PRGCL: PATH RECONSTRUCTION FOR GRAPH CONTRASTIVE LEARNING IN COLLABORATIVE FILTERING

Anonymous authors
Paper under double-blind review

ABSTRACT

Graph contrastive learning (GCL) has recently emerged as an effective self-supervised enhancement paradigm for alleviating data sparsity. However, existing methods still face three key challenges: (1) Stochastic structural perturbations may compromise local structure, making the constructed views unreliable; (2) The mixing of frequency domain signals can cause noise interference, leading to ambiguous optimization objectives; (3) Independently constructed views may induce different connectivity in the signal propagation path. To address these challenges, we propose a coupled reconstruction principle that leverages path reconstruction to identify task relevant structural subsets, guiding self-supervised view enhancement for graph collaborative filtering. Building on this, we propose **P**ath **R**econstruction for **G**raph **C**ontrastive **L**earning in collaborative filtering (PRGCL), which improves graph contrastive recommendation by constructing reliable self-supervised views, refining frequency domain signals, and enforcing connectivity along signal propagation paths. Extensive experiments on three highly sparse real-world datasets demonstrate that PRGCL consistently outperforms state-of-the-art baselines in recommendation accuracy.

1 INTRODUCTION

Personalized recommender systems aim to infer users' preferences from historical interactions and subsequently generate appropriate and suitable recommendations, constituting the cornerstone of contemporary e-commerce platforms(Kunaver & Požrl, 2017). A prevalent paradigm is collaborative filtering (CF)(Rendle et al., 2009; He et al., 2020b), which leverages user similarity and item attributes to construct high quality embedding representations of users and items for downstream tasks(Koren et al., 2021). Mainstream CF methods devised to accomplish this objective include matrix factorization (MF)(Rendle et al., 2009), multilayer perceptrons (MLP)(He et al., 2017), graph convolutional networks (GCN)(He et al., 2020b; Liu et al., 2026), and graph contrastive learning (GCL)(Wang et al., 2020; Lin et al., 2022; Cai et al.). The inherent heterogeneity of user-item interaction data naturally induces a bipartite graph structure (e.g., Fig. 1a). LightGCN(He et al., 2020b) performs convolution on the normalized adjacency matrix of the user-item graph, which has constituted a standard paradigm in graph-based recommender systems for a considerable period. Recent studies demonstrate that GCL enhances the expressiveness of user and item embedding representations(Li et al., 2024; Yu et al., 2023) because it propagates informative self-supervised signals, and this capability has gradually established GCL as a new mainstream methodology(Zhang et al., 2024).

The construction methodology of contrastive views classifies existing GCL methods into structure augmentation and data augmentation(He et al., 2023). Structure augmentation(Wu et al., 2021) generates augmented views by perturbing the graph structure and acquires information on central nodes by contrasting invariant elements during this process, extracting meaningful supervisory signals. Data augmentation(Yu et al., 2023; Zhang et al., 2024) disrupts the positional stability of node information through the injection of random noise or other mechanisms, aligning node embeddings with their optimal positions in the distribution to achieve superior representations. Although existing methods have demonstrated the potential of GCL, traditional augmentation techniques may reduce the intensity of self-supervised signal propagation, which leads to structural reliability degradation, different signal interference, and propagation topology inconsistency, undermining the overall performance of the contrastive learning method. Specifically, a node is randomly selected from the interaction graph

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

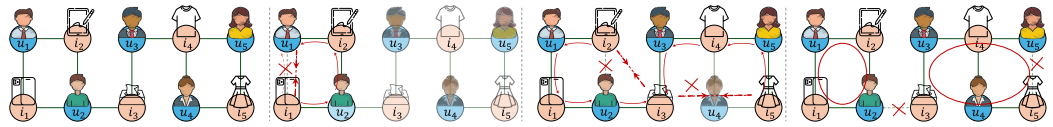

(a) user-item interaction bipartite graph (b) structural reliability degradation (c) different signal interference (d) propagation topology inconsistency

Figure 1: A simplified bipartite graph and three failure scenarios.

as the target node, and the path extending from this target node is designated as the central path. However, traditional augmentation techniques may introduce the following three failure scenarios:

- *Structural reliability degradation.* As shown in Fig. 1b, when the edge between u_1 and i_1 is removed, the signal propagation path becomes $i_1 \to u_2 \to i_2 \to u_1$. This operation disrupts the original local structure around the target node, weakens the structural dependency relationship between the key neighboring node i_1(Android phone) and the target node u_1(employees), and thus reduces the effectiveness of local signal propagation.

- *Different signal interference.* As shown in Fig. 1c, removing the edge between u_2 and i_2 changes the effective neighborhood of i_3. The original second-order propagation from i_2 (iPad) to i_3 is replaced by a fourth-order path, $i_2 \to u_1 \to i_1 \to u_2 \to i_3$, while i_5 (dress) also reaches i_3 through another fourth-order path, $i_5 \to u_5 \to i_4 \to u_3 \to i_3$. As a result, different order frequency domain signals are introduced into the same high order aggregation scope as the originally relevant signals. Since message passing aggregates signals according to topological proximity, these inconsistent frequency domain components can inject noisy signals into representation learning.

- *Propagation topology inconsistency*. As shown in Fig. 1d, the structural enhancement operation may alter the connected components of the original signal propagation paths $u_2 \to i_2 \to u_1 \to i_1 \to u_2$ and $i_3 \to u_3 \to i_4 \to u_5 \to i_5 \to u_4$. By removing the edges associated with u_2 and i_3, as well as those associated with u_5 and i_5, two distinct augmented views are constructed. Specifically, the first augmented view becomes disconnected after the edge removal, whereas the second augmented view still preserves connectivity but exhibits a different connected structure. Since these views contain connected components that are inconsistent with those in the original graph, they result in homotopically inconsistent views and undermine the topological consistency of signal propagation.

To address the above challenges, we first revisit the topological organization of self-supervised signal propagation paths in GCL. We observe that the quality of contrastive views is mainly characterized by three criteria: preservation of key shortest propagation path structures, separation of frequency domain signals across different orders, and homotopic consistency among contrastive views. From this perspective, constructing contrastive views solely through random perturbation is insufficient, and a coupled reconstruction principle guided by these three criteria is also needed.

Following this intuition, we propose a coupled reconstruction principle that leverages path reconstruction to identify task relevant structural subsets, guiding self-supervised view enhancement for graph collaborative filtering. Building on this, we propose **P**ath **R**econstruction for **G**raph **C**ontrastive **L**earning in collaborative filtering (PRGCL), which improves graph contrastive recommendation by constructing reliable self-supervised views, refining frequency domain signals, and enforcing connectivity along signal propagation paths, so that the resulting views remain connected through the necessary signal propagation structures. The main contributions of this paper are as follows:

- We re-examine GCL from the perspective of the topological structure of self-supervised signal propagation paths and establish the coupled reconstruction principle.

- We propose PRGCL, a path reconstruction graph contrastive recommendation method. PRGCL jointly enhances self-supervised views from the perspective of path reconstruction.

- We conduct extensive experiments on three highly sparse real-world datasets, and the results demonstrate that PRGCL consistently outperforms existing graph contrastive recommendation methods in terms of recommendation accuracy.

2 Related Work and Preliminaries

Graph Neural Network in Recommendation. In recent years, graph neural networks (GNNs) (Hamilton et al., 2017; Kipf & Welling, 2017) have increasingly complemented or replaced matrix factorization (MF), becoming a predominant approach in recommender systems. Among the diverse GNN variants, graph convolutional networks (GCNs) have demonstrated exceptional performance across various recommendation scenarios, surpassing previous approaches. GCMC(Berg et al., 2017) applies graph convolutions directly to the user-item interaction graph to enhance embedding propagation. NGCF(Wang et al., 2019) explicitly models high-order connectivity within user-item graphs to capture more intricate collaborative signals. LightGCN(He et al., 2020b) simplifies GCN by removing feature transformations and nonlinear activations, focusing solely on embedding propagation to achieve greater efficiency and performance. LR-GCCF(Chen et al., 2020a) further introduces low-rank constraints to reduce model complexity.

Self-Supervised Learning in Recommendation. Self-supervised learning (SSL)(Chen et al., 2020b; Grill et al., 2020; He et al., 2020a; Wu et al., 2018) seeks to derive meaningful representations from unlabeled data by constructing supervisory signals directly from the data itself. In this context, Contrastive Learning (CL) represents a prevalent paradigm within SSL and is predicated on the principle that similar samples should be brought closer while dissimilar samples are separated in the embedding space. GCL methods employ contrastive objectives over graph structured user and item data to enhance representation learning and consequently improve recommendation accuracy under sparse supervision. Building upon this foundation, numerous approaches have been proposed, including SGL(Wu et al., 2021), NCL(Lin et al., 2022), SimGCL(Yu et al., 2022), LightGCL(Cai et al.), CGCL(He et al., 2023) and MAWU(Park et al., 2023). Furthermore, DirectAU(Wang et al., 2022) and SCCF(Wu et al., 2024) demonstrate that the objectives of BPR loss(Rendle et al., 2009) and SSL loss align in terms of alignment and uniformity.

Preliminaries. Let $\mathcal{G} = (\mathcal{U} \cup \mathcal{I}, \mathcal{E})$ denote a user-item bipartite graph, where $\mathcal{U}$ and $\mathcal{I}$ represent the user and item sets, respectively, and $\mathcal{E}$ represents the interaction edges between users and items(Tran et al., 2018). The embedding dimension of each user and item is defined as d. The interaction between users and items can be represented by an interaction matrix $\mathbf{R} \in \mathbb{R}^{|\mathcal{U}| \times |\mathcal{I}|}$, where $R_{ui} = 1$ if user u has interacted with item i, and 0 otherwise. The adjacency matrix is defined as $\mathbf{A} \in \mathbb{R}^{(|\mathcal{U}|+|\mathcal{I}|) \times (|\mathcal{U}|+|\mathcal{I}|)}$. GNN-based methods primarily learn information from neighboring nodes through recursive neighborhood aggregation operations in the graph, obtaining embedding representations of users and items that incorporate both neighbor information and intrinsic features.

Definition 2.1 (GNN Learning Process (He et al., 2020b))**.** *At the k-th layer, the representation of the user node u is obtained by propagating and combining the embeddings from its neighboring nodes:*

$$E_u^{(k)} = \sum_{v \in \mathcal{N}(u)} \alpha_{uv}^{(k)} E_v^{(k-1)}, \quad E_u = \sum_{k=0}^{K} \beta_k E_u^{(k)}, \tag{1}$$

where $E_u^{(k)}$ denotes the embedding of user node u at layer k, $\mathcal{N}(u)$ denotes the set of neighboring nodes of node u, $\alpha_{uv}^{(k)}$ denotes the propagation weight from node v to node u at layer k, and β_k denotes the contribution weight of the k-th layer representation. The embedding process for item nodes can be formulated in an analogous manner.

Definition 2.2 (Signal Propagation Path in GCL)**.** *Write $V := \mathcal{U} \cup \mathcal{I}$ for the node set. For any $s, t \in V$ and integer $K \geq 1$, define the set of all admissible propagation paths of length at most K from s to t by $\mathcal{P}_K(s,t) = \{p_l = (v_0, \ldots, v_l) \mid v_0 = s,\ v_l = t,\ 1 \leq l \leq K,\ (v_{r-1}, v_r) \in \mathcal{E}\ \forall r = 1, \ldots, l\}$, and let $\mathcal{P}_K = \bigcup_{s,t \in V} \mathcal{P}_K(s,t)$.*

Definition 2.3 (Effective Self-supervised Signal)**.** *For each edge $(u,i) \in \mathcal{E}$, let $P_{ui} \in (0,1]$ be its propagation coefficient and $\kappa_{ui} \in (0,1]$ its frequency domain coefficient. For a path $p = (v_0, \ldots, v_l)$ define the transmission weight $w(p) = \prod_{r=1}^{l} P_{v_{r-1}v_r} \kappa_{v_{r-1}v_r}$. The effective self-supervised signal between v_a and v_b within K orders is $\mathcal{S}_K(v_a, v_b) = \sum_{p \in \mathcal{P}_K(v_a, v_b)} w(p)$.*

Based on Def. 2.1 to Def. 2.3, we observe that the optimization of node representations is closely tied to the specific paths along which self-supervised signals propagate during graph enhancement.

