CS886: Causal Inference in Machine Learning

Fall 2020

W 1:00 – 4:00 (AOE)

Web: https://cs.uwaterloo.ca/~y328yu/mycourses/886

Piazza: https://piazza.com/uwaterloo.ca/fall2020/cs886/home

Syllabus

Instructor

Yao-Liang Yu
DC 3617
Assistant Professor, School of Computer Science, University of Waterloo
CIFAR AI Chair, Vector Institute

Web: https://cs.uwaterloo.ca/~y328yu
Email: yaoliang.yu@uwaterloo.ca
Office hours: TBD or by email appointment

Course Description

Causal inference has gone through major progress in the past two decades and we see explosive growth and interest in the machine learning field recently. Many important questions in ML can be formulated and addressed through a causal lens. Examples include algorithmic fairness, interpretability of black-box models, transfer knowledge among different domains, robustness of deep models, and more. In this course we will learn the basics of causal inference (from a graphical model perspective) and the recent developments in the machine learning field (through students' presentations).

Course Objectives

The main objective of this course is to raise awareness of the current tools that can aid causal inference, and the many subtleties in dealing with observational data (what can and cannot be inferred).

The course will start from a basic introduction to the core concepts in causal inference (cause, effect, intervention and counterfactual), supplemented with many intriguing examples. Towards the end of the course, the students will be able to

- grasp the basic theory behind causal inference;
- learn the relevance and connections to applications;
- know the state of causal inference in machine learning;
- gain some research experience on causality.

Course Overview (tentative)

- Logistics
- Introduction

- Intervention
- Counterfactual
- Learning Cause-Effect
- Connections to ML
- Hidden Variables
- Student Presentation I
- Student Presentation II
- Student Presentation III
- Student Presentation IV
- Student Presentation V
- Summary

Prerequisites

Knowledge of basic probability and statistics (common distributions, independence, conditional independence, mean, etc.) and machine learning. Exposure to Bayesian networks and graphical models is a plus but not required.

Textbooks

There is no required textbook. We will pose lecture notes or slides before class. You are encouraged to check out the following excellent books. Our plan is to go through [PGJ16] and [PJS17], with some supplements from [J09].

- *Elements of Causal Inference: Foundations and Learning Algorithms*. Jonas Peters, Dominik Janzing, and Bernhard Schölkopf. MIT, 2017. *free access*
- *Causal Inference in Statistics: A Primer*. Judea Pearl, Madelyn Glymour, and Nicholas P. Jewell. Wiley, 2016.
- *Causal Inference for Statistics, Social, and Biomedical Sciences*. Guido W. Imbens and Donald B. Rubin. Cambridge, 2015.
- *Counterfactuals and Causal Inference (2nd)*. Stephen L. Morgan and Christopher Winship. Cambridge, 2015.
- *Causality: Models, Reasoning, and Inference (2nd)*. Judea Pearl. Cambridge, 2009.
- *Causation, Prediction, and Search (2nd)*. Peter Spirtes, Clark Glymour and Richard Scheines. MIT, 2000.

Grading

The assignments of this course consist of 3 parts:

- write weakly critic of classic papers (20% + 10%);
  - use the provided style and tex files
  - you may follow the instructions
- present a recent ML paper on causality (from the list; constantly updated) (20% + 10%)
- perform a course project about causal inference and write 8-page report (25% + 15%)

In each item above, you are also expected to review each other's work. The 1st percentage is decided by the review scores you (as author) obtain and the 2nd percentage is decided by author's evaluation of you (as reviewer). We use eduflow to handle all assignments.

As usual, it is OK to seek for help, but you must write your solutions independently and individually, and you should always acknowledge any help you get (book, friend, internet, etc.).

Late Policy

We do **NOT** accept any late homework submissions, unless you have a legitimate reason with formal proof (e.g. hospitalization, family urgency, etc.). Traveling, busy with other stuff, or simply forgetting to submit, are not considered legitimate.

Presentation (20% presenting + 10% reviewing)

Each student (**individually**) is expected to present a recent ML paper on causality (from the list; constantly updated). Create a video by recording your presentation at home. It is sufficient to create a set of slides and then to record a narration of your slides. You may use keynote, powerpoint, beamer or any other software to prepare your slides. **You must make your own slides.** Although you are free (provided explicit acknowledgement is given) to use any material you gather (including the authors' slides, if any), you are expected to deliver the talk in your own words and understanding.

Length: 20-30 minutes.

Here is a (non-exuahstive) list of affordable software to record and edit videos (thanks to Pascal):

- Keynote: you can record narrations directly in keynote. Advantage: you can record one slide at a time and keynote is free. Here is a tutorial for recording presentations in keynote.
- Powerpoint: you can record narrations directly in powerpoint. Advantage: you can record one slide at a time and you can change a slide without having to re-record the narration. Here is a tutorial for recording presentations in powerpoint. You will need a desktop version of powerpoint. At the moment, a personal subscription to Microsoft 365 costs $8/month.
- Zoom: you can record your desktop with narrations and a video of yourself directly in Zoom. Advantage: free. Here is a tutorial that explains how to record presentations.
- OpenShot: open source video editor where you can post-process any video (trim sections, combine various tracks, add overlays) on any platform (Windows, Mac, Linux). Advantage: free. Here is a tutorial that explains the basic features available in OpenShot.
- Screencast-o-matic: record and edit videos. Advantage: professional comprehensive software for both recording and editing. You can record your desktop with narrations and a video of yourself. You can also postprocess the resulting videos (trim sections, combine various tracks, add overlays). There are several tutorials that explain how to record and edit videos. While you can record videos of up to 15 minutes for free, a subscription of US$1.65/month allows you to record videos of any length and to edit them.

