

GRAPH NEURAL NETWORKS FOR CAUSAL INFERENCE UNDER NETWORK CONFOUNDING

Marc ae P Leun † Pante s Loupos ‡

Marc

Abstract. s p p s s | n n o s on | n W

A n n p o s n s e pos | o n e n n

o po n | o o s n s | on n o e n o p e o p

o s n e s e e e o on s nonp e s p n W

o s e s e e o s o | n o s on s o | s s | s n

o n s on | n W on o n n W e e n W n o e s o ||

n s on s o e s o s | on s e | e p o n n | s e s

on o n n n e s e now | o W e n s on | n on o i

e s o e s n s n | W s | on o | s s o on o e s

o i n on s W s o W p n | n W s e W || s e o

s o W e n s on | n W on o n n W s | s n W n

o o i pp o e sp s n e p e on on s on n e e n e s

e on s e s e o e | s | o W e n s on | s e e s s

1 Introduction

Treatment assignment is said to be unconfounded if it is as good as random within subpopulations of observationally equivalent units. In settings where the stable unit treatment value assumption (SUTVA) is plausible, units with identical covariates are naturally considered observationally equivalent. However, when units are connected through a network, they may differ on other observed dimensions that may confound causal inference if SUTVA is violated and interference is mediated by the network. These dimensions include, for example, the number of type- x neighbors, the number of type- x neighbors with m neighbors of type y , and so on through higher-order neighbors.

Existing formulations of unconfoundedness only utilize a small subset of these dimensions. For example, a common set of controls used in the literature is the vector consisting of own covariates, number of neighbors, and average covariates of neighbors. This choice may be difficult to justify in practice due to a lack of behavioral models of selection. Neighbor covariates may influence selection into treatment in more complex ways not adequately captured by the mean. Furthermore, this choice of controls implies no confounding from higher-order neighbors, which we show rules out economically interesting sources of interference in treatment selection, such as endogenous peer effects.

In this paper, we study estimation and inference for treatment and spillover effects under a fully nonparametric formulation of unconfoundedness motivated by a model of selection. To allow for peer effects, selection is governed by the reduced form of a simultaneous-equations model, which is a function of the entirety of X , the matrix of all units' covariates, and A , the network adjacency matrix. As a result, it is not generally possible to summarize confounding by a simple low-dimensional function of these objects. Our unconfoundedness condition therefore considers units observationally equivalent if they occupy identical positions in the network, meaning that they match on all observed neighborhood and higher-order neighborhood dimensions.

Existing methods that rule out complex forms of interference in selection may result in biased estimates of treatment and spillover effects. For example, consider the causal effect of vaccine adoption on illness. With peer effects in vaccine adoption, vaccinated individuals tend to have more vaccinated direct and indirect social contacts, and a simple comparison of adopters and nonadopters may overstate vaccine

GNNs for Network Confounding

Leung and Loupos

utilize the model of approximate neighborhood interference (ANI) proposed by [Leung \(2022a\)](#), which posits that interference in the outcome stage decays with network path distance. Leung shows that ANI allows for endogenous peer effects but focuses on a setting with randomized assignment. In observational settings, it stands to reason that peer effects in selection may be a possibility. We theref

GNNs for Network Confounding

We provide conditions under which the doubly robust estimator is approximately normally distributed as the network size grows large. This type of result is well known for i.i.d. data (e.g. [Farrell, 2018](#)), but it is nontrivial to extend to our setting since we allow for a complex form of network dependence. For example, asymptotically linearizing the doubly robust estimator requires a new argument due to dependence, and application of an appropriate CLT requires verification of a high-level weak dependence condition under a nonparametric model with outcome and selection stages both governed by simultaneous-equations models. For inference, we utilize a network HAC estimator due to [Kojevnikov et al. \(2021\)](#) and propose a new bandwidth that adjusts for estimation error in the first-stage machine learners.

We substantiate the theory in a simulation study and empirical application to microfinance diffusion. The simulations demonstrate that the use of GNNs can substantially reduce bias relative to conventional choices of network controls even with shallow architectures. The empirical illustration revisits the microfinance diffusion application of [He and Song \(2024\)](#). We show how our estimands can capture complementary aspects of diffusion relative to their “average diffusion at the margin” measure. Our theoretical framework allows for more complex diffusion processes without requiring the econometrician to prespecify the maximum number of within-period rounds of diffusion. Finally, by including richer controls that account for network confounding, we find more attenuated diffusion effects.

1.2 Related Literature

There is a large literature on interference, much of which focuses on randomized control trials (e.g. [Athey et al., 2018](#); [Li and Wager, 2022](#); [Toulis and Kao, 2013](#)). We contribute to a growing recent literature on unconfoundedness, much of which operates in a partial interference setting where units are partitioned into disjoint groups with no interference across groups (e.g. [Liu et al., 2019](#); [Qu et al., 2022](#)).

Studying a network interference setting, [Veitch et al. \(2019\)](#) propose to use “node embeddings” as network controls, which are learned functions of the graph. Since node embeddings can be obtained from a variety of methods, there remains the issue of justifying a particular choice of network controls. GNNs can be interpreted as a method of estimating node embeddings (see §3), and our behavioral model provides justification for their use. We defer to §2.1 a more detailed review of the literature

GNNs for Network Confounding

Let \mathcal{N}_i be the set of units connected through the network A . Each unit $i \in \mathcal{N}_n$ is endowed with unobservables $(\varepsilon_i, \nu_i) \in \mathbb{R}^{d_\varepsilon} \times \mathbb{R}^{d_\nu}$ and observables $\mathbf{X}_i \in \mathbb{R}^{d_x}$. The model primitives determine outcomes and treatments according to

$$Y_i = g_n(i, D, X, A, \varepsilon) \quad \text{and} \quad D_i = h_n(i, X, A, \nu), \quad (1)$$

respectively, where $X = (\mathbf{X}_i)_{i=1}^n$ is the matrix with i th row equal to \mathbf{X}_i' ; Y, D, ε , and ν are similarly defined; and $(g_n, h_n)_{n \in \mathbb{N}}$ is a sequence of function pairs such that each $g_n(\cdot)$ has range \mathbb{R} and $h_n(\cdot)$ has range $[0, 1]$. The econometrician observes (Y, D, X, A) . Our analysis treats (A, X, ε, ν) as random, but the asymptotic theory in §4 conditions on (X, A) to avoid imposing additional assumptions on its dependence structure.

We view the timing of the model as follows. First, nature draws the primitives (A, X, ε, ν) . Next, units select into treatment, potentially based on other units' decisions, and $h_n(\cdot)$ is the reduced-form outcome of that process. Finally, $g_n(\cdot)$ is the reduced form of the subsequent process that generates outcomes. Because $g_n(\cdot)$ and $h_n(\cdot)$ may depend on the primitives of all units, the setup allows Y_i and D_i to be outcomes of simultaneous-equations models with endogenous peer effects, as shown in the next examples.

Example 1 (Linear-in-Means) Consider the outcome model

$$Y_i = \alpha + \frac{\sum_{j=1}^n A_{ij} Y_j}{\sum_{j=1}^n A_{ij}} + \frac{\sum_{j=1}^n A_{ij} Z_j'}{\sum_{j=1}^n A_{ij}} + Z_i' + \varepsilon_i,$$

where $Z_i = (D_i, \mathbf{X}_i)'$ (Manski, 1993). The coefficient captures endogenous peer

denote the row-normalized adjacency matrix and $\mathbf{1}$ the n -dimensional vector of ones, if A is connected, the reduced form of the model can be written in matrix form as

$$Y = \frac{1}{1} \mathbf{1} + Z + \sum_{k=0}^{\infty} \tilde{A}^{k+1} Z + \sum_{k=0}^{\infty} \tilde{A}^k \epsilon.$$

This characterizes Y_i as a function $g_n(i, D, X, A, \epsilon)$.

