



Graph Neural Network with Two Uplift Estimators for Label-Scarcity Individual Uplift Modeling

Dingyuan Zhu* Daixin Wang* Zhiqiang Zhang Kun Kuang
 Ant Group, China Ant Group, China Ant Group, China Zhejiang University, China
 dingyuan.zdy@antgroup.com daixin.wdx@antgroup.com lingyao.zzq@antgroup.com kunkuang@zju.edu.cn

Yan Zhang Yulin Kang Jun Zhou[†]
 Ant Group, China Ant Group, China Ant Group, China
 zy118409@antgroup.com yulin.kyl@antgroup.com jun.zhoujun@antgroup.com

ABSTRACT

Uplift modeling aims to measure the incremental effect, which we call uplift, of a strategy or action on the users from randomized experiments or observational data. Most existing uplift methods only use individual data, which are usually not informative enough to capture the unobserved and complex hidden factors regarding the uplift. Furthermore, uplift modeling scenario usually has scarce labeled data, especially for the treatment group, which also poses a great challenge for model training. Considering that the neighbors' features and the social relationships are very informative to characterize a user's uplift, we propose a graph neural network-based framework with two uplift estimators, called GNUM, to learn from the social graph for uplift estimation. Specifically, we design the first estimator based on a class-transformed target. The estimator is general for all types of outcomes, and is able to comprehensively model the treatment and control group data together to approach the uplift. When the outcome is discrete, we further design the other uplift estimator based on our defined partial labels, which is able to utilize more labeled data from both the treatment and control groups, to further alleviate the label scarcity problem. Comprehensive experiments on a public dataset and two industrial datasets show a superior performance of our proposed framework over state-of-the-art methods under various evaluation metrics. The proposed algorithms have been deployed online to serve real-world uplift estimation scenarios.

CCS CONCEPTS

• **Mathematics of computing** → **Graph algorithms.**

KEYWORDS

uplift modeling, graph neural networks, partial label

*Equal contribution.

[†]Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](https://permissions.acm.org).

WWW '23, April 30–May 04, 2023, Austin, TX, USA

© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9416-1/23/04...\$15.00

<https://doi.org/10.1145/3543507.3583368>

ACM Reference Format:

Dingyuan Zhu, Daixin Wang, Zhiqiang Zhang, Kun Kuang, Yan Zhang, Yulin Kang, and Jun Zhou. 2023. Graph Neural Network with Two Uplift Estimators for Label-Scarcity Individual Uplift Modeling. In *Proceedings of the ACM Web Conference 2023 (WWW '23), April 30–May 04, 2023, Austin, TX, USA*. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3543507.3583368>

1 INTRODUCTION

Uplift modeling refers to the set of techniques used to estimate the effect of an action on a user's outcome. This technology can be applied to various fields, such as economics[14, 36], medicine[23, 24] and sociology[17, 37]. For example, an e-commerce company is preparing to send promotional coupons (i.e., action or treatment) to some users to attract them to purchase more products (i.e., outcome). Then estimating the individual uplift can help to find target users. Uplift modeling is a complicated problem because one needs to estimate the difference between two outcomes with and without the treatment which are mutually exclusive to individuals. For example, we can only observe the outcome of a user getting or not getting a promotional coupon. The outcome we can observe is called the factual outcome and the outcome we can not observe is called the counterfactual outcome. Therefore, uplift modeling also can be seen as a counterfactual inference problem[19, 29].

To model the counterfactual outcome, existing uplift methods mainly rely on randomized experiments or observational data: users will be assigned to either the treatment group or the control group and we can observe only one type of outcome. Based on these data, most existing methods [16, 19, 40] utilize user's individual features to estimate the user uplift [2, 5]. An accurate uplift modeling is required to capture the complex and the unobserved hidden factors (representations) within the user. However, many factors regarding the user's uplift are difficult to be captured by only using individual data for two reasons. Firstly, in real-world scenarios, the individual features of users, especially the new users, are missing. Secondly, some informative information for uplift estimation is hidden and difficult to be characterized by individual features, like social status and personality.

To resolve the aforementioned problems, we hope to introduce the social graph. On the one hand, considering that users with close social relations usually have similar behaviors and preferences, we can utilize the information from social neighbors as a supplement to the user's own features. On the other hand, social graphs can reveal some informative social information like social status, which

is also beneficial for uplift modeling [45]. From the experiments, we also demonstrate that the uplift difference between users with friend relationships is much smaller than the difference between random users. Therefore, it is very essential to incorporate social graphs when doing uplift estimation.

A few of the existing works propose graph-based uplift models for uplift estimation [17, 18, 45]: They first use graph neural network (GNN) model to learn the graph-based representations from the social graphs. Then with the graph-based representations, these works design uplift estimators, which consist of two separate pathways of models to predict the outcome with and without the treatment using the data of the treatment group and control group respectively. Although these works have achieved substantial improvements, they still have limitations. In real-world scenarios, since the effect of the treatment is unpredictable, imposing the treatment on samples may bring a bad effect like capital loss and user loss. In this way, it is common that only a little flow will be assigned to the treatment group. In this way, the number of labeled instances, particularly the instances from the treatment group, is commonly limited, especially in the scenario of randomized experiments we mainly focus on. Moreover, to model the relational data, graph-based models require learning more parameters. From Figure 4.3.3, compared with tree-based and NN-based models, the performance of existing graph-based models will drop more quickly when the labeled data is scarce, thereby labeled data scarcity is a more severe and unresolved problem for graph-based representation learning in uplift estimation.

To address aforementioned problems, we propose a general GNN-based framework with two uplift estimators for uplift modeling. First of all, a GNN-based model with breadth and depth aggregators is proposed to generate the graph-based representations for the following uplift estimation. And empirically we demonstrate the proposed uplift estimation framework is general for many existing GNN-based models. Furthermore, to address the label scarcity problem, we propose two uplift estimators to utilize the treatment and control group data and their corresponding social relations. Firstly, we design a class-transformed target, which we prove is equal to the uplift and is general for uplift scenarios. Unlike two separate pathways of estimators, our transformed target is able to utilize the training instances with and without the treatment simultaneously in a common model to approach the uplift. Furthermore, when the outcome is discrete, we design partial labels based on the user's treatment and the observed outcome. Then we design the other uplift estimator with two classifiers to learn partial labels. The two classifiers can focus on different facets of the uplift but all require two groups of data for training. In this way, our model can capture the relations between two groups of data and utilize the labeled data more effectively. Experimental results show our proposed framework can outperform the best-performing baseline method by an improvement of 5% to 10% in the regression setting and 12% to 25% in the classification setting.

The main contributions of the paper are summarized here:

- **Problem:** We point out the label scarcity problem is severe for uplift modeling, especially when trained with the graph-based model. To the best of our knowledge, it is the first work trying to solve the label scarcity problem in uplift estimation.
- **Methodology:** We introduce a novel GNN-based framework with two uplift estimators for uplift modeling, which can

utilize the two groups of data and the social graph comprehensively. Specifically, when the outcome is discrete, we are the first to introduce partial label learning to uplift estimation, which is able to utilize the labeled data more effectively to alleviate the labeled data scarcity problem.

- **Results:** Extensive experiments on a public dataset and two industrial datasets demonstrate the superiority of proposed GNN-based uplift model(GNUM) on different types of outcomes. Specifically, we find labeled data scarcity is indeed a serious problem for previous graph-based uplift methods and our methods are robust to the label scarcity problem.

2 RELATED WORK

2.1 Uplift modeling

Existing uplift methods can be classified into three categories [19], i.e. the Two-Model methods [35, 38], the Class-Transformation methods [24] and the methods that model uplift directly [39, 40, 43].

The Two-Model methods construct two independent models for the two groups of data. One model infers the label using the data from the treatment group and the other model is for the control group data. However, [39] points out that the Two-Model methods may miss the uplift signal. Then The Class Transformation methods are introduced by [1, 24], which aim to create a new target to approach the uplift. Then a single model is proposed to learn the new target. But Class Transformation methods usually require a balanced dataset between the control and treatment groups. The last type of uplift method aims to directly infer the uplift. The work [48] proposes a method based on a modification of the SVM model and the work [15] focuses on k-nearest neighbors to do the uplift estimation.

The aforementioned works assume that the data of the control group and treatment group are randomly collected. If the collecting data is naturally observed, besides uplift estimation, uplift modeling also needs to reduce the bias of the data from the treatment group and control group. The most popular methods of this type are the doubly robust learning methods [6, 9, 26, 54]. They usually adopt the Inverse Propensity Scoring (IPS) [21, 28] to re-weight each instance, aiming at making the uplift estimator unbiased. And some methods[32, 41] which estimate individual treatment effect can be used to estimate the uplift.