3 METHODOLOGY

3.1 PATH TOPOLOGY PERSPECTIVE OF SELF-UPERVISED SIGNAL PROPAGATION

Existing GCL methods implicitly assume that self-supervised signals remain reliable under stochastic view perturbations (Yu et al., 2022; 2023). However, the validity of self-supervised signals depends on whether the induced propagation structures preserve key shortest propagation paths, separate frequency-domain signals across different propagation orders, and maintain homotopic consistency among contrastive views. Therefore, the quality of contrastive views should be assessed not merely by the perturbation intensity over edges or nodes, but by whether the perturbed graph induces a propagation topology that preserves the reliability and consistency of self-supervised signal propagation. To formalize this perspective, we introduce a truncated path complex, which treats self-supervised propagation signals as weighted topological objects indexed by admissible propagation paths. Based on this formulation, we extend the path structures from the original interaction graph $\mathcal{G}$ to a contrastive view $\widetilde{\mathcal{G}} = (\widetilde{\mathcal{U}} \cup \widetilde{\mathcal{I}}, \widetilde{\mathcal{E}})$, and establish the constraints through the following definition and propositions.

Definition 3.1 (Truncated Propagation Path Complex)**.** *For a path complex $\mathfrak{P}_K(\mathcal{G})$ of path p_l, let $C_l(\mathfrak{P}_K(\mathcal{G}); \mathbb{R}) = \mathrm{span}_{\mathbb{R}}\{e_p \mid p \in \mathcal{P}_K,\ |p| = l\}$ be the real vector space generated by all admissible l-paths, where e_p denotes the basis element associated with path p. The boundary operator ∂_l is defined by $\partial_l(e_p) = \sum_{j=0}^{l}(-1)^j e_{p^{\hat{j}}}$, where $p^{\hat{j}} = (v_0, \ldots, \widehat{v_j}, \ldots, v_l)$ if this vertex-deleted sequence remains admissible in $\mathfrak{P}_K(\mathcal{G})$, and $e_{p^{\hat{j}}} = 0$ otherwise.*

Extending Definition 3.1 to a contrastive view $\widetilde{\mathcal{G}} = (\mathcal{U} \cup \mathcal{I}, \widetilde{\mathcal{E}})$ with similarly defined path structures $\mathfrak{P}_K(\widetilde{\mathcal{G}})$, $\mathcal{P}_K^{\widetilde{\mathcal{G}}}(s,t)$, and $\widetilde{\mathcal{S}}_K(s,t)$, we introduce Prop. 1 to Prop. 3, which formalize the three failure scenarios mentioned above and establish the corresponding constraints.

Proposition 1. *Fix $v_a, v_b \in \mathcal{U} \cup \mathcal{I}$, and define $d_{\mathcal{G}}(v_a, v_b) = \min_{p \in \mathcal{P}_K(v_a, v_b)} |p|$. Let the total weight of all shortest admissible paths between v_a and v_b in $\mathcal{G}$ be $M_K^{\mathrm{ess}}(v_a, v_b) = \sum_{\substack{p \in \mathcal{P}_K(v_a, v_b) \\ |p| = d_{\mathcal{G}}(v_a, v_b)}} w(p)$. If no shortest path in $\mathcal{P}_K(v_a, v_b)$ remains admissible in $\widetilde{\mathcal{G}}$, then*

$$\widetilde{\mathcal{S}}_K(v_a, v_b) \leq \mathcal{S}_K(v_a, v_b) - M_K^{\mathrm{ess}}(v_a, v_b). \tag{2}$$

Proposition 2. *Let $p = (v_0, \ldots, v_a, \ldots, v_b, \ldots, v_l) \in \mathcal{P}_K$ with $0 \leq a < b \leq l$ and $(v_a, v_b) \in \mathcal{E}$. The subpath $q = (v_a, \ldots, v_b)$ denotes the corresponding contiguous segment. The path $p' = (v_0, \ldots, v_a, v_b, \ldots, v_l)$ represents the elementary contraction of q to the chord (v_a, v_b). The operator $T : C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R}) \to C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R})$ is defined on basis paths by replacing each occurrence of q as a contiguous subpath with (v_a, v_b) while all other basis paths remain unchanged.*

Proposition 3. *A contiguous subpath $(v_a, \ldots, v_b)$ of an admissible path is called strictly dominated if $p_{v_a v_b}\kappa_{v_a v_b} > \prod_{j=a+1}^{b} p_{v_{j-1} v_j}\kappa_{v_{j-1} v_j}$. A path $p = (v_0, \ldots, v_l)$ is called ρ-*tailed *if there exists $m \in \{0, \ldots, l-1\}$ such that $\kappa_{v_{r-1} v_r} \leq \rho$ for every $r \in \{m+1, \ldots, l\}$ and $\rho \in (0,1)$. The set $\mathcal{P}_K^{\mathrm{rel}}(\mathcal{G}) \subseteq \mathcal{P}_K(\mathcal{G})$ consists of all admissible paths that contain no strictly dominated contiguous subpath and are not ρ-tailed.*

Proof. See Appendix A.3. □

Prop. 1 to Prop. 3 jointly formalize the three failure scenarios described above from the perspective of admissible propagation paths. Specifically, Prop. 1 shows that disrupting indispensable shortest path clusters will weaken the effective self-supervised signal between nodes, formalizing the necessity of preserving key shortest propagation paths. Prop. 2 characterizes the topological redundancy between the weak high order paths and their direct shortcuts through the contraction operator T, providing an operable criterion for separating frequency domain signals across different propagation orders and suppressing cross order interference. Prop. 3 further introduces the relativized path family $\mathcal{P}_K^{\mathrm{rel}}(\mathcal{G})$ by excluding strictly dominated subpaths and ρ-tailed paths, so that the remaining admissible paths preserve only those paths that are minimally reliable for self-supervised propagation and remain consistent across contrastive views. Node representations are optimized through signal propagation along admissible paths. Preserving key shortest propagation path structures, separating mixed frequency-domain signals across different orders, and maintaining homotopic consistency

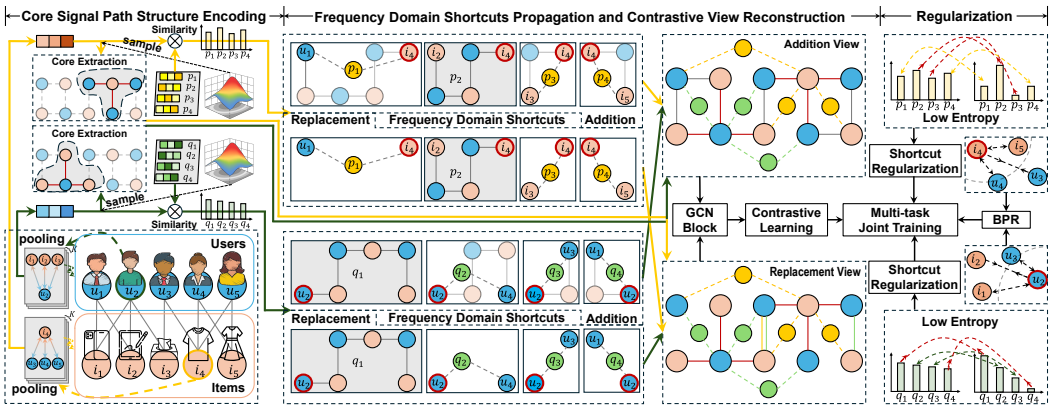

Figure 2: Overview of the proposed PRGCL framework. i) Core signal path structure encoding starts from the original user-item interaction graph and the GCN block, where path reconstruction is performed to extract key shortest propagation path structures while preserving indispensable signal propagation paths. ii) Frequency domain shortcuts refinement computes affinity assignments from the base embeddings and injects weighted reconstruction connections into the contrastive views to separate frequency domain signals across different propagation orders. CL is conducted on the addition view and replacement view to align node representations under homotopically consistent reconstructed propagation structures. iii) Multi-task joint training integrates the contrastive objective with the BPR objective and low entropy regularization to improve both recommendation performance and self-supervised signal reliability.

among contrastive views help reduce the unreliability induced by stochastic perturbation, making the contrastive objective more compatible with the downstream ranking objective.

3.2 PRGCL

In Section 3.1, we discuss the mathematical propositions corresponding to the three failure scenarios. These propositions characterize how contrastive views should be reorganized relative to the original interaction graph to preserve reliable self-supervised signal propagation structures. Based on this, we propose a coupled reconstruction principle that identifies task-relevant structural subsets to guide self-supervised view enhancement for downstream recommendation, providing a principled rationale within the graph-based collaborative filtering paradigm (Li et al., 2022a;b; Wu et al.). We summarize path contrastive reconstruction through a single abstract criterion that jointly enforces the preservation of key shortest propagation path structures, the separation of frequency domain signals across different propagation orders, and the homotopic consistency among contrastive views. Specifically, let $\Psi(\cdot)$ map an interaction graph to its induced multihop effective self-supervised signal field $\mathcal{S}_K$, let $\mathcal{D}$ quantify discrepancies in key shortest propagation path retention relative to the original topology, let $\mathcal{C}$ measure the residual cross-order frequency interference between the original propagation paths and their T-contracted counterparts, and let $\mathcal{I}$ score the homotopic incompatibility between relativized path families induced by two candidate views at the level of $\mathcal{P}_K^{\mathrm{rel}}(\cdot)$. For two contrastive views $\widetilde{\mathcal{G}}^{(1)}$ and $\widetilde{\mathcal{G}}^{(2)}$ derived from $\mathcal{G}$, the coupled reconstruction principle is formally expressed as

$$\min_{\widetilde{\mathcal{G}}^{(1)},\widetilde{\mathcal{G}}^{(2)}} \mathcal{D}(\Psi(\mathcal{G}),\{\Psi(\widetilde{\mathcal{G}}^{(m)})\}_{m=1}^{2}) + \mathcal{C}(\mathfrak{P}_K(\mathcal{G}),\{\mathfrak{P}_K(\widetilde{\mathcal{G}}^{(m)})\}_{m=1}^{2};T) + \mathcal{I}(\mathcal{P}_K^{\mathrm{rel}}(\widetilde{\mathcal{G}}^{(1)}),\mathcal{P}_K^{\mathrm{rel}}(\widetilde{\mathcal{G}}^{(2)})). \quad (3)$$

Building on the coupled reconstruction principle in Eq. 3, PRGCL reconstructs the path topology perspective in Section 3.1 by turning contrastive view construction into a path reconstruction procedure guided by reliable self-supervised signal propagation. As illustrated in Fig. 2, the reconstruction procedure in each training step is organized into three coupled stages: core signal path structure encoding, frequency domain shortcuts propagation and contrastive view reconstruction; and joint optimization for downstream recommendation. We associate each training step with two augmented graphs, $\widetilde{\mathcal{G}}^{(1)}$ and $\widetilde{\mathcal{G}}^{(2)}$, derived from $\mathcal{G}$, together with a set of learnable reconstruction affinities that serve as auxiliary connections for signal propagation. Let $\mathcal{F}_\Theta$ denote the stacked neighborhood aggregation and layer aggregation that maps any graph instance to user and item embeddings, where

Θ collects all trainable parameters. For a node $v \in \mathcal{U} \cup \mathcal{I}$, write $h_v = \mathcal{F}_\Theta(v\,;\,\mathcal{G})$ for embeddings obtained on the original interaction topology, and $h_v^{(m)} = \mathcal{F}_\Theta(v\,;\,\widetilde{\mathcal{G}}^{(m)})$ for embeddings under the m-th view ($m \in \{1,2\}$). The design of $\widetilde{\mathcal{G}}^{(m)}$ and the coupling between views are organized to address the three failure scenarios formalized from Fig. 1b to Fig. 1d. We provide a further analysis of the connections among the three propositions, the coupled reconstruction principle, and PRGCL in Appendix A.4.1. We also provide further discussion of PRGCL in Appendix A.5.

