

Detecting Beneficial Feature Interactions for Recommender Systems

Yixin Su¹, Rui Zhang¹, Sarah Erfani¹, Zhenghua Xu²

¹University of Melbourne, ²Hebei University of Technology

yixins1@student.unimelb.edu.au, rui.zhang@unimelb.edu.au, sarah.erfani@unimelb.edu.au, zhenghua.xu@hebut.edu.cn

Abstract

Feature interactions are essential for achieving high accuracy in recommender systems. Many studies take into account the interaction between every pair of features. However, this is suboptimal because some feature interactions may not be that relevant to the recommendation result, and taking them into account may introduce noise and decrease recommendation accuracy. To make the best out of feature interactions, we propose a graph neural network approach to effectively model them, together with a novel technique to automatically detect those feature interactions that are beneficial in terms of recommendation accuracy. The automatic feature interaction detection is achieved via edge prediction with an L_0 activation regularization. Our proposed model is proved to be effective through the information bottleneck principle and statistical interaction theory. Experimental results show that our model (i) outperforms existing baselines in terms of accuracy, and (ii) automatically identifies beneficial feature interactions.

1 Introduction

Recommender systems play a central role in addressing information overload issues in many Web applications, such as e-commerce, social media platforms, and lifestyle apps. The core of recommender systems is to predict how likely a user will interact with an item (e.g., purchase, click). An important technique is to discover the effects of features (e.g., contexts, user/item attributes) on the target prediction outcomes for fine-grained analysis (Shi, Larson, and Hanjalic 2014). Some features are correlated to each other, and the joint effects of these correlated features (i.e., *feature interactions*) are crucial for recommender systems to get high accuracy (Blondel et al. 2016; He and Chua 2017). For example, it is reasonable to recommend a user to use *Uber* on a *rainy* day at off-work hours (e.g., during *5-6pm*). In this situation, considering the feature interaction $\langle 5-6pm, rainy \rangle$ is more effective than considering the two features separately. Therefore, in recent years, many research efforts have been put in modeling the feature interactions (He and Chua 2017; Lian et al. 2018; Song et al. 2019). These models take into

account the interaction between every pair of features. However, in practice, not all feature interactions are relevant to the recommendation result (Langley et al. 1994; Siegmund et al. 2012). Modeling the feature interactions that provide little useful information may introduce noise and cause overfitting, and hence decrease the prediction accuracy (Zhang et al. 2017; Louizos, Welling, and Kingma 2018). For example, a user may use *Gmail* on a workday no matter what weather it is. However, if the interaction of workday and specific weather is taking into account in the model and due to the bias in the training set (e.g., the weather of the days when the Gmail usage data are collected happens to be cloudy), the interaction $\langle Monday, Cloudy \rangle$ might be picked up by the model and make less accurate recommendations on Gmail usage. Some work considers each feature interaction’s importance through the attention mechanism (Xiao et al. 2017; Song et al. 2019). However, these methods still take all feature interactions into account. Moreover, they capture each *individual* feature interaction’s contribution to the recommendation prediction, failing to capture the *holistic* contribution of a set of feature interactions together.

In this paper, we focus on identifying the set of feature interactions that together produce the best recommendation performance. To formulate this idea, we propose the novel problem of *detecting beneficial feature interactions*, which is defined as identifying the set of feature interactions that contribute most to the recommendation accuracy (see Definition 1 for details). Then, we propose a novel graph neural network (GNN)-based recommendation model, L_0 -SIGN, that detects the most beneficial feature interactions and utilizes only the beneficial feature interactions for recommendation, where each data sample is treated as a graph, features as nodes and feature interactions as edges. Specifically, our model consists of two components. One component is an L_0 edge prediction model, which detects the most beneficial feature interactions by predicting the existence of edges between nodes. To ensure the success of the detection, an L_0 activation regularization is proposed to encourage unbeneficial edges (i.e. feature interactions) to have the value of 0, which means that edge does not exist. Another component is a graph classification model, called **Statistical Interaction Graph neural Network (SIGN)**. SIGN takes nodes (i.e., features) and detected edges (i.e., beneficial feature interactions) as the input graph, and outputs recommendation pre-

*Rui Zhang and Zhenghua Xu are corresponding authors.
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

dictions by effectively modeling and aggregating the node pairs that are linked by an edge.

Theoretical analyses are further conducted to verify the effectiveness of our model. First, the most beneficial feature interactions are guaranteed to be detected in L_0 -SIGN. This is proved by showing that the empirical risk minimization procedure of L_0 -SIGN is a variational approximation of the Information Bottleneck (IB) principle, which is a golden criterion to find the most relevant information correlating to target outcomes from inputs (Tishby, Pereira, and Bialek 2000). Specifically, only the most beneficial feature interactions will be retained in L_0 -SIGN. It is because, in the training stage, our model simultaneously minimizes the number of detected feature interactions by the L_0 activation regularization, and maximizes the recommendation accuracy with the detected feature interactions. Second, we further show that the modeling of the detected feature interactions in SIGN is very effective. By accurately leveraging the relational reasoning ability of GNN, *iff* a feature interaction is detected to be beneficial in the L_0 edge prediction component, it will be modeled in SIGN as a statistical interaction (an interaction is called statistical interaction if the joint effects of variables are modeled correctly).

We summarize our contributions as follows:

- This is the first work to formulate the concept of beneficial feature interactions, and propose the problem of detecting beneficial feature interactions for recommender systems.
- We propose a model, named L_0 -SIGN, to detect the beneficial feature interactions via a graph neural network approach and L_0 regularization.
- We theoretically prove the effectiveness of L_0 -SIGN through the information bottleneck principle and statistical interaction theory.
- We have conducted extensive experimental studies. The results show that (i) L_0 -SIGN outperforms existing baselines in terms of accuracy; (ii) L_0 -SIGN effectively identifies beneficial feature interactions, which result in the superior prediction accuracy of our model.