However, aforementioned uplift methods assume that the uplift can be fully estimated by the individual features. As we have stated before, the social relationships between users are important for uplift estimation. Then Guo et. al.[17, 18, 33] first introduce the networked observational data to the problem of causal effects estimation. NetEst[25] formalize the networked causal effects estimation to a multi-task learning problem and HyperSCI[34] learning causal effects on hypergraphs. These methods prove that networked data is important for predicting causal effects. Although incorporating the graph data will violate the Stable Unit Treatment Value Assumption(SUTVA), following works [42] have pointed out that SUTVA is not plausible in real-world scenarios. We will follow these works [17, 18] to incorporate the graph data but are distinct from them by addressing the label scarcity problem for graph-based uplift estimation.

2.2 Partial Label Learning

Partial label learning deals with the problem that each sample is associated with a set of candidate labels, among which only one label is the ground-truth label to be predicted. Existing partial label learning methods can be classified into two categories: the averaged-based strategy and the identification-based strategy. The averaged-based strategy assumes that each candidate label contributes equally to the model training [13, 22], which may suffer from the problem that the real label is overwhelmed by other labels. To overcome this drawback, the identification-based methods give different confidence to different candidate labels by learning the topological information [50, 51]. Existing partial-label learning is often applied to automatic face naming, object detection, and web mining. As far as we know, this is the first work to apply partial label learning to uplift modeling.

Other typical types of weakly-supervised learning include incomplete label learning[53] and inaccurate label learning[46]. They are not closely related to our problem, which will not be discussed here.

2.3 Graph Neural Network

Graph neural networks (GNNs), aiming to generalize neural networks to deal with graph data, have drawn increasing research interest recently [3, 52]. Generally, current GNNs can be divided into two categories: spectral-based methods and spatial-based methods. Spectral-based GNNs are originated from signal processing and are commonly based on the Laplacian Matrix [4, 7, 27]. Spatial-based GNNs regard the graph convolution as the 'message-passing' framework in the spatial domain, i.e. defining the graph convolution as nodes aggregating information from neighborhoods [11, 20]. And [8, 10, 44] have explored the causal inference problems with GNN-based models. More GNN-based models can be referred in recent surveys [47, 52]. However, only a few works focus on uplift estimation using GNNs and they do not address some specific and important problems when using GNNs on uplift estimation.

3 MODEL FORMULATION

3.1 Notations and Preliminaries

Firstly, we describe the notations used in this paper. We denote a scalar with a letter (e.g., t), a vector with a boldface lowercase letter (e.g., \mathbf{x}), and a matrix with a boldface uppercase letter (e.g., \mathbf{A}).

In our uplift estimation problem, we assume there are N users in total. The data of each user i can be represented as $\{\mathbf{x}_i, t_i, y_i\}$, where $\mathbf{x}_i \in \mathbb{R}^d$ represents the individual feature vector, $t_i \in \{0, 1\}$ denotes the observed treatment and y_i denotes the observed outcome. Note that $y_i(1)$ denotes the outcome of user i when he receives the active treatment, and $y_i(0)$ denotes its outcome with the control treatment. In this paper, we focus on the scenario of a randomized experiment, which means that each user is randomly given the treatment or not. Then, the actual uplift of user i is defined as:

$$\tau_i = y_i(1) - y_i(0), \quad (1)$$

and our target in this paper is to estimate the uplift of each user.

However, for a specific user i , we can only observe $y_i(1)$ or $y_i(0)$. The one we can observe is called the factual outcome. And the other one we cannot observe is called the counterfactual outcome.

It is not difficult to find that the key and the challenging issue of uplift modeling is to do the counterfactual prediction. As we have stated before, the user's social relationships and his/her social neighbors' features contribute a lot to uplift estimation. Therefore, we introduce the social graph in our work.

We define the social graph as $G = (V, E)$ ¹, where $V = \{v_1, \dots, v_N\}$ denotes the set of nodes, $N = |V|$ is the number of nodes, and $E \subseteq V \times V$ is the set of edges between nodes. Here the node denotes a user and the edge denotes two users' social relationships. Let \mathbf{X} be a matrix of node attributes. We define $\mathbf{H}^{(l)} = [\mathbf{h}_1^{(l)}, \mathbf{h}_2^{(l)}, \dots, \mathbf{h}_N^{(l)}]$ as the hidden representations of nodes in the l^{th} layer of the graph neural networks where $\mathbf{h}_i^{(l)}$ is the representation of node v_i . And we use L as the number of layers for the GNN model. For convenience, we also denote \mathbf{X} as $\mathbf{H}^{(0)}$.

3.2 GNUM

To incorporate the social relations and the neighbors' attributes, we propose a novel GNN-based uplift model (GNUM). Figure 1 shows the overall framework of the proposed GNUM, which consists of two components, i.e. graph-based representation learning and uplift estimation. Specifically, we propose two GNN-based uplift estimators in this paper working for different scenarios to address the labeled data scarcity problem for uplift estimation scenarios.

3.2.1 Graph-based Representation Learning. This component aims to learn the graph-based node (user) representations by extracting the key information from the social relations and neighbors' attributes. Basically, most GNNs can be represented by:

$$\tilde{\mathbf{h}}_i^{(l+1)} = \text{AGGR}(\tilde{\mathbf{h}}_i^{(l)}, \tilde{\mathbf{h}}_{\mathcal{N}_i}^{(l)}), \quad (2)$$

where \mathcal{N}_i represents the neighbors of user i and AGGR represents the aggregation function of the target user's embedding and his neighbors' embeddings. Experimentally, we demonstrate our proposed uplift estimation framework works well with different aggregators like GCN and GAT.

Inspired by [31], we propose more comprehensive graph-based aggregators in GNUM, which consist of a breadth aggregator to learn information from social neighbors of the current layer and a depth aggregator to ensemble information from different layers. The l -th graph convolution layer of GNUM is defined as:

$$\tilde{\mathbf{h}}_i^{(l+1)} = \text{tanh} \left(\sum_{j \in \text{ne}(i) \cup \{i\}} \alpha(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}) \mathbf{h}_j^{(l)} \mathbf{W}^{(l)} \right). \quad (3)$$

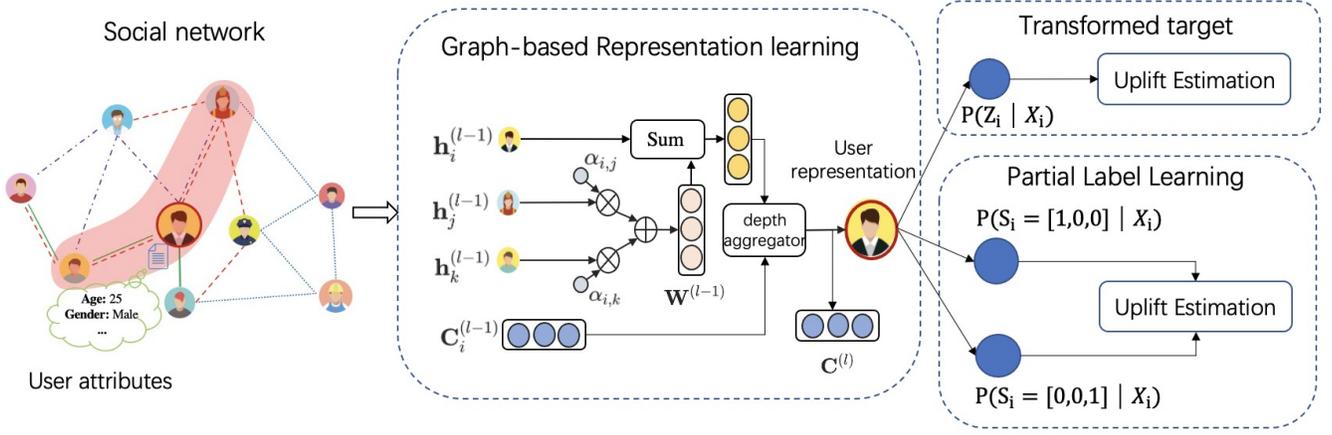
where $\alpha(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)})$ is the attention function of the breadth aggregator to measure the importance of user i and user j for uplift modeling, defined as:

$$\alpha(\mathbf{h}_i^{(l)}, \mathbf{h}_j^{(l)}) = \text{softmax}(\mathbf{v}^{(l)} \text{tanh}(\mathbf{W}_s \mathbf{h}_i^{(l)} + \mathbf{W}_d \mathbf{h}_j^{(l)})) \quad (4)$$

where $\mathbf{W}_s^{(l)}$ represents the weight for the source node, $\mathbf{W}_d^{(l)}$ represents the weight for the target node and $\mathbf{v}^{(l)}$ denotes a vector to map the representations to a value.