3.2.1 Preserving Key Shortest Propagation Path Structures

To preserve the integrity of the minimal path cluster structure, PRGCL aims to automatically extract salient graph structures from the interaction graph $\mathcal{G}$ through coupled reconstruction. Each reconstructed view in PRGCL retains the key path clusters within the original interaction edges, minimizing the probability of the attenuation effect described in Prop. 1. The PRGCL reconstruction graph is constructed through two fundamental operations: *addition* and *replacement*. These operations allow previously interrupted self-supervised signals to be reconnected through reconstruction affinities and restore critical connectivity under local perturbations induced by stochastic edge operations. Specifically, a sparse deletion set $\mathcal{R}^{(1)} \subset \mathcal{E}$ is constructed for the first view, while the second view preserves the complete edge set:

$$\begin{aligned} &\mathcal{E}_{\text{core}}^{(1)} = \mathcal{E} \setminus \mathcal{R}^{(1)}, \quad \mathcal{E}_{\text{core}}^{(2)} = \mathcal{E}, \\ &\widetilde{\mathcal{E}}^{(m)} = \mathcal{E}_{\text{core}}^{(m)} \cup \mathcal{E}_{\text{aux}}^{(m)}, \quad m \in \{1,2\}, \end{aligned} \tag{4}$$

where $\mathcal{E}_{\text{core}}^{(m)} \subseteq \mathcal{E}$ preserves key shortest propagation path structures and constitutes the structural backbone of the m-th reconstructed view, while $\mathcal{E}_{\text{aux}}^{(m)}$ is determined by the auxiliary edge construction strategy for separating frequency domain signals across different propagation orders in Section 3.2.2.

The explicit enumeration of all admissible propagation paths $p \in \mathcal{P}_K$ that pass through each node $n \in V$ is computationally prohibitive because $|\mathcal{P}_K|$ grows rapidly with graph size and truncation depth K and the shortest key path naturally exists in the form of a high probability within the randomly sampled paths. Therefore, PRGCL employs randomized subsampling of observed interactions, which is widely adopted in GCL(Wu et al., 2021). The training interactions are indexed by $\ell \in \{1,\dots,L\}$, and each interaction is denoted as $e_\ell = (u_\ell, i_\ell) \in \mathcal{E}$, where L is the length of the interaction list underlying $\mathcal{E}$. The hyperparameters $\alpha, \beta \in (0,1)$ denote hyperparameters for the addition and replacement views, respectively. For the replacement view $m = 1$, two independent uniformly sampled index sets $J_{\mathcal{U}}^{(1)}, J_{\mathcal{I}}^{(1)} \subseteq \{1,\dots,L\}$ are constructed, each with cardinality βL under sampling without replacement, and the set $\mathcal{R}^{(1)} = \{e_\ell \mid \ell \in J_{\mathcal{U}}^{(1)} \cup J_{\mathcal{I}}^{(1)}\}$ is defined as in Eq. 4. For the addition view $m = 2$, the set $\mathcal{R}^{(2)} = \emptyset$ is specified, and the index sets $J_{\mathcal{U}}^{(2)}, J_{\mathcal{I}}^{(2)}$ with cardinality αL are sampled using the same strategy and retained as candidate anchors for subsequent shortcut augmentation. The set $\mathcal{E}_{\text{core}}^{(m)}$ therefore contains only retained user-item interaction edges, which serve as the minimal path cluster backbone for the subsequent reconstruction.

3.2.2 Separating Cross Order Frequency Domain Signals

Given the two core views constructed in Section 3.2.1, PRGCL refines signal propagation by introducing weighted auxiliary shortcuts for separating frequency domain signals across different propagation orders. In this stage, we focus on the cross order signal intensity and separate the mixed frequency domain components in Eq. 3. We adopt GCN (He et al., 2020b) as the backbone, which propagates signals on the user-item graph via a simplified GCN without feature transformation. Let $\mathbf{P} \in \mathbb{R}^{P_u \times d}$ and $\mathbf{Q} \in \mathbb{R}^{Q_i \times d}$ denote learnable frequency domain affinity matrices for users and items, respectively, where P_u and Q_i are hyperparameters. Let $\mathbf{p}_r \in \mathbb{R}^d$ denote the r-th row of $\mathbf{P}$, and let $\mathbf{q}_s \in \mathbb{R}^d$ denote the s-th row of $\mathbf{Q}$. We first compute the base representations h_u for users and h_i for items on the original interaction graph $\mathcal{G}$. Following prior research (Zhang et al., 2025c), we employ cosine similarity ($\text{sim}(\cdot,\cdot)$) and the softmax function to compute the affinity intensity:

$$\pi_{u,r} = \frac{\exp \text{sim}(h_u, \mathbf{p}_r)}{\sum_{r'=1}^{P_u} \exp \text{sim}(h_u, \mathbf{p}_{r'})}, \qquad \pi_{i,s} = \frac{\exp \text{sim}(h_i, \mathbf{q}_s)}{\sum_{s'=1}^{Q_i} \exp \text{sim}(h_i, \mathbf{q}_{s'})}. \tag{5}$$

The quantities in Eq. 5 specify the associations between base user-item nodes and learnable frequency domain matrices, which are derived from the base aggregation on $\mathcal{G}$. To instantiate $\mathcal{E}_{\text{aux}}^{(m)}$ in Eq. 4, we

take the sampled interaction edges as anchors and aggregate the assignment scores of the opposite endpoints, constructing cross side auxiliary shortcuts induced by the sampled edges. Specifically, for each $m \in \{1, 2\}$, the sampled opposite endpoint multisets and the corresponding cross side aggregation scores are defined as:

$$\begin{aligned} \mathcal{N}_u^{(m)} &= \{\, i_\ell | \ell \in J_{\mathcal{U}}^{(m)},\ u_\ell = u \,\}, \quad \bar{\pi}_{u,s}^{(m)} = \frac{1}{|\mathcal{N}_u^{(m)}|} \sum_{i \in \mathcal{N}_u^{(m)}} \pi_{i,s}, \quad s = 1, \ldots, Q_i, \\ \mathcal{M}_i^{(m)} &= \{\, u_\ell | \ell \in J_{\mathcal{I}}^{(m)},\ i_\ell = i \,\}, \quad \bar{\pi}_{i,r}^{(m)} = \frac{1}{|\mathcal{M}_i^{(m)}|} \sum_{u \in \mathcal{M}_i^{(m)}} \pi_{u,r}, \quad r = 1, \ldots, P_u. \end{aligned} \tag{6}$$

Based on these aggregated scores, PRGCL retains at most q weighted cross side auxiliary shortcuts for each selected user and item node. Let $\mathrm{Top}_q(\cdot)$ represent the index set of the top q largest elements in the score matrices. Then, the auxiliary edge set is defined as:

$$\mathcal{E}_{\mathrm{aux}}^{(m)} = \bigcup_u \left\{ (u, \mathbf{q}_s, \bar{\pi}_{u,s}^{(m)}) \;\middle|\; s \in \mathrm{Top}_q(\bar{\pi}_{u,:}^{(m)}) \right\} \cup \bigcup_i \left\{ (\mathbf{p}_r, i, \bar{\pi}_{i,r}^{(m)}) \;\middle|\; r \in \mathrm{Top}_q(\bar{\pi}_{i,:}^{(m)}) \right\}. \tag{7}$$

These weighted cross side auxiliary shortcuts are then combined with the retained backbone to form the weighted connectivity of $\widetilde{\mathcal{G}}^{(m)}$. We then adjoin the learnable frequency domain affinity matrices $\mathbf{P}$ and $\mathbf{Q}$, symmetrize the union of these weighted auxiliary shortcuts with the unit weight core edges $\mathcal{E}_{\mathrm{core}}^{(m)}$, apply GCN degree normalization to the resulting adjacency, and instantiate a weighted graph $\widetilde{\mathcal{G}}^{(m)}$ on which the same aggregation $\mathcal{F}_\Theta$ is evaluated. Whenever contrastive views are refreshed, $\pi_{u,r}$ and $\pi_{i,s}$ are recomputed from the base aggregation $\{h_v\}_{v \in \mathcal{U} \cup \mathcal{I}}$ on $\mathcal{G}$.

3.2.3 Aligning Representations across Contrastive Views

Given the two shortcut augmented views $\widetilde{\mathcal{G}}^{(1)}$ and $\widetilde{\mathcal{G}}^{(2)}$, we next discuss how PRGCL scores the incompatibility between the relativized path families induced by the two reconstructed views. The explicit construction of $\mathcal{P}_K^{\mathrm{rel}}(\cdot)$ at scale is computationally prohibitive when K and the graph grow. Therefore, we instantiate $\mathcal{I}$ as a contrastive objective over the sampled interaction batch $\mathcal{B}$:

$$\begin{aligned} \mathcal{I}(a, b, c) &= -\mathbb{E}_{(a,b) \sim \mathcal{B}}[\mathrm{sim}(a, b)/\tau - \log \mathbb{E}_{(a,c) \sim \mathcal{B}}[\exp(\mathrm{sim}(a, c)/\tau)]], \\ \mathcal{L}_{\mathrm{ssl}} &= \mathcal{I}(h_u^{(0)}, h_u^{(1)}, h_u^{(2)}) + \mathcal{I}(h_i^{(0)}, h_i^{(1)}, h_i^{(2)}) + \mathcal{I}(h_u^{(0)}, h_i^{(0)}, h_i^{(0)}), \end{aligned} \tag{8}$$

where τ denotes the temperature coefficient, and the superscript denotes the base view and the contrastive views. Beyond cross view alignment, we regularize the learnable frequency domain affinity matrices induced by $\mathbf{P}$ and $\mathbf{Q}$ so that each learnable row exhibits anisotropy after softmax normalization. Specifically, for each user side affinity row $\mathbf{p}_r$, we apply softmax across the d coordinates to obtain $\sigma_{r,j} = \exp(\mathbf{p}_{r,j}) / \sum_{j'=1}^{d} \exp(\mathbf{p}_{r,j'})$, and we define the item probabilities from $\{\varsigma_{s,j}\}_{j=1}^{d}$ from $\mathbf{q}_s$ in the same strategy. Formally, the affinity regularizer aggregates the mean Shannon entropy(Lin, 2002) of these row anisotropy distributions on both sides:

$$\mathcal{L}_{\mathrm{car}} = -\frac{1}{P_u} \sum_{r=1}^{P_u} \sum_{j=1}^{d} \sigma_{r,j} \log \sigma_{r,j} - \frac{1}{Q_i} \sum_{s=1}^{Q_i} \sum_{j=1}^{d} \varsigma_{s,j} \log \varsigma_{s,j}. \tag{9}$$

3.2.4 Multi-Task Joint Objective

Although PRGCL couples discriminative recommendation with self-supervised consistency across views and a regularizer on affinity matrices, the lack of clear optimization signals may lead to suboptimal performance in executing the primary CF task. Following previous works(Yu et al., 2023; 2022), we introduce the Bayesian Personalized Ranking (BPR)(Rendle et al., 2009) loss as follows:

$$\mathcal{L}_{\mathrm{bpr}} = -\log(\sigma(\mathrm{sim}(e_u, e_i) - \mathrm{sim}(e_u, e_j))), \tag{10}$$

where σ denotes the sigmoid function, (u, i, j) denotes a BPR triplet, with u representing the user, i representing the positive item, and j representing the negative item. Through multi-task joint learning, we employ the BPR loss as the primary objective to guide the optimization of PRGCL, enhancing the recommendation quality as follows:

$$\mathcal{L} = \mathcal{L}_{\mathrm{bpr}} + \lambda_{\mathrm{reg}} \mathcal{L}_{\mathrm{reg}} + \lambda_{\mathrm{ssl}} \mathcal{L}_{\mathrm{ssl}} + \lambda_{\mathrm{car}} \mathcal{L}_{\mathrm{car}}, \tag{11}$$

where λ_{reg} is the regularization weight, $\mathcal{L}_{\mathrm{reg}}$ is the regularization loss, and λ_{ssl} and λ_{car} are the nonnegative coefficients.

