ECE 8843 / CS 8803 / ISYE 8803:
Mathematical Foundations of Machine Learning

Fall 2017 Syllabus

August 21, 2017

Summary

The purpose of this course is to provide first year PhD students in engineering and computing with a solid mathematical background for two of the pillars of modern data science: linear algebra and applied probability.

Prerequisites

There are no formal prerequisites. I expect student to have basic exposure to linear algebra and probability, and have basic programming skills in either MATLAB or Python.

Instructor

Justin Romberg
Email: jrom@ece.gatech.edu
Office: Centergy 5227
Office phone: 404-894-3930
Office hours: I am also available to meet in Centergy 5227 by appointment. I will also have regularly scheduled office hours; the times for these will be announced in the near future.

Teaching Assistant

A TA has not been assigned for this course yet. Hopefully this will be taken care of by the end of the first week of class.

Grading

30% Homework, 10 assignments
20% Quiz #1 (preliminary date: September 29)
20% Quiz #2 (early November)
25% Final exam
5% Attendance, see below

Homework

Homework will be assigned weekly (approximately). Homework will be turned in at the beginning of lecture. Late homework will get zero credit.

Students are encouraged to discuss homework problems with one another, however each student must write up and turn in their own solutions.

Unauthorized use of any previous semester course materials, such as tests, quizzes, homework, projects, and any other coursework, is prohibited in this course. Using these materials will be considered a direct violation of academic policy and will be dealt with according to the GT Academic Honor Code.

The homework assignments will be hard; many of them will require significant amounts of time and effort to complete. But this is really where most of the learning takes place. You will get out of the assignments what you put into them. Students who complete all of the assignments in full will be rewarded with a deep understanding of the role that linear algebra plays in modern signal processing (among other things).

Effectively, homework is worth much more than 30% of your grade. In teaching classes like this, **I have yet to see a case where a student does not put effort into the homework assignments but does well on the exams**.

Lecture

Lectures are Monday, Wednesday, and Friday from 12:20-1:10p in CoC 17.
Lecture attendance is mandatory and will count towards your grade. A sign in sheet (starting in week 2) will be passed around at every lecture; please sign next to your name and your name only.

Dead week

As per Institute policy, I am required to inform you on the syllabus that **there will be a homework assignment due during the last week of class**.

Web Page, T-square, and Piazza

The course webpage is located at
http://jrom.ece.gatech.edu/MLMATH-Fall-2017/. (The page will be created the afternoon after the first lecture.

Course information, notes for the lectures, homework assignments, and supplemental materials will be posted here.

Student can also review their grades on T-square.

This course will also make use of Piazza: http://piazza.com Please use this as a resource to post questions about lectures and homework assignments. I will also use it for general course announcements.

I also encourage students to answer questions. The enrollment in this class is fairly large, so there will be a great economy of scale here. **Extra-credit consideration will be given to notable Piazza contributors.**

Text

There is no required text. Extensive course notes will be provided. From time to time, I will point you to books that are good resources for the material we are talking about.

Outline

The outline below should be treated as an approximation; it is subject to (small) changes.

1. Vector space basics
   (a) linear vector spaces, linear independence
   (b) norms and inner products
   (c) bases and orthobases
   (d) examples: Bsplines, cosines/Fourier, radial basis functions, etc.
   (e) linear approximation (closest point in a subspace, least-squares I)
2. Linear Estimation
   (a) Examples: classical regression/recovering a function from point samples, imaging, etc.
   (b) The Singular Value Decomposition (SVD)
   (c) least-squares solutions and the pseudoinverse
   (d) stable inversion and regularization
   (e) kernels, Mercers theorem, RKHS, representer theorem
   (f) computing least-squares solutions

      i. matrix factorizations
      ii. steepest descent and conjugate gradients
      iii. low rank updates for online least-squares

3. Probability and random vectors
   (a) review: joint pdfs, random vectors, conditional probability, Bayes rule
   (b) multivariate Gaussian
   (c) random processes
4. Statistical estimation
   (a) best linear unbiased estimator and weighted least-squares
   (b) maximum likelihood
   (c) computing estimates: unconstrained optimization, stochastic gradient descent
   (d) Bayesian estimation
   (e) hypothesis testing
5. Modeling
   (a) probabilistic graphical models, HMMS, Gaussian graphical models
   (b) geometric models
      i. principal components analysis
      ii. low rank approximation (Eckart-Young theorem)
      iii. structured matrix factorization (e.g. NNMF, dictionary learning)
      iv. sparsity