Example 1 (Binary Game) Consider the binary analog of Example 1 but for selection into treatment:

$$D_i = 1 + \frac{\sum_{j=1}^n A_{ij} D_j}{\sum_{j=1}^n A_{ij}} + \frac{\sum_{j=1}^n A_{ij} Z'_j}{\sum_{j=1}^n A_{ij}} + Z'_i + \epsilon_i \quad (2)$$

Unlike Example 1, there may exist multiple equilibria. The equilibrium selection mechanism is a reduced-form mapping from the primitives (X, A, ν) to outcomes D and therefore characterizes D_i as a function $h_n(i, X, A, \nu)$. This formulation corresponds to a game of complete information. In a game of incomplete information, as modeled by [Xu \(2018\)](#) for instance, a unit i 's information set is (i, X, A) . Here an analog of (2) holds with each D_j replaced with $\hat{D}_j(X, A)$, the equilibrium belief that $D_j = 1$. This characterizes D_i as a function $h_n(i, X, A, \hat{D})$.

Example 2 (Diffusion) [He and Song \(2024\)](#) study the following two-period diffusion model. Let D_i denote i 's decision to adopt microfinance in period 0 and Y_i its decision in period 1. Their equations (2.4) and (3.6) posit that

$$Y_i = g_n(D_{\mathcal{N}(i, K)}, i) \quad \text{and} \quad D_i = 1 + W_i$$

where W_i is a known function of (X, A) and K is the maximum distance that adoption decisions can diffuse through the network between periods 0 and 1. We provide a more detailed comparison of our models in §7.

Given specification (1), we define potential outcomes as

$$Y_i(d) = g_n(i, d, X, A, \epsilon)$$

Confounding may arise first because $Y_i(d)$ is potentially correlated with D_i due to

GNNs for Network Confounding

the high-dimensional observables (X, A) and second because of dependence between unobservables that drive outcomes ε and those that drive selection ν . We restrict the second source of confounding.

Assumption 1 (Unconfoundedness) $\forall n \in \mathbb{N}, \varepsilon, \nu \perp (X, A)$

As discussed below, unconfoundedness conditions used in the existing literature additionally limit the first source of confounding to known summary statistics of (X, A) . Ours is analogous to standard formulations of unconfoundedness under SUTVA ($\varepsilon_i \perp X_i$) since we do not impose further restrictions on the nature of observed confounding.

Because the econometrician only observes a single network, a large-sample theory requires restrictions on interference in order to obtain some form of weak dependence. We next specify a nonparametric model of decaying interference that accommodates the previous examples. For any $S \subseteq \mathcal{N}_n$, let $D_S = (D_{i \in S})$, and similarly define X_S and other such submatrices. Let A_S

counterfactual s -neighborhood outcomes. The error from approximating the observed outcome with the s -neighborhood counterfactual is bounded by $\epsilon_n(s)$, which decays with the neighborhood radius s . This formalizes the idea that Y_i is primarily determined by units relatively proximate to i , so that the further a unit is from i , the less it influences i 's outcome. The second equation imposes the analogous requirement on D_i .

Example 1 For the linear-in-means model in Example 1, an argument similar to Proposition 1 of [Leung \(2022a\)](#) shows that (3) holds with $\sup_n \epsilon_n(s) = C|s|^\alpha$ for some $C > 0$. For the binary game in Example 2, an argument similar to Proposition 2 of [Leung \(2022a\)](#) establishes (4) with $\sup_n \epsilon_n(s)$ decaying at an exponential rate with s . Finally, for the [He and Song \(2024\)](#) diffusion model in Example 3, Y_i only depends on D through $D_{\mathcal{N}(i,K)}$, so (3) holds with $\epsilon_n(s) = c1/s^K$ for some universal constant c . In their empirical application, they use own covariates as controls, so $W_i = X_i$, in which case (4) holds with $\epsilon_n(s) = 0$ for all s .

2.1 Related Literature

The standard SUTVA model and unconfoundedness condition correspond to

$$Y_i = g(D_i, X_i, \epsilon_i) \text{ and } \epsilon_i \perp D_i | X_i. \quad (5)$$

To generalize this setup to allow for network interference, the typical approach in the existing literature is as follows. Define

$$T_i = f_n(i, D, A) \text{ and } W_i = q_n(i, X, A) \quad (6)$$

where $f_n(\cdot)$ and $q_n(\cdot)$ are known vector-valued functions. The $\epsilon_n(s)$ [\(Manski, 2013\)](#) or $\epsilon_n(s) = o_p(\epsilon_n(s))$ [\(Aronow and Samii, 2017, t9\(i\)-0.34923.463\(e\)-297.001000cmBT 2T](#)

GNNs for Network Confounding

o o o n n model and unconfoundedness condition

$$Y_i = g(T_i, W_i, \epsilon_i) \quad \text{and} \quad \epsilon_i \perp T_i | W_i, \quad (7)$$

which is a direct generalization of (5) (Emmenegger et al., 2022; Forastiere et al., 2021; Ogburn et al., 2022). Here T_i entirely summarizes interference while W_i summarizes confounding.

Common examples of T_i and W_i are

$$T_i = D_i, \sum_{j=1}^n A_{ij} D_j \quad \text{and} \quad W_i = X_i, \sum_{j=1}^n A_{ij} X_j, \frac{\sum_{j=1}^n A_{ij} X_j}{\sum_{j=1}^n A_{ij}}. \quad (8)$$

This choice of T_i implies that Y_i depends on D only through two statistics: own treatment and the number of treated neighbors. Variation in the first component identifies a direct treatment effect and variation in the second a spillover effect. Like most exposure mappings used in the literature, this only depends on $D_{\mathcal{N}(i,1)}$, so the outcome model (7) implies no interference beyond the 1-neighborhood. Likewise, this choice of W_i implies no confounding beyond 1-neighborhood covariates.

More generally, one could restrict the outcome model to depend only on the K -neighborhood treatments $D_{\mathcal{N}(i,K)}$ for some fixed threshold K . As shown by Leung (2022a), this rules out economically interesting forms of interference such as endogenous peer effects, which motivates the ANI condition (3). Furthermore, (7) assumes the econometrician can correctly specify the summary statistic T_i in the outcome model, which may be difficult to justify (Sävje, 2024).

Whereas Leung (2022a) and Sävsje (2024) focus on randomized experiments, we study observational data on economic agents that choose to select into treatment. It then becomes important to specify a behavioral model rationalizing the choice of controls W_i . Sánchez-Becerra (2022) is the first to provide such a model. Under neighborhood interference (7) and an exposure mapping similar to (8), he shows that it is sufficient to set $W_i = X_i$, that is, to solely control for own covariates. Since much of the literature utilizes controls such as (8), this raises the question of what model of selection justifies their use or more broadly the use of “network controls” that depend more generally on X and A .