Table 1: A comparison between the proposed PRGCL method and state-of-the-art baselines, with best results in bold and second-best underlined. The 'p-value' indicates statistically significant improvements (t-test $p < 0.05$) over the second-best baseline.

Settings		Baseline Methods								Ours	
Dataset	Metric	LightGCN	DirectAU	SimGCF	SGL	SimGCL	CGCL	DimCL	MixRec	**PRGCL**	p-Values
Douban-book	Recall@10	0.1025	0.1153	0.1178	0.1153	<u>0.1230</u>	0.1216	0.1141	0.1147	**0.1247**	5.66e-06
	NDCG@10	0.1359	0.1527	0.1585	0.1558	<u>0.1637</u>	0.1622	0.1417	0.1527	**0.1721**	1.63e-04
	Recall@20	0.1504	0.1660	0.1699	0.1633	<u>0.1772</u>	0.1741	0.1557	0.1668	**0.1811**	5.07e-05
	NDCG@20	0.1404	0.1568	0.1625	0.1585	0.1583	<u>0.1667</u>	0.1417	0.1578	**0.1757**	1.94e-04
Alibaba-iFashion	Recall@10	0.0669	<u>0.0800</u>	0.0653	0.0788	0.0795	0.0782	0.0784	0.0796	**0.0837**	1.05e-10
	NDCG@10	0.0382	<u>0.0468</u>	0.0365	0.0439	0.0465	0.0429	0.0474	0.0463	**0.0484**	1.28e-08
	Recall@20	0.1078	<u>0.1182</u>	0.0963	0.1065	0.1151	0.1101	0.1174	0.1162	**0.1223**	4.86e-12
	NDCG@20	0.0507	<u>0.0571</u>	0.0447	0.0519	0.0567	0.0523	0.0571	0.0559	**0.0586**	1.19e-09
Tmall	Recall@10	0.0435	0.0475	0.0487	0.0457	0.0557	0.0512	0.0538	<u>0.0558</u>	**0.0560**	4.58e-07
	NDCG@10	0.0406	0.0443	0.0454	0.0434	0.0535	0.0510	0.0500	<u>0.0536</u>	**0.0537**	4.54e-08
	Recall@20	0.0711	0.0752	0.0791	0.0738	<u>0.0840</u>	0.0820	0.0853	0.0875	**0.0890**	2.18e-05
	NDCG@20	0.0530	0.0576	0.0585	0.0556	<u>0.0615</u>	0.0619	0.0622	0.0673	**0.0677**	6.23e-08

4 EXPERIMENTS

In this section, we conduct comparisons with state of-the-art methods on three real-world datasets to demonstrate the effectiveness of PRGCL. Detailed settings, path diagnostics study, case study, sparsity study, robustness study and full parameter sensitivity study are provided in Appendix A.7.

4.1 EXPERIMENTAL SETTINGS

We utilize three widely used public benchmark datasets in our experiments: Douban-book(Yu et al., 2022), Alibaba-iFashion(Zhang et al., 2025b) andTmall(Ren et al., 2023). We randomly split each dataset into training and test sets with a ratio of 80% to 20%, respectively, and adapt two standard metrics: Recall@K and NDCG@K, where $K \in \{10, 20\}$. To evaluate performance differences, we conduct a comparison between PRGCL and eight state-of-the-art recommendation approaches: **AU- and GCN-based** methods: LightGCN(He et al., 2020b), DirectAU(Wang et al., 2022), and SimGCF(Liu et al., 2026); **GCL-based** methods: SGL(Wu et al., 2021), SimGCL(Yu et al., 2022), CGCL (He et al., 2023), DimCL(Zhang et al., 2025a), and MixRec(Zhang & Zhang, 2025).

4.2 PERFORMANCE COMPARISON WITH BASELINES

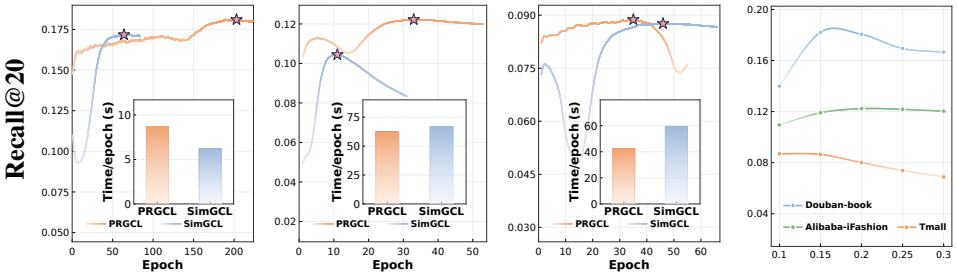

(a) Douban-book (b) Alibaba-iFashion (c) Tmall (d) τ

Figure 3: (a) - (c) Comparison of the training process and training time cost of SimGCL and PRGCL. The red star indicates the best performing on test sets; (d) hyperparameter study of coefficient τ.

To evaluate the effectiveness of PRGCL, we compare it with representative state-of-the-art baselines on three sparse real-world datasets, and the results in Table 1 show that PRGCL consistently achieves the best performance across all datasets and evaluation metrics, demonstrating its ability to learn informative user-item representations from sparse interaction graphs and confirming the effectiveness of the proposed path reconstruction mechanism. The superiority of PRGCL mainly comes from its path reconstruction strategy. Although PRGCL improves recommendation accuracy, the additional path reconstruction module may introduce extra computational costs. As shown in Fig. 3, PRGCL

and the base model exhibit different training-efficiency patterns across datasets. Although the relative training speed varies depending on dataset characteristics, PRGCL still achieves highly competitive final performance. Notably, PRGCL demonstrates strong performance in the early stages of training, indicating that it can learn more discriminative representations within fewer training epochs and effectively improve the quality of early stage convergence.

4.3 Ablation Study

Table 2: Ablation Experiments on the Douban-book, Alibaba-iFashion and Tmall datasets.

	Douban-book				Alibaba-iFashion				Tmall			
	Recall@10	**NDCG@10**	**Recall@20**	**NDCG@20**	**Recall@10**	**NDCG@10**	**Recall@20**	**NDCG@20**	**Recall@10**	**NDCG@10**	**Recall@20**	**NDCG@20**
(a) Contrastive Objective												
+w/o CL	0.0580	0.0790	0.0924	0.0832	0.0642	0.0357	0.0949	0.0431	0.0458	0.0429	0.0750	0.0555
(b) Reconstruction Regularization												
+w/o CR	0.1221	0.1613	0.1733	0.1649	0.0828	0.0475	0.1202	0.0566	0.0521	0.0501	0.0832	0.0634
+w/o PR	0.1188	0.1558	0.1701	0.1603	0.0805	0.0462	0.1168	0.0551	0.0522	0.0501	0.0832	0.0634
(c) Path Reconstruction Mechanisms												
+w/o RR	0.1029	0.1373	0.1472	0.1404	0.0825	0.0473	0.1190	0.0563	0.0519	0.0499	0.0830	0.0632
+w/o AR	0.1063	0.1369	0.1528	0.1427	0.0827	0.0474	0.1202	0.0566	0.0509	0.0488	0.0814	0.0618
+w/o AL	0.1222	0.1648	0.1758	0.1683	0.0831	0.0475	0.1201	0.0565	0.0522	0.0499	0.0832	0.0632
PRGCL	**0.1247**	**0.1721**	**0.1811**	**0.1757**	**0.0837**	**0.0483**	**0.1223**	**0.0586**	**0.0560**	**0.0537**	**0.0890**	**0.0677**

- *+w/o CL*: Remove the self-supervised contrastive objective from PRGCL.
- *+w/o CR*: Remove the coupled reconstruction regularization imposed on the learnable affinities.
- *+w/o PR*: Remove the path reconstruction module from both contrastive views.
- *+w/o RR*: Keep only the replacement reconstructed view as the effective path enhanced branch.
- *+w/o AR*: Keep only the addition reconstructed view as the effective path branch.
- *+w/o AL*: Replace learned reconstructed paths with random paths of the same form.

The ablation results in Table 2 verify the contribution of each component in PRGCL. Among all variants, *+w/o CL* causes the most severe performance drop, confirming the central role of the self-supervised contrastive. *+w/o CR* and *+w/o PR* also consistently degrade performance, showing that stable path affinity learning and task-relevant path reconstruction are both necessary for preserving useful propagation structures. In addition, the inferior results of *+w/o RR* and *+w/o AR* indicate that the replacement style and addition style reconstructed views provide complementary signals under the coupled reconstruction principle. Finally, the degradation of *+w/o AL* demonstrates that the improvement of PRGCL comes from learning meaningful reconstructed paths.

4.4 Parameter Sensitivity Study

By fixing other parameters, we change τ to a set of predetermined representative values presented in Fig. 3d. As can be observed, with the increase of τ, the performance of PRGCL starts to increase at the beginning and gradually reaches its peak when τ is 0.15 on the Douban-book dataset, 0.2 on the Alibaba-iFashion dataset, and 0.15 on the Tmall dataset. We suppose that a smaller τ will make PRGCL more sensitive to the data. Moreover, the relatively small training set of the Douban-book dataset makes it prone to overfitting, increasing the sensitivity to τ. We conduct a detailed parameter sensitivity study in Appendix A.7.6.

5 Conclusion

In this paper, we propose PRGCL. PRGCL revisits view generation from the perspective of self-supervised signal propagation paths. We identify three limitations of conventional view construction. To address these issues, PRGCL introduces a coupled reconstruction principle that preserves task relevant structural subsets, refines frequency domain signals, and maintains necessary connectivity across reconstructed views. Extensive experiments on three highly sparse real-world datasets show that PRGCL consistently outperforms state-of-the-art graph contrastive recommendation baselines, demonstrating its effectiveness in improving recommendation accuracy.

REFERENCES

Rianne van den Berg, Thomas N Kipf, and Max Welling. Graph convolutional matrix completion. *arXiv preprint arXiv:1706.02263*, 2017.

Xuheng Cai, Chao Huang, Lianghao Xia, and Xubin Ren. Lightgcl: Simple yet effective graph contrastive learning for recommendation. In *The Eleventh International Conference on Learning Representations*.

Lei Chen, Le Wu, Richang Hong, Kun Zhang, and Meng Wang. Revisiting Graph Based Collaborative Filtering: A Linear Residual Graph Convolutional Network Approach. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01):27–34, April 2020a. ISSN 2374-3468, 2159-5399. doi: 10.1609/aaai.v34i01.5330.

Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pp. 1597–1607. PmLR, 2020b.

Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pp. 249–256. JMLR Workshop and Conference Proceedings, 2010.

Jean-Bastien Grill, Florian Strub, Florent Altché, Corentin Tallec, Pierre Richemond, Elena Buchatskaya, Carl Doersch, Bernardo Avila Pires, Zhaohan Guo, Mohammad Gheshlaghi Azar, et al. Bootstrap your own latent-a new approach to self-supervised learning. *Advances in neural information processing systems*, 33:21271–21284, 2020.

Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. *Advances in neural information processing systems*, 30, 2017.

Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9729–9738, 2020a.

Wei He, Guohao Sun, Jinhu Lu, and Xiu Susie Fang. Candidate-aware Graph Contrastive Learning for Recommendation. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1670–1679, Taipei Taiwan, July 2023. ACM. ISBN 978-1-4503-9408-6. doi: 10.1145/3539618.3591647.

Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web*, pp. 173–182, Perth Australia, April 2017. International World Wide Web Conferences Steering Committee. doi: 10.1145/3038912.3052569.

Xiangnan He, Kuan Deng, Xiang Wang, Yan Li, YongDong Zhang, and Meng Wang. LightGCN: Simplifying and Powering Graph Convolution Network for Recommendation. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 639–648, Virtual Event China, July 2020b. ACM. doi: 10.1145/3397271.3401063.

Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations*, 2017.

Yehuda Koren, Steffen Rendle, and Robert Bell. Advances in collaborative filtering. *Recommender systems handbook*, pp. 91–142, 2021.

