

# Invariance Matters: Empowering Social Recommendation via Graph Invariant Learning

Yonghui Yang

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Hefei, China  
yyh.hfut@gmail.com

Le Wu\*

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Institute of Dataspace,  
Hefei Comprehensive National  
Science Center  
Hefei, China  
lewu.ustc@gmail.com

Yuxin Liao

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Hefei, China  
yuxinliao314@gmail.com

Zhuangzhuang He

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Hefei, China  
hyicheng223@gmail.com

Pengyang Shao

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Hefei, China  
shaopymark@gmail.com

Richang Hong

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Institute of Dataspace,  
Hefei Comprehensive National  
Science Center  
Hefei, China  
hongrc.hfut@gmail.com

Meng Wang

Key Laboratory of Knowledge  
Engineering with Big Data,  
Hefei University of Technology  
Hefei, China  
eric.mengwang@gmail.com

## ABSTRACT

Graph-based social recommender systems have demonstrated great potential in alleviating data sparsity by leveraging high-order user influence embedded in social networks. However, most existing methods rely heavily on the observed social graph, which is often noisy and includes spurious or task-irrelevant connections that can mislead user preference learning. Identifying and removing these noisy relations is crucial but challenging due to the lack of ground-truth annotations. In this paper, we approach the social denoising problem from the perspective of graph invariant learning and propose a novel approach, *Social Graph Invariant Learning (SGIL)*. Specifically, *SGIL* aims to uncover stable user preferences within the input social graph, thereby enhancing the robustness of

graph-based social recommendation systems. To achieve this goal, *SGIL* first simulates multiple noisy social environments through graph generators. It then seeks to learn environment-invariant user preferences by minimizing invariant risk across these environments. To further promote diversity in the generated social environments, we employ an adversarial training strategy to simulate more potential social noisy distributions. Extensive experimental results demonstrate the effectiveness of the proposed *SGIL*. The code is available at <https://github.com/yimutianyang/SIGIR2025-SGIL>.

## CCS CONCEPTS

• **Human-centered computing** → **Collaborative and social computing**.

## KEYWORDS

Social Recommendation, Invariant Learning, Graph Denoising

## ACM Reference Format:

Yonghui Yang, Le Wu, Yuxin Liao, Zhuangzhuang He, Pengyang Shao, Richang Hong, and Meng Wang. 2025. Invariance Matters: Empowering Social Recommendation via Graph Invariant Learning. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '25)*, July 13–18, 2025, Padua, Italy. ACM, New York, NY, USA, 10 pages. <https://doi.org/10.1145/3726302.3730013>

\* Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
*SIGIR '25, July 13–18, 2025, Padua, Italy*

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.  
ACM ISBN 979-8-4007-1592-1/2025/07  
<https://doi.org/10.1145/3726302.3730013>

## 1 INTRODUCTION

With the proliferation of social platforms, social recommendation has emerged as a pivotal technique for delivering personalized suggestions to users, harnessing both user-item interaction data and user-user social relationships [29]. The underlying premise is that users who are socially connected are more likely to exhibit similar preference patterns [26]. Consequently, much of the prior research has focused on how to effectively leverage the potential of social networks to enhance the learning of user preferences [17, 21]. Early work in this domain predominantly relied on first-order social connections to improve recommendation quality. For instance, user preferences were enhanced through social regularization [13] or by incorporating trust relationships into user preference factorization [8]. These methods paved the way for more sophisticated integration of social data.

In recent years, graph-based social recommendation methods have gained prominence due to their remarkable success in improving recommendation accuracy, particularly in addressing the challenge of data sparsity within user behaviors [9, 12, 36]. By leveraging the powerful representational capabilities of Graph Neural Networks (GNNs), these methods encode user preferences while simultaneously capturing higher-order social influences and collaborative signals [35, 36]. This has made graph-based approaches a compelling choice for incorporating social information into recommendation systems, enabling them to better model the complex relationships between users and their social networks.

However, despite the apparent benefits of aggregating all available information from social networks, many graph-based social recommendation methods fail to account for the inevitable noise inherent in these networks. Social networks are often fraught with task-irrelevant relationships that may obscure or distort the true preferences of users. Previous works [15, 27, 42] have empirically and statistically demonstrated the pervasiveness of social noise. For example, [15] defines the concept of preference-aware social homophily ratio and observes that social networks typically exhibit low homophily, meaning that social connections do not always align with user preferences. Furthermore, [42] investigates the direct use of original social networks in graph-based recommenders, revealing that such approaches result in only incremental improvements, or in some cases, a decline in performance. This highlights the importance of addressing the noise present in social data for effective recommendation learning.

Nevertheless, the task of identifying and eliminating these noisy social relations remains a substantial challenge. Current approaches typically follow a two-step paradigm: 1) filtering the social structure based on user preference signals, and 2) optimizing this filtering process using specific learning objectives, such as co-optimization [24, 27], contrastive learning [15], and graph bottleneck learning [42]. Although these methods have made notable strides in denoising social data, they are inherently limited by the lack of reliable mechanisms to identify and verify noise, as there are no ground truth labels available to assess the accuracy of the filtered relationships. This uncertainty raises an important question: *Given the difficulty in exhaustively removing noisy social connections, can we still learn stable and sufficient user preferences, even in the presence of noisy environments?*

To achieve this goal, we adopt the principle of invariant learning to address the social noise issue. Invariant learning [1, 18] posits that the input data contains both invariant features (stable) and spurious features (variable across different environments). The focus of invariant learning is to identify and utilize only the invariant features, thereby enhancing the model’s robustness. As illustrated in Figure 1, we compare existing filtering-based social denoising methods with our proposed invariance-based approach. Drawing inspiration from this principle, we hypothesize that a task-relevant structure exists within the social graph that is both sufficient and invariant for recommendation tasks. Although this invariant structure is not directly accessible, we propose narrowing the bounds of invariance by exploring diverse noisy environments. Specifically, for a given input social graph, we first simulate multiple social environments (which embody different types of noise), and then optimize the learning objectives to encourage invariance across these environments. This formulation presents two key challenges: i) How can we simulate diverse noisy environments based on the input social networks and interactions? ii) How can we achieve invariant user preference learning across different noisy environments?

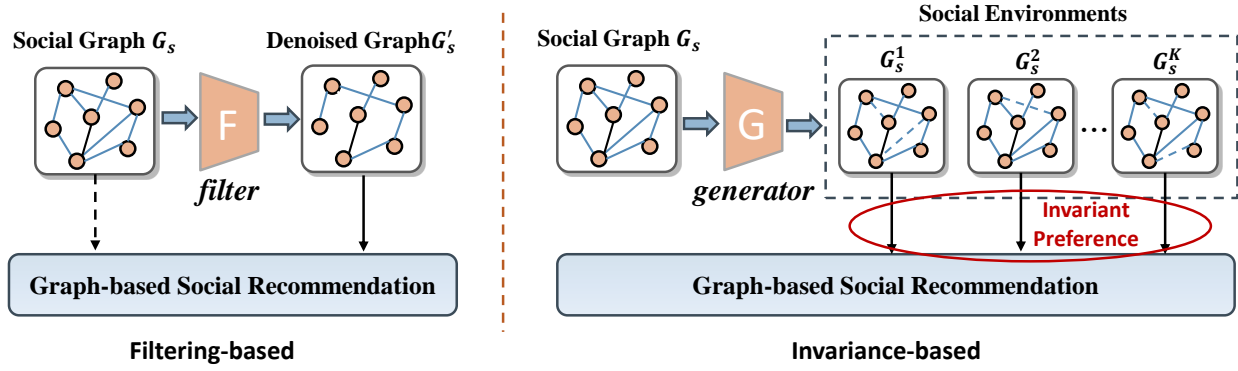
In this work, we propose a novel *Social Graph Invariant Learning (SGIL)*, to tackle the above challenges. First, we design preference-guided environment generators that simulate multiple noisy environments, thereby approximating real-world scenarios where social noise is prevalent. Next, we capture invariant user preferences by leveraging an invariance-based regularizer. Specifically, we derive the optimization objective for each environment, taking into account the characteristics of the recommendation task, and incorporate an invariance penalty term to enforce the stability of user preferences across environments. To further ensure the diversity of the simulated noisy environments, we employ an adversarial training strategy. This strategy challenges the model iteratively, forcing it to perform well in progressively more difficult and varied environments, which ultimately improves its generalization ability. The main contributions of this work are summarized as follows:

- We propose a novel social denoising approach *SGIL* for robust recommendation from the standpoint of invariant learning. Unlike traditional methods that attempt to remove noisy data directly, *SGIL* focuses on learning invariant user preferences across diverse environments, ensuring robustness despite noise.
- We derive an optimization objective tailored to each social environment and design an invariance-based optimization to learn stable user preferences. Additionally, we incorporate an adversarial training strategy to guarantee the diversity of the generated environments, further enhancing the model’s robustness.
- We conduct extensive experiments on three benchmark datasets. Empirical studies, including both real and semi-synthetic datasets, demonstrate the effectiveness of the proposed *SGIL* method in improving social recommendation performance and robustness.

## 2 RELATED WORKS

### 2.1 Graph-based Social Recommendation

Social recommendations have been widely deployed in online services, which leverage social networks as auxiliary information to



**Figure 1: Illustration of the existing filtering-based social denoising methods and our proposed invariance-based method. Instead of filtering the noisy part of the input social graph, we encourage learning the invariant yet sufficient user preferences across environments.**

capture users' behavior patterns [25, 29]. Following the social homophily [26] and social influence theory [25], social recommendations are devoted to characterizing social relation effects on user preferences, including trust-based preference learning [8], social regularization [13, 23], and social influence modeling [36]. Early works mainly exploit shallow social relations to enhance recommendation, ignoring the complex high-order social structure influences.

Users' interactions can be naturally formulated in graph forms [9, 44], many graph-based social recommendations have been proposed [36, 43, 45]. Particularly, graph-based social recommendations achieved an impressive process by borrowing the power of representation ability of GNNs. The core premise of this work is that incorporating high-order social neighbors and interacted items into user preference learning enriches the available information, thereby addressing the issue of data sparsity. Current studies mainly focus on modeling social structure in Euclidean space [35, 45], few attempts also explore the potential of Hyperbolic space learning [30, 43]. However, despite the effectiveness of social graph modeling, current research often overlooks the challenge of social noise. Despite the effectiveness of social graph modeling, current studies rarely notice the social noise problem. Directly using the observed social graph may lead to sub-optimal results. In this work, we propose a novel social denoising method to alleviate the effect of task-irrelevant social structures in recommendations.

## 2.2 Recommendation Denoising

Recommender systems are usually built on clean label assumption, but it's hard to guarantee in real-world datasets. Therefore, recommendation denoising methods are designed to collaborative filtering [10, 32, 41], multimodal recommendation [11, 22, 40], and social recommendation [15, 24, 42] scenarios. Specifically, researchers design social denoising methods to reduce the effect of redundant social structures [15, 27, 42, 46]. Among these, ESRF leverages the adversarial training technique to generate potential social relations and filter unstable part [46]. GDSMR [27] proposes a distilled social graph based on progressive preference-guided social denoising. ShaRe [15] designs a social graph rewriting method with social-enhanced contrastive learning optimization. GBSR [42] introduces

the information bottleneck principle to guide the social denoising process, which can preserve the minimal yet sufficient social structure. In summary, current solutions aim to filter out the noisy part of social networks, but identifying the noisy part is challenging due to a lack of labels. Instead of removing the noise part directly, we focus on learning stable user preferences across diverse environments (which are full of various noise). If the model can survive in various noisy environments, it indeed captures the invariant social structure for recommendation tasks. Thus, we introduce the invariant learning principle to tackle the social noise problem.

## 2.3 Invariant Learning and Applications

Invariant learning has emerged as a promising approach to address the challenges posed by Out-Of-Distribution (OOD) generalization in machine learning [1, 5, 18, 48]. The core idea behind invariant learning is to uncover stable feature representations that remain robust under varying environmental conditions, thereby enhancing the reliability of label prediction models. This framework is based on the assumption that the data contains both spurious and invariant features, where spurious features vary unpredictably across different environments, while invariant features remain stable and consistent. The primary objective of invariant learning methods is to optimize models using Invariant Risk Minimization (IRM) by introducing an invariance-based penalty during optimization, which discourages reliance on spurious features and encourages the model to learn stable, generalizable patterns. Invariant learning has been successfully applied across a wide range of machine learning tasks, including image classification [1, 5], graph learning [4, 19, 38], and recommender systems [2, 6, 31, 33, 49]. Especially for recommender systems, [2, 6] learn invariant multimedia representation to improve the robustness of the general and cold-start recommendation scenarios. Besides, [33, 49] leverage the invariant learning to address popularity debiasing, and [31] captures the invariant KG subgraph to enhance recommendations. In this work, we introduce the invariant learning principle to the social recommendation, intending to capture stable user preferences that persist across diverse social environments.

### 3 PRELIMINARY

#### 3.1 Graph-based Social Recommendation

In the fundamental social recommender, there are two kinds of entities: a userset  $U$  ( $|U| = M$ ) and an itemset  $V$  ( $|V| = N$ ). Users present two kinds of behaviors, user-user social behaviors (e.g., follow, favor, group) and user-item interaction behaviors (e.g., click, view, purchase). We use the matrix  $\mathbf{S} \in \mathbb{R}^{M \times M}$  to describe the user-user social relation matrix, where each element  $s_{ab} = 1$  if user  $b$  has social behavior with user  $a$ , otherwise  $s_{ab} = 0$ . Meanwhile, we use the matrix  $\mathbf{R} \in \mathbb{R}^{M \times N}$  to describe the user-item interaction behaviors, where each element  $r_{ai} = 1$  if user  $a$  interacted with item  $i$ . Given the userset  $U$ , itemset  $V$ , user-user social relations  $\mathbf{S}$  and user-item interactions  $\mathbf{R}$ , we construct the user-item interaction graph  $\mathcal{G}_r = \{U \cup V, \mathcal{E}_r = \{(a, i) | r_{ai} = 1\}\}$  and the user-user social graph  $\mathcal{G}_s = \{U, \mathcal{E}_s = \{(a, b) | s_{ab} = 1\}\}$ . Graph-based social recommendation aims to infer the unknown preference of user  $a$  to item  $i$ :  $\hat{r}_{ai} = f_\phi(a, i, \mathcal{G}_r, \mathcal{G}_s)$ , where  $f_\phi$  denotes the GNN formulation with parameters  $\phi$ . The traditional learning paradigm focuses on Empirical Risk Minimization (ERM):

$$\text{Min} : \mathbb{E}_{(a,i,r_{ai}) \sim \mathcal{P}} \mathcal{L}(r_{ai}; f_\phi(a, i, \mathcal{G}_r, \mathcal{G}_s)), \quad (1)$$

where  $\mathcal{P}$  denotes the distribution of training data. However, the observed social networks are inevitably noisy with task-irrelevant relations, and directly using them will lead to sub-optimal results [15, 42]. Therefore, tackling the social noise issue is an urgent need to enhance recommendations.

