MACRec: A Multi-View Subspace Alignment Framework for Contrastive Sampling Calibration in Recommendation

Junping Liu[1*], Mingchao Yu[1*], Xinrong Hu[1], Rui Yan[2], Wanqing Li[3], Jie Yang[3†], Yi Guo[4†]
[1]Wuhan Textile University, Wuhan, China
[2]Gaoling School of Artificial Intelligence, Renmin University of China, Beijing, China
[3]University of Wollongong, Wollongong, Australia
[4]Western Sydney University, Sydney, Australia
{jpliu, hxr}@wtu.edu.cn,kidkdd78@gmail.com, ruiyan@ruc.edu.cn, {wanqing, jiey}@uow.edu.au,
y.guo@westernsydney.edu.au

Abstract

Graph Contrastive Learning (GCL) has proven effective in mitigating data sparsity and enhancing representation learning for recommendation. Yet, most GCL frameworks indiscriminately treat all non-anchor nodes as negatives during contrastive sampling, often leading to the *false negative* problem where semantically similar nodes are incorrectly repelled. Previous attempts to mitigate this issue rely on predetermined heuristics or local neighborhood mining, which struggle to reliably identify false negatives. More critically, they often overlook authentic user-item interactions for anchoring sample relationships. As a result, this paper presents *MACRec*, a Multi-View subspace-Alignment framework designed to Calibrate contrastive sampling in GCL-based Recommendation. *MACRec* comprises three core components: (1) a Multi-View Affinity (MVA) module that captures consistent semantic relations across multiple augmentations via self-expression modeling; (2) a Cross-Subspace Alignment (CSA) mechanism that leverages authentic user-item behavioral interactions to enforce semantic consistency across user and item subspaces; and (3) a Calibration-based Contrastive Reweighting (CCR) strategy to dynamically down-weight potential false negatives during the contrastive learning process. Extensive experiments on three real-world benchmarks demonstrate that *MACRec* consistently improves performance across various augmentation backbones, achieving up to 14.55% relative gains.

Introduction

Recently, **Graph Contrastive Learning (GCL)** has emerged as a powerful self-supervised paradigm for modern recommender systems (Wu et al. 2021; Ren et al. 2024a; Xiong et al. 2025). By generating multiple graph views and maximizing agreement between corresponding node embeddings, GCL-based methods significantly improve recommendation performance across various benchmarks.

Despite their effectiveness, most of these models rely on the standard InfoNCE contrastive loss, which indiscriminately treats all non-anchor samples as negatives, overlooking their potential semantic affinity. This design often leads to the **false negative** problem, where semantically similar entities (*e.g.*, users with shared interests or items within the same niche) are incorrectly repelled, impairing representation quality and downstream recommendation performance.

While some recent works attempt to alleviate this issue using similarity-based heuristics (He et al. 2023; Jing et al. 2024) or local neighborhood mining (Sun et al. 2023; Niu, Pang, and Chen 2024), these strategies exhibit inherent limitations. On one hand, heuristic rules are predefined and inflexible, often failing to generalize across diverse graph structures or dynamic user behaviors (Sheng et al. 2025; Xiong et al. 2025). On the other hand, local neighborhoods primarily capture short-range dependencies, overlooking semantically relevant but distant nodes. More critically, these approaches overlook authentic user-item interactions, limiting their ability to effectively distinguish true positives from false negatives during contrastive learning.

To address these limitations, this paper proposes a Multi-view subspace-Alignment framework to Calibrate contrastive sampling in GCL-based Recommendation (*MACRec*). At its core, *MACRec* employs data driven technique, in particular subspace learning in multi-view setting, to automatically discover robust semantic affinities. These learned affinities are then reinforced via interaction-guided alignment using behavioral signals to anchor user-item subspaces. *MACRec* further incorporates a calibration strategy that dynamically adjusts weights of negative samples based on both semantic affinity and behavioral interaction, effectively suppressing potential false negatives and improving representation quality. Extensive experiments on three benchmarks demonstrate that *MACRec* consistently outperforms state-of-the-art methods. Comprehensive ablation studies and component analyses further validate the effectiveness and generality of the proposed design. Moreover, *MACRec* is empirically shown to be model-agnostic and can be seamlessly integrated into various GCL backbones. The key contributions of this work are summarized as follows:

- A novel contrastive sampling calibration algorithm, *MACRec*, is proposed to address the false negative problem in GCL-based recommendation.
- *MACRec* introduces (i) multi-view affinity modeling to capture consistent semantic proximity; (ii) interaction-

∗ Equal contribution to this work.
† Corresponding authors.

Method	View Generation	Contrastive Sampling Strategy
♣ **SGL** (Wu et al. 2021)	Node/edge dropout and Randomwalk	Standard in-batch negatives
♣ **SimGCL** (Yu et al. 2022)	Embedding noisification	Standard in-batch negatives
♣ **NESCL** (Sun et al. 2023)	Node/edge dropout	Structural and semantic neighbors as additional positives
♣ **AuGCL** (Niu et al. 2024)	Node/edge dropout	Uncertainty-based hard negative reweighting
♢ **DCCF** (Ren et al. 2023)	Intent-aware masking via learnable generators	Standard in-batch negatives
♢ **CGCL** (He et al. 2023)	Multi-layer GNN embeddings	Semantic similarity-based positive selection
♢ **GraphAug** (Zhang et al. 2024)	Denoised graph sampling through reparameterization	Standard in-batch negatives
♢ **Sterling** (Jing et al. 2024)	Multi encoders	Negative free and similarity-based positive selection
♠ **AlphaRec** (Sheng et al. 2025)	Textual encoding from LLMs	Interaction-driven positive sampling
♠ **RLMRec** (Ren et al. 2024a)	LLM-enhanced user/item descriptions	Standard in-batch negatives
♠ **IRLLRec** (Wang et al. 2025)	LLM-guided intention	Standard in-batch negatives

Table 1: Summary of existing contrastive learning methods grouped by augmentation strategy, including: ♣ Random Augmentation; ♢ Adaptive Augmentation; and ♠ Semantic Augmentation.

guided cross-subspace alignment; and (iii) contrastive reweighting to dynamically suppress noisy negatives.

