GENET: Unleashing the Power of Side Information for Recommendation via Hypergraph Pre-training

Yang Li[1], Qi'ao Zhao[2], Chen Lin[2(✉)], Zhenjie Zhang[3], and Xiaomin Zhu[4]

[1] Institute of Artificial Intelligence,Xiamen University, Xiamen, China
[2] School of Informatics,Xiamen University, Xiamen, China
[3] Neuron Mobility Pte.Ltd, Singapore
[4] Strategic Assessments and Consultation Institute,Academy of Military Science, Beijing, China

Abstract. Recommendation with side information has drawn significant research interest due to its potential to mitigate user feedback sparsity. However, existing models struggle with generalization across diverse domains and types of side information. In particular, three challenges have not been addressed, and they are (1) the diverse formats of side information, including text sequences. (2) The diverse semantics of side information that describes items and users from multi-level in a context different from recommendation systems. (3) The diverse correlations in side information to measure similarity over multiple objects beyond pairwise relations. In this paper, we introduce GENET (Generalized hypErgraph pretraiNing on sidE informaTion), that pre-trains user and item representations on feedback-irrelevant side information and fine-tunes the representations on user feedback data. GENET leverages pre-training as a means to prevent side information from overshadowing critical ID features and feedback signals. It employs a hypergraph framework to accommodate various types of diverse side information. During pre-training, GENET integrates tasks for hyperlink prediction and self-supervised contrast to capture fine-grained semantics at both local and global levels. Additionally, it introduces a unique strategy to enhance pre-training robustness by perturbing positive samples while maintaining high-order relations. Extensive experiments demonstrate that GENET exhibits strong generalization capabilities, outperforming the SOTA method by up to 38% in TOP-N recommendation and sequential recommendation tasks on various datasets with different side information.

Keywords: Side information · Hypergraph pre-training · TOP-N recommendation · Sequential recommendation
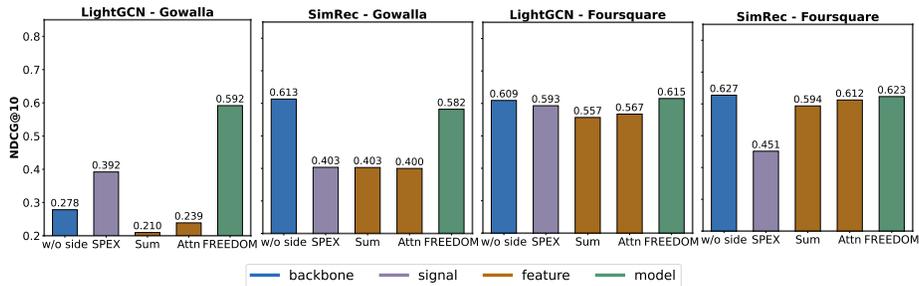
Fig. 1. NDCG@10 on the Gowalla and FourSquare dataset by incorporating side information at feature, model, and signal level upon backbone models LightGCN and SimRec

1 INTRODUCTION

A major and ongoing thrust of research on Recommendation Systems (RSs) is to leverage side information. Side information, a.k.a, auxiliary data, which refers to non-feedback data associated with users and items [6, 17, 23], is available in various forms. For example, most RSs provide rich demographic profiles of users [17, 42], e.g., gender, location, etc. Numerous item features [6, 35, 43] can be accumulated on an E-commerce RS, including the item's category, brand, price, textual descriptions, display images, and so on. Users of RSs also generate a wealth of data on item reviews [21, 41]. Due to the sparsity of user feedback data, side information has shown great potential in enhancing recommendation performance by complementing user feedback [17, 22], especially on cold-start users and items [22]. Furthermore, since side information (e.g., product reviews) explains detailed reasons why a user favors an item, it has been extensively exploited to develop explainable RSs [24, 29].

Existing methods of utilizing side information can be roughly classified into three types: (1) extract and combine the *feature* from side information by direct addition [28] or an attention mechanism [6, 35]; (2) utilize an additional *model*, e.g., a graph [43] or a hypergraph [38], to capture the relationships between users and items in the side information; (3) incorporate side information as a source of supervision *signal*, e.g., multi-task learning [15].

However, *existing methods can not generalize well to different domains and/or side information*. To demonstrate this, we select four recent studies that incorporate side information at feature (i.e., directly summing features [28] and combining features with attention [40]), model (i.e., FREEDOM[5] [43]), and signal (i.e., SPEX [15]) level, and apply them on two backbones, including the commonly adopted LightGCN [11] and the state-of-the-art SimRec [33]. Then, we implement these methods to perform POI recommendations by utilizing social network information on the Gowalla [37] and FourSquare [7] datasets. As shown

---

[5] We create an item-item graph for FREEDOM based on the similarity of POI, and we adapted the FREEDOM implementation to the SimRec framework.

in Figure 1, the inclusion of various side information in most cases reduces the performance of the backbone model. For example, FREEDOM was designed to utilize item textual and visual descriptions and had gained a 15.65% increase over LightGCN on Amazon Sports dataset as per [43]. But as shown in Figure 1, FREEDOM still can't match SimRec without side information, and even using SimRec as its backbone, FREEDOM is still inferior to SimRec itself. SPEX was designed to utilize social network information to improve social recommendation and had gained a 13.48% increase on the Weibo dataset as reported in [15]. Unfortunately, it has led to a significant performance decrease in backbone SimRec. Furthermore, existing methods with side information are inferior to SimRec, which is based solely on feedback data without side information. Therefore, *it is critical to develop a generalized model that can utilize various side information to enhance recommendation performance across domains.*

We argue that the fundamental reason for the malfunction of existing methods is that the side information directly participates in the learning of recommendation. The volume of side information is larger than the number of users/items and the user feedback data. Thus, treating side information as features can overshadow the crucial ID features. For methods that model side information separately, the imbalance of instances results in poorer performance of the recommendation model. Furthermore, some side information, such as the users' social network, is indirectly involved in the recommendation. It is risky to adopt side information as supervision signals because the noisy signals can cause negative transfer in multitasking.

Recently, pre-training-fine-tuning has achieved great success and has become a standard paradigm in many areas, e.g., Natural Language Processing [2]. The paradigm first learns representations from raw data in a self-supervision manner and then fine-tunes the representations on the downstream tasks. The goal is to benefit target tasks with knowledge acquired in pre-training and mitigate negative transfer.