Deadline: See the lecture schedule and sign up on Piazza before the reading week. We expect to have 5 students present in each week after the reading week. Sign up is first come first serve. Send a link to your video and slides to Yaoliang on Monday 11:59pm (latest) of the week you are presenting.

Suggested structure for the presentation

- Introduction
  - What is the paper about?
  - What is the problem tackled?
  - What is the solution proposed?

- Background
  - Explain the problem tackled
  - Explain necessary background
  - Discuss related work briefly
- Content of paper
  - Explain proposed solution at a high level
  - Explain advantages and disadvantages of the proposed solution
  - Compare the solution to other work
- Empirical evaluation
  - Summarize the empirical evaluation
  - Showcase some figures and tables
  - Highlight anything that's special
- Conclusion
  - What are the contributions?
  - What is the take home message?
  - What are possible extensions for future work?

Project (25% report + 15% review)

Students are expected to conduct a research project: your project could be a survey of a topic in causal inference, or an empirical comparison of several causal inference algorithms on an interesting dataset, or an application of causal inference to a new field/application, or designing a novel algorithm to address a need in causality, or theoretically analyzing an existing causal inference algorithm (new or old). The list is a good starting point to form some initial ideas.

You project should

- relate to the course (obviously)
- allow you to learn something new (and hopefully significant)
- be interesting and nontrivial, preferably publishable in a top conference/journal

You are allowed to form a team of *at most 2 members*. However, you need to justify why this is necessary and clearly describe who did what in your report. Each team member will get the same mark for the project (irrespective of what is actually accomplished by each member). So, choose your teammate wisely! It is the team's responsibility to make sure each member is on track. There will be no extra credit for doing everything alone.

The project report will be due shortly after the last day of the term (**11:59pm, Dec 9, 2020**). Please summarize all your findings (empirical, algorithmic, theoretical) in a scientific report. We expect there is an introduction section, a background section, a main result section, and a conclusion section. Depending on your project, you may include an experimental section and/or discussion section. Please always give proper citations to prior work or results. Be precise and concise. The report should be at most 8 pages (excluding references).

Your project report will be evaluated by your peers according to its **clarity, significance, rigor, presentation, and completeness**. The review will be due in a week after the report due (**11:59pm, Dec 16, 2020**).

Academic Integrity

In order to maintain a culture of academic integrity, members of the University of Waterloo community are expected to promote honesty, trust, fairness, respect and responsibility. Check the university website for more information.

Grievance

A student who believes that a decision affecting some aspect of hisher university life has been unfair or unreasonable may have grounds for initiating a grievance. Read Student Petitions and Grievances, Section 4. When in doubt please be certain to contact the department's administrative assistant who will provide further assistance.

Discipline

A student is expected to know what constitutes academic integrity to avoid committing an academic offence, and to take responsibility for hisher actions. A student who is unsure whether an action constitutes an offence, or who needs help in learning how to avoid offences (e.g., plagiarism, cheating) or about "rules" for group work/collaboration should seek guidance from the course instructor, academic advisor, or the undergraduate Associate Dean. For information on categories of offences and types of penalties, students should refer to Policy 71, Student Discipline. For typical penalties check Guidelines for the Assessment of Penalties.

Avoiding Academic Offenses

Most students are unaware of the line between acceptable and unacceptable academic behaviour, especially when discussing assignments with classmates and using the work of other students. For information on commonly misunderstood academic offenses and how to avoid them, students should refer to the Faculty of Mathematics Cheating and Student Academic Discipline Policy.

Appeals

A decision made or penalty imposed under Policy 70 (Student Petitions and Grievances) (other than a petition) or Policy 71 (Student Discipline) may be appealed if there is a ground. A student who believes he/she has a ground for an appeal should refer to Policy 72 (Student Appeals).

Note for Students with Disabilities

AccessAbility Services, located in the new addition to Needles Hall, Room 1401, collaborates with all academic departments to arrange appropriate accommodations for students with disabilities without compromising the academic integrity of the curriculum. If you require academic accommodations to lessen the impact of your disability, please register with the office at the beginning of each academic term.

Intellectual Property

Students should be aware that this course contains the intellectual property of their instructor, TA, and/or the University of Waterloo. Intellectual property includes items such as:

- Lecture content, spoken and written (and any audio/video recording thereof);
- Lecture handouts, presentations, and other materials prepared for the course (e.g., PowerPoint slides);
- Questions or solution sets from various types of assessments (e.g., assignments, quizzes, tests, final exams); and
- Work protected by copyright (e.g., any work authored by the instructor or TA or used by the instructor or TA with permission of the copyright owner).

Course materials and the intellectual property contained therein, are used to enhance a student's educational experience. However, sharing this intellectual property without the intellectual property owner's permission is a violation of intellectual property rights. For this reason, it is necessary to ask the instructor, TA and/or the University of Waterloo for permission before uploading and sharing the intellectual property of others online (e.g., to an online repository).

Permission from an instructor, TA or the University is also necessary before sharing the intellectual property of others from completed courses with students taking the same/similar courses in subsequent terms/years. In many cases, instructors might be happy to allow distribution of certain materials. However, doing so without expressed permission is considered a violation of intellectual property rights.

Please alert the instructor if you become aware of intellectual property belonging to others (past or present) circulating, either through the student body or online. The intellectual property rights owner deserves to know (and may have already given their consent).