- Empirical results show that *MACRec* consistently outperforms state-of-the-art baselines while remaining model-agnostic and augmentation-compatible.

Related Work

GCL for Recommendation

Graph Contrastive Learning (GCL) has advanced recommendation tasks by enhancing user and item embeddings through pulling semantically similar (positive) pairs closer while pushing dissimilar (negative) pairs apart in the representation space (Wu et al. 2021; Xia et al. 2022). GCL primarily consists of two core modules: graph augmentation and contrastive sampling. Table 1 overviews several mainstream GCL methods, characterized by their augmentation and sampling strategies.

Graph augmentation techniques, *though not the main focus of this study*, enhance data diversity via generating multiple views of the original graph. Existing GCL-based recommendation models typically adopt one of three augmentation strategies: (i) *random perturbation*, such as edge/node dropout and noise injection (Wu et al. 2021; Yu et al. 2022); (ii) *adaptive/learnable augmentation*, which generates task-specific views via adversarial or performance-aware mechanisms (Jiang, Huang, and Huang 2023; Ren et al. 2023); and (iii) *semantic augmentation*, which leverages LLMs to enrich node semantics with textual contexts or reasoning paths (Ren et al. 2024a; Wang et al. 2025; Sheng et al. 2025). On the other hand, **contrastive sampling** plays a pivotal role in GCL. Conventional GCL frameworks typically assume that the same node across two augmented views forms a positive pair, while all other nodes are treated as negative samples (You et al. 2020; Hassani and Khasahmadi 2020). However, this assumption overlooks the fact that some nodes may be semantically similar to the anchor node. Treating these related nodes as negatives and forcing them apart in the latent space introduces conflicting training results, ultimately leading to the **false negative** problem (Yang et al. 2023; Ji et al. 2024).

To mitigate this issue, several strategies have been proposed. Some works adopt **sample reweighting** mechanisms, which adjust the contribution of negative samples by estimating their semantic similarity to the anchor (Liu et al. 2022; Liu, Yang, and Li 2024). Others perform **false-negative correction** through graph homophily or neighborhood-based heuristics (Sun et al. 2023; Sheng et al. 2025), treating semantically similar nodes as positives. More recently, **negative-free** contrastive learning has gained attention, exemplified by Sterling (Jing et al. 2024), which entirely removes negatives but relies on positives.

Despite the empirical success of existing sampling methods, a fundamental limitation remains: they often rely on heuristics or local structures that fail to capture global semantic consistency. In response, we propose to jointly model multi-view affinities and interaction-guided structural cues to dynamically calibrate contrastive sampling.

Subspace Learning

Subspace learning assumes that high-dimensional data often lies in a union of multiple low-dimensional subspaces (Elhamifar and Vidal 2013). A key principle is that each data sample can be linearly reconstructed by other samples belonging to the same subspace (Tierney, Gao, and Guo 2014; Yang, Bouzerdoum, and Phung 2009). Various subspace learning methods have been developed based on different regularization strategies. **Sparse representation** (Guo, Gao, and Li 2014), for example, aims at efficient reconstruction by selecting a minimal subset of representative samples. Formally, given the observed data matrix $Y \in \mathbb{R}^{d\times n}$, sparse representation seeks a coefficient matrix $X \in \mathbb{R}^{n\times n}$ and a reconstruction error term E by solving the following:

$$\min_{X,E} \|X\|_{2,1} + \lambda_E \|E\|_2, \quad \text{s.t.} \quad Y = YX + E, \qquad (1)$$

where $\|X\|_{2,1} = \sum_q \|X_q\|_2$, X_q is the q-th row, and λ_E is a penalty parameter.

Beyond single-view scenarios, **multi-view subspace learning** has also been developed to model data from multiple complementary perspectives. It aims to uncover consistent subspace across heterogeneous views while preserving view-specific information (Yang et al. 2022; Yu et al. 2025).

Given a total of V views, the objective is formulated as:

$$\begin{aligned} \min_{X,\{E^{(v)}\}} \quad & \|X\|_{2,1} + \sum_{v=1}^{V} \lambda_E \|E^{(v)}\|_2, \\ \text{s.t.} \quad & Y^{(v)} = Y^{(v)}X + E^{(v)}, \quad \forall v \in [1, V], \end{aligned} \tag{2}$$

where $Y^{(v)}$ and $E^{(v)}$ denotes the input matrix and the reconstruction error for view v.

Proposed Method

Preliminary

In recommendation systems, the interactions between the user set $\mathcal{U}$ and the item set $\mathcal{I}$ can be naturally represented as a bipartite graph $\mathcal{G} = (\mathcal{U} \cup \mathcal{I}, \mathcal{E})$, where $\mathcal{E}$ denotes observed user-item interactions. Let $\mathbf{A}$ be the adjacency matrix of $\mathcal{G}$ and $\mathbf{D}$ the diagonal degree matrix. Graph-based recommendation models, such as LightGCN (He et al. 2020), update node embeddings through L layers of message passing:

$$\mathbf{H}^{(l)} = (\mathbf{D}^{-\frac{1}{2}}\mathbf{A}\mathbf{D}^{-\frac{1}{2}})\mathbf{H}^{(l-1)}, \quad \forall\, l \in [1, L],$$

where $\mathbf{H}^{(l)}$ is the embedding matrix at layer l, and $\mathbf{H}^{(0)}$ is the original node feature matrix. For a user u and an item i, the final representations $\mathbf{h}_u$ and $\mathbf{h}_i$ are typically obtained by aggregating embeddings from all L layers. These embeddings are optimized using the **Bayesian Personalized Ranking (BPR)** loss, which encourages the model to assign higher scores to observed items than to unobserved ones:

$$\mathcal{L}_{\text{BPR}} = -\sum_{(u,i,j)} \log\big(\sigma(\mathbf{h}_u^\top \mathbf{h}_i - \mathbf{h}_u^\top \mathbf{h}_j)\big), \tag{3}$$

where $\sigma(\cdot)$ is the sigmoid function, and $R(u, i, j)$ denotes that user u interacted with item i but not with item j.