In this paper, we propose to learn user- and item- representations by *pre-training on the side information* and *fine-tuning on the feedback data*. We focus on solving three challenges with regard to the generalization of pre-training on side information.

C1: diversity of format. Common side information includes text sequences (e.g., product reviews), numerical features (e.g., Longitude and Latitude of a POI), categorical features (e.g., item brand), and relation data (e.g., social networks). Previous pre-training methods are based on either sequence models [2] or graph neural networks [12], and they are natural for modeling one format but sub-optimal for others. How to develop a unified pertaining framework to accommodate heterogeneous side information?

C2: diversity of semantics. Side information describes users and items from multi-levels in a context different from RSs. For example, product reviews provide fine-grained user sentiment on specific items, while social networks express more coarse-grained user preferences (e.g., linked friends may have similar tastes). Prior studies often tailor pre-training tasks to different side information,

e.g., link prediction task in the social network learns node semantics from a global perspective of user collaboration [5], contrastive learning tasks with domain-specific data augmentations [3] learns item semantics from a local perspective of textual descriptions, and these pre-training tasks are not robust across domains. How to design the pre-training tasks to integrate complementary information from multiple levels to adapt to diverse side information?

C3: diversity of correlation. Side information measures similarity/dissimilarity over pairs, triples, or more objects, and these multi-faceted and/or high-order relationships are important. For example, a user's preference is affected by his/her social circle connecting multiple friends (i.e., multi-faceted) and his/her friend's friend (i.e., high-order). Former study emphasizes pair-wise relationships [16] and overlooks the high-order and multi-faceted relationship. How to derive a pre-training model that captures complex high-order and multi-faceted relationships as well as the pairwise relationship?

In this paper, we propose `GENET`, meaning `Generalized hypErgraph pretraiNing on sidE informaTion`. To address **C1**, `GENET` is based on a hypergraph, which is suitable to represent heterogeneous side information. To address **C2**, `GENET` presents three pre-training tasks, i.e., hyperlink prediction, global contrastive, and local contrastive, to reveal semantic relationships among users/items at different levels and combine fine-grained and coarse-grained representations. To address **C3**, `GENET` presents a novel strategy to corrupt the positive sample in the hyperlink prediction task, to increase the robustness of pertaining tasks while preserving the high-order and multi-faceted relation.

To verify the generalization of `GENET`, we perform extensive experiments on two recommendation tasks, i.e., TOP-N recommendation and sequential recommendation, on three different datasets with various side information, including social networks, POI locations, product reviews, and item brands. We compare `GENET` with 25 existing methods. The experimental results demonstrate that `GENET` can improve the SOTA SimRec by $5\% - 15\%$ in terms of NDCG@10 on TopN recommendations, $3\% - 38\%$ on sequential recommendations. `GENET` is significantly powerful in cold-start recommendations and increases NDCG@10 of cold-start users by at least 175% on all datasets.

In summary, the main contributions are three-fold. (1) A unified model `GENET` is presented to unleash the power of side information and enhance recommendation performance across different domains. (2) To the best of our knowledge, `GENET` is also the first pre-training framework on hypergraphs, which encodes high-order and multi-faceted (i.e., beyond pairwise) relationships with self-supervision. (3) Extensive experiments (i.e., on different real-world datasets, recommendation tasks, and side information) demonstrate that the proposed `GENET` significantly outperforms the SOTA competitor.

2 RELATED WORK

We briefly introduce relevant studies on recommendation models with side information, pre-training, and hypergraphs.

2.1 Recommendation Models with Side Information

Using side information to improve recommendation has been extensively studied [6, 15, 23, 36, 43]. This section discusses three methods of improving recommendation systems using side information. The first method involves combining features from side information, exemplified by HIRE [17] and FDSA [6] which integrate these details in different ways. The second approach employs additional models, like graph models, to capture relationships between users and items, with FREEDOM [43] and Flashback [36] constructing various item graphs. The third method uses side information as a supervisory signal, as seen in SPEX [15] which incorporates social network data through multitask learning. Although these methods show improvements, they may fall short in representing heterogeneous side information and in generalization capabilities.

2.2 Recommendation Models with Pre-training

Pre-training in recommendation models is mainly divided into two categories: one based on user-item feedback data and the other on side information. In the first category, models like SimRec [33] and BERT4Rec [25] enhance robustness and efficiency by pre-training on user-item feedback. The second category, exemplified by Graph-Flashback [23] and S^3-rec [42], employs various pre-training strategies to learn representations from side information, enriching user and item profiles. These methods overlook the diversity of data formats and lose high-order information.

2.3 Recommendation models based on Hypergraph

Recent studies increasingly focus on using hypergraphs for recommendations. For instance, HyperRec [27] employ hypergraph convolution networks to capture short-term correlations in sequence recommendation tasks and combine prediction tasks with contrastive tasks on hypergraphs. MHCN [38] introduces a multi-channel hypergraph convolutional network to enhance social recommendations. We identify a potential risk in previous methods: the direct use of side information in training recommendation tasks, leading to negative transfer effects.

Our method leverages pre-trained representations on a hypergraph, integrating diverse side information like sequences and graphs. This strategy alleviates the above issues, may prevent negative transfer effects, and offers a versatile solution for different side information types.

