

# Inferring Individual Level Causal Models from Graph-based Relational Time Series

Ryan Rossi,<sup>1</sup> Somdeb Sarkhel,<sup>1</sup> Nesreen Ahmed<sup>2</sup>

<sup>1</sup>Adobe Research

<sup>2</sup>Intel Labs

## Abstract

In this work, we formalize the problem of causal inference over graph-based relational time-series data where each node in the graph has one or more time-series associated to it. We propose causal inference models for this problem that leverage both the graph topology and time-series to accurately estimate local causal effects of nodes. Furthermore, the relational time-series causal inference models are able to estimate local effects for individual nodes by exploiting local node-centric temporal dependencies *and* topological/structural dependencies. We show that simpler causal models that do not consider the graph topology are recovered as special cases of the proposed relational time-series causal inference model. We describe the conditions under which the resulting estimate can be used to estimate a causal effect, and describe how the Durbin-Wu-Hausman test of specification can be used to test for the consistency of the proposed estimator from data. Empirically, we demonstrate the effectiveness of the causal inference models on both synthetic data with known ground-truth and a large-scale observational relational time-series data set collected from Wikipedia.

## Introduction

Causal inference aims to estimate the effect of interventions and is fundamentally important for decision-making in many application domains (Bakshy, Eckles, and Bernstein 2014; Bosch 2012). Most work on estimating causal effects (causal inference) assumes the data is independent and identically distributed (iid) (Pearl 2009). Under such assumptions, there has been a lot of work on using A/B testing, *i.e.*, randomized experimentation for detecting such causal inferences. However, most real-world data has non-iid structures, *e.g.*, web pages have links, social media influencers have friends as well as followers. Recently, there has been work that extends such randomized experimentation techniques for the case of such structured data where instances may be connected (Ugander et al. 2013). However, there are many practical settings where performing an experiment is impossible, impractical, expensive, or time-consuming. For this reason, there has also been recent work on estimating causal effects from observational iid data (Pearl 1998; Cochran and Rubin 1973; Rosenbaum and Rubin 1984; Rubin 2006). One particularly

appealing class of methods for observational causal inference occurs when the observed data comes in the form of a time-series (Gottman, Glass, and Kratochwill 1978; Abadie, Diamond, and Hainmueller 2010), where it is possible to infer the effect of an observed intervention on an individual unit when sufficient conditions are met.

In order for accurate inference of effects in time series, it is often necessary for practitioners to have access to a large number of pre-treatment time periods to ensure suitable accuracy of the employed models. In many settings this requirement cannot be met, either due to a lack of observations or the applicability of only local stationarity assumptions for a time series model. In settings where the observed system consists of a *network* of individuals this shortcoming is often addressed by employing the so-called templating assumption, where all nodes marginal and conditional structure are identical ((Arbour, Garant, and Jensen 2016)), and pooling observations for all units together. However, in many settings observed effects can be heterogeneous with respect to the network structure, *i.e.*, the causal effect of a treatment varies as a function of the position of the node in the network. As a motivating example, consider the case of Wikipedia page views after a major event such as an earthquake. The event is likely to have a dramatic effect on pages that are directly related to earthquakes, *e.g.* "earthquake" and "richter scale", but is unlikely to have an effect on distant pages in the Wikipedia page graph, *e.g.*, the cooking page is unlikely to experience greater traffic. Unfortunately, neither the individual estimation method nor the templated model provide much utility in this setting. There is likely only a small amount of relevant observations on either side of the event time series making individual treatment underpowered. On the other hand, the number of nodes where unaffected implies that using the templated model will result in a diluted estimate of the causal effect.

In this work, we provide a flexible alternative that allows for practitioners to strike a balance between individual level and fully templated models. Specifically, we assume that effects smoothly vary as a function of distance in the network to provide a mechanism that allows for weighted local pooling of observations which increases power. In this way, the model incorporates the fact that effects of a node are topologically dependent, *i.e.*, if measurements are taken at different regions in the graph, then we would obtain different effect measure-

ments. Given an event of interest, e.g., the earthquake in our running example, this model can be employed within a interrupted time series design (Gottman, Glass, and Kratochwill 1978) to infer causal effects.

The remainder of this paper is organized as follows. We first review related work. Next, we formally introduce the problem of relational time-series causal inference. We then describe our approach for local estimation of node effects under this setting, and show how simpler models are recovered as special cases of the proposed causal inference model. We examine the efficacy of our approach through a series of experiments on synthetically generated data. Finally, we provide a demonstration on Wikipedia data assessing the effect of an earthquake on page views.

## Related Work

Work related to ours can broadly be placed into the following categories: temporal relational learning, causal estimation in networks, and network diffusion processes.

The closest work to our own is (Marazopoulou, Maier, and Jensen 2015). However, the approach in (Marazopoulou, Maier, and Jensen 2015) does not focus on estimating causal effects for relational time-series, and only focused on estimating the structure, rather than estimating the effects from a given causal structure. Previous work that leverages graph-based relational time-series data has focused primarily on classification or regression, but do not provide a mechanism for causal inference in such relational time-series data. In particular, existing work focuses on leveraging temporal dependencies to improve predictive performance (Sharan and Neville 2008; Rossi and Neville 2012). However, this work focuses on classification and does not lend itself to producing causal quantities. Similarly, there has also been work on extending relational probability trees to the spatio-temporal domain for relational data that can vary in both space and time (McGovern et al. 2008), but do not provide a mechanism for producing causal estimates. More recently, there has been work on relational time-series regression (Rossi 2018), however, this work does not focus on causal inference.

