SGMT: Social Generating with Multiview-guided Tuning in Recommender Systems

Jianghong Ma[1*†], Changran He[1*†], Dezhao Yang[1], Tianjun Wei[2†], Haijun Zhang[1], Xiaofeng Zhang[1]

[1]Harbin Institute of Technology(Shenzhen), Shenzhen, China
[2]Nanyang Technological University, Singapore
majianghong@hit.edu.cn, 2023311c12@stu.hit.edu.cn, yangdezhaoomg@gmail.com, tjwei2-c@my.cityu.edu.hk, hjzhang@hit.edu.cn, zhangxiaofeng@hit.edu.cn

Abstract

The sparsity of user–item interactions remains a fundamental obstacle in collaborative filtering, limiting the ability of Graph Neural Network (GNN)-based recommender systems to capture high-order user relationships without incurring over-smoothing and computational overhead. Existing social recommendation approaches mitigate this by incorporating social networks, yet most rely on explicit ties and fail to construct informative links in their absence. Meanwhile, contrastive learning (CL) has shown promise in improving representation quality, but current view generation strategies, augmentation-based for robustness and nonaugmentation-based for semantic fidelity, are seldom combined, leaving their complementary potential underexplored. We propose Social Generating with Multiview-guided Tuning (SGMT), a unified framework that addresses both challenges. First, an interest-aware social generation mechanism constructs synthetic user–user links from shared interaction patterns, theoretically shown to compress collaborative paths and uncover latent high-order relations. Second, we present two complementary CL modules, Noise-augmented View and Semantic-explored View, which we theoretically prove to preferentially enhance uniformity and alignment, respectively, two fundamental objectives in CL. Experiments on three real-world datasets show that SGMT outperforms state-of-the-art baselines, validating both the theoretical analysis and the practical efficacy of our model.

Introduction

The exponential growth of online content has led to information overload, positioning recommender systems as vital tools for filtering relevant information (Han et al. 2024). Collaborative filtering (CF), based on the assumption that similar users share similar preferences (Li et al. 2024; Wu et al. 2024; Wang et al. 2022), remains a foundational approach. However, the sparsity of user-item interactions limits its effectiveness (Wu et al. 2021a). To mitigate this, recent efforts have incorporated social networks, leveraging user relationships to enhance recommendation performance (Wu et al. 2020a; Fan et al. 2020; Yang et al. 2022; Sun et al. 2025; Xiong et al. 2025), as social ties offer valuable signals for preference modeling (Sharma et al. 2022).

Graph Neural Networks (GNNs) have recently gained prominence for their ability to model higher-order dependencies in graph-structured data (Wu et al. 2022b). Given the inherent graph structure of recommender systems—such as user-item interactions and social networks—GNNs have been widely adopted in both collaborative filtering (Wu et al. 2020b; Ji et al. 2020; Wang et al. 2020) and social recommendation tasks (Sun et al. 2023; Ni et al. 2023). Despite their success, two key challenges remain:

Dataset		Flickr	Ciao	Yelp
Number of Common Interactions between Users	1	80.6%	65.9%	63.9%
	2	12.0%	17.4%	18.5%
	3	3.50%	7.01%	7.41%
Average Number of Common Interactions between Users		1.38	1.84	1.90

Table 1: Statistics w.r.t. users with common interactions in the collaborative domain.

The limitation of GNNs in sparse interactions. Despite their success, GNN-based recommendation systems face intrinsic limitations due to over-smoothing, which impairs representation quality as network depth increases and restricts the modeling of high-order dependencies within user-item graphs (Chen et al. 2020). Deeper architectures also incur prohibitive computational costs, with most models restricted to three layers (He et al. 2020). As shown in Table 1, real-world datasets such as Flickr, Ciao, and Yelp[1] exhibit high sparsity—over 60% of user pairs share only a single item interaction, with average common interactions per pair below two. Such sparsity places semantically informative connec-

[*]These authors contributed equally.
[†]Corresponding authors.

[1]Detailed descriptions of these datasets are provided in the Experiments section.

tions beyond the reach of shallow GNNs. Consider the example in Fig. 1, where predicting user u_2's preference for item i_2 depends on signals from user u_4, who alone has interacted with i_2. In the interaction graph, u_2 and u_4 are separated by a four-hop path, necessitating a fourth-order GNN, rendering the embeddings susceptible to over-smoothing. This observation motivates the need for auxiliary structural connections that reduce effective propagation depth. By generating social edges between u_2 and u_3, and between u_3 and u_4, the system can compress the original four-hop collaborative path into a two-hop path in the social graph, allowing a second-order GNN to capture high-order signals with improved efficiency and reduced degradation.

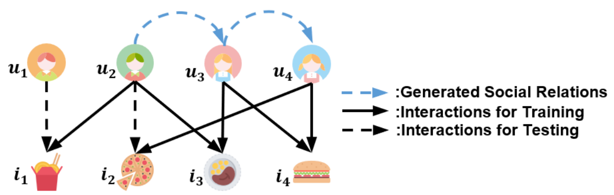

Figure 1: In the collaborative domain, predicting u_2's interaction with i_2 requires the learning of u_4's embedding via at least *fourth-order* GNN. In contrast, with the generated social graph that connects users sharing similar interaction patterns, u_2 can reach u_4 via a *second-order* user-user path.

The limitation of single view generation strategies in sparse data. To alleviate data sparsity in user-item interaction graphs, leveraging social neighbors introduces valuable user-user signals. In parallel, contrastive learning (CL) has emerged as a powerful self-supervised approach that improves representation quality by maximizing agreement between different views of the same user or item (Jing et al. 2023; Singh 2020; Cheng et al. 2022; Wei et al. 2022; Xie et al. 2022; Zang et al. 2023). Existing CL view generation strategies fall into two categories: *augmentation-based* (Yu et al. 2022; Wu et al. 2021a), which perturb the graph or embeddings to improve robustness; and *nonaugmentation-based* (He et al. 2023; Wu et al. 2024), which retain the original structure (e.g., via multi-layer representations) to preserve semantic fidelity. However, whether these two strategies can be combined in a complementary optimization manner to improve representation has not been explored.

To address the aforementioned challenges, we propose a novel framework: Social Generating with Multiview-guided Tuning (SGMT) for recommender systems. SGMT comprises three key modules, with the primary module, **Social Collaborative Filtering**, designed to address the **first** challenge. Central to this module is the interest-aware social generation component, which constructs synthetic social links between users by identifying shared interaction patterns within a predefined neighborhood. In addition, the second component, **social integrated propagation**, incorporates the generated social edges into the message-passing process, facilitating joint propagation over both social and collaborative structures. The third component, **cross-domain gated aggregation**, fuses user representations from the social and interaction domains using a gating mechanism that adaptively balances their contributions. To address the **second** challenge, SGMT further incorporates two auxiliary modules: a *Noise-augmented View*, introducing controlled perturbations in embeddings; and a *Semantic-explored View*, enforcing semantic consistency across GNN layers. We provide theoretical evidence that these views respectively promote the two key objectives in contrastive learning, uniformity and alignment, which are essential for learning high-quality representations (Wang and Isola 2020).

