

Revisiting Contrastive Learning in Collaborative Filtering via Parallel Graph Filters

Fang Kai¹, Yu Zhang², Kaibin Wang³, Lei Sang², Yiwen Zhang^{2*}

¹Faculty of Innovation Engineering, Macau University of Science and Technology, Macau, China

²School of Computer Science and Technology, Anhui University, Hefei, China

³Department of Computing Technologie, Swinburne University of Technology, Melbourne, Australia
fangk@student.must.edu.mo, zhangyu.ahu@gmail.com, kaibinwang@swin.edu.au, sanglei@ahu.edu.cn, zhangyiwen@ahu.edu.cn

Abstract

Graph Contrastive Learning (GCL) has recently emerged as a powerful paradigm for modeling user-item interactions and for learning high-quality representations in recommender systems. While existing GCL-based methods benefit from data augmentation and sampling strategies, they often overlook the inherent limitations of the contrastive objectives: 1) Stacking multiple Graph Convolutional Network layers to capture high-order information often causes the over-smoothing phenomenon, where node representations become overly similar. 2) Structurally similar negative sample pairs may exhibit high cosine similarity, causing gradient saturation during representation optimization. To address the above challenges, we revisit matrix factorization in recommendation models and uncover its implicit connection to a parallel graph filter bank. This perspective reveals how overly aggressive low-pass or high-pass filtering distorts feature distributions, contributing to gradient saturation. Building on this insight, we propose Light Cosine Similarity Collaborative Filtering (LightCSCF), a margin-constrained method that improves gradient optimization in contrastive learning by focusing on structurally hard examples, alleviating both gradient saturation and boundary over-smoothing. Extensive experiments on three real-world datasets demonstrate that LightCSCF consistently outperforms state-of-the-art baselines in recommendation accuracy and robustness to data sparsity.

Code — <https://github.com/qingzhuyanyue/LightCSCF>

Introduction

In the contemporary digital landscape, recommender systems play a vital role in addressing information overload by guiding users toward personalized content. By analyzing user preferences, recommender systems enhance user experiences across a wide range of domains, including e-commerce, media streaming, and content discovery (Gao et al. 2023; Singh et al. 2021). The primary goal of recommender systems is to accurately model the complex and high-dimensional interactions between users and items, enabling effective personalization at scale.

Traditional collaborative filtering methods (Koren 2009; Rao et al. 2015) suffer from many challenges in recom-

mender systems, including the absence of explicit collaborative signals and the cold-start problem (Lin et al. 2021; Liu et al. 2021). These challenges have motivated the exploration of graph-based models to learn representations. To better capture complex user-item interactions, Graph Convolutional Networks (GCN) have been introduced and have shown strong performance by modeling high-order connectivity patterns (He et al. 2020). However, GCN-based methods still face limitations, including data sparsity and long-tail distribution problems. To address these challenges, recent studies have turned to Contrastive Learning (CL) (Chuang et al. 2020; Khosla et al. 2020; Tian et al. 2020; Zhang and Wu 2024), a promising self-supervised learning paradigm. CL aims to construct a discriminative representation space by pulling positive pairs closer and pushing negative pairs apart. By deriving supervision signals directly from the data itself, CL enables the learning of generalizable representations and enhances recommendation performance.

Building on this paradigm, a diverse array of recommendation methods has adopted CL to improve representation quality (Ma, Lian, and Song 2025; Wang et al. 2024; Zhang et al. 2021). SGL (Wu et al. 2021) introduces a graph augmentation framework that perturbs the graph structure to generate informative self-supervised signals, while SimGCL (Yu et al. 2022), XSimGCL (Yu et al. 2023), and LightGCL (Cai et al. 2023) primarily explore augmentation by refining noise injection into the embeddings. Furthermore, methods such as CGCL (He et al. 2023), BIGCF (Zhang, Sang, and Zhang 2024) and SCCF (Wu et al. 2024) enhance the contrastive quality by further modeling user intention and multi-scale contrast. A growing line of research investigates the role of graph signal processing in CL (Liu et al. 2022a; Yu et al. 2023), constructing contrastive views that selectively emphasize low-pass or high-pass components through graph filter banks.

Despite continued advancements in CL methods, several fundamental obstacles remain hindering their representation quality (Yu et al. 2023): 1) Deepening GCN to capture high-order neighborhood relationships often results in over-smoothing, where node features become overly similar and lose their discriminative power; 2) The inclusion of hard negatives structurally similar to positives can cause gradient saturation by providing weak contrastive signals dur-

*Corresponding author

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

ing training. These limitations motivate us to reconsider the problem from a spectral perspective. Inspired by GAME (Liu et al. 2022b), which emphasizes that the discrepancy in high-frequency components between two augmented graphs should exceed that in low-frequency components, we explore new directions for contrastive learning.

To address these challenges, we first conduct a theoretical analysis to further reveal the origins of over-smoothing in GCN. Specifically, as the number of GCN layers increases, repeated propagation through the normalized adjacency matrix progressively suppresses high-frequency components in node features. This leads to overly smooth representations, making them nearly indistinguishable across nodes and weakening their discriminative power.

We then investigate the origins of gradient saturation in CL from both geometric and theoretical perspectives. The gradient updates in CL are affected by cosine similarity, which causes gradient saturation for extreme embeddings. Meanwhile, we reveal that the low-rank projection of Matrix Factorization (MF) acts as a low-pass graph filter that promotes smoothness. This insight inspires unifying embedding updates using identity, low-pass, and high-pass operators to efficiently capture both global trends and local variations.

In this paper, we propose **Light Cosine Similarity Collaborative Filtering (LightCSCF)**, a novel method designed to effectively address margin gradient saturation. Motivated by the limitations of conventional cosine similarity, we propose a margin-constrained method that not only increases the similarity of positive pairs but also suppresses negative pairs that are structurally similar. The main contributions of this work are as follows:

- We theoretically derive the over-smoothing phenomenon in GCN and incorporate a geometric perspective to investigate the gradient saturation problem in CL.
- We pinpoint the limitations of cosine similarity and propose LightCSCF, a lightweight yet effective collaborative filtering method to efficiently mitigate graph over-smoothing and alleviate gradient saturation.
- We demonstrate through comprehensive experiments on three highly sparse datasets that LightCSCF significantly outperforms existing CL-based methods, achieving enhanced recommendation accuracy and improve robustness against data sparsity.

