CSC 412/2506, Spring 2017: Probabilistic Learning and Reasoning

Instructor: David Duvenaud <duvenaud@cs.toronto.edu>
Office hours: Tuesdays 1pm - 2pm in Pratt room 384

Lectures: Tuesdays 3:00-5:00pm
Tutorials: Thursday 1:00-2:00

Teaching Assistants: Geoffrey Roeder, Jake Snell, Assimakis Kattis, David Madras

Course webpage: www.cs.toronto.edu/~duvenaud/courses/csc412

Piazza: https://piazza.com/utoronto.ca/winter2017/csc412/home

Marking Scheme:

Assignment 1: due Feb 10th worth 15%
Assignment 2: due March 3rd worth 15%
Assignment 3: due March 24th worth 20%
1-hour Midterm: Feb 16th worth 20%
Course project: due April 10th worth 30%

Collaboration policy. You are expected to work on the homeworks by yourself. The report you hand in should be entirely your own work and you may be asked to demonstrate how you got any results that you report.

Programming Assignments A typical assignment will require you to write (or modify) and use some Python code that implements a simple version of a learning procedure that has recently been covered in the course. You will have to submit a brief report (roughly one page plus figures) that describes the results you obtained.

Course project The idea of the project is to give you some experience trying to do a piece of original research in machine learning and writing up your results in a paper style format. What we expect to see is an idea/task that you describe clearly, relate to existing work, implement and test on a dataset. To do this you will need to write code, run it on some data, make some figures, read a few background papers, collect some references, and write a few pages describing your task, the algorithm(s) you used and the results you obtained. As a rough rule of thumb, spend about one week's worth of work (spread out over a longer time to allow the computers to do some work in the interim!), and about a day writing it up after that. Projects can be done individually, or in pairs. We encourage you to work in pairs, but of course, the expectations will be higher for pair projects.

Graduate students will be evaluated at the graduate level according to the University Assessment and Grading Practices Policy. Graduate students in CSC2506 may have some additional questions to do on assignments (these questions will not be required for undergraduates).

Course Outline:
This course covers some of the theory and methodology of statistical aspects of machine learning. The preliminary set of topics to be covered include:

Linear methods for regression + classification, Bayesian linear regression
Probabilistic Generative and Discriminative models, Regularization methods
Stochastic Optimization and Neural Networks
Graphical model notation and exact inference.
Mixture Models, Bayesian Networks
Model Comparison and marginal likelihood
Stochastic Variational Inference
Time series and recurrent models
Gaussian processes
Variational Autoencoders

Readings: There is no required textbook for the class. A few small readings may be assigned if the need arises. These required readings will all be available on the web, for free. There are also some relevant resources which are freely available online. We will try to provide links on a lecture-by-lecture basis.

- David MacKay's excellent textbook (2003) Information Theory, Inference, and Learning Algorithms is an excellent introduction to probabilistic reasoning and Bayesian models. http://www.inference.phy.cam.ac.uk/mackay/itila/
- Ian Goodfellow, Yoshua Bengio and Aaron Courville (2016) have a recent textbook on Deep Learning at http://www.deeplearningbook.org/
- Richard Socher's lecture notes, focusing on RNNs. http://cs224d.stanford.edu/syllabus.html

Assignments are due by 23:59 on the due date.

- Late penalty for assignments: 15% will be deducted from assignments each day after the due date (rounding up). Assignments won't be accepted after 48 hours after the due date.

Marking concerns

Any requests to have marked work re-evaluated must be made in writing within *one week* of the date the work was returned. The request must contain a justification for consideration.

Additional help

Need extra help with the coursework? Here are some options:

- For continued class discussion and questions outside of class, try posting on the piazza. The instructor and TAs will be monitoring them
- You can visit the instructor or teaching assistants during their office hours
- E-mail should only be used for emergencies or personal matters

How to communicate with your instructor

Questions about course material such as:

- How do I do question 3.7 in the textbook?
- What is standard deviation?
- When is the midterm?

should be posted on the discussion forums.

For private communication, such as "I missed the test because I was ill," e-mail your instructor. Use your utoronto.ca e-mail account to ensure that your message doesn't automatically go to a Junk folder and include your full name and student number.

Academic integrity

You are responsible for knowing the content of the University of Toronto's Code of Behaviour on Academic Matters at www.governingcouncil.utoronto.ca/policies/behaveac.htm. If you have any questions about what is or is not permitted in this course, please do not hesitate to contact your instructor.

Accessibility needs

The University of Toronto is committed to accessibility. If you require accommodations for a disability, or have any accessibility concerns about the course, the classroom, or course materials, please contact Accessibility Services as soon as possible: accessibility.services@utoronto.ca or http://accessibility.utoronto.ca.