There is a small but growing literature that concerns itself with modeling causal effects in relational data. (Arbour, Garant, and Jensen 2016) proposed relational covariate adjustment (RCA). RCA extends non-parametric adjustment to relational data, and allows the estimation of a wide range of functional dependencies. (Sherman and Shpitser 2018) generalized non-parametric identification theory to latent variable causal chain graph models. (Ogburn, Shpitser, and Lee 2018) proposed a parameterization for relational social network data, that corresponds to a particular family of graphical models known as chain graphs. They demonstrated the potential of using chain graphs under certain conditions to analyze data with contagion and interference when DAG models are intractable. (Ugander et al. 2013) proposed methods for A/B testing in the context of social networks where the treatment of individuals in the network spills over to neighboring individuals. Their proposed methods use graph clustering to partition the graph into clusters, and the estimate average treatment effects after adjustment under social interference. A key difference between this paper and prior work is that

here we explicitly take into account heterogeneity that is associated with the topology itself. Prior work either corresponds to modeling only individual attributes, or assuming the global (templated) model.

The final line of related work is work studying information diffusion processes, and contagion, e.g. (Leskovec, Adamic, and Huberman 2007; Bakshy et al. 2011; Kempe, Kleinberg, and Tardos 2003; Gomez-Rodriguez, Leskovec, and Krause 2012; Gleich and Rossi 2013). Our work differs from this line in three important aspects. First and foremost, none of the above work deals with relational (graph-based) time-series data. Second, the notion of heterogeneity based on topology can be seen as accounting for the effect of latent homophily in the networks—a distinct phenomenon from contagion. The second key distinction is in the estimand of interest. Finally, work in diffusion and contagion seeks to quantify aggregate measures of diffusion by considering individual effects, in contrast this work seeks to model individual effects as distinct estimands of interest.

## Problem Formulation

Throughout this work we employ the language of potential outcomes, where  $X(0)$  indicates the counterfactual quantity, i.e., the value of  $X$  that would have been observed had treatment been set to 0. We define the *individual effect* as the effect of an intervention on an individual’s outcome, and the *peer effect* as the effect of an intervention on an individual’s immediate neighbors on the individual’s outcome. We summarize the notation introduced in this paper in Table 1.

Assume  $G = \langle V, E \rangle$  is a graph with  $n = |V|$  nodes and  $m = |E|$  edges. Let  $\mathbf{X}$  be a  $n \times t_{\max}$  matrix consisting of  $n$  time-series of length  $t_{\max}$ .

$$\mathbf{X}' = \mathbf{A}\mathbf{X} \quad (1)$$

Hence,  $X'_{it} = \sum_{j \in \Gamma_i} X_{jt}$  for any  $t = 1, \dots, t_{\max}$  where  $\Gamma_i = \{j \mid (i, j) \in E\}$  is the set of neighbors of node  $i$ . Similarly, the relational time-series mean is:

$$\bar{\mathbf{X}} = \mathbf{D}^{-1}(\mathbf{A}\mathbf{X}) \quad (2)$$

where  $\mathbf{D} = \text{diag}(\mathbf{A}\mathbf{e}) = \text{diag}(|\Gamma_1|, \dots, |\Gamma_n|)$  is the diagonal degree matrix  $\mathbf{e}$  is the vector of all ones.

We will assume that our estimand is the individual average treatment effect observed after making an intervention at time  $t$ , i.e.

$$\mathbb{E} \left( \sum_{k=t+1}^{t_{\max}} X_{ik}(1) \right) - \mathbb{E} \left( \sum_{k=t+1}^{t_{\max}} X_{ik}(0) \right)$$

We further assume that nodes whom obey some definition of closeness (proximity) have time series which evolve in a similar fashion, i.e. if two nodes are close in the network their temporal auto-dependence functions are also similar. In this setting, the problem at hand is to leverage nearby nodes to improve the statistical efficiency of the estimate for a particular node, i.e., we would like to learn a function as follows,

$$\hat{X}_{ik} = g_i(X_{i,k-1}) + \sum_{j \in \Gamma_i} g_j(X_{j,k-1})$$

Table 1: Summary of notation.

