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
306 Lieb
samantha.kleinberg@stevens.edu
Office hours are Monday 2-2:45pm and by appointment.

Lecture

Monday 3-5:30pm

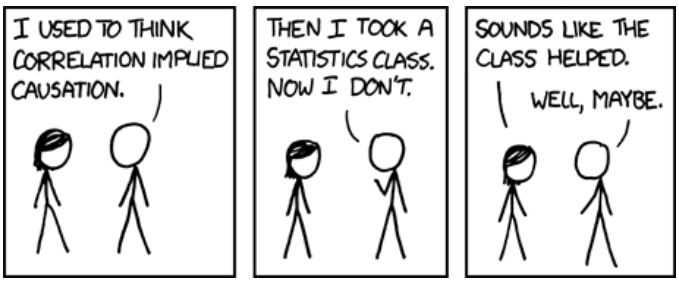

http://xkcd.com/552/

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 moodle. 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?

(1) Introduction to causality and causal inference: Why is it challenging? Why do we seek causes anyway?

(2) Regularities, counterfactuals and token causality (what normally happens, what would happen and what did happen)
- CPT, Section 2.1 – 2.2.2

Optional:
- D. Lewis. Causation. *The Journal of Philosophy*, 70(17):556–567, 1973 [ezproxy]
- J. L. Mackie. *The Cement of the Universe*. Clarendon Press, 1974

(3) Probabilistic causality
- CPT, 2.3

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

How can we find causes?

(4) Introduction to graphical models, probability review
- D. Heckerman. A tutorial on learning with Bayesian networks. *Innovations in Bayesian Networks*, pages 33–82, 2008 [pdf]

(5) Conditions for Inference, Bayesian networks
- CPT, 2.4.1
- R. Scheines. An introduction to causal inference. *Causality in Crisis*, pages 185–99, 1997 [pdf]

Optional:
- J. Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2000

(6) Causality in time series: dynamic Bayesian networks, logic-based methods
- CPT, 2.4.2

Optional:
- CPT, chapters 4 – 5

(7) Causality in time series: Granger causality and other methods
- C. W. J. Granger. Testing for Causality: A Personal Viewpoint. *Journal of Economic Dynamics and Control*, 2:329–352, 1980 [moodle]
- Discussion of case studies

Optional:

  - 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]

(8) Midterm

When can causes be inferred?

(9) Mechanisms, interventions, and randomized trials
  - P. W. Holland and D. B. Rubin. Causal inference in retrospective studies. *Evaluation Review*, 12(3):203, 1988 [moodle]
  - F. J. Mackenzie and J. Grossman. The randomized controlled trial: gold standard, or merely standard? *Perspectives in Biology and Medicine*, 48(4):516–534, 2005 [ezproxy]
  - 1-page proposal for final project due.

(10) Applications: epidemiology, economics, psychology
  - A. B. Hill. The Environment and Disease: Association or Causation? *Proceedings of the Royal Society of Medicine*, 58(5):295–300, 1965 [pdf]
  - C. T. Carlstrom and E. N. Gamber. Why we don't know whether money causes output. *Economic Review*, (Q III):27–39, 1989 [ezproxy]
  - M. Steyvers, J. B. Tenenbaum, E. J. Wagenmakers, and B. Blum. Inferring causal networks from observations and interventions. *Cognitive Science*, 27(3):453–489, 2003 [ezproxy]

(11) Assumptions and experiments
(selection bias, confounding, cross-sectional data)
  - S. S. Young and A. Karr. Deming, data and observational studies. *Significance*, 8(3):116–120, 2011 [ezproxy]
  - M. Joffe. The gap between evidence discovery and actual causal relationships. *Preventive Medicine*, 53(4 - 5):246 – 249, 2011 [ezproxy]

(12) Evidence and evaluation (journal club)
  - Readings TBA (will be recent experimental work on causal inference)
  - Brief presentation on and discussion of assigned papers.

(13) Presentations of final projects
  - Written project reports due

(14) Presentations of final projects