Research Topics in Statistical Machine Learning

Important Links

- **Course web site:** `https://www.cs.toronto.edu/~cmaddis/courses/sta4273_w21/`
- **Quercus:** `https://q.utoronto.ca/courses/197351`
- **Piazza:** See Quercus.
- **Zoom:** See Quercus.

Course materials (schedule, slides, readings, assignments) can be found on the course web site.

Overview

Decision-making under uncertainty is typically formalized as the problem of minimizing an expected cost (or maximizing an expected reward). The decision-maker takes an action by sampling from a distribution over actions, and it recieves a cost for that action. The problem is to find the action distribution that minimizes the decision-maker's expected cost.

This problem may seem rather specific, but it appears throughout machine learning and statistics. The most prominent example is Bayesian inference, which can be cast in this paradigm as a variational optimization problem. More broadly, progress on optimizing expected values would improve generative models of real-world data, neural network models with calibrated uncertainties, reinforcement learning algorithms, and many other application areas.

This seminar course introduces students to the various methodological issues at stake in the problem of optimizing expected values and leads them in a discussion of its recent developments in machine learning. The course emphasizes the interplay between reinforcement learning and Bayesian inference. The topics that this course will explore include the following:

- Iterative methods, including gradient descent, policy iteration, and value estimation.
- Gradient estimation, including the score function and reparameterization estimators.
- Variational objectives, including the evidence lower bound.
- Offline policy evaluation and policy optimization.
- Online policy optimization.
- The relationship between planning and policy optimization.
- The relationship between optimal control and probabilistic inference.

Teaching Staff

Instructor: Chris J. Maddison
Teaching Assistants: Cait Harrigan, Farnam Mansouri
Staff e-mail: `sta2473-min-expectations@cs.toronto.edu`

Please send course-related e-mails to the instructor/staff emails above. Please use Piazza for any questions about the course content and to form groups. If you have a private matter that you would like to discuss with the instructor, you can email `cmaddis@cs.toronto.edu`.

Schedule

Lecture Times

- LEC9101: Thursdays 13:00-15:00.

Lecture Structure With the exception of the first two weeks, each week students will be presenting recent papers from the literature. Every student will present a paper once during the course.

The weeks are organized into themes and associated with a list of recent references. No one is expected to read every paper on the list for each week, but there will be some recommended readings for the whole class. See the course web page for information about topics.

Online Delivery Class will be held synchronously online every week via Zoom. The lectures will be recorded for asynchronous viewing by enrolled students. All students are encouraged to attend class each week. Information on attending class, attending office hours, viewing recorded lectures, and using Piazza is available on Quercus.

Course videos and materials belong to your instructor, the University, and/or another source depending on the specific facts of each situation, and are protected by copyright. In this course you are permitted to download session videos and materials for your own academic use, but you should not copy, share, or use them for any other purpose without the explicit permission of the instructor. For questions about recording and use of videos in which you appear please contact your instructor.

Prerequisites

This is a graduate course designed to guide students in an exploration of the current state of the art, so while there are no formal prerequisites, a certain level of mathematical maturity is expected. In particular, a previous course in machine learning such as CSC311, CSC412, CSC413, STA414, or ECE521 is strongly recommended. In addition, it is strongly recommended that students have a strong background in linear algebra, multivariate calculus, probability, and computer programming.

Course Evaluation

- 25% – Paper presentation and code notebook
- 15% – Project proposal
- 60% – Final project

Details will be posted on the course web site.

Submission Policies

There are 4 assignments to submit: presentation slides, a code notebook, a project proposal, and a final project report.

Format. The presentation slides, project proposal, and final project write-up must be submitted in PDF format through Quercus. We encourage typesetting using LaTeX, but other formats are acceptable as long as they are legible.

The code notebook must be submitted through Quercus as a Jupyter notebook (`.ipynb` file format) that can be run on Google Colab (`https://colab.research.google.com`) or UToronto's Jupyter Hub (`https://jupyter.utoronto.ca`). We strongly encourage Python, but a Jupyter notebook running R is acceptable as long as it runs without bugs. We will not attempt to decipher a buggy Jupyter notebook running R.

Lateness. Assignments will be accepted up to 3 days late, but 10% will be deducted for each day late, rounded up to the nearest day. No credit will be given for assignments submitted after 3 days. Extensions will be granted only in special situations, and you will need a Student Medical Certificate or a written request approved by the course coordinator at least one week before the due date.

Collaboration policy. Collaboration on the presentation and final project is allowed. The teams do not need to be the same for the presentation and the final project. The report for the final project should list the contributions of each team member.

Remarks. Remark requests will be considered by the same TA who marked the assignment. The deadline for requesting a remark is one week after the marked assignments are returned. Remarks may result in a decrease in the grade.

Auditing

It is possible for non-enrolled persons to audit this course (sit in on the lectures) *only if the auditor is a student at U of T, and no University resources are to be committed to the auditor*. This means that students of other universities, employees of outside organizations, or any other non-students, are not permitted to be auditors.