 compares score distributions at an early training stage and after convergence.

At convergence, several competing methods assign relatively similar scores across interacted items. In contrast, COMPASS produces a wider range of preference values, indicating stronger differentiation among items within the same interaction history. Rather than treating historical interactions as equally informative, the learned representations preserve varying degrees of preference intensity.

To further examine this behavior at the population level, the standard deviation of preference scores over interacted items is computed for each user and averaged across all users on Yelp. COMPASS exhibits the largest average score dispersion (0.06), compared with 0.05 for RGCML, 0.04 for RecDiff and DSL, and 0.02 for LightGCL. Although larger score dispersion does not necessarily imply better performance by itself, it suggests that the proposed model preserves richer variations in user preferences. This behavior is consistent with the improved results observed for low-activity users in Fig. 6,




Fig. 7. Recall@20 under interaction-noise injection on (a) LastFM, (b) Douban, and (c) Yelp.




Fig. 8. Preference-score heatmaps at epoch 10 and at convergence.

where distinguishing subtle preference differences becomes particularly important.

4.5.3 Robustness to Interaction Noise

To evaluate robustness against imperfect feedback, interaction noise is injected into the training graph by randomly replacing a fraction of observed user-item interactions with non-interacted items while preserving the number of interactions per user. Noise levels range from 0% to

20%, and Fig. 7 reports Recall@20 under different corruption ratios.

As expected, recommendation accuracy decreases as the amount of noise increases. Nevertheless, COMPASS consistently maintains the best performance across all datasets and corruption levels. More importantly, its degradation is more gradual than that of competing approaches. While the relative ordering among baseline methods changes under severe corruption, COMPASS preserves a stable advantage throughout the entire noise range.

For example, on LastFM, Recall@20 decreases from 0.2725 to 0.2405 when the noise ratio increases from 0% to 20%, corresponding to a relative reduction of approximately 11.7%. Under the same setting, RGCML exhibits a larger reduction of approximately 15.7%. Similar trends are observed on Douban and Yelp, where competing methods experience more pronounced performance degradation as the noise level increases.

These observations indicate that the proposed framework is less sensitive to spurious interactions and can preserve useful collaborative information even when part of the training graph is corrupted. The complementary multi-resolution views provide representations

with different levels of structural detail, while social-guided regularization supplies additional supervision when interaction signals become unreliable. Furthermore, adaptive loss balancing helps maintain stable optimization during noisy training. Overall, the results suggest that COMPASS remains effective under realistic data imperfections and exhibits strong robustness to corrupted feedback.

5 Conclusions and Future Work

This paper presented COMPASS, a recommendation framework that combines multi-resolution low-rank decomposition with social-guided contrastive learning to improve robustness under sparse and noisy interaction environments. By constructing multiple low-rank views of the interaction graph and adaptively fusing them, the proposed framework captures collaborative information at different levels of structural detail while incorporating social relations as soft semantic guidance. Experimental results on three real-world datasets showed consistent improvements over representative graph-based, contrastive learning, and social recommendation methods. Additional analyses under varying sparsity levels and interaction noise further demonstrated that the proposed framework maintains stable performance under challenging conditions without introducing substantial computational overhead.

Beyond recommendation accuracy, the analyses provided insight into the behavior of the learned representations. The learned fusion weights revealed that users with limited interaction histories rely more heavily on coarse collaborative structure, whereas users with richer histories benefit from finer structural details. This

behavior suggests that different user populations require different levels of structural granularity and highlights the value of multi-resolution representations for recommendation.

Several directions remain for future research. One possible extension is to improve the quality of social supervision by incorporating trust reliability estimation, allowing the model to distinguish informative social relations from weak or noisy connections. Another promising direction is to extend the multi-resolution decomposition framework to the social graph itself, enabling collaborative and social information to be modeled at multiple structural scales. Finally, integrating semantic information from item content, such as textual descriptions encoded by large language models, may further enrich fine-grained representations and help distinguish genuine niche interests from noisy interactions, particularly in sparse recommendation scenarios.

Author contributions

Mohammad Aliabadi contributed to conceptualization, methodology, software, formal analysis, investigation, and writing of the original draft.

Alireza Abdollahpouri contributed to conceptualization, supervision, writing, review, and editing of the manuscript.

Parham Moradi contributed to supervision, writing, review, and editing of the manuscript.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Competing interests

The authors declare that they have no competing interests.

Data availability

The datasets used in this study are publicly available. The original data sources and corresponding preprocessing procedures are described in Section 5.1.2. The dataset repositories can be accessed at <https://github.com/lab4ml/socialrec-datasets>.