In summary, our main contributions are listed as follows:

- We propose a novel model that constructs social links by identifying shared user interests to compress collaborative paths. To our knowledge, this is the **first** work to ***theoretically*** show that the resulting social graph more effectively captures latent high-order relationships in the interaction graph, thereby enabling more efficient propagation of preferences among semantically aligned users.
- To refine the embeddings produced by the Social Collaborative Filtering module, we introduce two complementary views. The first uses augmentation to promote representational robustness, while the second preserves structural semantics. ***Theoretical analysis*** shows that these views preferentially promote uniformity and alignment, respectively, in the learned representations.
- We evaluate the performance of SGMT on three datasets. The experimental results reveal the superiority of our proposed model over state-of-the-art methods.

Method

Preliminaries

Let $\mathcal{U} = \{u_1, \ldots, u_n\}$ and $\mathcal{I} = \{i_1, \ldots, i_m\}$ be the user and item sets, with n and m being their respective sizes. User–item interactions are denoted as a bipartite graph $\mathcal{G}_r = (\mathcal{U} \cup \mathcal{I}, \mathcal{E}_r)$, where $\mathcal{E}_r$ is the observed interactions. The interactions can be denoted by a matrix $\mathbf{R} \in \mathbb{R}^{n \times m}$, where $r_{ui} = 1$ if user u interacts with item i, and 0 otherwise.

Overview

Fig. 2 illustrates the SGMT architecture, comprising a core *Social Collaborative Filtering* backbone and two auxiliary modules: *Noise-augmented View* and *Semantic-explored View*. The backbone includes: (1) *Interest-aware Generating*, which builds a synthetic social graph via shared interests and encodes it using Graph Attention Network (GAT) (Veličković et al. 2018); (2) *Social Integrated Propagation*, which fuses social and collaborative signals; and (3) *Cross-Domain Gated Aggregation*, which adaptively balances the two via a learnable gate. The *Noise-augmented View* introduces perturbations to encourage uniformity, while the *Semantic-explored View* enforces cross-layer consistency to enhance alignment.

Social Collaborative Filtering

Interest-aware Generating: To overcome the depth limitations imposed by over-smoothing in GNNs, we generate synthetic social links from shared user interests, allowing

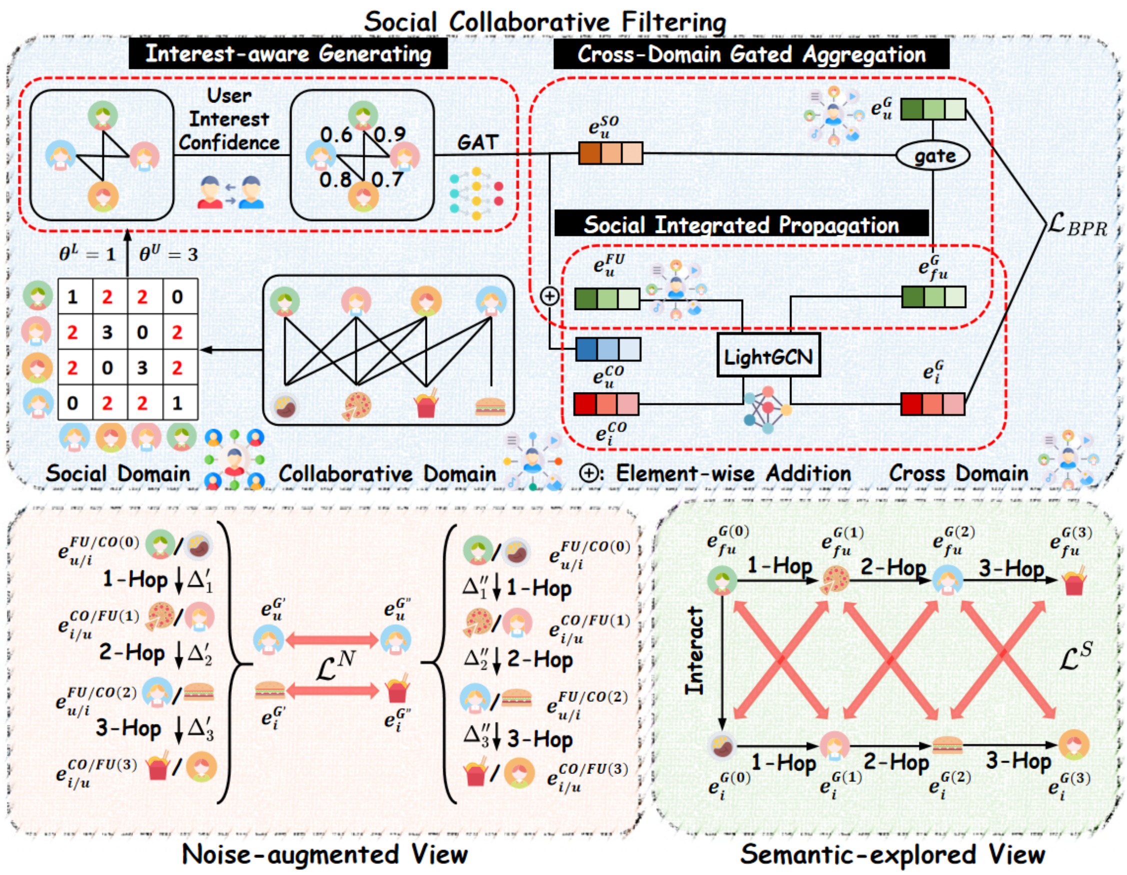

Figure 2: Illustrative of the proposed SGMT model, which comprises the core *Social Collaborative Filtering* backbone, which constructs and integrates interest-aware social links, and two auxiliary modules, the *Noise-augmented View* and the *Semantic-explored View*, which jointly refine the learned representations by balancing uniformity and alignment.

access to high-order collaborative signals without deeper propagation layers. We begin by computing the number of co-interactions between users via $\mathbf{S} = \mathbf{R} \cdot \mathbf{R}^\top$, where $\mathbf{S} \in \mathbb{R}^{n\times n}$ is the user co-interaction matrix. To form a social adjacency graph, we define two thresholds, θ^L and θ^U, and retain user pairs (u, v) with shared interactions as:

$$\mathbf{S}_{uv} = \begin{cases} 1, & \text{if } \theta^L < \mathbf{S}_{uv} < \theta^U \\ 0, & \text{if } \mathbf{S}_{uv} \leq \theta^L \quad \text{or} \quad \mathbf{S}_{uv} \geq \theta^U. \end{cases} \tag{1}$$

This filtering removes spurious connections (below θ^L) and redundant edges (above θ^U), preserving only moderately aligned user pairs likely to benefit from social augmentation.