2 Related Work

Feature Interaction based Recommender Systems Feature interaction modeling has attracted attention in recommender systems nowadays. Factorization machine (FM) (Rendle 2010) is one of the most popular algorithms in considering feature interactions. However, FM and its deep learning-based extensions (Xiao et al. 2017; He and Chua 2017; Guo et al. 2017) consider all possible feature interactions, while our model detects and models only the most beneficial feature interactions. Recent work considers the importance of feature interactions by giving each feature interaction an attention value (Xiao et al. 2017; Song et al. 2019), or select important interactions by using gates (Liu et al. 2020) or by searching in a tree-structured space (Luo et al. 2019). In these methods, the importance is not determined by the holistic contribution of the feature interactions, thus limit the performance. Our model detects beneficial feature interactions that together produce the best recommendation performance.

Graph Neural Networks (GNNs) GNNs can facilitate learning entities and their holistic relations (Battaglia et al. 2018). Existing work leverages GNNs to perform relational reasoning in various domains. For example, Duvenaud et al. (2015) and Gilmer et al. (2017) use GNNs to predict molecules’ property by learning their features from molecular graphs. Chang et al. (2016) use GNNs to learn object relations in dynamic physical systems. Besides, some relational reasoning models in computer vision such as (Santoro et al. 2017; Wang et al. 2018) have been shown to be variations of GNNs (Battaglia et al. 2018). Fi-GNN (Li et al. 2019) use GNNs for feature interaction modeling in CTR prediction. However, it still models all feature interactions. Our model theoretically connects the beneficial feature interactions in recommender systems to the edge set in graphs and leverages the relational reasoning ability of GNNs to model beneficial feature interactions’ holistic contribution to the recommendation predictions.

L_0 Regularization L_0 regularization sparsifies models by penalizing non-zero parameters. Due to the problem of non-differentiable, it does not attract attention previously in deep learning domains until Louizos, Welling, and Kingma (2018) solve this problem by proposing a hard concrete distribution in L_0 regularization. Then, L_0 regularization has been commonly utilized to compress neural networks (Tsang et al. 2018; Shi, Glocker, and Castro 2019; Yang et al. 2017). We explore to utilize L_0 regularization to limit the number of detected edges in feature graphs for beneficial feature interaction detection.

3 Problem Formulation and Definitions

Consider a dataset with input-output pairs: $D = \{(X_n, y_n)\}_{1 \leq n \leq N}$, where $y_n \in \mathbb{R}/\mathbb{Z}$, $X_n = \{c_k : x_k\}_{k \in J_n}$ is a set of categorical features (c) with their values (x), $J_n \subseteq J$ and J is an index set of all features in D . For example, in app recommendation, X_n consists of a user ID, an app ID and context features (e.g., *Cloudy*, *Monday*) with values to be 1 (i.e., recorded in this data sample), and y_n is a binary value to indicate whether the user will use this app. Our goal is to design a model $F(X_n)$ that detects the most beneficial *pairwise*¹ feature interactions and utilizes only the detected feature interactions to predict the true output y_n .

Beneficial Pairwise Feature Interactions Inspired by the definition of relevant feature by usefulness (Langley et al. 1994; Blum and Langley 1997), we formally define the beneficial feature interactions in Definition 1.

Definition 1. (Beneficial Pairwise Feature Interactions) Given a data sample $X = \{x_i\}_{1 \leq i \leq k}$ of k features whose corresponding full pairwise feature interaction set is $A = \{(x_i, x_j)\}_{1 \leq i, j \leq k}$, a set of pairwise feature interactions $I_1 \subseteq A$ is more beneficial than another set of pairwise feature interactions $I_2 \subseteq A$ to a model with X as input if the accuracy of the predictions that the model produces using I_1 is higher than the accuracy achieved using I_2 .

¹We focus on pairwise feature interactions in this paper, and we leave high-order feature interactions in future work.

The above definition formalizes our detection goal: find and retain only a part of feature interactions that together produce the highest prediction accuracy by our model.

Statistical Interaction Statistical interaction, or non-additive interaction, ensures a joint influence of several variables on an output variable is not additive (Tsang et al. 2018). Sorokina et al. (2008) formally define the pairwise statistical interaction:

Definition 2. (Pairwise Statistical Interaction) Function $F(X)$, where $X = \{x_i\}_{1 \leq i \leq k}$ has k variables, shows **no pairwise statistical interaction between variables x_i and x_j** if $F(X)$ can be expressed as the sum of two functions $f_{\setminus i}$ and $f_{\setminus j}$, where $f_{\setminus i}$ does not depend on x_i and $f_{\setminus j}$ does not depend on x_j :

$$F(X) = f_{\setminus i}(x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_k) + f_{\setminus j}(x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_k). \quad (1)$$

More generally, if using $\mathbf{v}_i \in \mathbb{R}^d$ to describe the i -th variable with d factors (Rendle 2010), e.g., variable embedding, each variable can be described in a vector form $\mathbf{u}_i = x_i \mathbf{v}_i$. Then, we define the pairwise statistical interaction in variable factor form by changing the Equation 1 into:

$$F(X) = f_{\setminus i}(\mathbf{u}_1, \dots, \mathbf{u}_{i-1}, \mathbf{u}_{i+1}, \dots, \mathbf{u}_k) + f_{\setminus j}(\mathbf{u}_1, \dots, \mathbf{u}_{j-1}, \mathbf{u}_{j+1}, \dots, \mathbf{u}_k). \quad (2)$$

The definition indicates that the interaction information of variables (features) x_i and x_j will not be captured by $F(X)$ if it can be expressed as the above equations. Therefore, to correctly capture the interaction information, our model should not be expressed as the above equations. In this paper, we theoretically prove that the interaction modeling in our model strictly follow the definition of pairwise statistical interaction, which ensures our model to correctly capture interaction information.

