Network Experiment Design for Estimating Direct Treatment Effects

Zahra Fatemi
zfatem2@uic.edu
University of Illinois at Chicago
Chicago, IL

Elena Zheleva
ezheleva@uic.edu
University of Illinois at Chicago
Chicago, IL

ABSTRACT

Network experiment design refers to the design of controlled experiments for interacting units with the goal of estimating a causal effect of interest. Estimating the effect of treatment alone on units' outcome, known as direct treatment effect, in network experiments is challenging due to information spillover between peers through shared edges. Prominent methods for network experiment design mostly focus on estimating total treatment effects, the combination of peer effects and direct treatment effects. Less focus has been given to approaches that provide an unbiased estimation of direct treatment effect. We present a framework that takes advantage of *independent sets* and assigns treatment and control only to a set of non-adjacent nodes in a graph, in order to disentangle peer effects from direct treatment effect estimation. Randomizing over independent set nodes removes peer effects between nodes in the experiment while canceling out the peer effects from nodes outside the experiment. Through a series of simulated experiments on synthetic and real-world network datasets, we show that our framework significantly increases the accuracy of direct treatment effect estimation in network experiments.

KEYWORDS

Causal inference, Network experiment design, Peer effect, Direct treatment effect

1 INTRODUCTION

Causal inference plays a central role in many disciplines, from medical sciences to economics. Randomized Controlled Trials (RCT) or A/B testing is considered the gold standard for estimating the effect of a treatment on a population. As a motivating example, consider the problem of assessing the effectiveness of a vaccine on providing immunity against disease. To conduct a randomized controlled trial, a group of individuals from the population are vaccinated at random (treatment) and the infection rate in this group is compared to the infection rate in a random group of people who are not vaccinated (control). The difference between infection rates of the two groups reflects the causal effect of the vaccine on immunity, that is if we can assume that individuals do not interact with one other.

Unfortunately, assuming that individuals do not interact is rarely realistic in real-world scenarios. Due to potential interference or spillover between connected individuals in a social network, measuring causal effects is challenging. The presence of interference violates a fundamental causal inference assumption known as the *Stable Unit Treatment Value Assumption (SUTVA)*. SUTVA states that in order for a causal effect estimate to be unbiased, the outcome of an individuals should be affected by their own treatment only and not by the treatment assignment of other individuals in the population [9]. In the running example, if vaccinated people in the treatment group interact with people in the control group, then people in the control group may be protected from the disease based on herd immunity and the infection rate can appear to be the same in both groups. In reality, for an effective vaccine the base rate of infection would have been higher in the control group in the absence of interactions with the treated group.

In the presence of peer effects, the measured causal estimand is the combination of *Direct Treatment Effect (DTE)* and *Peer Effects (PE)*, known as *Total Treatment Effect (TTE)*. Direct Treatment Effect is defined as the difference between average outcome of treated and untreated individuals due to the treatment alone. Fig .1 shows two different types of peer effects that exist in a fully randomized experiment and can make treatment effect estimation biased. Different studies focus on designing network experiments that focus on different types of causal effects of interest. For example, Cluster-based network experiments are popular approaches for estimating TTE. These approaches reduce interference between treatment groups by partitioning the network into clusters with dense connections within clusters and few edges across clusters [8, 20, 26]. Then, by randomizing treatment assignment at the cluster-level, spillover across treatment groups is reduced.

In contrast, the goal of this paper is to measure Direct Treatment Effects in network experiments. We propose a network experiment design using independent sets as an unbiased approach for measuring Direct Treatment Effects. In the proposed framework, we divide network nodes into two sets: 1) *independent set nodes*, and 2) graph nodes that are not in the independent set to which we refer as *bystander nodes*. By assigning the independent set nodes to treatment and control groups, we ensure that there are no peer effects between nodes participating in the experiment, regardless whether they are in different treatment groups or the same treatment group. Key to our proposed experiment design is the idea that in expectation, the peer effects of bystander nodes on the treatment group is the same as the peer effect of bystander nodes on the control group, thus canceling each other in the total treatment effect estimation.

MLG '20, Aug 24, 2020, San Diego, CA, US

ACM ISBN 978-1-4503-XXXX-X/18/06…$15.00
https://doi.org/10.1145/1122445.1122456

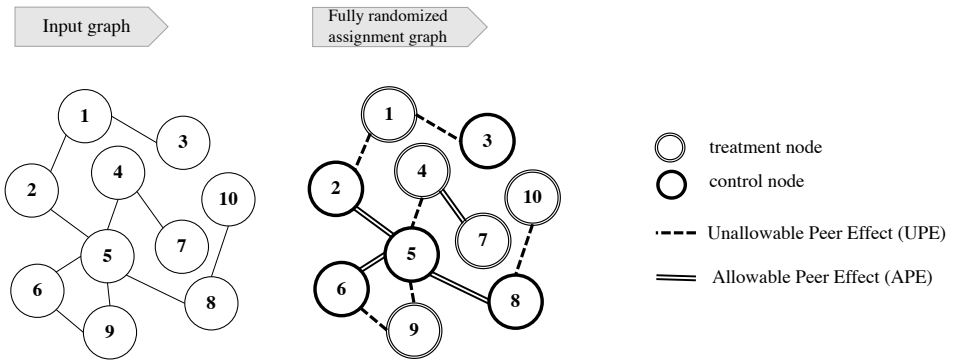

Figure 1: Illustration of fully randomized assignment in network experiments. *Input graph*: a graph of nodes and the connections between them. *Fully randomized assignment graph*: the output graph that represents fully randomized treatment assignment of nodes and peer effects that exists in the experiment.

Applying this framework to the vaccination example would ensure that individuals in treatment and control do not interact with each other and allow accurate estimation of the effect of vaccination on the infection rate.

The rest of the paper is structured as follows. In Section 2, we review background on causal effect estimation in network experiments. In Section 3, we define different causal estimands and the problem that we address in the paper. In Section 4, we present our proposed network experiment design that increases the accuracy of Direct Treatment effect estimation. In Section 5, we present experimental setup and simulation results in real-world and synthetic data. In Section 6, we conclude and discuss directions for future work.

