

Abstract

Name: Hans Jarett J. Ong

Thesis Title: Advancing Causal Discovery in Observed, Latent, and Few-shot Settings

Laboratory: Mathematical Informatics Laboratory

Supervisor: Professor Kazushi Ikeda

While machine learning excels at pattern recognition, its reliance on statistical correlations leaves it vulnerable to out-of-distribution shifts and incapable of predicting interventional outcomes. Graphical causality address this by modeling underlying cause-and-effect mechanisms, yet they require prior knowledge of the causal graph. This dissertation advances *causal discovery* to infer these structures from data across three increasingly complex frontiers. In the fully **observed** setting, a unified beam search framework and a novel data compression-based dependence measure are introduced to overcome the brittleness of standard greedy algorithms and successfully navigate noisy local optima. In the **latent** setting, where causal factors are embedded within unstructured data, the Latent Additive Noise Model Causal Autoencoder (LANCA) is proposed to recover hidden factors from purely unsupervised observational data by explicitly optimizing for residual noise independence. Finally, in the data-scarce **few-shot** setting where intervention targets across environments are unknown, MetaCaDI is presented. This meta-learning approach utilizes a closed-form interventional data likelihood model to bypass unstable gradient updates, enabling the rapid inference of both the shared causal structure and specific intervention targets from minimal samples.