Contrastive learning is often introduced as an auxiliary objective, wherein graph augmentations are typically applied to generate additional view of $\mathcal{G}$. For each node $v \in \mathcal{U} \cup \mathcal{I}$, this results in two corresponding embeddings, denoted as $\mathbf{h}_v^{(1)}$ and $\mathbf{h}_v^{(2)}$, derived from two distinct views. The **InfoNCE** loss, denoted as $\mathcal{L}_{\text{CL}}$, encourages embeddings from two views of the same node to remain close, while pushing them away from those of other nodes:

$$z_{\text{pos}} = \exp\left(\frac{\text{sim}(\mathbf{h}_v^{(1)}, \mathbf{h}_v^{(2)})}{\tau}\right), \tag{4}$$

$$z_{\text{neg}} = \sum_{k \neq v} \exp\left(\frac{\text{sim}(\mathbf{h}_v^{(1)}, \mathbf{h}_k^{(2)})}{\tau}\right), \tag{5}$$

$$\mathcal{L}_{\text{CL}} = -\sum_{v \in \mathcal{U} \cup \mathcal{I}} \log\left(\frac{z_{\text{pos}}}{z_{\text{pos}} + z_{\text{neg}}}\right), \tag{6}$$

where $\text{sim}(\cdot, \cdot)$ denotes cosine similarity and τ is a temperature hyperparameter. Accordingly, the final training objective of GCL jointly optimizes recommendation performance and representation consistency:

$$\mathcal{L}_{\text{GCL}} = \mathcal{L}_{\text{BPR}} + \lambda \mathcal{L}_{\text{CL}}, \tag{7}$$

where λ is a penalty hyperparameter.

Multi-View Affinity modeling

Traditional GCL treats all non-anchor nodes as contrasting samples, which inevitably misclassifies semantically related nodes as negatives, resulting in the *false-negative problem*. To mitigate this issue, we introduce a *Multi-View subspace-aware Affinity (MVA) modeling*. Specifically, each node embedding is sparsely reconstructed as a linear combination of the remaining nodes, and the resulting self-expression coefficients (SECs) are used to quantify semantic affinity. Intuitively, nodes that reside in the same latent subspace as the anchor tend to exhibit similar characteristics; for example, users may share interaction behaviours or interests, while items may share semantic attributes or functionalities. As a result, such nodes typically receive large self-expression coefficient (SEC) values. These nodes are semantically close to the anchor and are likely to co-occur in recommendations. Therefore, they should not be treated as negative examples in contrastive learning. To this end, nodes with high SEC scores can be down-weighted from the contrastive negative set, aiming to preserve meaningful relational structures and enhance representation quality.

Assume there are M ($M \geq 1$) augmentations of the original interaction graph $\mathcal{G}$. For the user set $\mathcal{U}$, a shared coefficient matrix $\mathbf{C} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{U}|}$ is learned by minimizing the multi-view reconstruction loss:

$$\begin{aligned} \mathcal{L}_{\text{rec}}^{\mathcal{U}} &= \sum_{m=0}^{M} \big\|\mathbf{H}_{\mathcal{U}}^{(m)} - \mathbf{C}\mathbf{H}_{\mathcal{U}}^{(m)}\big\|_F^2 + \beta\Omega(\mathbf{C}) \\ \text{s.t.} \quad & \mathbf{C}_{ii} = 0, \quad \mathbf{C}_{ij} \geq 0, \quad \forall i, j. \end{aligned} \tag{8}$$

where $\mathbf{H}_{\mathcal{U}}^{(m)} \in \mathbb{R}^{|\mathcal{U}| \times d}$ is the user embedding matrix from the m-th view, with m=0 corresponding to the original graph, $\Omega(\mathbf{C}) = \|\mathbf{C}\|_{2,1}$ is a sparsity-inducing regularization term, and β is the regularization hyperparameter. Similarly, the item set $\mathcal{I}$ is treated analogously, resulting in a coefficient matrix $\mathbf{Z}$ and reconstruction loss $\mathcal{L}_{\text{rec}}^{\mathcal{I}}$:

$$\begin{aligned} \mathcal{L}_{\text{rec}}^{\mathcal{I}} &= \sum_{m=0}^{M} \big\|\mathbf{H}_{\mathcal{I}}^{(m)} - \mathbf{Z}\mathbf{H}_{\mathcal{I}}^{(m)}\big\|_F^2 + \beta\Omega(\mathbf{Z}) \\ \text{s.t.} \quad & \mathbf{Z}_{ii} = 0, \quad \mathbf{Z}_{ij} \geq 0, \quad \forall i, j. \end{aligned} \tag{9}$$

Sharing $\mathbf{C}$ (or $\mathbf{Z}$) across all M views enforces **multi-view consistency**, ensuring that the captured *subspace structures* remains stable and meaningful across different augmentations, rather than being artifacts of any specific one.

Furthermore, to solve this problem efficiently, the *low-rank factorization* technique is employed to parameterize the coefficient matrix as $\mathbf{C} = \mathbf{C}_A\mathbf{C}_B^\top$, where $\mathbf{C}_A, \mathbf{C}_B \in \mathbb{R}^{|\mathcal{U}| \times d_r}$ and $d_r \ll |\mathcal{U}|$. This formulation provides two benefits. First, it reduces the number of optimization parameters from $|\mathcal{U}| \times |\mathcal{U}|$ to $2 \times |\mathcal{U}| \times d_r$, significantly improving computational efficiency. Second, the low-rank constraint introduces an implicit structural prior on the coefficient matrix, which further helps mitigate overfitting and enhances generalization. An analogous factorization is also applied to $\mathbf{Z}$, *i.e.*, $\mathbf{Z} = \mathbf{Z}_A\mathbf{Z}_B^\top$, where $\mathbf{Z}_A, \mathbf{Z}_B \in \mathbb{R}^{|\mathcal{I}| \times d_r}$.

With the optimized $\mathbf{C}$ and $\mathbf{Z}$, symmetrization and normalization are applied to produce the following indicators:

$$\begin{cases} \mathbf{S}^{\mathcal{U}} = \text{Norm}\left(\frac{1}{2}(\mathbf{C} + \mathbf{C}^\top)\right), \\ \mathbf{S}^{\mathcal{I}} = \text{Norm}\left(\frac{1}{2}(\mathbf{Z} + \mathbf{Z}^\top)\right), \end{cases} \tag{10}$$

where $\text{Norm}(\cdot)$ denotes row-wise normalization. These matrices serve as **subspace affinity indicators**, quantifying the degree to which node pairs share the same subspace. They are subsequently used to regulate contrastive learning, that is, high-affinity nodes with anchors are down-weighted from the negative set to mitigate false-negative issue.