References

- Berg, R.v.d., Kipf, T.N., Welling, M.: Graph convolutional matrix completion. arXiv preprint arXiv:1706.02263 (2017)
- Cai, X., Huang, C., Xia, L., Ren, X.: Lightgcl: Simple yet effective graph contrastive learning for recommendation. In: Proceedings of the International Conference on Learning Representations (2023)
- Eckart, C., Young, G.: The approximation of one matrix by another of lower rank. *Psychometrika* **1**(3), 211–218 (1936)
- Finn, C., Abbeel, P., Levine, S.: Model-agnostic meta-learning for fast adaptation of deep networks. In: International Conference on Machine Learning (ICML) (2017)
- Fan, W., Ma, Y., Li, Q., He, Y., Zhao, E., Tang, J., Yin, D.: Graph neural networks for social recommendation. In: The World Wide Web Conference, pp. 417–426 (2019)
- Glorot, X., Bengio, Y.: Understanding the difficulty of training deep feedforward neural networks. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, pp. 249–256 (2010). JMLR Workshop and Conference Proceedings
- He, L., Cheng, D., Zhang, G., Zhang, S.: Leveraging long-range nodes in multi-view graph contrastive learning. *Information Fusion* **122**, 103186 (2025)
- He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., Wang, M.: Lightgcn: Simplifying and powering graph convolution network for recommendation. In: Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 639–648 (2020)
- Halko, N., Martinsson, P.-G., Tropp, J.A.: Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review* **53**(2), 217–288 (2011)
- Hammond, D.K., Vandergheynst, P., Gribonval, R.: Wavelets on graphs via spectral graph theory. *Applied and Computational Harmonic Analysis* **30**(2), 129–150 (2011)
- Jang, E., Gu, S., Poole, B.: Categorical reparameterization with gumbel-softmax. arXiv preprint arXiv:1611.01144 (2016)
- Jiang, W., Gao, X., Xu, G., Chen, T., Yin, H.: Challenging low homophily in social recommendation. In: Proceedings of the ACM Web Conference 2024, pp. 3476–3484 (2024)
- Kendall, A., Gal, Y., Cipolla, R.: Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In: Proceedings of the IEEE Conference on Computer Vision and Pattern

- Recognition (CVPR) (2018)
- Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. arXiv preprint arXiv:1609.02907 (2016)
- Liu, H., Simonyan, K., Yang, Y.: DARTS: Differentiable architecture search. In: International Conference on Learning Representations (2019)
- Lin, Z., Tian, C., Hou, Y., Zhao, W.X.: Improving graph collaborative filtering with neighborhood-enriched contrastive learning. In: Proceedings of the ACM Web Conference 2022, pp. 2320–2329 (2022)
- Li, Z., Xia, L., Huang, C.: Recdiff: Diffusion model for social recommendation. In: Proceedings of the 33rd ACM International Conference on Information and Knowledge Management, pp. 1346–1355 (2024)
- Li, J., Yang, C., Ye, G., Nguyen, Q.V.H.: Graph neural networks with deep mutual learning for designing multi-modal recommendation systems. *Information Sciences* **654**, 119815 (2024) <https://doi.org/10.1016/j.ins.2023.119815>
- Mirsky, L.: Symmetric gauge functions and unitarily invariant norms. *The Quarterly Journal of Mathematics* **11**(1), 50–59 (1960)
- Miao, R., Yang, Y., Ma, Y., Juan, X., Xue, H., Tang, J., Wang, Y., Wang, X.: Negative samples selecting strategy for graph contrastive learning. *Information Sciences* **613**, 667–681 (2022) <https://doi.org/10.1016/j.ins.2022.09.024>
- Ma, H., Zhou, D., Liu, C., Lyu, M.R., King, I.: Recommender systems with social regularization. In: Proceedings of the Fourth ACM International Conference on Web Search and Data Mining, pp. 287–296 (2011)
- Meng, S., Zhang, X., Liu, N., Tu, L., Li, Q.: Knowledge&social-based collaborative method with contrastive graph structure learning for explainable recommendation. *Information Sciences* **709**, 122077 (2025) <https://doi.org/10.1016/j.ins.2025.122077>
- Quan, Y., Ding, J., Gao, C., Yi, L., Jin, D., Li, Y.: Robust preference-guided denoising for graph based social recommendation. In: Proceedings of the ACM Web Conference 2023, pp. 1097–1108 (2023)
- Shirali, A., Kazemi, R., Amini, A.: Collaborative filtering with representation learning in the frequency domain. *Information Sciences* **681**, 121240 (2024) <https://doi.org/10.1016/j.ins.2024.121240>
- Sakong, D., Vu, V.H., Huynh, T.T., Nguyen, P.L., Yin, H., Nguyen, Q.V.H., Nguyen, T.T.: Higher-order knowledge-enhanced recommendation with heterogeneous hypergraph multi-attention. *Information Sciences* **680**, 121165 (2024) <https://doi.org/10.1016/j.ins.2024.121165>
- Sun, M., Xu, Y., Tan, Z., You, D., Chen, Z.: Multi-level graph contrastive learning for cold-start recommendation in mashup development. *Information Sciences* **717**, 122319 (2025) <https://doi.org/10.1016/j.ins.2025.122319>
- Tang, J., Hu, X., Liu, H.: Social recommendation: a review. *Social network*