Our model (1) provides an answer. The presence of complex interference in both the outcome and treatment stages induces selection on (X, A) , so that it is generally

GNNs for Network Confounding

with and without at least one treated neighbor, which captures a spillover effect. For $t = (1, 0)$ and $t' = (0, 0)$, it compares the average outcomes of treated and untreated units with no treated neighbors, which captures a treatment effect. For overlap, we need to exclude units with zero degree since a treated neighbor occurs with probability zero for such units. This is accomplished by choosing \mathcal{M}_n to be the subset of units whose degree $n(i, 1) = |\mathcal{N}(i, 1)|$ lies in some desired set excluding zero. That is, choose some $\mathbb{R}_+ \ni 0$ and

$$\mathcal{M}_n = \{i \mid n(i, 1) \in \mathbb{R}_+ \setminus \{0\}\}$$

can be a complex function of (X, A, ν) , and both T_i and Y_i can be complex functions of D , which makes it difficult to characterize the dependence structure necessary for the application of a central limit theorem without additional structure.

Identification results in [Leung \(2024\)](#) provide conditions under which (t, t') has a causal interpretation. The focus of our paper is estimation, so we provide only a brief discussion here. Under the neighborhood interference model (7), (t, t') has a transparent causal interpretation. In settings where (7) fails to hold, [Leung \(2024\)](#) shows that (t, t') retains a causal interpretation under restrictions on interference either in potential outcomes or selection into treatment. For example, suppose treatment adoption follows a nonparametric game of incomplete information where ν_i is unit i 's private information, so that $D_i = h_n(i, X, A, \nu_i)$ (see Example 2). If private information is independent across units conditional on (X, A) , as is typically assumed in structural analyses of the model (e.g. [Lin and Vella, 2021](#); [Xu, 2018](#)), then by Theorem 1 of [Leung \(2024\)](#) (t, t') can be written as a non-negatively weighted average of certain unit-level effects.

Returning to the vaccine adoption example in §1, recall that Adopt_i is

GNNs for Network Confounding

where

$$\hat{\tau}_i(\mathbf{t}, \mathbf{t}') = \frac{1}{\hat{\rho}_t(\mathbf{i}, \mathbf{X}, \mathbf{A})} \left(Y_i - \hat{\mu}_t(\mathbf{i}, \mathbf{X}, \mathbf{A}) \right) - \hat{\mu}_t(\mathbf{i}, \mathbf{X}, \mathbf{A}) + \frac{1}{\hat{\rho}_{t'}(\mathbf{i}, \mathbf{X}, \mathbf{A})} \left(Y_i - \hat{\mu}_{t'}(\mathbf{i}, \mathbf{X}, \mathbf{A}) \right) - \hat{\mu}_{t'}(\mathbf{i}, \mathbf{X}, \mathbf{A}) .$$

\mathbf{t} To estimate

3.1 Architecture

The standard GNN architecture consists of nested, paramete

GNNs for Network Confounding

Leung and Loupos

aggregation no longer holds when the support of \mathbf{X}_i is uncountable, so using multiple aggregators can result in more powerful architectures (Corso et al., 2020).

For an example of $\Gamma(\cdot)$, let $\mu(\cdot)$, $\sigma(\cdot)$, $\Sigma(\cdot)$, $\min(\cdot)$, and $\max(\cdot)$ be respectively the mean, standard deviation, sum, min, and max functions, defined component-wise for multisets of vectors. Then setting $\Gamma(\cdot) = \Gamma_1(\cdot)$ for

$$\Gamma_1(\cdot) = \left(\mu(\cdot), \sigma(\cdot), \Sigma(\cdot), \min(\cdot), \max(\cdot) \right)$$

results in an architecture utilizing five aggregation functions.

The authors combine multiple aggregators with “degree scalars” that multiply each aggregation function by a function of the size of the multiset input $n(i, 1)$. The simplest example is the identity scalar, which maps any multiset to unity. This trivially multiplies each aggregation function in $\Gamma_1(\cdot)$ above, but it is useful to consider non-identity scalars. Let $|\cdot|$ be the function that takes as input a multiset and outputs its size. Corso et al. (2020) define logarithmic amplification and attenuation scalars

$$\mathbf{S}(\cdot) = \left(\frac{\log(|\cdot| + 1)}{\log 2}, \frac{1}{n} \prod_{i=1}^n \log_2 \prod_{j=1}^n \mathbf{A}_{ij} + 1, \dots, |\cdot|, 1, 1 \right)$$

The choice of $\mathbf{S}(\cdot)$ defines whether the scalar “amplifies” ($\alpha > 1$) or “attenuates” ($\alpha < 1$) the aggregation function, and $\alpha = 1$ is the identity scalar. The purpose of the logarithm is to prevent small changes in degree from amplifying gradients in an exponential manner with each successive GNN layer. Thus, an aggregation function that augments $\Gamma_1(\cdot)$ with logarithmic amplification and attenuation is

$$\Gamma_2(\cdot) = \mathbf{S}(\Gamma_1(\cdot))$$

GNNs for Network Confounding

function (:

$$\hat{\mathbf{f}}_{\text{GNN}} = \underset{\mathbf{f} \in \mathcal{F}_{\text{GNN}}(\mathbf{L})}{\text{argmin}} \mathbf{y}_{\text{e}}$$

Leung and Loupos

As discussed in §3.1, any $f \in \mathcal{F}_{\text{GNN}}(\mathcal{L})$ is invariant, so by using GNNs to esti-

GNNs for Network Confounding

that does not depend on i , so evaluating i 's propensity score is now only a matter of supplying the correct i -specific inputs $(i(X), i(A))$.

We close this section with a result demonstrating that invariance is an extremely weak requirement. In particular, it holds under minimal exchangeability conditions on the structural primitives.

Proposition 1. Let $\mathcal{N} = (N, E)$ be a network with n nodes and m edges.

$f_n(i, D, A) = f_n(i, (D, (A, \dots)))$

$g_n(i, D, X, A, \epsilon) = g_n(i, (D, (X, (A, (\epsilon, \dots))))$

Leung and Loupos

conditions imposed below. Define

$$\mu_{t,t'}(i) = \frac{1}{p_t(i, X, A)} \left(T_i \left(Y_i - \mu_t(i, X, A) \right) + \mu_t(i, X, A) \right) + \frac{1}{p_{t'}(i, X, A)} \left(T_i \left(Y_i - \mu_{t'}(i, X, A) \right) + \mu_{t'}(i, X, A) \right) \quad (t, t')$$

whose average over $i \in \mathcal{M}_n$ is the doubly robust moment condition, and let

$$\frac{2}{n} \text{Var} \left(\frac{1}{m_n} \sum_{i \in \mathcal{M}_n} \mu_{t,t'}(i, X, A) \right)$$

Assumption (Moments)

$n \in \mathbb{N}$, $i \in \mathcal{N}_n$, $n \rightarrow \infty$, $d \in \{0, 1\}^n$, $E[Y_i(d) | X, A] = M$
 $\hat{p}_t(i, X, A) \in (0, 1)$, $\hat{p}_{t'}(i, X, A) \in (0, 1)$

GNNs for Network Confounding

The next assumption is used to show that $\mathbb{P}_{t,t'}(\mathbf{i}_{i=1}^n)$ is ϵ -dependent (see Definition C.1, which is due to [Kojevnikov et al., 2021](#)) to apply [m728\(o\)-313.551\(a\)-.823-5.1.\(o\)-313.551\(a\)-.823-5.1.\(](#)

Leung and Loupos

Parts (a) and (b) are used to establish that $\Lambda_{n,t}(i_{i=1}^n)$ is ϵ -dependent. Part (b) is a Lipschitz condition that holds if potential outcomes are uniformly bounded. In particular we can take $\Lambda_n(s) = 2M$ where M is the uniform bound on the ranges

GNNs for Network Confounding

strengthens Assumption 5(b) but only mildly since we nonparametrically estimate both nuisance functions. Since it does not require uniform convergence, it is more readily verifiable for machine learning estimators. Part (c) is Assumption 4.1(iii) of [Kojevnikov et al. \(2021\)](#). Parts (d)–(f) correspond to Assumptions 7(b)–(d) of [Leung \(2022a\)](#), which are used to characterize the bias of the variance estimator. We discuss verification of (c)–(f) in §B.2; the derivations there show that (f) is closely related to (c).

$$E[Y_i | T_i = t, X] = \mu(t, X) + \int \tau(t, t', i) dP_{t'}(i)$$

case $\hat{\epsilon}^2$ would be asymptotically conservative. This can be formalized under additional weak dependence conditions on the superpopulation as in §A of [Leung \(2022b\)](#).