2 RELATED WORK

The assumption of no interference is at the core of most causal studies [13]. With the rise of interest in social networks, experiments revealed that this assumption is violated in studies with non i.i.d data where SUTVA does not hold [26]. By attracting attention toward network experiments, dependent on the assumptions made in the study different estimands have been proposed. Several methods have been considered to infer causal effects from observational studies [2, 15, 18, 22], but observational studies are not the focus of this paper. Here, we review relevant work on three main causal estimands for designing randomized controlled trials in networks.

Direct Treatment Effect (DTE) Estimation. Estimating the effect of treatment alone has been studied in the context of network experiment design. Jagadeesan et al. [10] proposed an approach to reduce the bias of Neymanian estimator of DTE under interference and homophily. In this approach, treatment assignment is considered as a quasi-coloring on a graph and every treated node is tried to be matched with a control node with identical number of treated and control neighbors to create a balanced interference in network experiments. In networks where a perfect quasi- coloring is not possible, nodes are ordered by degree and then nodes with similar degree are paired and assigned to treatment or control. The accuracy of causal effect estimation in this method depends on the network structure, degree distribution of the nodes and approaching perfect quasi-coloring to a perfect quasi-coloring. Sussman and Airoldi [24] propose an approach to estimate DTE considering a fixed design for potential outcomes. Similar to these approaches, we focus on estimating DTE in the presence of peer effects, but our approach can be applied in networks with different structural properties.

Peer Effect (PE) Estimation. Different network experiments designs have been proposed to tease out the effects of peers from the total treatment effects. Eckles et al. [6] conducted a randomized experiment to measure peer effects by assigning individual's peers to encouragements to behaviors that affect the outcome of individuals. Saint-Jacques et al. propose an ego-network randomization approach to measure peer effects where ego is a focal node and alter is her first-hop neighbors [19]. To measure peer effects, the outcome of a treated ego when all of its alters are treated is compared to the outcome of a treated ego when all of its alters are untreated. Toulis and Kao [25] defined a new causal estimand for peer influence effects and proposed two sequential randomization and model-base approaches to measure the estimand. In sequential randomization approach, nodes are randomly assigned to a k-level exposure status where exactly k neighbors of nodes are treated or non-exposure status where all neighbors are untreated.

Total Treatment Effect (TTE) Estimation. TTE is one of the most popular causal estimands in network experiments, especially in cluster-based randomization approaches [5, 6, 16]. Using a Horvitz-Thompson estimator, Ugander et al. [26] proposed a cluster-based approach to estimate TTE in network experiment design. Saveski et al. [20], by stratification of balanced clusters attempted to present an unbiased TTE estimator. Fatemi and Zheleva [8] presented *CMatch*, a framework that uses weighted graph clustering technique to minimize interference and selection bias in network experiment design. In this approach, by matching clusters with more homogeneous nodes and assigning matched clusters to treatment and control groups, they are able to reduce the error in TTE estimation in cluster-based network experiment design.

3 PROBLEM DEFINITION

In this section, we formally define the data model, the causal estimand and the problem that we address in this paper, following the notation and terminology of Fatemi and Zheleva [8].

3.1 Data model

Let $G = (\mathbf{V}, \mathbf{E})$ be an undirected graph of n nodes where $\mathbf{V}$ denotes the nodes and $\mathbf{E} = \{e_{ij}\}$ the edges, such that e_{ij} corresponds to an edge between node $v_i \in \mathbf{V}$ and node $v_j \in \mathbf{V}$. For each node v_i there is a vector of attributes denoted by $v_i.\mathbf{X}$, an outcome denoted by $v_i.Y$, and a set of neighbors denoted by $\mathbf{N}_i$. Let $v_i.T \in \{0, 1, 2\}$ denote the treatment assignment of node v_i such that for a treated node $v_i.T = 1$, for an untreated node $v_i.T = 0$ and for a node excluded from the experiments $v_i.T = 2$. Let $\mathbf{Z} \in \{0, 1, 2\}^n$ be the treatment assignment vector of a population of size n.

3.2 Causal effect estimation

Total Treatment Effect (TTE) is defined as the difference between the outcomes of individuals in a population when all are treated $\mathbf{Z_1} = \{1\}^n$ and all are not ($\mathbf{Z_0} = \{0\}^n$) [26]:

$$TTE = \frac{1}{n} \sum_{v_i \in \mathbf{V}} (v_i.Y(\mathbf{Z_1}) - v_i.Y(\mathbf{Z_0})). \tag{1}$$

where $v_i.Y(\mathbf{Z_1})$ and $v_i.Y(\mathbf{Z_0})$ are the potential outcomes of node v_i under the treatment assignment vectors $\mathbf{Z_1}$ and $\mathbf{Z_0}$, respectively.

However, the fundamental problem of causal inference is that it is impossible to observe $v_i.Y(T = 1)$ and $v_i.Y(T = 0)$ simultaneously. In real-world scenarios, TTE is estimated by averaging outcomes over randomized treatment and control groups via difference-in-means: $TTE \hat{=} \overline{V_1.Y} - \overline{V_0.Y}$ [23].

In network experiments and in the presence of interference, total treatment effects measure the contribution of both direct treatment effects (DTE) and peer effects (PE):

$$\hat{TTE} = \overline{V_1.Y} - \overline{V_0.Y} = DTE(\mathbf{V}) + PE(\mathbf{V}_1) - PE(\mathbf{V}_0). \tag{2}$$

Direct Treatment Effect (DTE) is the difference between average outcome of treated and untreated individuals due to the treatment alone. Under the assumption that PE(V_1)-PE(V_0)=0, TTE=DTE.