Matevž Kunaver and Tomaž Požrl. Diversity in recommender systems – A survey. *Knowledge-Based Systems*, 123:154–162, May 2017. ISSN 09507051. doi: 10.1016/j.knosys.2017.02.009.

Haoyang Li, Ziwei Zhang, Xin Wang, and Wenwu Zhu. Learning invariant graph representations for out-of-distribution generalization. *Advances in Neural Information Processing Systems*, 35: 11828–11841, 2022a.

Sihang Li, Xiang Wang, An Zhang, Yingxin Wu, Xiangnan He, and Tat-Seng Chua. Let invariant rationale discovery inspire graph contrastive learning. In *International conference on machine learning*, pp. 13052–13065. PMLR, 2022b.

Yang Li, Kangbo Liu, Ranjan Satapathy, Suhang Wang, and Erik Cambria. Recent Developments in Recommender Systems: A Survey [Review Article]. *IEEE Computational Intelligence Magazine*, 19(2):78–95, May 2024. ISSN 1556-603X, 1556-6048. doi: 10.1109/MCI.2024.3363984.

Jianhua Lin. Divergence measures based on the shannon entropy. *IEEE Transactions on Information theory*, 37(1):145–151, 2002.

Zihan Lin, Changxin Tian, Yupeng Hou, and Wayne Xin Zhao. Improving graph collaborative filtering with neighborhood-enriched contrastive learning. In *Proceedings of the ACM web conference 2022*, pp. 2320–2329, 2022.

Feng Liu, Hao Cang, Huanhuan Yuan, Jiaqing Fan, Yongjing Hao, Fuzhen Zhuang, Guanfeng Liu, and Pengpeng Zhao. How do graph signals affect recommendation: Unveiling the mystery of low and high-frequency graph signals. In *Proceedings of the 32nd ACM SIGKDD Conference on Knowledge Discovery and Data Mining V. 1*, pp. 831–842, 2026.

Seongmin Park, Mincheol Yoon, Jae-woong Lee, Hogun Park, and Jongwuk Lee. Toward a better understanding of loss functions for collaborative filtering. In *Proceedings of the 32nd ACM international conference on information and knowledge management*, pp. 2034–2043, 2023.

Xubin Ren, Lianghao Xia, Jiashu Zhao, Dawei Yin, and Chao Huang. Disentangled Contrastive Collaborative Filtering. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1137–1146, Taipei Taiwan, July 2023. ACM. ISBN 978-1-4503-9408-6. doi: 10.1145/3539618.3591665.

Steffen Rendle, Christoph Freudenthaler, Zeno Gantner, and Lars Schmidt-Thieme. Bpr: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence*, pp. 452–461, 2009.

Thanh Tran, Kyumin Lee, Yiming Liao, and Dongwon Lee. Regularizing Matrix Factorization with User and Item Embeddings for Recommendation. In *Proceedings of the 27th ACM International Conference on Information and Knowledge Management*, pp. 687–696, Torino Italy, October 2018. ACM. ISBN 978-1-4503-6014-2. doi: 10.1145/3269206.3271730.

Chenyang Wang, Yuanqing Yu, Weizhi Ma, Min Zhang, Chong Chen, Yiqun Liu, and Shaoping Ma. Towards Representation Alignment and Uniformity in Collaborative Filtering. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 1816–1825, Washington DC USA, August 2022. ACM. ISBN 978-1-4503-9385-0. doi: 10.1145/3534678.3539253.

Xiang Wang, Xiangnan He, Meng Wang, Fuli Feng, and Tat-Seng Chua. Neural graph collaborative filtering. In *Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval*, pp. 165–174, 2019.

Xiang Wang, Hongye Jin, An Zhang, Xiangnan He, Tong Xu, and Tat-Seng Chua. Disentangled Graph Collaborative Filtering. In *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1001–1010, Virtual Event China, July 2020. ACM. ISBN 978-1-4503-8016-4. doi: 10.1145/3397271.3401137.

Jiancan Wu, Xiang Wang, Fuli Feng, Xiangnan He, Liang Chen, Jianxun Lian, and Xing Xie. Self-supervised graph learning for recommendation. In *Proceedings of the 44th international ACM SIGIR conference on research and development in information retrieval*, pp. 726–735, 2021.

Yihong Wu, Le Zhang, Fengran Mo, Tianyu Zhu, Weizhi Ma, and Jian-Yun Nie. Unifying Graph Convolution and Contrastive Learning in Collaborative Filtering. In *Proceedings of the 30th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pp. 3425–3436, August 2024. doi: 10.1145/3637528.3671840.

Yingxin Wu, Xiang Wang, An Zhang, Xiangnan He, and Tat-Seng Chua. Discovering invariant rationales for graph neural networks. In *International Conference on Learning Representations*.

Zhirong Wu, Yuanjun Xiong, Stella X Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 3733–3742, 2018.

Junliang Yu, Hongzhi Yin, Xin Xia, Tong Chen, Lizhen Cui, and Quoc Viet Hung Nguyen. Are Graph Augmentations Necessary?: Simple Graph Contrastive Learning for Recommendation. In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 1294–1303, Madrid Spain, July 2022. ACM. doi: 10.1145/3477495.3531937.

Junliang Yu, Xin Xia, Tong Chen, Lizhen Cui, Nguyen Quoc Viet Hung, and Hongzhi Yin. Xsimgcl: Towards extremely simple graph contrastive learning for recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 36(2):913–926, 2023.

Chi Zhang, Qilong Han, Qiaoyu Tan, Shengjie Wang, Xiangyu Zhao, and Rui Chen. DimCL: Dimension-Aware Augmentation in Contrastive Learning for Recommendation. In *Proceedings of the 31st ACM SIGKDD Conference on Knowledge Discovery and Data Mining V.1*, pp. 1913–1923, Toronto ON Canada, July 2025a. ACM. ISBN 979-8-4007-1245-6. doi: 10.1145/3690624.3709200.

Yi Zhang and Yiwen Zhang. MixRec: Individual and Collective Mixing Empowers Data Augmentation for Recommender Systems. In *Proceedings of the ACM on Web Conference 2025*, pp. 2198–2208, April 2025. doi: 10.1145/3696410.3714565.

Yi Zhang, Lei Sang, and Yiwen Zhang. 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*, pp. 1253–1262, 2024.

Yi Zhang, Yiwen Zhang, Lei Sang, and Victor S. Sheng. Simplify to the Limit! Embedding-Less Graph Collaborative Filtering for Recommender Systems. *ACM Transactions on Information Systems*, 43(1):1–30, January 2025b. ISSN 1046-8188, 1558-2868. doi: 10.1145/3701230.

Yu Zhang, Yiwen Zhang, Yi Zhang, Lei Sang, and Yun Yang. 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*, pp. 1985–1994, 2025c.

A Technical Appendices and Supplementary Material

A.1 Notation

Table 3: Symbols and descriptions used in the complexity analysis of PRGCL.

Symbol	Description
U	Set of users
I	Set of items
G	Original user-item interaction graph
E	Set of observed interaction edges
d	Embedding dimension
K	Number of GCN layers
$\mathbf{P}$	User frequency domain affinity matrix
$\mathbf{Q}$	Item frequency domain affinity matrix
P_u	Number of user frequency domain affinities
Q_i	Number of item frequency domain affinities
$\bar{\pi}_{u,s}^{(m)}$	Assignment cross side probability of user u to the s-th item affinity assignments
$\bar{\pi}_{i,r}^{(m)}$	Assignment cross side probability of item i to the r-th user affinity assignments
q	Maximum number of retained affinity shortcut edges per node
$\widetilde{\mathcal{G}}^{(1)}, \widetilde{\mathcal{G}}^{(2)}$	Two reconstructed contrastive views

The list of notations in this paper is in Table 3.

A.2 PSEUDOCODE

The pseudocode for PRGCL is shown in Algorithm 1.

Algorithm 1 PRGCL Training Algorithm

Input: User-item interaction graph $\mathcal{G} = (\mathcal{U}, \mathcal{I}, \mathcal{E})$,
Embedding dimension d, number of LightGCN layers K,
Temperature τ, regularization coefficients $\lambda_{reg}, \lambda_{ssl}, \lambda_{car}$,
Augmentation ratios α, β, edge number q,
Learning rate η, total training epochs T.
Output: User embeddings $\mathbf{H}_U$, item embeddings $\mathbf{H}_I$.

```
Initialize H_U^0 and H_I^0 randomly;
Initialize affinity matrices P and Q;
for epoch = 1 to T do
    Compute (H_U^G, H_I^G) via K-layer GCN on G;
    Compute assignment distributions π̄_{u,s}^{(m)}, π̄_{i,r}^{(m)} through Eq. 5 to Eq. 6.
    Construct G̃_1 using replacement strategy with ratio β through Eq. 4 and Eq. 7;
    Construct G̃_2 using addition strategy with ratio α through Eq. 4 and Eq. 7;
    Compute (H_U^1, H_I^1) on G̃_1;
    Compute (H_U^2, H_I^2) on G̃_2;
    for each triplet (u, i, j) do
        Compute loss L = L_bpr + λ_reg L_reg + λ_ssl L_ssl + λ_car L_car through Eq. 8 to Eq. 11.
        Update parameters using optimizer with learning rate η;
    end for
end for
return (H_U^G, H_I^G); =0
```

A.3 PROPOSITION PROOF

The three propositions in Section 3.1 serve to formalize three distinct failure scenarios. In particular, Prop 2 is used only to establish that the contraction rule induces a linear operator on the truncated path complex, while Prop 3 is used only to derive the explicit membership criterion for the relativized path family.

Proposition 1. *Fix $v_a, v_b \in \mathcal{U} \cup \mathcal{I}$, and define $d_{\mathcal{G}}(v_a, v_b) = \min_{p \in \mathcal{P}_K(v_a, v_b)} |p|$. Let the total weight of all shortest admissible paths between v_a and v_b in $\mathcal{G}$ be*

$$M_K^{\mathrm{ess}}(v_a, v_b) = \sum_{\substack{p \in \mathcal{P}_K(v_a, v_b) \\ |p| = d_{\mathcal{G}}(v_a, v_b)}} w(p). \tag{12}$$

If no shortest path in $\mathcal{P}_K(v_a, v_b)$ remains admissible in $\widetilde{\mathcal{G}}$, then

$$\widetilde{\mathcal{S}}_K(v_a, v_b) \le \mathcal{S}_K(v_a, v_b) - M_K^{\mathrm{ess}}(v_a, v_b). \tag{13}$$

Proof. Following the local structural collapse setting in Section 3.1, we consider a perturbative view $\widetilde{\mathcal{G}}$ obtained by removing admissible propagation routes from the original graph without introducing new routes between v_a and v_b. Under this setting,

$$\mathcal{P}_K^{\widetilde{\mathcal{G}}}(v_a, v_b) \subseteq \mathcal{P}_K(v_a, v_b). \tag{14}$$

We then decompose the original admissible path family into the shortest part and the remaining part:

$$\begin{aligned} \mathcal{P}_K^{\mathrm{ess}}(v_a, v_b) &= \{p \in \mathcal{P}_K(v_a, v_b) \mid |p| = d_{\mathcal{G}}(v_a, v_b)\}, \\ \mathcal{P}_K^{\mathrm{rem}}(v_a, v_b) &= \mathcal{P}_K(v_a, v_b) \setminus \mathcal{P}_K^{\mathrm{ess}}(v_a, v_b). \end{aligned} \tag{15}$$

By Definition 2.3, the effective self-supervised signal on $\mathcal{G}$ can be written as

$$\mathcal{S}_K(v_a, v_b) - M_K^{\mathrm{ess}}(v_a, v_b) = \sum_{p \in \mathcal{P}_K^{\mathrm{rem}}(v_a, v_b)} w(p). \tag{16}$$