Given a user u_i , the breadth aggregator in each layer will adaptively gather the information from his neighbors and his own representations of the previous layer. Then we further stack multiple

¹For simplicity, we assume the social graph is a directed unweight graph


Figure 1: Framework of GNUM

convolution layers and utilize a memory-based depth aggregator to aggregate the user embedding, defined as:

$$\begin{aligned} i_i &= \sigma(W_i^{(l)\top} \tilde{\mathbf{h}}_i^{(l+1)}), & f_i &= \sigma(W_f^{(l)\top} \tilde{\mathbf{h}}_i^{(l+1)}) \\ o_i &= \sigma(W_o^{(l)\top} \tilde{\mathbf{h}}_i^{(l+1)}), & \tilde{C} &= \tanh(W_c^{(l)\top} \tilde{\mathbf{h}}_i^{(l+1)}) \\ C_i^{(l+1)} &= f_i \odot C_i^{(l)} + i_i \odot \tilde{C}, & \mathbf{h}_i^{(l+1)} &= o_i \odot \tanh(C_i^{(l+1)}) \end{aligned}$$

In this way, our proposed method can extract both local and global structural information from the social graph to obtain the graph-based representations for each user, which facilitates the learning of the following uplift estimators.

3.2.2 Transformed Target-based Uplift Estimator. With the graph-based hidden representations, an intuitive way is to build two pathways of models to project the node representations into the label space, aiming to approach the outcome with and without treatment using the treatment group data and control group data respectively [17, 18]. The process can be formulated as $\hat{y}_i(t) = f^{(t)}(\mathbf{h}_i^{(L)})$, $t \in \{0, 1\}$, where $f^{(t)}(\cdot)$ is one pathway of the model for treatment or control group data and $\hat{y}_i(t)$ represents the prediction of the user's outcome with the treatment or not. Note that $\hat{y}_i(t)$ can be continuous as the regression task or discrete as the multi-classification task.

By optimizing the regression loss or classification loss, the uplift of user i can be estimated as $\hat{\tau}_i = \hat{y}_i(1) - \hat{y}_i(0)$. However, any one pathway of their model can only utilize one group of data to learn, which will suffer from the labeled data scarcity problem. Furthermore, as previous works of literature state [19], Two-Model methods cannot well capture the relationship between data of the treatment group and the control group, which limits their performance on uplift modeling.

To address the problem, we propose a class-transformed target and prove that by using the transformed target as the learning objective for both treatment and control group data, the model is equal to do uplift estimation. We first define the observed outcome of user i as:

$$y_i^{obs} = t_i y_i(1) + (1 - t_i) y_i(0). \quad (5)$$

Then the class-transformed target can be defined as follows:

$$z_i = y_i^{obs} \cdot \frac{t_i - p}{p(1 - p)}, \quad (6)$$

where $p = \mathbb{E}(t_i | \mathbf{x}_i, G) = \mathbb{E}(t_i) = p(t_i = 1)$ is defined as the probability that user i receives the treatment. Since we focus on the randomized experiment, it is a constant.

We have the following important proposition:

PROPOSITION 1. *The uplift of user i can be estimated in the following form: $\hat{\tau}_i = \mathbb{E}(z_i | \mathbf{x}_i, G)$.*

Inspired by [1], the proof of the proposition can be found in the Appendix A.1.

Based on the proposition, we can project the node representations $\mathbf{H}^{(L)}$ to form the transformed target. The loss for the GNUM with class-transformed target (GNUM-CT) is:

$$\mathcal{L}_{CT} = \sum_{i \in V} (z_i - \hat{z}_i)^2 + \lambda L_{reg} \quad (7)$$

where z_i has been given in Eq. (6), $\hat{z}_i = \sigma(\mathbf{W}_{CT} \mathbf{h}_i^{(L)} + \mathbf{b}_{CT})$ is the prediction of the target, \mathbf{W}_{CT} and \mathbf{b}_{CT} are learnable parameters. L_{reg} is defined as the L2 regularized loss on all parameters of the proposed model and λ is set to 0.0005 in this paper.

In summary, both the treatment group and control group data are utilized to learn and optimize the proposed transformed uplift estimator, which can well alleviate the label scarcity problem and capture the inherent relationship between the two groups of data. Moreover, we do not make any assumption regarding the proposed uplift estimator with the class-transformed target. It can work well for any types of outcome, i.e. continuous and discrete outcome. Additionally, the newly proposed target is also general to balanced and imbalanced data of two groups. If the two groups of data are biased, we can also replace the p in Eq. (6) with the sample's propensity score. Therefore, it is a very general uplift estimator which can comprehensively utilize the two groups of data for modeling individual uplift values.

3.2.3 Partial-Label-based Uplift Estimator. When the outcome is discrete, i.e. a multi-classification outcome prediction problem, we propose a partial-label-based uplift estimator to further utilize more labeled data for uplift estimation. For simplicity, we will assume that the outcome is binary and introduce our solution. It can be

generalized to a multi-classification problem by transforming the problem into multiple binary scenarios.

In detail, we divide the whole users into three groups:

- Group A: The group of users who give positive outcomes regardless of whether they receive the treatment.
- Group B: The group of users who give positive outcomes only when they receive the treatment.
- Group C: The group of users who do not give positive outcomes regardless of whether they receive the treatment.

Without loss of generality, in this paper, we consider the cases where the treatment has a positive impact on the user's outcome (e.g., promotional coupon). Thus the situation where a user gives a negative outcome with the treatment but gives a positive outcome without the treatment does not exist.

Next, we generate a 3-bits coding as the partial label for each user. Each bit of code corresponds to whether the user belongs to the above-mentioned group. According to the treatment and observed outcome, the partial label of $S_i = [s_i^A, s_i^B, s_i^C]$ for user i is defined as:

$$S_i = \begin{cases} [1, 0, 0] & \text{if } t_i = 0 \text{ and } y_i^{obs} = 1 \\ [0, 1, 1] & \text{if } t_i = 0 \text{ and } y_i^{obs} = 0 \\ [1, 1, 0] & \text{if } t_i = 1 \text{ and } y_i^{obs} = 1 \\ [0, 0, 1] & \text{if } t_i = 1 \text{ and } y_i^{obs} = 0. \end{cases} \quad (8)$$

Taking a user belonging to $t_i = 1, y_i^{obs} = 1$ as an example: if we give the treatment to the user i , we can observe his positive outcome. Then based on our aforementioned definition, the user may belong to Group A or Group B. Therefore, based on our definition, the partial label is $[1, 1, 0]$ for user i .

We build two binary classifiers to help estimate the uplift. The first classifier gives the probability that the user belongs to group A, i.e. $P(S_i = [1, 0, 0] | \mathbf{x}_i)$. Then the instances whose partial labels are $S_i = [1, 0, 0]$ are regarded as positive samples, and the instances with the partial label $S_i \in \{[0, 1, 1], [0, 0, 1]\}$ are negative samples for the first classifier. The second classifier gives the probability that the user belongs to group C, i.e. $P(S_i = [0, 0, 1] | \mathbf{x}_i)$. Thus the instances with the partial label $S_i = [0, 0, 1]$ are positive samples, and the instances with $S_i \in \{[1, 1, 0], [1, 0, 0]\}$ are negative samples.

Then we project the node representations to the partial label space:

$$\begin{aligned} \hat{y}_i^{pl1} &= \sigma(\mathbf{W}^{pl1} \mathbf{h}_i^{(L)} + \mathbf{b}^{pl1}) \\ \hat{y}_i^{pl2} &= \sigma(\mathbf{W}^{pl2} \mathbf{h}_i^{(L)} + \mathbf{b}^{pl2}), \end{aligned} \quad (9)$$

where \hat{y}_i^{pl1} and \hat{y}_i^{pl2} represent the prediction of the two partial labels, \mathbf{W}^{pl1} , \mathbf{W}^{pl2} , \mathbf{b}^{pl1} and \mathbf{b}^{pl2} are learnable parameters.