#### 3.2 Graph Invariant Learning

Invariant learning is the classic branch of Out-Of-Distribution (OOD) generalization methods, tackling the spurious correlations in machine learning tasks [1, 5, 18]. Graph invariant learning focuses on Graph OOD generalization, targeting the pursuit of invariant subgraph across diverse noisy environments [4, 19, 37, 39]. Typically, graph invariant learning methods have the following assumption:

**ASSUMPTION 1.** *Given the input graph  $\mathcal{G}$ , there exists a subgraph  $\mathcal{G}^I \in \mathcal{G}$  satisfy the invariance assumption: (1) Sufficiency condition:  $y = f(\mathcal{G}^I) + \epsilon$ , where  $y$  is the prediction target and  $\epsilon$  is an independent noise. (2) Invariance condition: for any two environments  $e, e' \in \mathcal{E}$ ,  $p(y|\mathcal{G}^I, e) = p(y|\mathcal{G}^I, e')$ , where  $\mathcal{E}$  denotes all potential environments.*

This assumption means that the invariant subgraph involves sufficient information from the input graph, which can build a robust prediction model for unknown distributions. To achieve this goal, an Invariant Risk Minimization (IRM) based optimization has been proposed [1, 18], which is defined as follows:

$$\text{Min} : \mathbb{E}_{e \in \mathcal{E}} \mathcal{L}_e + \beta \text{Var}(\{\mathcal{L}_e, e \in \mathcal{E}\}), \quad (2)$$

where the first term focuses on the empirical risk minimization in each environment (sufficiency condition), and the second term encourages the invariant risk minimization (invariance condition). The parameter  $\beta$  is set to find the trade-off between sufficiency and invariance. Inspired by graph invariant learning, we propose a novel *Social Graph Invariant Learning (SGIL)* approach to tackle the social noise issue. Imaging that there exists an invariant social subgraph for recommendation, while other redundant relations construct various noisy environments. Following the invariance

assumption, can we learn the invariant social subgraph to facilitate the recommendation task? Let's begin with the technical details.

### 4 METHODOLOGY

As illustrated in Figure 2, we present the overall framework of our proposed *Social Graph Invariant Learning (SGIL)* approach. Essentially, *SGIL* aims to learn stable yet sufficient user preferences across diverse social environments. To achieve this goal, *SGIL* elaborates two modules: Environment Simulation, and Invariant Preference Learning. Specifically, the environment simulation module aims to generate multiple social environments, and the invariant preference learning module targets to achieve the invariant user preference across environments. Next, we introduce each module in detail.

#### 4.1 Environment Simulation

Given the user-user social graph  $\mathcal{G}_s$  and user-item interaction graph  $\mathcal{G}_r$  as input, we first simulate diverse noisy environments. Here, we define  $K$  social environments  $\mathcal{E} = \{e_k : 1 \leq k \leq K\}$ . Consequently, we define  $K$  independent environment generators  $\{g_\theta^k\}_{k=1}^K$  with parameters  $\{\theta^1, \theta^2, \dots, \theta^K\}$ . We take  $k^{th}$  social environment as an example to illustrate how to construct it. Formally, constructing  $k^{th}$  social environment equals extracting  $k^{th}$  social subgraph  $\mathcal{G}_s^k$ , each edge  $(a, b)$  will be preserved with a probability  $\rho_{ab}$ . Thus, we have the extracted subgraph structure  $\mathcal{S}^k$ :

$$\mathcal{S}^k = g_\theta^k(U, \mathcal{G}_r, \mathcal{G}_s) = \{s_{ab} \odot \rho_{ab}^k\}, \quad (3)$$

where  $\rho_{ab}^k \sim \text{Bern}(w_{ab}^k) + \epsilon$  denotes that each edge  $\langle u_a, u_b \rangle$  will be dropped with the probability  $1 - w_{ab}^k + \epsilon$ .

Following the previous work [42], we also add the bias term  $\epsilon = 0.5$  to guarantee that the confidences of observed social relations are higher than those non-observed. Following the previous studies [42], we feed users' collaborative singles into parameter  $w_{ab}^k$  calculation. Let  $\mathbf{E}^U \in \mathbb{R}^{M \times d}$  and  $\mathbf{E}^I \in \mathbb{R}^{N \times d}$  denote user and item preference representations, which learned from the input data  $\{\mathcal{G}_r, \mathcal{G}_s\}$ . We parameterize the distribution parameter  $w_{ab}^k$  as follows:

$$w_{ab}^k = (g_\theta^k(\mathbf{e}_a, \mathbf{e}_b)), \quad (4)$$

where  $\mathbf{e}_a$  and  $\mathbf{e}_b$  denote user  $a$ 's and user  $b$ 's preference representations, respectively. Besides, we implement the generator  $g_\theta^k(\cdot)$  through two-layer MLPs with parameters  $\theta^k$ . After the above learning process, we can obtain  $k^{th}$  social subgraph  $\mathcal{S}^k$ . However, the discrete Bernoulli distribution is not differentiable. Gumbel-Softmax operation is the popular solution to relax the learning process:

$$\text{Bern}(w_{ab}) = \text{sigmoid}(\log(\delta/(1 - \delta) + w_{ab})/t), \quad (5)$$

where  $\delta \sim U(0, 1)$ , and  $t \in \mathbb{R}^+$  is the temperature parameter (we set  $t=0.2$  in our experiments).

After the re-parameterization operation, the discrete Bernoulli distribution is transferred to a differentiable function. Repeating  $K$  times subgraph extraction  $\{\mathcal{S}^1, \mathcal{S}^2, \dots, \mathcal{S}^K\}$ , we can simulate  $K$  social environments  $\{\mathcal{G}_s^1, \mathcal{G}_s^2, \dots, \mathcal{G}_s^K\}$ .

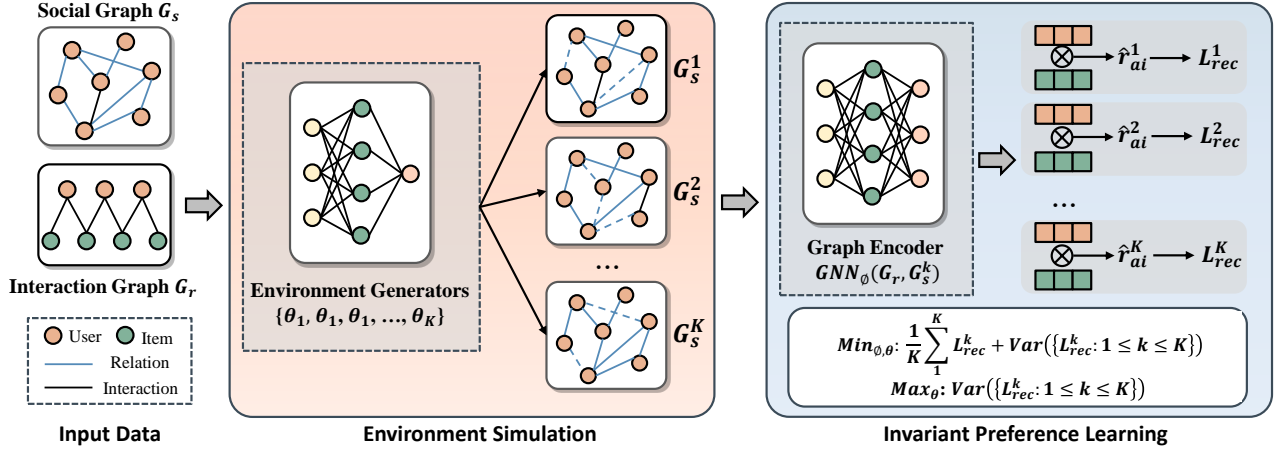


Figure 2: Overview of our proposed SGIL framework, consists of two elaborated modules. (1) Environment Simulation: aiming to generate multiple social environments; (2) Invariant Preference Learning: targeting to achieve the invariant user preference across environments.

## 4.2 Invariant Preference Learning

Given the simulated social environments, SGIL then learns the invariant preference across diverse environments. The invariant learning module consists of three components: graph-based preference learning, empirical risk minimization, and overall optimization with invariance.

