

Causal Inference I (COMS W4775)

(Intro to Causal Inference)

Time: MW 4:10pm - 5:25pm

Location: TBD - Online

Instructor: Elias Bareinboim [eb@cs.columbia.edu]

Webpage: <https://causalai.net/ci20f/>

Summary

You run your favorite machine learning algorithm and obtain a strong correlation between two variables, say X and Y . Some weeks later, you implement a new policy that increases the value of X . Surprisingly, nothing happens with the value of Y . Was the algorithm incorrectly implemented? Do you need to collect more data? Would another algorithm perform differently? Those are challenging questions that appear in everyday data analysis. This course will try to address this and other common questions that relate different ways for explaining the data and data collection strategies, which usually comes under the general name of causal inference.

The emergence of causal inference in AI, Machine Learning, and the data sciences does not come as a surprise since the interest in knowing that X is (probabilistically) correlated with Y is not rarely devoid of meaningful interpretation or practical implications. There are plenty of real world examples demonstrating that correlation does not imply causation, hence not suitable to substantiate causal claims and principled decision making. For instance, no one expects that forcing the rooster to crow in the middle of the night will make the sun rise. Still, these two events are strongly correlated.

In this course, we will introduce concepts, principles, and algorithms necessary to solve modern, large-scale problems in scientific inferences, business, and engineering. Emphasis will be given to the tradeoff between assumptions (delineated by current scientific knowledge) and conclusions for standard types of queries, including associational, causal, and counterfactual. In other words, we will consider the problem of providing different types of “explanation” for the vast amount of data collected in different fields of human inquiry, including engineering, medicine, and the empirical sciences. This problem lies at the heart of current discussions in artificial intelligence, machine learning, and statistics.

Prerequisites

In order to be successful in this course, you should have a basic knowledge of:

- Discrete Math (proof techniques, search algorithms, and graph theory)
- Calculus (find min/max of functions)
- Statistics (basic probability, modeling, experimental design)
- Machine Learning (graphical models)-Some programming experience (with special understanding of complexity analysis)

References

We will use material selected from different sources, including chapters of the following books:

[C] Causality: Models, Reasoning, and Inference.

J. Pearl.

Cambridge Press, 2000.

[W] The Book of Why

J. Pearl, D. Mackenzie

Basic Books, 2018.

[PCH] 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.

[P] Causal Inference in Statistics: A primer

J. Pearl, M. Glymour, N. Jewel

Cambridge Press, 2016.

[DF] Causal Inference and the Data-fusion Problem

E. Bareinboim, J. Pearl

Proc. National Academy of Sciences, 2016.

Material

Subject	Material	References
Introduction	Logistics, Motivation, Machine Learning, Pearl's Causal Hierarchy.	[PCH] Sec. 1.1; [W] Ch. 1; [P] Ch. 1
Structural Causal Models	Structural Causal Models, Graphs, and d-Separation.	[PCH] Sec. 1.2; [P] Ch. 2; [C] Sec. 1.4; PGM-Koller, Ch. 3.
Identification of Causal Effects - Basics	Definition of Causal Effects, the Truncated Product, and the Identification Problem.	[C] Sec. 1.3, 3.1, 3.2; [PCH] Sec. 1.4.
The Problem of Confounding and the Backdoor Criterion	Identifiable and non-identifiable effects. Confounding Bias. The Backdoor Criterion. IPW and Propensity Score.	[PCH] Sec. 1.4; [C] Sec. 3.3; [P] Ch. 3.
The Algorithmic Backdoor Criterion	The Conditional Backdoor Criterion. Adjustment-Backdoor Criterion. Poly-time delay Backdoor.	J. Textor & M. Liskiewicz, UAI, 2011; Yung, Tian, Bareinboim, 2020.
The Causal Calculus	The Generalized Truncated Product. The Front-door Case. The Causal Calculus. Identifiable and Non-Identifiable Effects.	[PCH] Sec. 1.4.3; [C] Sec. 3.4-3.5; Pearl, Biometrika, 1995.

Subject	Material	References
Algorithmic Approach for Identification	C-factors, Q-factorization, Identify Algorithm.	Tian, Ch. 5, 2002; Huang & Valtorta (Annals of Math & AI, 2008); Lee et al (UAI, 2019).
Soft Interventions and Sigma Calculus	Non-atomic interventions. Sigma-node. Sigma-calculus. Comparison with do-calculus. Examples.	[C] Sec 4.2; Correa and Bareinboim (AAAI, 2020).
Partial Identification and Generalized Off-Policy Learning	Causal Bounding. Generalized Off-Policy Learning (Reinforcement Learning).	[C] Ch. 8; Zhang and Bareinboim (IJCAI, 2017, NeurIPS, 2019).
Causal Data Science	The Data-Fusion Framework. Generalized Identification (or, non-parametric IVs). Transportability. Selection Bias.	Bareinboim and Pearl (PNAS, 2016); Correa et al. (ICML, 2019, AAAI, 2019); Lee et al. (UAI 2019, AAAI, 2020).
Linear Structural Causal Models	Wright's Rules. Regression vs Structural Coefficients. Single- and Back-door Criteria. (Conditional) IVs. AV-Cutsets.	[C] Ch. 5; Pearl (JCI, 2013, link); Kumor et al (NeurIPS, 2019).
Counterfactuals - Basics	Causal Hierarchy Theorem. Potential Response. 3-step Procedure. Twin Networks. Axioms. Identification. ETT. Mediation.	[C] Ch. 7; [PCH] 1.3, 1.5.

Grading

The course will have both a theoretical and practical components.

- Midterm exam: 25%
- Homeworks: 25%
- Final Exam: 50%
- Attendance and participation: bonus

Please review the Columbia honor code. While working on assignments in small teams is okay, your homework solutions must be your own.

Course Policies

You are expected to attend lectures and participate in class in the designated modality be it virtual or in person. If present, taking notes on laptops and having small snacks is permitted, please make sure you are quiet and respectful of those around you, including those behind you who might be distracted by your snacks/devices. You are expected to come prepared to class, and to participate in class discussion, especially when something is not clear. If you are too shy to ask in class, please post on piazza or attend the office hours.