3 Pre-training on Side Information

To pre-train on side information, we first construct a hypergraph on the side information (Section 3.1), perform propagation on the hypergraph to obtain node embeddings for users and items (Section 3.2) and optimize the node embeddings via a set of pre-training tasks (Section 3.3 to Section 3.4).
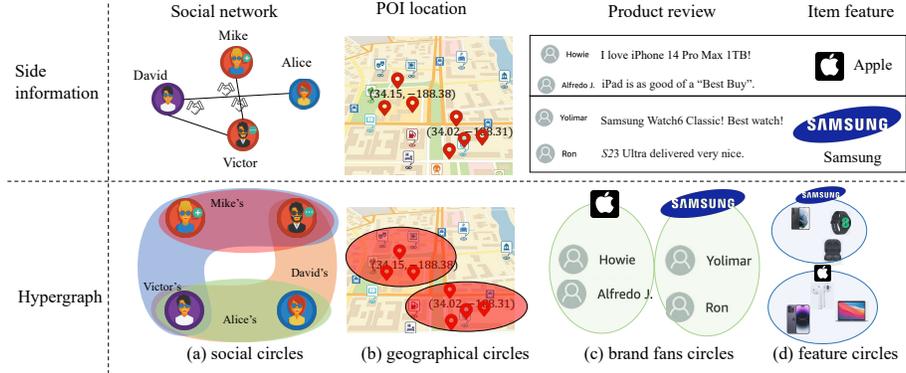
Fig. 2. Different types of side information and the construction of hypergraphs

Formally, a hypergraph $\mathcal{G} = (\mathcal{X}, \mathcal{E})$ consists of a set of vertices $\mathcal{X}$ and a set of hyperedges $\mathcal{E}$, where $\mathcal{X} = \{x_1, x_2, \cdots, x_{|\mathcal{X}|}\}$, each node can be either a user or an item. $e = \{e_1, e_2, \cdots, e_{|\mathcal{E}|}\}$, each edge connects multiple nodes. A node x_i and a hyperedge e_p are incident if $x_i \in e_p$. The incidence matrix $\mathbf{H} \in \{0, 1\}^{|\mathcal{X}| \times |\mathcal{E}|}$ describes the structure of $\mathcal{G}$, where $h(x_i, e_p) = 1$ if x_i and e_p are incident. Two nodes x_i, x_j are adjacent if there exists at least one hyperedge e_p that $x_i \in e_p, x_j \in e_p$.

3.1 Construction of Hypergraphs

In constructing hypergraphs ($\mathcal{G}$), we focus on four types of side information: social networks, POI locations, product reviews, and item brands. These encompass various data formats such as relational data, categorical and numerical features, and text sequences, suitable for common recommendation scenarios like POI and e-commerce recommendations.

- **Social Network Hypergraph:** Constructs hyperedges to represent social circles, capturing high-order and multifaceted social connections. For example, a hyperedge represents David's social circle, connecting individuals like Victor and Alice if they are friends of David.
- **POI Location Hypergraph:** Utilizes geographical data (longitude and latitude) to segment POIs into regions. POIs are clustered into regions using k-means clustering, and a hyperedge is constructed to represent each region.
- **Product Review Hypergraph:** Builds hyperedges based on brand-aware sentiments. We analyze sentiments of reviews under a brand to connect users who exhibit brand loyalty.
- **Item Feature Hypergraph:** Handles item-related metadata, such as brands and categories. Hyperedges are constructed for each brand and category. An item is connected to the corresponding hyperedges if it belongs to a specific brand or category.
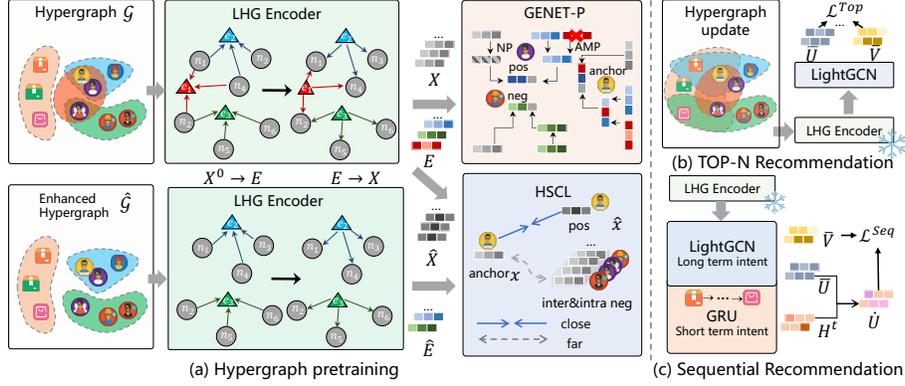
Fig. 3. The framework of `GENET` pre-training phase and fine-tuning stage.

3.2 Node Representation via Hypergraph Convolution

Input of `GENET` is a hypergraph $\mathcal{G}$, a sparse matrix of the initial node embeddings $\mathbf{X}^0 \in \mathbb{R}^{|\mathcal{X}| \times |\mathcal{X}|}$, where each node is represented by one-hot encoding, the degree matrix of hyperedges $\mathbf{D}^e$ and the degree matrix of $\mathbf{D}^v$.

LHG Encoder Inspired by the general Hypergraph Neural Network HGNN$^+$ [10], we derive the node embeddings by introducing an edge embedding vector for each hyperedge and propagating edge embeddings to incident nodes. As shown in Figure 3(a), the propagation is performed from nodes to hyperedges and from hyperedges to nodes, as shown below:

$$
\begin{aligned}
\mathbf{E} &= \mathbf{D}^{e^{-1}} \mathbf{H}^{\top} \mathbf{X}^0 \mathbf{\Theta}^0, \\
\mathbf{X} &= \mathbf{D}^{v^{-1}} \mathbf{H} \mathbf{I} \mathbf{E}.
\end{aligned}
\tag{1}
$$

where $\mathbf{E}$ is the edges embedding matrix. $\mathbf{\Theta}$ is the trainable matrix. $\mathbf{X}$ is the nodes embedding matrix. $\mathbf{I}$ is an identity matrix. Noted that our encoder is Light HyperGraph encoding (LHG), because different from HGNN$^+$, we remove the non-linear activation function, which may affect the aggregation of information in the hypergraph.

3.3 Hyperlink Prediction

One pre-training task on $\mathcal{G}$ is hyperlink prediction, i.e., to predict if two nodes are adjacent given a hyperedge. The motivation for using hyperlink prediction is that adjacent nodes in a hyperedge are more likely to be mutually related. We phrase this property as mutual connectivity.

Mutual connectivity highlights the importance of predicting adjacent nodes. We propose the hyperlink prediction task, where given a hyperedge e_h, an anchor node x_i which is incident with e_h, the positive node x_j which is also incident

with e_h is discriminated from a negative node x_k which is not incident with e_h. Furthermore, we corrupt the representation of the positive sample x_j to increase the difficulty of discrimination and the robustness of pre-training tasks.