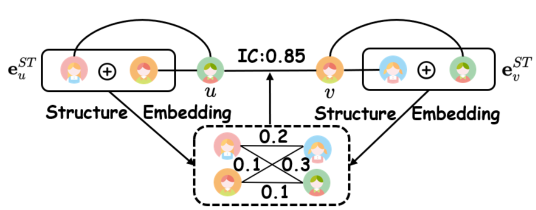

Figure 3: Illustrative example for IC of user u and v.

To refine the generated links, we assess the interest consistency of each user pair by structure-aware embeddings:

$$\mathbf{E}_U^{ST} = \mathbf{S}\mathbf{E}_U^{CO}, \tag{2}$$

where $\mathbf{E}_U^{CO} \in \mathbb{R}^{n\times d}$ denotes user embeddings from the collaborative domain. Here, $\mathbf{E}_U^{CO}$ carries user interest information, while $\mathbf{S}$ encodes the social structure.

To quantify the reliability degree of the social connection between user u and user v, we then evaluate the cosine similarity between structural embeddings and normalize it to yield an *Interest Confidence* (IC) score $\text{IC}_{u,v}$ as:

$$\text{IC}_{u,v} = \left(\frac{(\mathbf{e}_u^{ST})^\top \mathbf{e}_v^{ST}}{\|\mathbf{e}_u^{ST}\|_2 \, \|\mathbf{e}_v^{ST}\|_2} + 1\right)/2. \tag{3}$$

Figure 3 depicts the construction of interest-aware social connections. For each candidate user pair (u, v), we compute the cosine similarity between their structure-aware embeddings $\mathbf{e}_u^{ST}$ and $\mathbf{e}_v^{ST}$. This similarity is then normalized to yield the IC score. In the illustrated case, the computed IC value of 0.85 indicates a high degree of alignment in their structural preferences within the social space.

The IC is calculated as a pruning criterion: only social edges with IC above a certain threshold γ are retained, ensuring that the augmented graph includes only high-confidence social connections. This process is

$$\widetilde{\mathcal{E}}_s = \{(u, v) \in \mathcal{E} \mid \mathrm{IC}_{u,v} \geq \gamma\}, \tag{4}$$

where $\mathcal{E}$ stands for edge set of the generated social gragh $\mathcal{G}_s$ and $\widetilde{\mathcal{E}}_s$ denotes the pruned social edge set.

Subsequently, we employ a GAT to compute user embeddings in the social domain:

$$\mathbf{E}_U^{SO} = \mathrm{GAT}(\mathbf{E}_U^{ST}, \widetilde{\mathcal{G}}_s), \tag{5}$$

where $\widetilde{\mathcal{G}}_s$ is the pruned social graph, and $\mathbf{E}_U^{SO} \in \mathbb{R}^{n\times d}$ inherently captures the social relationships between users. The GAT selectively aggregates information from structurally and semantically aligned neighbors.

Social Integrated Propagation: Unlike conventional social recommendation methods that separate social and interaction signals (Sharma et al. 2024), we adopt a unified propagation scheme that embeds social information directly into collaborative filtering, by fusing user embeddings from both domains:

$$\mathbf{E}_U^{FU} = \mathbf{E}_U^{SO} \oplus \mathbf{E}_U^{CO}, \tag{6}$$

where $\oplus$ denotes element-wise addition, $\mathbf{E}_U^{FU} \in \mathbb{R}^{n\times d}$ represents the fused user embedding, which is learned through fusing information from social and interaction graphs.

To encode user-item interactions, we employ LightGCN(He et al. 2020), a simplified yet effective variant of GCNs specifically tailored for recommendation tasks. The fused user representations $\mathbf{E}_U^{FU}$ are subsequently propagated through the interaction graph:

$$\mathbf{E}_{FU}^G, \mathbf{E}_I^G = \mathrm{LightGCN}(\mathbf{E}_U^{FU}, \mathbf{E}_I^{CO}, \mathcal{G}_r), \tag{7}$$

yielding the global user and item embeddings $\mathbf{E}_{FU}^G$ and $\mathbf{E}_I^G$. This propagation captures both collaborative signals and social influence, improving representations in sparse settings.

Cross-Domain Gated Aggregation: While $\mathbf{E}_{FU}^G$ captures collaborative patterns, it is largely shaped by interaction-graph signals. To more comprehensively model preferences, we integrate the social embedding $\mathbf{E}_U^{SO}$ via a gating mechanism (Wu et al. 2022a), which adaptively balances collaborative and social contributions. The final user embedding $\mathbf{E}_U^G$ is given by:

$$\mathbf{E}_U^G = \mathrm{Gate}\left(\mathbf{E}_{FU}^G, \mathbf{E}_U^{SO}\right), \tag{8}$$

where the gating unit dynamically weighs the two inputs:

$$\begin{aligned} \mathbf{g}_u &= \sigma\left(\mathbf{W}_G^1\, \mathbf{e}_{fu}^G + \mathbf{W}_G^2\, \mathbf{e}_u^{SO}\right), \\ \mathbf{e}_u^G &= \mathbf{g}_u \odot \mathbf{e}_{fu}^G + (1 - \mathbf{g}_u) \odot \mathbf{e}_u^{SO}, \end{aligned} \tag{9}$$

where $\mathbf{g}_u \in \mathbb{R}^d$ is the gate vector, $\odot$ is element-wise multiplication, and $\sigma(\cdot)$ is the sigmoid function. The learnable matrices $\mathbf{W}_G^1$ and $\mathbf{W}_G^2$ guide the domain-specific weighting. This adaptive fusion enables $\mathbf{E}_U^G$ to combine personalized preferences with social signals, yielding more expressive and robust user representations.