4 Our Proposed Model

In this section, we formally describe L_0 -SIGN’s overview structure and the two components in detail. Then, we provide theoretical analyses of our model.

L_0 -SIGN

Model Overview Each input of L_0 -SIGN is represented as a graph (without edge information), where its features are nodes and their interactions are edges. More specifically, a data sample n is a graph $G_n(X_n, E_n)$, and $E_n = \{(e_n)_{ij}\}_{i,j \in X_n}$ is a set of edge/interaction values², where $(e_n)_{ij} \in \{1, 0\}$, 1 indicates that there is an edge (beneficial feature interaction) between nodes i and j , and 0 otherwise. Since no edge information are required, $E_n = \emptyset$.

While predicting, the L_0 edge prediction component, $F_{ep}(X_n; \omega)$, analyzes the existence of edges on each pair of nodes, where ω are parameters of F_{ep} , and outputs the predicted edge set E'_n . Then, the graph classification component, SIGN, performs predictions based on $G(X_n, E'_n)$.

²In this paper, nodes and features are used interchangeably, and the same as edges and feature interactions.




Figure 1: An overview of L_0 -SIGN.

Specifically, SIGN firstly conducts interaction modeling on each pair of initial node representation that are linked by an edge. Then, each node representation is updated by aggregating all of the corresponding modeling results. Finally, all updated node representations are aggregated to get the final prediction. The general form of SIGN prediction function is $y'_n = f_S(G_n(X_n, E'_n); \theta)$, where θ is SIGN’s parameters and the predicted outcome y'_n is the graph classification result. Therefore, the L_0 -SIGN prediction function f_{LS} is:

$$f_{LS}(G_n(X_n, \emptyset); \theta, \omega) = f_S(G_n(X_n, F_{ep}(X_n; \omega)); \theta). \quad (3)$$

Figure 1 shows the structure of L_0 -SIGN³. Next, we will show the two components in detail. When describing the two components, we focus on one input-output pair, so we omit the index “ n ” for simplicity.

L_0 Edge Prediction Model A (neural) matrix factorization (MF) based model is used for edge prediction. MF is effective in modeling relations between node pairs by factorizing the adjacency matrix of a graph into node dense embeddings (Menon and Elkan 2011). In L_0 -SIGN, since we do not have the ground truth adjacency matrix, the gradients for optimizing this component come from the errors between the outputs of SIGN and the target outcomes.

Specifically, the edge value, $e'_{ij} \in E'$, is predicted by an edge prediction function $f_{ep}(\mathbf{v}_i^e, \mathbf{v}_j^e) : \mathbb{R}^{2 \times b} \rightarrow \mathbb{Z}_2$, which takes a pair of node embeddings for edge prediction with dimension b as input, and output a binary value, indicating whether the two nodes are connected by an edge. $\mathbf{v}_i^e = \mathbf{o}_i \mathbf{W}^e$ is the embedding of node i for edge prediction, where $\mathbf{W}^e \in \mathbb{R}^{|J| \times b}$ are parameters and \mathbf{o}_i is the one-hot embedding of node i . To ensure that the detection results are identical to the same pair of nodes, f_{ep} should be invariant to the order of its input, i.e., $f_{ep}(\mathbf{v}_i^e, \mathbf{v}_j^e) = f_{ep}(\mathbf{v}_j^e, \mathbf{v}_i^e)$. For example, in our experiments, we use a multilayer neural network (MLP) with the input as the element-wise product result of \mathbf{v}_i^e and \mathbf{v}_j^e (details are in Section 5). Note that $e'_{ii} = 1$ can be regarded as the feature i being beneficial.

While training, an L_0 activation regularization is performed on E' to minimize the number of detected edges, which will be described later in the section about the empirical risk minimization function of L_0 -SIGN.

³Section D of Appendix lists the pseudocodes of our model and the training procedures.

SIGN In SIGN, each node i is first represented as an initial node embedding \mathbf{v}_i of d dimensions for interaction modeling, i.e., each node has node embeddings \mathbf{v}_i^e and \mathbf{v}_i for edge prediction and interaction modeling, respectively, to ensure the best performance of respective tasks. Then, the interaction modeling is performed on each node pair (i, j) that $e'_{ij} = 1$, by a non-additive function $h(\mathbf{u}_i, \mathbf{u}_j) : \mathbb{R}^{2 \times d} \rightarrow \mathbb{R}^d$ (e.g., an MLP), where $\mathbf{u}_i = x_i \mathbf{v}_i$. The output of $h(\mathbf{u}_i, \mathbf{u}_j)$ is denoted as \mathbf{z}_{ij} . Similar to f_{ep} , h should also be invariant to the order of its input. The above procedure can be reformulated as $\mathbf{s}_{ij} = e'_{ij} \mathbf{z}_{ij}$, where $\mathbf{s}_{ij} \in \mathbb{R}^d$ is called the statistical interaction analysis result of (i, j) .

Next, each node is updated by aggregating all of the analysis results between the node and its neighbors using a linear aggregation function ψ : $\mathbf{v}'_i = \psi(\zeta_i)$, where $\mathbf{v}'_i \in \mathbb{R}^d$ is the updated embedding of node i , ζ_i is a set of statistical interaction analysis results between node i and its neighbors. Note that ψ should be invariant to input permutations, and be able to take inputs with variant number of elements (e.g., element-wise summation/mean). While training, node embedding will be updated iteratively by replacing initial embedding with last updated embedding: $(\mathbf{v}_i)_t \leftarrow (\mathbf{v}'_i)_{t-1}$, where t is the iteration index while training.