**4.2.1 Graph-based Preference Learning.** Due to the advanced representation ability of GNNs, here we use GNNs as encoders to capture users' preferences. Specifically, we employ LightGCN-S [9, 42] as the encoder to obtain user preference embeddings. We first formulate the user-item interactions and the user-user social relations in  $k^{th}$  environment as a heterogeneous graph  $\mathcal{G}^k = \{U \cup V, A^k\}$ , where  $U \cup V$  denotes the set of nodes, and  $A^k$  is the adjacent matrix defined as follows:

$$A^k = \begin{bmatrix} S^k & R \\ R^T & \mathbf{0}^{N \times N} \end{bmatrix}. \quad (6)$$

Let  $\mathbf{P}^0 \in \mathbb{R}^{M \times d}$  and  $\mathbf{Q}^0$  denote the initialized user and item embeddings, where  $d$  is the embedding dimension. Then, we iteratively update node embeddings through graph convolutions:

$$\begin{bmatrix} \mathbf{P}^{l+1} \\ \mathbf{Q}^{l+1} \end{bmatrix} = \mathbf{D}_k^{-\frac{1}{2}} \mathbf{A}^k \mathbf{D}_k^{-\frac{1}{2}} \times \begin{bmatrix} \mathbf{P}^l \\ \mathbf{Q}^l \end{bmatrix}, \quad (7)$$

where  $\mathbf{D}_k$  is the degree matrix of graph  $\mathcal{G}^k$ . When stacking  $L$  graph convolution layers, we have the final node embeddings through the avg-pooling operation:

$$\mathbf{P} = \frac{1}{L+1} \sum_{l=0}^L \mathbf{P}^l, \mathbf{Q} = \frac{1}{L+1} \sum_{l=0}^L \mathbf{Q}^l. \quad (8)$$

After obtaining the learned node representations in the  $k^{th}$  social environment, we can infer the propensity that user  $a$  will interacts with item  $i$  by an inner product:  $\hat{r}_{ai} = \langle \mathbf{p}_a^k, \mathbf{q}_i^k \rangle$ . All the above process are summarized as  $\hat{r}_{ai}^k = f_\phi^k(a, i, \mathcal{G}_r, \mathcal{G}_s^k)$ , where  $\phi = \{\mathbf{P}^0, \mathbf{Q}^0\}$  denote the initialized embedding parameters.

**4.2.2 Empirical Risk Minimization.** Following the above use preference learning procedure, we obtain  $K$  predictions across all environments. Next, we implement the sufficiency condition in Assumption 1. Given the predictions in  $k^{th}$  environment, the empirical risk loss is defined as follows:

$$\mathcal{L}_{rec}^k = \sum_{a \in U, i \in V} \mathcal{L}(\sigma(f_\phi^k(a, i, \mathcal{G}_r, \mathcal{G}_s^k)), \mathbf{r}_{ai}). \quad (9)$$

Although this point-wise loss satisfies the sufficiency condition, in practice, we find that it's hard to guarantee the invariance condition. The reason is that the user's interaction behaviors are not IID data [3], directly aligning the user's preference to a specific item under various environments can not guarantee invariant user preference. In other words, we must align each user preference distribution across all environments. Therefore, we rewrite the empirical risk loss as:

$$\mathcal{L}_{rec}^k = \sum_{a \in U} \mathcal{L}(\sigma(f_\phi^k(a, V, \mathcal{G}_r, \mathcal{G}_s^k)), \mathbf{R}_a). \quad (10)$$

Next, we introduce how to optimize the above loss function. Let  $V_a(|V_a| = d_a)$  denote each user  $a$ 's interacted items, then we have the prior information:  $p(r = 1) = \frac{d_a}{N}$  and  $p(r = 0) = \frac{N - d_a}{N}$ . Next, we analyze the interaction probability as follows:

$$\begin{aligned} p(r = 1|a, i) &= \frac{p(r = 1)p(a, i|r = 1)}{\sum_{r \in \{0, 1\}} p(r)p(a, i|r)} \\ &= \frac{\frac{d_a}{N} p(a, i|r = 1)}{\frac{d_a}{N} p(a, i|r = 1) + \frac{N - d_a}{N} p(a, i|r = 0)} \\ &= \frac{p(a, i|r = 1)}{p(a, i|r = 1) + \frac{N - d_a}{d_a} p(a, i|r = 0)} \\ &= \frac{p(a, i|r = 1)}{p(a, i|r = 1) + Cp(a, i|r = 0)}, \end{aligned} \quad (11)$$

where  $C = \frac{N - d_a}{d_a}$  is the negative sample ratio. In  $k^{th}$  social environment, we infer the probability  $p(r = 1|a, i) = \delta(f_\phi^k(a, i, \mathcal{G}_r, \mathcal{G}_s^k))$ ,

where  $\delta(\cdot)$  is the sigmoid function. For the sake of clarity, we use  $f_\phi^k(a, i)$  to replace  $f_\phi^k(a, i, \mathcal{G}_r, \mathcal{G}_s^k)$ , then we define the density ratio as:

$$\frac{p(a, i|r=1)}{p(a, i|r=0)} = C \frac{p(r=1|a, i)}{1 - p(r=1|a, i)} = C \exp(f_\phi^k(a, i)). \quad (12)$$

In practice, directly optimizing the scores of full items is consuming, so we employ the popular batch training strategy. Previous works demonstrate that the more negative samples, the better the conditional probability estimation  $p(r=1|a, i)$ . For batch training data with  $\mathcal{I}_b$  items, supposing there are one positive sample  $(a, i)$  and  $|\mathcal{I}_b| - 1$  negative samples  $(a, j)$  in the training data, the goal is to maximize the following probability:

$$\begin{aligned} P &= \frac{p(a, i|r=1) \prod_{j \neq i, j \in \mathcal{I}_b} p(a, j|r=0)}{\sum_{i \in \mathcal{I}_b} p(a, j|r=1) \prod_{w \neq j, w \in \mathcal{I}_b} p(a, w|r=0)} \\ &= \frac{p(a, i|r=1)/p(a, i|r=0)}{\sum_{j \in \mathcal{I}_b} p(a, j|r=1)/p(a, j|r=0)} \\ &= \frac{\exp(f_\phi^k(a, i))}{\sum_{j \in \mathcal{I}_b} \exp(f_\phi^k(a, j))}. \end{aligned} \quad (13)$$

We use the popular scaled cosine similarity [34] to implement  $f_\phi^k(a, i)$ , thus we can obtain the final optimization objective of ERM in each environment:

$$\mathcal{L}_{rec}^k = -\frac{1}{|\mathcal{D}|} \sum_{(a, i) \in \mathcal{D}} \frac{\exp(\langle p_a^k, q_i \rangle / \tau)}{\sum_{j \in \mathcal{I}_b} \exp(\langle p_a^k, q_j^k \rangle / \tau)}. \quad (14)$$

**4.2.3 Overall Optimization with Invariance.** Given the above optimization objective of ERM in each social environment, we then combine the sufficiency condition and invariance condition for optimization. Specifically, we achieve the overall optimization objective with a variance-based regularizer:

$$\text{Min} : \frac{1}{K} \sum_{k=1}^K \mathcal{L}_{rec}^k + \beta \text{Var}(\{\mathcal{L}_{rec}^k : 1 \leq k \leq K\}), \quad (15)$$

where the first term focuses on the empirical risk minimization in each social environment (sufficiency condition), and the second term encourages the invariant risk minimization (invariance condition), and  $\beta$  is the trade-off parameter. To enhance the robustness of *SGIL* in more diverse noisy environments, we further conduct the environment exploration based on an adversarial training strategy. Specifically, we iteratively optimize the following two objectives:

$$\phi^*, \theta^* = \arg \min_{\phi, \theta} \frac{1}{K} \sum_{k=1}^K \mathcal{L}_{rec}^k + \beta \text{Var}(\{\mathcal{L}_{rec}^k : 1 \leq k \leq K\}) \quad (16)$$

$$\theta^* = \arg \max_{\theta} \text{Var}(\{\mathcal{L}_{rec}^k : 1 \leq k \leq K\}). \quad (17)$$