Preliminaries

Problem Definition. The fundamental goal of a recommender system is to model user preferences based on historical interactions, in order to accurately predict future interests. We represent the user-item interaction system as a bipartite graph $\mathcal{G} = (\mathcal{U} \cup \mathcal{I}, \mathcal{E})$, where $\mathcal{U} = \{u_1, \dots, u_M\}$ and $\mathcal{I} = \{i_1, \dots, i_N\}$ denote the sets of users and items, respectively, and \mathcal{E} denotes the set of interactions between users and items. Let $\mathbf{A} \in \{0, 1\}^{M \times N}$ be the adjacency matrix. We define $\mathbf{D}_u \in \mathbb{R}^{M \times M}$ and $\mathbf{D}_i \in \mathbb{R}^{N \times N}$ as the diagonal degree matrices for users and items, respectively.

Collaborative Filtering. Matrix Factorization (MF) (Lee and Seung 1999) serves as a foundational technique in collaborative filtering, where the goal is to approximate the

user-item interaction matrix $\mathbf{R} \in \mathbb{R}^{M \times N}$ by learning low-dimensional latent embeddings. Specifically, MF defines this approximation as the factorization $\mathbf{R} \approx \mathbf{E}_u \mathbf{E}_i^T$, where $\mathbf{E}_u \in \mathbb{R}^{M \times d}$ and $\mathbf{E}_i \in \mathbb{R}^{N \times d}$ are the user and item embedding matrices, respectively, and the predicted interaction score is the inner product $\hat{r}_{ui} = \langle \mathbf{e}_u, \mathbf{e}_i \rangle$. In contrast to learning a static set of embeddings, a GCN (Kipf and Welling 2017) layer refines node representations iteratively by aggregating information from their neighbors. This process begins with an initial learnable embedding matrix $\mathbf{E}^{(0)}$ for all users and items. At each layer l , the network updates these representations to produce a new matrix $\mathbf{E}^{(l)} \in \mathbb{R}^{(M+N) \times d}$.

Contrastive Learning. Contrastive Learning (Chen et al. 2020) (CL) is a self-supervised learning paradigm that aims to learn representations by pulling positive pairs closer and pushing negative pairs apart in the embedding space. Formally, given an anchor embedding \mathbf{e}_u , a positive embedding \mathbf{e}_i , and a set of negative embeddings $\{\mathbf{e}_j\}$, the contrastive loss encourages \mathbf{e}_u to be more similar to \mathbf{e}_i than to any negative \mathbf{e}_j (Oord, Li, and Vinyals 2018):

$$\mathcal{L}_{cl} = -\log \frac{\exp(\text{sim}(\mathbf{e}_u, \mathbf{e}_i)/\tau)}{\exp(\text{sim}(\mathbf{e}_u, \mathbf{e}_i)/\tau) + \sum_j \exp(\text{sim}(\mathbf{e}_u, \mathbf{e}_j)/\tau)}, \quad (1)$$

where τ is the temperature coefficient controlling the concentration level, $\text{sim}(\mathbf{e}_u, \mathbf{e}_i)$ represents cosine similarity.

Investigation of GCN and CL

In this section, we primarily analyze the over-smoothing in GCN and the gradient saturation problem in CL.

Over-smoothing in GCN

GCN iteratively propagates information through the normalized adjacency matrix $\tilde{\mathbf{A}} = \mathbf{D}^{-\frac{1}{2}} \mathbf{A} \mathbf{D}^{-\frac{1}{2}}$. Under the simplification of removing non-linearities and setting all weight matrices to identity, this process reduces to the propagation rule $\mathbf{E}^{(K)} = \tilde{\mathbf{A}}^K \mathbf{E}^{(0)}$ (He et al. 2020), where repeated applications of $\tilde{\mathbf{A}}$ progressively smooth the node embeddings.

From a spectral perspective, the eigenvalues $\{\lambda_i\}$ of $\tilde{\mathbf{A}}$ lie within $[-1, 1]$. For a connected and non-bipartite graph, the largest eigenvalue $\lambda_1 = 1$ is simple, ensuring that $\tilde{\mathbf{A}}^K$ converges to a rank-one matrix as $K \rightarrow \infty$, with each row proportional to the principal eigenvector (Li, Han, and Wu 2018). This implies that high-frequency components associated with smaller eigenvalues $|\lambda_i| < 1$ diminish exponentially, causing node embeddings to asymptotically align along the same global direction. Consequently, the ability of the model to distinguish nodes degenerates with depth, which limits its effectiveness on tasks requiring local structural information. This trade-off between model depth and local discriminability in GCNs necessitates the design of a parallel graph filter bank to mitigate over-smoothing effects and preserve the diversity of node representations.

Gradient Saturation in CL

CL relies on gradient signals, which are derived from similarity loss functions, to effectively shape the embedding

space (Khosla et al. 2020). A commonly used choice is the cosine similarity, which facilitates the learning of expressive representations by encouraging closeness between positive pairs and enforcing separation between negative pairs, including structurally similar ones.

Definition 1 Given two embedding vectors \mathbf{z}_u and \mathbf{z}_i that are projected onto the unit n -sphere via ℓ_2 -normalization, their cosine similarity is defined as:

$$\cos(\mathbf{z}_u, \mathbf{z}_i) = \frac{\mathbf{z}_u^\top \mathbf{z}_i}{\|\mathbf{z}_u\| \cdot \|\mathbf{z}_i\|} = \mathbf{z}_u^\top \mathbf{z}_i. \quad (2)$$

Definition 2 Given two unit vectors $\mathbf{a}, \mathbf{b} \in \mathbb{R}^d$ with angle θ between them, a circular cone in \mathbb{R}^d with axis along \mathbf{a} and aperture angle θ is defined as:

$$\mathcal{K}_\theta(\mathbf{a}) = \left\{ \mathbf{x} \in \mathbb{R}^d \setminus \{\mathbf{0}\} \mid \frac{\mathbf{a}^\top \mathbf{x}}{\|\mathbf{x}\|} = \cos(\mathbf{a}, \mathbf{b}) \right\}. \quad (3)$$

Intersecting this cone with the unit $(d-1)$ -sphere $\mathbb{S}^{d-1} = \{\mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\| = 1\}$ yields:

$$\begin{aligned} \mathcal{L}_\theta(\mathbf{a}) &= \mathcal{K}_\theta(\mathbf{a}) \cap \mathbb{S}^{d-1} \\ &= \{ \mathbf{x} \in \mathbb{R}^d \mid \|\mathbf{x}\| = 1, \mathbf{a}^\top \mathbf{x} = \cos(\mathbf{a}, \mathbf{b}) \}. \end{aligned} \quad (4)$$

