

Course Syllabus

Advanced Topics in High-Performance Computing (MCSC 6230G/7230G)

Faisal Qureshi

<http://faculty.uoit.ca/qureshi>

Computer Science, Faculty of Science
University of Ontario Institute of Technology

Fall 2017



Description

This is an introductory graduate course in machine learning. This course will focus on both supervised and un-supervised learning methods, covering both theory and practice. The course is geared towards students who wish to develop a working knowledge of the recent advances in machine learning, and how these are applied in various domains.

Machine learning deals with how to design computer programs that learn from “experience.” Residing at the intersection of computer science and statistics, machine learning aims to extract useful information from data (often referred to as the *training data*) and leverages this information to create computer models capable of carrying out useful, non-trivial tasks, such as designing cars that can drive on their own, filters for blocking junk email, diagnostics tools for disease discovery, etc. By many accounts machine learning is the “greatest export” of computer science (and statistics) to other disciplines.

The course will cover the following topics:

- regression;
- classification;
- clustering;
- dimensionality reduction;
- mixture-models; and
- neural networks and deep learning.

Pre-requisites

The course assumes that students are comfortable with statistics, basic linear algebra, and programming.

I recommend reading Part 1 of “Deep Learning” by I. Goodfellow, Y. Bengio and A. Courville to brush up on linear algebra and statistics. The book is available at [here](#)

We will be using Python for the programming part of this course. For Python, I recommend the [Anaconda](#) distribution, which comes pre-loaded for nearly all the packages that we will be using in this course. Of course you are welcome to use any variant/distribution of Python that suits you.

[Here](#) you’ll find a number of tutorials showcasing Python use in machine learning. I strongly recommend that you become comfortable with the following four Python packages/environment:

- [numpy](#);
- [scipy](#);
- [matplotlib](#); and
- [jupyter notebook](#).

Instructor

Faisal Qureshi
Email: faisal.qureshi@uoit.net
Office: UA4032
Web: <http://faculty.uoit.ca/qureshi>

Lectures

Wednesday, 12:40-3:30 pm, ERC3027
First day of lectures, Wednesday, September 13, 2017.

Office hours

Tuesday, 1-2 pm in UA4032, or by appointment.

Grading

- Assignments, 40%
 - A1, Sep. 27 - Oct. 11
 - A2, Oct. 18 - Nov. 1
 - A3, Nov. 8 - Nov. 22
- Presentation, 10%
- Project, 50%

A student needs to get 60% marks in the project to successfully complete the course.

Project/presentation expectations for students enrolled in MCSC 7230G are different from those enrolled in MCSC 6230G. E.g., level of difficulty for the course project is higher for those enrolled in MCSC 7230G.

Lectures

Attendance is mandatory. We will discuss topics in class that are not easily found in any single textbook. Most lectures will contain a programming activity. You are asked to bring your laptops to the lectures.

Assignments

There will be three assignments. These assignments will involve implementing machine learning algorithms and testing these algorithms on test datasets. The assignments follow the philosophy that the best way to learn machine learning is to program it yourself and experiment with it. You will also be required to submit a written report for each assignment. The writeup will be between 3 to 4 pages long. The report will discuss your results. It is strongly recommended that you use Python programming language for your assignments.

Presentation

Each student will give a presentation on a machine learning topic. The topic will be decided in consultation with the course instructor.

Project

Course project is worth 50% of the marks. A project can be implementation oriented—where a student implements a machine learning algorithm—or application oriented—where a student attempts to solve a problem (of suitable difficulty) by applying machine learning techniques. Project topic will be decided in consultation with the instructor.

One page project proposals are due by Oct. 31. The project report, in the form of a paper, is due by Dec. 10. Each student will also give a short project presentation.

Project grade will depend on the ideas, how well you present them in the report, how well you position your work in the related literature, how thorough are your experiments and how thoughtful are your conclusions.

