



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

Ryan A. Rossi¹, Somdeb Sarkhel¹, Nesreen K. Ahmed²

¹Adobe Research, San Jose, CA

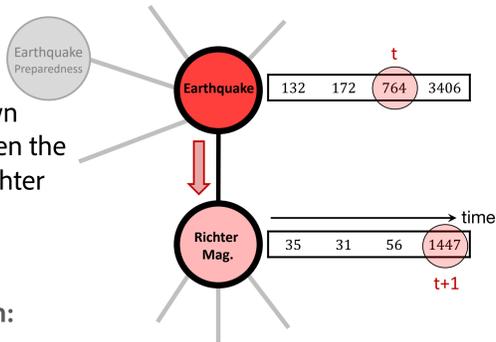
²Intel Labs, Santa Clara, CA

Overview

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.

Causal Inference in Relational Time Series

Let $G=(V,E)$ denote a graph and \mathbf{X} be a n -by- t_{\max} matrix of node time-series. Further, assume that there is a known intervention that occurs at some time and that the causal effect of the intervention is a smoothly varying process centered at some node i . The problem is to infer the individual and peer effects of node i .



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

Basic Model Formulation:

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

Peer Model Formulation:

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

γ provides the ability to control the extent to which information from other nodes in the network contribute to the estimation of β .

- At one extreme, as γ approaches zero, we recover an i.i.d. estimate.
- At the other, as γ approaches 1, the estimate will pool all instances and a global model is recovered.

Individual (IID) Model:

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

Individual (IID) Peer Model:

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

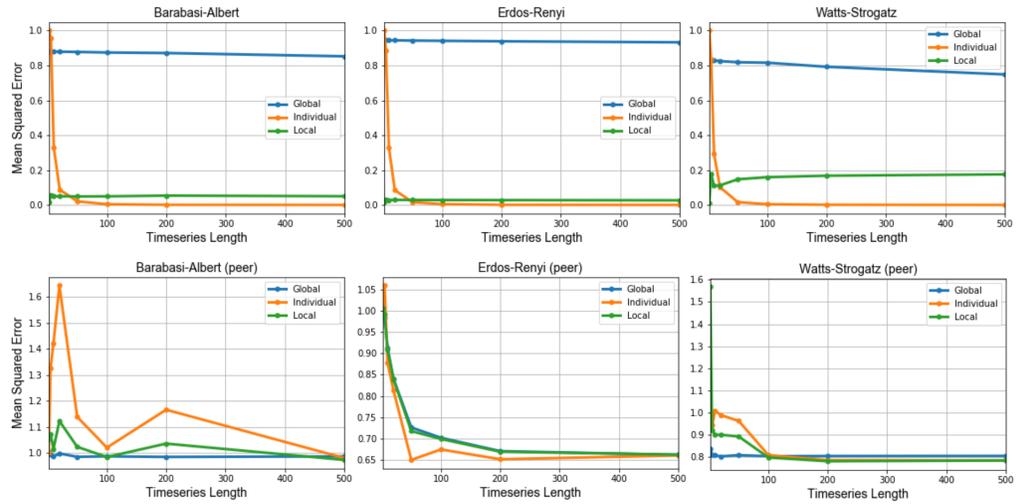
Global Models:

Global Model:

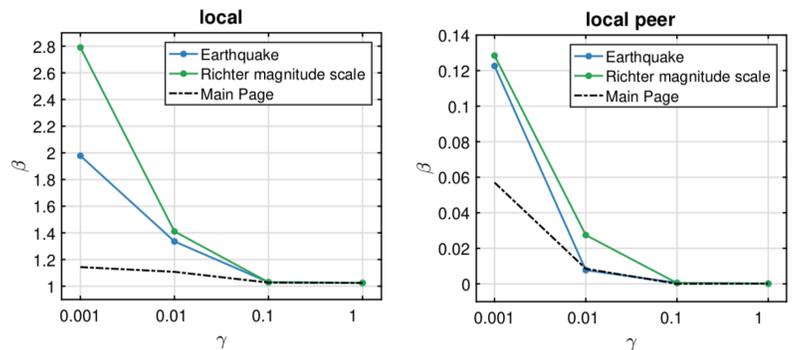
$$\min_{\beta} \sum_i^n \sum_{t \in 2 \dots t_{\max}} (X_{i,t} - \beta X_{i,t-w})^2$$