5 Approximate Sparsity

As discussed in §3, the number of layers L in a GNN determines its neighborhood $(\mathcal{X}_{\mathcal{N}(i,L)})$,

GNNs for Network Confounding

$(W_i)_{i=1}^n$ and $\nu = (\nu_i)_{i=1}^n$, define

$$V_i(W, \nu; \beta) = \beta + \frac{\sum_{j=1}^n A_{ij} W_j}{\sum_{j=1}^n A_{ij}} + \frac{\sum_{j=1}^n A_{ij} X_j}{\sum_{j=1}^n A_{ij}} + X_j + \nu_i + \frac{\sum_{j=1}^n A_{ij} \nu_j}{\sum_{j=1}^n A_{ij}}$$

where $(\beta, \nu) \in \mathbb{R} \times \mathbb{R}^n$. We generate $(Y_i)_{i=1}^n$ from the linear-in-means model, where $Y_i = V_i(Y, \varepsilon; \beta)$ and $\varepsilon = (0.5, 0.8, 10, 1)$. We generate $(D_i)_{i=1}^n$ according to Example 2, so that $D_i = 1 - V_i(D, \nu; \alpha)$ with $\alpha = (0.5, 1.5, 1, 1)$. The equilibrium selection mechanism is myopic best-response dynamics starting from the initial condition $(D_i^0)_{i=1}^n$ for $D_i^0 = 1 - V_i(0, \nu; \alpha)$.

The design induces a greater degree of dependence than what our assumptions allow. The error term $\nu_i + \frac{\sum_{j=1}^n A_{ij} \nu_j}{\sum_{j=1}^n A_{ij}}$ is not conditionally independent across units unlike what Assumption 6(a) requires. Also, back-of-the-envelope calculations indicate that peer effects are sufficiently large in magnitude that Assumption 6(d) is violated.

We use the estimand in Example 5 whose true value is zero. About 57 percent of units select into treatment, so the effective sample size used to estimate the outcome regressions is around $n/2$ since $E(Y_i | T_i = t, X, A)$ is estimated only with observations for which $T_i = t$. We report results for $n = 1000, 2000, 4000$.

6.2 Nonparametric Estimators

The GNNs use the PNA architecture in Example 9 with aggregator $\Gamma_2(\cdot)$ defined in the example and $L = 1$,

(2022) and Forastiere et al. (2021). For these, we estimate the nuisance functions using GLMs (logistic and linear regression) with polynomial sieves of order 1, 2, or 3. Recall that a GNN with $L = 1$ corresponds to a receptive field that only encompasses the ego's 1-neighborhood. This is the same as the implied receptive field of the GLM estimators.

Table 1: Simulation results for random geometric graph

		$L = 1$			$L = 2$			$L = 3$		
n	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
H	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI
	e	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI	$\hat{\tau}(1,0)$	CI	CI

5 In situations where the estimator is $\hat{\tau}(1,0) = 0$, we treated $\hat{\tau}(1,0) \approx 0$.

Leung and Loupos

A choice of $L = 2$ is not unusual in the literature. [Zhou et al. \(2021\)](#) compute the prediction error of GNNs on several different datasets with $L = 2, 4, 8, \dots$ and find that $L = 2$ has the best performance across several architectures. The fact that GNN performance often fails to improve (and indeed can worsen) with larger L is well known in the GNN literature, and we survey different expl

GNNs for Network Confounding

study or later.

7.1 Comparison with He and Song (2024)

Leung and Loupos

For the first (second) choice, (t, t') measures the effect of going from 0 to 1 (more than 1) adopting neighbor(s). This sheds light on a different dimension of diffusion relative to the ADM. Whereas the ADM measures how many others are affected by the ego's adoption, our estimands quantify the effect of having multiple adopting neighbors on the ego's adoption. We find below that having multiple adopting neighbors has a much larger effect than having only one.

As previously stated, [He and Song \(2024\)](#) define the treatment in two ways. One is a binary indicator for having a leader in the household, the idea being that all leaders were initially informed about microfinance and told to spread the word. However, not all leaders adopted in the first trimester, which perhaps

GNNs for Network Confounding

To compute the estimates, we concatenate the village networks into a single adjacency matrix of size $n = 4413$. For the GNN and GLM estimates, we trim observations with propensity scores outside of $[0.01, 0.99]$. Standard errors are obtained from the network HAC variance estimator defined in §2.3.

Table 3: Exposure Mapping $T_i^{(1)}$

	ADM	GNN			GLM		
		Layer	Layer	Layer	Order	Order	Order
Leader case							
G_{ee}	0.03	0.00	0.00	0.00	0.00	0.00	0.00
G_{sc}	0.00	0.00	0.00	0.00	0.00	0.00	0.00
G_{all}	0.03	0.00	0.00	0.00	0.00	0.00	0.00
Leader adopter case							
G_{ee}	0.03	0.00	0.00	0.03	0.00	0.00	0.03
G_{sc}	0.00	0.00	0.03	0.00	0.00	0.00	0.00
G_{all}	0.03	0.00	0.00	0.00	0.00	0.00	0.03
Adopter case							
G_{ee}	0.00	0.00	0.00	0.03	0.03	0.00	0.03
G_{sc}	0.00	0.00	0.03	0.00	0.00	0.00	0.00
G_{all}	0.03	0.00	0.00	0.00	0.00	0.00	0.03

$n = 4413$

the leader case.

First consider the estimand using $T_i^{(1)}$, which contrasts microfinance adoption rates for units with 1 versus 0 initially adopting neighbors. The GNN results are consistent across L. For the leader case, we obtain precise zeros for almost all estimates, including the ADM. For the leader-adopter case, the GNN estimates are substantially smaller in magnitude than the ADM with an effect size of at most 10 percentage points compared to the smallest ADM estimate of 20 percentage points. This may be attributed to the use of richer network controls. For the adopter case, the contrast is even starker. Our estimates are an order of magnitude smaller than the corresponding ADM estimates. The GLM estimates are typically slightly smaller than the GNN estimates except for the order-3 polynomials, which are outliers in terms of magnitude.

The estimand using $T_i^{(2)}$ contrasts units with 2+ versus 0 initially adopting neighbors. The estimates for the leader case are similar to those of $T_i^{(1)}$. We find sizeable effects for the leader-adopter case, almost of the same order as ADM, but the robustness of the result is partly tempered by the large amount of trimming discussed below. The adopter case sees estimates of around 20 percentage points, whereas the ADM is double or triple that. Once again the GLM estimates are often slightly smaller relative to the GNN estimates, except for the order-3 polynomials.