$G$	a graph $G = (V, E)$ where $V$ is the set of nodes and $E$ is the edge set
$n$	number of nodes in the graph, $n =  V $
$\mathbf{A}$	adjacency matrix of the graph $G$
$\mathbf{X}$	node time-series matrix where each row is a node and columns represent time-series observations
$\mathbf{x}_t$	$n$ -dimensional vector of time-series values for nodes at time $t$
$X_{i,t}$	time-series value of node $i$ at time $t$
$\mathbf{D}$	diagonal degree matrix where $D_{ii} =  \Gamma_i , \forall i \in V$
$w$	temporal lag (window size)
$d(i, j)$	shortest path distance between nodes $i$ and $j$ in $G$
$\gamma$	a decay factor that determines the weight given to nodes (and their time-series) further away in the graph
$\beta$	estimated regression coefficient
$\Gamma_i$	set of neighbors of node $i$
$t_{\max}$	the last observation in the time-series

In this work, we formalize the problem of estimating causal effects from *relational time-series* data (Rossi 2018) as follows:

**Definition 1.** Let  $G = \langle V, E \rangle$  be a graph with  $n = |V|$  nodes and  $m = |E|$  edges; and let  $\mathbf{X}$  be a  $n \times t_{\max}$  matrix of node time-series such that for each  $i \in V$  there is an associated time-series  $X_{i,1}, \dots, X_{i,t_{\max}}$  of length  $t_{\max}$ . We further assume that there is a known intervention that occurs at time  $0 < t_{\text{int}} < t_{\max}$ , and that the causal effect of the intervention is a smoothly varying process centered at some node  $i \in V$ . The problem is to infer the individual and peer effects of node  $i$ .

Unless otherwise mentioned, we assume the following:

**A 1.** The maximum degree  $\Delta(G) = \max\{|\Gamma_1|, \dots, |\Gamma_n|\}$  is bounded by some constant  $c < \infty$

**A 2.** For any three nodes,  $i, j, k \in V$  if  $i$  and  $j$  share  $k$  as a neighbor then the support of the conditional distributions  $p(X_i|X_k)$  and  $p(X_j|X_k)$  overlap, i.e. there is a shared support of the conditional distributions.

**A 3.** No feedback cycles within time-steps

**A 4.** There are no unobserved confounding variables.

## Framework

In this section, we describe our proposed approach, then we describe how current alternative (individual and global/template models) can be viewed as special cases of the framework. Without loss of generality, we will assume linear models throughout this section. Consider the following least squares estimation problem for a single node  $j$ :

$$\min_{\beta} \sum_i^n \sum_{t=2 \dots t_{\max}} (X_{i,t} - \beta X_{i,t-w})^2 \gamma^{d(i,j)} \quad (3)$$

where  $w$  is the temporal lag (window size) and  $d(i, j)$  is defined as the shortest path distance between nodes  $i$  and  $j$  in

the graph. In this problem  $\gamma$  provides the ability to control the extent to which information from other nodes in the network contribute to the estimation of  $\beta$ . At one extreme, as  $\gamma$  approaches zero, the estimate will recover an i.i.d. estimate. At the other, as  $\gamma$  approaches 1, the estimate will pool all instances and a global model is recovered.

A similar definition can be made in the case of measuring peer-effects. Here, we will assume that the estimand of interest is the global average treatment effect, i.e., the counterfactual is that *both* the individual and their peers are treated. We can now define an analogous quantity for the global treatment effect where we have defined  $\beta^I$  as the coefficient for individual effects and  $\beta^P$  as the coefficient for peer values.

$$\min_{\beta} \sum_{i \in \Gamma_j} \sum_{t=2 \dots t_{\max}} (X_{i,t} - \beta X_{i,t-w})^2 \quad (4)$$

where  $\Gamma_j = \{i \in V | (i, j) \in E\}$ .

$$\min_{\beta^I, \beta^P} \sum_i^n \sum_{t=2 \dots t_{\max}} \left( X_{i,t} - \left( \beta^I X_{i,t-w} + \beta^P [\mathbf{D}^{-1} \mathbf{A} \mathbf{X}]_{i,t-w} \right) \right)^2 \gamma^{d(i,j)} \quad (5)$$

where  $[\mathbf{D}^{-1} \mathbf{A} \mathbf{X}]_{i,t-w}$  is the average peer value at time  $t-w$ . Both optimization problems in Eq. 4 and Eq. 5 are easily implemented with off-the-shelf software packages by giving each sample a weight defined by  $\gamma^{d(\cdot, \cdot)}$ .

Existing work modeling causal effects in temporal non-relational (iid) data are recovered as special cases of the proposed approach. We now address each in turn.

**Individual Model.** This model estimates causal inference using only the time-series of the specific node. This approach reduces the problem to standard multivariate time-series modeling (Akaike and Kitagawa 2012). The advantage of this approach is that it allows the practitioner to use long-studied methodology and is essentially assumption free with respect to the form of relational dependence. Within the framework described above, this model corresponds to setting  $\gamma = 0$ . The corresponding effect models are given by

$$\min_{\beta} \sum_{t=2 \dots t_{\max}} (X_{j,t} - \beta X_{j,t-w})^2 \quad (6)$$

**Individual (IID) Peer Model:** Similarly, if we set  $\gamma = 1$ , then the local peer model (Eq. 5) reduces to the individual (iid) peer model:

$$\min_{\beta^I, \beta^P} \sum_{t=2 \dots t_{\max}} \left( X_{j,t} - \left( \beta^I X_{j,t-w} + \beta^P [\mathbf{D}^{-1} \mathbf{A} \mathbf{X}]_{j,t-w} \right) \right)^2 \quad (7)$$

The cost of this approach, however, is that the effective sample size of the data is limited by the number of observations. This trade-off is often untenable for practitioners who work with moderately sized time series and/or are seeking to measure phenomenon which are characterized with small effect sizes.