Cross-Subspace Alignment

Although the self-expression coefficients $\mathbf{C}$ and $\mathbf{Z}$ capture intra-subspace semantic affinities among users and items respectively, the two subspaces remain disconnected. Without an explicit alignment mechanism, the model may develop inconsistent or incompatible relational patterns across users and items, which can impair its ability to reason about interactions that span either user or item spaces. In other words, mismatched user and item subspaces may hinder the integration of user preferences with item characteristics.

To address this issue, a *Cross-Subspace semantic Alignment (CSA)* strategy is introduced to bridge the two sides. Specifically, the self-expressive relationships among users, encoded by the coefficient matrix $\mathbf{C}$, are encouraged to propagate through the user-item interaction graph (represented by the adjacency matrix $\mathbf{A}$) and remain consistent with the item-side relationships captured in $\mathbf{Z}$. As such, the proposed alignment leverages user–item interaction as a guiding signal to match the semantic information learned from the user subspace with that from the item subspace.

To this end, a learnable weight matrix $\mathbf{W} \in \mathbb{R}^{|\mathcal{U}|\times|\mathcal{I}|}$ is applied element-wise to the original interaction $\mathbf{A}$ to form the *subspace propagation matrix* $\mathbf{Q} = \mathbf{A} \odot \mathbf{W}$ ($\in \mathbb{R}^{|\mathcal{U}|\times|\mathcal{I}|}$). Rather than replicating original user-item interactions, $\mathbf{Q}$ serves to adjust the interaction strength based on subspace-level proximity, guiding semantic alignment between user and item subspaces. For example, if user u interacts with item i, *and* the self-expression coefficients indicate that user u is semantically close to u' ($C_{uu'} > 0$) and item i is close to i' ($Z_{ii'} > 0$), then the weights $W_{ui'}$ and $W_{u'i}$ are increased. This allows the original interaction (u, i) to be softly propagated to (u, i') and (u', i) within $\mathbf{Q}$. Consequently, to encourage semantic consistency across subspaces, the following alignment loss is proposed:

$$\mathcal{L}_{\text{ali}} = \left\| \mathbf{Z} - \mathbf{Q}^\top \mathbf{C} \mathbf{Q} \right\|_F^2 + \left\| \mathbf{C} - \mathbf{Q} \mathbf{Z} \mathbf{Q}^\top \right\|_F^2. \tag{11}$$

This bidirectional consistency ensures that the self-expression structures learned from both user and item subspaces are mutually aligned, aiming to support the recommendation reasoning across the user-item interaction space.

Notably, the u-th row vector from $\mathbf{Q}$ encodes user u's weighted preferences across items, while the i-th column vector from $\mathbf{Q}$ reflects item i's weighted attractiveness to users. Therefore, **interaction-aware behavioral indicators** can be derived from $\mathbf{Q}$ to provide an additional signal for identifying semantically aligned nodes, which again can be used to guide contrastive sampling. Specifically, user-user and item-item similarities are computed as $\mathbf{B}^{\mathcal{U}} = \mathbf{Q}\mathbf{Q}^\top$ and $\mathbf{B}^{\mathcal{I}} = \mathbf{Q}^\top\mathbf{Q}$, respectively. Given a predefined threshold θ, the following behavioral indicators are obtained as:

$$\begin{cases} \mathbf{R}^{\mathcal{U}} = \text{Norm}\left(\mathbf{B}^{\mathcal{U}} \odot \mathbb{I}_{\{\mathbf{B}^{\mathcal{U}} > \theta\}}\right), \\ \mathbf{R}^{\mathcal{I}} = \text{Norm}\left(\mathbf{B}^{\mathcal{I}} \odot \mathbb{I}_{\{\mathbf{B}^{\mathcal{I}} > \theta\}}\right), \end{cases} \tag{12}$$

where $\mathbb{I}_{\{\cdot\}}$ is the element-wise indicator function. These behavioral indicators complement the previously introduced subspace affinity indicators (*i.e.*, $\mathbf{S}^{\mathcal{U}/\mathcal{I}}$). While $\mathbf{S}^{\mathcal{U}/\mathcal{I}}$ capture latent proximity derived from multi-view subspaces, $\mathbf{R}^{\mathcal{U}/\mathcal{I}}$ characterize similarity based on interaction behaviors. Together, they offer orthogonal yet complementary perspectives for identifying semantically related node pairs.

Overall Objective

The subspace affinity scores $\mathbf{S}^{\mathcal{U}/\mathcal{I}}$ and behavioral indicators $\mathbf{R}^{\mathcal{U}/\mathcal{I}}$ are fused to estimate the false-negative risk via:

$$\mathbf{P}^{\mathcal{X}} = \max(\mathbf{S}^{\mathcal{X}}, \mathbf{R}^{\mathcal{X}}),$$

where $\mathcal{X} \in \{\mathcal{U}, \mathcal{I}\}$, and $\max(\cdot,\cdot)$ denote the element-wise maximum operator. For each anchor node v, the **calibration-based contrastive reweighting** for another node k ($\neq v$) is then introduced as:

$$w_{vk} = \exp\left(-\mathbf{P}^{\mathcal{X}}_{vk}\right), \tag{13}$$

which adaptively suppresses the influence of semantically similar negatives. Following Eq. (4) and (5), the calibration-aware contrastive learning loss is formulated as:

$$\mathcal{L}_{\text{CL-Cali}} = -\sum_{v \in \mathcal{U}\cup\mathcal{I}} \log\left(\frac{z_{\text{pos}}}{z_{\text{pos}} + w_{vk} z_{\text{neg}}}\right). \tag{14}$$

Finally, the proposed *MACRec* is formulated as the following multi-task learning:

$$\mathcal{L} = \mathcal{L}_{\text{BPR}} + \lambda_{\text{rec}}(\mathcal{L}^{\mathcal{U}}_{\text{rec}} + \mathcal{L}^{\mathcal{I}}_{\text{rec}}) + \lambda_{\text{ali}}\mathcal{L}_{\text{ali}} + \lambda_{\text{CL}}\mathcal{L}_{\text{CL-Cali}}, \tag{15}$$

where λ_{rec}, λ_{ali}, and λ_{CL} are hyperparameters.