Existing graph pre-training methods corrupt node representations by cutting down some edges and creating a subgraph, and then using a graph encoder on the subgraph [18,19] to derive a representation of the target node. This strategy is infeasible in hypergraphs because hyperedges are indescomposable [26]. If we cut down the adjacency between the positive node and the anchor node, the hyperedge is broken, the meaning of the hyperedge is incomplete, and high-order/multi-faceted information is discarded.

To avoid disrupting the high-order/multi-faceted information, we propose to corrupt the positive sample by combining direct node perturbation (i.e., NP) and incidence matrix perturbation (IMP).

Node perturbation randomly draws a node embedding $\mathbf{x}_j^g$ from:

$$\mathbf{x}_j^g = \mathcal{N}(\mathbf{x}_j, \lambda \mathbf{I}), \tag{2}$$

where $\mathcal{N}$ is the Gaussian distribution with the original node representation as its mean, $\mathbf{I}$ is the identity matrix, and λ is the covariance scalar.

Incidence matrix perturbation. To avoid information leakage, i.e., the node embedding may contain information about the condition hyperedge e_h, we remove the connection and rewrite the incidence matrix by letting $h(x_j, e_h) = 0$. The corrupted incidence matrix is denoted as $\hat{\mathbf{H}}$. We then propagate the filtered hyperedges back to nodes by:

$$\mathbf{x}_j^a = (\mathbf{D}^{v^{-1}} \hat{\mathbf{H}} \mathbf{W} \mathbf{E})_j, \tag{3}$$

where $(\cdot)_j$ represents the j-th row elements.

Then we merge $\mathbf{x}_i^g$ and $\mathbf{x}_i^a$ and obtain the corrupted node representation $\widetilde{\mathbf{x}}_j$, which is shown below:

$$\widetilde{\mathbf{x}}_j = \mathbf{x}_j^g + \mathbf{x}_j^a. \tag{4}$$

Given the hyperedge e_h, the anchor node x_i, the positive node x_j and the negative node x_k. We use the ranking loss $\mathcal{L}^P$ as the training objective of the hyperlink prediction task (GENET-P), which is shown below:

$$\mathcal{L}^P = - \sum_{i,j,k} \sigma(\mathbf{x}_i \cdot \widetilde{\mathbf{x}}_j - \mathbf{x}_i \cdot \mathbf{x}_k), \tag{5}$$

where $\sigma(\cdot)$ is the sigmoid function.

3.4 Hypergraph contrastive learning

Furthermore, as label information is unavailable during the pre-training stage, we introduce Hypergraph Self-Contrastive Learning (HSCL) to better the global information of the hypergraph and the local variability among nodes within a hyperedge.

First, we propose self-contrastive inter-hyperedges. Formally, given a node set $\mathcal{X}$ in a batch, we treat any node embedding $\mathbf{x}_i$ as the anchor, the positive sample is the anchor's augmented embedding $\mathbf{x}_i^a$ from the augmented hypergraph $\hat{\mathcal{G}}$ by Equation 3, and negative samples are other node embeddings $\mathbf{x}_j$ in $\mathcal{X}$. We denote the inter-hyperedge contrastive loss as $\mathcal{L}^{inter}$.

$$\mathcal{L}^{inter} = \sum_{x_i \in \mathcal{X}} -\log \frac{\exp\left(sim(\mathbf{x}_i, \mathbf{x}_i^a)/\tau\right)}{\sum_{x_j \in \mathcal{X}} \exp\left(sim(\mathbf{x}_i, \mathbf{x}_j)/\tau\right)}, \tag{6}$$

where τ is a temperature parameter.

Second, we propose intra-hyperedge self-contrastive learning. Formally, given a hyperedge set $\mathcal{E}$ in a batch, we sample a node set $\mathcal{X}_h$ from each hyperedge $e_h \in \mathcal{E}$ and $|\mathcal{X}_h| = K$. Given an anchor node $x_i \in \mathcal{X}_h$, its positive sample $\mathbf{x}_i^a$ is obtained by Equation 3, and its negative samples are other nodes embedding $x_j \in e_h$. We denote the intra- hyperedge contrastive loss as $\mathcal{L}^{intra}$.

$$\mathcal{L}^{intra} = \frac{1}{|\mathcal{E}||\mathcal{X}_h|} \sum_{e_h \in \mathcal{E}} \sum_{x_i \in \mathcal{X}_h} -\log \frac{\exp\left(sim(\mathbf{x}_i, \mathbf{x}_i^a)/\tau)\right)}{\sum_{x_j \in \mathcal{X}_h} \exp\left(sim(\mathbf{x}_i, \mathbf{x}_j)/\tau)\right)}, \tag{7}$$

where τ is a temperature parameter.

The overall loss in the pre-training stage includes the objective loss in hyperlink prediction $\mathcal{L}^P$ and the hypergraph self-contrastive from global and local levels, as shown below

$$\mathcal{L}^{Pre} = \mathcal{L}^P + \beta_1 \mathcal{L}^{intra} + \beta_2 \mathcal{L}^{inter}. \tag{8}$$

where β_1 and β_2 are hyper-parameters.

4 Fine-tuning on User Feedback

In recommendation systems, TOP-N recommendation and sequence recommendation are two of the most common recommendation tasks. GENET designs a simple yet effective downstream fine-tuning approach for both tasks.

4.1 Top-N recommendation

In the pre-training phase, users and items are isolated, and there is no direct connection between them. Therefore, as shown in Figure 3(b), in the fine-tuning stage, we first establish connections between users and items.

Hypergraph updating. We update the incidence matrix according to the interaction between users and items. Formally, if user u_i is connected to hyperedge e_p and item v_j is connected to hyperedge e_q, and u_i interacts with v_j (e.g., buy or click), we set $h(u_i, e_q) = 1$ and $h(v_j, e_p) = 1$.

The node embedding $\check{\mathbf{X}}$ is updated by:

$$\check{\mathbf{X}} = \mathbf{X} + \mathbf{D}^{v^{-1}} \tilde{\mathbf{H}} \mathbf{W} \mathbf{E}. \tag{9}$$

where $\mathbf{E}$ is obtained in the pre-training stage and $\mathbf{W}$ is the parameters obtained in the pre-training stage, $\tilde{\mathbf{H}}$ is the updated hypergraph.