**Global Model** On the other side of the spectrum is assuming a global, i.e., templated model. In this approach all nodes

are assumed to have the same marginal and conditional distributions, all observations are pooled and treated as if they were generated from a single node and it's neighbors (in the case of peer effect measurements), which corresponds to setting  $\gamma = 1$  in our proposed framework. The corresponding individual and peer effect models are given by

$$\min_{\beta} \sum_i^n \sum_{t \in 2 \dots t_{\max}} (X_{i,t} - \beta X_{i,t-w})^2 \quad (8)$$

**Global Peer Model:** Similarly, if we set  $\gamma = 0$ , then the local peer model (Eq. 5) reduces to the global peer model:

$$\min_{\beta^I, \beta^P} \sum_i^n \sum_{t \in 2 \dots t_{\max}} \left( X_{i,t} - \left( \beta^I X_{i,t-w} + \beta^P [\mathbf{D}^{-1} \mathbf{A} \mathbf{X}]_{i,t-w} \right) \right)^2 \quad (9)$$

The benefit of this approach is that the number of observations available for a single model can be used to infer parameters. Unfortunately, the templating approach comes at a significant cost in terms of the necessary assumptions, which are opaque and difficult to reason over for practitioners. It is straightforward to see that the global and individual (iid) models are special cases of the proposed *local causal inference models*.

### Test of Specification

The consistency of the proposed local modeling assumption hinges on the level of correlation between the target node series and the series of the rest of the network and the value of  $\gamma$ . If  $\gamma$  is set to be too high with respect to the level of correlation between the series the resulting estimate can be very biased.

Since the individual model is consistent, the Durbin-Wu-Hausman test (Durbin 1954) can be used to test correct specification of the proposed model as follows:

$$H(\beta_{local}) = (\beta_{local} - \beta_{ind})^T (\text{Var}(\beta_{ind}) - \text{Var}(\beta_{local}))^{-1} (\beta_{local} - \beta_{ind}) \quad (10)$$

A test of significance can be performed by using the Chi-Square distribution with the degrees of freedom given by the rank of  $(\text{Var}(\beta_{ind}) - \text{Var}(\beta_{local}))$ . By having a test of specification practitioners can make a principled decision on choosing between the local and individual level data without relying on opaque assumptions.

### Inferring Causal Effects

Thus far we have focused on modeling the relational time series, but have yet to address how a causal quantity can be obtained. Given the proposed setting of a relational time series and an observed intervention, a interrupted time series design (Hahn, Todd, and Van der Klaauw 2001), can be employed to estimate the effect under the following assumption:

**A 5.** *For all nodes with  $\gamma > 0$ , the effect of treatment is constant across units.*

We note that while assumption 5 is strong, the condition can be checked by applying the Hausmann test as described

previously. The interrupted time series design in the individual model corresponds to,

$$X_{i,t} \sim \beta^I X_{i,t-w} + \beta^C c X_{i,t-w} + \beta^{C'} c + \epsilon$$

where  $c$  is an indicator function of treatment and  $\beta^C, \beta^{C'}$  is the post treatment offset and slope, respectively. Modeling of both peer and individual effects simultaneously can be achieved by considering an interaction between treatment status and the first order time series differences.

## Experiments

In this section, we first validate the causal models using synthetic data with known ground-truth and then investigate using the causal inference models for estimating local node effects from a real large-scale relational time-series data set.

### Synthetic Graph Experiments

We first evaluate the efficacy of the proposed approach on synthetic data. Throughout we consider networks generated using the following graph models:

1. Erdos-Renyi (Random) with probability of edge existence varied between [0.1, 0.3].
2. Watts-Strogatz (Small-World) with neighborhood size of 5 and rewiring probability set within to [0.15, 0.3]
3. Barabasi-Albert (Scale Free) with power of preferential attachment varied from 1 to 5.

We consider two ground-truth scenarios. In the first, we seek to measure the individual effect, with each observation at time  $i$  generated as

$$x_i \sim \beta x_{i-1} + \epsilon$$

where

$$\epsilon \sim \mathcal{N}(0, 1)$$

Each node's  $\beta$  is given as a draw from a multivariate normal with mean 1, and covariance given by transition probability from nodes  $i$  to  $j$  after 2 steps. Finally, we consider the case of heterogeneous effects with the addition of peer effects. In the second scenario we consider the case of peer effects, with each observation at time  $i$  generated as

$$x_i \sim \beta^I x_{i-1} + \beta^P \mathbf{D}^{-1} \mathbf{A} x_{i-1} + \epsilon$$

where

$$\epsilon \sim \mathcal{N}(0, 1)$$

For all methods we ran 5000 trials and report the root mean squared error of the estimated causal effect from ground truth, i.e.  $\|\beta - \hat{\beta}\|^2$ , where  $\hat{\beta}$  is the inferred coefficient. We compare against a model that models each node's series independently (**ind**) and a templated model (**relational**). For the local model we consider  $\gamma = 0.05$ , which we have found works well across a variety of settings, using cross-validated estimates.