The number of observations trimmed for $T_i^{(1)}$ is negligible in the leader and adopter cases. In (f)-373.392792(g)-20792(b)1.82158(le)3.46346(a)-2.372961296(f)-2.013g

in the lasso literature, which posit that a high-dimensional regression function is well-approximated by a function of a relatively small number of covariates.

A Additional Results on GNNs

Primitive conditions for Assumption 5 appear to be beyond the scope of the existing GNN literature, but we provide some potentially useful intermediate results. Consider the problem of establishing a rate of convergence for the propensity score:

$$\frac{1}{m_n} \sum_{i \in \mathcal{M}_n} \left(\hat{p}_t(i, X, A) - p_t(i, X, A) \right)^2 = o_p(n^{-1/2}).$$

Under network approximate sparsity (Definition 1), the problem simplifies to showing

$$\frac{1}{m_n} \sum_{i \in \mathcal{M}_n} \left(\hat{p}_t(i, X, A) - p_t(i, X_{\mathcal{N}(i,L)}, A_{\mathcal{N}(i,L)}) \right)^2 = o_p(n^{-1/2}). \quad (\text{A.1})$$

Since $\hat{p}_t(i, X, A)$ is an L-layer GNN, which only uses information from $(X_{\mathcal{N}(i,L)}, A_{\mathcal{N}(i,L)})$, this should well approximate $p_t(i, X_{\mathcal{N}(i,L)}, A_{\mathcal{N}(i,L)})$ under appropriate conditions, so (A.1) should be more feasible to verify directly.

Farrell et al. (2021) provide a bound analogous to (A.1) for MLPs, which, were it applicable to our setting, would be of the form

$$\frac{1}{n} \sum_{i=1} \left(\hat{p}_t(i, X, A) - p_t(i, X_{\mathcal{N}(i,L)}, A_{\mathcal{N}(i,L)}) \right)^2 \leq C \frac{WL \log R}{n} \log n + \frac{\log \log n}{n} + \epsilon^2 \quad (\text{A.2})$$

with probability at least $1 - e^{-n}$. Here W is the number of GNN parameters, C is a constant that does not depend on n , R depends on the architecture through the number of hidden neurons, and ϵ is the function approximation error, a measure of the ability of the neural network to approximate any function in a desired class. Establishing a corresponding result for GNNs requires an analog of Lemma 6 of Farrell et al. (2021), which is a bound on the pseudo-dimension of the GNN class, and concentration inequalities for ϵ -dependent data. Jegelka (2022) surveys the few available complexity and generalization bounds for GNNs. These are not sufficiently general for our setup and only apply to settings where the sample consists of many independent networks.

GNNs for Network Confounding

To use a bound of the form (A.2) to verify Assumption 5, we require knowledge of how \mathcal{F} varies with key aspects of the architecture, such as W, R, L, n . As a first step toward obtaining such a result, it is necessary to characterize the function class that GNNs can approximate. Our next result, which draws heavily from existing results in the GNN literature, shows that an additional shape restriction on the function class beyond invariance (§3.3) is required.

A.1 WL Function Class

MLPs can approximate any measurable function (Hornik et al., 1989), so given the discussion in §3.3, a natural question is whether GNNs can approximate any measurable, $n \times n$ function of graph-structured inputs. In other words, is it enough to require invariance (and regularity conditions), or are stronger restrictions on the function class necessary? For reasons related to the graph isomorphism problem, it turns out stronger restrictions are necessary. We next motivate the need for such restrictions and then state our function approximation result.

Chen et al. (2019) show that, for a function class such as GNNs to approximate any invariant function, some element of the class must be able to separate any pair of non-isomorphic graphs. By “separate,” we mean that for any non-isomorphic “labeled graphs” $(X, A), (X', A')$, the function f satisfies $f(X, A) \neq f(X', A')$. Hence, a function with separating power of this sort solves the graph isomorphism problem, a problem for which no known polynomial-time solution exists (Kobler et al., 2012; Morris et al., 2021). Since GNNs can be computed in polynomial time, this suggests that approximating any invariant function is too demanding of a requirement.

To define the subclass of invariant functions that GNNs can approximate, we need to take a detour and discuss graph isomorphism tests. The subclass will be defined by a weaker graph separation criterion than solving the graph isomorphism problem, in particular one defined by the W -isomorphism test. This is a (generally imperfect) test for graph isomorphism on which almost all practical graph isomorphism solvers are based (Morris et al., 2021).

Given a labeled graph (X, A) , the WL test outputs a graph coloring (a vector of

Leung and Loupos

labels for each unit) according to the following recursive procedure, whose definition

GNNs for Network Confounding

H^2 such that

$$(h, h') \in \mathcal{F} \text{ if and only if } f(h) \neq f(h') \text{ for all } f \in \mathcal{F}.$$

For any two sets of functions \mathcal{E}, \mathcal{F} with domain H , we say that \mathcal{E} is *at most as separating as* \mathcal{F} if $\mathcal{F} \subseteq \mathcal{E}$.

This is essentially Definition 2 of [Azizian and Lelarge \(2021\)](#). Intuitively, if \mathcal{E} is at most as separating as \mathcal{F} , the latter is more complex in the sense that some function in \mathcal{F} can separate weakly more elements of H than any function in \mathcal{E} .

Let $f_{L,L}$ denote the function of (X, A) with range Σ^n that outputs the vector of node colorings from the WL test run for L iterations. Let $\mathcal{C}(H)$ be the set of continuous functions with domain H . For any $L \in \mathbb{N}$, define the *Weakly Separating*

$$\mathcal{F}_L(L) = \{ f \in \mathcal{C}(H) : (f_{L,L} \in \mathcal{F}) \}.$$

This is the set of continuous functions of (X, A) that are at most as separating as the WL test with L iterations.

The next result says that $\rho_t(\cdot)$ and $\mu_t(\cdot)$ can be approximated by

Leung and Loupos

In other words, any function in the class $\mathcal{F}_L(\mathcal{L})$ can be approximated by L-layer GNNs in $\mathcal{F}_{\text{GNN}}(\mathcal{L})$. The result is a consequence of a Stone-Weierstrauss theorem due to [Azizian and Lelarge \(2021\)](#) and a version of the [Morris et al. \(2019\)](#) and [Xu et al. \(2018\)](#) result on the equivalent separation power of GNNs and the WL test. The proof is given below.

The result is essentially Theorem 4 of [Azizian and Lelarge \(](#)

GNNs for Network Confounding

T in place of L) and add an additional MLP layer that implements their equation (26). Here we use the additional MLP layer added to the output of our architecture (see the paragraph prior to the statement of Theorem A.1). In particular, for any $f \in \mathcal{F}_{\text{GNN}}(L)$, consider the mapping

$$(X, A) \mapsto \left(\tilde{y} f(i, X, A), \dots, \tilde{y} f(i, X, A) \right) \in \mathbb{R}^n$$

n times

(their (26) in our notation) corresponds to adding a linear output layer $L+1$ that is implementable by an MLP of the form $L+1(h_i^{(L)}, h_j^{(L)} : j \in \mathcal{N}_n)$. The mapping remains an element of $\mathcal{F}_{\text{GNN}}(L)$, which completes the argument for (A.5).