Then we define the cross-entropy loss as the classification loss for partial labels. The overall loss function for GNUM with partial-label-based estimator (GNUM-PL) is defined as:

$$\mathcal{L} = \sum_{i \in \mathcal{V}} (\mathcal{F}(y_i^{pl1}, \hat{y}_i^{pl1}) + \mathcal{F}(y_i^{pl2}, \hat{y}_i^{pl2})) + \lambda L_{reg}, \quad (10)$$

where y_i^{pl1} and y_i^{pl2} are the ground truth for the first and second partial labels we defined. $\mathcal{F}(y_i, \hat{y}_i)$ is defined as the cross-entropy loss, where $\mathcal{F}(y_i, \hat{y}_i) = -(y_i \log \hat{y}_i - (1 - y_i) \log(1 - \hat{y}_i))$. L_{reg} is

defined as the L_2 regularized loss on all the parameters of the model and λ is set to 0.0005 in this paper.

Using the above loss function, the proposed graph model with partial labels can be trained. Then the final uplift of user i can be defined in the following ways:

$$\begin{aligned} \hat{\tau}_i &= p(y_i(1) | \mathbf{x}_i, G) - p(y_i(0) | \mathbf{x}_i, G) \\ &= 1 - P(S_i = [0, 0, 1] | \mathbf{x}_i, G) - P(S_i = [1, 0, 0] | \mathbf{x}_i, G) \\ &= 1 - \hat{y}_i^{pl1} - \hat{y}_i^{pl2} \end{aligned}$$

Now we introduce why the proposed model with the partial label can further improve the performance. Although the graph-based model with the transformed target we just proposed in Section 3.2.2 utilizes a general target for both treatment group data and control group data to capture the relationship between two groups of data, each sample can only be utilized once in each epoch of training. However, in our partial-label-based uplift estimator, both classifiers will utilize both the treatment data and the control data but focus on different facets of information. Specifically, the data with $S_i = [1, 1, 0]$ and $S_i = [0, 1, 1]$ can be utilized by both classifiers. In this way, on the one hand, more labeled data can be utilized to train each classifier, ensuring a better performance especially when the labeled data is scarce in uplift estimation scenarios. On the other hand, since the two classifiers can both obtain the two groups of data, the classifiers can better capture the relations and learn useful information from the two groups to achieve better results.

The pseudo-code and complexity analysis of GNUM-CT and GNUM-PL can be found in Appendix A.2.

4 EXPERIMENTS

In this section, we conduct experimental results to answer the following three questions:

- **Q1:** How our method performs compared with all the baseline methods on a public dataset and two industrial datasets? (Answered in Section 4.2.1 and Section 4.2.2.)
- **Q2:** How can our proposed estimators generalize to other GNNs? (Answered in Section 4.3.1.)
- **Q3:** What is the relationship between the user uplift and the social relation? (Answered in Section 4.3.2.)
- **Q4:** How do our proposed graph-based methods and comparing methods perform with different amounts of labeled data? (Answered in Section 4.3.3.)

4.1 Experiment settings

4.1.1 Dataset. We evaluate our method on one public dataset and two real-world industrial datasets.

Firstly, we follow [18] to build a semi-synthetic dataset base on BlogCatalog: the node features and network structures are collected from the BlogCatalog. The treatments and outcomes are synthesized. There are confounders in this dataset and the control and test datasets are biased. In detail, the node represents a blogger and the edge denotes their social relationships. The node features are bag-of-words representations of keywords to describe the bloggers. We synthesize (1) the outcomes as the rating of readers on the bloggers and (2) the treatments as whether the blog contents are shown on mobile devices or desktops. The detailed synthetic process for the synthetic process can be found in Appendix A.3.1. There are

Table 1: The experimental result on public dataset BlogCatalog. The result of the best performance is in bold and the result of the second best performance is underlined.

κ_2	Outcome Type	Metric	GNUM-PL	GNUM-GCN	GNUM-GAT	GNUM-CT	Two-Model	CTM	Uplift-RF	NetDeconf	DML	DRL
0.5	Continuous	$\sqrt{\epsilon_{PEHE}}$	/	/	/	4.164	9.215	8.448	6.760	<u>4.496</u>	5.312	5.407
		ϵ_{ATE}	/	/	/	0.935	4.172	3.317	2.629	<u>0.970</u>	1.244	1.360
	Binary	$\sqrt{\epsilon_{PEHE}}$	0.529	0.642	<u>0.586</u>	0.613	1.253	0.964	0.740	0.621	0.689	0.707
2	Continuous	$\sqrt{\epsilon_{PEHE}}$	0.121	0.147	<u>0.129</u>	0.138	0.470	0.364	0.210	0.135	0.159	0.172
		ϵ_{ATE}	/	/	/	9.337	23.348	17.277	15.945	<u>9.623</u>	12.134	13.060
	Binary	$\sqrt{\epsilon_{PEHE}}$	/	/	/	2.102	10.920	8.624	8.004	<u>2.243</u>	6.530	7.191
	Continuous	$\sqrt{\epsilon_{PEHE}}$	0.353	0.430	<u>0.392</u>	0.411	1.047	0.820	0.699	0.423	0.554	0.598
		ϵ_{ATE}	0.105	0.136	<u>0.114</u>	0.126	0.417	0.255	0.202	0.131	0.164	0.160

Table 2: Statistics of datasets. $|V|$ denotes the number of users, $|E|$ denotes the number of edges and $|F|$ denotes the number of attributes. V_t and V_c represent the set of users in the treatment group and the set of users in the control group, respectively.

Datasets	$ V $	$ E $	$ F $	$ V_t $	$ V_c $
Industry-A	505.2K	2.2M	652	252.6K	252.6K
Industry-B	573.8K	2.6M	915	286.9K	286.9K
Blogcatalog	5.2K	173.5K	8189	1.6K	3.6K

three different parameters C, κ_1, κ_2 that control the synthetic results of the dataset. Following the experimental settings in previous work[18], we set $C = 5, \kappa_1 = 10, \kappa_2 \in \{0.5, 2\}$.

The two real-world industrial datasets are collected from an internet company², denoted as Industry-A and Industry-B. The company has two products and plans to send discount coupons to users who have not purchased the products. Because the discount coupons are limited, we need to find the users who are most likely to purchase the product when receiving the coupons. This can be regarded as a problem of uplift modeling. For each dataset, the users are randomly split into the treatment group and the control group. The users of the treatment group will receive the coupons and the users of the control group will not. Then we observe whether they will purchase the product for the following 30 days. The nodes represent users and the edges represent users' friendships. The user features mainly consist of statistical features regarding the user's profile and behaviors in our platform. Note that the collection of data removes the user's sensitive information and obtains the user's privacy authorization.

The detail statistics of three datasets are shown in Table 2.

4.1.2 Baseline Methods. To evaluate the effectiveness of GNUM-PL and GNUM-CT, we compare them with three lines of state-of-the-art uplift methods, including NN-based methods (Two-Model [38] and CTM [24]), a tree-based method (Uplift-RF [16]), Causal-effect-based methods (DML [6] and DRL [54]) and graph-based uplift methods (NetDeconf [18]). GNUM-GCN and GNUM-GAT are implemented by using the GCN or GAN as the GNN backbones but

²The data set does not contain any Personal Identifiable Information. The data set is desensitized and encrypted. Adequate data protection was carried out during the experiment to prevent the risk of data copy leakage, and the data set was destroyed after the experiment. The data set is only used for academic research, it does not represent any real business situation.

still use partial-label-based uplift estimator. More details of the comparing methods can be found in Appendix A.3.2. And the parameter settings can be found in Appendix A.3.3.

4.2 Overall Performance

4.2.1 Results on BlogCatalog. Since we have both factual and counterfactual outcomes for each user in BlogCatalog, we can measure the performance of different methods by comparing the predicted average treatment effect (ATE) with the ground-truth ATE, where ATE is defined as the average uplift over the users as $ATE = \frac{1}{n} \sum_{i=1}^n \tau_i$. Specifically, we use two evaluation metrics, i.e. the Rooted Precision in Estimation of Heterogeneous Effect ($\epsilon_{PEHE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\hat{\tau}_i - \tau_i)^2}$) and Mean Absolute Error on ATE ($\epsilon_{ATE} = \left| \frac{1}{n} \sum_{i=1}^n (\hat{\tau}_i) - \frac{1}{n} \sum_{i=1}^n (\tau_i) \right|$). Furthermore, we also extend the setting to the scenario of the discrete outcome, by setting the outcome of each sample greater than the mean ATE value as 1, otherwise as 0. It is worth noting that these two metrics require both factual and counterfactual outcomes. Therefore, we only report the results on semi-synthetic dataset BlogCatalog in Table 1. Then we have the following observations:

- In the regression setting, our proposed method GNUM-CT outperforms other baseline methods by 5% to 10%, which demonstrates that the proposed class-transformed target is able to utilize the graph-based data more effectively.
- In the classification setting, our proposed method GNUM-PL outperforms GNUM-CT and other baseline methods by 12% to 25%, which demonstrates that the proposed partial-label learning can further utilize the labeled graph data to improve the overall performance.
- All the graph-based methods achieve a substantial gain over non-graph-based methods, which demonstrates the importance of graph data for uplift modeling.
- The result that the Two-Model achieves bad results demonstrates that it is very important to use a common uplift estimator to utilize the treatment and control group of data together.