Now assume that no shortest path in $\mathcal{P}_K(v_a, v_b)$ remains admissible in $\widetilde{\mathcal{G}}$. Then every path in $\mathcal{P}_K^{\widetilde{\mathcal{G}}}(v_a, v_b)$ must belong to $\mathcal{P}_K^{\mathrm{rem}}(v_a, v_b)$, namely $\mathcal{P}_K^{\widetilde{\mathcal{G}}}(v_a, v_b) \subseteq \mathcal{P}_K^{\mathrm{rem}}(v_a, v_b)$. Since each propagation coefficient and frequency domain affinity lies in $(0, 1]$, we have $w(p) \geq 0$ for admissible path p:

$$\widetilde{\mathcal{S}}_K(v_a, v_b) = \sum_{p \in \mathcal{P}_K^{\widetilde{\mathcal{G}}}(v_a, v_b)} w(p) \leq \sum_{p \in \mathcal{P}_K^{\mathrm{rem}}(v_a, v_b)} w(p) = \mathcal{S}_K(v_a, v_b) - M_K^{\mathrm{ess}}(v_a, v_b). \tag{17}$$

□

Proposition 2. *Let $p = (v_0, \ldots, v_a, \ldots, v_b, \ldots, v_l) \in \mathcal{P}_K$ with $0 \leq a < b \leq l$ and $(v_a, v_b) \in \mathcal{E}$. The subpath $q = (v_a, \ldots, v_b)$ denotes the corresponding contiguous segment. The path $p' = (v_0, \ldots, v_a, v_b, \ldots, v_l)$ represents the elementary contraction of q to the chord (v_a, v_b). The operator $T : C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R}) \to C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R})$ is defined on basis paths by replacing each occurrence of q as a contiguous subpath with (v_a, v_b) while all other basis paths remain unchanged.*

Proof. We verify the proposition in two steps. We first show that the contraction rule maps a basis path of the truncated path complex to another admissible basis path, and then extend this rule linearly. Consider a path $p = (v_0, \ldots, v_a, \ldots, v_b, \ldots, v_l) \in \mathcal{P}_K$ that contains the contiguous segment $q = (v_a, \ldots, v_b)$ and satisfies $(v_a, v_b) \in \mathcal{E}$. By replacing the whole segment q with the chord (v_a, v_b), we obtain the contracted path $p' = (v_0, \ldots, v_a, v_b, \ldots, v_l)$. All consecutive pairs of vertices outside the segment q are unchanged, and the newly inserted pair (v_a, v_b) is an edge of $\mathcal{G}$ by assumption. Therefore p' is still an admissible path. Moreover,

$$|p'| = |p| - (b - a - 1) \leq |p| \leq K, \tag{18}$$

so p' remains in $\mathcal{P}_K$. Hence the basis rule

$$T(e_p) = \begin{cases} e_{p'}, & \text{if } p \text{ contains } q \text{ as a contiguous subpath,} \\ e_p, & \text{otherwise} \end{cases} \tag{19}$$

is well defined on basis elements of $C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R})$.

Next consider an arbitrary chain $c = \sum_i \alpha_i e_{p_i}$ with $\alpha_i \in \mathbb{R}$. Since $C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R})$ is the real vector space spanned by the basis paths $\{e_p\}$, the above basis rule extends uniquely by linearity:

$$T(\sum_i \alpha_i e_{p_i}) = \sum_i \alpha_i T(e_{p_i}). \tag{20}$$

Accordingly, $T : C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R}) \to C_*(\mathfrak{P}_K(\mathcal{G}); \mathbb{R})$ is a well-defined linear operator. □

Proposition 3. *A contiguous subpath $(v_a, \ldots, v_b)$ of an admissible path is called strictly dominated if $p_{v_a v_b} \kappa_{v_a v_b} > \prod_{j=a+1}^{b} p_{v_{j-1} v_j} \kappa_{v_{j-1} v_j}$. A path $p = (v_0, \ldots, v_l)$ is called ρ-*tailed *if there exists $m \in \{0, \ldots, l-1\}$ such that $\kappa_{v_{r-1} v_r} \leq \rho$ for every $r \in \{m+1, \ldots, l\}$ and $\rho \in (0, 1)$. The set $\mathcal{P}_K^{\mathrm{rel}}(\mathcal{G}) \subseteq \mathcal{P}_K(\mathcal{G})$ consists of all admissible paths that contain no strictly dominated contiguous subpath and are not ρ-tailed.*

This proposition introduces the filtering rule that defines the relativized path family. To write this rule explicitly, let

$$\begin{aligned} \mathcal{D}_K(\mathcal{G}) &= \{\, p \in \mathcal{P}_K(\mathcal{G}) \mid p \text{ contains a strictly dominated contiguous subpath} \,\}, \\ \mathcal{T}_{K,\rho}(\mathcal{G}) &= \{\, p \in \mathcal{P}_K(\mathcal{G}) \mid p \text{ is } \rho\text{-tailed} \,\}. \end{aligned} \tag{21}$$

Then the set described in the proposition is precisely obtained by removing these two excluded classes from the full admissible path family:

$$\mathcal{P}_K^{\mathrm{rel}}(\mathcal{G}) = \mathcal{P}_K(\mathcal{G}) \setminus (\mathcal{D}_K(\mathcal{G}) \cup \mathcal{T}_{K,\rho}(\mathcal{G})). \tag{22}$$

Therefore, for any admissible path $p \in \mathcal{P}_K(\mathcal{G})$, the membership relation $p \in \mathcal{P}_K^{\mathrm{rel}}(\mathcal{G})$ holds if and only if the following two conditions are simultaneously satisfied: p contains no strictly dominated contiguous subpath and p is not ρ-tailed.

A.4 Model Analysis

A.4.1 Propositions and Coupled Reconstruction Principles in PRGCL

PRGCL instantiates the coupled reconstruction principle in Eq. 3 and starts from the original interaction graph, introduces frequency domain shortcuts as auxiliary propagation paths, and finally optimizes two reconstructed views through multi-task joint training.

Applying propositions in PRGCL. Section 3.1 provides the motivation for PRGCL through three propositions, which are further instantiated as core signal path structure encoding, frequency domain shortcut propagation, and contrastive view reconstruction. For local path cluster integrity, PRGCL adopts an asymmetric backbone construction,$\mathcal{E}_{\text{core}}^{(1)} = \mathcal{E} \setminus \mathcal{R}^{(1)}$ and$\mathcal{E}_{\text{core}}^{(2)} = \mathcal{E}$, so that one view provides a contrastive perturbation while the other preserves the full interaction backbone as a stable structural reference. Thus, the shortest admissible propagation paths supporting the effective signal field in Prop. 1 remain anchored to the original topology. For frequency domain signal interference, PRGCL builds shortcut edges computed from the original graph and retains only core frequency domain signals $\Gamma_u^{(m)} = \sum_{s \in \text{Top}_q(\bar{\pi}_{u,:}^{(m)})} \bar{\pi}_{u,s}^{(m)}$ and $\Gamma_i^{(m)} = \sum_{r \in \text{Top}_q(\bar{\pi}_{i,:}^{(m)})} \bar{\pi}_{i,r}^{(m)}$. This concentrates propagation on reliable shortcuts, compressing weak high order paths into compact frequency domain shortcuts, in line with the contraction intuition in Prop. 2. For contrastive view reconstruction and multi-task joint training, PRGCL inputs the two reconstructed views into a shared encoder and aligns their node representations with a contrastive objective, which couples the two propagation paths and encourages structurally stable paths, aligned with Prop. 3.

Instantiating the coupled reconstruction principle in PRGCL. The abstract criterion in Eq. 3 is instantiated in PRGCL by the following reconstruction stages:

$$(h_U, h_I) = \mathcal{F}_\Theta(\mathcal{G}), \qquad \{\pi_{u,r}, \pi_{i,s}\} \to \{\bar{\pi}_{u,s}^{(m)}, \bar{\pi}_{i,r}^{(m)}\} \to \widetilde{\mathcal{G}}^{(m)} \to (h_U^{(m)}, h_I^{(m)}), \quad m \in \{1, 2\}. \tag{23}$$

Specifically, PRGCL first derives base representations on the original graph and obtains information from different frequency domains via Eq. 5. We then aggregate the assignments of opposite endpoints over sampled anchors according to Eq. 6, producing the shortcut scores $\bar{\pi}_{u,s}^{(m)}$ and $\bar{\pi}_{i,r}^{(m)}$. Finally, the top-q frequency domain shortcuts are materialized as weighted shortcut edges by Eq. 7 and combined with the retained core backbone in Eq. 4, yielding the reconstructed view $\widetilde{\mathcal{G}}^{(m)}$. Under this implementation, the discrepancy term $\mathcal{D}$ is controlled by comparing the core path differences between the constructed view and the original view, the contraction term $\mathcal{C}$ is realized by constructing frequency-domain shortcuts to optimize the propagation of self-supervised signals, and the incompatibility term $\mathcal{I}$ is optimized by aligning the two reconstructed views with a shared encoder and contrastive learning. Therefore, Eq. 3 stabilizes local propagation via the core skeleton, compresses multi-order frequency signals through shortcut edges, and aligns the two contrastive views toward a compatible propagation pattern, instantiating the PRGCL method.

A.4.2 Complexity Analysis

In this section, we analyze the computational complexity of PRGCL, focusing on its main components: the GCN backbone, core view construction, affinity shortcut augmentation, and cross view contrastive optimization.

Complexity of GCN Backbone PRGCL adopts GCN to compute base user and item representations on the sparse user-item interaction graph. Since the interaction graph is sparse, each propagation layer is implemented by sparse matrix multiplication, followed by layer-wise aggregation over the propagated embeddings.

Let the embedding matrix at each layer have dimension $(|\mathcal{U}| + |\mathcal{I}|) \times d$. For one LightGCN layer, the cost per layer is approximately:

$$\mathcal{O}(|\mathcal{E}|d). \tag{24}$$

After K propagation layers, the total sparse propagation cost is:

$$\mathcal{O}(K|\mathcal{E}|d). \tag{25}$$

PRGCL further averages the K propagated embeddings for all users and items. The layer aggregation cost is:

$$\mathcal{O}(K(|\mathcal{U}| + |\mathcal{I}|)d). \tag{26}$$

Thus, the total complexity is:

$$\mathcal{O}(K|\mathcal{E}|d + K(|\mathcal{U}| + |\mathcal{I}|)d). \tag{27}$$

Complexity of Core View Construction The core view construction stage preserves the minimal user-item interaction backbone for the two reconstructed views. To avoid the prohibitive cost caused by the rapid growth of $|\mathcal{P}_K|$ with graph size and depth, PRGCL adopts randomized interaction sampling to determine the retained core edges and anchor interactions. Specifically, for the two reconstructed views, PRGCL samples two index sets of size βL for the replacement view and two index sets of size αL for the addition view. Therefore, the cost of sampling these index sets is approximately:

$$\mathcal{O}(2(\alpha + \beta)L). \tag{28}$$

For the replacement view, forming $\mathcal{R}^{(1)}$ and applying the corresponding edge mask over the interaction list require a linear pass over the sampled indices and the interaction backbone. The cost is $\mathcal{O}(L)$. Constructing $\mathcal{E}^{(1)}_{\text{core}}$ and $\mathcal{E}^{(2)}_{\text{core}}$ involves retaining the masked edge list for the replacement view and the full edge list for the addition view. Considering both views, the total complexity is approximately $\mathcal{O}((2 + \alpha + \beta)L) = \mathcal{O}(L)$.

Hence, the overall complexity of this module is linear in the number of observed interactions. In PRGCL, the reconstruction cost is governed by interaction-level sampling and sparse edge retention, without requiring explicit enumeration over the path space.

Complexity of Affinity Shortcut Augmentation The affinity shortcut augmentation stage assigns users and items to learnable affinities, aggregates cross side affinity scores over sampled anchors, and retains only top-q shortcuts for each selected node. As a result, shortcut construction is confined to sampled interactions, with only a limited number of shortcuts retained for each selected node, avoiding dense affinity connectivity across all node pairs.

