1/27/21

Syllabus
CS 6923 Machine Learning
Spring 2021

Instructor: Prof. Lisa Hellerstein
Email: lisa.hellerstein@nyu.edu
Office Hours:

Course format

This course will be given on zoom. This is a synchronous course and you are expected to attend lectures and to have your cameras on during class. A small amount of class participation is expected and will be factored into your grade. Exams will be given during class time.

Grading policy

Your grade will be calculated as follows: midterm 30%, final 40%, class partic-ipation 10%, and the homework/project 20%. No late homework assignments will be accepted.

Textbooks

Many of the lectures will follow the approach used in the lecture notes from cs220 at Stanford: http://cs229.stanford.edu/syllabus.html.

There is no required textbook. Additional resources specific to each lecture will be posted on NYU Classes.

Optional textbooks suitable for all students:

- Alpaydin: Introduction to Machine Learning, by Ethem Alpaydin, Third Edition. Published by MIT Press. (available online through the NYU library). [Good coverage of topics, but terse.]
- Daume: A Course in Machine Learning by Hal Daume (unfinished book draft), `http://ciml.info/` [A more intuitive and accessible approach to some of the course topics.]
- Mitchell: Machine Learning, by Tom Mitchell. Published by McGraw Hill. [Classic text, only covers some of our topics.]

Optional textbooks if you have a strong math/statistics background:

- HTF: Hastie, Tibshirani, Friedman, Elements of Statistical Learning, Second Edition Published by Springer. (available online through the NYU library and `https://web.stanford.edu/~hastie/ElemStatLearn/`)
- Murphy: Machine Learning: a Probabilistic Perspective by Kevin Patrick Murphy
- Bishop: Pattern Recognition and Machine Learning, by Christopher Bishop. Published by Springer.

Optional background reading:

Stanford's machine learning course has a nice review of:

- linear algebra: `http://cs229.stanford.edu/summer2019/cs229-linalg.pdf` although we are unlikely to use sections 3.4, 3.10, 4.5, 4.6, and will approach the topics in sections 3.12, 3.13 in a more intuitive way.
- probability theory: http://cs229.stanford.edu/summer2019/cs229-prob.pdf

A very nice (gentle) introduction to some of the linear algebra concepts can be found at: `https://davetang.org/file/Singular_Value_Decomposition_Tutorial.pdf` and `https://mml-book.github.io/book/mml-book.pdf`.

Prerequisites

The official prerequisites are graduate status and an undergraduate-level course in probability and statistics.

However, it is also important to note Machine Learning is more mathematical than most other graduate CS courses. Students often have difficulty with this course (and risk getting a grade of C) if they know how to program but have not taken much math, or have not done well in their math courses. You should know about the probability density function (pdf), cumulative density function (cdf), continuous probability distributions, conditional probability, and expected values. It is also assumed that you know first-year calculus (derivatives and integrals) and the basics of linear algebra. You should also know about partial derivatives, gradients, and the chain rule (or be prepared to learn about them on your own). You need to know the basics of linear algebra.

Homework:

Homeworks will include written exercises as well as hands-on work involving datasets and programming. All programming must be done with Python/numpy. There will be a homework approximately every two weeks.

Exams:

The course will have a midterm and a final exam.

Attendance at exams is mandatory. Make-up exams will only be given in the case of a emergency, such as illness, which must be documented, e.g. with a doctor's note. In such cases, you must notify me as early as possible, preferably before the exam is given. If you miss an exam without a valid excuse, you will receive a grade of zero for that exam. The exams will be closed book and no notes.

Code of Conduct:

Students are responsible for following the rules in the School of Engineering Code of Conduct (`https://engineering.nyu.edu/campus-and-community/student-life/office-student-affairs/policies/student-code-conduct`)

Absences:

The NYU School of Engineering policies and procedures on excused absences are available here: `https://engineering.nyu.edu/campus-and-community/student-life/office-student-affairs/policies#chapter-id-30199`

Moses Center:

If you are student with a disability who is requesting accommodations, please contact New York University's Moses Center for Students with Disabilities (CSD) at 212-998-4980 or mosescsd@nyu.edu. You must be registered with CSD to receive accommodations. Information about CSD can be found at `www.nyu.edu/csd`.

Tentative Schedule:

Below is a *tentative* schedule of lecture topics:

Lecture	Date	Topics
1	1/28	Introduction, Bayesian Decision Theory, Maximum Likelihood Estimation
2	2/4	QDA, Naive Bayes, k-NN, ML Experiments
3	2/11	Linear Regression
4	2/25	Bias/Variance, Confidence bounds, Regularization
5	3/4	Logistic Regression
6	3/11	Decision Trees and Random Forests
7	3/18	Data Preprocessing, Dimensionality Reduction
8	3/25	Midterm
9	4/1	Perceptrons and Neural Nets
10	4/8	Deep Learning
11	4/15	Kernel Machines, SVMs
12	4/22	Clustering
13	4/29	Ensemble Methods
14	5/6	Reinforcement Learning
15	5/13	FINAL EXAM