



Computer Science Department (EN.601.177)
Should Susan Smoke: An Introduction To Causal Inference
Intersession 2020

Instructors

Rohit Bhattacharya
Computer Science
Johns Hopkins University
Office: Malone Hall 340
rbhattacharya@jhu.edu

Razieh Nabi
Computer Science
Johns Hopkins University
Office: Malone Hall 340
rnabi@jhu.edu

Meetings

MTuThu: 2pm-4pm, Hodson 211

Prerequisites

Exposure to the phrase “correlation is not causation”.

Course Description

The fear of granting credence to spurious correlations has caused modern journalism, scientific reporting, and even everyday conversation to steer clear of causal language – a new miracle diet is found to be *associated* with lower risk for cardiac disease. However, when it comes to prescribing solutions (approving a new medication, passing new laws to protect against climate change), it is clear that no one wants to act on mere associations. In this course, we describe how and when we can think of these associations as being causal relations. We delve into the use of causal tools to reason about issues across disciplines, such as the efficacy of new treatments in medicine, policy evaluation in economics, education, and governance, and disparity in wages, hiring, and prison time across gender and race. Throughout the course, we will use case studies to illustrate problems where even the best scientists in their respective field were unsure if an association could be considered causal, and how these issues were eventually resolved by advances in causal inference.

Learning Objectives

By the end of this course, students should be able to:

- Be able to distinguish observations from interventions.
- Draw a causal diagram and reason about its implications.
- Be able to evaluate the merits of causal and associational claims in news or research articles.
- Understand the use and limitations of randomized controlled trials, and their role in causal analysis.
- Understand the use and limitations of purely observational data, and their role in causal analysis.

Schedule

- **Week 1 Lectures: Correlation vs. Causation – What is the Difference?**
 - What do causal questions look like?
 - Mathematical definition of correlation.
 - Pros and cons of past definitions of causation from philosophers, econometricians, doctors, and computer scientists.

- The ladder of causation – seeing vs. doing.
 - Elements of the modern definition of causality – probability theory, interventions, counterfactual reasoning, and a new notation (do-operator).
 - Randomized controlled trials – where correlation meets causation.
 - Simpson’s paradox and confounding bias.
 - Causal inference as a missing data problem – individual level causal effects.
 - Defining different causal targets of inference (for e.g., the average causal effect).
 - Linking counterfactuals to factials through assumptions – consistency, and (conditional) ignorability.
- **Week 1 Case Study 1: Chocolate Consumption & Nobel Laureates**
 - How to refute frivolous claims using precise causal language and reasoning.
 - Reviewing the methodological flaws in Messerli (2012) in the New England Journal of Medicine.
 - Reichenbach’s common-cause principal as an explanation of the observed correlation.
- **Week 1 Case Study 2: The Smoking & Lung Cancer Controversy**
 - Defining the causal target of inference.
 - Review of the first case-control study by Hill and Doll (1950).
 - Criticism from Fisher as letters to the British Medical Journal.
 - Drawbacks of case-control studies and the issue of selection bias in causal inference.
 - The complication of potential unmeasured confounders.
- **Week 2 Lectures: Causal Models of Our World - Assumptions and Representations**
 - The dryness of conditional independence statements in probability theory.
 - Representation of complex joint probability distributions via graphical models.
 - Introducing directed acyclic graphs (DAGs) and genealogical relations between nodes in the graph.
 - Interpretation of the graphical model using a family tree analogy.
 - Information flow in DAGs: chains, forks, and collider paths.
 - d-separation criterion for blocking information flow.
 - Examples of d-separation + demo.
 - Linking DAGs to causal language.
 - Structural equation models (SEMs) of a DAG.
 - Interventions in SEMs.
 - Structural causal models of a DAG.
 - Independence statements between factials and counterfactuals.
 - Single world intervention graphs (SWIGs).
 - Generalizing conditional ignorability by finding adjustment sets in SWIGs.
 - Beyond adjustment: the front-door criterion.
 - Inverse probability weighting to calculate the average causal effect in a representative pseudo-population.
- **Week 2 Case Study: Estrogens & Endometrial Cancer**
 - Berkson’s bias – what happens if you accidentally open up a collider path?
 - Mack et al study on affluent post-menopausal women implicating estrogens as a cause of endometrial cancer.
 - Horwitz and Feinstein follow up study claiming that the correlation is due to collider bias.
 - Flaws in the Horwitz and Feinstein study – trading one collider bias for another.
 - How causal DAGs help us to formulate the issues surrounding the controversy and suggest testable hypotheses.
- **Week 3 Lecture 1: Formalizing (Un)Fairness Using The Language of Causal Inference**
 - Name-swapping experiments in hiring discrimination cases.