By Theorems VIII.1 and VIII.4 of [Grohe \(2021\)](#), which use finiteness of the support of X_i , $(\mathcal{F}_{\text{GNN}}(L)) = (\mathcal{F}_{L,L})$. That is, L-layer GNNs have the same separation power as the WL test run for L iterations. ■

A.2 Disadvantages of Depth

The receptive field is the main consideration when selecting L, but Theorem A.1 provides a second consideration, which is imposing a weaker implicit shape restriction. It shows that, for GNNs to approximate a target function well, the target must satisfy a shape restriction stronger than invariance, namely that it is at most as separating as the WL test with L iterations. The larger the choice of L, the weaker the shape restriction imposed. However, there are several reasons why shallow architectures remain preferable.

Now it turns to the point. A natural question is how many iterations are required for the WL test to converge for a given graph, which corresponds to the choice of L for which the shape restriction is weakest. Unfortunately, the answer is not generally known, being determined by the topology of the input graph in a complex manner. However, there is a range of results bounding the number of iterations required for convergence. For instance, [Kiefer and McKay \(2020\)](#) construct graphs for which the WL test requires $n-1$ iterations to converge, so such graphs require $n-1$ layers to obtain the weakest shape restriction. This makes the estimation problem extremely high-dimensional, requiring substantially more layers than what is typically required

GNNs for Network Confounding

than the standard architecture (13) (Dwivedi et al., 2022). These disadvantages may partly explain the common use in practice of the standard architecture with few layers.

B Verifying §8 Assumptions

Leung (2022a), §A, verifies analogs of Assumptions 6(d) and 7(c) from an older working paper version of Kojevnikov et al. (2021). This section repeats the exercise for Assumptions 6(d) and 7(c) and (d). We assume throughout that $\max_n (s^2, n(s \exp(c(1 - 4p)^{-1}s$ for some $c > 0$ and p in Assumption 4(a). As in Leung (2022a), we say a sequence of networks exhibits polynomial neighborhood growth if

$$\sup_n \max_{i \in \mathcal{N}_n} |\mathcal{N}_A(i, s)| \leq Cs^d$$

for some $C > 0, d > 1$

Leung and Loupos

and $D'_B = (D'_j)_{j \in B}$ for any $B \in \mathcal{N}_n$. Using the first equality of (C.3),

$$\begin{aligned}
 p_t(i, X, A) &= P(D'_i = (D_i, D'_i) \mid a, b, V'_i = (V_i, V'_i) \mid \cdot, X, A) \\
 &= P(D'_i \mid a, b, V'_i \mid \cdot, X, A) \\
 &= P(D_i = D'_i \mid X, A) + P(V_i = V'_i \mid X, A) \cdot R_0
 \end{aligned}$$

By (C.3), the right-hand side equals

$$P(D'_i \mid a, b)$$

GNNs for Network Confounding

Combining (C.6) and (C.7) and using the law of iterated expectations,

$$|\mathbf{p}_t(\mathbf{i}, \mathbf{X}, \mathbf{A}) - \mathbf{p}_t(\mathbf{i}, \mathbf{X}_{\mathcal{N}(\mathbf{i}, r_\lambda(s+1))}, \mathbf{A}_{\mathcal{N}(\mathbf{i}, r_\lambda(s+1))})| \leq \frac{1}{n} (\mathbf{s} + 1) \leq 2\mathbf{R}_0.$$

From (C.2) Noting that $\mu_t(\mathbf{i}, \mathbf{X}, \mathbf{A}) = \mathbb{E}[Y_i \mathbf{1}_i(t) | \mathbf{X}, \mathbf{A}] = \mathbf{p}_t(\mathbf{i}, \mathbf{X}, \mathbf{A})$, we first bound the numerator. For $\mathbf{B} \in \mathcal{N}(\mathbf{i}, s)$, define $Y_i' = g_{n(\mathbf{i}, s)}(\mathbf{i}, D_{\mathbf{B}}, X_{\mathbf{B}}, A_{\mathbf{B}}, \epsilon_{\mathbf{B}})$. By Lemma C.2,

$$|\mathbb{E}[Y_i \mathbf{1}_i(t) | \mathbf{X}, \mathbf{A}] - \mathbb{E}[Y_i' \mathbf{1}_i(t) | \mathbf{X}, \mathbf{A}]| \leq \frac{1}{n} (\mathbf{s} + \Lambda_n(\mathbf{i}, s)) \leq \frac{1}{n} (\mathbf{s} + \mathbf{s})$$

Leung and Loupos

where, using Assumption 4(b),

$$\begin{aligned}
 |R_1| &= n(2s + 2) \Lambda_n(\mathbf{i}, s) n(\mathbf{i}, s) n(s) + C'(1 + n(\mathbf{i}, 1) n(s)), \\
 |R_2| &= C n(2s + 1) n(\mathbf{i}, 1) n(2s - 1), \text{ and} \\
 |R_3| &= C''(|R_1| + |R_2|)
 \end{aligned}$$

for some universal $C'' = 0$. Substituting $s = 2$ for s yields the result. ■

$$a C' = n \sum_{\mathbf{B}_i} \mathcal{N}(\mathbf{i}, s, D'_j) h_{n(\mathbf{j}, s)}(\mathbf{j}, X_{\mathbf{B}_j}, A_{\mathbf{B}_j}, \nu_{\mathbf{B}_j}, D'_{\mathbf{B}_i}) (D'_j)$$

GNNs for Network Confounding

for $|R_1| \leq n(s + \Lambda_n(i, s) n(i, s) n(s) \dots$ ■

$a \leq C \sum_{j=1}^n Y_j' D_j' \dots$
 $n \mathbb{N}, i \in \mathcal{N}_n, n \geq s > 0,$

$$E|Y_i|1_i(t = 1_i(t' | X, A) \leq C(1 + n(i, 1) n(s).$$

Proof. Recall the definition of a, b, \dots prior to (C.3). Define $V_i = \sum_{j=1}^n A_{ij} D_j,$
 $V_i' = \sum_{j=1}^n A_{ij} D_j',$ and $\mathcal{C} = \{|D_i - D_i'| \leq \dots, |V_i - V_i'| \leq \dots\}.$ Then

$$E|Y_i|1_i(t = 1_i(t' | X = x, A = a) \leq a + C P(\mathcal{C}^c | X = x, A = a) \quad (C.11)$$

for some universal $C > 0$ by Assumptions 1 and 4(a). By Assumption 3,

$$1_i(t = 1 | D_i \leq |a, b, V_i|, \dots \quad \text{and} \quad 1_i(t' = 1 | D_i' \leq |a, b, V_i'|, \dots$$

Under event $\mathcal{C},$

$$1 | D_i \leq |a, b, V_i|, \dots \subset \left\{ 1 | D_i' \leq (|D_i - D_i'| + |a, b, V_i'|) + (|V_i - V_i'| \vee |D_i'|) \leq 1 | D_i' \leq |a, b, V_i'| + \dots \right\}$$

Leung and Loupos

Lipschitz functions on \mathbb{R}^d , $\text{Lip}(f)$ be the Lipschitz constant of $f \in \mathcal{L}_d$, and

$$\mathcal{P}_n(\mathbf{h}, \mathbf{h}'; \mathbf{s}) = \{(\mathbf{H}, \mathbf{H}') : \mathbf{H}, \mathbf{H}' \in \mathcal{N}_n, \|\mathbf{H} - \mathbf{h}, \mathbf{H}' - \mathbf{h}'\|_A(\mathbf{H}, \mathbf{H}') > \mathbf{s}\}.$$