The results can be seen in Figures 1 and 2. We see that across graph topologies, the local estimate provides the lowest error when a small number of observations ( $< 50$ ) is available. When the time series length increases, the individual

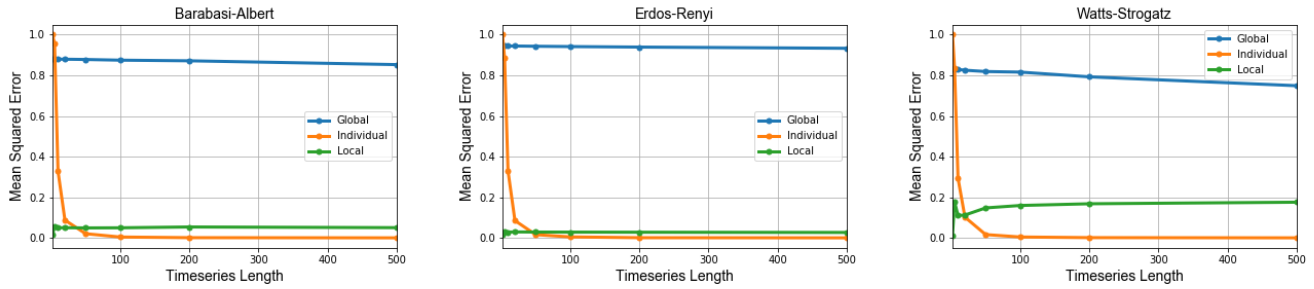


Figure 1: Mean squared error  $((\beta - \hat{\beta})^2)$  under different network models for inferring an individual effect given heterogeneously affected neighbors. The local model provides the lowest error estimates in small sample sizes, and significantly less biased estimates as the sample size is increased. See text for more discussion.



Figure 2: Mean squared error  $((\beta - \hat{\beta})^2)$  under different network models for inferring a peer effect given heterogeneously affected neighbors. While all models converge to the same level of error with larger number of time steps, the local model provides an estimate with lower error than the individual only model at smaller lengths. This, in concert with the results in figure 1, indicate that the local model provides a robust and flexible approach for inference. See text for a full discussion.

model begins to perform better, which is expected given the simplicity of the model and the known consistency of the individual model for large samples. In the case of peer effects we see that all estimates eventually converge, but the individual level model has much poorer performance in moderate sample sizes. This demonstrates the advantages of local pooling: the increased power associated with incorporating local observations decreases the bias substantially in small sample sizes, and by considering only nearby observations we avoid the substantial bias that can result from global pooling.

### Observational Network Data Experiments

We also investigate the causal inference models using real-world network data extracted from Wikipedia consisting of 4,143,840 Wikipedia pages (nodes) with 72,718,664 hyperlinks (edges) between those pages. In addition to the large Wikipedia hyperlink graph described above, we also obtained a time-series for each page (node) representing the hourly page views, *i.e.*, the number of times a page was viewed in a given hour. In other words, each node in the graph has an associated (relational) time-series of hourly page views. There are a total of 48 hours of page view time-series data starting from March 6, 2009 and moving forward in time. In that time period, the average page views is 1.42 whereas the maximum page views is 353,799 for any page at any time. The time-

series of page views for each node can be interpreted as a measure of external interest in Wikipedia pages.

In these experiments, we explore using the approach to infer causality between nodes that would otherwise be impossible to do without considering the graph structure and the time-series associated with each node. One such example is shown in Figure 3, where the effect of interest is over small timescale. However, with the graph structure we can hope to detect and quantify such changes in near real-time. In particular, the page views of the Earthquake page spike at

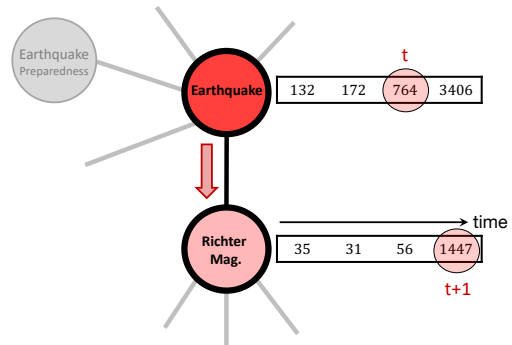


Figure 3: Example with known causal effect between the Earthquake and Richter Magnitude pages.