4.2.2 Results on Industrial Datasets. Without the counterfactual outcome, we use commonly accepted metrics, i.e. uplift curve and the Qini Coefficient value to evaluate the performance of different methods on our industrial datasets.

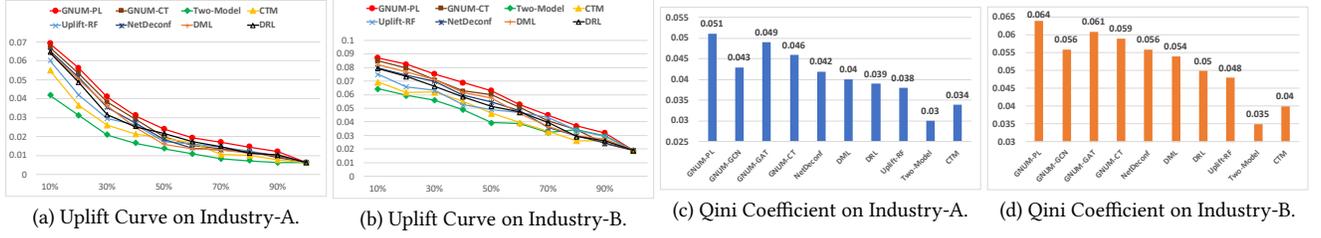


Figure 2: The uplift curve on (a) Industry-A and (b) Industry-B: The ordinate represents the uplift value defined in Eq. 11, and the abscissa represents users with top $k\%$ largest predicted uplift values. The Qini Coefficient on (c) Industry-A and (d) Industry-B.

Table 3: The detailed comparison of the uplift curve. The result of the best performance is in bold and the result of the second best performance is underlined.

Datasets	Quantiles	GNUM-PL	GNUM-GCN	GNUM-GAT	GNUM-CT	Two-Model	CTM	Uplift-RF	NetDeconf	DML	DRL
Industry-A	Top10%	6.92%	6.58%	<u>6.79%</u>	6.67%	4.18%	5.50%	6.00%	6.54%	6.50%	6.44%
	Top20%	5.52%	5.22%	<u>5.45%</u>	5.34%	3.12%	3.63%	4.19%	5.19%	5.02%	4.86%
Industry-B	Top10%	8.72%	8.38%	<u>8.51%</u>	8.47%	6.44%	6.92%	7.51%	8.35%	8.21%	7.98%
	Top20%	8.20%	7.85%	<u>8.11%</u>	7.97%	5.95%	6.21%	6.58%	7.74%	7.67%	7.32%

In detail, we sort all users in the treatment group and control group based on their predicted uplift values. Then we will select the top- $k\%$ users to get their uplift Y_k as:

$$Y_k = \sum_{i \in V_t^k} \hat{\tau}_i / |V_t^k| - \sum_{i \in V_c^k} \hat{\tau}_i / |V_c^k| \quad (11)$$

where V_t^k and V_c^k are the set of the top $k\%$ samples of the treatment and control group, and $\hat{\tau}_i$ is the predicted uplift value by the model. By changing the k from 10% to 100%, we can obtain the curve on Y_k . Note that the leftmost point corresponds to the top 10% users which the models predict as the most sensitive to the treatment and the following right part corresponds to the top 20% users. Since the ATE of the dataset is fixed, all the curves will converge to the same value. A well-performing model will show a curve with a larger slope. The left part has larger values than other methods. The results of uplift curve is shown in Figure 2(a)(b). In addition, we introduce the Qini metric to measure the overall performance of uplift methods [38]. Similar to the AUC value, the Qini metric measures the distance between the Qini curve and the random curve. The detail of the calculation process of Qini curve and Qini Coefficient can be referred in [38]. We show the results of the Qini Coefficient in Figure 2(c)(d).

From Figure 2, we find that the proposed method GNUM-PL and GNUM-CT consistently outperform the baseline methods on two metrics. Specifically, GNUM-PL improves the best performing baseline method NetDeconf by an improvement of 21% and 14% on two datasets in terms of Qini Coefficient. It demonstrates that our proposed methods have a better ranking performance regarding the user’s uplift. Additionally, the result that GNUM-PL performs better than GNUM-CT demonstrates that in the classification scenario, the proposed uplift estimators based on partial label learning can further improve the overall performance because partial label learning can utilize the labeled data more effectively. The result

that CTM outperforms Two-Model demonstrates that a common uplift estimator to model both the treatment and control group data is very essential.

In many real-world scenarios, we only focus on samples with uplift values ranking ahead because we will only give actions to users with larger uplift values. Therefore, we give a detailed comparison of the users with the largest 20% uplift values in Table 3. We can find that GNUM-CT and GNUM-PL also achieve better results than other baseline methods, which demonstrates that our proposed methods with the two estimators can find users with larger uplift.

4.3 In-Depth Analysis

4.3.1 Generality to different backbones of GNN models. We further replace our GNN models with GAT and GCN to demonstrate the generality of our proposed uplift estimators. The results are shown in Figure 4. We find in most cases, our proposed graph-based methods GNUM-PL, GNUM-GCN and GNUM-GAT perform better than NetDeconf and DML. It demonstrates that our proposed uplift estimators can be adapted to different GNN-based representation learning methods effectively. Furthermore, GNUM-PL still achieves the best performance because our breadth and depth aggregators can utilize more informative information from the social graph.

4.3.2 Correlation between Uplift and Social Relationship. Previously, we claim that the uplift difference between users with social relationships is smaller than the uplift difference between random users. To verify this, for each user we first calculate the average of the inferred uplift value of his neighbors. Then, according to the number of his neighbors, we randomly sample the same number of users to calculate the average of their inferred uplift values. Finally, we compare the uplift difference between the above two averaged values and the user’s own inferred uplift value. Specifically, we use mean-squared error (MSE) as the evaluation metric.

Table 4: The mean-squared error of the inferred uplift between different users.

Datasets	Sampling strategy	GNUM-PL	GNUM-GCN	GNUM-GAT	GNUM-CT	NetDeconf	Two-Model	CTM	Uplift-RF	DML	DRL
Industry-A	Neighbors	1.05×10^{-3}	0.97×10^{-3}	1.07×10^{-3}	1.11×10^{-3}	1.07×10^{-3}	1.29×10^{-3}	1.30×10^{-3}	1.22×10^{-3}	1.14×10^{-3}	1.18×10^{-3}
	Random	1.40×10^{-3}	1.38×10^{-3}	1.40×10^{-3}	1.39×10^{-3}	1.37×10^{-3}	1.35×10^{-3}	1.41×10^{-3}	1.39×10^{-3}	1.36×10^{-3}	1.38×10^{-3}
Industry-B	Neighbors	1.44×10^{-3}	1.32×10^{-3}	1.42×10^{-3}	1.49×10^{-3}	1.48×10^{-3}	1.75×10^{-3}	1.68×10^{-3}	1.63×10^{-3}	1.50×10^{-3}	1.55×10^{-3}
	Random	1.95×10^{-3}	1.94×10^{-3}	1.93×10^{-3}	1.95×10^{-3}	1.91×10^{-3}	1.87×10^{-3}	1.90×10^{-3}	1.88×10^{-3}	1.93×10^{-3}	1.92×10^{-3}

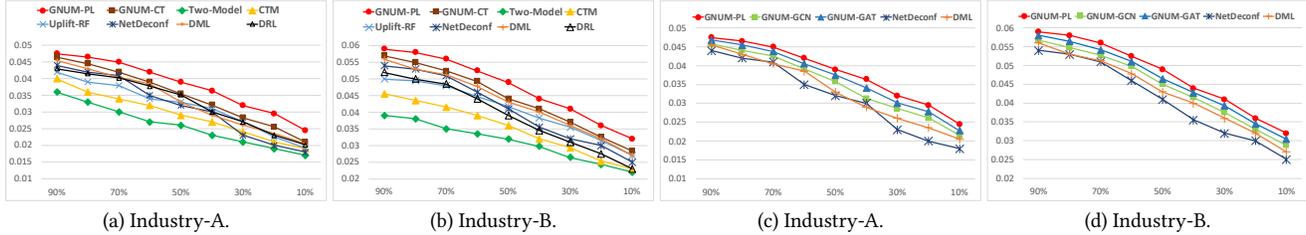


Figure 3: The Qini Coefficient of the different uplift methods under different percentage of labeled users. The ordinate represents the value of Qini Coefficient, and the abscissa represents different percentage of labeled users. (c)(d) More comparisons by replacing the GNN layers in GNUM-PL with GCN/GAT.