LightGCN fine-tuning. To comprehensively understand the relationship between users and items, we build a user-item bipartite graph $\mathcal{A}$ and employ the classical LightGCN to finetune the representations

$$\overline{\mathbf{u}}_i, \overline{\mathbf{v}}_j = LightGCN(\breve{\mathbf{u}}_i, \breve{\mathbf{v}}_j, \mathcal{A}). \tag{10}$$

where $\breve{\mathbf{u}}_i$, and $\breve{\mathbf{v}}_j$ are the user embedding and the item embedding in $\breve{\mathbf{X}}$, respectively.

Then, we introduce $\mathcal{L}^{Top}$ as the training objective in the fine-tuning stage

$$\mathcal{L}^{Top} = -\sum_{i,j,k} \sigma(\overline{\mathbf{u}}_i \overline{\mathbf{v}}^j - \overline{\mathbf{u}}_i \overline{\mathbf{v}}^k). \tag{11}$$

where the user i interacts with item j and not interacts with item k, $\sigma(\cdot)$ is the sigmoid function.

4.2 Sequential recommendation

Unlike the TOP-N recommendation, sequential recommendation focuses on not only long-term user intent but also short-term user intent to make item predictions for the next moment. We believe that LightGCN only captures the long-term user intent, therefore we extra employ GRU [1] to capture short-term user intent.

Formally, given a user i 's interaction sequence on most recent s items, $\mathcal{V}_i = \{v_{t-s+1}, \cdots, v_t\}$, our model GENET learns the user's short-term intents at moment t by GRU [1]

$$\begin{aligned} \mathbf{h}_i^t &= GRU(\mathcal{V}_i), \\ \dot{\mathbf{u}}_i^t &= \overline{\mathbf{u}_i} + \mathbf{h}_i^t. \end{aligned} \tag{12}$$

where $\mathbf{h}_i^t$ is the hidden state of user i at time t. The user representation $\dot{\mathbf{u}}_i^t$ at time t is an ensemble of short-term user intent and long-term user preference.

Then, We introduce the sequential recommendation loss $\mathcal{L}^{Seq}$.

$$\mathcal{L}^{Seq} = -\sum_{i,j,k} \sigma(\dot{\mathbf{u}}_i^t \overline{\mathbf{v}}_j - \dot{\mathbf{u}}_i^t \overline{\mathbf{v}}_k). \tag{13}$$

where the user i interacts with item j and not interacts with item k, $\sigma(\cdot)$ is the sigmoid function.

5 EXPERIMENT

In this section, we conduct experiments to study the following questions:

RQ1: Can GENET generalize well to different side information?
RQ2: How does each component in GENET contribute to the overall performance?
RQ3: Is GENET suitable for pre-training hypergraphs?
RQ4: Does GENET alleviate the users and items cold start problems?

Table 1. Statistics of the datasets

Dataset	#User	#Item	#Interactions	#Relationship	#Side Information	#Density
Gowalla	10,671	48,587	1,195,532	93,260	Friendship,POI location	0.23%
Foursquare	22,065	9,891	182,044	79,704	Friendship,POI location	0.08%
Books	28,898	85,188	3,233,028	160,448,633	Product review, category	0.13%

5.1 Experimental Setup

Datasets: Our experiments are conducted on three real-world datasets: Gowalla [37], Foursquare [7], and Books [20]. We evaluate `GENET` on these datasets because they are evaluated frequently and cover a wide range of common types of side information in recommendation systems, including item brand, item category, item review, user social networks, and POI geo-location. These datasets include timestamped feedback. Thus, they can be used for Top-N recommendations and sequential recommendations. Statistics of the datasets are shown in Table 1.

Evaluation Metrics We evaluated the TOP-N recommendation performance and Sequential recommendation performance of all models on three datasets using two widely-used metrics, N@K(NDCG@K) and R@K(Recall@K), where K=[10,20]. We adopt the widely used leave-one-out evaluation method, similar to SPEX [15].

Competitors. We compare `GENET` to an extensive list of competitors. We classify based on whether competitors use side information and pre-training. For the TOP-N task, the competitors are (1).without side information and pre-training: NGCF [30], LightGCN [11],SGL [32], HCCF [34], (2). with pre-training but without side information: SimRec [33], (3). with side information but without pre-training: SPEX [15], HIRE [17], FREEDOM [43], Flashback [36], HGNN$^+$ [10], (4). with side information and pre-training: Graph-Flashback [23]. There are three categories in (3): Feature [17], Signal [15] and Model [8,10,36,43].

For the Sequential task, the competitors are (1). without side information and pre-training: SASRec [13], ContrastVAE [31], BRET4Rec [25], CBiT [4], (2). with side information but without pre-training: Flashback [36],HGNN$^+$ [10], (3). with side information and pre-training: S^3-Rec [42],Graph-Flashback [23]. More details about competitors can be found in the related works 2.1 section.

Implementations. The proposed `GENET` framework is implemented in Pytorch library and Deep HyperGraph [9,10]. We adopt the Adam [14] optimizer in both the pre-training and fine-tuning stages. In the pre-training stage, the embedding sizes of the nodes and hyperedges are both 64, the batch size is 4096, the learning rate is 0.0005, and the epochs are 500. The noise intensity λ is 0.1. The hypergraph contrastive learning hyperparameter β_1 and β_2 are 0.005 and 0.01, respectively. For the updated hypergraph, we only use the training dataset. In the fine-tuning stage, The epochs are 10, and the learning rate is 0.0005. During the first three epochs, we amplify the learning rate by ten times. Models are in the same setting for fairness. The purpose is to expedite the transfer of the pre-trained universal representations to downstream tasks. The LightGCN

Table 2. Comparative performance for TOP-N recommendation, bold fonts for best results, underlined scores for the second best results.