Multi-View Regularization

Noise-augmented View: To improve embedding uniformity, we adopt a lightweight noise-injection strategy inspired by SimGCL (Yu et al. 2022). Instead of perturbing graph structures, we inject controlled noise into embeddings to produce two distinct views, mitigating semantic loss. For a representation matrix $\mathbf{E} \in \mathbb{R}^{x\times d}$, noise is generated as:

$$\Delta = \bar{\Delta} \odot sign\left(\mathbf{E}\right) \odot \epsilon, \text{with } \bar{\Delta} = Rand\_like\left(\mathbf{E}\right), \tag{10}$$

where $Rand\_like(\cdot)$ produces random noise matching the dimensions of $\mathbf{E}$ and is L_2-normalized. The $sign(\cdot)$ function assigns $+1$ or -1 based on the sign of each element, while ϵ controls the noise magnitude. This operation preserves the directional semantics of the original embeddings. Two perturbed views of user and item embeddings are:

$$\begin{aligned} \mathbf{E}_U^{G'}, \mathbf{E}_I^{G'} &= Encoder\left(\mathbf{E}_U^{FU} + \Delta_u', \mathbf{E}_I^{CO} + \Delta_i', \mathcal{G}_r\right), \\ \mathbf{E}_U^{G''}, \mathbf{E}_I^{G''} &= Encoder\left(\mathbf{E}_U^{FU} + \Delta_u'', \mathbf{E}_I^{CO} + \Delta_i'', \mathcal{G}_r\right), \end{aligned} \tag{11}$$

where $Encoder(\cdot)$ comprises social integration and gating components. $\mathbf{E}_U^{G'}, \mathbf{E}_U^{G''} \in \mathbb{R}^{n\times d}$ and $\mathbf{E}_I^{G'}, \mathbf{E}_I^{G''} \in \mathbb{R}^{m\times d}$ are user and item embeddings from the two generated views.

To perform contrastive learning across views, we treat $\{\mathbf{e}_u^{G'}, \mathbf{e}_u^{G''}\}$ and $\{\mathbf{e}_i^{G'}, \mathbf{e}_i^{G''}\}$ as positive pairs and $\{\mathbf{e}_u^{G'}, \mathbf{e}_v^{G''}\}$ and $\{\mathbf{e}_i^{G'}, \mathbf{e}_j^{G''}\}$ as negative samples, and the noise-augmented contrastive loss is defined by adopting the InfoNCE loss (Gutmann and Hyvärinen 2010) as:

$$\begin{aligned} \mathcal{L}^N =& \frac{1}{2}\sum_{u\in\mathcal{U}} -\log\frac{\exp(\mathbf{e}_u^{G'}\cdot\mathbf{e}_u^{G''}/\tau)}{\sum_{v\in\mathcal{U}}\exp(\mathbf{e}_u^{G'}\cdot\mathbf{e}_v^{G''}/\tau)} + \\ & \frac{1}{2}\sum_{i\in\mathcal{I}} -\log\frac{\exp(\mathbf{e}_i^{G'}\cdot\mathbf{e}_i^{G''}/\tau)}{\sum_{j\in\mathcal{I}}\exp(\mathbf{e}_i^{G'}\cdot\mathbf{e}_j^{G''}/\tau)}, \end{aligned} \tag{12}$$

where τ being the temperature parameter.

Semantic-explored View: To complement the geometric regularization of the noise-augmented view, we propose a contrastive strategy that preserves semantic fidelity without altering the input structure. This Semantic-explored View exploits the hierarchical nature of GNNs, wherein different layers encode distinct semantics (He et al. 2023): even-numbered layers capture user–user similarities, while odd-numbered layers model user–item interactions. Thus, cross-layer embeddings naturally form semantically meaningful positive pairs. Given an observed interaction (u, i), we pair the fused user embedding at the k-th layer, $\mathbf{e}_{fu}^{G^{(k)}}$, with the item embedding at the $(k+1)$-th layer, $\mathbf{e}_i^{G^{(k+1)}}$, and vice versa, forming symmetric positive pairs. In contrast, negative pairs are formed via random unobserved interactions, and the semantic contrastive loss $\mathcal{L}^S$ is formulated as:

$$\begin{aligned} \mathcal{L}^S =& \frac{1}{2}\sum_{(u,i)\in\mathbf{R}} -\log\frac{\exp(\mathbf{e}_{fu}^{G^{(k)}}\cdot\mathbf{e}_i^{G^{(k+1)}}/\tau)}{\sum_{j\in\mathcal{I}}\exp(\mathbf{e}_{fu}^{G^{(k)}}\cdot\mathbf{e}_j^{G^{(k+1)}}/\tau)} + \\ & \frac{1}{2}\sum_{(u,i)\in\mathbf{R}} -\log\frac{\exp(\mathbf{e}_i^{G^{(k)}}\cdot\mathbf{e}_{fu}^{G^{(k+1)}}/\tau)}{\sum_{v\in\mathcal{U}}\exp(\mathbf{e}_i^{G^{(k)}}\cdot\mathbf{e}_{fv}^{G^{(k+1)}}/\tau)}, \end{aligned} \tag{13}$$

Model Training: For the main recommendation task, we adopt the Bayesian Personalized Ranking (BPR) loss:

$$\mathcal{L}_{BPR} = \sum_{(u,i,j)\in O} -\log \sigma\left(\hat{y}_{ui} - \hat{y}_{uj}\right), \tag{14}$$

where $O = \{(u, i, j) \mid r_{ui} = 1, r_{uj} = 0\}$ denotes the triplets sampled, and $\hat{y}_{ui} = \mathbf{e}_u^{G\top}\mathbf{e}_i^{G}$. For each user u, i is a positive item with which u has interacted, and j is a negative item randomly sampled from the set of items that u has *not* interacted with. For each observed interaction (u, i), a non-interacted item is randomly selected as the negative sample.

To jointly optimize the recommendation and contrastive losses, we use a multi-task learning framework. The loss is:

$$\mathcal{L} = \mathcal{L}_{BPR} + \lambda_1\mathcal{L}^N + \lambda_2\mathcal{L}^S + \lambda_3\|\Theta\|_2, \tag{15}$$

where Θ is the set of learned user and item embeddings, λ_1 and λ_2 control the weights of the noise-based and semantic contrastive losses, and λ_3 is the regularization coefficient.

Theoretical Analysis

Theoretical Analysis of Semantic Path Coverage in Social Graph Generation

Let $\mathcal{G}_s = (\mathcal{U}, \mathcal{E}_s)$ be the generated social graph containing only explicit user–user edges; and let $d_{\mathcal{G}}^{(u,v)}$ denote the shortest-path length between users u and v in graph $\mathcal{G}$. We define the **Semantic Path Coverage** from k-hop in the collaborative graph to h-hop in the social graph as:

$$\eta_{k\to h} = \frac{|\{(u,v) \mid d_{\mathcal{G}_r}^{(u,v)} > k,\ d_{\mathcal{G}_s}^{(u,v)} \le h\}|}{|\{(u,v) \mid d_{\mathcal{G}_r}^{(u,v)} \le k\}|}. \tag{16}$$

The numerator counts user pairs whose distance in $\mathcal{G}_r$ is at least k but is compressed to at most h in $\mathcal{G}_s$, while the denominator enumerates all pairs within k hops in $\mathcal{G}_r$. A higher $\eta_{k\to h}$ indicates that the social graph more effectively exposes high-order relationships embedded in the collaborative graph, thereby facilitating more efficient propagation of preferences among semantically aligned users.