Peer Effect (PE) measures the average influence of neighbors on the outcome of individuals. PE is estimated as:

$$PE(V) = \mathop{\mathbb{E}}_{v_i \in V}[v_i.Y | v_i.T = t, N_i.\boldsymbol{\pi}] - \mathop{\mathbb{E}}_{v_i \in V}[v_i.Y | v_i.T = t, N_i = \emptyset]. \tag{3}$$

where $N_i.\boldsymbol{\pi}$ denotes the vector of treatment assignments to node v_i's neighbors N_i. Let t show the treatment group of each node v_i. Peer effect can be divided into two different types of effects:

- *Allowable Peer Effect (APE)*, defined as the peer effect between neighbors N_i^t in the same treatment class as node v_i where their treatment assignment is denoted by $N_i^t.\boldsymbol{\pi}$ and is estimated as:

$$APE(V) = \mathop{\mathbb{E}}_{v_i \in V}[v_i.Y | v_i.T = t, N_i^t.\boldsymbol{\pi}] - \mathop{\mathbb{E}}_{v_i \in V}[v_i.Y | v_i.T = t, N_i^t = \emptyset]. \tag{4}$$

- *Unallowable Peer Effect (UPE)* is the peer effect of neighbors $N_i^{\bar{t}}$ ($\bar{t} \neq t$) from different treatment class $N_i^{\bar{t}}.\boldsymbol{\pi}$ on the outcome $v_i.Y$ of individual v_i. *UPE* is formally defined as:

$$UPE(V) = \mathop{\mathbb{E}}_{v_i \in V}[v_i.Y | v_i.T = t, N_i^{\hat{t}}.\boldsymbol{\pi}] - \mathop{\mathbb{E}}_{v_i \in V}[v_i.Y | v_i.T = t, N_i^{\hat{t}} = \emptyset]. \tag{5}$$

The fully randomized assignment graph in Fig. 1 shows these two types of peer effects between connected nodes. For example, the peer effect between control node 2 and treatment node 1 is different from the peer effect that exists between node 2 and another control node, node 5.

The question we are interested to answer is: What is the causal effect of the treatment alone? This question has many practical applications for estimating the effectiveness of different policy interventions. Some examples include: What is the individual protection from a disease due to vaccination alone (and not herd immunity)? What is the effect of advertisement on motivating a person to buy a new phone? In network experiments, it is challenging to disentangle *DTE* from *PE* and this is the main focus of our paper.

Now, we are ready to define the problem we try to address in this paper:

Problem 1. *Network experiment design for direct treatment effect estimation. Given an undirected graph $G = (\boldsymbol{V}, \boldsymbol{E})$, and a set of attributes $\boldsymbol{V}.\boldsymbol{X}$ associated with each node. Find a treatment assignment vector $\mathbf{Z}$ of a population with three different subsets of nodes, the treatment nodes $\boldsymbol{V}_1 \in \boldsymbol{V}$, the control nodes $\boldsymbol{V}_0 \in \boldsymbol{V}$, and nodes excluded from the experiment $\boldsymbol{V}_2 \in \boldsymbol{V}$, such that:*

a. $\mathbf{V}_0 \cap \mathbf{V}_1 \cap \mathbf{V}_2 = \emptyset$
b. $|\mathbf{V}_0| + |\mathbf{V}_1|$ *is maximized*
c. $PE(V_1) - PE(V_0) \approx 0$.

The first component aims to choose treatment, control and bystander nodes excluded from the experiments that do not overlap. The second component ensures to choose as many nodes as possible from $\mathbf{V}$ to be assigned to treatment and control groups. The third component removes peer effects from causal effect estimation.

4 SOLUTION METHOD

In this section, we define an objective function corresponding to the problem of this paper and describe our proposed framework which we refer as *CauseIS* for estimating causal effect in network experiments.

Typically, total treatment effect estimation includes both APE and UPE. In a randomized approach TTE is estimated as:

$$\hat{TTE}(\mathbf{V}) = DTE(\mathbf{V}) + (APE(\mathbf{V}_1) - APE(\mathbf{V}_0)) + (UPE(\mathbf{V}_1) - UPE(\mathbf{V}_0)) \tag{6}$$

In this work, we propose an approach that makes APE($\mathbf{V}_1$)=0 and APE($\mathbf{V}_0$)=0 and in expectation makes UPE($\mathbf{V}_1$)-UPE($\mathbf{V}_0$)=0, thus making the estimated TTE correspond to DTE. We first define an objective function that addresses the goals specified in *Problem 1.*

4.1 Objective function

The goal of the objective function is to find a subset of $\mathbf{V}$ with maximum cardinality (*Problem 1.b*) such that by randomizing treatment

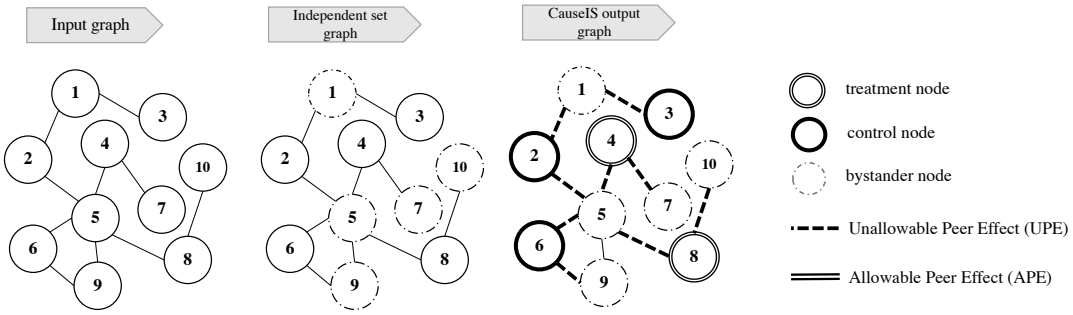

Figure 2: Illustration of *CauseIS* frameworks in network experiments. *Input graph*: a graph of nodes and the connection between them. *Independent set graph*: a graph of bystander and independent set nodes selected by the independent set algorithm. *CauseIS output graph*: the output graph that represents randomized treatment assignment of independent set nodes and peer effects that exists in the experiment.

assignment over the selected subset, the allowable peer effects from the experiment are removed (*Problem 1.c*). We define $s \in \{0, 1\}$ such that $s_i = 1$ if node v_i is in the set of selected nodes, else $s_i = 0$.

$$\begin{aligned} \text{maximize} \quad & \sum_{i=1}^{|V|} s_i \\ \text{subject to} \quad & s_i + s_j \leq 1 \quad \forall e_{i,j} \in E \\ & s_i \in \{0, 1\} \quad \forall v_i \in V \end{aligned}$$

The two constraints together guarantee that adjacent nodes are not included in the in our network experiment design. This optimization can be solved by reducing our problem to the maximum independent set problem in graph theory [7] such that nodes in the independent set correspond to the nodes selected for the network experiment.

