Discovering Heterogeneous Causal Effects in Relational Data

Shishir Adhikari
University of Illinois Chicago
Chicago, IL, USA
sadhik9@uic.edu

Abstract

Causal inference in relational data should account for the non-IID nature of the data and the interference phenomenon, which occurs when a unit's outcome is influenced by the treatments or outcomes of others. Existing solutions to causal inference under interference consider either homogeneous influence from peers or specific heterogeneous influence contexts (e.g., local neighborhood structure). This thesis investigates causal reasoning in relational data and the automated discovery of heterogeneous causal effects under arbitrary heterogeneous peer influence contexts and effect modification.

Introduction

Causal inference is essential for artificial intelligence systems for informed decision-making by anticipating the consequences of actions or policies and understanding underlying mechanisms. Causal inference, in domains like healthcare and public policy, concerns measuring the impact of hypothesized causes or treatments on the outcome as causal effects after analyzing necessary assumptions (Rubin 1974; Pearl 2009). Average causal effect measures the causal impact of a treatment policy for a population while heterogeneous causal effect (HCE) measures the variation in causal impact across subpopulations or contexts, which aids in understanding policy generalization and designing targeted interventions. Significant advancements have been made in the identification and estimation of HCE (Athey and Imbens 2016; Künzel et al. 2019). However, the identification and estimation of HCE from relational data, e.g., social networks, pose unique challenges, and my thesis focuses on investigating and addressing such challenges.

Relational data is abstracted by a *schema* with entity classes (e.g., User or Post), relationship classes (e.g., Friend or Share), and attributes associated with each class. The instantiation of the schema is an *attributed network* with nodes representing entities or units (e.g., social network users) and edges representing their relationships. The major challenge with causal inference in relational data is to account for the *interference* phenomenon, which occurs when treatments or outcomes of other units influence a unit's outcome. For example, the stances of individuals interacting in the social network can be influenced by their political affiliation as well as the affiliations and/or stances of their peers. Even in settings where interference is absent or negligible, causal inference in relational data should accommodate the non-IID nature of the data and the aggregation of relevant covariates.

Under interference, we need to distinguish *individual (or direct) effects* induced by a unit's own treatment, *peer (or indirect) effects* induced by the treatments of other units, and *total effects* induced by both unit's and others' treatments (Hudgens and Halloran 2008). Network experiment designs (Aral 2016) and observational studies (Arbour, Garant, and Jensen 2016; Forastiere, Airoldi, and Mealli 2021) define and control the exposure of treatment from other units for estimating direct, peer, or total effects. Most methods assume homogeneous peer influence and control for peer exposure using the number or fraction of treated peers. However, the influence of peers on an individual can be heterogeneous depending on ego (i.e., individual) and peer traits (e.g., demographic similarity), relationship characteristics (e.g., friendship duration), and network properties (e.g., degree and diversity of friendship). Similarly, an individual's susceptibility to peer influence or response to treatment (e.g., vaccination stance) may vary based on individual characteristics (e.g., age group) and network properties (e.g., friend count). Some recent methods are designed to address heterogeneous peer influence only due to specific contexts such as local neighborhood structure (Yuan, Altenburger, and Kooti 2021), roles (e.g., parent-child, spouses, or siblings) (Qu et al. 2021), and group interactions (Ma et al. 2022). *This thesis aims to characterize sources of heterogeneity and estimate robust heterogeneous causal effects by addressing arbitrary heterogeneous peer influence contexts, different susceptibilities to peer influence, and variable responses to treatment, i.e., effect modifications.*

Current Research Progress

Understanding heterogeneous effects of policy changes on health outcomes from online social network data. This work (Adhikari et al. 2021) presents a quasi-experimental study design, under minimal interference assumption, to test a causal hypothesis and analyze the heterogeneity of causal effects in known contexts using public opinions from Twitter. We test whether recreational cannabis legalization promotes the development of pro-cannabis attitudes for the population in favor of tobacco vaping, as well as analyze how causal effects vary depending on the states implementing the policy and the time since legalization. Here, the causal hypothesis is for the user entity, and we flatten the data from the user's perspective by aggregating tweets created and retweeted by the user to define opinions and stances. The collaborators helped in forming the causal hypothesis, identifying heterogeneous contexts, and collecting the Twitter data while I contributed to the data analysis pipeline and matching-based causal estimation framework.

Discovering causes and effect modifiers in practice. My second work (Adhikari et al. 2022) concentrates on automatically discovering multiple causal hypotheses as well as sources of heterogeneity using retrospective observational data without interference. I investigate a flexible framework that incorporates an ensemble of causal structure learning algorithms for practically discovering causes and effect modifiers that trigger or inhibit an outcome. The framework is applied to the problem of discovering preventive and risk factors as well as vulnerable subpopulations for repeat emergency room visits using medical insurance claims. I provide empirical evidence that the ensemble approach and confidence measures derived from the ensemble approach, as opposed to model selection with untestable assumptions, provide robustness in the discovery task. My collaborators helped in collecting data, defining aggregations, and analyzing discovered causes and effect modifiers.