	Gowalla				Foursquare				Book			
	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20
NGCF	0.5091	0.5393	0.7653	0.8839	0.5513	0.5735	0.6894	0.7767	0.6593	0.6748	<u>0.8894</u>	<u>0.9501</u>
LightGCN	0.2782	0.3095	0.4257	0.5502	0.6091	0.6276	0.7709	0.8439	0.5792	0.5994	0.8303	0.9097
SGL	0.5607	0.5875	0.7501	0.8566	0.6192	0.6336	0.7597	0.8164	<u>0.6689</u>	<u>0.6845</u>	0.8433	0.9047
HCCF	0.6080	0.6281	0.8624	0.9406	0.5678	0.5882	0.7204	0.8011	0.6005	0.6200	0.8550	0.9318
SimRec	<u>0.6132</u>	<u>0.6336</u>	<u>0.8682</u>	<u>0.9478</u>	<u>0.6274</u>	<u>0.6474</u>	<u>0.7984</u>	<u>0.8772</u>	0.6026	0.6231	0.8552	0.9354
NGCF &SPEX	0.3542	0.3919	0.5571	0.7063	0.5357	0.5665	0.6670	0.7882	0.5665	0.5881	0.8215	0.9065
LightGCN &SPEX	0.4469	0.4810	0.6802	0.8147	0.5929	0.6119	0.7350	0.8278	0.6573	0.6723	0.8801	0.9403
HIRE	0.2335	0.2662	0.3825	0.5120	0.5679	0.5889	0.7249	0.8081	0.2887	0.3365	0.5056	0.6952
FREEDOM	0.5916	0.6160	0.8536	0.9496	0.6151	0.6404	0.7740	0.8741	N/A	N/A	N/A	N/A
Flashback	0.1354	0.1663	0.2427	0.3655	0.5278	0.5490	0.6819	0.7663	N/A	N/A	N/A	N/A
HGNN$^+$	0.6256	0.6630	0.7244	0.8711	0.3859	0.4289	0.5013	0.6699	0.1914	0.2241	0.2904	0.4207
Graph-Flashback	0.2052	0.2376	0.3384	0.4672	0.5398	0.5621	0.6999	0.7879	N/A	N/A	N/A	N/A
GENET	**0.7193**	**0.7289**	**0.9412**	**0.9785**	**0.7081**	**0.7241**	**0.8835**	**0.9463**	**0.7009**	**0.7129**	**0.9159**	**0.9626**

layer number K is 2. To potentially address oversmoothing, we also explore the possibility of incorporating contrastive learning and denoising techniques [39,43].

5.2 Generalization of GENET

To answer **RQ1**, we evaluate GENET on three datasets (Gowalla, Foursquare, Books) for two recommendation tasks (Top-N and sequential recommendation), showing superior performance across all metrics and tasks. For instance, in Top-N recommendation, it increased NDCG@10 by up to 14.98% and Recall@10 by up to 10.66%. In sequential recommendation, the increase was up to 37.88% for NDCG@10 and 19.58% for Recall@10. Existing methods demonstrated weaker generalization, especially across different datasets. GENET consistently achieved a recommendation performance of over 0.6000 for NDCG@10 across all datasets and tasks, indicating strong generalization capability. Moreover, pre-training positively impacted recommendation performance. Models like SimRec and Graph-Flashback significantly improved after pre-training. S^3-rec, by incorporating pre-training into SASRec, achieved improvements of 119.56% and 131.28% in Recall@10 and NDCG@10, respectively, on the Gowalla dataset. Overall, GENET demonstrated robust generalization across three datasets and two tasks, outperforming 25 competitors in various side information scenarios.

5.3 Ablation Study

To answer **RQ2**, we assessed the effectiveness of various variants of GENET by removing components such as the node perturbation method NP, incidence ma-

Table 3. Comparative performance on Sequential recommendation

	Gowalla				Foursquare				Book			
	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20	N@10	N@20	R@10	R@20
SASRec	0.2516	0.2810	0.3784	0.4954	0.6425	0.6605	0.7906	0.8620	0.2474	0.2904	0.4489	0.6192
BRET4Rec	0.2693	0.3031	0.4188	0.5531	0.6988	0.7140	0.8424	0.9024	0.4329	0.4624	0.6288	0.7455
ContrastVAE	0.5403	0.5680	0.8129	0.9211	0.6926	0.7095	0.8472	0.9138	0.3868	0.4234	0.6295	0.7734
CBiT	0.6591	0.6766	0.8592	0.9273	0.6460	0.6639	0.8485	0.9083	0.4356	0.4663	0.6440	0.7649
HGNN$^+$	0.6641	0.6955	0.7618	0.8846	0.2550	0.3019	0.3841	0.5701	0.2380	0.2531	0.3301	0.3892
Flashback	0.7087	0.7223	0.8873	0.9406	0.7270	0.7392	0.8427	0.8909	N/A	N/A	N/A	N/A
S^3-Rec	0.5819	0.6021	0.8308	0.9097	0.6634	0.6801	0.8241	0.8897	0.4239	0.4573	0.6738	0.8051
Graph-Flashback	0.7282	0.7408	0.8939	0.9432	0.7368	0.7479	0.8508	0.8987	N/A	N/A	N/A	N/A
GENET	**0.7505**	**0.7578**	**0.9468**	**0.9752**	**0.7561**	**0.7690**	**0.9040**	**0.9547**	**0.6006**	**0.6211**	**0.8057**	**0.8866**

trix perturbation IMP, and hypergraph self-contrastive learning HSCL in pre-training. As is shown in Figure 4(a), the results indicated that each component is crucial for performance. On the Gowalla dataset, removing NP(i.e.,"w/o" NP), IMP(i.e.,"w/o" IMP), and HSCL(i.e.,"w/o" HSCL) led to declines in Recall@10 and NDCG@10 by 0.0124 and 0.0185, 0.0181 and 0.0332, 0.0103 and 0.0196, respectively. Particularly, removing IMP had the most significant impact as it is a core component of the pre-training stage. Its removal left only NP for node perturbation, oversimplifying the link prediction task and affecting the model's ability to learn intrinsic relationships and higher-order dependencies within the data.

5.4 Pre-training on Hypergraphs

To study whether the proposed pre-training tasks are suitable for pre-training on hypergraphs (**RQ3**), we studied the suitability of pre-training tasks for hypergraphs (**RQ3**), comparing the pre-training phase (GENET-P) of GENET with two common graph-based pretext tasks: Link Prediction (LP) and Node Feature Reconstruction (NR). We utilized two popular graph pre-training strategies: Contrastive and Generative. Consequently, we obtained four pre-training tasks: LP-C, LP-G, NR-C, and NR-G. As a baseline, we also implemented random initialization of node embeddings. These pre-training methods were evaluated by directly computing user-item similarity for recommendation (No-tuning) and performance after fine-tuning with downstream models, including Recall@10 and NDCG@10, which are shown in 4(b)(c).