Remark 1. *The weights in Eq.* (13) *allow MACRec to dynamically adjust the contribution of negatives, which is lower-bounded by the affinity automatically discovered by subspace learning and reinforced by the interaction strength. Specifically, when $\mathbf{P}^{\mathcal{X}}_{vk}$ is large, indicating strong evidence that node k is semantically close to the anchor, w_{vk} becomes nearly zero (*e.g.*, $\exp(-18) \approx 10^{-8}$), effectively equivalent to removing node k from the set of negatives in GCL. Otherwise, node k stays as negative to contrast the anchor's representation. Our design circumvents the **ad hoc** nature of the heuristics mentioned earlier, providing a hands-off solution to mitigate false negatives. Additionally, the concept of subspace learning has been well studied. Several works, such as (Guo, Tierney, and Gao 2021), have proposed subspace discovery guarantees that*

can effectively handle large datasets under certain conditions. These studies provide a solid foundation for the effectiveness of MACRec. Finally, to further increase the flexibility, one could further introduce a temperature hyperparameter t as $w_{vk} = \exp\left(-\mathbf{P}^{\mathcal{X}}_{vk} / t\right)$ to control the sensitivity. To avoid the determination of the value of t, in this paper, we fix $t = 1$ for simplicity.

Experiments

Experimental Setup

Datasets. Three popular benchmarks, *i.e.*, Amazon-Book, Yelp (Ren et al. 2024a), and Amazon-Movie (Wang et al. 2025), are considered in this paper. Table 2 summarizes key statistics of these datasets.

Dataset	#Users	#Items	#Interactions	Density
Amazon-Book	11,000	9,332	200,860	2.0×10^{-3}
Yelp	11,091	11,010	277,535	2.3×10^{-3}
Amazon-Movie	16,994	9,370	168,243	1.1×10^{-3}

Table 2: Statistics of the employed benchmarks.

Baselines. `MACRec` focuses on the contrastive sampling to mitigate the impact of false negatives. As such, **to ensure fair comparison, we retain the original graph augmentation strategies of the baselines and vary only the contrastive sampling component.** This design isolates other impacts but focus on the contrastive sampling. Various graph augmentations are included: **random** (NESCL, AuGCL, SGL, SimGCL), **adaptive** (GraphAug, CGCL, Sterling, DCCF), and **semantic** (AlphaRec, IRLLRec, RLMRec). See Table 1 for more details. *Their results are either directly sourced from original papers or reproduced using released source codes with default configurations.*

Implementation Details. The proposed method is implemented using **SSLRec** (Ren et al. 2024b). Specifically, user and item embeddings are randomly initialized with a latent dimension of 32. The Adam optimizer is employed with a learning rate of 10^{-3} and a batch size of 4096. Early stopping is applied based on Recall@10 performance on the validation set, with a patience of 5 epochs.

Overall Performance

Table 3 presents a detailed performance comparison across different methods and datasets, with all results averaged over five independent runs. The experimental results demonstrate that *MACRec* consistently achieves superior or comparable performance across all three benchmarks. Specifically: (1) when applied to existing models, our contrastive sampling strategy significantly boosts their performance (*e.g.*, SGL+ outperforms SGL by 6.84% on the Amazon-Book dataset); (2) compared to prior sampling methods, our approach also delivers notable gains (*e.g.*, SGL+ outperforms AuGCL by 6.36% on the Amazon-Movie dataset. This can be attributed to the limitations of existing strategies, which often suffer from false negative noise due to reliance on ad hoc heuristics or local neighborhoods. In contrast, our method leverages structural and semantic cues to better distinguish false negatives, promoting more discriminative and informative embedding and enhanced recommendation accuracy.

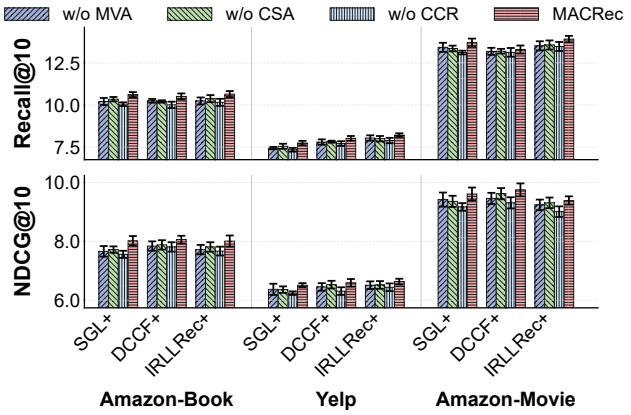

Figure 1: Impact from the proposed MVA, CSA and CCR components across three GCL backbones.

Model Breakdown

The proposed *MACRec* is composed of three core components: (1) Multi-view Affinity (**MVA**) modeling, (2) interaction-guided Cross-Subspace Alignment (**CSA**), and (3) Calibration-based Contrastive Reweighting (**CCR**). To assess the individual impact of each module, ablation studies are conducted by selectively removing one component at a time (*i.e.*, removing the relevant loss from model fine-tuning). To demonstrate the flexibility of *MACRec*, it is further integrated into three representative GCL backbones (SGL, DCCF, and IRLLRec) while retaining their original graph augmentation strategies. This yields the enhanced variants **SGL+**, **DCCF+**, and **IRLLRec+**, while their corresponding results are reported in Fig. 1. As observed, all components positively contribute to the overall performance, with noticeable drops observed when any single component is removed. For instance, removing either MVA or CSA leads to performance degradation, indicating that both multi-view affinity modeling and cross-subspace alignment play complementary roles in capturing richer semantic and structural relations. By contrast, excluding CCR results in the most significant performance decline, highlighting its critical role in alleviating the impact of false negatives. This is likely because CCR ultimately integrates MVA and CSA, enabling them to work jointly in identifying false negatives.

