

CIS 700: Causal Inference

Instructor: Junzhe Zhang (izhan403@syr.edu)

Time: TuTh 2-3:20pm

Location: CST 3-212

Website: <https://junzhez.com/ci26f/>

Summary

Imagine a scenario where you're analyzing electronic healthcare records from US hospitals. After applying your favorite Reinforcement Learning algorithm, you identify a strong correlation between a specific treatment (X) and a positive impact on the patient's health condition (Y). Eager to act on this discovery, you implement a new treatment policy recommending treatment X to patients. However, instead of getting better, most patients' health conditions deteriorate a few weeks after receiving treatment X. Were there issues with how the algorithm was implemented? Is the sample size adequate? Should an alternative algorithm be considered? These are the real challenges encountered when reasoning about optimal decision-making for real-world applications. This course provides guidance to understand and address these challenges, especially in the context of data bias, by leveraging causal knowledge of the data-generating process.

The rise of Causal Inference in Artificial Intelligence and Data Science is not surprising. Many people are interested in understanding if X is statistically related to Y, but this correlation doesn't always have a meaningful interpretation or practical implication. There are many real-world examples showing that correlation does not imply causation. Therefore, it is not suitable for supporting causal claims and principled decision-making. For example, just because the rooster crows before sunrise doesn't mean the crowing causes the sun to rise.

Throughout this course, we will cover concepts, principles, and algorithms essential for solving modern, large-scale decision-making problems in consequential domains such as scientific research, business, and engineering. In the first half of the course, we will focus on the tradeoff between assumptions about the data collection process and deriving conclusions for different types of queries, including associational, causal, and counterfactual, from various data sources. In the latter half of the course, we will use this knowledge to develop novel decision-making algorithms that improve state-of-the-art methods such as online learning, off-policy learning, and imitation learning. This

challenge is central to current conversations in artificial intelligence, machine learning, reinforcement learning, and data science.

Prerequisites

To be successful in this course, students should have a basic knowledge of:

- Discrete math
- Calculus
- Statistics and probability theory
- Analysis of algorithms
- Some programming experience

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.

[PGM] Probabilistic graphical models: principles and techniques. Koller, Daphne, and Nir Friedman. MIT press, 2009.

[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.

[DF] Causal Inference and the Data-fusion Problem E. Bareinboim, J. Pearl Proc. National Academy of Sciences, 2016

Outline

Subject	Materials	References
Introduction	Logistics, Motivation, Machine Learning, Pearl's Causal Hierarchy	[W] Ch. 1; [PCH] Ch. 1;

Structural Causal Models	Structural Causal Models, Causal Diagram, d-separation	[PGM] Ch. 3 [C] Ch. 1.4 [PCH] Sec. 1.2
Identification of Causal Effects	Definition of Causal Effects and the Identification Problem.	[C] Sec. 1.3, 3.1, 3.2; [PCH] Sec. 1.4.
The Problem of Confounding and the Backdoor Criterion	Challenges of Confounding Bias. The Backdoor Criterion. IPW and Propensity Score	[PCH] Sec. 1.4; [C] Sec. 3.3;
The Causal Calculus	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.
Partial Identification of Causal Effects	Natural Bounds, Instrumental Variables, Canonical Models	[C] Ch. 8
Causal Discovery	Markov Equivalence, Inferred Causation, Partial Ancestral Graph	[C] Ch. 2
Counterfactual Inference	Potential Response. 3-step Procedure. Twin Networks. Axioms. Identification. ETT.	[C] Ch. 7; [PCH] 1.3, 1.5
Causal Data Science	The Data-Fusion Framework. Generalized Identification (or, non-parametric IVs). Transportability. Selection Bias	[DF]; Correa et al. (ICML, 2019, AAAI, 2019); Lee et al. (UAI 2019, AAAI, 2020)
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)
Generalized Off-Policy Evaluation	Off-Policy Evaluation, Sigma Calculus	[C] Sec 4.2; Correa and Bareinboim (AAAI, 2020).
Generalized Partial Identification	Generalized Canonical Models, Bayesian Approach for Bounding	Zhang, Tian & Bareinboim (ICML, 2022)
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);
Causal Explanation	Mediation Formula, Causal Explanation Formula, Path Specific Effects	Pearl (AAAI, 2001); Zhang & Bareinboim (AAAI, 2018)

Grading

Students will be evaluated based on a final project and class participation. Your grade is decomposed as follows:

- Project Proposal: 15%
- Mid-semester Progress Report: 15%
- Final Report: 60%

- Attendance and participation: 10%

All assignments must be submitted as PDF documents compiled using [TeX, LaTeX, or similar systems](#) with bibliographic references (e.g., using [BibTeX](#)) included as necessary.

Academic Integrity Policy

As a pre-eminent and inclusive student-focused research institution, Syracuse University considers academic integrity at the forefront of learning, serving as a core value and guiding pillar of education. Syracuse University's Academic Integrity Policy provides students with the necessary guidelines to complete academic work with integrity throughout their studies. Students are required to uphold both course-specific and university-wide academic integrity expectations such as crediting your sources, doing your own work, communicating honestly, and supporting academic integrity. The full Syracuse University Academic Integrity Policy can be found by visiting class.syr.edu, selecting, "Academic Integrity," and "Expectations and Policy."

Upholding Academic Integrity includes the protection of faculty's intellectual property. Students should not upload, distribute, or share instructors' course materials, including presentations, assignments, exams, or other evaluative materials without permission. Using websites that charge fees or require uploading of course material (e.g., Chegg, Course Hero) to obtain exam solutions or assignments completed by others, which are then presented as your own violates academic integrity expectations in this course and may be classified as a Level 3 violation. All academic integrity expectations that apply to in-person assignments, quizzes, and exams also apply online.

Students found in violation of the policy are subject to grade sanctions determined by the course instructor and non-grade sanctions determined by the School or College where the course is offered. Students may not drop or withdraw from courses in which they face a suspected violation. Any established violation in this course may result in course failure regardless of violation level.

Policy on Artificial Intelligence Use

Based on the specific learning outcomes and assignments in this course, artificial intelligence is permitted on the following: **Grammarly**. See each assignment, quiz, or exam instructions for more information about what artificial intelligence tools are permitted and to what extent, as well as citation requirements. If no instructions are provided for a specific assignment, then no use of any artificial intelligence tool is

permitted. Any AI use beyond that which is detailed in course assignments is explicitly prohibited except when documented permission is granted.