Based on the results, we draw the following conclusions: (1) GENET-P outperforms all other pre-training tasks on hypergraphs. For instance, on the FourSquare dataset, it increases Recall@10 by 28.24% and NDCG@1 by 30.86% compared to the best competitor (i.e., LP-C). (2) GENET-P is already very powerful without fine-tuning on feedback data, especially on the BOOK dataset, where all competitors' Recall@10 is under 0.5379 (0.2792 in NDCG@10), and GENET-P's
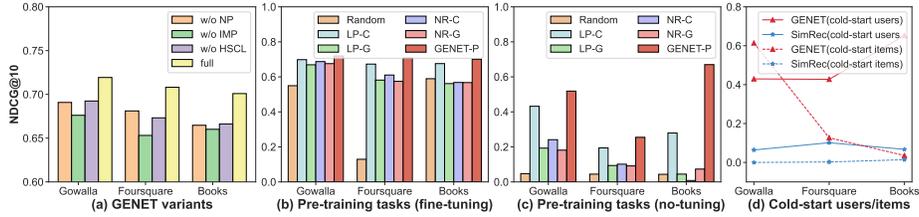
Fig. 4. Performances of GENET variants, different pre-training tasks and cold start.

performance achieves 0.8946(0.6700 in NDCG@10). (3) Compared to not pre-training and directly training the downstream model (Random), all pre-training strategies lead to significant performance improvement. (4) The link prediction (LP) pre-training strategy performs better than the node feature reconstruction (NR) strategy. (5) The contrastive strategy is superior to the generative strategy in pre-training generative strategy, resulting in lower uncertainty.

5.5 Cold-start

To answer **RQ4**, we evaluate the performance of `GENET` in scenarios where it encounters user cold-start and item cold-start situations.

To create a user cold-start scenario, we select the last 1% of users based on their interaction counts and designate them as cold-start users. We remove these cold-start users from the training dataset and retain only cold-start users in the test dataset. We create an item cold-start scenario in the same manner.

As shown in Figure 4(d), `GENET` demonstrates a remarkable capability to deal with cold-start users and items. It has a 321.50% to 869.53% improvement in NDCG@10 and a 258.98% to 538.89% improvement in Recall@10 compared with SimRec in the users cold-start scenario, and it has a 33.33% to 4069.35% improvement in Recall@10 and 152.11% to 5191.67% improvement in NDCG@10 in the items cold-start scenario. This highlights the effectiveness of the hypergraph constructed from social networks and user reviews in accurately modeling users' preferences and the effectiveness of the hypergraph constructed from POI location and item features in accurately modeling items' properties.

6 CONCLUSION

In this work, we proposed a novel framework called Generalized hypErgraph pretraiNing on sidE informaTion(`GENET`). It aims to enhance recommendation performances by integrating side information. `GENET` is based on pretraining-finetuning, where heterogeneous side information is constructed as a unified hypergraph. We propose novel pre-training tasks tailored explicitly for hypergraphs, which can effectively capture high-order and multi-faceted relationships in hypergraphs. The finetuning is conducted with simple and straightforward

approaches. Extensive experimental results show that `GENET` has excellent generalization. It outperforms SOTA competitors on two recommendation tasks on three public datasets with varying side information.

References

1. Chung, J., Gulcehre, C., Cho, K., Bengio, Y.: Empirical evaluation of gated recurrent neural networks on sequence modeling. arXiv preprint arXiv:1412.3555 (2014)
2. Devlin, J., Chang, M.W., Lee, K., Toutanova, K.: BERT: Pre-training of deep bidirectional transformers for language understanding. In: HAACL-HLT. vol. 1, pp. 4171–4186 (2019)
3. Dong, X., Zhan, X., Wu, Y., Wei, Y., Kampffmeyer, M.C., Wei, X., Lu, M., Wang, Y., Liang, X.: M5product: Self-harmonized contrastive learning for e-commercial multi-modal pretraining. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 21252–21262 (2022)
4. Du, H., Shi, H., Zhao, P., Wang, D., Sheng, V.S., Liu, Y., Liu, G., Zhao, L.: Contrastive learning with bidirectional transformers for sequential recommendation. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. pp. 396–405 (2022)
5. El-Kishky, A., Markovich, T., Park, S., Verma, C., Kim, B., Eskander, R., Malkov, Y., Portman, F., Samaniego, S., Xiao, Y., et al.: Twhin: Embedding the twitter heterogeneous information network for personalized recommendation. In: Proceedings of the 28th ACM SIGKDD conference on knowledge discovery and data mining. pp. 2842–2850 (2022)
6. Fang, Y., Si, L.: Matrix co-factorization for recommendation with rich side information and implicit feedback. In: Proceedings of the 2nd International Workshop on Information Heterogeneity and Fusion in Recommender Systems. pp. 65–69 (2011)
7. Feng, J., Li, Y., Zhang, C., Sun, F., Meng, F., Guo, A., Jin, D.: Deepmove: Predicting human mobility with attentional recurrent networks. In: Proceedings of the 2018 world wide web conference. pp. 1459–1468 (2018)
8. Feng, Y., You, H., Zhang, Z., Ji, R., Gao, Y.: Hypergraph neural networks. In: AAAI. pp. 3558–3565 (2019)
9. Feng, Y., You, H., Zhang, Z., Ji, R., Gao, Y.: Hypergraph neural networks. In: Proceedings of the AAAI conference on artificial intelligence. vol. 33, pp. 3558–3565 (2019)
10. Gao, Y., Feng, Y., Ji, S., Ji, R.: Hgnn$^+$: General hypergraph neural networks. IEEE Transactions on Pattern Analysis and Machine Intelligence (2022)
11. He, X., Deng, K., Wang, X., Li, Y., Zhang, Y., Wang, M.: 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 (2020)
12. Hu, Z., Dong, Y., Wang, K., Chang, K.W., Sun, Y.: Gpt-gnn: Generative pre-training of graph neural networks. In: Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining. pp. 1857–1867 (2020)
13. Kang, W.C., McAuley, J.: Self-attentive sequential recommendation. In: 2018 IEEE international conference on data mining (ICDM). pp. 197–206. IEEE (2018)