Given a graph $G = (\mathbf{V}, \mathbf{E})$, $\mathbf{IS} \subseteq \mathbf{V}$ is a subset of nodes such that for each pair of nodes $v_i \in \mathbf{IS}$ and $v_j \in \mathbf{IS}$ there is no shared edge between them ($e_{i,j} \notin \mathbf{E}$). A *maximal independent set* is an independent set that is not a subset of any other independent sets of the graph. Using a greedy sequential approach, a maximal independent set of a graph can be found in $O(|E|)$ [3] but there are parallel algorithms that can solve this problem much faster in $O(log(N))$ [12, 28]. A maximal independent set with the largest possible size for a given graph is known as a *maximum independent set*. Finding maximum independent sets in graphs is known to be NP-hard. There are exact algorithms that can find maximum independent sets in $O(1.1996^n n^{O(1)})$ [27] and also approximation algorithms that can find it in $O(n/(logn)^2)$ [4].

4.2 *CauseIS* Framework

We propose *CauseIS*, a network experiment design for robust estimation of Direct Treatment Effects by disentangling peer effects from DTE. *CauseIS* has two main steps:

(1) Finding a maximum independent set of the graph (*Independent set graph* in Fig.2)
(2) Assigning nodes of maximum independent set to treatment and control in a randomized fashion (*CauseIS output graph* in Fig.2).

In this framework, we find the treatment assignment vector $\mathbf{Z}$ of nodes by dividing the population to treatment, control and bystander nodes. Considering the proposed objective function, we first use an algorithm to choose the maximum independent set of the given graph that partitions the graph into two set of nodes: 1) nodes in the maximum independent set denoted by *MIS* ($MIS \subseteq \mathbf{V}$) where by randomizing treatment assignment over these nodes, we achieve treatment ($\mathbf{V_1}$) and control ($\mathbf{V_0}$) groups, and 2) bystander nodes ($\mathbf{V_2}$) that are not in *MIS* denoted by $\mathbf{B}$ where $\mathbf{B} \subseteq \mathbf{V}$, $\mathbf{B} \cap \mathbf{MIS} = \emptyset$, and $\mathbf{B} \cup \mathbf{MIS} = \mathbf{V}$. The main idea is to assign nodes of *MIS* to treatment and control at random and ensure that there is no peer effect across treatment and control nodes.

Fig. 2 represents the pipeline of *CauseIS* framework. *Input graph* shows the graph of the network that the network experiment is conducted on. After using an independent set algorithm on the Input graph, independent set and bystander nodes are selected from the graph that are shown in *Independent set graph*. Finally, by randomizing treatment assignment over independent set nodes, treatment and control nodes are selected. *CauseIS output graph* shows the assignment of Input graph nodes to three treatment groups where APE is removed from the experiment.

We remove bystander nodes from the randomized treatment assignment because of the interaction within these nodes which leads to APE in treatment effect estimation. However, it is still possible that information flows from peers in $\mathbf{B}$ to $\mathbf{V}_0$ and $\mathbf{V}_1$, leading to undesired peer effects (nodes 1, 5, 7, 9, 10 in Fig. 2). In the running example, an infected person in $\mathbf{B}$ may infect his peers in $\mathbf{V}_0$ and $\mathbf{V}_1$.

By removing APE from Eq. 6, we have $\hat{TTE}(\mathbf{V}) = DTE(\mathbf{V}) + (UPE(\mathbf{V}_1) - UPE(\mathbf{V}_0))$. By randomizing the treatment assignment over *MIS* nodes, we aim to provide a chance for treatment and control nodes to have the same number of peers in $\mathbf{B}$. We expect that UPE in treatment and control from neighbors in $\mathbf{B}$ cancel each other and by removing APE and UPE from treatment effect estimation, we achieve an unbiased estimation of DTE where TTE=DTE.

5 EXPERIMENTS

In this section, we evaluate the performance of *CauseIS* in direct treatment effect estimation compared to the baselines. We first

describe datasets used in our experiments and then discuss experimental setup and results.

5.1 Data generation

Since existing network datasets do not have ground truth for treatment and its causal effect on outcome, we use synthetic and real-world data structures and simulate the outcome and causal effect in the experiments.

5.1.1 Synthetic data. For generating synthetic networks, we use two graph generation models:

- *Barabási-Albert (BA)* model: This model generates random scale-free networks using preferential attachment model. At the beginning, the network is constructed from m_0 connected nodes. Then, new nodes are connected to m existing nodes with a probability that is proportional to the number of edges that the existing nodes already have [1]. We set $m = 3$ in all experiments.
- *Forest Fire (FF)* model: In this model, a new node v_i attaches to an existing node v_j and then links to nodes connected to v_j with forward and backward burning probabilities denoted by p_f and p_b, respectively. Leskovec et al. [11] show that synthetic network generated by this model can mimic most of real-world structure characteristics. In the experiments, we generate all the graphs with forward burning probability $p_f = 0.3$ and backward burning probability $p_b = 0.3$.

After generating the network structure, we generate 10 attributes for each node with a uniform distribution where the values varies in $[-1, 1]$.

5.1.2 Real-world data. We use five real-world datasets in our experiments. The *50 Women* dataset [14] includes sport, smoking, drug and alcohol habits of 50 students with 74 friendship connections. *Cora* and *Citeseer* datasets [21] incorporate the citation networks of 2,708 and 3312 papers with 4,675 and 5278 edges, respectively. *Hamsterster* dataset [29] includes the online friendship network of 2,059 hamsters with 10,943 edges. *Hateful users* dataset [17] is a a sample of Twitter's retweet graph containing 100,386 users with 1024 attributes and more than two millions retweet edges. In *hateful users* dataset, we remove singletons and nodes with degree 1 from the graph.