To model this phenomenon, we adopt the random intersection graph framework (Godehardt and Jaworski 2003), where $\mathcal{G}_r$ is a sparse bipartite graph. The probability that an edge exists between users u and v is then given by:

$$p_{\text{edge}} = 1 - \sum_{k=0}^{m} P_k \cdot \frac{\mathbb{E}_k[m - X^*]}{(m)_k}. \tag{17}$$

Here, p_{edge} is the probability of a social edge. The random variable X^* is the number of attributes a given object selects in the underlying bipartite model. The notation $\mathbb{E}[\cdot]$ is the expectation, and $(\cdot)_k$ is the k-th falling factorial, e.g., $(m)_k = m(m-1)...(m-k+1)$. For analytical tractability, we assume X^* is a discrete non-negative random variable:

$$\Pr(X^* = j) = P_j, \quad j = 0, 1, \ldots, m. \tag{18}$$

Considering the expectation term in p_{edge}, by using suitable approximations, when $m \gg j, k$, we can express it as:

$$\mathbb{E}\big[(m - X^*)_k\big] \approx \mathbb{E}\Big[(m)_k\Big(1 - \frac{kj}{m}\Big)\Big] = (m)_k \sum_{j=0}^{m} P_j\Big(1 - \frac{kj}{m}\Big). \tag{19}$$

Since $\sum_{j=0}^{m} P_j = 1$ and $\sum_{j=0}^{m} P_j \cdot j = \mathbb{E}[X^*]$, it follows:

$$\mathbb{E}\big[(m - X^*)_k\big] \approx (m)_k\Big(1 - \frac{k\,\mathbb{E}[X^*]}{m}\Big), \tag{20}$$

Based on our derivation, when $\mathcal{G}_r$ satisfies the above conditions, the probability that a social edge is formed between any two users u and v is:

$$p_{\text{edge}} \approx 1 - \sum_{k=0}^{m} P_k\Big(1 - \frac{k\,\mathbb{E}[X^*]}{m}\Big) = \frac{\mathbb{E}[X^*]^2}{m}. \tag{21}$$

Under this approximation, the social graph can be viewed as an Erdős–Rényi random graph $\mathcal{G}(n, p_{edge})$. In the sparse limit, its h-hop neighborhood size can be approximated by a Galton–Watson branching process, yielding:

$$c = (n-1)\,p_{\text{edge}}, \qquad N_l = \sum_{i=0}^{\ell} c^i, \tag{22}$$

where c is the mean degree of the social graph $\mathcal{G}_s$, N_l is the expected number of distinct users reachable from a given starting user in at most l hops. This allows a closed-form approximation for the **Semantic Path Coverage**:

$$\eta_{k\to h} = \frac{\big|\{(u,v) \mid d_{G_r}^{(u,v)} > k,\ d_{G_s}^{(u,v)} \le h\}\big|}{\big|\{(u,v) \mid d_{G_r}^{(u,v)} \le k\}\big|} \approx \frac{N_h - N_{\lfloor k/2\rfloor}}{N_{\lfloor k/2\rfloor}}. \tag{23}$$

Empirical Verification. To further demonstrate the extensive semantic coverage of the social graph intuitively, we take the dataset Flickr as an example. Although X^* is integer-valued, fitting a continuous density by maximum likelihood estimation is standard in large-sample, heavy-tailed settings and introduces negligible bias. Given Flickr's sparsity and long-tail interaction patterns, we fit the data using a log-normal distribution (Clauset, Shalizi, and Newman 2009): $X^* \sim \text{Log-normal}(\mu, \sigma^2)$, With $\mathbb{E}[X^*] = e^{\mu + \frac{1}{2}\sigma^2}$, we come that $p_{\text{edge}} \approx 0.02083$.

In typical GNNs, over-smoothing limits depth to three layers (He et al. 2020; Chen et al. 2020). For $k = h = 3$, we have $\eta_{3\to3} = c^2$, implying that the semantic interaction radius in $\mathcal{G}_s$ is c^2 times larger than in $\mathcal{G}_r$. Substituting the estimated parameters yields $\eta_{3\to3} = 1.38 \times 10^4$, indicating that **the generated social graph vastly expands the reach of semantic propagation—bridging distant but related users across orders of magnitude shorter paths**.

Theoretical Analysis of Alignment and Uniformity in Dual-View Contrastive Learning

Following advances in visual representation learning (Wang and Isola 2020), contrastive loss optimization has been shown to improve two fundamental properties of representation learning, alignment and uniformity, defined as:

$$\begin{aligned} l_{\text{align}} &\triangleq \mathbb{E}_{(x,x^+)\sim p_{\text{pos}}}\big[\|f(x) - f(x^+)\|^2\big], \\ l_{\text{uniform}} &\triangleq \log \mathbb{E}_{x,y\sim p_{\text{data}}}\big[e^{-2\|f(x)-f(y)\|^2}\big], \end{aligned} \tag{24}$$

where $p_{data}(\cdot)$ denotes for the distribution of data and $p_{pos}(\cdot,\cdot)$ stands for the distribution of positive pairs. $f(\cdot)$ indicates l_2-normalized representations, then:

$$\|f(x) - f(y)\|^2 = 2 - 2\,\text{sim}(x,y),\ \text{sim}(x,y) = f(x)^\top f(y). \tag{25}$$

By substituting Eq.25 into the original definitions, both losses can be written purely in terms of pairwise similarity:

$$l_{\text{align}} = 2 - 2\,\mathbb{E}_{(x,x^+)\sim p_{\text{pos}}}\big[\text{sim}(x,x^+)\big], \\ l_{\text{uniform}} = -4 + \log \mathbb{E}_{x,y\sim p_{\text{data}}}\big[e^{4\,\text{sim}(x,y)}\big]. \tag{26}$$

Here, l_{align} decreases linearly as the positive-pair similarity grows, directly promoting semantic closeness, while l_{uniform} imposes an exponential penalty on globally high similarity, encouraging an even spread of embeddings. Eq.26 reveals that both properties are governed by the similarity distribution across the dataset. Building on recent analyses of contrastive loss landscapes (Chuang et al. 2022), we derive the partial derivatives of the InfoNCE loss w.r.t. positive and negative similarities, thereby elucidating how our dual-view contrastive framework differentially modulates gradients to balance alignment and uniformity.