Considering the instability of adversarial training, we select a cross-batch training strategy. We will conduct the exploration process every  $T$  batch. After cross-environment optimization, we can obtain a stable user preference for the recommendation. In the inference stage, we first compute the mean representations of users and items

---

**Algorithm 1:** The algorithm of *SGIL*


---

- 1: **Input:** Userset  $U$ , Itemset  $V$ , user-item interactions  $R$ , user-user social relations  $S$ , the number of social environments  $K$ , and observation bias  $\epsilon$ ;
  - 2: **Output:** Optimal graph-based social recommender  $\mathcal{G}_{\theta, \phi}^*(\cdot)$ , where  $\theta$  denote MLPs used in environment generators and  $\phi = \{\mathbf{P}^0, \mathbf{Q}^0\}$  denote embedding matrices;
  - 3: **while** not converged **do**
  - 4:   Sample a batch training data  $\mathcal{D}$ ;
  - 5:   **# Environment Simulation #**
  - 6:   Compute the edge drop probability (Eq. (4));
  - 7:   Obtain edge drop distribution by Gumbel-Softmax (Eq.(5));
  - 8:   Extract  $K$  social subgraphs (Eq. (3));
  - 9:   **# Invariant Preference Learning #**
  - 10:   Learn user preference in each environment (Eq.(6)-Eq.(8));
  - 11:   Obtain ERM loss in each environment (Eq.(14));
  - 12:   Obtain overall loss across all environments (Eq.(15));
  - 13:   **# Adversarial Training #**
  - 14:   Update model parameters  $\theta, \phi$  according to invariance based loss minimization (Eq.(16));
  - 15:   Update environment generator parameters  $\theta$  according to invariance loss maximization (Eq.(17));
  - 16: **end while**
  - 17: **return** the optimal  $\mathcal{G}_{\theta, \phi}^*(\cdot)$
- 

across  $K$  social environments:

$$\mathbf{U} = \frac{1}{K} \sum_{k=1}^K \mathbf{P}^k, \mathbf{V} = \frac{1}{K} \sum_{k=1}^K \mathbf{Q}^k. \quad (18)$$

Then, we infer the interaction score based on the inner product:  $\hat{r}_{ai} = \langle \mathbf{u}_a, \mathbf{v}_i \rangle$ . We summarize the learning process of *SGIL* as Algorithm 1.

### 4.3 Model Discussion

In this section, we analyze the proposed *SGIL* from both space and time complexity.

**4.3.1 Space Complexity.** As illustrated in Algorithm 1, the parameters of *SGIL* consist of two components: the embedding parameters  $\phi$  and the environment generator parameters  $\theta$ . The embedding parameters  $\phi \in \mathbb{R}^{(M+N) \times d}$  are the same as those used in general social recommenders (such as DiffNet[36] and GBSR [42]). In addition, the environment generator parameters  $\theta$  consist of  $K$  individual MLPs ( $\mathbb{R}^{2d \times d} + \mathbb{R}^{d \times 1}$ ). Since the embedding size  $d$  is much smaller than node number  $M + N$ , the additional parameters introduced by *SGIL* is negligible.

**4.3.2 Time Complexity.** The computational overhead introduced by *SGIL* mainly consists of three components: environment simulation, invariant preference learning, and adversarial training. For environment simulation, we first obtain the encoded node representations using GCNs, and then employ  $K$  individual MLPs to generate diverse environments. For preference learning, we learn user and item representations within each environment and optimize them through an invariance-based objective. Finally, we adopt



an across-batch adversarial training strategy to iteratively update the parameters. Compared to the backbone model (LightGCN-S), *SGIL* incurs approximately  $K$  times the computational cost during the preference learning stage, and the employed softmax-based loss [34] is computationally more intensive than the traditional BPR loss [28]. Fortunately, in our experiments, the optimal value of  $K$  is typically less than 5, and *SGIL* converges much faster than LightGCN-S. Overall, the total training time of *SGIL* is approximately 1 to 3 times that of LightGCN-S.

**Table 1: The statistics of three datasets.**

Dataset	Douban-Book	Yelp	Epinions
Users	13,024	19,593	18,202
Items	22,347	21,266	47,449
Interactions	792,062	450,884	298,173
Social Relations	169,150	864,157	381,559
Interaction Density	0.048%	0.034%	0.035%
Relation Density	0.268%	0.206%	0.115%

## 5 EXPERIMENTS

### 5.1 Experimental Settings

**5.1.1 Datasets.** Following [42], we select three public social recommendation datasets to conduct our empirical studies: Douban-Book, Yelp, and Epinions. Each dataset includes user-user social links and user-item interactions. For our experiments, we randomly sample 80% of the interactions as training data, with the remaining 20% reserved as test data. The detailed statistics of all datasets are summarized in Table 1.

**5.1.2 Baselines and Evaluation Metrics.** We compare *SGIL* with state-of-the-art methods. The baselines can be divided into two groups: 1) Graph-based social recommendations: GraphRec [7], DiffNet++ [35], SocialLGN [20], LightGCN-S [9]; 2) Social graph denoising methods: ESRF [46], GDMSR [27], SEPT [47], ShaRe [15], GBSR [42], which are list as follows:

- **LightGCN** [9]: is the SOTA graph-based collaborative filtering method, which simplifies GCNs by removing the redundant feature transformation and non-linear activation components for ID-based recommendation.
- **LightGCN-S** [9, 42]: extends LightGCN to graph-based social recommendation, that each user’s neighbors include their interacted items and linked social users.
- **GraphRec** [7]: is a classic graph-based social recommendation method, it incorporates user opinions and user two kinds of graphs for preference learning.
- **DiffNet++** [35]: is the SOTA graph-based social recommendation method, it recursively formulates user interest propagation and social influence diffusion process with a hierarchical attention mechanism.
- **SocialLGN** [20]: propagates user representations on both user-item interactions graph and user-user social graph with light graph convolutional layers, and fuses them for recommendation.
- **ESRF** [46]: generates alternative social neighbors and further performs neighbor denoising with adversarial training.

- **GDMSR** [27]: designs the robust preference-guided social denoising to enhance graph-based social recommendation, it only remains the informative social relations according to preference confidences.
- **SEPT** [47]: is the SOTA self-supervised social recommendation method, it incorporates social relations to contrastive learning.
- **SHaRe** [15]: challenges the low social homophily and proposes a social graph rewriting strategy with relation-based contrastive learning enhancement.
- **GBSR** [42]: devises a model-agnostic social denoising method, which leverages the information bottleneck principle to remove the redundant social structure for recommendation tasks.

The evaluation metrics are Recall@N and NDCG@N [14], which have been widely used in ranking-based recommender systems. Following the mainstream evaluation process [50], we adopt a full-ranking protocol that all non-interacted items are candidates for evaluation.

**5.1.3 Implementation Details.** All embedding-based methods are initialized by a normal initializer with a mean value of 0 and a standard variance of 0.01, the embedding size is fixed to 64. For model optimization, we use the Adam optimizer [16] with a learning rate of 0.001 and a batch size of 2048. We empirically set the GCN layer  $L = 3$  and the temperature coefficient  $\tau = 0.2$  for all datasets. For the edge observation bias, we follow the previous work [42] and set the bias  $\epsilon = 0.5$  for all datasets. For the environment number  $K$ , the variance-based penalty coefficient  $\beta$ , we carefully search the best parameters for each dataset and report the detailed comparisons. For the adversarial training step, we set  $T = 20$  on both the Douban-Book and Yelp datasets, and  $T = 3$  on the Epinions datasets. For all baselines, we refer to the original parameters and carefully fine-tune them for fair comparisons.