This set $\mathcal{L}_\theta(\mathbf{a})$ consists of all unit vectors that share the same cosine similarity with \mathbf{a} as \mathbf{b} does. Geometrically, it forms the intersection between a cone centered at \mathbf{a} and the unit n -sphere. Specifically, treating \mathbf{a} as the anchor and \mathbf{b} as the positive sample, any vector $\mathbf{c} \in \mathcal{L}_\theta(\mathbf{a})$ lies on the intersection of the unit n -sphere and the cone centered at \mathbf{a} with aperture angle $\theta = \arccos(\mathbf{a}^\top \mathbf{b})$. By construction, \mathbf{c} satisfies $\mathbf{a}^\top \mathbf{c} = \mathbf{a}^\top \mathbf{b}$ and thus shares the same cosine similarity with \mathbf{a} as the positive sample \mathbf{b} . This makes \mathbf{c} indistinguishable from \mathbf{b} despite representing a structurally negative sample. As the angle θ decreases, the cone narrows and $\mathcal{L}_\theta(\mathbf{a})$ collapses toward the direction of \mathbf{a} , causing the cosine similarity $\mathbf{a}^\top \mathbf{b}$ to approach 1 and placing the anchor-positive pair in the saturation region of the loss function. This saturation effect weakens the optimization signal, making the model insensitive to structurally negatives residing in $\mathcal{L}_\theta(\mathbf{a})$, limiting the discriminative ability of learned representations.

From theoretical analysis, the optimization signal in CL can be defined as $\text{sim}(\mathbf{e}_u, \mathbf{e}_i) = e^{\cos(\mathbf{e}_u, \mathbf{e}_i)}$, which is influenced by the derivative of the cosine similarity between the embedding vectors \mathbf{e}_u and \mathbf{e}_i . When the angle $\theta_{ui} \in [0, \pi]$ between the two embeddings approaches the endpoints, the derivative vanishes and the gradient becomes negligible:

$$\text{At } \theta_{ui} = 0 \text{ or } \pi, \frac{\partial \text{sim}(\mathbf{e}_u, \mathbf{e}_i)}{\partial \theta_{ui}} = \frac{\partial e^{\cos(\theta_{ui})}}{\partial \theta_{ui}} = 0. \quad (5)$$

As a result, when two embeddings are either nearly identical or diametrically opposed, the learning signal degenerates. This saturation effect restricts the model's ability to make meaningful adjustments, limiting the refinement of representation learning in later stages of training.

Methodology

In this section, we analyze the gradient update process in CL through the lens of MF, demonstrating that this process can be interpreted as a parallel graph filter bank. Based on this insight, we further propose a margin-constrained method.

CL with MF as Parallel Graph Filter

The update process in MF implicitly acts as a graph convolution by smoothing node embeddings over the interaction graph. To clarify this connection, we analyze the mathematical structure of the low-rank projection operator and its effect on embedding smoothness.

Definition 3 Given an embedding matrix $\mathbf{P} \in \mathbb{R}^{n \times d}$, the matrix $\mathbf{M} = \mathbf{P}\mathbf{P}^\top \in \mathbb{R}^{n \times n}$ is called a low-rank projection operator, which maps any vector $x \in \mathbb{R}^{n \times d}$ to its projection $x' = \mathbf{M}x = \mathbf{P}(\mathbf{P}^\top x)$ and lies in the column space of \mathbf{P} .

Definition 4 The spectral norm $\|\cdot\|_2$ satisfies the sub-multiplicative property: for any compatible matrices \mathbf{A} and \mathbf{B} , $\|\mathbf{A}\mathbf{B}\|_2 \leq \|\mathbf{A}\|_2 \cdot \|\mathbf{B}\|_2$.

The projection matrix $\mathbf{M} = \mathbf{P}\mathbf{P}^\top$ maps input signals onto a low-dimensional subspace spanned by the learned embeddings. This operation suppresses high-frequency components in the original signal, acting as a low-pass filter. In the context of graph signal processing, such filtering promotes smoothness by reducing variations across connected nodes. To quantify this effect, we examine how the smoothness of the transformed signal $\mathbf{M}x$ relates to the original signal.

Theorem 1 Given a graph Laplacian matrix \mathbf{L} , the smoothness of a signal \mathbf{x} is measured by $S(\mathbf{x}) = \mathbf{x}^\top \mathbf{L} \mathbf{x}$.

proof 1 Let $S(\mathbf{x}) = \mathbf{x}^\top \mathbf{L} \mathbf{x} = \mathbf{x}^\top (\mathbf{D} - \mathbf{A}) \mathbf{x} = \mathbf{x}^\top \mathbf{D} \mathbf{x} - \mathbf{x}^\top \mathbf{A} \mathbf{x}$, where $\mathbf{D}_{ii} = \sum_{j=1}^N \mathbf{A}_{ij}$. This expands as:

$$\begin{aligned} S(\mathbf{x}) &= \sum_{i=1}^N \left(\sum_{j=1}^N \mathbf{A}_{ij} \right) x_i^2 - \sum_{i=1}^N \sum_{j=1}^N \mathbf{A}_{ij} x_i x_j \\ &= \sum_{i=1}^N \sum_{j=1}^N \mathbf{A}_{ij} (x_i^2 - x_i x_j). \end{aligned} \quad (6)$$

By symmetry, the quadratic form can be rewritten as:

$$S(\mathbf{x}) = \frac{1}{2} \sum_{i=1}^N \sum_{j=1}^N \mathbf{A}_{ij} (x_i - x_j)^2. \quad (7)$$

Here, $(x_i - x_j)^2$ represents the squared difference between signal values and \mathbf{A}_{ij} is the corresponding edge weight.

