COMP597: Applications of Machine Learning in Real World Systems

General Information

Prerequisites:
Math courses: Calculus (Math 141), Linear Algebra (Math 223), Matrix Numerical Analysis (Math 327), Probability Theory (Math 323 or ECSE 205 or ECSE 305)
Computer Systems Courses: Operating Systems (COMP 310).
Machine Learning courses: COMP551 (or COMP-652/ ECSE-608)
Course Quota: 30.

Please Note: **These prerequisites are strict**, unless you get special approval from the course instructor. Please bring your transcript(s) and a detailed CV with detailed project and research experiences.

Class web page: myCourses
Instructor: Prof. Xue (Steve) Liu
Office: Room 326, McConnell Engineering Building
Office Hours: TBA
Location: TBA

Email: solve the reCAPTCHA on my webpage (Recommended means for communication. Do NOT use the phone.)

Note: Emails should be sent from your official McGill email address in order to be responded. Emails not from your official McGill email address will NOT be properly filtered hence will not be responded. Emails should be started with the title "COMPXXX: *** ". For every email communication, please make sure to use "COMPXXX:" as a single word (with no spaces) as the start of the title and replace *** with your topic/questions.
Please note: Due to the LARGE number of emails (together with all the spams) we receive every day, emails not started with this title may be categorized as spams by the spam filter and will not be responded. We will NOT use WebCT email. Thank you for your understanding.

Academic Integrity: ***"McGill University values academic integrity. Therefore, all students must understand the meaning and consequences of cheating, plagiarism, and other academic offences under the Code of Student Conduct and Disciplinary Procedures (see www.mcgill.ca/integrity for more information)."***

Course Overview

This is an advanced undergraduate level / junior graduate-level course for students who are interested in learning how to apply machine learning algorithms on solving real-world problems.

We will first begin with a general introduction for machine learning basics and then focused on selected on a few selected interested topics including computer systems, communication networks, transportation system, smart grid, and go games. We will also discuss some high-impact industry machine learning products and the research problems behind their successes. The class consists of instructor lectures, student presentations, class discussions, and class project(s). he main topics that are covered in this course are listed as follows.

Course Syllabus[1]

1. Review of Machine Learning (week 1)
2. Review of Machine Learning (Week 2)
3. Machine Learning for Computer Systems (Weeks 3-4)
4. Machine Learning for Computer Systems (Week 5)
5. **Project topic selection – Phase 0 (Week 6)**
6. Machine Learning for Communication Networks (Week 7)
7. Machine Learning for Communication Networks (Week 8)
8. **Project progress presentation — Phase I (week 9)**
9. Machine Learning for Transportation System (Week 9)
10. Machine Learning for Transportation System (Week 10)
11. **Project progress presentations/initial demo – Phase II (week 11)**
12. Machine Learning for Smart Grid (Week 11)
13. Overview of research challenges to bring machine learning to real-world (Week 12)
14. **Project report due/final demo (week 13)**

Textbooks: Due to the nature of this course, many topics are not available in any textbook yet. Hence there is no required textbook. Instead, we will use many publications, technical notes or materials available from the Internet.

Course evaluation:

[1] Subject to change based on the course progress and student presentation sign ups

Class participations/presentations: 40%
Project, project progress presentations, and final report: 60%
(First presentation 5%, 2nd presentation 15%, final report and deliverables: 40%)
Total 100%

Final project: Two intermediate (milestone) presentations are needed for your final project. These presentations/demos are expected to receive proper feedbacks on your project. Also, these milestone presentations/demos could help you pace the project for timely completion. Every final project must have a final report (and other associated materials, such as software programs) submitted for final grading.

Late project policy: There will be a strict deadline for the final project (the week before the final exams – subject to change). Please pay close attention to the announcements during the course. It is your responsibility to make sure that the final project report (and its associated materials, if any) is properly submitted via myCourses.