Finally, each updated node embedding will be transformed into a scalar value by a linear function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, and all scalar values are linearly aggregated as the output of SIGN. That is: $y' = \phi(\nu)$, where $\nu = \{g(\mathbf{u}'_i) \mid i \in X\}$, $\mathbf{u}'_i = x_i \mathbf{v}'_i$ and $\phi : \mathbb{R}^{|\nu| \times 1} \rightarrow \mathbb{R}$ is an aggregation function having similar properties to ψ . Therefore, the prediction function of SIGN is:

$$f_S(G; \theta) = \phi(\{g(\psi(\{e'_{ij} h(\mathbf{u}_i, \mathbf{u}_j)\}_{j \in X}))\}_{i \in X}). \quad (4)$$

In summary, we formulate the L_0 -SIGN prediction function of Equation 3 with the two components in detail ⁴:

$$f_{LS}(G; \omega, \theta) = \phi(\{g(\psi(\{f_{ep}(\mathbf{v}_i^e, \mathbf{v}_j^e) h(\mathbf{u}_i, \mathbf{u}_j)\}_{j \in X}))\}_{i \in X}). \quad (5)$$

Empirical Risk Minimization Function

The empirical risk minimization function of L_0 -SIGN minimizes a loss function, a reparameterized L_0 activation regularization⁵ on E'_n , and an L_2 activation regularization on \mathbf{z}_n . Instead of regularizing parameters, activation regularization regularizes the output of models (Merity, McCann, and Socher 2017). We leverage activation regularization to link our model with the IB principle to ensure the success of the interaction detection, which will be discussed in the theoretical analyses. Formally, the function is:

$$\begin{aligned} \mathcal{R}(\theta, \omega) &= \frac{1}{N} \sum_{n=1}^N (\mathcal{L}(F_{LS}(G_n; \omega, \theta), y_n) \\ &\quad + \lambda_1 \sum_{i, j \in X_n} (\pi_n)_{ij} + \lambda_2 \|\mathbf{z}_n\|_2), \quad (6) \\ \theta^*, \omega^* &= \arg \min_{\theta, \omega} \mathcal{R}(\theta, \omega), \end{aligned}$$

⁴The time complexity analysis is in Section E of Appendix.

⁵Section F of Appendix gives detailed description about L_0 regularization and its reparameterization trick.




Figure 2: The interaction analysis results \mathbf{s} from L_0 -SIGN are like the relevant part (bottleneck) in the IB principle.

where $(\pi_n)_{ij}$ is the probability of $(e'_n)_{ij}$ being 1 (i.e., $(e'_n)_{ij} = \text{Bern}((\pi_n)_{ij})$), $G_n = G_n(X_n, \emptyset)$, λ_1 and λ_2 are weight factors for the regularizations, $\mathcal{L}(\cdot)$ corresponds to a loss function and θ^* , ω^* are final parameters.

A practical difficulty of performing L_0 regularization is that it is non-differentiable. Inspired by (Louizos, Welling, and Kingma 2018), we smooth the L_0 activation regularization by approximating the Bernoulli distribution with a hard concrete distribution⁶ so that e'_{ij} is differentiable.

Theoretical Analyses

We conduct theoretical analyses to verify our model's effectiveness, including how L_0 -SIGN satisfies the IB principle to guarantee the success of beneficial interaction detection, the relation of the statistical interaction analysis results with the spike-and-slab distribution (the golden standard in sparsity), and how SIGN and L_0 -SIGN strictly follow the statistical interaction theory for effective interaction modeling.

Satisfaction of Information Bottleneck (IB) Principle

IB principle (Tishby, Pereira, and Bialek 2000) aims to extract the most relevant information that input random variables \mathbf{X} contains about output variables \mathbf{Y} by considering a trade-off between the accuracy and complexity of the process. The relevant part of \mathbf{X} over \mathbf{Y} denotes \mathbf{S} . The IB principle can be mathematically represented as:

$$\min (I(\mathbf{X}; \mathbf{S}) - \beta I(\mathbf{S}; \mathbf{Y})), \quad (7)$$

where $I(\cdot)$ denotes mutual information between two variables and $\beta > 0$ is a weight factor.

The empirical risk minimization function of L_0 -SIGN in Equation 6 can be approximately derived from Equation 7 (Section A of Appendix gives a detailed derivation):

$$\min \mathcal{R}(\theta, \omega) \approx \min (I(\mathbf{X}; \mathbf{S}) - \beta I(\mathbf{S}; \mathbf{Y})). \quad (8)$$

Intuitively, the L_0 regularization in Equation 6 minimizes a Kullback–Leibler divergence between every e'_{ij} and a Bernoulli distribution $\text{Bern}(0)$, and the L_2 regularization minimizes a Kullback–Leibler divergence between every \mathbf{z}_{ij} and a multivariate standard distribution $\mathcal{N}(\mathbf{0}, \mathbf{I})$. As illustrated in Figure 2, the statistical interaction analysis results, \mathbf{s} , can be approximated as the relevant part \mathbf{S} in Equation 7.

Therefore, through training L_0 -SIGN, \mathbf{s} is the most compact (beneficial) representation of the interactions for recommendation prediction. This provides a theoretical guarantee that the most beneficial feature interactions will be detected.

⁶Section G of Appendix gives details about the approximation.