5.1.3 Synthetic causal effect. To generate causal effect in our datasets, we generate the outcome of nodes in the real-world and synthetic datasets. We activate (change the outcome of a node from 0 to 1) a treated node with 0.4 probability and a control node with 0.2 probability that makes $DTE = 0.2$. After generating the base activations, we add the contagion process in the next time step and generate UPE and APE for each model. Every inactive node can be activated with some predetermined spillover probability by any of its active peers. A node has as many chances of being activated as the number of neighbors it has. We consider three different possible spillover probabilities for every edge from an activated peer: 0.1, 0.5 and edge weight. Spillover probability denoted by $e.p$ indicates the likelihood of information flow between two peers. For each design, we simulate this process according to the design and measure TTE that is the combination of DTE, UPE and APE. Then we calculate how close the measured TTE is to DTE. To estimate edge weights, we calculate node pairwise similarity of edges' endpoints using one minus the normalized L2 norm: $1 - L_2(v_i.x, v_j.x)$.

5.2 Baselines

We compare the performance of four different approaches in our experiments.

- **Randomized**: In this method, we assign all population nodes to treatment and control groups in a randomized fashion.
- **Match**: In this method, we match nodes using *Best Node Match (BNM)* technique [8]. Node v_i is matched with the most similar unmatched node in the graph. Then, nodes of matched pairs are assigned to treatment or control at random. To estimate the similarity of two nodes, we calculate the pairwise similarity of nodes based on their attributes.
- **CauseIS**: In our proposed framework, we use an algorithm to find the maximum independent set *MIS* and then assign nodes of the set to treatment or control at random.
- **CauseIS_match**: This method uses *CauseIS* framework, but it matches nodes of *MIS* and then assigns nodes of matched pairs to treatment or control at random.

The goal of comparing our method with *Match* and *CauseIS_Match* is to show whether our method has selection bias. In our context, selection bias refers to the difference between the attribute distribution of treatment and control nodes and matching is one of the prominent methods to mitigate this bias in causal studies [23]. Using matching for RCT is unusual, but in small datasets altering the randomization process by posing structural constraints on the graph may lead to worse randomization and matching can mitigate this problem.

5.3 Experimental setup

Our experimental setup follows the experimental setup of previous work [8]. We run a number of experiments to evaluate the accuracy of different methods in estimating treatment effect and the selection bias of these approaches. To measure the estimation error of different methods, we compute *Root Mean Squared Error* (RMSE) of the estimated DTE as:

$$RMSE = \sqrt{\frac{1}{S}\sum_{s=1}^{S}((\hat{\tau}_s - \tau_s)^2)}$$

where S is the number of runs and τ_s and $\hat{\tau}_s$ are true DTE and estimated TTE in run s, respectively. We set $S = 10$ in all experiments. To evaluate the attribute distribution differences between treatment and control nodes, we calculate the Euclidean distance between the attribute vector mean of treated and that of untreated nodes. To show the strength of UPE imposed by bystander nodes in *CauseIS* framework, we calculate the difference between the percentage of edges from bystander nodes to treatment and control nodes as:

$$\frac{1}{|\mathbf{E}|}\Big(\sum_{\substack{e_{i,j}\in\mathbf{E}\\ v_i\in\mathbf{T}\\ v_j\in\mathbf{B}}} d_{i,j} - \sum_{\substack{e_{i,j}\in\mathbf{E}\\ v_i\in\mathbf{C}\\ v_j\in\mathbf{B}}} d_{i,j}\Big) \times 100 \tag{7}$$

where $d_{i,j} = 1$ if there is an edge between node v_i and v_j. T and C shows the vector of treatment and control nodes.

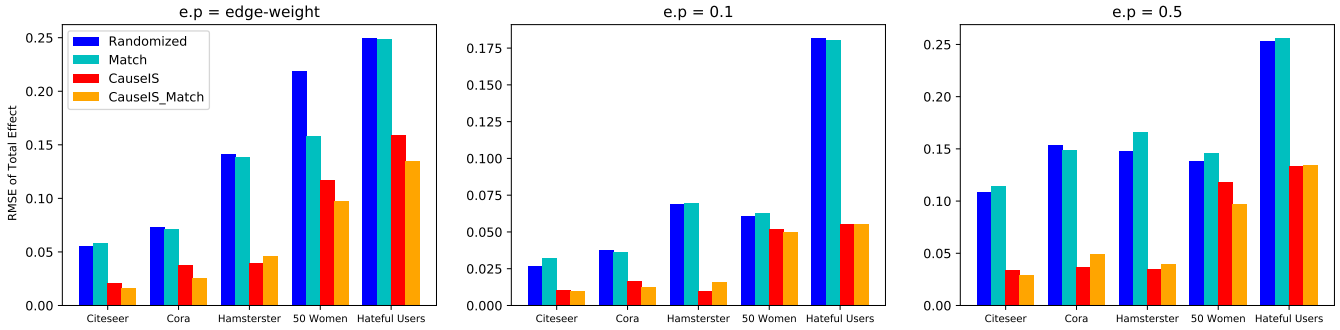

Figure 3: RMSE of direct treatment effect in real-world datasets considering different unallowable peer effect probabilities.