Theorem 2 Let \mathbf{M} be a linear projection operator and \mathbf{L} be the graph Laplacian. For any signal \mathbf{x} , the smoothness of the transformed signal $\mathbf{M}\mathbf{x}$ is upper-bounded by: $S(\mathbf{M}\mathbf{x}) \leq \|\mathbf{M}\|_2^2 \cdot \|\mathbf{L}\|_2 \cdot \|\mathbf{x}\|_2^2$.

proof 2 The smoothness of the signal $\mathbf{M}\mathbf{x}$ is defined as:

$$S(\mathbf{M}\mathbf{x}) = (\mathbf{M}\mathbf{x})^\top \mathbf{L} (\mathbf{M}\mathbf{x}) = \mathbf{x}^\top \mathbf{M}^\top \mathbf{L} \mathbf{M} \mathbf{x}. \quad (8)$$

By applying the sub-multiplicative property of the spectral norm, we have:

$$\mathbf{x}^\top \mathbf{M}^\top \mathbf{L} \mathbf{M} \mathbf{x} \leq \|\mathbf{M}^\top \mathbf{L} \mathbf{M}\|_2 \cdot \|\mathbf{x}\|_2^2. \quad (9)$$

Since $\|\mathbf{M}^\top \mathbf{L} \mathbf{M}\|_2 \leq \|\mathbf{M}\|_2^2 \cdot \|\mathbf{L}\|_2$, we obtain the result:

$$S(\mathbf{M}\mathbf{x}) \leq \|\mathbf{M}\|_2^2 \cdot \|\mathbf{L}\|_2 \cdot \|\mathbf{x}\|_2^2. \quad (10)$$

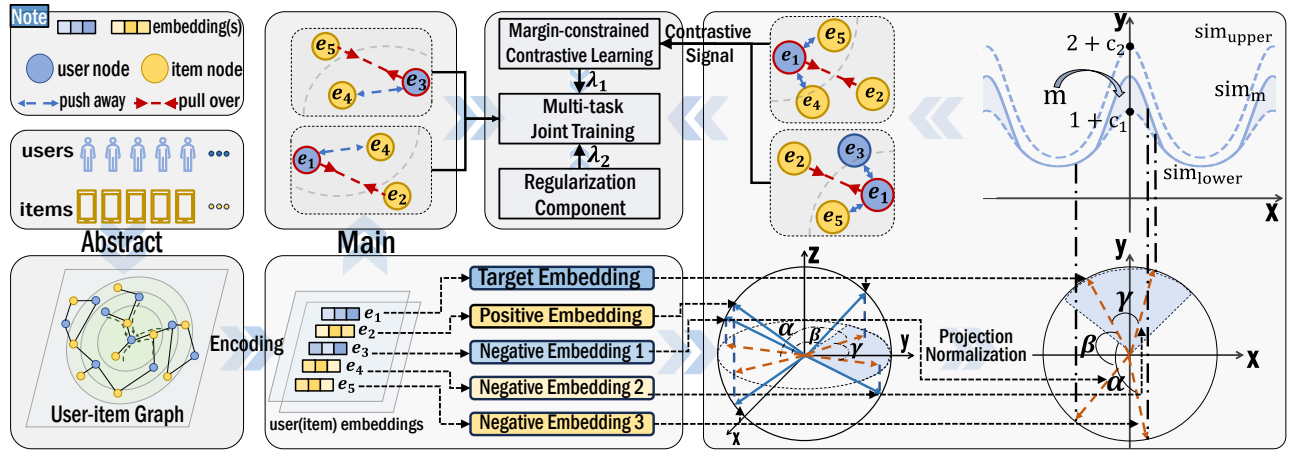


Figure 1: The LightCSCF method encodes user-item interactions via a bipartite graph, projects and normalizes the embeddings, then applies margin-constrained CL with multi-negative sampling for joint optimization of recommendation and representation.

Theorem 2 provides an upper bound on the smoothness of the transformed signal, showing that the projection operator \mathbf{M} acts as a low-pass filter. Leveraging this property, we aim to guide the embedding \mathbf{E} toward representations with desirable frequency characteristics. A natural method (Lee and Seung 2000) is to update embeddings by reinforcing low-frequency components via \mathbf{M} and suppressing high-frequency components via the complementary high-pass operator $\mathbf{N} = \mathbf{I} - \mathbf{M}$. This leads to the embedding update rule:

$$\mathbf{E}^{(t+1)} = \mathbf{E}^{(t)} + \mu\mathbf{M}\mathbf{E}^{(t)} - \nu\mathbf{N}\mathbf{E}^{(t)}, \quad (11)$$

where μ and ν are coefficients. Unlike sequential message passing, this formulation applies the three filters in parallel, resembling a residual filtering mechanism that simultaneously captures smooth trends and discriminative variations.

We formalize the CL process through a parallel graph filter bank composed of three components: the identity operator \mathbf{I} , a low-pass filter \mathbf{M} that promotes smoothness, and a high-pass filter \mathbf{N} that captures high-frequency residuals. This unified design allows the contrastive gradient to be decomposed into interpretable signals, where the attractive gradient satisfies $-\eta\nabla_{\mathbf{E}}\mathcal{L}_{\text{pull}} = \mu\mathbf{M}\mathbf{E}^{(t)}$ and the repulsive gradient satisfies $-\eta\nabla_{\mathbf{E}}\mathcal{L}_{\text{push}} = -\nu\mathbf{N}\mathbf{E}^{(t)}$. Consequently, the embedding update follows $\mathbf{E}^{(t+1)} = (\mathbf{I} + \mu\mathbf{M} - \nu\mathbf{N})\mathbf{E}^{(t)}$, blending identity, low-pass, and high-pass operations within a unified filter bank framework. Furthermore, the repulsive signal from negative samples implicitly enforces low-rank constraints on the similarity matrix, acting mainly in the spectral domain. This connection underscores the importance of selecting structurally meaningful negative samples: while low-pass filters promote semantic smoothness, high-pass filters accentuate structural differences, thus enabling better discrimination between topologically similar yet semantically distinct nodes.

LightCSCF

To investigate how cosine similarity influences the CL loss in recommender systems, we analyze the loss function

adopted by CL-based methods. Despite the growing interest in contrastive loss functions, we observe that few studies explicitly address the inherent limitations of the standard cosine similarity (Wu et al. 2021; Lin et al. 2022; Yu et al. 2022; Cai et al. 2023; He et al. 2023; Yang et al. 2023; Zhang, Sang, and Zhang 2024; Zhang et al. 2024b, 2025; Zhang and Zhang 2025). Among the existing efforts, VGCL (Yang et al. 2023) introduces a probabilistic framework which leverages the probability of two users (or items) being assigned to the same prototype to assist the cosine similarity in making similarity judgments, without altering the similarity function directly. Additionally, SCCF (Wu et al. 2024) enhances CL by modifying the similarity function $\text{sim}(\cdot)$. However, our analysis reveals a drawback of SCCF: the squared term becomes positive when cosine similarity becomes -1 , which signals a false similarity.