### 5.2 Performance Comparisons

Following the current SOTA method [42], we implement all social denoising methods on the LightGCN-S backbone. As shown in Table 2, we have the following observations:

- **Compared to LightGCN, graph-based social recommendation methods show modest improvements on most datasets.** For instance, DiffNet++ achieves a 2.24% gain in NDCG@20 on the Yelp dataset. However, this is not universally observed; all graph-based social recommendation methods experience performance degradation on the Douban-Book dataset. These results highlight an important observation: while social graphs provide additional information, graph-based social recommendation methods do not consistently outperform LightGCN in terms of performance. This suggests that directly incorporating social graphs may sometimes hinder recommendation performance, emphasizing the need to eliminate redundant social relations to improve outcomes.
- **Social denoising methods outperform traditional graph-based methods.** Compared to the backbone model (LightGCN-S), we observe that social denoising methods achieve stable improvements across all three datasets. Whether leveraging adversarial training (ESRF), preference-guided graph filtering (GDMSR), self-supervised signals (SEPT and SHaRe), or the graph information

**Table 2: Overall performance comparisons on three benchmarks ( $p$ -value < 0.05). The best performance is highlighted in bold and the second is highlighted by underlines.**


Models	Douban-Book				Yelp				Epinions			
	R@10	N@10	R@20	N@20	R@10	N@10	R@20	N@20	R@10	N@10	R@20	N@20
LightGCN	0.1039	0.1195	0.1526	0.1283	0.0698	0.0507	0.1081	0.0623	0.0432	0.0314	0.0675	0.0385
GraphRec	0.0971	0.1145	0.1453	0.1237	0.0672	0.0485	0.1077	0.0607	0.0436	0.0315	0.0681	0.0387
DiffNet++	0.1010	0.1184	0.1489	0.1270	0.0707	0.0516	0.1114	0.0640	0.0468	0.0329	0.0727	0.0406
SocialLGN	0.1034	0.1182	0.1527	0.1274	0.0681	0.0507	0.1059	0.0620	0.0416	0.0307	0.0634	0.0371
LightGCN-S	0.1021	0.1187	0.1506	0.1281	0.0714	0.0529	0.1126	0.0651	0.0477	0.0347	0.0716	0.0417
ESRF	0.1042	0.1199	0.1534	0.1301	0.0718	0.0526	0.1123	0.0645	0.0462	0.0329	0.0727	0.0406
GDMSR	0.1026	0.1001	0.1538	0.1245	0.0739	0.0535	0.1148	0.0658	0.0461	0.0326	0.0721	0.0414
SEPT	0.1094	0.1300	0.1592	0.1382	0.0749	0.0553	0.1176	0.0682	0.0457	0.0341	0.0724	0.0416
SHaRe	0.1050	0.1243	0.1544	0.1325	0.0724	0.0534	0.1138	0.0658	0.0490	0.0354	0.0771	0.0438
GBSR	<u>0.1189</u>	<u>0.1451</u>	<u>0.1694</u>	<u>0.1523</u>	<u>0.0805</u>	<u>0.0592</u>	<u>0.1243</u>	<u>0.0724</u>	<u>0.0529</u>	<u>0.0385</u>	<u>0.0793</u>	<u>0.0464</u>
SGIL	<b>0.1303</b>	<b>0.1567</b>	<b>0.1809</b>	<b>0.1627</b>	<b>0.0872</b>	<b>0.0654</b>	<b>0.1323</b>	<b>0.0787</b>	<b>0.0565</b>	<b>0.0409</b>	<b>0.0844</b>	<b>0.0489</b>
Impro.	9.59%	7.99%	6.79%	6.83%	8.93%	8.32%	6.44%	8.70%	6.81%	6.23%	6.43%	5.39%

bottleneck principle (GBSR), these methods demonstrate remarkable performance enhancements. Notably, GBSR stands out as the strongest baseline among them, effectively preserving the minimal yet sufficient social structure necessary for recommendation tasks. By eliminating redundant information while retaining valuable social connections, GBSR effectively mitigates the impact of social noise, thereby enhancing recommendation accuracy. In summary, learning a denoised social graph significantly boosts the effectiveness of social recommendation, underscoring the critical role of social denoising in practical applications.

- **Our proposed SGIL consistently outperforms all baselines by a substantial margin.** Specifically, SGIL improves upon the strongest baseline (GBSR) to NDCG@10 by 7.99%, 10.47%, and 5.97% gains on the Douban-Book, Yelp, and Epinions datasets, respectively. These experimental results strongly validate the superiority of our proposed SGIL model. Compared to the backbone model (LightGCN-S), SGIL achieves significant improvements of approximately 32.01%, 23.63%, and 17.58% in terms of NDCG@20 across three benchmark datasets. Experimental results verify the effectiveness of our proposed SGIL. Compared to traditional graph-based social recommendation methods, SGIL effectively alleviates the adverse effects of social noise and enhances recommendation performance. Compared to current social denoising methods, SGIL simulates diverse social noise environments and learns invariant user preferences, providing stronger guarantees under various potential noise scenarios. This innovative approach not only improves the robustness of the recommendation system but also offers a novel solution for complex social recommendation contexts, ensuring stability and reliability when faced with social noise.

### 5.3 Ablation Study

To further analyze the effect of each component of the proposed SGIL, we conduct ablation studies on three datasets. As shown in Table 3, we compare SGIL with three corresponding variants: **SGIL-w/o EG** denotes that without the environment generator, then environment-inspired invariant learning disappears and SGIL degenerates to the backbone model (LightGCN-S+ERM). **SGIL-w/o**

**Figure 3: Recommendation performance (NDCG@20) under different degrees of noise.**

IL denotes that without the invariant learning regularizer, we only adopt the mean representations of all environments for inference. **SGIL-w/o EE** denotes that without environment exploration, we only perform graph invariant learning on the initialized environments. From Table 3, we have the following observations. First, SGIL-w/o EG exhibits the most performance degeneration, which indicates the importance of environment simulation. Environment simulation is the basis of SGIL, as the task-irrelevant relations are unknown, so SGIL simulates diverse environments and encourages preference consistency across environments to alleviate the effect of potential noise. Second, SGIL-w/o IL also performs significant degradation, which demonstrates the effectiveness of the invariance-based regularizer. Finally, SGIL-w/o EE shows a consistent performance drop. Because social noise is not directly perceived, we use an adversarial training strategy to encourage more diverse environments. The more diverse the environments the better the robustness of the proposed method. All ablation results demonstrate the effectiveness of each technical component of our proposed SGIL.

### 5.4 Detailed Analysis of SGIL


In this section, we further analyze SGIL from the following aspects: robustness to different scale noises, sparsity analysis, and parameter sensitivities.