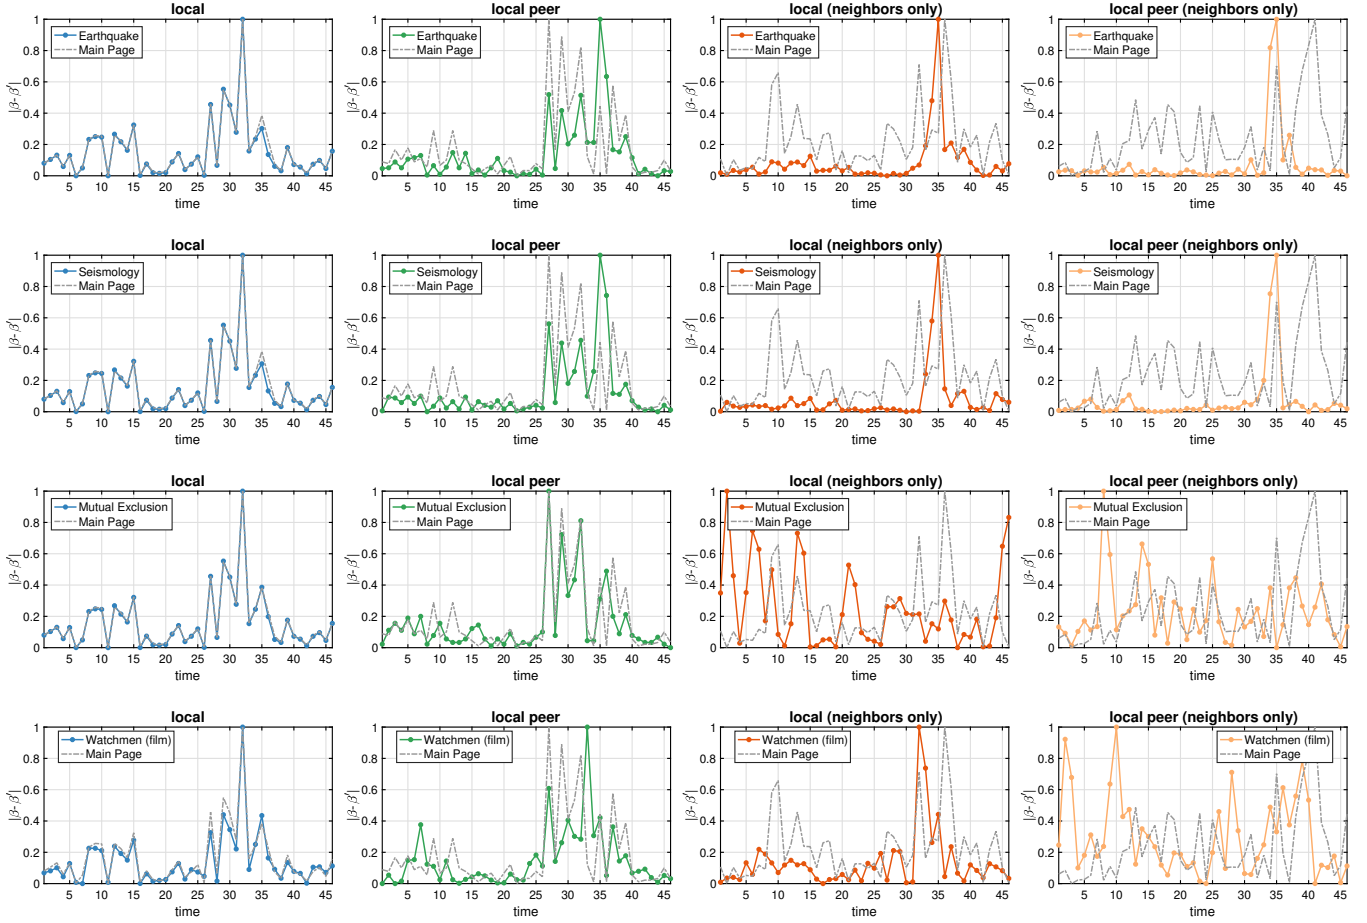


Figure 4: Comparing four different local causal inference models for a variety of Wikipedia pages, namely, the “Main Page”, Earthquake, Seismology, Mutual Exclusion and Watchmen (film) page. The “Main Page”, Mutual Exclusion, and Watchmen (film) pages are unrelated to the Earthquake event that occurred off the coast of Australia. Note all time-series  $\mathbf{x}$  are scaled between 0 and 1 via  $\frac{\mathbf{x} - \min(\mathbf{x})}{\max(\mathbf{x}) - \min(\mathbf{x})}$ . In this experiment, we use 1 hour of data to estimate  $\beta$  and the next hour to estimate  $\beta'$ , and repeat this for all 48 hours by sliding the model to obtain a time-series of  $\beta$ 's (and  $\beta'$ ). See text for discussion.

time  $t$  whereas the page views of Richter Mag. appear normal at time  $t$ . In the next hour, we find that the page views of Richter Magnitude spike, as shown in Figure 3. This implies that immediately after the March 6, 2009 earthquake that occurred in Victoria, Australia users first visited the Earthquake page, and then over time, users clicked on other important and highly related pages associated with this natural disaster, such as Richter Magnitude. This particular case indicates the users were also interested in knowing how big the Earthquake was, a natural question after such an event.

In Figure 4, we compare four different local causal inference models for a variety of Wikipedia pages including the Earthquake, Seismology, Mutual Exclusion (computer science), Watchmen (film), and the “Main Page”. The last three pages are unrelated to the Earthquake that occurred near Victoria, Australia on March 6, 2009. Unless otherwise mentioned, we set  $w = 1$ ,  $\gamma = 0.1$ , and a single time-step is used. As expected, the individual causal inference model that leverages time-series from all nodes in  $G$ , appears very

similar to the Main Page for both Earthquake and Seismology as shown in Figure 4. Interestingly, in both cases where we estimate the causal effect using neighbors only (as opposed to the time-series from all nodes in  $G$ ), we observe a significant difference between the time-series of  $|\beta - \beta'|$  as shown in the two rightmost plots in Figure 4.