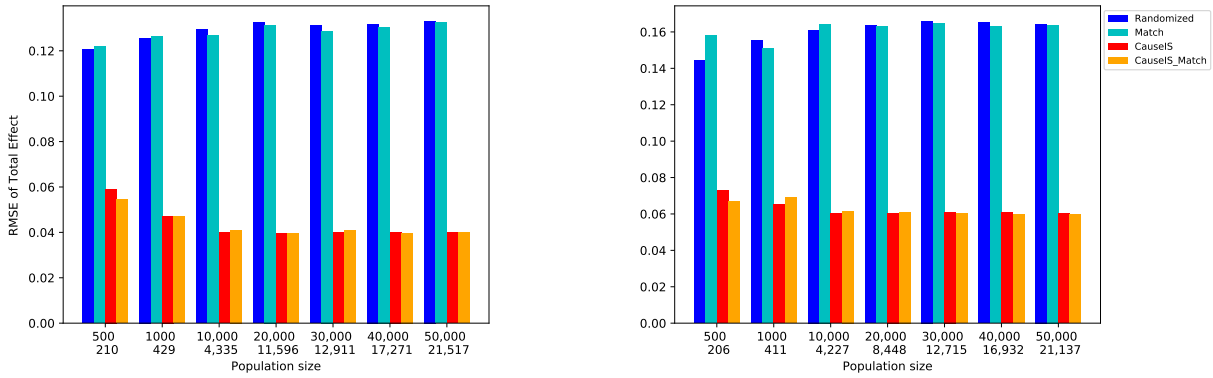

(a) Forest Fire model **(b) Barabasi-Albert model**

Figure 4: RMSE of direct treatment effect in synthetic data with different number of nodes and edges. Numbers in the first row of x-axis shows the number of nodes in graphs, and the second row represents the size of MIS.

in our experiments, we use the *maximal_independent_set* function from the *NetworkX* Python library to find a maximal independent set of each graphs which implements the approach by Blelloch et al. [3]. All the results of synthetic datasets are averaged over 10 runs.

5.4 Results

Evaluation of Direct Treatment Effect estimation: To assess the accuracy of *CauseIS* in estimating DTE compared to the baselines, we measure causal effect estimation error for different unallowable peer effect probabilities. Fig. 3 shows RMSE of DTE in real-world data sets. In all five datasets, *CauseIS* and *CauseIS_Match* get lower estimation error, compared to *Randomized* and *Match*, especially in *Hamsterster* with 72.1% and 76.6% estimated error reduction for $e.p$ = *edge_weight* and $e.p$ = 0.5 and *Hateful Users* with 69.4% estimated error reduction for $e.p$ = 0.1. By increasing the spillover probability from 0.1 to 0.5, we get higher estimation errors because the probability of changing treatment and control outcomes through peer effects increases.

Synthetic data experiments depict a similar picture. Fig 4 shows the stronger performance of *CauseIS* and *CauseIS_Match* over *Randomized* and *Match* methods in reducing causal effect estimation error. For example, *CauseIS*'s error is more than half of the error of *Randomized* approach (0.04 vs. 0.12 for graphs with 10,000 nodes, 0.13 vs. 0.035 for graphs with 20,000 nodes in Forest Fire model). In graphs with 50,000 nodes, *CauseIS* obtains 63.4% and 69.9% estimation error reduction in Forest Fire and Barabasi-Albert models respectively, compared to other graphs.

In both synthetic and real-world datasets, *Randomized* and *Match* in one hand and *CauseIS* and *CauseIS_Match* in other hand show similar performances. This is intuitive, because they use similar randomization techniques. While MIS size is approximately half of population size in all datasets, by increasing the size of MIS the estimation error of *CauseIS* is still significantly lower than *Randomized* methods with smaller population size.

Selection bias evaluation: In this experiment, we evaluate the selection bias of different methods by comparing the Euclidean distance between treatment and control nodes' attributes in real-world and synthetic datasets with different population sizes. Fig. 5 shows this comparison on real-world and synthetic data. It is not surprising that *Match* method gets the lowest selection bias in all datasets, because it matches most similar treatment and control nodes based on the similarity of attributes. *CauseIS_Match* have higher selection bias than *Match*, because the number of nodes matched in this approach is less that *Match* method. Although

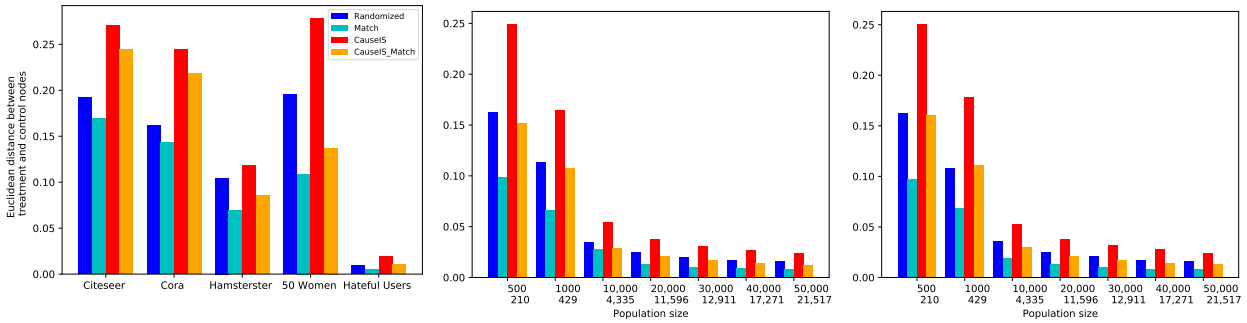

(a) Real-world data (b) Forest Fire model (c) Barabasi-Albert model

Figure 5: Euclidean distance between the attribute vector means of treatment and control nodes in real-world and synthetic datasets. In synthetic dataset plots, numbers in the first row of x-axis show the number of nodes in graphs, and in the second row show the size of MIS.

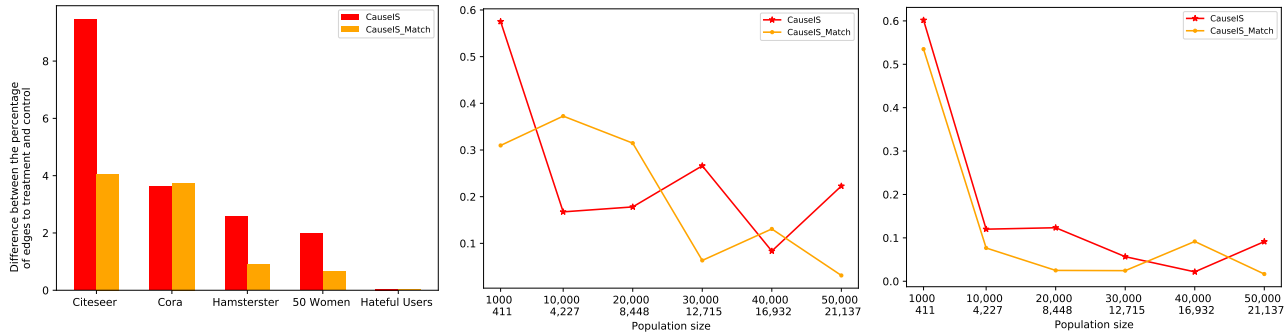

(a) Real-world data (b) Forest Fire model (c) Barabasi-Albert model

Figure 6: Difference between the percentage of edges to treatment and control nodes in real-world and synthetic datasets with different number of nodes and edges. In synthetic dataset plots, numbers in the first row of x-axis show the number of nodes in graphs, and in the second row show the size of MIS.