Define $\mathcal{D}_n(\mathbf{t}, \mathbf{t}') \subset \mathcal{C}^1$. A triangular array $\mathbf{Z}_i^n_{i=1}$ is *on* $\mathcal{D}_n(\mathbf{t}, \mathbf{t}')$ if there exist $C > 0$, and an \mathcal{F}_n -measurable sequence $\mathbf{s}_n(\mathbf{s})_{\mathbf{s} \in \mathbb{N}^d}$ with $\mathbf{s}_n(0) = 1$ for all n such that

$$|\text{Cov}(f(\mathbf{Z}_H), f'(\mathbf{Z}_{H'}) | \mathcal{H}_n) \leq C \mathbf{h} \mathbf{h}' (\|\mathbf{f}\|_\infty + \text{Lip}(f) (\|\mathbf{f}'\|_\infty + \text{Lip}(f')) \mathbf{s}_n(\mathbf{s}) \text{ a.s. (C.12)}$$

for all $n, \mathbf{h}, \mathbf{h}' \in \mathbb{N}^d; \mathbf{s} \geq 0; f \in \mathcal{L}_h; f' \in \mathcal{L}_{h'}$; and $(\mathbf{H}, \mathbf{H}') \in \mathcal{P}_n(\mathbf{h}, \mathbf{h}'; \mathbf{s})$. We call $\mathbf{s}_n(\mathbf{s})$ the *on* $\mathcal{D}_n(\mathbf{t}, \mathbf{t}')$ of $\mathbf{Z}_i^n_{i=1}$.

Let $\mathcal{C}_n(\mathbf{t}, \mathbf{t}') \subset \mathcal{A}$ be *on* $\mathcal{D}_n(\mathbf{t}, \mathbf{t}')$ if $\mathbf{Z}_i^n_{i=1} \in \mathcal{C}_n(\mathbf{t}, \mathbf{t}')$ and $\mathbf{s}_n(\mathbf{s}) \geq 0$ for all $\mathbf{s} \in \mathbb{N}^d$. Define $\mathcal{D}_n(\mathbf{t}, \mathbf{t}') = \{(\mathbf{Z}_i^n_{i \in H}, \mathbf{Z}_i^n_{i \in H'}) : (\mathbf{Z}_i^n_{i \in H}, \mathbf{Z}_i^n_{i \in H'}) \in \mathcal{C}_n(\mathbf{t}, \mathbf{t}')\}$.

Proof. Let \mathcal{F}_n be the σ -algebra generated by $(\mathbf{X}, \mathcal{A})$, $(\mathbf{h}, \mathbf{h}') \in \mathbb{N}^d \times \mathbb{N}^d$, $(f, f') \in \mathcal{L}_h \times \mathcal{L}_{h'}$, $\mathbf{s} \geq 0$, and $(\mathbf{H}, \mathbf{H}') \in \mathcal{P}_n(\mathbf{h}, \mathbf{h}'; \mathbf{s})$. Define $\mathbf{Z}_i = (\mathbf{Z}_i^{\mathbf{t}, \mathbf{t}'})_{i \in H}$, $\mathbf{Z}_i = (\mathbf{Z}_i^{\mathbf{t}, \mathbf{t}'})_{i \in H'}$, and

$$\mathbf{D}_i^{(\mathbf{s})} = \mathbf{h}_{n(i, \mathbf{s})}(\mathbf{i}, \mathbf{X}_{\mathcal{N}(i, \mathbf{s})}, \mathcal{A}_{\mathcal{N}(i, \mathbf{s})}, \nu_{\mathcal{N}(i, \mathbf{s})}).$$

For $\mathbf{D}_{\mathcal{N}(i, \mathbf{s}^1)}^{(\mathbf{s})} = (\mathbf{D}_j^{(\mathbf{s})})_{j \in \mathcal{N}(i, \mathbf{s}^1)}$, let

$$\begin{aligned} \mathbf{1}_i^{(\mathbf{s})}(\mathbf{t}) &= \mathbf{1}_{\mathbf{f}_{n(i, \mathbf{s}/2)}(\mathbf{i}, \mathbf{D}_{\mathcal{N}(i, \mathbf{s}/2)}^{(\mathbf{s}/2)}, \mathcal{A}_{\mathcal{N}(i, \mathbf{s}/2)}, \nu_{\mathcal{N}(i, \mathbf{s}/2)})}(\mathbf{t}), \\ \mathbf{Y}_i^{(\mathbf{s})} &= \mathbf{g}_{n(i, \mathbf{s}/2)}(\mathbf{i}, \mathbf{D}_{\mathcal{N}(i, \mathbf{s}/2)}^{(\mathbf{s}/2)}, \mathbf{X}_{\mathcal{N}(i, \mathbf{s}/2)}, \mathcal{A}_{\mathcal{N}(i, \mathbf{s}/2)}, \nu_{\mathcal{N}(i, \mathbf{s}/2)}), \\ \mathbf{Z}_i^{(\mathbf{s})} &= \frac{\mathbf{1}_i^{(\mathbf{s})}(\mathbf{t}) (\mathbf{Y}_i^{(\mathbf{s})}) \mu_{\mathbf{t}}(\mathbf{i}, \mathbf{X}, \mathcal{A})}{\rho_{\mathbf{t}}(\mathbf{i}, \mathbf{X}, \mathcal{A})} + \mu_{\mathbf{t}}(\mathbf{i}, \mathbf{X}, \mathcal{A}) \\ &= \frac{\mathbf{1}_i^{(\mathbf{s})}(\mathbf{t}') (\mathbf{Y}_i^{(\mathbf{s})}) \mu_{\mathbf{t}'}(\mathbf{i}, \mathbf{X}, \mathcal{A})}{\rho_{\mathbf{t}'}(\mathbf{i}, \mathbf{X}, \mathcal{A})} + \mu_{\mathbf{t}'}(\mathbf{i}, \mathbf{X}, \mathcal{A}) \quad \mathbf{1}_i(\mathbf{t}, \mathbf{t}'). \end{aligned}$$

Finally, let $\mathbf{Z}_i^{(\mathbf{s})} = (\mathbf{Z}_i^{(\mathbf{s})})_{i \in H}$ and $\mathbf{Z}_i^{(\mathbf{s})} = (\mathbf{Z}_i^{(\mathbf{s})})_{i \in H'}$.

GNNs for Network Confounding

By Assumption 6(a), $(\mathbf{Z}_i^{(s/2, \cdot)})_{i \in H}$ and $(\mathbf{Z}_j^{(s/2, \cdot)})_{j \in H^c}$ are independent given \mathcal{F}_n , so

$$\begin{aligned}
 |\text{Cov}(\mathbf{f}, \mathbf{f}' | \mathcal{F}_n)| &= |\text{Cov}(\mathbf{f}^{(s/2)}, \mathbf{f}'^{(s/2)} | \mathcal{F}_n)| + |\text{Cov}(\mathbf{f}^{(s/2)}, \mathbf{f}'^{(s/2)} | \mathcal{F}_n)| \\
 &\leq 2\|\mathbf{f}'\|_\infty \mathbf{E}|\mathbf{f}^{(s/2)} - \mathbf{f}| + 2\|\mathbf{f}\|_\infty \mathbf{E}|\mathbf{f}'^{(s/2)} - \mathbf{f}'| \\
 &\leq 2\mathbf{h}\|\mathbf{f}'\|_\infty \text{Lip}(\mathbf{f}) + \mathbf{h}'\|\mathbf{f}\|_\infty \text{Lip}(\mathbf{f}') \max_{i \in \mathcal{N}}
 \end{aligned}$$

