UCLA Department of Statistics

Stats 212 Graphical Models

Instructor: Qing Zhou [zhou@stat.ucla.edu], OH: Thursday 2–3pm (zoom).
Zoom link (OHs): https://ucla.zoom.us/j/97572469743.
Requisites: Stats 200A (Applied Probability).

Overview

Graphical models are widely used in statistical modeling, machine learning and causal discovery. This course gives a comprehensive introduction to graphical models and their applications, covering classical topics and recent literature. It is a useful preparation for graduate students interested in this research area.

Course Description

Introduction to graphical models with applications in statistical modeling, machine learning and causal inference. Common graphical models, such as undirected graphs, directed acyclic graphs and ancestral graphs, for modeling conditional independence and causality. Methods and theory for structure learning of graphical models from observational and experimental data.

Learning Outcomes

At the completion of this course, students are expected to be able to:

- Understand the basic concepts of common graphical models and correctly interpret graphical models under different application context.
- Use graphical models for modeling and inference on complex data.
- Apply existing and develop novel structure learning methods to estimate graphical models in real applications.
- Perform causal discovery and inference with graphical models on real-world data.

Topics (tentative)

1. Conditional independence: graphoid axioms, conditional independence test.
2. Undirected graphs: graph separations, perfectness, Gaussian graphical models, Ising models for discrete data, neighborhood regression, graphical lasso.
3. Directed acyclic graphs (DAGs): d-separations, Markov equivalence class, faithfulness, causal modeling, structural equation models (SEMs), chain graphs.
4. Constraint-based learning of DAGs: PC algorithm, consistency, local structure learning.
5. Score-based learning: hill climbing, continuous optimization methods, order-based search, large-sample guarantees.

6. Causal inference and discovery: intervention, causal effect identification and estimation, learning causal DAGs from experimental data.
7. Causal bandits and sequential intervention design.
8. Identifiable causal graphs: linear non-Gaussian acyclic models, nonlinear SEMs, additive noise models, generalized linear causal graphs.
9. Ancestral graphs, hidden confounders, selection bias, FCI algorithm, graphical models for latent factors.
10. Mixed graphical models, unifying Markov properties, neighborhood lattices.

Grading Structure

S/U or letter grading:

1. Reading papers and write a review (50%)
2. Presentation and slides (50%)

References

- Lauritzen S.L., *Graphical Models* (1996), Oxford Science Publications: Chapters 1–3.
- Spirtes, P., Glymour, C. and Scheines, R., *Causation, Prediction, and Search* (second edition, 2000), MIT Press: Chapters 2–6.
- Selected papers (to be provided during the course).