Recognizing that hard negatives can generate misleading gradient signals and that similarity values in CL often saturate near their extremes, we reformulate the similarity function from a spectral perspective. Our similarity function consists of two exponential terms: the first corresponds to the conventional low-pass filtered cosine similarity, while the second introduces a margin-constrained rectification that activates when the similarity exceeds a threshold m , emphasizing high-frequency differences, thus aligning with the structure of the update operator $\mathbf{H} = \mathbf{I} + \mu\mathbf{M} - \nu\mathbf{N}$.

To better understand and control the gradient behavior introduced by the margin-constrained high-pass term, we analyze the theoretical bounds of our similarity formulation. First, the high-pass filter is designed to be sensitive to high-frequency signals and to amplify alignment signals between positive pairs. We therefore estimate the lower bound:

$$\text{sim}_{\text{lower}} = \exp\left(\frac{\cos(e_u, e_i)}{\tau}\right) + c_1, \quad (12)$$

where c_1 is a positive constant that ensures a non-vanishing contribution from the low-pass term while providing a consistent gain across positive samples. Empirically, we choose $c = 1$, consistent with the design of the similarity function

used in SCCF. Next, we consider an upper bound beyond which the high-pass filter may dominate overall similarity, contradicting our design goal of balanced contributions from both frequency bands. We define the upper bound as

$$\text{sim}_{\text{upper}} = \gamma \cdot \exp\left(\frac{\cos(e_u, e_i)}{\tau}\right) + c_2, \quad (13)$$

where c_2 is a positive constant and $\gamma \leq 2$ ensures that the low-frequency structure remains the primary contributor to smoothness, while the high-frequency component acts as a targeted rectifier. Then we define the final margin-constrained similarity as

$$\text{sim}_m = \exp\left(\frac{\cos(e_u, e_i)}{\tau}\right) + G(e_u, e_i; \tau), \quad (14)$$

where $G(e_u, e_i; \tau)$ is the high-pass filter that satisfies

$$\text{sim}_{\text{lower}} \leq \text{sim}_m \leq \text{sim}_{\text{upper}}. \quad (15)$$

We instantiate G using the constraint function with a margin parameter m as $\exp(\max(\cos(e_u^\top e_i) - m, 0)/\tau)$ that introduces piecewise non-linearity while preserving gradient flow. The operational process of our method is shown in Fig. 1. Based on the bounded formulation, we design a loss function that optimizes the margin-constrained similarity:

$$\mathcal{L}_{\text{LightCSCF}} = -\log \text{sim}_m(e_u, e_i) + \log\left(\sum_{j \in \mathcal{B}_u} \text{sim}_m(e_u, e_j)\right). \quad (16)$$

Here, e_u and e_i denote the embedding vectors for user u and positive item i , respectively. In the loss formulation, \mathcal{B}_u denotes the set of other users and items in the same training batch, which serve as implicit negative samples. When the cosine similarity between the anchor and a positive sample exceeds the margin m , sim_m increases the contribution of this pair through the second term, effectively pulling the positive sample closer. On the other hand, if the similarity between the anchor and a sample from \mathcal{B}_u does not exceed the margin, the second term degenerates into a constant, reducing the overall similarity and pushing apart structurally similar but undesired negative samples.

Our proposed margin-constrained method offers several key advantages: (1) the margin constraint selectively emphasizes highly similar pairs, enhancing the model’s ability to separate positives from negatives in a margin-constrained metric learning manner; (2) to address gradient saturation, an additional term provides refinement signals when similarities exceed the margin, alleviating vanishing gradients for aligned pairs; (3) the piecewise-smooth activation function helps mitigate over-smoothing by preserving informative gradients and introducing non-linearity at zero.

Multi-task Joint Training

Although $\mathcal{L}_{\text{LightCSCF}}$ alleviates the problems of gradient saturation and over-smoothing, its implicit supervision may be insufficient for optimizing on the core CF task. To address this, we adopt a multi-task learning framework that jointly optimizes the model using both the BPR loss (Rendle et al. 2009) and the LightCSCF loss. The overall training objective is defined as:

$$\mathcal{L} = \mathcal{L}_{\text{main}} + \lambda_1 \mathcal{L}_{\text{LightCSCF}} + \lambda_2 \mathcal{L}_{\text{reg}}, \quad (17)$$

Dataset	#Users	#Items	#Interactions	Density
Amazon-book	52643	91599	2984108	0.06%
Douban-book	13024	22347	792062	0.27%
Tmall	47939	41390	2619389	0.13%

Table 1: Statistics of the datasets.

where λ_1 balances the impact of the LightCSCF loss, λ_2 denotes the strength of the regularization loss \mathcal{L}_{reg} .

Theoretical Analysis

Let $s = \cos(e_u, e_i)$ denote the cosine similarity between node embeddings filtered by the parallel graph filter bank $\mathbf{H} = \mathbf{I} + \mu\mathbf{M} - \nu\mathbf{N}$, which serves as the basis for defining our similarity function. The similarity function is defined as

$$\text{sim}_m(s) = \exp\left(\frac{s}{\tau}\right) + \exp\left(\frac{\max(s - m, 0)}{\tau}\right). \quad (18)$$

The first and second derivatives with respect to s are

$$\text{sim}_m^{(1,2)}(s) = \begin{cases} \frac{1}{\tau} \exp\left(\frac{s}{\tau}\right), & \text{if } s \leq m, \\ \frac{1}{\tau} \left[\exp\left(\frac{s}{\tau}\right) + \exp\left(\frac{s-m}{\tau}\right) \right], & \text{if } s > m, \end{cases} \quad (19)$$

This piecewise smooth function causes the gradient magnitude $\text{sim}'_m(s)$ to increase more rapidly when $s > m$, imposing stronger penalties on hard negative pairs with high similarity. Moreover, by amplifying the gradients for high-similarity negative pairs, the margin-based similarity function implicitly preserves discriminative features across similar nodes, mitigating the over-smoothing effect that often occurs in deep GCN-based CL. Simultaneously, the increased curvature $\text{sim}''_m(s)$ in this regime may improve optimization stability and convergence, while also enhancing sensitivity to subtle differences among highly similar embeddings by providing more informative gradient signals.