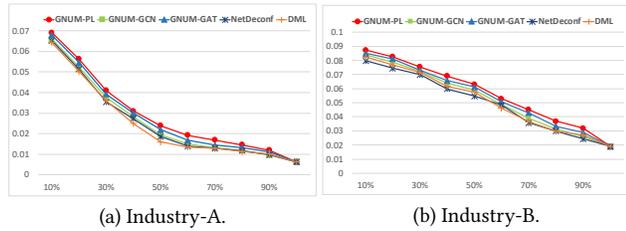


Figure 4: The uplift curve of different graph-based methods.

The result of the uplift analysis is shown in Table 4. We can see that for the graph-based methods, the uplift difference between neighbors is significantly smaller than the difference between random users compared with non-graph-based methods. Compared with the random sampling strategy, the MSE of inferred uplift calculated from neighbors drops by more than 30%. It demonstrates that the graph-based method can indeed learn the similarity information from neighbors which is useful for uplift modeling. It is worth mentioning that the smaller difference of uplift is not representing that the model performs better necessarily. The result of GNUM-GCN is worse than that of GNUM-PL and GNUM-GAT because GNUM-GCN does not capture the attentional weight of different neighbors. The uplift differences between the neighbors of Two-Model and CTM are largest, since they can not capture structural information at all. Therefore, this result demonstrates that it is very essential to utilize social relations to do the uplift estimation. And the performance gain can be achieved by proposing the graph-based model to mine the social relationships between users.

4.3.3 Data Scarcity Analysis. Due to the labeled data scarcity problem for uplift modeling, we propose GNUM-PL and GNUM-CT in this paper to alleviate the problem. To prove this point, we randomly sample 10% to 90% of the labeled users to train different uplift models and compare their performance on the same test set. Note that we will not sample the edges of the graph. We use the Qini Coefficient to evaluate the performance and note that the uplift curve

is consistent with the Qini coefficient. The experimental results are shown in Figure 3(a)(b) and we have the following observations:

- The performance of GNUM-PL and GNUM-CT are consistently above the curves of baseline methods, which demonstrates that both the partial-label-based uplift estimator and the class-transformed uplift estimator are robust to the label scarcity. Furthermore, in the classification setting, the partial-label-based uplift estimator can further improve the performance by utilizing more labeled data explicitly.
- Comparing Netdeconf and other baseline methods, we find that NetDeconf drops more quickly than other baseline methods, which proves our assumption that data scarcity is a more severe problem for the graph-based uplift method because of the more parameters. Therefore, our solution for alleviating the labeled data scarcity problem is very critical.

Moreover, as is shown in Figure 3(c)(d), GNUM-GCN and GNUM-GAT are still robust to the data scarcity problem. It demonstrates that our partial-label-based uplift estimator can alleviate the data scarcity problem when using different types of GNN-based models.

5 CONCLUSION

In this paper, we propose GNUM, a novel and general GNN-based framework with two uplift estimators for user uplift modeling. The first uplift estimators are general to different uplift scenarios, which can utilize the two groups of data together to estimate the user uplift. Specifically, when the outcome is discrete, we further propose an uplift estimator based on our designed partial labels to address the labeled data scarcity problem. Experimental results demonstrate that our proposed methods outperform state-of-the-art uplift methods under various evaluation metrics. We also give an analysis of the relationship between the uplift and social relations. And we further demonstrate the robustness of our proposed model when labeled data is limited. In the future, we expect to utilize more types of graphs to estimate the user uplift.

ACKNOWLEDGMENTS

Kun Kuang's research was supported by National Natural Science Foundation of China (62006207), Young Elite Scientists Sponsorship Program by CAST (2021QNRC001), and the Fundamental Research Funds for the Central Universities (226-2022-00142)