**5.4.1 A-Robustness to different noises.** To evaluate the robustness under different degrees of noise, we conduct experiments on the



**Table 3: Ablation study of *SGIL* on three datasets.**

Models	Douban-Book				Yelp				Epinions			
	R@10	N@10	R@20	N@20	R@10	N@10	R@20	N@20	R@10	N@10	R@20	N@20
<i>SGIL</i> -w/o EG	0.1243	0.1445	0.1672	0.1495	0.0795	0.0591	0.1201	0.0713	0.0475	0.0348	0.0708	0.0417
<i>SGIL</i> -w/o IL	0.1272	0.1503	0.1748	0.1562	0.0813	0.0612	0.1229	0.0735	0.0499	0.0361	0.0749	0.0436
<i>SGIL</i> -w/o EE	0.1292	0.1562	0.1804	0.1626	0.0850	0.0643	0.1290	0.0773	0.0556	0.0410	0.0825	0.0487
<b><i>SGIL</i></b>	<b>0.1303</b>	<b>0.1567</b>	<b>0.1809</b>	<b>0.1627</b>	<b>0.0872</b>	<b>0.0654</b>	<b>0.1323</b>	<b>0.0787</b>	<b>0.0565</b>	<b>0.0409</b>	<b>0.0844</b>	<b>0.0489</b>

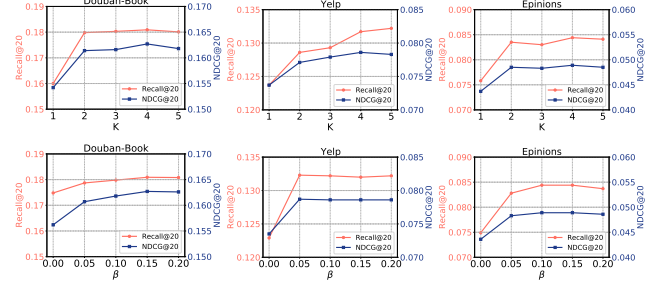


**Figure 4: Recommendation performance (NDCG@20) across different sparsity users.**

semi-synthetic datasets, in which we inject different proportions of fake social relations into the original social graph. As illustrated in Figure 3, we exhibit the recommendation performances under different degrees of noise, where the horizontal axis  $\sigma$  represents the injection fake relation ratio, the vertical axes denote NDCG@20 and relative improvements compared with the backbone model. We can find that with the increase in noise degree, all models show a performance decrease. However, social denoising methods perform higher improvements than the backbone model when the noise degree increases. This verifies that social denoising methods have better robustness for highly noisy environments. Furthermore, our proposed *SGIL* consistently shows significant improvements with various settings, especially improving the backbone model about 28% gains when the noise ratio  $\sigma = 2.0$ . Empowering the ability of graph invariant learning, *SGIL* can extract invariant social structure for robust preference learning.

**5.4.2 B-Sparsity Analysis.** Here, we investigate how the proposed *SGIL* performs for different sparsity users. As depicted in Figure 4, we report recommendation performances of *SGIL* and other social denoising methods on different sparsity user groups. Specifically, we divide all users into 'Low, Medium, High' three groups based on their historical interaction records. Then, we evaluate all methods across different groups. From Figure 4, we observe that with the increase in interaction density, each model shows significant performance improvements. Compared with the backbone model, our proposed *SGIL* consistently presents better performances in all settings. Especially, *SGIL* performs higher improvements on the sparser scenarios, because their long-tail users are more dependent on social relations, our proposed *SGIL* provides more stable user preferences in the noisy environment.

**5.4.3 C-Parameter Sensitivities.** We next analyze the influence of two core parameters in *SGIL*: the environment number  $K$  and the invariance penalty coefficient  $\beta$ . We search these two parameters carefully and report the performances with different parameters



**Figure 5: Impact of the number of environments  $K$  and invariance coefficient  $\beta$ .**

on three datasets. The parameter  $K$  controls the exploration range and the coefficient  $\beta$  determines the weight of the invariance constraint. As shown in Figure 5, we search  $K \in \{1, 2, 3, 4, 5\}$  and  $\beta \in \{0, 0.05, 0.10, 0.15, 0.20\}$ . In which,  $K = 1$  denotes that we only perform a single environment generator, then the invariance constraint disappears. Besides,  $\beta = 0$  also denotes the invariance constraint disappears, the difference compared with  $K = 0$  is the number of environments is not the same. We can observe that *SGIL* achieves the optimal results when ( $K = 4, \beta = 0.15$ ) on the Douban-Book dataset, ( $K = 4, \beta = 0.05$ ) on the Yelp dataset, and ( $K = 4, \beta = 0.10$ ) on the Epinions dataset.

## 6 CONCLUSION

In this paper, we investigate social denoising recommendations from a novel invariant learning perspective. Technically, we propose a novel *Social Graph Invariant Learning (SGIL)* approach, aiming to learn the invariant user preference across diverse noisy environments. To achieve this goal, we first simulate multiple social environments based on user preferences guidance, and then learn invariant user preferences from invariance-based optimization. Particularly, we derive the empirical risk minimization objective in each environment when considering the characterization of recommendation tasks. Furthermore, *SGIL* fully explores diverse environments with an adversarial training strategy. This strategy allows the model to adapt to a wide range of challenging conditions, thereby improving its generalization and robustness. Extensive experiments on three benchmarks verify the effectiveness of the proposed *SGIL*.

## ACKNOWLEDGMENTS

This work was supported in part by grants from the National Key Research and Development Program of China (Grant No. 2021ZD0111802), the National Natural Science Foundation of China (Grant No. U23B2031, 721881011).