First, PRGCL computes affinity assignment scores between user embeddings and $\mathbf{P} \in \mathbb{R}^{P_u \times d}$, and between item embeddings and $\mathbf{Q} \in \mathbb{R}^{Q_i \times d}$. The cost of score computation is:

$$\mathcal{O}(|\mathcal{U}|P_u d + |\mathcal{I}|Q_i d). \tag{29}$$

Next, for sampled anchor interactions, PRGCL aggregates opposite-endpoint assignment scores. For the replacement and addition views together, considering all sampled interactions, the total complexity is approximately:

$$\mathcal{O}((\alpha + \beta)L(P_u + Q_i)). \tag{30}$$

Let $\mathcal{U}^{(m)}_{\text{s}}$ and $\mathcal{I}^{(m)}_{\text{s}}$ denote the selected users and items incident to the sampled anchors in view m. Retaining top-q shortcuts for each selected node requires ranking over Q_i or P_u affinity scores. With a sorting-based top-q operator, the cost is:

$$\mathcal{O}(\sum_{m=1}^{2}(|\mathcal{U}^{(m)}_{\text{s}}|Q_i \log Q_i + |\mathcal{I}^{(m)}_{\text{s}}|P_u \log P_u)). \tag{31}$$

After top-q retention, the number of auxiliary shortcuts is bounded by q times the number of selected nodes. Thus, the cost of constructing $\mathcal{E}^{(m)}_{\text{aux}}$ is:

$$\mathcal{O}(q\sum_{m=1}^{2}(|\mathcal{U}^{(m)}_{\text{s}}| + |\mathcal{I}^{(m)}_{\text{s}}|)) \leq \mathcal{O}(q(\alpha + \beta)L). \tag{32}$$

If the normalized adjacency of the weighted reconstructed views is rebuilt after shortcut insertion, the cost is linear in the number of retained core edges, auxiliary shortcuts, and involved nodes. Therefore, the cost of rebuilding the normalized augmented views is:

$$\mathcal{O}(2|\mathcal{E}| + q(\alpha + \beta)L + |\mathcal{U}| + |\mathcal{I}| + P_u + Q_i). \tag{33}$$

Combining these terms, the overall complexity of this module is:

$$\mathcal{O}(|\mathcal{U}|P_u d + |\mathcal{I}|Q_i d + (\alpha+\beta)L(P_u + Q_i) + \sum_{m=1}^{2}(|\mathcal{U}^{(m)}_{\text{s}}|Q_i \log Q_i + |\mathcal{I}^{(m)}_{\text{s}}|P_u \log P_u) + q(\alpha+\beta)L + 2|\mathcal{E}|). \tag{34}$$

Complexity of Cross View Contrastive Alignment and Joint Optimization This stage evaluates the original graph and the two reconstructed views, and then optimizes the contrastive, ranking, and affinity matrix regularization objectives jointly. The main cost arises from sparse propagation on three graph instances and similarity computation in the contrastive losses. Let $|\widetilde{\mathcal{E}}^{(m)}|$ denote the number of edges in the m-th reconstructed view after adding affinity shortcuts. Evaluating the backbone on the original graph and the two reconstructed views requires K layers of sparse propagation on each graph. Thus, the propagation cost per training step is approximately:

$$\mathcal{O}(K(|\mathcal{E}| + |\widetilde{\mathcal{E}}^{(1)}| + |\widetilde{\mathcal{E}}^{(2)}|)d). \tag{35}$$

For the contrastive objective, each similarity between two d-dimensional embeddings costs $\mathcal{O}(d)$. Under contrastive learning with batch $|\mathcal{B}|$, the matrix multiplication cost of $\mathcal{O}(|\mathcal{B}|^2 d)$. Since PRGCL uses three types of contrastive terms, the total complexity is:

$$\mathcal{O}(3|\mathcal{B}|^2 d). \tag{36}$$

The affinity regularization over $\mathbf{P}$ and $\mathbf{Q}$ applies softmax and entropy evaluation:

$$\mathcal{O}((P_u + Q_i)d). \tag{37}$$

The BPR loss over sampled triplets requires two inner products per triplet. Therefore, the BPR optimization cost per batch is:

$$\mathcal{O}(2|\mathcal{B}|d). \tag{38}$$

Thus, the total optimization cost per training step is:

$$\mathcal{O}(K(|\mathcal{E}| + |\widetilde{\mathcal{E}}^{(1)}| + |\widetilde{\mathcal{E}}^{(2)}|)d + 3|\mathcal{B}|^2 d + 2|\mathcal{B}|d + (P_u + Q_i)d). \tag{39}$$

Summary Overall, the dominant computational cost of PRGCL arises from three main components: sparse GCN propagation, affinity assignment with respect to the learnable affinity matrices during reconstructed-view refresh, and contrastive similarity computation. Given typical sparse user-item graphs, q, α, β, P_u, and Q_i are small hyperparameters, and the contrastive objective is evaluated in minibatches, these costs remain well controlled. Therefore, PRGCL preserves efficient sparse training dynamics on large-scale user-item graphs while avoiding the prohibitive cost associated with explicit path enumeration.

A.5 Discussion

A.6 Path Reconstruction Interpretation of the Contrastive Objective

We further show why the representation contrastive objective has a tractable upper bound connection with view path consistency. Let $\widetilde{\mathbf{A}}^{(m)}$ be the normalized weighted adjacency matrix of the m-th reconstructed view, and let $\mathbf{e}_t$ denote the trainable initial embedding of node t. For a K layer GCN encoder, the representation of node v in view m can be written as

$$h_v^{(m)} = \frac{1}{K} \sum_{k=1}^{K} \sum_{t \in V} \left[(\widetilde{\mathbf{A}}^{(m)})^k \right]_{vt} \mathbf{e}_t. \tag{40}$$

Since $\left[(\widetilde{\mathbf{A}}^{(m)})^k \right]_{vt}$ equals the accumulated normalized transmission weight of all walks with length k from v to t, we define the path signal coefficient of view m as

$$\Phi_K^{(m)}(v,t) = \frac{1}{K} \sum_{k=1}^{K} \left[(\widetilde{\mathbf{A}}^{(m)})^k \right]_{vt}. \tag{41}$$

Then the representation discrepancy across two views can be decomposed as

$$h_v^{(1)} - h_v^{(2)} = \sum_{t \in V} \left(\Phi_K^{(1)}(v,t) - \Phi_K^{(2)}(v,t) \right) \mathbf{e}_t. \tag{42}$$

If the initial embeddings are bounded by $\|\mathbf{e}_t\|_2 \leq B$, the representation discrepancy satisfies the following upper bound:

$$\|h_v^{(1)} - h_v^{(2)}\|_2 \leq B \sum_{t \in V} \left| \Phi_K^{(1)}(v,t) - \Phi_K^{(2)}(v,t) \right|. \tag{43}$$

Therefore, the mismatch between the two path signal fields provides an upper bound for the discrepancy between their induced node representations. This shows that representation alignment is mathematically connected to the control of multihop propagation inconsistency between the two reconstructed views.

Let $z_v^{(m)} = h_v^{(m)}/\|h_v^{(m)}\|_2$ be the normalized representation. For the positive pair term in InfoNCE (Chen et al., 2020b), we have

$$-\mathrm{sim}(z_v^{(1)}, z_v^{(2)}) = \frac{1}{2}\|z_v^{(1)} - z_v^{(2)}\|_2^2 - 1. \tag{44}$$

Thus, up to the denominator term involving negative samples, minimizing the positive contrastive term directly minimizes the squared distance between the normalized embeddings. When the representation norms are bounded away from zero, this normalized discrepancy is controlled by the discrepancy between the corresponding unnormalized representations. Combining this identity with Eq. 43, the contrastive objective minimizes a representation discrepancy whose upper bound is determined by the mismatch between the path induced signal fields of $\widetilde{\mathcal{G}}^{(1)}$ and $\widetilde{\mathcal{G}}^{(2)}$. This explains why PRGCL can align the replacement and addition views at the representation level and avoid explicit enumeration of $\mathcal{P}_K^{\mathrm{rel}}(\cdot)$ while preserving a tractable upper bound connection with multihop path consistency.

A.7 Experiment Study

A.7.1 Experimental Settings

Table 4: Statistics of the datasets.

Datasets	#Users	#Items	#Interactions	Density
Douban-book	13024	22347	792062	0.27%
Tmall	47939	41390	2619389	0.13%
Alibaba-iFashion	300000	81614	1607813	0.01%

The statistics of the datasets are summarized in Table. 4. To ensure a fair comparison, all experiments are conducted using a unified framework. PyTorch[1] and Xavier Initialization (Glorot & Bengio, 2010) are applied to all models, and the embedding size is fixed at 64. For GCN-based methods, the number of layers is set to 3. All models are trained for up to 500 epochs with an early stopping step of 20 on a Linux server with two Intel(R) Xeon(R) Platinum 8481C CPUs and an NVIDIA RTX 4090 GPU. For DimCL, which used excessive GPU memory on the Alibaba-iFashion dataset, we conducted the experiment on an RTX A6000 GPU. The batch size is set to 4,096 for all datasets and DimCL on Alibaba-iFashion dataset is set to 1024. The temperature coefficient τ is set within $\{0.1, 0.15, 0.2, 0.25, 0.3\}$. We compare PRGCL with following methods:

- **LightGCN** (He et al., 2020b) is a classic graph convolution model in recommender systems.
- **DirectAU** (Wang et al., 2022) is a method that directly uses uniformity loss and alignment loss.
- **SimGCF** (Liu et al., 2026) uses a frequency signal scaler to adjust the graph signal filtering function, fine-tuning the smoothness between user–item pairs in a GNN model.
- **SimGCL** (Yu et al., 2022) discards graph augmentations and adds noise to the embedding space to create contrastive views.
- **CGCL** (He et al., 2023) explores the relationship between the user and the candidate item in the embedding at different layers and uses similar semantic embeddings to construct contrastive pairs.
- **DimCL** (Zhang et al., 2025a) is a dimension-aware augmentation approach for contrastive learning in recommendation, designed based on robust and curriculum learning.
- **MixRec** (Zhang & Zhang, 2025) is a dual mixing-based recommendation framework that empowers data augmentation.

[1] https://pytorch.org

Table 5: Path Diagnostics Study. (a) Shortest-path survival and recovery under different reconstructed views. 'Rep.' and 'Add.' denote the replacement and addition views, respectively. (b) Recommendation performance of shortest-path-aware variants. R/N@10 and R/N@20 denote Recall/NDCG at 10 and 20, respectively.

(a) Path level diagnostics.

Dataset	Diagnostic	Value
Douban-book	Pairs (Critical / Deleted)	157,695 / 365
	Strict Survival	99.77%
	Recovery (Rep. / Add.)	99.92% / 100.00%
	Shortcut Recovery	65.48%
Tmall	Pairs (Critical / Deleted)	307,577 / 65,015
	Strict Survival	78.86%
	Recovery (Rep. / Add.)	99.70% / 100.00%
	Shortcut Recovery	98.60%
iFashion	Pairs (Critical / Deleted)	198,069 / 139,068
	Strict Survival	29.79%
	Recovery (Rep. / Add.)	90.67% / 100.00%
	Shortcut Recovery	86.71%

(b) Three hop shortest path view ablation.

Dataset	Model	R/N@10	R/N@20
Douban-book	PRGCL	0.1247/ 0.1721	0.1811/ 0.1757
	Both	0.1232/ 0.1681	0.1807/ 0.1729
	Addition	0.1241/ 0.1686	0.1808/ 0.1730
	Replacement	0.0974/ 0.1157	0.1414/ 0.1242
Tmall	PRGCL	0.0560/ 0.0537	0.0890/ 0.0677
	Both	0.0553/ 0.0517	0.0843/ 0.0649
	Addition	0.0548/ 0.0525	0.0865/ 0.0655
	Replacement	0.0430/ 0.0409	0.0635/ 0.0497
iFashion	PRGCL	0.0837/ 0.0484	0.1223/ 0.0586
	Both	0.0806/ 0.0471	0.1206/ 0.0533
	Addition	0.0761/ 0.0434	0.1118/ 0.0528
	Replacement	0.0769/ 0.0415	0.1152/ 0.0517

A.7.2 Path Diagnostics Study

For the path level diagnostic, we select user-item pairs within three hops of positive interactions. Critical and Deleted Pairs denote the three hop user item pairs identified before perturbation and those disrupted by edge removal, Strict Survival measures the proportion of critical pairs that remain connected in the replacement skeleton, Recovery measures the proportion recovered in the replacement and addition views, and Shortcut Recovery measures the proportion of deleted critical pairs reconnected through affinity shortcuts. For the three hop shortest path view ablation, 'Both' replaces both reconstructed views with fixed three hop shortest path views, while 'Addition' and 'Replacement' replace only the addition or replacement view, respectively.