Relation to Spike-and-slab Distribution The spike-and-slab distribution (Mitchell and Beauchamp 1988) is the golden standard in sparsity. It is defined as a mixture of a delta spike at zero and a continuous distribution over the real line (e.g., a standard normal):

$$\begin{aligned} p(a) &= \text{Bern}(\pi), & p(\theta | a = 0) &= \delta(\theta), \\ p(\theta | a = 1) &= \mathcal{N}(\theta | 0, 1). \end{aligned} \quad (9)$$

We can regard the spike-and-slab distribution as the product of a continuous distribution and a Bernoulli distribution. In L_0 -SIGN, the predicted edge value vector e' (the vector form of E') is a multivariate Bernoulli distribution and can be regarded as $p(a)$. The interaction modeling result z is a multivariate normal distribution and can be regarded as $p(\theta)$. Therefore, L_0 -SIGN’s statistical interaction analysis results, s , is a multivariate spike-and-slab distribution that performs edge sparsification by discarding unbeneficial feature interactions. The retained edges in the spike-and-slab distribution are sufficient for L_0 -SIGN to provide accurate predictions.

Statistical Interaction in SIGN The feature interaction modeling in SIGN strictly follows the definition of statistical interaction, which is formally described in Theorem 1 (Section B of Appendix gives the proof):

Theorem 1. (Statistical Interaction in SIGN) Consider a graph $G(X, E)$, where X is the node set and $E = \{e_{ij}\}_{i,j \in X}$ is the edge set that $e_{ij} \in \{0, 1\}$, $e_{ij} = e_{ji}$. Let $G(X, E)$ be the input of SIGN function $f_S(G)$ in Equation 4, then the function flags pairwise statistical interaction between node i and node j if and only if they are linked by an edge in $G(X, E)$, i.e., $e_{ij} = 1$.

Theorem 1 guarantees that SIGN will capture the interaction information from node pairs *iff* they are linked by an edge. This ensures SIGN to accurately leverage the detected beneficial feature interactions for both inferring the target outcome and meanwhile providing useful feedback to the L_0 edge prediction component for better detection.

Statistical Interaction in L_0 -SIGN L_0 -SIGN provides the same feature interaction modeling ability as SIGN, since we can simply extend Theorem 1 to Corollary 1.1 (Section C of Appendix gives the proof):

Corollary 1.1. (Statistical Interaction in L_0 -SIGN) Consider a graph G that the edge set is unknown. Let G be the input of L_0 -SIGN function $F_{LS}(G)$ in Equation 5, the function flags pairwise statistical interaction between node i and node j if and only if they are predicted to be linked by an edge in G by L_0 -SIGN, i.e., $e'_{ij} = 1$.

5 Experiments

We focus on answering three questions: (i) how L_0 -SIGN performs compared to baselines and whether SIGN helps detect more beneficial interactions for better performance? (ii) How is the detection ability of L_0 -SIGN? (iii) Can the statistical interaction analysis results provide potential explanations for the recommendations predictions?

Table 1: Dataset statistics. All datasets are denoted in graph form. Each node represents one feature. Twitter and DBLP datasets are used for answering question (ii).

DATASET	#FEATURES	#GRAPHS	#NODES/GRAPH
FRAPPE	5,382	288,609	10
MOVIELENS	90,445	2,006,859	3
TWITTER	1,323	144,033	4.03
DBLP	41,324	19,456	10.48

Experimental Protocol

Datasets We study two real-world datasets for recommender systems to evaluate our model:

Frappe (Baltrunas et al. 2015) is a context-aware recommendation dataset that records app usage logs from different users with eight types of contexts (e.g, weather). Each log is a graph (without edges), and nodes are user ID, app ID, or the contexts.

MovieLens-tag (He and Chua 2017) focuses on the movie tag recommendation (e.g., “must-see”). Each data instance is a graph, with nodes as user ID, movie ID, and a tag that the user gives to the movie.

To evaluate the question (ii), we further study two datasets for graph classification, which will be discussed later. The statistics of the datasets are summarized in Table 1.

Baselines We compare our model with recommender system baselines that model all feature interactions:

FM (Koren 2008): It is one of the most popular recommendation algorithms that models every feature interactions. **AFM** (Xiao et al. 2017): Addition to FM, it calculates an attention value for each feature interaction. **NFM** (He and Chua 2017): It replaces the dot product procedure of FM by an MLP. **DeepFM** (Guo et al. 2017): It uses MLP and FM for interaction analysis, respectively. **xDeepFM** (Lian et al. 2018): It is an extension of DeepFM that models feature interactions in both explicit and implicit way. **AutoInt** (Song et al. 2019): It explicitly models all feature interactions using a multi-head self-attentive neural network. We use the same MLP settings in all baselines (if use) as our interaction modeling function h in SIGN for fair comparison.

Experimental Set-up In the experiments, we use element-wise mean as both linear aggregation functions $\psi(\cdot)$ and $\phi(\cdot)$. The linear function $g(\cdot)$ is a weighted sum function (i.e., $g(\mathbf{u}'_i) = \mathbf{w}_g^T \mathbf{u}'_i$, where $\mathbf{w}_g \in \mathbb{R}^{d \times 1}$ are the weight parameters). For the interaction modeling function $h(\cdot)$, we use a MLP with one hidden layer after element-wise product: $h(\mathbf{u}_i, \mathbf{u}_j) = \mathbf{W}_2^h \sigma(\mathbf{W}_1^h(\mathbf{u}_i \odot \mathbf{u}_j) + \mathbf{b}_1^h) + \mathbf{b}_2^h$, where $\mathbf{W}_1^h, \mathbf{W}_2^h, \mathbf{b}_1^h, \mathbf{b}_2^h$ are parameters of MLP and $\sigma(\cdot)$ is a Relu activation function. We implement the edge prediction model based on the neural collaborative filtering framework (He et al. 2017), which has a similar form to $h(\cdot)$: $f_{ep}(\mathbf{v}_i^e, \mathbf{v}_j^e) = \mathbf{W}_2^e \sigma(\mathbf{W}_1^e(\mathbf{v}_i^e \odot \mathbf{v}_j^e) + \mathbf{b}_1^e) + \mathbf{b}_2^e$. We set node embedding sizes for both interaction modeling and edge prediction to 8 (i.e., $b, d = 8$) and the sizes of hidden layer for both h and f_{ep} to 32. We choose the weighting factors λ_1 and λ_2 from $[1 \times 10^{-5}, 1 \times 10^{-1}]$ that produce the best performance in each dataset.