- analysis and mining **3**(4), 1113–1133 (2013)
- Tao, Q., Liu, C., Xia, Y., Xu, Y., Li, L.: Adaptive multi-graph contrastive learning for bundle recommendation. *Neural Networks* **181**, 106832 (2025)
- Von Luxburg, U.: A tutorial on spectral clustering. *Statistics and computing* **17**(4), 395–416 (2007)
- Wang, X., He, X., Wang, M., Feng, F., Chua, T.-S.: Neural graph collaborative filtering. In: *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 165–174 (2019)
- Wang, M., Lin, Y., Lin, G., Yang, K., Wu, X.-m.: M2grl: A multi-task multi-view graph representation learning framework for web-scale recommender systems. In: *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2349–2358 (2020)
- Wu, L., Li, J., Sun, P., Hong, R., Ge, Y., Wang, M.: Diffnet++: A neural influence and interest diffusion network for social recommendation. *IEEE Transactions on Knowledge and Data Engineering* **34**(10), 4753–4766 (2020)
- Wang, X., Li, Q., Yu, D., Huang, W., Li, Q., Xu, G.: Neural causal graph collaborative filtering. *Information Sciences* **677**, 120872 (2024) <https://doi.org/10.1016/j.ins.2024.120872>
- Won, H., Oh, B., Yang, H., Lee, K.-H.: Cross-modal contrastive learning for aspect-based recommendation. *Information Fusion* **99**, 101858 (2023)
- Wu, L., Sun, P., Fu, Y., Hong, R., Wang, X., Wang, M.: 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 (2019)
- Wu, S., Sun, F., Zhang, W., Xie, X., Cui, B.: Graph neural networks in recommender systems: a survey. *ACM Computing Surveys* **55**(5), 1–37 (2022)
- Wu, J., Wang, X., Feng, F., He, X., Chen, L., Lian, J., Xie, X.: Self-supervised graph learning for recommendation. In: *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 726–735 (2021)
- Wen, M., Wang, H., Li, W., Fan, Z., Yu, X.: Contrastive graph semantic learning via prototype for recommendation. *Information Sciences* **699**, 121799 (2025) <https://doi.org/10.1016/j.ins.2024.121799>
- Wei, Z., Wang, K., Li, F., Ma, Y.: M3KGR: A momentum contrastive multi-modal knowledge graph learning framework for recommendation. *Information Sciences* **676**, 120812 (2024) <https://doi.org/10.1016/j.ins.2024.120812>
- Wang, T., Xia, L., Huang, C.: Denoised self-augmented learning for social recommendation. In: *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, pp. 2324–2331 (2023)
- Wu, Q., Zhang, H., Gao, X., He, P., Weng, P., Gao, H., Chen, G.: Dual graph attention networks for deep latent representation of multifaceted social effects

- in recommender systems. In: The World Wide Web Conference, pp. 2091–2102 (2019)
- Xu, F., Zhu, Z., Fu, Y., Wang, R., Liu, P.: Collaborative denoised graph contrastive learning for multi-modal recommendation. *Information Sciences* **679**, 121017 (2024) <https://doi.org/10.1016/j.ins.2024.121017>
- Xiong, F., Zhang, T., Pan, S., Luo, G., Wang, L.: Robust graph based social recommendation through contrastive multi-view learning. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 39, pp. 12890–12898 (2025)
- Yang, X., Steck, H., Liu, Y.: Circle-based recommendation in online social networks. In: Proceedings of the 18th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, pp. 1267–1275 (2012)
- Yang, Y., Wu, L., Wang, Z., He, Z., Hong, R., Wang, M.: Graph bottlenecked social recommendation. In: Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, pp. 3853–3862 (2024)
- Yang, X.-J., Wu, Y.-H., Zhang, Z.-H., Wang, J., Nie, F.-P.: Multitype view of knowledge contrastive learning for recommendation. *Neural Networks* **181**, 106690 (2025)
- Yu, J., Xia, X., Chen, T., Cui, L., Hung, N.Q.V., Yin, H.: Xsingcl: Towards extremely simple graph contrastive learning for recommendation. *IEEE Transactions on Knowledge and Data Engineering* **36**(2), 913–926 (2023)
- Yu, J., Yin, H., Li, J., Wang, Q., Hung, N.Q.V., Zhang, X.: Self-supervised multi-channel hypergraph convolutional network for social recommendation. In: Proceedings of the Web Conference 2021, pp. 413–424 (2021)
- Yu, J., Yin, H., Xia, X., Chen, T., Cui, L., Nguyen, Q.V.H.: Are graph augmentations necessary? simple graph contrastive learning for recommendation. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval, pp. 1294–1303 (2022)
- Zhang, C., Xue, S., Li, J., Wu, J., Du, B., Liu, D., Chang, J.: Multi-aspect enhanced graph neural networks for recommendation. *Neural Networks* **157**, 90–102 (2023)
- Zhou, T., Ye, H., Cao, F.: Node-personalized multi-graph convolutional networks for recommendation. *Neural Networks* **173**, 106169 (2024)