As shown in Table 5, structural perturbation substantially disrupts shortest aggregation paths, while the reconstructed views recover most of the affected connectivity and the affinity shortcuts further restore disrupted critical relations. Consistent with this path level recovery, replacing the learned views with fixed shortest path views consistently reduces recommendation performance, with the largest degradation occurring when the learnable replacement view is removed. These results show that PRGCL preserves reliable paggregation structures through the reconstructed views while using the learnable replacement view and affinity shortcuts to recover disrupted connectivity.

A.7.3 Case Study

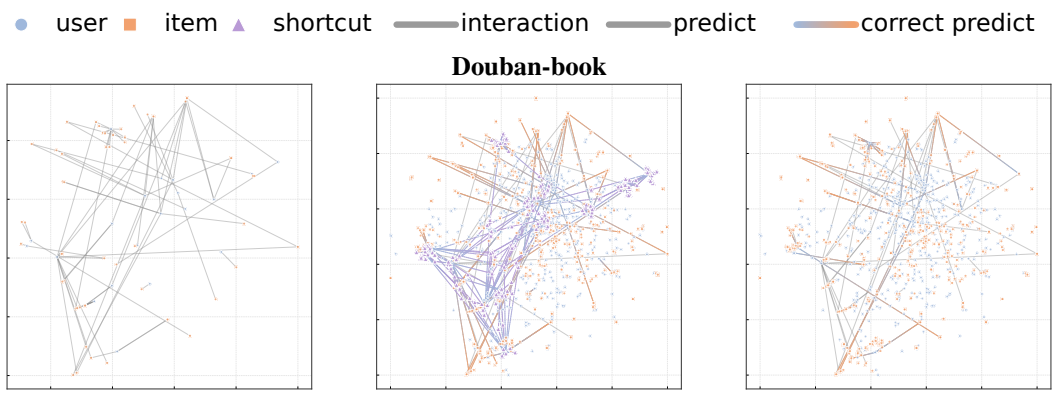

(a) Original (b) PRGCL (c) SimGCL

Figure 4: Case visualization of PRGCL and SimGCL on Douban-book dataset.

In the case study, we first select at most 20 target user-item interactions as examples. Around these interactions, we further expand the local subgraph by incorporating the Top-K recommendation results of PRGCL and SimGCL, together with the neighboring nodes from the original interaction graph, while keeping the total number of nodes within 1,000. The two-dimensional coordinates in the visualization are obtained by t-SNE. Specifically, we first extract the learned embeddings from the trained models and then perform dimensionality reduction on the selected local nodes. Therefore, points that are closer in the figure indicate stronger semantic similarity or collaborative relevance in the learned embedding space. The label *predict* denotes a recommendation edge generated by the model in the user's Top-K recommendation list, whereas *correct predict* indicates that the recommended item is also a ground-truth interacted item for the corresponding user in the evaluation split, namely a successful hit.

As shown in the figure, the correct predictions made by SimGCL are mainly concentrated around users and items with high frequency interactions, while it shows limited capability in accurately predicting users and items with sparse interactions. In contrast, PRGCL can better identify sparse user-item interactions by leveraging affinity shortcuts to disentangle propagation signals of different orders from the paths associated with high-frequency users and items. This observation further demonstrates the effectiveness of the proposed coupled reconstruction principle and the affinity shortcut mechanism.

A.7.4 Sparsity Study

To further evaluate whether PRGCL can alleviate the data sparsity issue, we divide users into four equally sized groups according to the number of user–item interactions: *normal*, *frequent*, *occasional*, and *sparse*. We then report the NDCG@20 score for each group. As shown in Figure 5, PRGCL already achieves strong performance on the *normal* group, indicating that it can learn effective user representations even when user interactions are relatively sufficient. On the Douban-book and Tmall datasets, the performance degradation of PRGCL from the *normal* group to the *sparse* group is relatively mild. This suggests that sparse users in these datasets still contain informative latent collaborative signals, which can be effectively exploited by PRGCL. On the Alibaba-iFashion dataset, the performance of PRGCL first decreases and then increases, whereas LightGCN exhibits a consistent downward trend. This indicates that PRGCL is more capable of identifying medium- and low-frequency users whose preferences are more diverse and difficult to model. Moreover, its performance recovery on sparse users further demonstrates its ability to capture useful collaborative signals from truly low-interaction users.

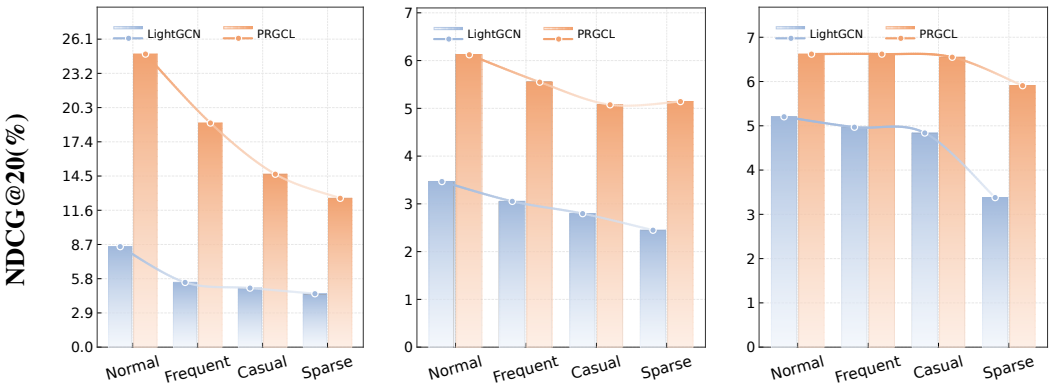

(a) Douban-book (b) Alibaba-iFashion (c) Tmall

Figure 5: Performance comparison of PRGCL with other model under different sparsity levels.

A.7.5 Robustness Study

To investigate the robustness of PRGCL under different levels of noise interference, we injected noisy user-item interactions into the training data at ratios of 10%, 20%, and 30%. As shown in Fig. 6, increasing the noise ratio from 0 to 0.3 leads to a decline in Recall@20 for both SimGCL and

PRGCL, indicating that noisy interactions weaken the ability of recommendation models to capture genuine user preferences. Nevertheless, compared with SimGCL, PRGCL consistently achieves higher Recall@20 on all three datasets and exhibits a more gradual performance degradation, with particularly clear advantages on Tmall and iFashion. Interestingly, we discover that the performance of PRGCL shows a slight increase rather than a further decline on the Douban-book dataset when the noise ratio reaches 0.2 and the same in the Alibaba-iFashion dataset when the noise ratio reaches 0.1. We suppose that the proposed different order frequency domain affinity shortcut can more effectively identify and separate noise signals under stronger perturbations. This suggests that PRGCL is more resistant to noisy interactions and can preserve informative user-item relations more stably when the training data are disturbed. These results further validate the effectiveness of PRGCL in robust recommendation scenarios.

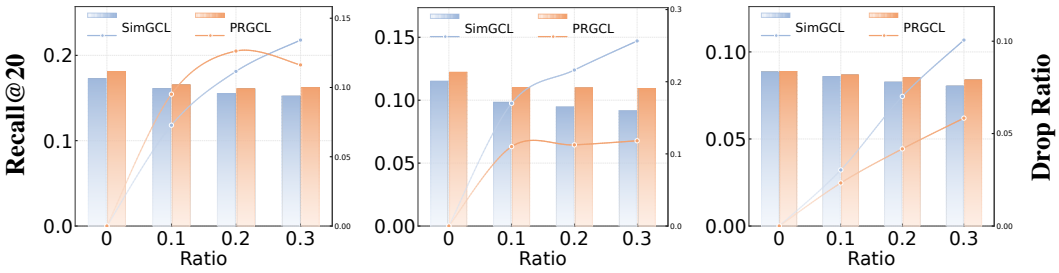

(a) Douban-book (b) Alibaba-iFashion (c) Tmall

Figure 6: Robustness comparison on the three datasets. The bar chart represents the Recall@20 values under different noise injections, and the line chart represents the rate of performance degradation.

A.7.6 Parameter Sensitivity Study

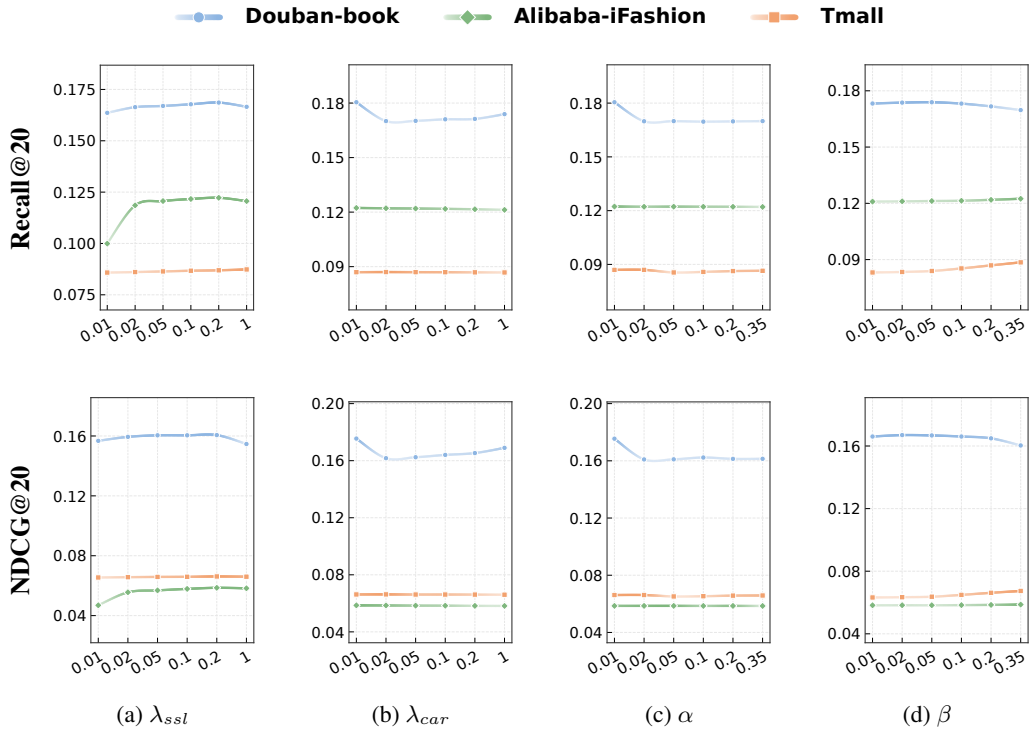

(a) λ_{ssl} (b) λ_{car} (c) α (d) β

Figure 7: Parameter sensitivity study of PRGCL on three datasets.

As shown in Fig. 7, we tune one hyperparameter at a time by assigning it a set of predefined representative values while keeping the others fixed. The results show that the performance of PRGCL initially improves as λ_{ssl} increases, reaching its optimum at $\lambda_{ssl} = 0.2$ on the Douban-book and Alibaba-iFashion datasets and at $\lambda_{ssl} = 1$ on the Tmall dataset. After that, the performance starts to decline. For λ_{car}, PRGCL achieves the best performance on all three datasets when $\lambda_{car} = 0.01$. We further observe that larger values of α and β lead to stronger view reconstruction, which may be detrimental to preserving signals from different frequency domains. Therefore, PRGCL achieves the best performance on the three datasets when $\alpha = 0.01$, and on the Douban-book dataset when $\beta = 0.01$. Interestingly, when $\beta = 0.35$, PRGCL obtains better performance on the Tmall and Alibaba-iFashion datasets. A possible explanation is that these two datasets contain more interaction edges, and thus PRGCL requires stronger reconstruction to learn better signal propagation paths.