Ablation Study

Evaluating Alternative Designs. This subsection investigates alternative variants, via replacing each component (*i.e.*, MVA, CSA, and CCR) with one substitute in *MACRec*. Specifically, for MVA, we employ a conventional similarity-based approach (labeled **Alt-Aff**), where $\mathbf{S}^{\mathcal{U}/\mathcal{I}}$ are computed using cosine similarity. For CSA, we leverage the original adjacency matrix $\mathbf{A}$ to derive $\mathbf{R}^{\mathcal{U}/\mathcal{I}}$ (labeled **Alt-Align**). For

Method	Amazon-Book				Yelp				Amazon-Movie			
	Recall		NDCG		Recall		NDCG		Recall		NDCG	
	@5	@10	@5	@10	@5	@10	@5	@10	@5	@10	@5	@10
NESECL♣	6.42	10.01	6.46	7.53	4.53	7.55	5.32	6.19	9.21	12.67	7.45	8.71
AuGCL♣	6.38	10.05	6.33	7.58	4.66	7.62	5.41	6.30	9.27	12.88	7.34	8.62
SGL♣	6.37	9.94	6.32	7.56	4.32	7.22	5.01	5.92	9.17	12.39	7.22	8.35
SGL+♣	**6.85** (± 0.11)	**10.62** (± 0.15)	**6.88** (± 0.12)	**8.02** (± 0.16)	4.73 (± 0.10)	7.74 (± 0.13)	5.55 (± 0.05)	6.51 (± 0.07)	10.29 (± 0.21)	**13.70** (± 0.25)	**8.52** (± 0.18)	9.61 (± 0.22)
SimGCL♣	6.18	9.92	6.19	7.49	4.67	7.72	5.46	6.38	9.25	12.76	7.38	8.61
SimGCL+♣	6.57 (± 0.09)	10.25 (± 0.12)	6.55 (± 0.08)	7.85 (± 0.11)	**4.85** (± 0.07)	**7.95** (± 0.09)	**5.62** (± 0.06)	**6.55** (± 0.08)	**10.38** (± 0.15)	13.62 (± 0.18)	8.45 (± 0.12)	**9.62** (± 0.15)
GraphAug♢	6.64	10.17	6.54	7.83	4.61	7.68	5.47	6.32	9.49	12.97	7.58	8.81
CGCL♢	6.39	10.09	6.51	7.61	4.58	7.61	5.32	6.23	9.46	12.65	7.45	8.67
Sterling♢	6.42	10.23	6.37	7.68	4.64	7.70	5.33	6.42	9.37	12.87	7.49	8.87
DCCF♢	6.62	10.19	6.58	7.80	4.68	7.78	5.43	6.40	9.65	13.17	7.63	9.05
DCCF+♢	**6.93** (± 0.14)	**10.51** (± 0.18)	**6.84** (± 0.09)	**8.06** (± 0.13)	**4.82** (± 0.09)	**8.02** (± 0.14)	**5.60** (± 0.08)	**6.59** (± 0.13)	**10.61** (± 0.21)	**13.30** (± 0.24)	**8.74** (± 0.19)	**9.75** (± 0.22)
RLMRec♠	6.08	9.69	6.06	7.34	4.45	7.54	5.18	6.14	8.34	11.93	6.52	7.78
AlphaRec♠	5.98	9.41	6.05	7.21	4.31	7.26	4.93	5.86	9.09	12.13	7.06	8.27
IRLLRec♠	6.43	10.09	6.38	7.65	4.60	7.81	5.42	6.42	9.50	13.41	7.52	8.88
IRLLRec+♠	**6.68** (± 0.11)	**10.65** (± 0.18)	**6.62** (± 0.09)	**8.01** (± 0.19)	**4.86** (± 0.08)	**8.21** (± 0.11)	**5.55** (± 0.07)	**6.63** (± 0.10)	**10.26** (± 0.18)	**13.92** (± 0.19)	**8.12** (± 0.11)	**9.39** (± 0.14)

Table 3: Performance comparison of various GCL methods, grouped by argumentation strategy: ♣ (Random), ♢ (Adaptive), and ♠ (Semantic). The **best** and second-best results are highlighted. Methods with "**+**" adopt our sampling strategy.

CCR, we remove potentially false-negative pairs whose calibration scores exceed a fixed threshold (**Alt-Reweight**, set to 0.5). Furthermore, we also employ a false-negative correction strategy (labeled **Alt-Pos**) to re-label identified false-negative samples as positives. Similarly, three backbones (SGL, DCCF, and IRLLRec) are employed.

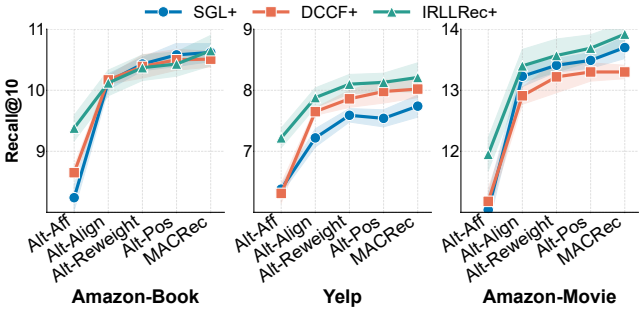

Figure 2: Performance comparison of *MACRec* against its four alternative variants across three datasets.

As shown in Fig. 2, Alt-Aff consistently yields the poorest performance. This is likely because cosine similarity primarily captures local feature proximity, which often identifies immediate neighbors as similar nodes. Alt-Align ranks second-worst in performance. This suggests that constructing structural relations purely based on raw adjacency misses meaningful semantic cues propagated through actual user-item interactions. Finally, both Alt-Reweight and Alt-Pos perform slightly worse than CCR. Simply discarding high-calibration pairs (Alt-Reweight) may result in information loss by removing potentially informative negatives. Conversely, reassigning them as positives (Alt-Pos) may introduce semantic noise, as not all high-scoring pairs are truly similar. These findings highlight the limitations of heuristic false-negative handling and demonstrate the necessity of calibrated reweighting, which adaptively preserves sample diversity while mitigating the impact of false negatives.