Table 2: Summary of results in comparison with baselines.

	FRAPPE		MOVIELENS	
	AUC	ACC	AUC	ACC
FM	0.9263	0.8729	0.9190	0.8694
AFM	0.9361	0.8882	0.9205	0.8711
NFM	0.9413	0.8928	0.9342	0.8903
DEEPM	0.9422	0.8931	0.9339	0.8895
xDEEPM	0.9435	0.8950	0.9347	0.8906
AUTOINT	0.9432	0.8947	0.9351	0.8912
SIGN	0.9448	0.8974	0.9354	0.8921
L_0 -SIGN	0.9580	0.9174	0.9407	0.8970

Each dataset is randomly split into training, validation, and test datasets with a proportion of 70%, 15%, and 15%. We choose the model parameters that produce the best results in validation set when the number of predicted edges being steady. We use accuracy (ACC) and the area under a curve with Riemann sums (AUC) as evaluation metrics.

Model Performance

Table 2 shows the results of comparing our model with recommendation baselines, with the best results for each dataset in bold. The results of SIGN are using all feature interactions as input, i.e., the input is a complete graph.

Through the table, we observe that: (i) L_0 -SIGN outperforms all baselines. It shows L_0 -SIGN’s ability in providing accurate recommendations. Meanwhile, SIGN solely gains comparable results to competitive baselines, which shows the effectiveness of SIGN in modeling feature interactions for recommendation. (ii) L_0 -SIGN gains significant improvement from SIGN. It shows that more accurate predictions can be delivered by retaining only beneficial feature interactions and effectively modeling them. (iii) FM and AFM (purely based on dot product to model interactions) gain lower accuracy than other models, which shows the necessity of using sophisticated methods (e.g., MLP) to model feature interactions for better predictions. (iv) The models that explicitly model feature interactions (xDeepFM, AutoInt, L_0 -SIGN) outperform those that implicitly model them (NFM, DeepFM). It shows that explicit feature interaction analysis is promising in delivering accurate predictions.

Comparing SIGN with Other GNNs in Our Model

To evaluate whether our L_0 edge prediction technique can be used on other GNNs and whether SIGN is more suitable than other GNNs in our model, we replace SIGN with existing GNNs: GCN (Kipf and Welling 2017), Chebyshev filter based GCN (Cheby) (Defferrard, Bresson, and Vandergheynst 2016) and GIN (Xu et al. 2019). We run on two datasets for graph classification since they contain heuristic edges (used to compare with the predicted edges):

Twitter (Pan, Wu, and Zhu 2015) is extracted from twitter sentiment classification. Each tweet is a graph with nodes being word tokens and edges being the co-occurrence between two tokens in each tweet.

DBLP (Pan et al. 2013) consists of papers with labels indicating either they are from DBDM or CVPR field. Each

Table 3: The results in comparison with existing GNNs. The model names without “ L_0 -” use heuristic edges, and those with “ L_0 -” automatically detect edges via our L_0 edge prediction technique.

	TWITTER		DBLP	
	AUC	ACC	AUC	ACC
GCN	0.7049	0.6537	0.9719	0.9289
L_0 -GCN	0.7053	0.6543	0.9731	0.9301
CHEBY	0.7076	0.6522	0.9717	0.9291
L_0 -CHEBY	0.7079	0.6519	0.9719	0.9297
GIN	0.7149	0.6559	0.9764	0.9319
L_0 -GIN	0.7159	0.6572	0.9787	0.9328
SIGN	0.7201	0.6615	0.9761	0.9316
L_0 -SIGN	0.7231	0.6670	0.9836	0.9427




Figure 3: The comparison of accuracy improvement via the L_0 edge prediction technique to using heuristic edges in Table 3.

paper is a graph with nodes being paper ID or keywords and edges being the citation relationship or keyword relations.

Table 3 shows the accuracies of each GNN that runs both on using the given (heuristic) edges (without “ L_0 ” in name) and on predicting edges with our L_0 edge prediction model (with “ L_0 ” in name). We can see that for all GNNs, L_0 -GNNs gain competitive results comparing to corresponding GNNs. It shows that our model framework lifts the requirement of domain knowledge in defining edges in order to use GNNs in some situations. Also, L_0 -SIGN outperforms other GNNs and L_0 -GNNs, which shows L_0 -SIGN’s ability in detecting beneficial feature interactions and leveraging them to perform accurate predictions. Figure 3 shows each GNN’s improvement from the results from “GNN” to “ L_0 -GNN” in Table 3. It shows that among all the different GNN based models, SIGN gains the largest improvement from the L_0 edge prediction technique v.s. SIGN without L_0 edge prediction. It shows that SIGN can help the L_0 edge prediction model to better detect beneficial feature interactions for more accurate predictions.

Evaluation of Interaction Detection

We then evaluate the effectiveness of beneficial feature interaction detection in L_0 -SIGN.