For convenience, let $\mathbf{s} = \{s^+, \{s_i^-\}_{j=1}^K\}$, where $s^+ = \text{sim}(z, z^+)$ denote the similarity between the anchor z and its positive sample z^+ and $s_j^- = \text{sim}(z, z_j^-)$ the similarity between z and the j-th negative sample z_j^-. Hence, the InfoNCE loss can be derived by the normalization term Z as:

$$\mathcal{L}_{\text{InfoNCE}}(\mathbf{s}) = \log Z - \frac{s^+}{\tau}, Z = \exp(\tfrac{s^+}{\tau}) + \sum_{j=1}^{K} \exp(\tfrac{s_j^-}{\tau}). \tag{27}$$

Taking the partial derivative of L w.r.t. the positive similarity s^+ yields:

$$\frac{\partial L}{\partial s^+} = \frac{\partial}{\partial s^+}\Big(-\frac{s^+}{\tau} + \log Z\Big) = -\frac{1}{\tau}(1 - p^+), \tag{28}$$

where $p^+ = \exp(s^+/\tau)/Z$ is the softmax probability of the positive pair. Similarly, differentiating L w.r.t. a negative similarity s_j^- gives:

$$\frac{\partial L}{\partial s_j^-} = \frac{\partial}{\partial s_j^-} \log Z = \frac{1}{\tau}\, p_j^-, \tag{29}$$

where $p_j^- = \exp(s_j^-/\tau)/Z$ is the softmax probability assigned to the j-th negative pair.

By applying Eqs.28 and 29, the gradients admit a direct interpretation in terms of their softmax weights, linking the distribution of pairwise similarities to the forces driving alignment and uniformity. As the positive similarity s^+ increases, p^+ grows, causing $-1/\tau(1-p^+)$—and thus the attractive term $\partial L/\partial s^+$—to shrink. In contrast, as p_j^- increases with s_j^-, so does repulsive push $\partial L/\partial s_j^-$.

From this perspective, *the noise-augmented view*, pairs are sampled within the same semantic category (e.g. user–user or item–item). This leads to consistently high pairwise similarity $\mathbf{s}$, with highly similar positive pairs and hard negative pairs, yielding large softmax weights p^+ and p_j^-. The resulting imbalance in gradients, namely the small magnitude of $\partial L/\partial s^+$ on positive pairs and the large magnitude of $\partial L/\partial s_j^-$ on hard negatives, applies strong repulsion to ambiguous samples, dispersing them in the embedding space and enhancing *uniformity*.

The *semantic-explored view* forms cross-layer user–item pairs, where fused user embeddings also encode social relations absent from item embeddings. This yields positive pairs with low initial similarity s^+ and easy negative pairs with low similarity s_j^-. Reversing the previous gradient pattern, the imbalance now drives optimization to pull semantically related positives closer, enhancing alignment.

In combination, the two views modulate attraction and repulsion in a complementary manner: the *noise-augmented view* enforces **global separation** to prevent collapse, while the *semantic-explored view* sharpens **local semantic cohesion**. The dual-view framework thus yields embeddings that are both evenly distributed and semantically coherent.

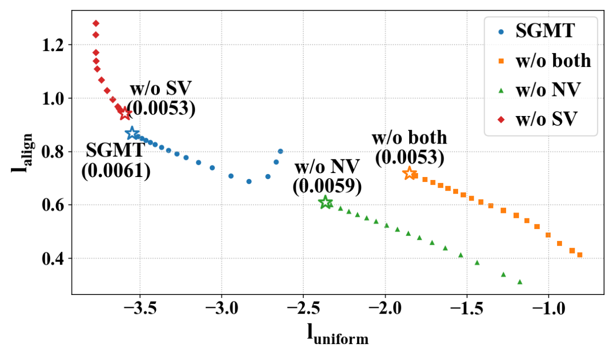

Figure 4: Training trajectories of l_{align} and l_{uniform} for different ablation settings on the Flickr dataset. Values of Hit Ratio@5 for each variant are indicated in parentheses. For both l_{align} and l_{uniform}, lower scores indicate better performance.

Empirical Verification. The training trajectories in Fig. 4 on Flickr reveal complementary roles of the two views. Without the Noise-augmented View (w/o NV), uniformity starts low and improves slowly; without the Semantic-explored View (w/o SV), uniformity is high from the outset but alignment remains weak. SGMT integrates both, with SV driving rapid alignment gains and NV maintaining uniformity for balanced optimization. Removing both views (w/o both) yields neither separation nor cohesion. In later epochs, the full model slightly trades uniformity for alignment, reflecting a balance between the two objectives.

Experiments

Datasets

Experiments are conducted on three datasets, Flickr (Cha, Mislove, and Gummadi 2009), Ciao, and Yelp (Jiang et al. 2024; Sun et al. 2025), using a 5-core setting that retains

Data	Flickr	Ciao	Yelp
# Users	5,642	5,836	4,846
# Items	21,176	10,708	5,695
# Interaction	199,500	142,805	142,950

Table 2: Statistics of datasets.

Dataset	Metric	BPR	LightGCN	SGL	SimGCL	DirectAU	CGCL	SHaRe	RGCML	SGMT	Improve
Flickr	H@5	0.0018	0.0048	0.0050	0.0056	0.0053	0.0056	0.0051	0.0055	**0.0061**	8.9%
	P@5	0.0027	0.0075	0.0077	0.0086	0.0082	0.0087	0.0077	0.0083	**0.0094**	8.0%
	R@5	0.0024	0.0057	0.0059	0.0069	0.0071	0.0069	0.0068	0.0072	**0.0084**	16.7%
	N@5	0.0033	0.0090	0.0091	0.0105	0.0103	0.0105	0.0098	0.0103	**0.0118**	12.3%
Ciao	H@5	0.0250	0.0420	0.0452	0.0458	0.0464	0.0471	0.0435	0.0465	**0.0496**	5.3%
	P@5	0.0281	0.0454	0.0489	0.0496	0.0502	0.0511	0.0471	0.0485	**0.0538**	5.2%
	R@5	0.0281	0.0452	0.0484	0.0512	0.0513	0.0526	0.0489	0.0510	**0.0555**	5.5%
	N@5	0.0358	0.0600	0.0643	0.0655	0.0673	0.0670	0.0624	0.0645	**0.0713**	5.9%
Yelp	H@5	0.0202	0.0359	0.0381	0.0399	0.0400	0.0389	0.0375	0.0387	**0.0418**	4.5%
	P@5	0.0258	0.0459	0.0487	0.0511	0.0512	0.0498	0.0479	0.0499	**0.0535**	4.4%
	R@5	0.0236	0.0426	0.0444	0.0463	0.0458	0.0470	0.0453	0.0460	**0.0505**	7.4%
	N@5	0.0315	0.0580	0.0608	0.0633	0.0640	0.0645	0.0613	0.0639	**0.0678**	5.1%

Table 3: Overall Top-5 recommendation performance comparison of different recommender models.

