Applied Causality

David M. Blei
Columbia University
Spring 2022

Updated January 17, 2022

Description

We will study applied causality, especially as it relates to Bayesian modeling and probabilistic machine learning. Topics include causal graphical models, potential outcomes, and counterfactual inference. We will study applications, such as in economics, genetics, and healthcare. Each student will embark on a semester-long project around applied causal inference.

Topics

Some of the topics we will study include the following:

- The basics of applied causality
- Probabilistic machine learning, Bayesian statistics, and causality
- Causality and representation learning
- Causality and invariance
- Causality and algorithmic fairness
- Multiple causality
- Applied sensitivity analysis
- Instrumental variables (and other natural experiments)
- Synthetic controls, difference-in-differences, factorization models, and panel data

Prerequisites

The course is open to doctoral students; no auditors and no pass/fail.

Additional requirements:

- You have taken STCS6701 *Foundations of Graphical Models* (or have studied the material).
- You have a causal inference research problem in hand.

Organization

Our meetings are based around discussion. The discussion will focus on the readings and on your projects, including "best practices" for research.

In the first part of each session, we will discuss the reading. I will lead the discussion, but I hope for active participation.

In the second part of each session, one or two students will discuss their projects (so far). For example this discussion can include:

- a tutorial about a tool you like (e.g., Stan, Jupyter, RStudio)
- background about your project and a technical issue
- an important paper or idea that we are not otherwise covering.

Broadly, I want us to share with each other the tools and techniques that we use to develop ideas, implement software, and produce reports. At a higher level, this is a course about doing research.

Coursework and grade

The coursework is the following:

- weekly readings
- weekly responses to the readings
- a weekly commit to your final project repository (see below)
- a final project report

You are graded on completing the responses (15%), working consistently on your project (15%), class participation (15%), and the final project (55%).

Final project report

The project has two components: a report and a git repository. The report is the final output of your project. It should be at most 15 pages long and prepared with the provided latex template. The report is graded on both content and writing quality.[1]

Final project repository

The repository is a record of your research for the semester. It contains all the files that are used to produce the report (except files that are too large). Commit often, at least every week. You are graded on the quality of the project and the path that you took to get there.

Your repository should be organized like this:

- `readme.md`
- `journal.md`
- `doc/`
- `src/`
- `etc/`

The file `readme.md` contains an abstract of the project. At first you will write an "aspirational abstract," one that describes the project you want to complete. As the semester continues, you will refine this abstract to accurately describe your project.

The file `journal.md` is a diary of your progress. It contains dated entries with a description of what you are doing, what you found, what you are thinking about, and so on. It is mainly a resource for you, but I will glance at it too (at the end of the semester). Update it at least once per week.

[1] Three good books about writing are Strunk and White (1979); Williams (1981); Francis-Noel and Turner (2011).

The `doc/` directory contains latex documents that you are writing, a subdirectory for each one.

The `src/` directory contains the code you are writing.

The `etc/` directory contains anything else—materials, notes, photos of whiteboards, and so on—that you want to keep track of.

References

Francis-Noel, T. and Turner, M. (2011). *Clear and Simple as the Truth: Writing Classic Prose, Second Edition*. Princeton University Press.

Strunk, W. and White, E. (1979). *Elements of Style*. Longman Press.

Williams, J. (1981). *Style: Towards Clarity and Grace*. University of Chicago Press.