Course Outline

Overview

This research-oriented course will explore the theoretical underpinnings of graphical modeling and causality. A graphical model is a mathematical structure that describes complex dependencies between random variables. More precisely, given a (directed or undirected) graph, we envision one random variable at each vertex. The graphical structure gives rise to conditional independence statements and, in the directed case, to functional relationships among the variables. Given data from a graphical model, we will discuss model selection: the problem of finding the graph that the data arose from, and inference: the problem of estimating the distribution assuming we know the graph. We will explore different types of algorithms used to solve these questions as well as the mathematical theory involved. Building on the theory of graphical models, we will study causal discovery. Here, we are interested in finding a directed graph that depicts the causal relationships among the observed random variables (e.g., X –¿ Y if X causes Y). We will discuss how to solve this problem in both the observational and interventional (e.g. randomized control trials) settings. We will conclude with theory and algorithms for the case of hidden variables as well as directed cycles in the graph.

List of Lectures

I. Undirected graphical models

   (a) Definition: factorization and Markov properties
   (b) Equivalence of definitions, Hammersley-Clifford Theorem
   (c) Faithfulness; Semi-graphoids, graphoids
   (d) Maximum likelihood estimation: Gaussian graphical models
   (e) Maximum likelihood estimation: discrete graphical models
   (f) The Sum-product and junction-tree algorithms
   (g) Algorithms for learning the graph structure I
   (h) Algorithms for learning the graph structure II

II. Causal models

   (a) Directed graphical models: factorization and Markov properties
   (b) Equivalence of definitions; Markov equivalent graphs
   (c) Algorithms for learning the graph structure I
   (d) Algorithms for learning the graph structure II
   (e) Structural equation models, Interventions, Counterfactuals
   (f) Linear causal models: Gaussian and non-Gaussian
   (g) Additive noise models
   (h) Learning the graph structure using interventions
   (i) Potential outcomes

(j) Hidden variable models: Simpson's paradox, instrumental variables
(k) Hidden variable models: Markov properties, m-separation, trek-separation
(l) Hidden variable models: other constraints; tensor decomposition
(m) Time series causal models
(n) Causal discovery for time series models

Grades

Final project, 50%: The final project should involve reading one or more research papers related to the course, identifying an open research problem, and making progress towards resolving the problem. Around the middle of the semester, a project proposal consisting of one page describing the problem and all references will be due.
Homework, 40%: There will be 3 homework sets throughout the semester. You are encouraged to work together with others on the problems, but you have to write down the solutions on your own.
Attendance and participation, 10%: Part of the grade will be based on your presence in class as well as your participation in discussions.