Number of Views. This subsection examines the impact from the number of augmented views (M from Eq. (8)). Specifically, for each training instance, we randomly select one augmentation type from the followings: (1) edge/node dropout, (2) random walk, and (3) embedding perturbation, for comparison. As shown in Fig. 3, using the SGL backbone, incorporating diverse augmentation views improves Recall@10 by introducing more complementary contrastive signals. Yet, performance worsens beyond five views, likely due to overfitting or inconsistent view diversity.

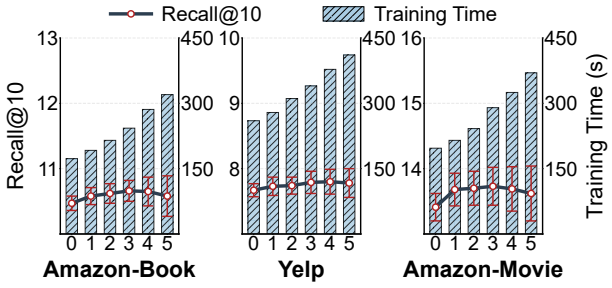

Figure 3: Effect of the number of augmented views(M) on recommendation performance and training time.

Additionally, we also observe that training time grows approximately linearly with the number of views, mainly due to the MVA module, which requires computing self-expression coefficient for multi view. On the other hand, the computational cost of CSA and CCR remains largely stable regardless of the increasing M. Despite the added overhead, the overall cost is a reasonable trade-off for the notable performance gains in reducing false negatives.

Impact of Low-rank Factorization. *MACRec* employs low-rank decomposition to parameterize the self-expression coefficient matrix (see Eq. (8) and (9)). This section presents an in-depth analysis of alternative parameterization strategies to validate the performance and computational cost. Specifically, the DCCF model is used as the backbone, and the following variants are evaluated: **Dense**, *i.e.*, removes the low-rank constraint and employs one fully dense coefficient matrix; **No-Sparsity**, *i.e.*, removes the sparsity regularization term; and **Neighborhood**, *i.e.*, limits self-expression to nodes within a K-hop neighborhood ($K = 2$).

Strategy	Amazon-Book	Yelp	Amazon-Movie
Dense	10.67 (302)	8.11 (434)	13.27 (591)
No-Sparsity	10.03 (153)	7.44 (226)	12.88 (271)
Neighborhood	10.44 (261)	7.92 (322)	13.11 (408)
Low-Rank	10.51 (160)	8.02 (209)	13.30 (328)

Table 4: Impact of self-expression parameterization strategies, in terms of Recall@10 and training time.

As observed from Table 4, the Dense variant achieves the best performance, benefiting from its unrestricted expressiveness; however, it suffers from substantial training costs. The No-Sparsity variant achieves the fastest training speed; yet, without sparsity regularization, it compromises representational capacity, resulting in the poor performance. The Neighborhood variant restricts representations to 2-hop neighbors, reducing training time but missing global dependencies, leading to degraded performance. In contrast, the Low-Rank strategy factorizes the coefficient matrix into two low-dimensional matrices, retaining key expressiveness while greatly reducing complexity. As such, it achieves comparable performance with much lower computational cost than the Dense variant.

Hyperparameter Analysis. We also consider a sensitivity analysis by tuning one hyperparameter at a time, while keeping others constant. To begin with, λ_{rec}, λ_{ali}, λ_{cl} are varied across $\{0.01, 0.05, 0.1, 0.2, 0.5, 0.8, 1.0\}$ to examine their influence. As illustrated in Fig. 4 (a), the model performance improves with increasing λ_{rec}, reaching a peak around moderate values (*e.g.*, 0.2 or 0.5), before slightly declining. This indicates that a moderate value of λ_{rec} is crucial for optimal performance. A similar trend is observed for λ_{ali} in Fig. 4 (b), where performance peaks at $\lambda_{ali} = 0.1$ before declining. In contrast, the optimal value of λ_{cl} varies: larger values benefit SGL, while smaller values are preferred for DCCF and IRLLRec, likely due to differences in their original contrastive setting.

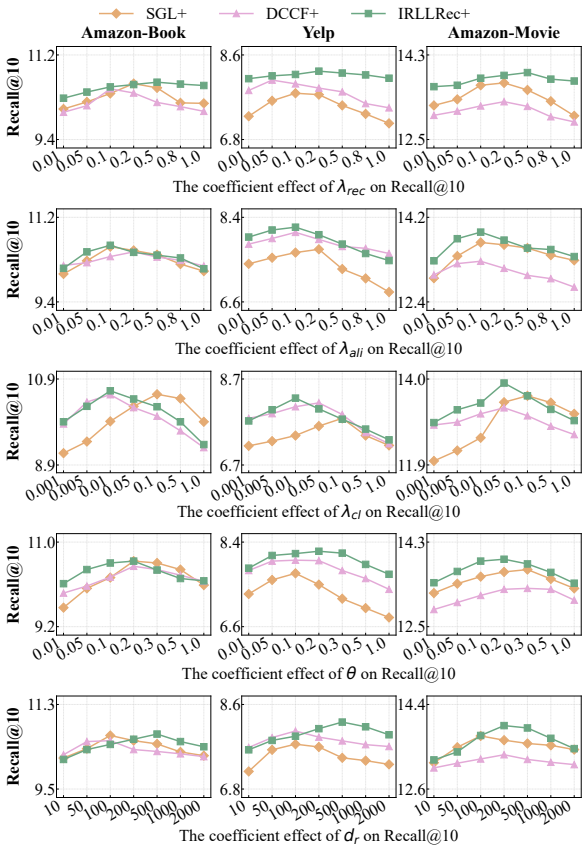

Figure 4: Hyperparameter sensitivity analysis.

For θ in Eq. (12), we consider values in the range of $\{0.01, 0.05, 0.1, 0.2, 0.3, 0.5, 1.0\}$. As depicted in Fig. 4 (d), performance generally declines as θ increases, since a larger threshold results in fewer false negatives being identified. Conversely, a smaller θ may introduce many noisy negatives, degrading contrastive learning quality as well. Finally, the low-rank dimension d_r is set to $\{10, 50, 100, 200, 500, 1000, 2000\$. We observe that increasing d_r leads to consistent performance gains, with the most effective range found to be between 100 and 1000 (Fig. 4(e)), as this range provides sufficient representation capacity to preserve informative features while avoiding over-parameterization and noise amplification with smaller or larger d_r.