Experiments

In this section, we conduct extensive experiments on three real-world datasets and compare LightCSCF with several state-of-the-art methods to evaluate its effectiveness. The statistics of the datasets are summarized in Table 1.

Experimental Setting

Datasets and Evaluation Metrics. We utilize three widely used public benchmark datasets in our experiments: *Amazon-book* (Sang et al. 2025), *Douban-book* (Yu et al. 2022), and *Tmall* (Ren et al. 2023). And we randomly split each dataset into training and test sets with a ratio of 80% to 20%, respectively, and adopt two standard metrics: Recall@ K and NDCG@ K , where $K \in \{10, 20\}$.

Baselines. The baseline models used for comparison can be categorized into four groups: (1) **BPR-, AU- and GCN-based methods:** MF-BPR (Koren, Bell, and Volinsky 2009), LightGCN (He et al. 2020) and DirectAU (Wang et al. 2022). (2) **Generative-based methods:** Mult-VAE (Liang et al. 2018), CVGA (Zhang et al. 2023), DiffRec (Wang

Dataset	Amazon-book				Tmall				Douban-book			
Method	R@10	N@10	R@20	N@20	R@10	N@10	R@20	N@20	R@10	N@10	R@20	N@20
BPR-MF (UAI'09)	0.0170	0.0182	0.0308	0.0239	0.0312	0.0287	0.0547	0.0400	0.0849	0.1079	0.1292	0.1147
DirectAU (KDD'22)	0.0296	0.0297	0.0506	0.0401	0.0475	0.0443	0.0752	0.0576	0.1153	0.1527	0.1660	0.1568
LightGCN (SIGIR'20)	0.0232	0.0245	0.0402	0.0312	0.0435	0.0406	0.0711	0.0530	0.1035	0.1369	0.1541	0.1422
MultVAE (WWW'18)	0.0224	0.0239	0.0407	0.0315	0.0467	0.0423	0.0740	0.0552	0.1156	0.1532	0.1670	0.1604
CVGA (TOIS'23)	0.0290	0.0302	0.0492	0.0379	0.0540	0.0517	0.0854	0.0648	0.1229	0.1670	0.1763	0.1699
DiffRec (SIGIR'23)	0.0310	0.0333	0.0514	0.0418	0.0485	0.0473	0.0792	0.0612	0.1102	0.1545	0.1619	0.1661
SGL (SIGIR'21)	0.0263	0.0281	0.0478	0.0379	0.0457	0.0434	0.0738	0.0556	0.1153	0.1558	0.1633	0.1585
SimGCL (SIGIR'22)	0.0313	0.0334	0.0515	0.0414	0.0559	0.0536	0.0884	0.0674	0.1230	0.1637	0.1772	0.1583
LightGCL (ICLR'23)	0.0303	0.0318	0.0506	0.0397	0.0531	0.0508	0.0833	0.0637	0.1069	0.1393	0.1570	0.1455
BIGCF (SIGIR'24)	0.0294	0.0320	0.0500	0.0398	0.0547	0.0524	0.0876	0.0664	0.1199	0.1642	0.1741	0.1682
MixRec (WWW'25)	0.0325	0.0351	0.0534	0.0427	0.0481	0.0458	0.0757	0.0577	0.1149	0.1531	0.1668	0.1578
LightCCF (SIGIR'25)	0.0340	0.0367	0.0569	0.0455	0.0592	0.0564	0.0930	0.0706	0.1321	0.1851	0.1863	0.1854
Improv.1%	5.00%	4.09%	4.39%	3.74%	4.39%	4.96%	4.52%	4.96%	2.04%	1.67%	1.93%	1.46%
SCCF (KDD'24)	0.0287	0.0294	0.0491	0.0399	0.0478	0.0453	0.0772	0.058	0.1252	0.1611	0.1711	0.1639
RecDCL (WWW'24)	0.0311	0.0318	0.0525	0.0407	0.0527	0.0492	0.0853	0.0632	0.1151	0.1452	0.1664	0.1526
LightCSCF (Ours)	0.0357*	0.0382*	0.0594*	0.0472*	0.0618*	0.0592*	0.0972*	0.0741*	0.1348*	0.1882*	0.1899*	0.1881*
Improv.2%	14.79%	20.13%	13.14%	15.97%	17.27%	20.33%	13.95%	17.25%	7.67%	16.82%	10.99%	14.77%

Table 2: A comparison between LightCSCF and state-of-the-art baselines is shown, with best results in **bold** and second-best underlined. 'Improv.1%' and 'Improv.2%' indicate relative improvements over CL-based baseline and the RCL-based baseline. An asterisk (*) marks statistically significant improvements (t-test $p < 0.05$) over the second-best.

et al. 2023). (4) **CL-based methods**: SGL (Wu et al. 2021), SimGCL (Yu et al. 2022), LightGCL (Cai et al. 2023), BIGCF (Zhang, Sang, and Zhang 2024), LightCCF (Zhang et al. 2025), MixRec (Zhang and Zhang 2025). (5) **RCL-based methods**: SCCF (Wu et al. 2024), RecDCL (Zhang et al. 2024a).

Implementation Details. To ensure a fair comparison, all experiments were conducted using a unified framework. Xavier initialization was applied to all models, and the embedding size was fixed at 64. An exception is made for RecDCL, which requires the embedding size to be set to 2048 due to its architectural design. For GCN-based methods, the number of layers is set to 3. All models are trained for up to 1000 epochs with an early stopping patience of 20. The temperature coefficient τ is set within $\{0.1, \dots, 1.0\}$ and the hyperparameter m is set within $\{0.1, \dots, 1\}$. The loss weight λ_1 in LightCSCF is set within $\{0.1, \dots, 10\}$. The regularization weight λ_2 and the learning rate are fixed at 0.0001 and 0.001, respectively.

Performance Compare with Baselines

As shown in Table 2, LightCSCF consistently outperforms all baselines across three datasets. Specifically, it improves Recall@20 over the second-best model by 4.39%, 4.52%, and 1.93% on Amazon-Book, Tmall, and Douban-book, respectively. Compared with RCL-based methods, LightCSCF achieves even greater gains of 13.14%, 13.95%, and 10.99%. These results highlight the ability of LightCSCF to identify structurally similar negative samples.