CauseIS has high selection bias, *CauseIS_Match* reduces selection bias to some extent.

Next, we look at how sample size impacts selection bias. We expect that asymptotically, there would be no selection bias with randomization for any design. Fig. 5 shows that independent from the network generating model, by increasing the population size the similarity between treatment and control nodes' attributes reduces and the value of matching decreases and disappears. For example, in graphs with 500 nodes generated by Forest Fire model, the difference between Euclidean distance of treatment and control nodes in *CauseIS* is 0.24, while in graphs with 50,000 nodes this difference decreases to 0.024. These results confirm the advantage of matching technique in small datasets.

Peer effect evaluation: To measure the extent to which UPE($\mathbf{V}_0$) and UPE($\mathbf{V}_1$) can cancel each other out, we consider the percentage of edges from bystander nodes to treatment and control nodes. Fig. 6 shows this quantity in real-world and synthetic datasets using *CauseIS* and *CauseIS_Match* methods. As expected, results show that for graphs with fewer number of nodes, the difference between the number of edges to treatment and control nodes is higher compared to larger graphs, 2.5 vs. 0.04 in *50 Women* vs. *Hateful Users* dataset. In synthetic data with higher population sizes (40,000 and 50,000), the difference between the percentages of edges to treatment and control is close to zero.

In both synthetic and real-world datasets, we observe that by increasing the sample size, the causal effect estimation error decreases because by increasing the density of the graph edges the percentage of edges from bystander nodes to treatment and control nodes becomes more similar and UPE($\mathbf{V_1}$) - UPE($\mathbf{V_0}$) goes to zero.

Degree distribution evaluation: To assess the extent to which the maximal independent set chosen by *CauseIS* biases the degree distribution of selected treatment and control nodes, we compare the degree distributions of treatment and control nodes selected by *CauseIS* and *Randomized*. Fig.7 shows that *CauseIS* selects treatment and control groups with roughly similar degree distribution in all datasets, except in *50 Women* dataset where the assignment looks more biased, likely due to its small size. *CauseIS* removes high degree nodes from the experiment which results in incorporating treatment and control groups with more balanced degree distribution in the experiments.

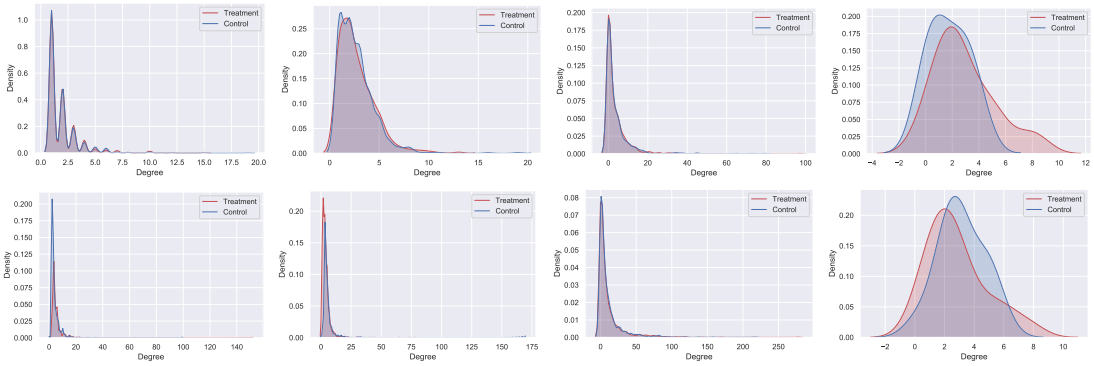

(a) Citeseer **(b) Cora** **(c) Hamsterster** **(d) 50 Women**

Figure 7: Degree distribution of treatment and control nodes selected by *CauseIS* (first row) and *Randomized* (second row).

6 CONCLUSION

We proposed *CauseIS*, a framework that uses independent set explicitly to disentangles peer effects from direct treatment effect estimation and increase the accuracy of direct treatment effect estimation. Our experiments on synthetic and real-world datasets confirms that this approach decreases direct treatment effect estimation error from 61.1% to 76.6% compared to the baselines. We observe that by increasing the population size 1) matching matters less, 2) selection bias is reduced, and 3) unallowable peer effects cancellation is more likely. This work opens many avenues for future research such as accounting for multi-hop contagion in network experiment design.

REFERENCES

[1] Réka Albert and Albert-László Barabási. 2002. Statistical mechanics of complex networks. *Rev. Mod. Phys.* 74 (Jan 2002), 47–97. Issue 1. https://doi.org/10.1103/RevModPhys.74.47

[2] D. Arbour, D. Garant, and D. Jensen. 2016. Inferring network effects in observational data. In *KDD*.

[3] Guy E. Blelloch, Jeremy T. Fineman, and Julian Shun. 2012. Greedy Sequential Maximal Independent Set and Matching Are Parallel on Average. In *Proceedings of the Twenty-Fourth Annual ACM Symposium on Parallelism in Algorithms and Architectures (SPAA âĂŹ12)*. Association for Computing Machinery, New York, NY, USA, 308âĂŞ317. https://doi.org/10.1145/2312005.2312058

[4] R. Boppana and M. M. Halldórsson. 1990. Approximating Maximum Independent Sets by Excluding Subgraphs. In *Proceedings of the Second Scandinavian Workshop on Algorithm Theory (SWAT âĂŹ90)*. Springer-Verlag, Berlin, Heidelberg, 13âĂŞ25.

