IMPORTANT:
Please refer to the USC Center for Excellence in Teaching for current best practices in syllabus and course design. This document is intended to be a customizable template that primarily includes the technical elements required for the the purpose of central review by UCOC.

EE 599: Causal Learning
Units: 3 units
Fall 2022—Day—Time: Monday/Wednesday 10:00am – 11:20am

Location: Physical address and/or course-related URLs, etc.

Instructor: Urbashi Mitra
Office: EEB 536
Office Hours: Monday, Wednesday 2pm
Contact Info: ubli@usc.edu, 213 740 4667
email responses within 48 hours

Teaching Assistant: TBD
Office: TBD
Office Hours: TBD
Contact Info:TBD

Course Description

We will examine methods for quantifying and analyzing notions of causality. This topic is increasingly important in data science and machine learning. We will define causal models and develop methods to learn these models from data. After explaining the need for causal models and discussing some of the principles underlying causal inference, we will investigate how to use causal models: how to compute intervention distributions, how to infer causal models from observational and interventional data, and how causal ideas could be exploited for classical machine learning problems. All of these topics are discussed first in terms of two variables and then in the more general multivariate case. The bivariate case turns out to be a particularly hard problem for causal learning because there are no conditional independences as used by classical methods for solving multivariate cases. We will analyze statistical asymmetries between cause and effect. Novel topics in modern causal learning will be explored through a course project and the reading of recent works.

Learning Objectives

- Define causal effects using potential outcomes
- Describe the difference between association and causation and the effect of selection bias
- Express assumptions with causal graphs
- Identify different types of causal inference problems in real-world applications
- Implement several types of causal inference methods (e.g. matching, instrumental variables, inverse probability of treatment weighting)
- Analyze causal inference problems and estimate the plausibility of solving them under certain conditions
- Identify which causal assumptions are necessary for each type of statistical method

Prerequisite(s): EE 503 or equivalent (a graduate course in probability and statistics) and **one** from the following or equivalent (EE 562 Random Processes, EE 563 Estimation Theory, EE 592 Computational Methods for Inverse Problems, EE 546-Mathematics of High-dimensional Data , EE 541 A Computational Introduction to Deep Learning, CS 573 Advanced Artificial Intelligence, Math 547 Statistical Learning Theory).

Co-Requisite(s): none

Concurrent Enrollment: none

Recommended Preparation: Basic computer skills (i.e. programming, plotting, random variable generation, familiarity with Matlab is helpful although not necessary), knowledge of convex functions and their properties, simple optimization, stationary points of functions, limits of sequences and series, etc.

Course Notes

The grading type is Letter. Materials will be posted on-line via Blackboard (assignments, solutions, reading materials, supplementary materials and any relevant lecture material such as power point slides). It is anticipated that the course will be held in person.

Technological Proficiency and Hardware/Software Required

Basic computer skills (i.e. programming, plotting, random variable generation, familiarity with Matlab is helpful although not necessary).

Required Readings and Supplementary Materials

*Required Text (**tentative**)*: Elements of Causal Inference, J Peters, D Janzing and B. Scholkopf, MIT Press, Cambridge MA, 2017 (PJS)

On Pearl's Hierarchy and the Foundations of Causal Inference E. Bareinboim, J. Correa, D. Ibeling, T. Icard In: "Probabilistic and Causal Inference: The Works of Judea Pearl", ACM Turing Series, 2020 (PCH) (available online)

Introduction to Causal Inference (https://www.bradyneal.com/causal-inference-course#course-textbook), Brady Neal, 2020 (BN)

Suggested Supplementary Materials:
Causal Inference in Statistics: A Primer, Judea Pearl, Madelyn Glymour and Nicholas. P. Jewell, Wiley, 2016.

Causal Inference for Statistics, Social, and Biomedical Sciences, Imbens, Guido W. and Donald B. Rubin, Cambridge University Press. 2010.

Causal Inference: What If, Hernán, Miguel A. and James M. RobinsChapman & Hall/CRC. 2020. (PDF available at: https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/)

Probabilistic graphical models: principles and techniques. D. Koller and N. Friedman, N MIT press, 2009.

Causation, prediction, and search. P. Spirtes, C. N. Glymour, R. Scheines, and D. Heckerman. MIT press, 2000.

Description and Assessment of Assignments
There will be approximately 5 Homework assignments throughout the course.
• Late HW will not be accepted. A late assignment results in a zero grade. Please have your homework turned in by the beginning of lecture on the date that it is due.
• Homeworks will be assigned on Thursdays and collected the following week in class (on Thursdays at the beginning of the lecture)
• Show your work in your homework solution; the correct answer alone is worth only partial credit.
• Homework collaboration is encouraged. This is discussing problems and solution strategies with your classmates, the TA, and/or the instructor and is to be distinguished from copying solutions of others which is prohibited.
• For computer-based assignments no code can be shared or copied from the internet. The only exception is code provided to the entire class by the instructor or TA.

Grading Breakdown
• 10% participation
• 20% Homework
• 30% Midterm Exam
• 40% Final Project (including proposal, presentation, report)

Assignment Submission Policy

• Late HW will not be accepted. A late assignment results in a zero grade. Please have your homework turned in by the beginning of lecture on the date that it is due.
• Homeworks will be assigned on Thursdays and collected the following week in class (on Thursdays at the beginning of the lecture)

Grading Timeline

Homework will be graded and returned within two weeks of submission.

Course Schedule: A Weekly Breakdown

	Topics/Daily Activities	Readings/Preparation	Deliverables
Week 1	Probability review Correlation does not imply causation Independent Mechanisms	PJS 1, PJS 2, BN 1, PCH 1.1	HW 1 assigned
Week 2	Cause-Effect models	PJS 3, BN2, PCH 1.4	
Week 3	Learning Cause-Effect Models Connections to Machine Learning 1	PJS 4, PJS 5	HW 1 due; HW 2 assigned
Week 4	Graphical Model principles Structural Causal Models	BN 3, PJS 6	project proposal due
Week 5	Causal Models	BN 4	HW 2 due; HW 3 assigned
Week 6	Multivariate Causal Models	PJS 6, PCH 1.2	
Week 7	Randomized Experiments	BN 5, BN 6	HW 3 due
Week 8	Learning Multivariate Causal Models Connections to Machine Learning II	PJS 7, PJS 8	Midterm
Week 9	Estimation (Grouped) Conditional Outcome Modeling	BN 7, PCH 1.4 handouts	HW4 assigned
Week 10	Unobserved Confounding and Sensitivity Analysis	BN 8 handouts	interim project report due
Week 11	Hidden and Instrumental Variables	PJS 9, BN 9	HW4 due; HW 5 assigned
Week 12	Time Series Analysis, Differences and Differences	PJS 10, BN 10	
Week 13	Causal Discovery from Observational Data Causal Discovery from Interventional Data	BN 11, BN12, PCH 1.3 PCH 1.5	HW 5 due
Week 14	Counterfactuals Redux Project presentations	handouts	

Week 15	Project presentations		
FINAL	**No final exam**		final project report due

Statement on Academic Conduct and Support Systems

[Paste most recent version of the statement here; see the CCO Resources page.]