Conclusion

This paper introduces *MACRec*, a novel contrastive sampling calibration framework for recommendation. By leveraging Multi-View Affinity (MVA), Cross-Subspace Alignment (CSA), and Calibration-based Contrastive Reweighting (CCR), *MACRec* effectively mitigates the false negative problem and enhances semantic alignment. Extensive experiments on several recommendation benchmarks demonstrate that *MACRec* consistently outperforms existing GCL-based approaches. Future work may extend *MACRec* to incorporate other sampling strategies, explore applications in heterogeneous graphs and large-scale scenarios.

Acknowledgements

This work is partially supported by Chinese Ministry of Education Humanities and Social Sciences General project (Grant No. 23YJAZH082), Hubei Provincial Natural Science Foundation (Grant No. 2025AFC097, China), Hubei Province Education Science Planning Key Project (Grant No. 2022GA046, China), Australian Research Council Discovery Project (DP210101426), and Telstra-UOW Hub for AIOT Solutions Seed Funding (2024, 2025).

References

Elhamifar, E.; and Vidal, R. 2013. Sparse subspace clustering: Algorithm, theory, and applications. *IEEE transactions on pattern analysis and machine intelligence*, 35(11): 2765–2781.

Guo, Y.; Gao, J.; and Li, F. 2014. Spatial subspace clustering for drill hole spectral data. *Journal of Applied Remote Sensing*, 8(1): 083644.

Guo, Y.; Tierney, S.; and Gao, J. 2021. Efficient sparse subspace clustering by nearest neighbour filtering. *Signal Processing*, 185: 108082.

Hassani, K.; and Khasahmadi, A. H. 2020. Contrastive multi-view representation learning on graphs. In *International conference on machine learning*, 4116–4126. PMLR.

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, C.; Huang, Z.; Sun, Q.; Peng, H.; Fu, X.; Li, Q.; and Li, J. 2024. ReGCL: Rethinking message passing in graph contrastive learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 8544–8552.

Jiang, Y.; Huang, C.; and Huang, L. 2023. Adaptive graph contrastive learning for recommendation. In *Proceedings of the 29th ACM SIGKDD conference on knowledge discovery and data mining*, 4252–4261.

Jing, B.; Yan, Y.; Ding, K.; Park, C.; Zhu, Y.; Liu, H.; and Tong, H. 2024. STERLING: Synergistic representation learning on bipartite graphs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 12976–12984.

Liu, C.; Yang, J.; and Li, W. 2024. Extractive Question Answering with Contrastive Puzzles and Reweighted Clues. In *Document Analysis and Recognition - ICDAR 2024*, 97–112.

Liu, J.; Mei, S.; Hu, X.; Yao, X.; Yang, J.; and Guo, Y. 2022. Seeing the wood for the trees: a contrastive regularization method for the low-resource Knowledge Base Question Answering. In *Findings of the Association for Computational Linguistics: NAACL 2022*, 1085–1094.

Niu, C.; Pang, G.; and Chen, L. 2024. Affinity uncertainty-based hard negative mining in graph contrastive learning. *IEEE Transactions on Neural Networks and Learning Systems*, 35(9): 11681–11691.

Ren, X.; Wei, W.; Xia, L.; Su, L.; Cheng, S.; Wang, J.; Yin, D.; and Huang, C. 2024a. Representation learning with large language models for recommendation. In *Proceedings of the ACM web conference 2024*, 3464–3475.

Ren, X.; Xia, L.; Yang, Y.; Wei, W.; Wang, T.; Cai, X.; and Huang, C. 2024b. SSLRec: A self-supervised learning framework for recommendation. In *Proceedings of the 17th ACM international conference on web search and data mining*, 567–575.

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.

Sheng, L.; Zhang, A.; Zhang, Y.; Chen, Y.; Wang, X.; and Chua, T. 2025. Language Representations Can be What Recommenders Need: Findings and Potentials. In *ICLR*.

Sun, P.; Wu, L.; Zhang, K.; Chen, X.; and Wang, M. 2023. Neighborhood-enhanced supervised contrastive learning for collaborative filtering. *IEEE Transactions on Knowledge and Data Engineering*, 36(5): 2069–2081.

Tierney, S.; Gao, J.; and Guo, Y. 2014. Subspace Clustering for Sequential Data. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, 1019–1026. IEEE.

Wang, Y.; Sang, L.; Zhang, Y.; and Zhang, Y. 2025. Intent Representation Learning with Large Language Model for Recommendation. In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval*, 1870–1879.

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.

Xia, L.; Huang, C.; Xu, Y.; Zhao, J.; Yin, D.; and Huang, J. 2022. Hypergraph contrastive collaborative filtering. In *Proceedings of the 45th International ACM SIGIR conference on research and development in information retrieval*, 70–79.

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, J.; Bouzerdoum, A.; and Phung, S. L. 2009. A Neural Network pruning approach based on Compressive Sampling. In *2009 International Joint Conference on Neural Networks*, 3428–3435.

Yang, J.; Ma, J.; Win, K. T.; Gao, J.; and Yang, Z. 2022. Low-rank and sparse representation based learning for cancer survivability prediction. *Information Sciences*, 582: 573–592.

Yang, X.; Liu, Y.; Zhou, S.; Wang, S.; Tu, W.; Zheng, Q.; Liu, X.; Fang, L.; and Zhu, E. 2023. Cluster-guided

contrastive graph clustering network. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, 10834–10842.

You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems*, 33: 5812–5823.

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.

Yu, X.; Jiang, Y.; Chao, G.; and Chu, D. 2025. Deep Contrastive Multi-View Subspace Clustering With Representation and Cluster Interactive Learning. *IEEE Transactions on Knowledge and Data Engineering*, 37(1): 188–199.

Zhang, Q.; Xia, L.; Cai, X.; Yiu, S.-M.; Huang, C.; and Jensen, C. S. 2024. Graph augmentation for recommendation. In *2024 IEEE 40th International Conference on Data Engineering (ICDE)*, 557–569. IEEE.