Global Peer Model:

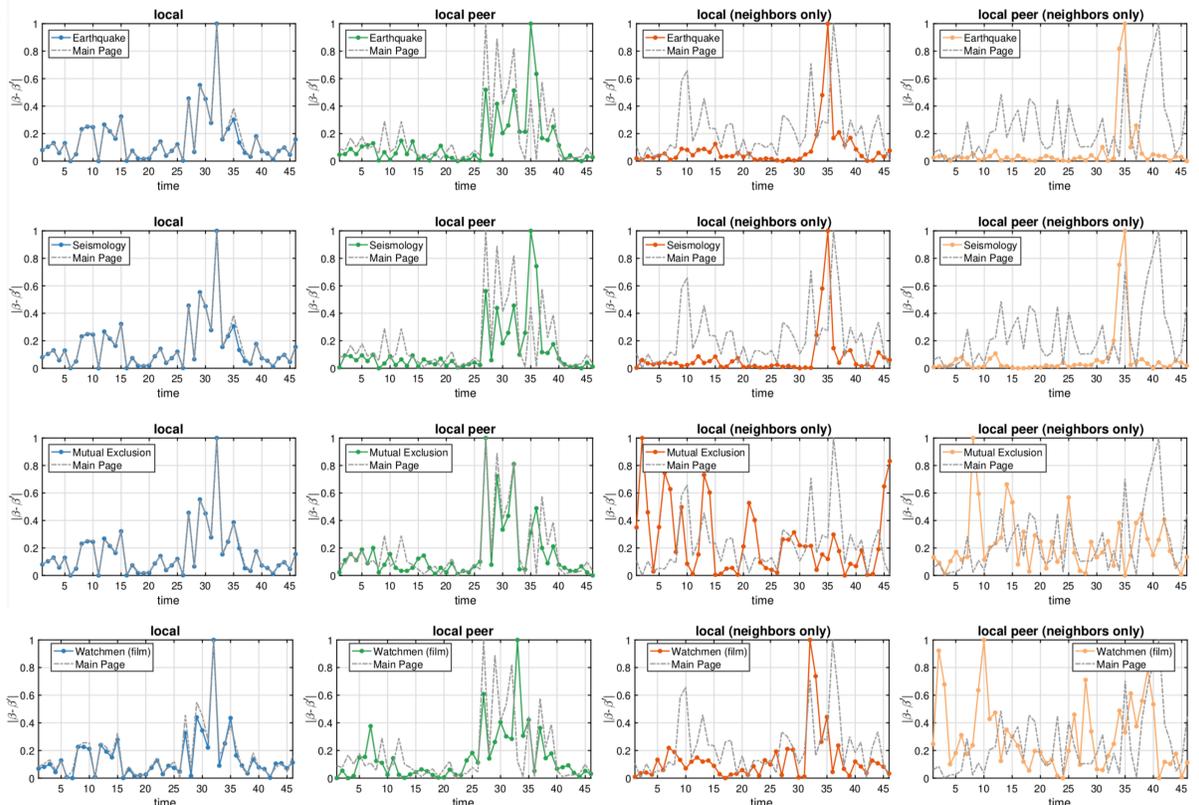
$$\min_{\beta^I, \beta^P} \sum_i^n \sum_{t \in 2 \dots t_{\max}} \left(X_{i,t} - \left(\beta^I X_{i,t-w} + \beta^P [\mathbf{D}^{-1} \mathbf{A} \mathbf{X}]_{i,t-w} \right) \right)^2$$



Mean squared error between actual β and estimated under different network models for inferring a **individual effect** (top) or **peer effect** (bottom) given heterogeneously affected neighbors.



Causal effects as γ varies.



Comparing four different local causal inference models for a variety of Wikipedia pages. 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 x are scaled between 0 and 1 via $x - \min(x) / (\max(x) - \min(x))$. In this experiment, we use 1 hour of data to estimate β and the next hour to estimate β' , and repeat this for all 48 hours by sliding the model to obtain a time-series of β 's (and β').

Main Findings & Contributions

1. We formalized 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.
2. We proposed causal inference models for this problem that leverage both the graph topology and time-series to accurately estimate local causal effects of nodes.
3. We showed 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.