Prediction Accuracy vs. Numbers of Edges Figure 4 shows the changes in prediction accuracy and the number of edges included while training. The accuracy first increases while the number of included edges decreases dramatically. Then, the number of edges becomes steady, and the accuracy reaches a peak at similar epochs. It shows that our model can recognize unbeneficial feature interactions and remove them for better prediction accuracy.

Predicted Edges vs. Reversed Edges We evaluate how predicted edges (the retained edges) and reversed edges (the




Figure 4: The changes in prediction accuracy and number of edges (in percentage) while training.





Figure 5: Evaluating different number of edges. “Pred” is the predicted edges and “Rev” is the reversed edges.

excluded edges) influence the performance. Specifically, we generate 5 edge sets with a different number of edges by randomly selecting from 20% predicted edges (ratio 0.2) to 100% predicted edges (ratio 1.0). We generate another 5 edge sets similarly from reversed edges. Then, we run SIGN on each edge set for 5 times, and show the results in average in Figure 5. It shows that using predicted edges always gets higher accuracy than using reversed edges. Meanwhile, the accuracy stops increasing when using reversed edges since the ratio 0.6, while it continually improves when using more predicted edges. According to Definition 1, the detected edges are proved beneficial since considering more of them can provide further performance gain and the performance is the best when using all predicted edges, while considering the reversed edges cannot. Note that the increment of reversed edges from 0.2 to 0.6 may come from covering more nodes since features solely can provide some useful information for prediction.



$$\text{Prediction: } 0.636 + 0.3626 + 0.3161 \dots - 0.1013 - 0.1159 = 1.313 > 0$$




Figure 6: (a) Highest beneficial interaction values with *Gmail*. The number features are city indexes (e.g., 1014). (b) The prediction that the user *u349* will use *Gmail* as the prediction value is $1.313 > 0$. Darker edges are higher interaction values.

Case Study

Another advantage of interaction detection is to automatically discover potential explanations about recommendation predictions. We conduct case studies on the Frappe dataset to show how L_0 -SIGN provides the explanations.

We first show the beneficial interactions that have the highest interaction values with *Gmail* in Figure 6a. We can see that *Workday* has a high positive value, while *Home* has a high negative value. It may show that *Gmail* is very likely to be used in workdays since people need to use *Gmail* while working, whereas *Gmail* is not likely to be used at home since people may not want to deal with email while resting. We then show how L_0 -SIGN provides potential explanations for the prediction on each data sample. Figure 6b visualizes a prediction result that a user (*u349*) may use *Gmail*. The graph shows useful information for explanations. Despite the beneficial interactions such as $\langle \textit{Gmail}, \textit{Workday} \rangle$ that have discussed above, we can also see that *Cloudy* and *Morning* have no beneficial interaction, meaning that whether it is a cloudy morning does not contribute to decide the user’s will of using *Gmail*.

6 Conclusion and Future Work

This is the first work to propose and formulate the problem of detecting beneficial feature interactions for recommender systems. We propose a model, named L_0 -SIGN, to detect the beneficial feature interactions via a graph neural network approach and L_0 regularization. Theoretical analyses and extensive experiments show the ability of L_0 -SIGN in detecting and modeling beneficial feature interactions for accurate recommendations. In future work, we will model high-order feature interactions in GNNs with theoretical foundations to effectively capture high-order feature interaction information in graph form.

Acknowledgments

This work is supported by the China Scholarship Council.

