CSC 411
MACHINE LEARNING and DATA MINING

Lectures:	Monday, Wednesday 12-1 (section 1), 3-4 (section 2), Thursday 6-8 (section 3)
Lecture Room:	MP134 (section 1), SS2106 (section 2), BA1200 (section 3)
Instructosr:	Raquel Urtasun, Ruslan Salakhutdinov <csc411prof@cs.toronto.edu>
Office hours:	TBA, and by appointment
Teaching Assistants:	TBA
TA email:	<csc411ta@cs.toronto.edu>
Tutorials:	Friday 12-1 (section 1), 3-4 (section 2), Thursday 8-9 (section 3)
Tutorial Room:	MP134 (section 1), SS2106 (section 2), BA1200 (section 3)
Class URL:	www.cs.toronto.edu/~urtasun/courses/CSC411/CSC411_Fall15.html

Overview

Machine learning research aims to build computer systems that learn from experience. Learning systems are not directly programmed by a person to solve a problem, but instead they develop their own program based on examples of how they should behave, or from trial-and-error experience trying to solve the problem. These systems require learning algorithms that specify how the system should change its behavior as a result of experience. Researchers in machine learning develop new algorithms, and try to understand which algorithms should be applied in which circumstances.

Machine learning is an exciting interdisciplinary field, with historical roots in computer science, statistics, pattern recognition, and even neuroscience and physics. In the past 10 years, many of these approaches have converged and led to rapid theoretical advances and real-world applications.

This course will focus on the machine learning methods that have proven valuable and successful in practical applications. This course will contrast the various methods, with the aim of explaining the circumstances under which each is most appropriate. We will also discuss basic issues that confront any machine learning method.

Pre-requisites

You should understand basic probability and statistics, (STA 107, 250), and college-level algebra and calculus. For example it is expected that you know about standard probability distributions (Gaussians, Poisson), and also how to calculate derivatives. Knowledge of linear algebra is also expected, and knowledge of mathematics underlying probability models (STA 255, 261) will be useful. For the programming assignments, you should have some background in programming (CSC 270), and it would be helpful if you know Matlab or Python. Some introductory material for Matlab will be available on the course website as well as in the first tutorial.

Readings

There is no required textbook for this course. There are several recommended books. We will recommend specific chapters from two books: *Introduction to Machine Learning* by Ethem Alpaydin, and *Pattern Recognition and Machine Learning* by Chris Bishop. We will also recommend other readings.

Course requirements and grading

The format of the class will be lecture, with some discussion. I strongly encourage interaction and questions. There are assigned readings for each lecture that are intended to prepare you to participate in the class discussion for that day.

The grading in the class will be divided up as follows:

Assignments	40%
Mid-Term Exam	25%
Final Exam	35%

There will be three assignments; the first two are worth 12.5% each, and the last one 15% of your grade.

Homework assignments

The best way to learn about a machine learning method is to program it yourself and experiment with it. So the assignments will generally involve implementing machine learning algorithms, and experimentation to test your algorithms on some data. You will be asked to summarize your work, and analyze the results, in brief (3-4 page) write ups. The implementations may be done in any language, but Matlab or Python is recommended. A brief tutorial on Matlab is available from the course web-site.

Collaboration on the assignments is not allowed. Each student is responsible for his or her own work. Discussion of assignments and programs should be limited to clarification of the handout itself, and should not involve any sharing of pseudocode or code or simulation results. Violation of this policy is grounds for a semester grade of F, in accordance with university regulations.

The schedule of assignments is included in the syllabus. Assignments are due at the beginning of class/tutorial on the due date. Because they may be discussed in class that day, it is important that you have completed them by that day. Assignments handed in late but before 5 pm of that day will be penalized by 5% (i.e., total points multiplied by 0.95); a late penalty of 10% per day will be assessed thereafter. Extensions will be granted only in special situations, and you will need a Student Medical Certificate or a written request approved by the instructor at least one week before the due date.

For the final assignment, we will have a *bake-off*: a competition between machine learning algorithms. We will give everyone some data for training a machine learning system, and you will try to develop the best method. We will then determine which system performs best on some unseen test data.

Exams

There will be a mid-term in class on October 26^{th}, which will be a closed book exam on all material covered up to that point in the lectures, tutorials, required readings, and assignments.

The final will not be cumulative, except insofar as concepts from the first half of the semester are essential for understanding the later material.

The exams will cover material presented in lectures, tutorials, and assignments. You will not be responsible for topics in the reading not covered in any of these.

Attendance

We expect students to attend all classes, and all tutorials. This is especially important because we will cover material in class that is not included in the textbook. Also, the tutorials will not only be for review and answering questions, but new material will also be covered.

Electronic Communication

If you have questions about the assignments, you should send email to the TA account, and cc me on it. You should include your full name in the email, and it will also be useful to include your CDF account name and/or student number. Feel free to email me with questions or comments about the material covered in the course, or other related questions.

For questions about marks on the assignments, please first contact the TA. Questions about the exams should be addressed to me.

CLASS SCHEDULE, Part 1

Shown below are the topics for lectures and tutorials (in italics), as are the dates that each assignment will be handed out and is due. The notes from each lecture and tutorial will be available on the class web-site the day of the class meeting. The assigned readings are specific sections from the book. All of these are subject to change.

Date	Topic	Assignments
Sep 14	Introduction	
Sep 16	Linear Regression	
Sep 18	*Probability for ML & Linear regression*	
Sep 21	Linear Classification	
Sep 23	Logistic Regression	
Sep 25	*Optimization for ML*	
Sep 28	Nonparametric Methods	
Sep 30	Decision Trees	
Oct 2	*kNN & Decision Trees*	Asst 1 Out
Oct 5	Multi-class Classification	
Oct 7	Probabilistic Classifiers	
Oct 9		
[Oct 12]	Thanksgiving: No class	
Oct 14	Probabilistic Classifiers II	
Oct 16	*Naive Bayes and Gaussian Bayes Classifier*	
Oct 19	Neural Networks I	Asst 1 In
Oct 21	Neural Networks II	
Oct 23	*Mid-term review*	
Oct 26	MIDTERM	

CLASS SCHEDULE, Part 2

Date	Topic	Assignments
Oct 28	Clustering	Assit 2 Out
Oct 30	*Clustering*	
Nov 2	Mixture of Gaussians & EM	
Nov 4	PCA & Autoencoders	
Nov 6	*PCA Tutorial*	
[Nov 9]	Mid-term break: No class	
Nov 11	Kernels and Margins	Asst 2 In
Nov 13	*SVM Tutorial*	Asst3 Out
Nov 16	Support Vector Machines	
Nov 18	Ensemble Methods I	
Nov 20	*Bagging & Boosting*	
Nov 23	Ensemble Methods II	
Nov 25	Bayesian Methods	
Nov 27		
Nov 30	Reinforcement Learning I	
Dec 2	Reinforcement Learning II	
Dec 4		
Dec 7	Final & Wrap-up	Ass 3 In