Leung and Loupos

$$\mathcal{L}_h = \mathcal{L}_{h^1, s} = 0, (\mathbf{H}, \mathbf{H}') = \mathcal{P}_n(\mathbf{h}, \mathbf{h}'; \mathbf{s},$$

$$\mathbf{Y}_i^{(s)} = \mathbf{g}_{n(i,s)}(\mathbf{i}, \mathbf{D}_{\mathcal{N}(i,s)}, \mathbf{X}_{\mathcal{N}(i,s)}, \mathbf{A}_{\mathcal{N}(i,s)}, \boldsymbol{\varepsilon}_{\mathcal{N}(i,s)}),$$

$\mathbf{f}(\mathbf{Y}_i \in \mathbf{H}, \mathbf{f}'(\mathbf{Y}_i \in \mathbf{H}^1, \mathbf{f}^{(s)}(\mathbf{Y}_i^{(s)} \in \mathbf{H}, \text{ and } \mathbf{f}'^{(s)}(\mathbf{Y}_i^{(s)} \in \mathbf{H}^1. \text{ By Assumption 6(a),}$

$$\begin{aligned} |\text{Cov}(\cdot, \mathcal{F}'_n) &= |\text{Cov}(\cdot^{(s/2)}, \mathcal{F}'_n) + |\text{Cov}(\cdot^{(s/2)}, \cdot^{(s/2)} \mathcal{F}'_n) \\ &\leq 2\|\mathbf{f}'\|_\infty \mathbf{E}|\cdot^{(s/2)}| \mathcal{F}'_n + 2\|\mathbf{f}\|_\infty \mathbf{E}|\cdot^{(s/2)}| \mathcal{F}'_n \\ &\leq 2\|\mathbf{h}\| \|\mathbf{f}'\|_\infty \text{Lip}(\mathbf{f}) + \|\mathbf{h}'\| \|\mathbf{f}\|_\infty \text{Lip}(\mathbf{f}') \max_{i \in \mathcal{N}_n} \mathbf{E}|\mathbf{Y}_i - \mathbf{Y}_i^{(s/2)}| \mathcal{F}'_n \\ &\leq 2\|\mathbf{h}\| \|\mathbf{f}'\|_\infty \text{Lip}(\mathbf{f}) + \|\mathbf{h}'\| \|\mathbf{f}\|_\infty \text{Lip}(\mathbf{f}') \sqrt{n} \mathbf{s}^{-2}, \end{aligned}$$

the last line using Assumption 2. Given \mathbf{s} -dependence, the claim follows from Corollary A.2 of [Kojevnikov et al.](#) (

Leung and Loupos

$\hat{\rho}_t(i, X, A)$. For some universal constants $C, C' > 0$, $E|\mathbf{R}_{1t}^2$ equals

$$\frac{1}{m_n} \sum_{i \in \mathcal{M}_n} \sum_{j \in \mathcal{M}_n} E \left[\left(Y_i - \mu_i(Y_j - \mu_j) \right)^2 \right] D, X, A \mathbf{1}_i(t) \mathbf{1}_j(t)$$

GNNs for Network Confounding

$\hat{\rho}_t(\mathbf{i}, \mathbf{X}, \mathbf{A})$, and

$$\Delta_i(t) = (\hat{\rho}_t(\mathbf{i}) - \mu_t(\mathbf{i})) \mathbf{p}_t(\mathbf{i}) \mathbf{1}_i(t)$$

References

Ailon and E Yaav, "On the Bottleneck of Graph Neural Networks and its Practical Implications," in "International Conference on Learning Representations" 2021.

Aronow and Casella, "Estimating Average Causal Effects Under General Interference, with Application to a Social Network Experiment," *Annals of Applied Statistics*, 2017, (4), 1912–1947.

Atkeson, DeSantis and Gins, "Exact p-Values for Network Interference," *Journal of the American Statistical Association*, 2018, (521), 230–240.

— Gins, J. J. and E. J. "Using Wasserstein Generative Adversarial Networks for the Design of Monte Carlo Simulations," *Journal of the American Statistical Association*, 2021.

GNNs for Network Confounding

- Bronstein, M., "Do We Need Deep Graph Neural Networks?," [https://towardsdatascience.com/do we need deep graph neural networks be2dec7c7](https://towardsdatascience.com/do-we-need-deep-graph-neural-networks-be2dec7c7) 2020. Accessed: 2022-07-02.
- J. Bruna, C. Conant, and M. Courville, "Geometric Deep Learning: Grids, Groups, Graphs, Geodesics, and Gauges," *Transactions on Machine Learning Research*, 2021.
- Conzatti, M., and J. Bruna, "On the Equivalence Between Graph Isomorphism Testing and Function Approximation with GNNs," in "Advances in Neural Information Processing Systems," Vol. 32 2019.
- Cornu, C., D. C. Durr, E. Dur, and C. Hansmann, "Double/Debiased Machine Learning for Treatment and Structural Parameters," *Journal of the Royal Statistical Society*, 2018, pp. C1–C68.
- Corso, G., C. Cavallari, D. Baccarelli, and M. Courville, "Principal Neighbourhood Aggregation for Graph Nets," in "Advances in Neural Information Processing Systems," Vol. 33 2020, pp. 13260–13271.
- D. Groot, "Reaching a Consensus," *Journal of the American Statistical Association*, 1974, (345), 118–121.
- D. Raiffa, F. C. Garman, J. J. Donovan, and A. Anand, "Identifying Causal Effects in Experiments with Spillovers and Non-compliance," *Journal of the Royal Statistical Society*, 2023, 5 (2), 1589–1624.
- Dwivedi, C. Joshi, Y. Ben-David, and X. Branson, "Benchmarking Graph Neural Networks," *Transactions on Machine Learning Research*, 2022.
- E. N. C. Poon, and B. J. Mann, "Treatment Effect Estimation from Observational Network Data Using Augmented Inverse Probability Weighting and Machine Learning," *Transactions on Machine Learning Research*, 2022.
- Farr, M., "Robust Inference on Average Treatment Effects with Possibly more Covariates than Observations," *Transactions on Machine Learning Research*, 2018.
- , and M. Sra, "Deep Neural Networks for Estimation and Inference," *Journal of the Royal Statistical Society*, 2021, (1), 181–213.

Leung and Loupos

Franco, J. and J. Lesk, "Fast Graph Representation Learning with PyTorch Geometric," *X*, vol. *n*, no. *X*, pp. *0*–*0*, 2019.

Forastieri, E. and F. A. P. S., "Identification and Estimation of Treatment and Interference Effects in Observational Studies on Networks," *Journal of the American Statistical Association*, 2021, vol. (534), pp. 901–918.

Grover, "The Logic of Graph Neural Networks," in "2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science" IEEE 2021, pp. 1–17.

Huang, "Measuring Diffusion over a Large Network," *Journal of the American Statistical Association*, 2024.

Hornik, S., "Multilayer Feedforward Networks are Universal Approximators," *Journal of the American Statistical Association*, 1989,

GNNs for Network Confounding

Y. R. and B. A., "The Iteration Number of Colour Refinement," in "47th International Colloquium on Automata, Languages, and Prog

GNNs for Network Confounding

ann
 X



“Causal Inference for Spatial Treatments,”

X

n

Leung and Loupos

Zop X , "1-WL Expressiveness Is (Almost) All You Need," X n
 X θ θ 5, 2022.