When Causal Inference Meets Graph Machine Learning

Jing Ma
Case Western Reserve University, Cleveland, Ohio, USA
jing.ma5@case.edu

Graphs (i.e., networks) are effective tools to model many real-world systems with connected units, such as social networks and biological networks. Recent years have witnessed rapid development in graph-based machine learning (GML) in various high-impact domains. Currently, the mainstream GML methods are based on statistical learning, e.g., utilizing the statistical correlations between node features, graph structure, and labels for node classification. However, statistical learning has been widely criticized for only capturing the superficial relations between variables, and consequently, rendering the lack of trustworthiness in real-world applications. Therefore, it is crucial to understand the *causality* in the data system and the learning process. Causal inference is the discipline that investigates the causality, for example, to identify and estimate the causal effect of a certain treatment (e.g., wearing a face mask) on an outcome (e.g., COVID-19 infection). Involving causal inference in ML methods is often considered significant for human-level intelligence and can serve as the foundation of artificial intelligence (AI). However, traditional causal inference often relies on strong assumptions, and focuses on independent and identically distributed (i.i.d.) data, while causal inference on graphs faces many barriers. Therefore, there is still a gap between causal inference and GML.

To bridge the gap, we need to address the following challenges: (C1): *Hidden confounders* - Classical causal inference often assumes that hidden confounders (unobserved variables that influence both treatment and outcome) do not exist. However, in practice, hidden confounders widely exist and are often time-varying. This phenomenon can cause serious confounding biases in causal effect estimation. (C2): *Network interference* - Traditional causal inference methods assume that different units do not interfere with each other, or interference only exists between pairs of units. These assumptions are often impractical. Especially, when there are group interactions with more than two units involved, high-order interference may exist in the group. (C3): *Fairness* - The predictions of ML models are often biased towards certain demographic groups w.r.t. sensitive attributes (e.g., age, gender). In graphs, such biases can also be induced by one's neighboring nodes, as well as the causal relations between variables. These make existing fair ML algorithms incapable of mitigating such biases without consideration of graph connections and causality. (C4): *Explainability* - Most GML predictions are made by opaque models without explainability from a causal view. As shown in Fig. 1, my research works mainly cover different aspects of the aforementioned two directions by tackling the challenges (C1, C2, C3, and C4). Specifically, some of them study the following topics and develop corresponding principled approaches: (1) **GML for causal inference**, including a) causal effect estimation under hidden confounders (Ma 2021, 2022d) and its application for COVID-19 policy assessment (Ma 2022a); and b) causal effect estimation under interference (Ma 2022c); (2) **Causality for GML**, including c) counterfactual fairness in graph node representation learning (Ma 2022e); and d) counterfactual explanation for GML models (Ma 2022b).

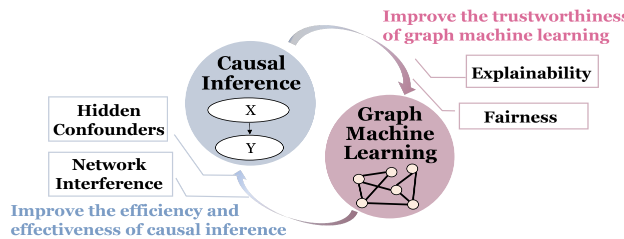

Figure 1: An overview of my research work.

References

Ma, J. e. a. 2021. Deconfounding with networked observational data in a dynamic environment. In *ACM WSDM*, 166–174.

Ma, J. e. a. 2022a. Assessing the causal impact of COVID-19 related policies on outbreak dynamics: A case study in the US. In *WWW*, 2678–2686.

Ma, J. e. a. 2022b. Clear: Generative counterfactual explanations on graphs. *NeurIPS*, 35: 25895–25907.

Ma, J. e. a. 2022c. Learning Causal Effects on Hypergraphs. In *ACM SIGKDD*.

Ma, J. e. a. 2022d. Learning causality with graphs. *AI Magazine*, 43(4): 365–375.

Ma, J. e. a. 2022e. Learning fair node representations with graph counterfactual fairness. In *ACM WSDM*, 695–703.