REFERENCES

- [1] Susan Athey and Guido W Imbens. 2015. Machine learning methods for estimating heterogeneous causal effects. *stat* 1050, 5 (2015), 1–26.
- [2] Leo Breiman. 2001. Random forests. *Machine learning* 45, 1 (2001), 5–32.
- [3] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. 2017. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* 34, 4 (2017), 18–42.
- [4] Joan Bruna, Wojciech Zaremba, Arthur Szlam, and Yann Lecun. 2014. Spectral networks and locally connected networks on graphs. In *Proceedings of the 3rd International Conference on Learning Representations*.
- [5] Tianqi Chen and Carlos Guestrin. 2016. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*. 785–794.
- [6] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. 2018. Double/debiased machine learning for treatment and structural parameters.
- [7] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. 2016. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems*. 3844–3852.
- [8] Sihao Ding, Fuli Feng, Xiangnan He, Yong Liao, Jun Shi, and Yongdong Zhang. 2022. Causal incremental graph convolution for recommender system retraining. *IEEE Transactions on Neural Networks and Learning Systems* (2022).
- [9] Richard Emsley, Mark Lunt, Andrew Pickles, and Graham Dunn. 2008. Implementing double-robust estimators of causal effects. *The Stata Journal* 8, 3 (2008), 334–353.
- [10] Fuli Feng, Weiran Huang, Xiangnan He, Xin Xin, Qifan Wang, and Tat-Seng Chua. 2021. Should graph convolution trust neighbors? a simple causal inference method. In *Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 1208–1218.
- [11] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. 2017. Neural Message Passing for Quantum Chemistry. In *International Conference on Machine Learning*. 1263–1272.
- [12] Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*. 249–256.
- [13] Chen Gong, Tongliang Liu, Yuanyan Tang, Jian Yang, Jie Yang, and Dacheng Tao. 2017. A regularization approach for instance-based superset label learning. *IEEE transactions on cybernetics* 48, 3 (2017), 967–978.
- [14] Robin Marco Gubela, Stefan Lessmann, Johannes Haupt, Annika Baumann, Tillmann Radmer, and Fabian Gebert. 2017. Revenue uplift modeling.
- [15] Leo Guelman, Montserrat Guillén, and Ana Maria Pérez Marin. 2014. Optimal personalized treatment rules for marketing interventions: A review of methods, a new proposal, and an insurance case study. *UB Riskcenter Working Paper Series*, 2014/06 (2014).
- [16] Leo Guelman, Montserrat Guillén, and Ana M Pérez-Marín. 2015. Uplift random forests. *Cybernetics and Systems* 46, 3-4 (2015), 230–248.
- [17] Ruo Cheng Guo, Jundong Li, and Huan Liu. 2020. Counterfactual evaluation of treatment assignment functions with networked observational data. In *Proceedings of the 2020 SIAM International Conference on Data Mining*. SIAM, 271–279.
- [18] Ruo Cheng Guo, Jundong Li, and Huan Liu. 2020. Learning individual causal effects from networked observational data. In *Proceedings of the 13th International Conference on Web Search and Data Mining*. 232–240.
- [19] Pierre Gutierrez and Jean-Yves Gérardy. 2017. Causal inference and uplift modelling: A review of the literature. In *International Conference on Predictive Applications and APIS*. 1–13.
- [20] Will Hamilton, Zhitao Ying, and Jure Leskovec. 2017. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*. 1024–1034.
- [21] Keisuke Hirano, Guido W Imbens, and Geert Ridder. 2003. Efficient estimation of average treatment effects using the estimated propensity score. *Econometrica* 71, 4 (2003), 1161–1189.
- [22] Eyke Hüllermeier and Jürgen Beringer. 2006. Learning from ambiguously labeled examples. *Intelligent Data Analysis* 10, 5 (2006), 419–439.
- [23] Szymon Jaroszewicz and Piotr Rzepakowski. 2014. Uplift modeling with survival data. In *ACM SIGKDD Workshop on Health Informatics (HI-KDD-14)*, New York City.
- [24] Maciej Jaskowski and Szymon Jaroszewicz. 2012. Uplift modeling for clinical trial data. In *ICML Workshop on Clinical Data Analysis*.
- [25] Song Jiang and Yizhou Sun. 2022. Estimating Causal Effects on Networked Observational Data via Representation Learning. In *Proceedings of the 31st ACM International Conference on Information & Knowledge Management*. 852–861.
- [26] Yonghan Jung, Jin Tian, and Elias Bareinboim. 2021. Estimating identifiable causal effects through double machine learning. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*.
- [27] Thomas N Kipf and Max Welling. 2017. Semi-supervised classification with graph convolutional networks. In *Proceedings of the 6th International Conference on Learning Representations*.
- [28] Toru Kitagawa and Aleksey Tetenov. 2018. Who should be treated? empirical welfare maximization methods for treatment choice. *Econometrica* 86, 2 (2018), 591–616.
- [29] Kun Kuang, Peng Cui, Hao Zou, Bo Li, Jianrong Tao, Fei Wu, and Shiqiang Yang. 2020. Data-Driven Variable Decomposition for Treatment Effect Estimation. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [30] Weibo Liu, Zidong Wang, Xiaohui Liu, Nianyin Zeng, Yurong Liu, and Fuad E Alsaadi. 2017. A survey of deep neural network architectures and their applications. *Neurocomputing* 234 (2017), 11–26.
- [31] Ziqi Liu, Chaochao Chen, Longfei Li, Jun Zhou, Xiaolong Li, Le Song, and Yuan Qi. 2019. Geniepath: Graph neural networks with adaptive receptive paths. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 33. 4424–4431.
- [32] Christos Louizos, Uri Shalit, Joris M Mooij, David Sontag, Richard Zemel, and Max Welling. 2017. Causal effect inference with deep latent-variable models. *Advances in neural information processing systems* 30 (2017).
- [33] Jing Ma, Ruo Cheng Guo, Chen Chen, Aidong Zhang, and Jundong Li. 2021. Deconfounding with Networked Observational Data in a Dynamic Environment. In *Proceedings of the 14th ACM International Conference on Web Search and Data Mining*. 166–174.
- [34] Jing Ma, Mengting Wan, Longqi Yang, Jundong Li, Brent Hecht, and Jaime Teevan. 2022. Learning causal effects on hypergraphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1202–1212.
- [35] Houssam Nassif, Finn Kuusisto, Elizabeth S Burnsise, and Jude W Shavlik. 2013. Uplift Modeling with ROC: An SRL Case Study.. In *ILP (Late Breaking Papers)*. 40–45.
- [36] Diego Olaya, Kristof Coussement, and Wouter Verbeke. 2020. A survey and benchmarking study of multitreatment uplift modeling. *Data Mining and Knowledge Discovery* 34, 2 (2020), 273–308.
- [37] Diego Olaya, Jonathan Vásquez, Sebastián Maldonado, Jaime Miranda, and Wouter Verbeke. 2020. Uplift Modeling for preventing student dropout in higher education. *Decision Support Systems* 134 (2020), 113320.
- [38] Nicholas J Radcliffe. 2007. Using control groups to target on predicted lift: Building and assessing uplift models. *Direct Marketing Analytics Journal* 1, 3 (2007), 14–21.
- [39] Nicholas J Radcliffe and Patrick D Surry. 2011. Real-world uplift modelling with significance-based uplift trees. *White Paper TR-2011-1, Stochastic Solutions* (2011), 1–33.
- [40] Piotr Rzepakowski and Szymon Jaroszewicz. 2012. Decision trees for uplift modeling with single and multiple treatments. *Knowledge and Information Systems* 32, 2 (2012), 303–327.
- [41] Uri Shalit, Fredrik D Johansson, and David Sontag. 2017. Estimating individual treatment effect: generalization bounds and algorithms. In *International Conference on Machine Learning*. PMLR, 3076–3085.
- [42] E. Michael SOBEL. 2006. What Do Randomized Studies of Housing Mobility Demonstrate?: Causal Inference in the Face of Interference. *J. Amer. Statist. Assoc.* (2006).
- [43] Michał Sołtys, Szymon Jaroszewicz, and Piotr Rzepakowski. 2015. Ensemble methods for uplift modeling. *Data mining and knowledge discovery* 29, 6 (2015), 1531–1559.
- [44] Yongduo Sui, Xiang Wang, Jiancan Wu, Min Lin, Xiangnan He, and Tat-Seng Chua. 2022. Causal attention for interpretable and generalizable graph classification. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*. 1696–1705.
- [45] Victor Veitch, Yixin Wang, and David Blei. 2019. Using embeddings to correct for unobserved confounding in networks. *Advances in Neural Information Processing Systems* 32 (2019).
- [46] Xiang Wu, Ran He, Zhenan Sun, and Tieniu Tan. 2018. A Light CNN for Deep Face Representation With Noisy Labels. *IEEE Transactions on Information Forensics and Security* (2018), 2884–2896.
- [47] Zonghan Wu, Shirui Pan, Fengwen Chen, Guodong Long, Chengqi Zhang, and S Yu Philip. 2020. A comprehensive survey on graph neural networks. *IEEE transactions on neural networks and learning systems* 32, 1 (2020), 4–24.
- [48] Łukasz Zaniewicz and Szymon Jaroszewicz. 2013. Support vector machines for uplift modeling. In *2013 IEEE 13th International Conference on Data Mining Workshops*. IEEE, 131–138.
- [49] Dalong Zhang, Xin Huang, Ziqi Liu, Zhiyang Hu, Xianzheng Song, Zhibang Ge, Zhiqiang Zhang, Lin Wang, Jun Zhou, and Yuan Qi. 2020. AGL: a Scalable System for Industrial-purpose Graph Machine Learning. *arXiv preprint arXiv:2003.02454* (2020).

- [50] Min-Ling Zhang and Fei Yu. 2015. Solving the partial label learning problem: An instance-based approach. In *IJCAI*. 4048–4054.
- [51] Min-Ling Zhang, Bin-Bin Zhou, and Xu-Ying Liu. 2016. Partial label learning via feature-aware disambiguation. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1335–1344.
- [52] Ziwei Zhang, Peng Cui, and Wenwu Zhu. 2020. Deep learning on graphs: A survey. *IEEE Transactions on Knowledge and Data Engineering* (2020).
- [53] Feipeng Zhao and Yuhong Guo. 2015. Semi-Supervised Multi-Label Learning with Incomplete Labels. *IJCAI* (2015), 4062–4068.
- [54] Ying-Qi Zhao, Donglin Zeng, Eric B Laber, Rui Song, Ming Yuan, and Michael Rene Kosorok. 2015. Doubly robust learning for estimating individualized treatment with censored data. *Biometrika* 102, 1 (2015), 151–168.

A APPENDIX

A.1 Proposition Proof

PROOF.

$$\begin{aligned} \mathbb{E}(z_i | \mathbf{x}_i, G) &= \mathbb{E}(y_i^{obs} \cdot \frac{t_i - p}{p(1-p)} | \mathbf{x}_i, G) \\ &= \mathbb{E}(t_i y_i^{obs} \cdot \frac{t_i - p}{p(1-p)} + (1 - t_i) y_i^{obs} \cdot \frac{t_i - p}{p(1-p)} | \mathbf{x}_i, G) \end{aligned} \quad (12)$$

Given $t_i \in \{0, 1\}$, we have $t_i^2 = t_i$ and $(1 - t_i)^2 = (1 - t_i)$. Simultaneously with the definition in Eq. (5), we have:

$$\begin{cases} t_i y_i^{obs} = t_i^2 y_i(1) + t_i(1 - t_i) y_i(0) = t_i^2 y_i(1) = t_i y_i(1) \\ (1 - t_i) y_i^{obs} = t_i(1 - t_i) y_i(1) + (1 - t_i)^2 y_i(0) = (1 - t_i) y_i(0). \end{cases}$$

Then we can rewire Eq. 12 as ³:

$$\begin{aligned} &\mathbb{E}(t_i y_i^{obs} \cdot \frac{t_i - p}{p(1-p)} + (1 - t_i) y_i^{obs} \cdot \frac{t_i - p}{p(1-p)} | \mathbf{x}_i, G) \\ &= \mathbb{E}(t_i y_i(1) \frac{t_i - p}{p(1-p)} + (1 - t_i) y_i(0) \frac{t_i - p}{p(1-p)} | \mathbf{x}_i, G, t_i^1) * p(t_i^1 | \mathbf{x}_i, G) \\ &+ \mathbb{E}(t_i y_i(1) \frac{t_i - p}{p(1-p)} + (1 - t_i) y_i(0) \frac{t_i - p}{p(1-p)} | \mathbf{x}_i, G, t_i^0) * p(t_i^0 | \mathbf{x}_i, G) \\ &= \mathbb{E}(\frac{y_i(1)}{p} | \mathbf{x}_i, G) * p(t_i = 1 | \mathbf{x}_i, G) - \mathbb{E}(\frac{y_i(0)}{1-p} | \mathbf{x}_i, G) * p(t_i = 0 | \mathbf{x}_i, G) \\ &= \mathbb{E}(y_i(1) | \mathbf{x}_i, G) \frac{p(t_i = 1 | \mathbf{x}_i, G)}{p} - \mathbb{E}(y_i(0) | \mathbf{x}_i, G) \frac{p(t_i = 0 | \mathbf{x}_i, G)}{1-p} \\ &= \mathbb{E}(y_i(1) | \mathbf{x}_i, G) - \mathbb{E}(y_i(0) | \mathbf{x}_i, G) \\ &= \hat{\tau}_i \end{aligned}$$

A.2 Pseudo Code and Complexity Analysis

The pseudo-code of GNUM-CT is given in Algorithm 1 and GNUM-PL in Algorithm 2. Our proposed method can be trained by an end-to-end back-propagation, and thus we can use gradient descent to optimize the model.