LightCSCF also shows significant advantages over various CL-based models such as SimGCL, MixRec, and LightCCF, which depend on complex augmentations or neighborhood aggregation. These methods are still constrained by standard cosine similarity. In contrast, LightCSCF surpasses them without extra augmentation, demon-

strating the effectiveness of its boundary-based cosine similarity method. Against advanced RCL models like SCCF and RecDCL, LightCSCF achieves substantial gains, including a 17.25% improvement in NDCG@20 on Tmall. Overall, LightCSCF adopts a margin-constrained method that improves the discrimination of structurally similar negatives while reducing reliance on graph convolutions.

In-depth Studies of LightCSCF

Embedding Size Analysis. As shown in Fig. 2, model performance improves as the embedding dimension increases, but tends to saturate after a certain point. This trend reflects the trade-off between representational capacity and efficiency. Beyond computational considerations, the effectiveness of lower-dimensional embeddings can be explained by classical MF methods, which are based on the assumption that a small number of leading singular components can sufficiently approximate the interaction matrix (Wu et al. 2024).




Figure 2: Performance of different embedding sizes.

Robustness Analysis. As illustrated in Fig. 3, the bar chart shows the changes of two models under different noise levels, while the line chart depicts the performance drop ratio relative to the performance without noise injection. We observe that NDCG@20 increases as more noise is introduced.

This counterintuitive trend can be attributed to the mitigation of overfitting, which in turn enhances the generalization ability of the model. Furthermore, we find that LightCSCF exhibits minimal performance degradation across Douban-book datasets, highlighting the strong discriminative capacity of its loss function $\mathcal{L}_{\text{LightCSCF}}$, which effectively alleviates gradient saturation and over-smoothing.




Figure 3: Robustness comparison of LightCSCF.

t-SNE Visualization Analysis. To evaluate the representation of LightCSCF, we randomly sampled 5,000 user and item embeddings and projected them onto a 2D surface using t-SNE, as shown in Fig. 4. The resulting visualization reveals a reduced clustering effect among users and items, with user embeddings exhibiting a more uniform distribution. This suggests that the margin-constrained method effectively identifies hard negative samples with structural similarity, mitigating the formation of clusters.




Figure 4: t-SNE visualization of LightCSCF (Tmall).

Sparsity Analysis. To further investigate the data sparsity issue, we conducted sparsity-based experiments on LightCSCF and competitive baselines by dividing the dataset into four user interaction frequency categories: frequent, normal, occasional, and sparse, then measuring performance via NDCG@20. As shown in Fig. 5, LightCSCF consistently outperforms baselines across all sparsity levels thanks to its margin-constrained cosine similarity method. Notably, in the Amazon-Book dataset, the sparse subset achieves higher scores than the other categories, which can be explained by the prevalence of false-positive samples weakening the effect of \mathcal{L}_{BPR} . Removing \mathcal{L}_{BPR} allows the influence of $\mathcal{L}_{\text{LightCSCF}}$ to stand out more clearly.

Hyperparameter Analysis. LightCSCF introduces only two hyperparameters, m and λ_1 , besides τ . To better un-

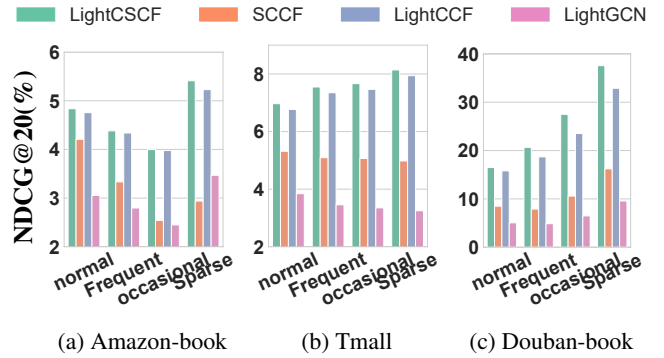


Figure 5: Performance comparison of LightCSCF with other models under different sparsity levels on three datasets.

derstand the influence of m and τ , we conducted a grid search over selected values to identify the optimal setting for three benchmark datasets, which are presented in Fig. 6. The best Recall@20 scores achieved were 0.0594, 0.0972, and 0.1898, respectively. As predicted by our theoretical analysis, the proposed margin-constrained cosine similarity method $\text{sim}_{\text{LightCSCF}}$ modifies the standard cosine similarity to some extent. While τ remains a key factor in CL loss, significantly impacting the uniformity of the learned embedding distribution (Zhang et al. 2024b).




Figure 6: Grid searches of impact on m and τ of LightCSCF on three datasets. Star represents optimal performance.

Conclusion

In this paper, we theoretically show that deep graph convolution intensifies the over-smoothing problem. We analyze gradient saturation in CL from both geometric and theoretical perspectives, revealing its connection to MF. Based on these insights, we propose LightCSCF, a novel recommendation method that alleviates both over-smoothing and gradient saturation via a margin-constrained method. This approach suppresses hard negatives that are structurally similar to the anchor, encouraging clearer separation between positive and negative pairs in the embedding space. Extensive experiments on three real-world datasets confirm the performance and robustness of LightCSCF.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No.62272001) and the Science and Technology Development Fund of Macau Project 0096/2023/RIA2.