In the next set of experiments, we use the causal inference models to estimate the causal effects for 1000 nodes selected uniformly at random as shown in Figure 5 (top). For comparison, we also select the top 1000 nodes given by the difference rank (Eq. 11) and use the models to estimate the causal effect for each of these individual nodes. The difference rank of a node is the difference between its maximum and minimum value in a time-series or a restricted time window  $W$  of that time-series:

$$\mathbf{d} = \max_t(\mathbf{x}_t) - \min_t(\mathbf{x}_t), \quad \mathbf{d}_W = \max_{t \in W}(\mathbf{x}_t) - \min_{t \in W}(\mathbf{x}_t) \quad (11)$$

where  $\mathbf{x}_t \in \mathbb{R}^n$  is the  $n$ -dimensional vector of page views

at time  $t$ , *i.e.*,  $\mathbf{x}_t$  is the  $t$ -th column of  $\mathbf{X}$ . The top 1000 nodes from the difference rank are the nodes with time-series that fluctuate the most. Hence, these nodes often correspond to pages related to recent events. This is in contrast to the majority of other Wikipedia pages with a time series that is relatively stationary with minor fluctuations. In Figure 5, nodes are ordered by  $|\beta - \beta'|$ , and for each node we estimate a  $\beta$  and  $\beta'$ . As expected,  $|\beta - \beta'|$  is typically smaller for nodes selected uniformly at random compared to nodes selected from the difference rank as shown in Figure 5. This is especially evident for the local and local peer (neighbors only) models where  $|\beta - \beta'|$  is very small ( $\approx 10^{-5}$ ) for the last 400 nodes in Figure 5 (top; selected uniformly at random). In contrast, only the last 100 nodes in Figure 5 (bottom) from the difference rank have  $|\beta - \beta'|$  of similar magnitude.

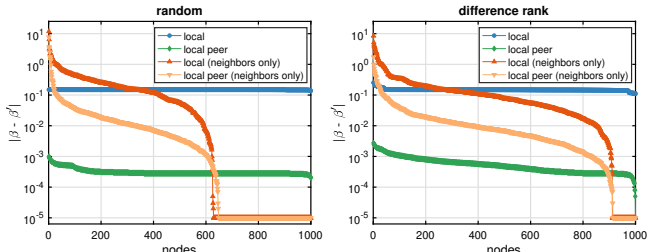


Figure 5: Causal effect estimated for 1000 nodes selected uniformly at random vs. 1000 nodes with largest difference rank (Eq. 11). See text for discussion.

### Effect of $\gamma$

We study the effect of varying  $\gamma$  for the local and local peer effect causal models. To understand the effect of  $\gamma$ , we set  $\gamma \in \{10^{-3}, 10^{-2}, 10^{-1}, 10^0\}$ . Results are shown in Figure 6. In particular, we study two pages with known causal effects (*i.e.*, Earthquake and Richter magnitude scale) and the “Main Page” that has relatively stationary page views. Notice that as  $\gamma \rightarrow 1$ , the  $\beta$ ’s converge to the same quantity regardless of the pages and neighbors/connectivity. Hence, as  $\gamma \rightarrow 1$ , the different causal models all weight the immediate neighbors the same as more distant neighbors or even nodes in different disconnected components in the graph. Conversely as  $\gamma \rightarrow 0$ , the relative difference in weight between nodes that are 1-hop away compared to nodes  $k$ -hops away becomes larger as shown in Figure 6. As an aside, when  $\gamma > 1$ , then nodes further away in the graph from a node  $i$  (larger  $d(i, j)$ ) are given more weight than neighbors close to  $i$  (*e.g.*, immediate neighbors of  $i$  are given less weight than neighbors 2-hops away and so on).

### Conclusion

Given the ubiquity of networks in modern society, developing methods for efficient inference in relational data is critical to understanding and decision making. The vast majority of work produced thus far ignore heterogeneity that can arise as a function of network structure. In this work, we studied the problem of consistent pooling observations to estimate local effects for individual nodes from relational time-series

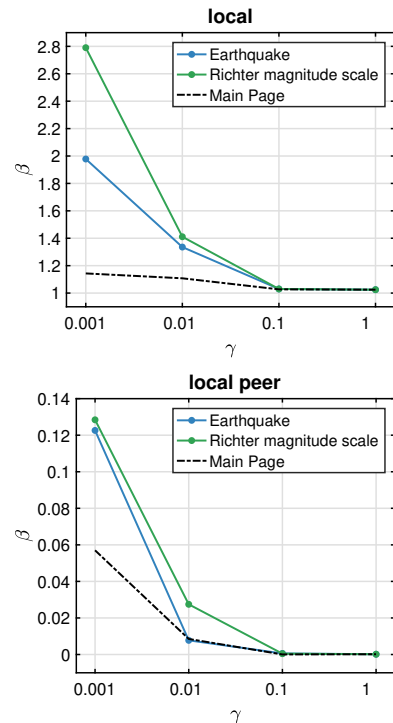


Figure 6: Causal effects as  $\gamma$  varies.

data consisting of a graph (network) where every node is associated with one or more time-series. For this problem, we described a general approach that exploits local node-centric temporal dependencies *and* topological/structural dependencies to accurately estimate effects for a given node. We show that simpler models that do not consider the graph topology are recovered as special cases of the proposed model. We provided a test of specification that allows practitioners to verify the consistency of the local model. This provides practitioners the ability to verify whether the increase in power comes at the cost of a biased model. The experiments demonstrated the effectiveness of the relational time-series causal inference models on both synthetic data with known ground-truth and a large-scale observational relational time-series data set collected from Wikipedia.