References

- Alemi, A. A.; Fischer, I.; Dillon, J. V.; and Murphy, K. 2017. Deep Variational Information Bottleneck. In *ICLR*, 1–16.
- Baltrunas, L.; Church, K.; Karatzoglou, A.; and Oliver, N. 2015. Frappe: Understanding the usage and perception of mobile app recommendations in-the-wild. *arXiv preprint arXiv:1505.03014*.
- Battaglia, P. W.; Hamrick, J. B.; Bapst, V.; Sanchez-Gonzalez, A.; Zambaldi, V.; Malinowski, M.; Tacchetti, A.; Raposo, D.; Santoro, A.; and Faulkner, R. 2018. Relational inductive biases, deep learning, and graph networks. *arXiv preprint arXiv:1806.01261*.
- Blondel, M.; Ishihata, M.; Fujino, A.; and Ueda, N. 2016. Polynomial Networks and Factorization Machines: New Insights and Efficient Training Algorithms. In *ICML*, 850–858.
- Blum, A. L.; and Langley, P. 1997. Selection of Relevant Features and Examples in Machine Learning. *AI* 97(1-2): 245–271.
- Chang, M. B.; Ullman, T.; Torralba, A.; and Tenenbaum, J. B. 2016. A compositional object-based approach to learning physical dynamics. In *ICLR*, 1–14.
- Defferrard, M.; Bresson, X.; and Vandergheynst, P. 2016. Convolutional Neural Networks on Graphs with Fast Localized Spectral Filtering. In *NeurIPS*, 3844–3852.
- Duvenaud, D. K.; Maclaurin, D.; Iparraguirre, J.; Bombarell, R.; Hirzel, T.; Aspuru-Guzik, A.; and Adams, R. P. 2015. Convolutional networks on graphs for learning molecular fingerprints. In *NeurIPS*, 2224–2232.
- Gilmer, J.; Schoenholz, S. S.; Riley, P. F.; Vinyals, O.; and Dahl, G. E. 2017. Neural message passing for quantum chemistry. In *ICML*, 1263–1272.
- Guo, H.; Tang, R.; Ye, Y.; Li, Z.; and He, X. 2017. Deepfm: a factorization-machine based neural network for ctr prediction. In *IJCAI*, 1725–1731.
- He, X.; and Chua, T.-S. 2017. Neural factorization machines for sparse predictive analytics. In *SIGIR*, 355–364.
- He, X.; Liao, L.; Zhang, H.; Nie, L.; Hu, X.; and Chua, T.-S. 2017. Neural Collaborative Filtering. In *WWW*, 173–182.
- Kipf, T. N.; and Welling, M. 2017. Semi-supervised classification with graph convolutional networks. In *ICLR*, 1–14.
- Koren, Y. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *SIGKDD*, 426–434.
- Langley, P.; et al. 1994. Selection of Relevant Features in Machine Learning. In *AAAI*, volume 184, 245–271.
- Li, Z.; Cui, Z.; Wu, S.; Zhang, X.; and Wang, L. 2019. Fignn: Modeling Feature Interactions via Graph Neural Networks for CTR Prediction. In *CIKM*, 539–548.
- Lian, J.; Zhou, X.; Zhang, F.; Chen, Z.; Xie, X.; and Sun, G. 2018. xdeepfm: Combining Explicit and Implicit Feature Interactions for Recommender Systems. In *SIGKDD*, 1754–1763.
- Liu, B.; Zhu, C.; Li, G.; Zhang, W.; Lai, J.; Tang, R.; He, X.; Li, Z.; and Yu, Y. 2020. AutoFIS: Automatic Feature Interaction Selection in Factorization Models for Click-Through Rate Prediction. In *SIGKDD*, 2636–2645. ACM.
- Louizos, C.; Welling, M.; and Kingma, D. P. 2018. Learning Sparse Neural Networks through L_0 Regularization. In *ICLR*, 1–11.
- Luo, Y.; Wang, M.; Zhou, H.; Yao, Q.; Tu, W.-W.; Chen, Y.; Dai, W.; and Yang, Q. 2019. Autocross: Automatic Feature Crossing for Tabular Data in Real-world Applications. In *SIGKDD*, 1936–1945.
- MacKay, D. J. 2003. *Information Theory, Inference and Learning Algorithms*. Cambridge university press.
- Maddison, C. J.; Mnih, A.; and Teh, Y. W. 2017. The Concrete Distribution: A Continuous Relaxation of Discrete Random Variables. In *ICLR*, 1–12.
- Menon, A. K.; and Elkan, C. 2011. Link Prediction via Matrix Factorization. In *ECML PKDD*, 437–452.
- Merity, S.; McCann, B.; and Socher, R. 2017. Revisiting Activation Regularization for Language RNNs. In *ICML Workshop*.
- Mitchell, T. J.; and Beauchamp, J. J. 1988. Bayesian variable selection in linear regression. *ASA* 83(404): 1023–1032.
- Pan, S.; Wu, J.; and Zhu, X. 2015. CogBoost: Boosting for Fast Cost-sensitive Graph Classification. In *TKDE*, 2933–2946.
- Pan, S.; Zhu, X.; Zhang, C.; and Philip, S. Y. 2013. Graph Stream Classification Using Labeled and Unlabeled Graphs. In *ICDE*, 398–409.
- Rendle, S. 2010. Factorization Machines. In *ICDM*, 995–1000.
- Santoro, A.; Raposo, D.; Barrett, D. G.; Malinowski, M.; Pascanu, R.; Battaglia, P.; and Lillicrap, T. 2017. A simple neural network module for relational reasoning. In *NeurIPS*, 4967–4976.
- Shi, C.; Glocker, B.; and Castro, D. C. 2019. PVAE: Learning Disentangled Representations with Intrinsic Dimension via Approximated L_0 Regularization. In *PMLR*, 1–6.
- Shi, Y.; Larson, M.; and Hanjalic, A. 2014. Collaborative Filtering Beyond the User-item Matrix: A Survey of the State of the Art and Future Challenges. *CSUR* 47(1): 1–45.
- Siegmund, N.; Kolesnikov, S. S.; Kästner, C.; Apel, S.; Batory, D.; Rosenmüller, M.; and Saake, G. 2012. Predicting Performance via Automated Feature-interaction Detection. In *ICSE*, 167–177.
- Song, W.; Shi, C.; Xiao, Z.; Duan, Z.; Xu, Y.; Zhang, M.; and Tang, J. 2019. AutoInt: Automatic Feature Interaction Learning via Self-attentive Neural Networks. In *CIKM*, 1161–1170.

- Sorokina, D.; Caruana, R.; Riedewald, M.; and Fink, D. 2008. Detecting Statistical Interactions with Additive Groves of Trees. In *ICML*, 1000–1007.
- Tishby, N.; Pereira, F. C.; and Bialek, W. 2000. The Information Bottleneck Method. *arXiv preprint physics/0004057*.
- Tsang, M.; Liu, H.; Purushotham, S.; Murali, P.; and Liu, Y. 2018. Neural Interaction Transparency (NIT): Disentangling Learned Interactions for Improved Interpretability. In *NeurIPS*, 5804–5813.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *CVPR*, 7794–7803.
- Xiao, J.; Ye, H.; He, X.; Zhang, H.; Wu, F.; and Chua, T.-S. 2017. Attentional factorization machines: Learning the weight of feature interactions via attention networks. In *IJCAI*, 3119–3125.
- Xu, K.; Hu, W.; Leskovec, J.; and Jegelka, S. 2019. How Powerful are Graph Neural Networks? *ICLR* 1–13.
- Yang, C.; Bai, L.; Zhang, C.; Yuan, Q.; and Han, J. 2017. Bridging collaborative filtering and semi-supervised learning: a neural approach for poi recommendation. In *SIGKDD*, 1245–1254.
- Zhang, C.; Bengio, S.; Hardt, M.; Recht, B.; and Vinyals, O. 2017. Understanding Deep Learning Requires Rethinking Generalization. In *ICLR*, 1–11.