## REFERENCES

- [1] Martin Arjovsky, Léon Bottou, Ishaan Gulrajani, and David Lopez-Paz. 2019. Invariant risk minimization. *arXiv preprint arXiv:1907.02893* (2019).
- [2] Haoyue Bai, Le Wu, Min Hou, Miaomiao Cai, Zhuangzhuang He, Yuyang Zhou, Richang Hong, and Meng Wang. 2024. Multimodality Invariant Learning for Multimedia-Based New Item Recommendation. In *SIGIR*. 677–686.
- [3] Longbing Cao. 2016. Non-iid recommender systems: A review and framework of recommendation paradigm shifting. *Engineering* 2, 2 (2016), 212–224.
- [4] Yongqiang Chen, Yonggang Zhang, Yatao Bian, Han Yang, MA Kaili, Binghui Xie, Tongliang Liu, Bo Han, and James Cheng. 2022. Learning causally invariant representations for out-of-distribution generalization on graphs. *NeurIPS* 35 (2022), 22131–22148.
- [5] Elliot Creager, Jörn-Henrik Jacobsen, and Richard Zemel. 2021. Environment inference for invariant learning. In *ICML*. PMLR, 2189–2200.
- [6] Xiaoyu Du, Zike Wu, Fuli Feng, Xiangnan He, and Jinhui Tang. 2022. Invariant representation learning for multimedia recommendation. In *MM*. 619–628.
- [7] Wenqi Fan, Yao Ma, Qing Li, Yuan He, Eric Zhao, Jiliang Tang, and Dawei Yin. 2019. Graph neural networks for social recommendation. In *TheWebConf*. 417–426.
- [8] Guibing Guo, Jie Zhang, and Neil Yorke-Smith. 2015. Trustsvd: Collaborative filtering with both the explicit and implicit influence of user trust and of item ratings. In *AAAI*, Vol. 29.
- [9] Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, Yongdong Zhang, and Meng Wang. 2020. Lightgcn: Simplifying and powering graph convolution network for recommendation. In *SIGIR*. 639–648.
- [10] Zhuangzhuang He, Yifan Wang, Yonghui Yang, Peijie Sun, Le Wu, Haoyue Bai, Jinqi Gong, Richang Hong, and Min Zhang. 2024. Double Correction Framework for Denoising Recommendation. *KDD* (2024).
- [11] Jun Hu, Bryan Hooi, Shengsheng Qian, Quan Fang, and Changsheng Xu. 2024. MGDCF: Distance learning via Markov graph diffusion for neural collaborative filtering. *IEEE Transactions on Knowledge and Data Engineering* 36, 7 (2024), 3281–3296.
- [12] Jun Hu, Bryan Hooi, Shengsheng Qian, Quan Fang, and Changsheng Xu. 2024. MGDCF: Distance learning via Markov graph diffusion for neural collaborative filtering. *IEEE Transactions on Knowledge and Data Engineering* 36, 7 (2024), 3281–3296.
- [13] Mohsen Jamali and Martin Ester. 2010. A matrix factorization technique with trust propagation for recommendation in social networks. In *Recsys*. 135–142.
- [14] Kalervo Järvelin and Jaana Kekäläinen. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)* 20, 4 (2002), 422–446.
- [15] Wei Jiang, Xinyi Gao, Guandong Xu, Tong Chen, and Hongzhi Yin. 2024. Challenging Low Homophily in Social Recommendation. In *TheWebConf*. 3476–3484.
- [16] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [17] Ioannis Konstas, Vassilios Stathopoulos, and Joemon M Jose. 2009. On social networks and collaborative recommendation. In *SIGIR*. 195–202.
- [18] David Krueger, Ethan Caballero, Joern-Henrik Jacobsen, Amy Zhang, Jonathan Binas, Dinghui Zhang, Remi Le Priol, and Aaron Courville. 2021. Out-of-distribution generalization via risk extrapolation (rex). In *ICML*. PMLR, 5815–5826.
- [19] Haoyang Li, Ziwei Zhang, Xin Wang, and Wenwu Zhu. 2022. Learning invariant graph representations for out-of-distribution generalization. *NeurIPS* 35 (2022), 11828–11841.
- [20] Jie Liao, Wei Zhou, Fengji Luo, Junhao Wen, Min Gao, Xiuhua Li, and Jun Zeng. 2022. SocialLGN: Light graph convolution network for social recommendation. *Information Sciences* 589 (2022), 595–607.
- [21] Hao Ma, Haixuan Yang, Michael R Lyu, and Irwin King. 2008. Sorec: social recommendation using probabilistic matrix factorization. In *CIKM*. 931–940.
- [22] Haokai Ma, Yimeng Yang, Lei Meng, Ruobing Xie, and Xiangxu Meng. 2024. Multimodal conditioned diffusion model for recommendation. In *Companion Proceedings of the ACM Web Conference 2024*. 1733–1740.
- [23] Hao Ma, Dengyong Zhou, Chao Liu, Michael R Lyu, and Irwin King. 2011. Recommender systems with social regularization. In *WSDM*. 287–296.
- [24] Wenze Ma, Yuejian Wang, Yanmin Zhu, Zhaobo Wang, Mengyuan Jing, Xuhao Zhao, Jiadi Yu, and Feilong Tang. 2024. MADM: A Model-agnostic Denoising Module for Graph-based Social Recommendation. In *WSDM*. 501–509.
- [25] Peter V Marsden and Noah E Friedkin. 1993. Network studies of social influence. *Sociological Methods & Research* 22, 1 (1993), 127–151.
- [26] Miller McPherson, Lynn Smith-Lovin, and James M Cook. 2001. Birds of a feather: Homophily in social networks. *Annual review of sociology* 27, 1 (2001), 415–444.
- [27] Yuhan Quan, Jintao Ding, Chen Gao, Lingling Yi, Depeng Jin, and Yong Li. 2023. Robust Preference-Guided Denoising for Graph based Social Recommendation. In *TheWebConf*. 1097–1108.
- [28] Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *UAI*. 452–461.
- [29] Jiliang Tang, Xia Hu, and Huan Liu. 2013. Social recommendation: a review. *Social Network Analysis and Mining* 3 (2013), 1113–1133.
- [30] Hao Wang, Defu Lian, Hanghang Tong, Qi Liu, Zhenya Huang, and Enhong Chen. 2021. Hypersorec: Exploiting hyperbolic user and item representations with multiple aspects for social-aware recommendation. *TOIS* 40, 2 (2021), 1–28.
- [31] Shuyao Wang, Yongduo Sui, Chao Wang, and Hui Xiong. 2024. Unleashing the Power of Knowledge Graph for Recommendation via Invariant Learning. In *TheWebConf*. 3745–3755.
- [32] Wenjie Wang, Fuli Feng, Xiangnan He, Liqiang Nie, and Tat-Seng Chua. 2021. Denoising implicit feedback for recommendation. In *WSDM*. 373–381.
- [33] Zimu Wang, Yue He, Jiahuo Liu, Wencho Zou, Philip S Yu, and Peng Cui. 2022. Invariant preference learning for general debiasing in recommendation. In *KDD*. 1969–1978.
- [34] Jiancan Wu, Xiang Wang, Xingyu Gao, Jiawei Chen, Hongcheng Fu, and Tianyu Qiu. 2024. On the effectiveness of sampled softmax loss for item recommendation. *TOIS* 42, 4 (2024), 1–26.
- [35] Le Wu, Junwei Li, Peijie Sun, Richang Hong, Yong Ge, and Meng Wang. 2020. Diffnet++: A neural influence and interest diffusion network for social recommendation. *TKDE* (2020).
- [36] Le Wu, Peijie Sun, Yanjie Fu, Richang Hong, Xiting Wang, and Meng Wang. 2019. A neural influence diffusion model for social recommendation. In *SIGIR*. 235–244.
- [37] Qitian Wu, Hengrui Zhang, Junchi Yan, and David Wipf. 2022. Handling distribution shifts on graphs: An invariance perspective. *ICLR* (2022).
- [38] Ying-Xin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. 2022. Discovering invariant rationales for graph neural networks. *ICLR* (2022).
- [39] Donglin Xia, Xiao Wang, Nian Liu, and Chuan Shi. 2024. Learning invariant representations of graph neural networks via cluster generalization. *NeurIPS* 36 (2024).
- [40] Yonghui Yang, Le Wu, Zhuangzhuang He, Zhengwei Wu, Richang Hong, and Meng Wang. 2025. Less is More: Information Bottleneck Denoised Multimedia Recommendation. *arXiv preprint arXiv:2501.12175* (2025).
- [41] Yonghui Yang, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. 2021. Enhanced graph learning for collaborative filtering via mutual information maximization. In *SIGIR*. 71–80.
- [42] Yonghui Yang, Le Wu, Zihan Wang, Zhuangzhuang He, Richang Hong, and Meng Wang. 2024. Graph Bottlenecked Social Recommendation. *KDD* (2024).
- [43] Yonghui Yang, Le Wu, Kun Zhang, Richang Hong, Hailin Zhou, Zhiqiang Zhang, Jun Zhou, and Meng Wang. 2023. Hyperbolic Graph Learning for Social Recommendation. *IEEE TKDE* (2023).
- [44] Yonghui Yang, Zhengwei Wu, Le Wu, Kun Zhang, Richang Hong, Zhiqiang Zhang, Jun Zhou, and Meng Wang. 2023. Generative-contrastive graph learning for recommendation. In *SIGIR*. 1117–1126.
- [45] Junliang Yu, Hongzhi Yin, Min Gao, Xin Xia, Xiangliang Zhang, and Nguyen Quoc Viet Hung. 2021. Socially-aware self-supervised tri-training for recommendation. In *SIGKDD*. 2084–2092.
- [46] Junliang Yu, Hongzhi Yin, Jundong Li, Min Gao, Zi Huang, and Lizhen Cui. 2020. Enhancing social recommendation with adversarial graph convolutional networks. *IEEE TKDE* 34, 8 (2020), 3727–3739.
- [47] Junliang Yu, Hongzhi Yin, Jundong Li, Qinyong Wang, Nguyen Quoc Viet Hung, and Xiangliang Zhang. 2021. Self-supervised multi-channel hypergraph convolutional network for social recommendation. In *TheWebConf*. 413–424.
- [48] Linan Yue, Qi Liu, Ye Liu, Weibo Gao, and Fangzhou Yao. 2025. Learning from shortcut: a shortcut-guided approach for explainable graph learning. *Frontiers of Computer Science* 19, 8 (2025), 198338.
- [49] An Zhang, Jingnan Zheng, Xiang Wang, Yancheng Yuan, and Tat-Seng Chua. 2023. Invariant collaborative filtering to popularity distribution shift. In *TheWebConf*. 1240–1251.
- [50] Wayne Xin Zhao, Junhua Chen, Pengfei Wang, Qi Gu, and Ji-Rong Wen. 2020. Revisiting alternative experimental settings for evaluating top-n item recommendation algorithms. In *CIKM*. 2329–2332.