Individual effects estimation under heterogeneous peer influence. While the first two studies made the simplifying assumption of negligible interference between units, the third work (Adhikari and Zheleva 2023) deals with the complexities of discovering individual effects under heterogeneous interference between units. This work focuses on causal modeling in network experiments or prospective observational studies to identify *relational random variables* that capture arbitrary contexts responsible for heterogeneous peer interference or effect modifications. Then, graph neural networks (GNNs) are utilized for representation learning of these relational variables for robust individual effect estimation controlling heterogeneous peer exposure. To model causal relationships and heterogeneity contexts for interacting units, I propose a flexible and expressive *Network Structural Causal Model (NSCM)* that encodes assumptions about interference conditions, network structure, and causal dependences. NSCM extends the existing relational causal model (RCM) (Maier 2014; Lee 2018) to incorporate relationship existence in networks (e.g., homophily), latent variables, and selection bias. Moreover, the *Network Abstract Ground Graph (NAGG)*, derived from NSCM, enables reasoning about the identification of heterogeneous treatment effects, under arbitrary network interventions, similar to the structural causal model (SCM) (Pearl 2009). The evaluation shows that ignoring heterogeneity of influence introduces bias, and the proposed framework is robust in estimating individual effects in networks.

Anticipated Progress

As the next step, I am working on interference and total effect estimation under heterogeneous peer influence. This work will investigate representation learning of peer exposures to deal with potentially large counterfactual neighborhoods under heterogeneous influence. Furthermore, I plan to discover subpopulations with heterogeneous direct, interference, or total effects in retrospective observational data utilizing NSCM/NAGG. This work will address the challenge of capturing and mapping representations to high-level concepts like diversity, echo chamber, and polarity and identifying subpopulations for targeted interventions.

References

Adhikari, S.; Muscioni, G.; Shapiro, M.; Petrov, P.; and Zheleva, E. 2022. On Discovering Causes and Effect Modifiers in Practice. *Under Submission*.

Adhikari, S.; Uppal, A.; Mermelstein, R.; Berger-Wolf, T.; and Zheleva, E. 2021. Understanding the Dynamics between Vaping and Cannabis Legalization Using Twitter Opinions. In *Proceedings of the International AAAI Conference on Web and Social Media*, volume 15, 14–25.

Adhikari, S.; and Zheleva, E. 2023. Inferring Causal Effects Under Heterogeneous Peer Influence. *arXiv preprint arXiv:2305.17479*.

Aral, S. 2016. Networked experiments. *The Oxford handbook of the economics of networks*, 376–411.

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.

Athey, S.; and Imbens, G. 2016. Recursive partitioning for heterogeneous causal effects. *Proceedings of the National Academy of Sciences*, 113(27): 7353–7360.

Forastiere, L.; Airoldi, E. M.; and Mealli, F. 2021. Identification and estimation of treatment and interference effects in observational studies on networks. *Journal of the American Statistical Association*, 116(534): 901–918.

Hudgens, M. G.; and Halloran, M. E. 2008. Toward causal inference with interference. *Journal of the American Statistical Association*, 103(482): 832–842.

Künzel, S. R.; Sekhon, J. S.; Bickel, P. J.; and Yu, B. 2019. Metalearners for estimating heterogeneous treatment effects using machine learning. *Proceedings of the National Academy of Sciences*, 116(10): 4156–4165.

Lee, S. 2018. *Causal Discovery from Relational Data: Theory and Practice*. PhD thesis, The Pennsylvania State University.

Ma, J.; Wan, M.; Yang, L.; Li, J.; Hecht, B.; and Teevan, J. 2022. Learning causal effects on hypergraphs. In *Proceedings of the 28th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, 1202–1212.

Maier, M. 2014. Causal discovery for relational domains: Representation, reasoning, and learning. *PhD thesis, University of Massachusetts, Amherst*.

Pearl, J. 2009. *Causality*. Cambridge university press.

Qu, Z.; Xiong, R.; Liu, J.; and Imbens, G. 2021. Efficient Treatment Effect Estimation in Observational Studies under Heterogeneous Partial Interference. *arXiv preprint arXiv:2107.12420*.

Rubin, D. B. 1974. Estimating causal effects of treatments in randomized and nonrandomized studies. *Journal of Educational Psychology*, 66(5): 688.

Yuan, Y.; Altenburger, K.; and Kooti, F. 2021. Causal Network Motifs: Identifying Heterogeneous Spillover Effects in A/B Tests. In *Proceedings of the Web Conference 2021*, 3359–3370.