ECE521 Inference Algorithms and Machine Learning

Winter 2017 Syllabus Draft 1

Teaching Team:

Instructor:	Jimmy Ba	**Session:**	M 10 – 11 TH 9 – 11
Email:	jimmy@psi.toronto.edu	**Location:**	BA1130
Office hours:	Mondays 11am-12pm in BA4161		
Instructor:	Mark Ebden	**Session:**	M 11 – 12 TH 12 – 14
Email:	mark.ebden@utoronto.ca	**Location:**	RS211 LM161
Office hours:	Thursdays 3-4 pm in SS6026C		

Course Pages:

1. Website: `http://ece521.github.io`
2. Piazza forum: `https://piazza.com/utoronto.ca/winter2017/ece521/home`

Teaching Assistants:

Min Bai
Sinisa Colic
Kaustav Kundu
Renjie Liao
Mengye Ren
Eleni Triantafillou
Shenlong Wang
Yuhuai Wu
Tianrui Xiao
Haiyan Xu

TA Contact: Please post your questions related to tutorials, assignments and exams on Piazza. Do NOT send emails about the class to the personal emails of the TAs directly. We will not answer.

Office Hours: After class, or by appointment, or post your questions in the forum provided for this purpose on Piazza.

Course Objectives: The twenty-first century has seen a series of breakthroughs in statistical machine learning and inference algorithms that allow us to solve many of the most challenging scientific and engineering problems in artificial intelligence, self-driving vehicles, robotics and DNA sequence analysis. In the past few years, machine learning applications in search engines, wearable devices and social networks have broadly impacted our daily life. These algorithms adapt to the data at hand and are tolerant to noisy observations. The goal of this course is to provide principled mathematical tools to solve statistical inference problems you may encounter later. The first half of the course covers the fundamentals of statistical machine learning and supervised learning models. The second half of the course focuses on probabilistic inference and unsupervised learning. The examples of the course include object recognition; image search, document retrieval; sequence filtering and alignment; and data compression. This course reviews state-of-the-art algorithms and models for probabilistic inference and machine learning.

Prerequisites: You must be familiar with the fundamentals of probability theory as per STA286 or ECE302. An undergraduate-level understanding of linear algebra and calculus is assumed. Knowledge of (or being eager to learn) Python programming is essential.

Main References: There is no required textbook for this course. You may find the following textbooks useful to gain additional insight:

- Christopher M. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006.
- Kevin P. Murphy, *Machine Learning: A Probabilistic Perspective*, MIT Press, 2012.
- David J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*, Cambridge University Press, 2003. (free from: http://www.inference.phy.cam.ac.uk/itprnn/book.pdf)
- Jerome H. Friedman, Robert Tibshirani, and Trevor Hastie, *The Elements of Statistical Learning*, Springer, 10th ed., 2013. (free from: http://statweb.stanford.edu/~tibs/ElemStatLearn/printings/ESLII_print10.pdf)
- Ian Goodfellow, Yoshua Bengio and Aaron Courville, *Deep Learning*, MIT Press, 2016. (free from: http://www.deeplearningbook.org)

If you prefer more animated learning materials, here is a list of excellent video lectures you can find online for free:

- Video lectures from Hugo Larochelle. http://info.usherbrooke.ca/hlarochelle/neural_networks/content.html
- Video lectures from Geoffrey Hinton. https://www.coursera.org/learn/neural-networks
- Metacadamy, an online website which helps you construct personalized learning plans and which has links to lots of resources relevant to particular concepts. http://www.metacademy.org

Grading Policy: Four assignments (40%), Midterm (25%), Final (35%).

Exams:

- All exams (midterm and final) are closed book and you are allowed to bring a double-sided aid sheet.
- The midterm exam will be held in class on Feb. 16th from 9-11 am. The details will be announced later.
- In the final exam, the ratio of pre-midterm material to post-midterm material will be approximately 1:4.
- Missed tests will get a score of 0 except in the case of an official Student Medical Certificate or a written (not emailed) request submitted at least one week before the test date and approved by the instructor.

Assignments and Computing:

- This course involves deriving algorithms mathematically, writing programs to implement them, and performing experiments with data sets. While the software you'll need to write will not be very complex, you will need to write your own programs, debug them, and use them to conduct various experiments, plot curves, classify images.

- You will complete the assignments by submitting a report including your solutions, experimental findings and programming codes. You are encouraged to typeset your assignment write-ups, for example using LaTeX, but this is not required.
- The programming part of the assignments will all be done in Python using TensorFlow library https://www.tensorflow.org/, but prior knowledge of TensorFlow is not required. Basic TensorFlow will be taught in a tutorial.
  - You can install Python and TensorFlow yourself on your own machine. For most of you, this will be the most convenient option. Alternatively, you may choose to use ECE lab machines in GB243, GB251E, SF2102, SF2204
  - You can install TensorFlow by following: https://www.tensorflow.org/get_started/os_setup
  - TensorFlow Python API reference is your best friend to complete the assignment: https://www.tensorflow.org/api_docs/python/
  - You may want to use Matplotlib to plot and visualize your experiments: http://matplotlib.org/
- All homework solutions, programming code, etc., must be submitted electronically. For electronic submission, email ece521ta@gmail.com with your .PDF files and code files as attachments. If there is something that you would like the TAs to know while grading your assignment, please write it in the body of the email. Do not send your homework solutions, code, etc to instructors or TAs directly.
- You may discuss assignments with other students and write up your assignments in a group of two. Please indicate your group members **and their contribution percentage** at the beginning of your assignment reports.
- Don't post solved assignments, or significant parts of a solved assignment on Piazza.
- Late assignment submission, except in the case of an official Student Medical Certificate, will be accepted with 25% penalty every 24 hours from the deadline. So, you will get 0% on the assignment if it is submitted 4 days late. Each group has 48 hours "grace budget" for the entire semester that can be used across the assignments without penalty. Assignments not submitted electronically before the deadline will be considered late. We will use the time-stamp of the electronic submissions to count for lateness. The time past the deadline will be rounded up.
- The tentative assignment deadlines are: Wednesday 25 January, Wednesday 15 February, Wednesday 15 March, and Wednesday 5 April.

Tutorials: Tutorials start after Jan 12th. It is important that you attend the tutorials. The tutorial instructors will review topics that were recently covered in lectures, go over material that is directly relevant to the assignments, and answer your questions about the material.

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.

Tentative Course Outline:

09-12 Jan	A little of probability theory; types of loss functions
16-19 Jan	Maximum likelihood estimation; optimization algorithms; regularization
23-26 Jan	Models: the nearest neighbor methods, linear regression
30 Jan-02 Feb	Models: logistic regression; neural networks and deep learning
06-09 Feb	Learning algorithms: backpropagation; principle component analysis; K-means
13-16 Feb	Review and Midterm
20-24 Feb	**Reading week**
27 Feb-02 Mar	Latent variables; Bayesian inference; marginal log likelihood
06-09 Mar	Probabilistic models: mixture models; probabilistic PCA
13-16 Mar	Graphical models: Bayesian networks; Markov random fields; factor graph
20-23 Mar	Graphical models: hidden Markov models; particle filtering
27-30 Mar	Inference algorithms: belief propagation
03-06 Apr	Learning algorithms: generalized EM
10-13 Apr	Review for final exam