CAUSAL INFERENCE: PREDICTION, EXPLANATION, AND INTERVENTION

What is causality and why is it useful? Causes are what allow us to predict what will happen in the future (that a stock price will rise based on a news report), explain why something happened in the past (what actually led to a patient's seizure), and intervene to produce particular outcomes (crafting political speeches to influence voter opinion). Whether you want to buy stocks, develop effective treatments, or manipulate elections, you need to know that you are not acting on a mere surrogate but rather the true culprit.

Overview

This course covers the practical tools needed for evaluating causal claims and making causal inferences. We will explore two primary questions – 1) what is causality? 2) how can we find it? After covering the conceptual and theoretical underpinnings of causal inference, we discuss how causal inference is handled by different fields and how to responsibly test it in real-world cases.

Prerequisites

None. The course is intended for advanced undergraduate and graduate students from computer science and other disciplines.

Text

There is no required textbook. Readings are journal articles or book chapters that will be provided. Optional supplementary readings are also listed.

Evaluation

Discussion of the readings is an important part of the course and will count towards the final grade. There will be a final project (and presentation), which may be theoretical or experimental in nature (for example, applying causal inference methods to data, writing a critique of a methodology or study, proposing a new inference method). There will also be one presentation for the journal club assignment, which will count toward participation grades.

Grades will be: 5% homework, 15% participation, 30% midterm exam, 50% final project.

Contact

Samantha Kleinberg
208 North samantha.kleinberg@stevens.edu
Office hours are Monday 2-2:45pm and by appointment.

Lecture

Monday 3-5:30pm

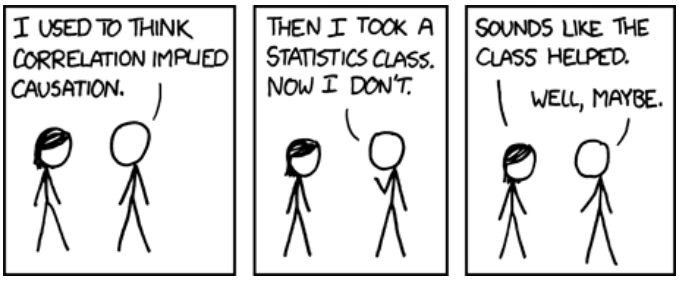

http://xkcd.com/552/

Schedule

Below is the weekly schedule of topics, readings for each session, and deadlines. All journal articles will be available electronically through the library or in canvas. Readings in CPT are from *Causality, Probability, and Time* (Cambridge, 2012). Other readings may change, but in that case an announcement will be made in class.

What is causality?

8/27 Introduction to causality and causal inference: Why is it challenging? Why do we seek causes anyway?

9/10 Regularities, counterfactuals and token causality (what normally happens, what would happen and what did happen)
- Discussion paper: A. Olteanu, O. Varol, and E. Kiciman. Distilling the outcomes of personal experiences: A propensity-scored analysis of social media. In *CSCW*, 2017 [pdf]

Optional:
- CPT, Section 2.1 – 2.2.2
- D. Lewis. Causation. *The Journal of Philosophy*, 70(17):556–567, 1973 [pdf]

9/17 Probabilistic causality
- Discussion paper: T. Kushnir and A. Gopnik. Young children infer causal strength from probabilities and interventions. *Psychological Science*, 16(9):678–683, 2005 [pdf]

Optional:
- CPT, 2.3
- P. Suppes. *A Probabilistic Theory of Causality*. North-Holland, 1970 [pdf]

How can we find causes?

9/24 Introduction to graphical models, Bayesian networks
- Discussion paper: M. J. Kusner, J. Loftus, C. Russell, and R. Silva. Counterfactual fairness. In I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems 30*, pages 4066–4076. Curran Associates, Inc., 2017 [pdf]

Optional:
- CPT, 2.4.1
- R. Scheines. An introduction to causal inference. *Causality in Crisis*, pages 185–99, 1997 [pdf]

10/1 Causality in time series: dynamic Bayesian networks, logic-based methods
- Discussion paper: N. Heintzman and S. Kleinberg. Using Uncertain Data from Body-Worn Sensors to Gain Insight into Type 1 Diabetes. *Journal of Biomedical Informatics*, 63:259–268, 2016 [pdf]

Optional:
- CPT, 2.4.2
- CPT, chapters 4 – 5

10/9 Causality in time series: Granger causality and other methods
- Discussion paper: B. Elbel, J. Gyamfi, and R. Kersh. Child and adolescent fast-food choice and the influence of calorie labeling: a natural experiment. *International Journal of Obesity*, 35(4):493–500, 2011 [pdf]

Optional:
- C. W. J. Granger. Testing for Causality: A Personal Viewpoint. *Journal of Economic Dynamics and Control*, 2:329–352, 1980
- W. N. Thurman and M. E. Fisher. Chickens, Eggs, and Causality, or Which Came First. *American Journal of Agricultural Economics*, 70(2):237–238, 1988 [ezproxy]

10/15 Midterm

When can causes be inferred?

10/22 Psychology of causation
- Discussion paper L. Rozenblit and F. Keil. The misunderstood limits of folk science: An illusion of explanatory depth. *Cognitive Science*, 26(5):521–562, 2002[pdf]
- 1-page proposal for final project due.

10/29 Mechanisms, interventions, and randomized trials
- Discussion paper N. Daskalova, K. Desingh, A. Papoutsaki, D. Schulze, H. Sha, and J. Huang. Lessons learned from two cohorts of personal informatics self-experiments. *Proc. ACM Interact. Mob. Wearable Ubiquitous Technol.*, 1(3):46:1–46:22, Sept. 2017[pdf]

11/5 Cancelled

11/12 Journal club
- Readings TBA (will be recent experimental work on causal inference)
- Brief presentation on and discussion of assigned papers.

11/19 Assumptions and experiments
(selection bias, confounding, cross-sectional data)
- Discussion paper A. F. Ward, K. Duke, A. Gneezy, and M. W. Bos. Brain drain: the mere presence of one's own smartphone reduces available cognitive capacity. *Journal of the Association for Consumer Research*, 2(2):140–154, 2017[pdf]

11/26 Presentations of final projects
- Written project reports due

12/3 Presentations of final projects