Tentative Schedule

- Week 1 - Introduction and non-parametric learning
- Week 2 - Clustering (k-means, k-medoids, meanshift, agglomerative)
- Week 3 - Linear regression, MLE view of linear regression
- Week 4 - Logistic regression, softmax
- Week 5 - Neural networks
- Week 6 - Convolutional neural networks
- Week 7 - Convolutional neural networks
- Week 8 - Bayesian learning
- Week 9 - Gaussian processes
- Week 10 - Decision trees
- Week 11 - Random forests
- Week 12 - Project presentations

Course work submission

Unless otherwise instructed, all course work should be submitted using Blackboard.

Partial marks

Assignments will primarily be evaluated based on the correctness of solutions; however, partial credit may be assigned for documentation, discussion, etc.

Remarking

It is extremely important that all work is fairly graded. Please submit a remark request by email within 5 days of receiving the grade. The email must contain the reasons for which you think the work should be remarked. *Please note that a remark may result in a lower grade.*

Late submission policy

The penalty for a late submission is 10% per day. An assignment or project will get a **zero** if submitted more than 48 hours after the submission deadline. A doctor's note will be needed to avoid late submission penalty.

Email traffic

The instructor and the TA will make every effort to respond to emails in a timely manner; however, it may take up to two working days to respond to an email. It simply means that emails sent right before a deadline may not be answered in time. Urgent emails may be sent to "faisal.queeshi@uoit.net" with the subject line "mcsc ml - fall 2017".

Discussions

Appropriate use of discussion groups include clarification of lecture material and assignments and other concerns and comments about the course that might of general interest to course participants. Please do not post assignment solutions to the discussion groups.

Collaboration

I encourage you to work together when discussing assignments/projects; however, it doesn't mean that you should share your written solutions or that you submit someone else's work as your own.

Course evaluation

It is important that every student participates in course evaluations. Course evaluations, which are completely anonymous, provide extremely useful feedback to the instructor and the TA, helping improve the course.

Important dates

UOIT academic calendar that lists important dates (and deadlines) is available at [here](#).

Academic integrity

Assignments and tests must be strictly individual work. UOIT takes academic dishonesty very seriously. Please read and understand UOIT's policy on academic integrity available [here](#)

Accessibility

Students with disabilities may request to be considered for formal academic accommodation in accordance with the Ontario Human Rights Code. Students seeking accommodation must make their requests through the Centre for Students with Disabilities in a timely manner, and provide relevant and recent documentation to verify the effect of their disability and to allow the University to determine appropriate accommodations. More information about Student Accessibility Services (SAS) is available [here](#).

Freedom of Information and Protection of Privacy Act

UOIT is governed by the Freedom of Information and Protection of Privacy Act ("FIPPA"). In addition to providing a mechanism for requesting records held by the university, this legislation also requires that UOIT not disclose the personal information of its students without their consent. FIPPA's definition of "personal information" includes, among other things, documents that contain both your name and your Banner ID. To ensure that your rights to privacy are protected, I encourage you to use only your Banner ID on assignments or test papers being submitted for grading (the exception to this rule are midterm and final exams, since these are returned individually). This policy is intended to prevent the inadvertent disclosure of your information where graded papers are returned to groups of students at the same time. If you still wish to write both your name and your Banner ID on your tests and assignments, please be advised that UOIT will interpret this as an implied consent to the disclosure of your personal information in the normal course of returning graded materials to students. Please contact the UOIT Chief Privacy Officer at accessandprivacy@uoit.ca for more information.

Sexual Violence Policy

UOIT is committed to the prevention of sexual violence in all its forms. For any UOIT student who has experienced Sexual Violence, **UOIT can help**. UOIT will make accommodations to cater to the diverse backgrounds, cultures, and identities of students when dealing with individual cases.

If you think you have been subjected to or witnessed sexual violence:

- Reach out to a Support Worker, who are specially trained individuals authorized to receive confidential disclosures about incidents of sexual violence. Support Workers can offer help and resolutions options which can include safety plans, accommodations, mental health support, and more. To make an appointment with a Support Worker, call 905.721.3392 or email supportworker@uoit.ca
- Learn more about your options at: www.uoit.ca/sexualviolence