14. Kinga, D., Adam, J.B., et al.: A method for stochastic optimization. In: International conference on learning representations (ICLR). vol. 5, p. 6. San Diego, California; (2015)
15. Li, H., Li, L., Xv, G., Lin, C., Li, K., Jiang, B.: Spex: A generic framework for enhancing neural social recommendation. ACM Transactions on Information Systems (TOIS) **40**(2), 1–33 (2021)
16. Li, X., Chen, H.: Recommendation as link prediction in bipartite graphs: A graph kernel-based machine learning approach. Decision Support Systems **54**(2), 880–890 (2013)
17. Liu, T., Wang, Z., Tang, J., Yang, S., Huang, G.Y., Liu, Z.: Recommender systems with heterogeneous side information. In: The world wide web conference. pp. 3027–3033 (2019)
18. Liu, Z., Yu, X., Fang, Y., Zhang, X.: Graphprompt: Unifying pre-training and downstream tasks for graph neural networks. In: Proceedings of the ACM Web Conference 2023. pp. 417–428 (2023)
19. Lu, Y., Jiang, X., Fang, Y., Shi, C.: Learning to pre-train graph neural networks. In: Proceedings of the AAAI conference on artificial intelligence. vol. 35, pp. 4276–4284 (2021)
20. McAuley, J., Targett, C., Shi, Q., Van Den Hengel, A.: Image-based recommendations on styles and substitutes. In: Proceedings of the 38th international ACM SIGIR conference on research and development in information retrieval. pp. 43–52 (2015)
21. Ning, X., Karypis, G.: Sparse linear methods with side information for top-n recommendations. In: Proceedings of the sixth ACM conference on Recommender systems. pp. 155–162 (2012)
22. Pfadler, A., Zhao, H., Wang, J., Wang, L., Huang, P., Lee, D.L.: Billion-scale recommendation with heterogeneous side information at taobao. In: 2020 IEEE 36th International Conference on Data Engineering (ICDE). pp. 1667–1676. IEEE (2020)
23. Rao, X., Chen, L., Liu, Y., Shang, S., Yao, B., Han, P.: Graph-flashback network for next location recommendation. In: Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining. pp. 1463–1471 (2022)
24. Shimizu, R., Matsutani, M., Goto, M.: An explainable recommendation framework based on an improved knowledge graph attention network with massive volumes of side information. Knowledge-Based Systems **239**, 107970 (2022)
25. Sun, F., Liu, J., Wu, J., Pei, C., Lin, X., Ou, W., Jiang, P.: Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer. In: Proceedings of the 28th ACM international conference on information and knowledge management. pp. 1441–1450 (2019)
26. Tu, K., Cui, P., Wang, X., Wang, F., Zhu, W.: Structural deep embedding for hyper-networks. In: Proceedings of the AAAI conference on artificial intelligence. vol. 32 (2018)
27. Wang, J., Ding, K., Hong, L., Liu, H., Caverlee, J.: Next-item recommendation with sequential hypergraphs. In: SIGIR. pp. 1101–1110 (2020)
28. Wang, J., Huang, P., Zhao, H., Zhang, Z., Zhao, B., Lee, D.L.: Billion-scale commodity embedding for e-commerce recommendation in alibaba. In: Proceedings of the 24th ACM SIGKDD international conference on knowledge discovery & data mining. pp. 839–848 (2018)
29. Wang, X., He, X., Feng, F., Nie, L., Chua, T.S.: Tem: Tree-enhanced embedding model for explainable recommendation. In: Proceedings of the 2018 world wide web conference. pp. 1543–1552 (2018)

30. Wang, X., He, X., Wang, M., Feng, F., Chua, T.S.: Neural graph collaborative filtering. In: Proceedings of the 42nd international ACM SIGIR conference on Research and development in Information Retrieval. pp. 165–174 (2019)

31. Wang, Y., Zhang, H., Liu, Z., Yang, L., Yu, P.S.: Contrastvae: Contrastive variational autoencoder for sequential recommendation. In: Proceedings of the 31st ACM International Conference on Information & Knowledge Management. pp. 2056–2066 (2022)

32. Wu, J., Wang, X., Feng, F., He, X., Chen, L., Lian, J., Xie, X.: 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)

33. Xia, L., Huang, C., Shi, J., Xu, Y.: Graph-less collaborative filtering. In: Proceedings of the ACM Web Conference 2023. pp. 17–27 (2023)

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

35. Xie, Y., Zhou, P., Kim, S.: Decoupled side information fusion for sequential recommendation. In: Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval. pp. 1611–1621 (2022)

36. Yang, D., Fankhauser, B., Rosso, P., Cudre-Mauroux, P.: Location prediction over sparse user mobility traces using rnns. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence. pp. 2184–2190 (2020)

37. Yin, H., Cui, B., Chen, L., Hu, Z., Zhang, C.: Modeling location-based user rating profiles for personalized recommendation. ACM Transactions on Knowledge Discovery from Data (TKDD) **9**(3), 1–41 (2015)

38. Yu, J., Yin, H., Li, J., Wang, Q., Hung, N.Q.V., Zhang, X.: Self-supervised multi-channel hypergraph convolutional network for social recommendation. In: WWW. pp. 413–424 (2021)

39. Yu, J., Yin, H., Xia, X., Chen, T., Cui, L., Nguyen, Q.V.H.: 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 (2022)

40. Zhang, T., Zhao, P., Liu, Y., Sheng, V.S., Xu, J., Wang, D., Liu, G., Zhou, X., et al.: Feature-level deeper self-attention network for sequential recommendation. In: IJCAI. pp. 4320–4326 (2019)

41. Zhao, F., Xiao, M., Guo, Y.: Predictive collaborative filtering with side information. In: IJCAI. pp. 2385–2391 (2016)

42. Zhou, K., Wang, H., Zhao, W.X., Zhu, Y., Wang, S., Zhang, F., Wang, Z., Wen, J.R.: S3-rec: Self-supervised learning for sequential recommendation with mutual information maximization. In: Proceedings of the 29th ACM international conference on information & knowledge management. pp. 1893–1902 (2020)

43. Zhou, X.: A tale of two graphs: Freezing and denoising graph structures for multimodal recommendation. arXiv preprint arXiv:2211.06924 (2022)