- Understanding causal mechanisms using nested counterfactuals.
 - Mediation analysis: decomposing the average causal effect into direct effect and indirect effect mediated through a third random variable.
 - Generalizing mediation effects to arbitrary causal path-specific effects.
 - Characterizing unfairness with respect to a sensitive random variable (sex, gender, ...) and an outcome as the presence of an effect along impermissible causal pathways.
 - Shifting the inference problem on the observed data distribution to an inference problem in a distribution close to the observed distribution but where the unfair effect is absent.
 - Well-defined interventions in causal reasoning.
- **Week 3 Lecture 2: Causal Discovery**
 - What do we do if we have insufficient knowledge to draw the causal DAG?
 - The faithfulness assumption for causal discovery.
 - Testable implications of a causal DAG.
 - The maximality condition and absence of edges.
 - Conditional vs. marginal independence and implications on edge orientation.
 - Observational equivalence and representations of equivalent structures.
 - Applications of causal discovery: protein-protein interaction networks, discovering causes of climate change, and constructing causal connections in the brain.
 - Demo of causal discovery algorithms.

Textbooks

There is no single comprehensive textbook on causal inference, and in this course, we only provide a high-level overview of the field. For students interested in diving deeper, the following is a good collection of books to start with.

- *Causal Inference in Statistics: A Primer* by Judea Pearl, Madelyn Glymour and Nicholas P. Jewel.
- *Causality: Models, Reasoning, and Inference* by Judea Pearl. ([3]).
- *Causation, Prediction, and Search* by Peter Spirtes, Clark Glymour, and Richard Scheines. ([4]).
- *Causal Inference* (draft form) by Miguel A. Hernan, James M. Robins. Found at: <https://www.hsph.harvard.edu/miguel-hernan/causal-inference-book/>.
- *Statistical Methods for Dynamic Treatment Regimes (reinforcement learning, causal inference, and personalized medicine)* by Bibhas Chakraborty and E.M. Moodie ([1]).
- *Probabilistic Reasoning In Intelligent Systems* by Judea Pearl. ([2]).
- *All Of Statistics (a Concise Course in Statistical Inference)* by Larry Wasserman ([5]).

Ethics

The strength of the university depends on academic and personal integrity. In this course, you must be honest and truthful, abiding by the *Computer Science Academic Integrity Policy*.

You can find more information about university misconduct policies on the web at these urls:

- Undergraduates: e-catalog.jhu.edu/undergrad-students/student-life-policies/
- Graduate students: e-catalog.jhu.edu/grad-students/graduate-specific-policies/

Students with Disabilities

Any student with a disability who may need accommodations in this class must obtain an accommodation letter from Student Disability Services, 385 Garland, (410) 516–4720, studentdisabilityservices@jhu.edu.

REFERENCES

- [1] Bibhas Chakraborty and Erica E. M. Moodie. *Statistical Methods for Dynamic Treatment Regimes (Reinforcement Learning, Causal Inference, and Personalized Medicine)*. Springer, New York, 2013.
- [2] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems*. Morgan and Kaufmann, San Mateo, 1988.

- [3] Judea Pearl. *Causality: Models, Reasoning, and Inference*. Cambridge University Press, 2 edition, 2009.
- [4] Peter Spirtes, Clark Glymour, and Richard Scheines. *Causation, Prediction, and Search*. Springer Verlag, New York, 2 edition, 2001.
- [5] Larry Wasserman. *All Of Statistics (a Concise Course in Statistical Inference)*. Springer, New York, 2005.