[5] Dean Eckles, Brian Karrer, and Johan Ugander. 2016. Design and Analysis of Experiments in Networks: Reducing Bias from Interference. *Journal of Causal Inference* 5, 1 (2016), 20150021. https://www.degruyter.com/view/journals/jci/5/1/article-20150021.xml

[6] D. Eckles, R. Kizilcec, and E. Bakshy. 2016. Estimating peer effects in networks with peer encouragement designs. *PNAS* (2016).

[7] Friedrich Eisenbrand, Stefan Funke, and et al. 2003. A Combinatorial Algorithm for Computing a Maximum Independent Set in a t-perfect Graph. In *ACM-SIAM Symp. on Discrete Algorithms*. 517âĂŞ–522.

[8] Zahra Fatemi and Elena Zheleva. 2020. Minimizing Interference and Selection Bias in Network Experiment Design. *Proceedings of the International AAAI Conference on Web and Social Media* 14, 1 (May 2020), 176–186. https://www.aaai.org/ojs/index.php/ICWSM/article/view/7289

[9] G. Imbens and D. Rubin. 2015. *Causal Inference in Statistics, Social and Biomedical Sciences: An Introduction.* Cambridge Univ Press.

[10] Ravi Jagadeesan, Natesh Pillai, and Alexander Volfovsky. 2017. Designs for estimating the treatment effect in networks with interference. (2017). arXiv:math.ST/1705.08524

[11] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph Evolution: Densification and Shrinking Diameters. *ACM Trans. Knowl. Discov. Data* 1, 1 (March 2007), 41. https://doi.org/10.1145/1217299.1217301

[12] M Luby. 1985. A Simple Parallel Algorithm for the Maximal Independent Set Problem. In *Proceedings of the Seventeenth Annual ACM Symposium on Theory of Computing (STOC 85)*. Association for Computing Machinery, New York, NY, USA, 1âĂŞ–10. https://doi.org/10.1145/22145.22146

[13] Charles F. Manski. 1990. Nonparametric Bounds on Treatment Effects. *The American Economic Review* 80, 2 (1990), 319–323. http://www.jstor.org/stable/2006592

[14] Lynn Michell and Amanda Amos. 1997. Girls, pecking order and smoking. *Social Science & Medicine* 44, 12 (1997), 1861 – 1869. https://doi.org/10.1016/S0277-9536(96)00295-X

[15] Elizabeth L Ogburn, Tyler J VanderWeele, et al. 2014. Causal diagrams for interference. *Statistical science* 29, 4 (2014), 559–578.

[16] Jean Pouget-Abadie, Vahab Mirrokni, David C. Parkes, and Edoardo M. Airoldi. 2018. Optimizing Cluster-Based Randomized Experiments under Monotonicity. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining (KDD âĂŹ18)*. Association for Computing Machinery, New York, NY, USA, 2090âĂŞ–2099. https://doi.org/10.1145/3219819.3220067

[17] Manoel Horta Ribeiro, Pedro H Calais, Yuri A Santos, Virgílio AF Almeida, and Wagner Meira Jr. 2018. Characterizing and Detecting Hateful Users on Twitter. *ICWSM* (2018).

[18] Donald B Rubin. 2005. Causal Inference Using Potential Outcomes. *J. Amer. Statist. Assoc.* 100, 469 (2005), 322–331. https://doi.org/10.1198/016214504000001880 arXiv:https://doi.org/10.1198/016214504000001880

[19] Guillaume Saint-Jacques, Maneesh Varshney, Jeremy Simpson, and Ya Xu. 2019. Using Ego-Clusters to Measure Network Effects at LinkedIn. (2019). arXiv:cs.SI/1903.08755

[20] M. Saveski, J. Pouget-Abadie, G. Saint-Jacques, W. Duan, S. Ghosh, Y. Xu, and E. Airoldi. 2017. Detecting network effects: Randomizing over randomized experiments. In *KDD*.

[21] Prithviraj Sen, Galileo Mark Namata, Mustafa Bilgic, Lise Getoor, Brian Gallagher, and Tina Eliassi-Rad. 2008. Collective Classification in Network Data. *AI Magazine* 29, 3 (2008), 93–106.

[22] Uri Shalit, Fredrik D. Johansson, and David Sontag. 2017. Estimating Individual Treatment Effect: Generalization Bounds and Algorithms. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70 (ICML âĂŹ17)*. JMLR.org, 3076âĂŞ3085.

[23] Elizabeth A Stuart. 2010. Matching methods for causal inference: A review and a look forward. *Statistical science* 25, 1 (2010), 1.

[24] Daniel L. Sussman and Edoardo M. Airoldi. 2017. Elements of estimation theory for causal effects in the presence of network interference. (2017). arXiv:stat.ME/1702.03578

[25] Panos Toulis and Edward Kao. 2013. Estimation of causal peer influence effects. In *ICML*. 1489–1497.

[26] J. Ugander, B. Karrer, L. Backstrom, and J. Kleinberg. 2013. Graph cluster randomization: Network exposure to multiple universes. In *KDD*.

[27] Mingyu Xiao and Hiroshi Nagamochi. 2017. Exact algorithms for maximum independent set. *Information and Computation* 255 (2017), 126 – 146. https://doi.org/10.1016/j.ic.2017.06.001

[28] Métivier Yves, John Michael Robson, Saheb-Djahromi Nasser, and Akka Zemmari. 2009. An Optimal Bit Complexity Randomized Distributed MIS Algorithm (Extended Abstract). In *Proceedings of the 16th International Conference on Structural Information and Communication Complexity (SIROCCOâĂŹ09)*. Springer-Verlag, Berlin, Heidelberg, 323âĂŞ337. https://doi.org/10.1007/978-3-642-11476-2_25

[29] E. Zheleva, J., U. Kuter, and L. Getoor. 2008. Using Friendship Ties and Family Circles for Link Prediction. *SNAKDD* (2008).