### References

- Abadie, A.; Diamond, A.; and Hainmueller, J. 2010. Synthetic control methods for comparative case studies: Estimating the effect of californias tobacco control program. *Journal of the American statistical Association* 105(490):493–505.
- Akaike, H., and Kitagawa, G. 2012. *The practice of time series analysis*. Springer Science & Business Media.
- Arbour, D.; Garant, D.; and Jensen, D. 2016. Inferring network effects from observational data. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 715–724. ACM.
- Bakshy, E.; Hofman, J. M.; Mason, W. A.; and Watts, D. J. 2011. Everyone’s an influencer: quantifying influence on



- twitter. In *Proceedings of the Fourth ACM International Conference on Web Search and Data Mining*, 65–74. ACM.
- Bakshy, E.; Eckles, D.; and Bernstein, M. S. 2014. Designing and deploying online field experiments. In *Proceedings of the 23rd International Conference on World Wide Web*, 283–292.
- Bosch, J. 2012. Building products as innovation experiment systems. In *International Conference of Software Business*, 27–39. Springer.
- Cochran, W. G., and Rubin, D. B. 1973. Controlling bias in observational studies: A review. *Sankhyā: The Indian Journal of Statistics, Series A* 417–446.
- Durbin, J. 1954. Errors in variables. *Revue de l'institut International de Statistique* 23–32.
- Gleich, D. F., and Rossi, R. A. 2013. A dynamical system for pagerank with time-dependent teleportation. *Internet Mathematics* 1–30.
- Gomez-Rodriguez, M.; Leskovec, J.; and Krause, A. 2012. Inferring networks of diffusion and influence. *ACM Transactions on Knowledge Discovery from Data (TKDD)* 5(4):21.
- Gottman, J. M.; Glass, G. V.; and Kratochwill, T. 1978. Analysis of interrupted time-series experiments. *Single subject research: Strategies for evaluating change* 197–234.
- Hahn, J.; Todd, P.; and Van der Klaauw, W. 2001. Identification and estimation of treatment effects with a regression-discontinuity design. *Econometrica* 69(1):201–209.
- Kempe, D.; Kleinberg, J.; and Tardos, É. 2003. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 137–146.
- Leskovec, J.; Adamic, L. A.; and Huberman, B. A. 2007. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)* 1(1):5.
- Marazopoulou, K.; Maier, M.; and Jensen, D. 2015. Learning the structure of causal models with relational and temporal dependence. In *Proceedings of the UAI 2015 Conference on Advances in Causal Inference*, 66–75.
- McGovern, A.; Hiers, N. C.; Collier, M.; Gagne II, D. J.; and Brown, R. A. 2008. Spatiotemporal relational probability trees: An introduction. In *International Conference on Data Mining (ICDM)*, 935–940.
- Ogburn, E. L.; Shpitser, I.; and Lee, Y. 2018. Causal inference, social networks, and chain graphs. *arXiv:1812.04990*.
- Pearl, J. 1998. Graphs, causality, and structural equation models. *Sociological Methods & Research* 27(2):226–284.
- Pearl, J. 2009. *Causality: models, reasoning and inference*, volume 2nd Edition. Cambridge University Press.
- Rosenbaum, P. R., and Rubin, D. B. 1984. Reducing bias in observational studies using subclassification on the propensity score. *Journal of the American statistical Association* 79(387):516–524.
- Rossi, R., and Neville, J. 2012. Time-evolving relational classification and ensemble methods. In Tan, P.-N.; Chawla, S.; Ho, C.; and Bailey, J., eds., *Advances in Knowledge Discovery and Data Mining*, volume 7301 of *Lecture Notes in Computer Science*. Springer Berlin Heidelberg. 1–13.
- Rossi, R. A. 2018. Relational time series forecasting. *Knowledge Engineering Review (KER)* 33:e1.
- Rubin, D. B. 2006. *Matched sampling for causal effects*. Cambridge University Press.
- Sharan, U., and Neville, J. 2008. Temporal-relational classifiers for prediction in evolving domains. In *International Conference on Data Mining (ICDM)*, 540–549.
- Sherman, E., and Shpitser, I. 2018. Identification and estimation of causal effects from dependent data. In *Advances in neural information processing systems*, 9424–9435.
- Ugander, J.; Karrer, B.; Backstrom, L.; and Kleinberg, J. 2013. Graph cluster randomization: Network exposure to multiple universes. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, 329–337. ACM.