Dataset	Flickr				Ciao				Yelp			
Method	H@5	P@5	R@5	N@5	H@5	P@5	R@5	N@5	H@5	P@5	R@5	N@5
w/o both	0.0053	0.0082	0.0068	0.0099	0.0438	0.0475	0.0482	0.0623	0.0371	0.0475	0.0431	0.0597
w/o SV	0.0059	0.0091	0.0081	0.0114	0.0461	0.0499	0.0513	0.0667	0.0392	0.0502	0.0462	0.0623
w/o NV	0.0059	0.0091	0.0084	0.0113	0.0468	0.0506	0.0522	0.0671	0.0399	0.0511	0.0465	0.0644
SGMT	**0.0061**	**0.0094**	**0.0084**	**0.0118**	**0.0496**	**0.0538**	**0.0555**	**0.0713**	**0.0418**	**0.0535**	**0.0505**	**0.0678**

Table 4: Experimental results of ablation study.

users and items with a minimum of five interactions. Summary statistics are provided in Table 2.

Baselines and Evaluation Metrics

We benchmark SGMT against four categories of baselines: a traditional method (BPR (Rendle et al. 2009)); collaborative filtering approaches (LightGCN (He et al. 2020), DirectAU (Wang et al. 2022)); contrastive learning-based methods (SGL (Wu et al. 2021b), SimGCL (Yu et al. 2022), CGCL (He et al. 2023)); and social recommendation methods SHaRe (Jiang et al. 2024), RGCML(Xiong et al. 2025). Top-K recommendation performance is evaluated by four standard metrics: Hit Ratio (H), Precision (P), Recall (R), and NDCG (N), with K set to 5.

Performance Comparison

Table 4 presents the comparative analysis, with the following key observations: **(1) Overall superiority.** SGMT achieves the highest scores across all datasets and metrics, with relative gains over the strongest baseline ranging from 4.4% to 16.7%. **(2) Strength in extreme sparsity.** The largest improvements occur on the highly sparse Flickr dataset, with +8.9% in H@5 and +16.7% in R@5, showing the benefit of interest-aware social link generation in compressing long collaborative paths and enabling higher-order signal propagation. **(3) Robustness in moderate sparsity.** On the moderately sparse Ciao and Yelp datasets, SGMT still outperforms all baselines, achieving 5.2–5.9% and 4.4–7.4% improvements in top-5 metrics, indicating its effectiveness beyond extreme sparsity. **(4) Advantage over social recommendation** SGMT also surpasses SHaRe, highlighting that its gains stem from both leveraging social structures and refining embeddings via its dual-view CL.

Ablation Study

To evaluate the contribution of each auxiliary view in fine-tuning the main module, we perform ablation studies on all three datasets. As shown in Table 4, "w/o both", "w/o SV", and "w/o NV" indicate the removal of both auxiliary views, the *Semantic-explored View*, and the *Noise-augmented View*, respectively. The results confirm that each view offers a measurable performance gain, underscoring their complementary roles in enhancing model effectiveness.

Conclusion

In summary, we tackle two challenges in GNN-based recommendation: (1) the inability of shallow networks to capture high-order dependencies under sparse interactions, and (2) the underuse of complementary CL views for representation optimization. We propose SGMT, which constructs an interest-aware synthetic social graph, theoretically proven to compress collaborative paths and improve access to high-order signals. Additionally, we design two complementary contrastive modules, the Noise-augmented View and the Semantic-explored View, with theoretical guarantees that they enhance uniformity and alignment, respectively. Extensive experiments on three real-world datasets confirm SGMT's superiority over state-of-the-art baselines.

Acknowledgments

This work was partially supported by the National Natural Science Foundation of China (Project No. 62202122 and 62073272), the Guangdong Basic and Applied Basic Research Foundation under Grant No. 2024A1515011949, the Shenzhen Science and Technology Program under Grant No. SYSPG20241211173609009, GXWD20231130110308001, JCYJ20250604145617023, JCYJ20240813104843058, and JCYJ20240813104837050, the Shenzhen Education Science "14th Five-Year Plan" 2023 Annual Project on Artificial Intelligence Special Project under Grant No. rgzn23001, the Guangdong Province Higher Education Research and Reform Project under Grant No. YueJiaoGaoHan(2024) No.9 (1227), and the Guangdong Province General Colleges and Universities Innovation Team Project under No. 2022KCXTD038.

References

Cha, M.; Mislove, A.; and Gummadi, K. P. 2009. A measurement-driven analysis of information propagation in the flickr social network. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, 721–730. New York, NY, USA: Association for Computing Machinery. ISBN 9781605584874.

Chen, D.; Lin, Y.; Li, W.; Li, P.; Zhou, J.; and Sun, X. 2020. Measuring and relieving the over-smoothing problem for graph neural networks from the topological view. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 3438–3445.

Cheng, Z.; Zhong, T.; Zhang, K.; Walker, J.; and Zhou, F. 2022. Learning Contrastive Multi-View Graphs for Recommendation (Student Abstract). In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, 12927–12928. AAAI Press. ISBN 2374-3468.

Chuang, C.-Y.; Hjelm, R. D.; Wang, X.; Vineet, V.; Joshi, N.; Torralba, A.; Jegelka, S.; and Song, Y. 2022. Robust Contrastive Learning against Noisy Views. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 16649–16660.

Clauset, A.; Shalizi, C. R.; and Newman, M. E. J. 2009. Power-Law Distributions in Empirical Data. *SIAM Review*, 51(4): 661–703.

Fan, W.; Ma, Y.; Li, Q.; Wang, J.; Cai, G.; Tang, J.; and Yin, D. 2020. A graph neural network framework for social recommendations. *IEEE Transactions on Knowledge and Data Engineering*.

Godehardt, E.; and Jaworski, J. 2003. Two Models of Random Intersection Graphs for Classification. In Schwaiger, M.; and Opitz, O., eds., *Exploratory Data Analysis in Empirical Research*, 67–81. Berlin, Heidelberg: Springer.

Gutmann, M.; and Hyvärinen, A. 2010. Noise-contrastive estimation: A new estimation principle for unnormalized statistical models. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, 297–304. JMLR Workshop and Conference Proceedings.

Han, J.; Tang, Y.; Tao, Q.; Xia, Y.; and Zhang, L. 2024. Dual Homogeneity Hypergraph Motifs with Cross-view Contrastive Learning for Multiple Social Recommendations. *ACM Transactions on Knowledge Discovery from Data*, 18(6): 1–24.

He, W.; Sun, G.; Lu, J.; and Fang, X. S. 2023. Candidate-aware Graph Contrastive Learning for Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1670–1679.

He, X.; Deng, K.; Wang, X.; Li, Y.; Zhang, Y.; and Wang, M. 2020. 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*, 639–648.

Ji, S.; Feng, Y.; Ji, R.; Zhao, X.; Tang, W.; and Gao, Y. 2020. Dual channel hypergraph collaborative filtering. In *Proceedings of the 26th ACM SIGKDD international conference on knowledge discovery & data mining*, 2020–2029.