References

- Cai, X.; Huang, C.; Xia, L.; and Ren, X. 2023. LightGCL: Simple Yet Effective Graph Contrastive Learning for Recommendation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*. OpenReview.net.
- Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, 1597–1607. PmLR.
- Chuang, C.-Y.; Robinson, J.; Lin, Y.-C.; Torralba, A.; and Jegelka, S. 2020. Debaised contrastive learning. In *Advances in neural information processing systems*, volume 33, 8765–8775.
- Gao, C.; Zheng, Y.; Li, N.; Li, Y.; Qin, Y.; Piao, J.; Quan, Y.; Chang, J.; Jin, D.; He, X.; et al. 2023. A survey of graph neural networks for recommender systems: Challenges, methods, and directions. *ACM Transactions on Recommender Systems*, 1(1): 1–51.
- 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.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised contrastive learning. In *Advances in neural information processing systems*, volume 33, 18661–18673.
- Kipf, T. N.; and Welling, M. 2017. Semi-Supervised Classification with Graph Convolutional Networks. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Koren, Y. 2009. Collaborative filtering with temporal dynamics. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, 447–456.
- Koren, Y.; Bell, R.; and Volinsky, C. 2009. Matrix factorization techniques for recommender systems. *Computer*, 42(8): 30–37.
- Lee, D.; and Seung, H. S. 2000. Algorithms for non-negative matrix factorization. In *Advances in neural information processing systems*, volume 13.
- Lee, D. D.; and Seung, H. S. 1999. Learning the parts of objects by non-negative matrix factorization. *nature*, 401(6755): 788–791.
- Li, Q.; Han, Z.; and Wu, X.-M. 2018. Deeper insights into graph convolutional networks for semi-supervised learning. In *Proceedings of the AAAI conference on artificial intelligence*, volume 32, 3538–3545.
- Liang, D.; Krishnan, R. G.; Hoffman, M. D.; and Jebara, T. 2018. Variational autoencoders for collaborative filtering. In *Proceedings of the 2018 world wide web conference*, 689–698.
- Lin, X.; Wu, J.; Zhou, C.; Pan, S.; Cao, Y.; and Wang, B. 2021. Task-adaptive neural process for user cold-start recommendation. In *Proceedings of the web conference 2021*, 1306–1316.
- Lin, Z.; Tian, C.; Hou, Y.; and Zhao, W. X. 2022. Improving graph collaborative filtering with neighborhood-enriched contrastive learning. In *Proceedings of the ACM web conference 2022*, 2320–2329.
- Liu, J.; Shi, C.; Yang, C.; Lu, Z.; and Yu, P. S. 2022a. A survey on heterogeneous information network based recommender systems: Concepts, methods, applications and resources. *AI Open*, 3: 40–57.
- Liu, N.; Wang, X.; Bo, D.; Shi, C.; and Pei, J. 2022b. Revisiting graph contrastive learning from the perspective of graph spectrum. In *Advances in Neural Information Processing Systems*, volume 35, 2972–2983.
- Liu, W.; Su, J.; Chen, C.; and Zheng, X. 2021. Leveraging distribution alignment via stein path for cross-domain cold-start recommendation. In *Advances in Neural Information Processing Systems*, volume 34, 19223–19234.
- Ma, R.; Lian, Y.; and Song, C. 2025. Sub-Interest-Aware Representation Uniformity for Recommender System. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 12346–12354.
- Oord, A. v. d.; Li, Y.; and Vinyals, O. 2018. Representation learning with contrastive predictive coding. arXiv:1807.03748.
- Rao, N.; Yu, H.-F.; Ravikumar, P. K.; and Dhillon, I. S. 2015. Collaborative Filtering with Graph Information: Consistency and Scalable Methods. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Ren, X.; Xia, L.; Zhao, J.; Yin, D.; and Huang, C. 2023. Disentangled contrastive collaborative filtering. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*, 1137–1146.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian Personalized Ranking from Implicit Feedback. In Bilmes, J. A.; and Ng, A. Y., eds., *UAI 2009, Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, Montreal, QC, Canada, June 18-21, 2009*, 452–461. AUAI Press.
- Sang, L.; Zhang, Y.; Zhang, Y.; Li, H.; and Zhang, Y. 2025. Towards similar alignment and unique uniformity in collaborative filtering. *Expert Systems with Applications*, 259: 125346.

- Singh, P. K.; Pramanik, P. K. D.; Dey, A. K.; and Choudhury, P. 2021. Recommender systems: an overview, research trends, and future directions. *International Journal of Business and Systems Research*, 15(1): 14–52.
- Tian, Y.; Sun, C.; Poole, B.; Krishnan, D.; Schmid, C.; and Isola, P. 2020. What makes for good views for contrastive learning? In *Advances in neural information processing systems*, volume 33, 6827–6839.
- 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, K.; Zhu, Y.; Zang, T.; Wang, C.; and Jing, M. 2024. Enhanced hierarchical contrastive learning for recommendation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 9107–9115.
- Wang, W.; Xu, Y.; Feng, F.; Lin, X.; He, X.; and Chua, T.-S. 2023. Diffusion recommender model. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*, 832–841.
- Wu, J.; Wang, X.; Feng, F.; He, X.; Chen, L.; Lian, J.; and Xie, X. 2021. 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, 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*, 3425–3436.
- Yang, Y.; Wu, Z.; Wu, L.; Zhang, K.; Hong, R.; Zhang, Z.; Zhou, J.; and Wang, M. 2023. Generative-contrastive graph learning for recommendation. In *Proceedings of the 46th international ACM SIGIR Conference on Research and Development in Information Retrieval*, 1117–1126.
- Yu, J.; Xia, X.; Chen, T.; Cui, L.; Hung, N. Q. V.; and Yin, H. 2023. XSimGCL: Towards extremely simple graph contrastive learning for recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 36(2): 913–926.
- 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*, 1294–1303.
- Zhang, D.; Geng, Y.; Gong, W.; Qi, Z.; Chen, Z.; Tang, X.; Shan, Y.; Dong, Y.; and Tang, J. 2024a. Recdcl: Dual contrastive learning for recommendation. In *Proceedings of the ACM Web Conference 2024*, 3655–3666.
- Zhang, P.; and Wu, M. 2024. Multi-Label Supervised Contrastive Learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 16786–16793.
- Zhang, Y.; Sang, L.; and Zhang, Y. 2024. Exploring the individuality and collectivity of intents behind interactions for graph collaborative filtering. In *Proceedings of the 47th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1253–1262.
- Zhang, Y.; Yin, C.; Wu, Q.; He, Q.; and Zhu, H. 2021. Location-Aware Deep Collaborative Filtering for Service Recommendation. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 51(6): 3796–3807.
- Zhang, Y.; and Zhang, Y. 2025. MixRec: Individual and Collective Mixing Empowers Data Augmentation for Recommender Systems. In *Proceedings of the ACM on Web Conference 2025*, 2198–2208.
- Zhang, Y.; Zhang, Y.; Sang, L.; and Sheng, V. S. 2024b. Simplify to the limit! embedding-less graph collaborative filtering for recommender systems. *ACM Transactions on Information Systems*, 43(1): 1–30.
- Zhang, Y.; Zhang, Y.; Yan, D.; Deng, S.; and Yang, Y. 2023. Revisiting graph-based recommender systems from the perspective of variational auto-encoder. *ACM Transactions on Information Systems*, 41(3): 1–28.
- Zhang, Y.; Zhang, Y.; Zhang, Y.; Sang, L.; and Yang, Y. 2025. Unveiling Contrastive Learning’s Capability of Neighborhood Aggregation for Collaborative Filtering. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1985–1994.