For each sample, it will go through the GNUM layer first and then perform the uplift prediction. The complexity of the whole process is $O(M \sum_{l=0}^L f_l + N \sum_{l=1}^L f_{l-1} f_l)$ where N denotes the total number of users, including the users in the treatment and control group, $M = |E|$ is the number of relationships, L is number of hidden layers and f_l is the dimensionality of the l^{th} hidden layer. In practice, L and f_l are often bounded within a small constant. The proposed method is linear to the number of users and number of relationships in the dataset respectively. Therefore, the overall model is scalable.

³Due to the limit of space, t_i^j represents $t_i = j$ user.

Algorithm 1 Graph Neural Network for Uplift Modeling (GNUM) using Class-Transformed Target

Require: Set of user features, treatment assignment and observed outcome $\{(\mathbf{x}_i, t_i, y_i)\}_{i=1}^N$, social graph $G = (V, E)$ and number of layers L ,

Ensure: Prediction of user uplift and GNUM parameters Θ .

- 1: Initialize all parameters Θ and using \mathbf{X} as $\mathbf{H}^{(0)}$
 - 2: **while** \mathcal{L} does not converge **do**
 - 3: **for** $i \leftarrow 1$ to L **do**
 - 4: Calculate graph-based representation $\mathbf{H}^{(l)}$.
 - 5: **end for**
 - 6: Calculate \hat{y}_i^{pl1} and \hat{y}_i^{pl2} from learned representations $\mathbf{H}^{(L)}$ using Eq. (9).
 - 7: Calculate the class-transformed target z_i using Eq. (6).
 - 8: Calculate the loss \mathcal{L}_{CT} using Eq. (7).
 - 9: Update Θ using back-propagation.
 - 10: **end while**
 - 11: Using the trained model parameters to estimate the user uplift as: $\hat{\tau}_i = \mathbb{E}(z_i | \mathbf{x}_i, G)$.
-

Algorithm 2 Graph Neural Network for Uplift Modeling (GNUM) using Partial Label Learning

Require: Set of user features, treatment assignment and observed outcome $\{(\mathbf{x}_i, t_i, y_i)\}_{i=1}^N$, social graph $G = (V, E)$ and number of layers L ,

Ensure: Prediction of user uplift and GNUM parameters Θ .

- 1: Initialize all parameters Θ and using \mathbf{X} as $\mathbf{H}^{(0)}$
 - 2: **while** \mathcal{L} does not converge **do**
 - 3: **for** $i \leftarrow 1$ to L **do**
 - 4: Calculate graph-based representation $\mathbf{H}^{(l)}$.
 - 5: **end for**
 - 6: Calculate \hat{y}_i^{pl1} and \hat{y}_i^{pl2} from learned representations $\mathbf{H}^{(L)}$ using Eq. (9).
 - 7: Construct the sample's partial labels based on y_i and t_i .
 - 8: Calculate the loss \mathcal{L} using Eq. (10).
 - 9: Update Θ using back-propagation.
 - 10: **end while**
 - 11: Using the trained model parameters to estimate the user uplift as: $\hat{\tau}_i = 1 - P(S_i = [1, 0, 0] | \mathbf{x}_i, G) - P(S_i = [0, 0, 1] | \mathbf{x}_i, G)$
-

$O(N d_{max}^2 S_{max})$, where N denotes the total number of users, including the users in the treatment and control group, d_{max} denotes the maximum number of dimensionality among different layers and S_{max} denotes in the social graph the maximum degree among all the users. In practice, S_{max} is often bounded within a constant. For example, in social networks like Facebook, the number of friends a user can have has an upper bound.

A.3 More Experimental Details

A.3.1 More Details about Dataset. Blogcatalog is a social blog directory that manages bloggers and their blogs. The dataset contains the features of bloggers and the social relationships between bloggers listed on the BlogCatalog website. The features are bag-of-words representations of keywords in bloggers' descriptions. We follow

the assumptions and procedures of synthesizing the outcomes and treatments assignments in [18], in which the outcomes represent the opinions of readers on each blogger and the treatments represent whether contents created by a blogger receive more views on mobile devices or desktops. In the treatment group, the blogger’s blogs are read more on mobile devices. In the control group, the blogger’s blogs are read more on desktops. We assume that the social relationships of bloggers can causally affect their treatment assignments and readers’ opinions of them. We trained the LDA topic model to get the device preference of the readers of the i -th blogger’s content as:

$$\begin{aligned} \Pr(t = 1 \mid \mathbf{x}_i, A) &= \frac{\exp(p_1^i)}{\exp(p_1^i) + \exp(p_0^i)}; \\ p_1^i &= \kappa_1 r(\mathbf{x}_i)^T r_1^c + \kappa_2 \sum_{j \in \mathcal{N}(i)} r(\mathbf{x}_j)^T r_1^c \\ &= \kappa_1 r(\mathbf{x}_i)^T r_1^c + \kappa_2 (Ar(\mathbf{x}_j))^T r_1^c; \\ p_0^i &= \kappa_1 r(\mathbf{x}_i)^T r_0^c + \kappa_2 \sum_{j \in \mathcal{N}(i)} r(\mathbf{x}_j)^T r_0^c \\ &= \kappa_1 r(\mathbf{x}_i)^T r_0^c + \kappa_2 (Ar(\mathbf{x}_j))^T r_0^c, \end{aligned}$$

where $\kappa_1, \kappa_2 \geq 0$ signifies the magnitude of the confounding bias resulting from a blogger’s topics and her neighbors’ topics, respectively. $\kappa_1 = 0, \kappa_2 = 0$ means the treatment assignment is random and there is no selection bias, and greater κ_1, κ_2 means larger selection bias. The factual outcome and the counterfactual outcome of the i -th blogger are simulated as:

$$\begin{aligned} y^F(\mathbf{x}_i) &= y_i = C(p_0^i + t_i p_1^i) + \epsilon \\ y^{CF}(\mathbf{x}_i) &= C[p_0^i + (1 - t_i) p_1^i] + \epsilon, \end{aligned}$$

where C is a scaling factor and the noise is sampled as $\epsilon \sim \mathcal{N}(0, 1)$. In this work, we set $C = 5, \kappa_1 = 10, \kappa_2 \in \{0.5, 2\}$.

A.3.2 Details about the comparing methods. The detail descriptions of the comparing methods can be found here.

- Two-Model: This method infers the labels in the treatment group and the control group respectively.
- CTM: The Class-Transformation model (CTM) creates the new target variable to estimate the uplift.
- Uplift-RF: This model modifies existing random forest algorithms to directly infer the uplift.
- DML: Double Machine Learning (DML) method uses the Neyman orthogonal score and cross-fitting to construct the uplift estimator.
- DRL: Doubly Robust Learning (DRL) method combines the error imputation and the inverse propensity score estimator to address the bias problem.
- NetDeconf: It uses the graph to minimize confounding bias in the task of estimating treatment effects.
- GNUM-GCN/GAT: To show the effectiveness of partial label learning, we replace the graph representation learning method with GCN and GAT respectively for comparison.

A.3.3 Parameter Settings. We adopt the deep neural networks [30] with two 64-unit hidden layers as the building block for the Two-Model and CTM methods. For uplift random forest, we set the number of estimators as 50 and the depth of the tree as 5. For

NetDeconf and our proposed method, we use two layers of GNN with units of 128 – 64. All the weight matrices are initialized using Xavier initialization [12]. We train the model for 5 epochs with a learning rate of 0.0001 and batch size of 256. Note that we use grid search to get the best hyper-parameters. For each dataset, we randomly sampled 70% of the users as the training set, 10% as the validation set, and 20% as the test set. The models are trained on a cluster of 10 Dual-CPU servers with AGL [49] framework. For the largest dataset, containing 573.8K nodes and 2.6M edges, the proposed method took about 17 minutes to train on a cluster of 10 Dual-CPU servers. We run each algorithm 5 times and report the average result.