Jiang, W.; Gao, X.; Xu, G.; Chen, T.; and Yin, H. 2024. Challenging Low Homophily in Social Recommendation. In *Proceedings of the ACM on Web Conference 2024*, 3476–3484.

Jing, M.; Zhu, Y.; Zang, T.; and Wang, K. 2023. Contrastive self-supervised learning in recommender systems: A survey. *ACM Transactions on Information Systems*, 42(2): 1–39.

Li, J.; Shi, J.; Zhang, J.; Lu, Y.; Li, Q.; Yu, C.; and Zhang, S. 2024. Quantum nearest neighbor collaborative filtering algorithm for recommendation system. *ACM Transactions on Knowledge Discovery from Data*, 18(8): 1–28.

Ni, X.; Xiong, F.; Pan, S.; Wu, J.; Wang, L.; and Chen, H. 2023. Community preserving social recommendation with Cyclic Transfer Learning. *ACM Transactions on Information Systems*, 42(3): 1–36.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, UAI '09, 452–461. Arlington, Virginia, USA: AUAI Press. ISBN 9780974903958.

Sharma, K.; Lee, Y.-C.; Nambi, S.; Salian, A.; Shah, S.; Kim, S.-W.; and Kumar, S. 2022. A survey of graph neural networks for social recommender systems. *arXiv preprint arXiv:2212.04481*.

Sharma, K.; Lee, Y.-C.; Nambi, S.; Salian, A.; Shah, S.; Kim, S.-W.; and Kumar, S. 2024. A Survey of Graph Neural Networks for Social Recommender Systems. *ACM Computing Surveys*, 56(10): 1–34.

Singh, M. 2020. Scalability and sparsity issues in recommender datasets: a survey. *Knowledge and Information Systems*.

Sun, Y.; Sun, Z.; Du, Y.; Zhang, J.; and Ong, Y. S. 2025. Model-Agnostic Social Network Refinement with Diffusion Models for Robust Social Recommendation. In *Proceedings of the ACM on Web Conference 2025*, WWW '25, 370–378. New York, NY, USA: Association for Computing Machinery. ISBN 9798400712746.

Sun, Y.; Sun, Z.; Sha, X.; Zhang, J.; and Ong, Y. S. 2023. Disentangling Motives behind Item Consumption and Social Connection for Mutually-enhanced Joint Prediction. In *Proceedings of the 17th ACM Conference on Recommender Systems*.

Veličković, P.; Cucurull, G.; Casanova, A.; Romero, A.; Liò, P.; and Bengio, Y. 2018. Graph Attention Networks. In *International Conference on Learning Representations*.

Wang, C.; Yu, Y.; Ma, W.; Zhang, M.; Chen, C.; Liu, Y.; and Ma, S. 2022. Towards representation alignment and uniformity in collaborative filtering. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1816–1825.

Wang, T.; and Isola, P. 2020. Understanding Contrastive Representation Learning through Alignment and Uniformity on the Hypersphere. In III, H. D.; and Singh, A., eds., *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, 9929–9939. PMLR.

Wang, X.; Jin, H.; Zhang, A.; He, X.; Xu, T.; and Chua, T.-S. 2020. Disentangled graph collaborative filtering. In *Proceedings of the 43rd international ACM SIGIR conference on research and development in information retrieval*.

Wei, W.; Huang, C.; Xia, L.; Xu, Y.; Zhao, J.; and Yin, D. 2022. Contrastive meta learning with behavior multiplicity for recommendation. In *Proceedings of the fifteenth ACM international conference on web search and data mining*, 1120–1128.

Wu, B.; Zhong, L.; Yao, L.; and Ye, Y. 2022a. EAGCN: An Efficient Adaptive Graph Convolutional Network for Item Recommendation in Social Internet of Things. *IEEE Internet of Things Journal*, 9(17): 16386–16401.

Wu, J.; Wang, X.; Feng, F.; He, X.; Chen, L.; Lian, J.; and Xie, X. 2021a. Self-supervised graph learning for recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, 726–735.

Wu, J.; Wang, X.; Feng, F.; He, X.; Chen, L.; Lian, J.; and Xie, X. 2021b. Self-Supervised Graph Learning for Recommendation. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*, SIGIR '21, 726–735. New York, NY, USA: Association for Computing Machinery. ISBN 9781450380379.

Wu, L.; Li, J.; Sun, P.; Hong, R.; Ge, Y.; and Wang, M. 2020a. Diffnet++: A neural influence and interest diffusion network for social recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 34(10): 4753–4766.

Wu, L.; Yang, Y.; Zhang, K.; Hong, R.; Fu, Y.; and Wang, M. 2020b. Joint item recommendation and attribute inference: An adaptive graph convolutional network approach. In *Proceedings of the 43rd International ACM SIGIR conference on research and development in Information Retrieval*.

Wu, S.; Sun, F.; Zhang, W.; Xie, X.; and Cui, B. 2022b. Graph neural networks in recommender systems: a survey. *ACM Computing Surveys*, 55(5): 1–37.

Wu, Y.; Zhang, L.; Mo, F.; Zhu, T.; Ma, W.; and Nie, J.-Y. 2024. Unifying Graph Convolution and Contrastive Learning in Collaborative Filtering. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, KDD '24, 3425–3436. New York, NY, USA: Association for Computing Machinery. ISBN 9798400704901.

Xie, X.; Sun, F.; Liu, Z.; Wu, S.; Gao, J.; Zhang, J.; Ding, B.; and Cui, B. 2022. Contrastive learning for sequential recommendation. In *2022 IEEE 38th international conference on data engineering (ICDE)*, 1259–1273. IEEE.

Xiong, F.; Zhang, T.; Pan, S.; Luo, G.; and Wang, L. 2025. Robust Graph Based Social Recommendation Through Contrastive Multi-View Learning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(12): 12890–12898.

Yang, L.; Liu, Z.; Wang, Y.; Wang, C.; Fan, Z.; and Yu, P. S. 2022. Large-scale personalized video game recommendation via social-aware contextualized graph neural network. In *Proceedings of the ACM Web Conference 2022*.

Yu, J.; Yin, H.; Xia, X.; Chen, T.; Cui, L.; and Nguyen, Q. V. H. 2022. 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*, SIGIR '22, 1294–1303. New York, NY, USA: Association for Computing Machinery. ISBN 9781450387323.

Zang, T.; Zhu, Y.; Zhang, R.; Wang, C.; Wang, K.; and Yu, J. 2023. Contrastive Multi-View Interest Learning for Cross-Domain Sequential Recommendation. *ACM Transactions on Information Systems*, 42(3): 1–30.