ECE 6250: Advanced Topics in Digital Signal Processing
Fall 2016 Syllabus

August 22, 2016

Summary

ECE 6250 is a general purpose, advanced DSP course designed to follow an introductory DSP course. The central theme of the course is the application of tools from linear algebra to problems in signal processing.

Prerequisites

An introductory course in digital signal processing. Students should be familiar with the fundamentals of linear algebra and have had exposure to basic probability and statistics. Students should also have basic MATLAB programming skills.

Instructor

Justin Romberg
Email: jrom@ece.gatech.edu
Office: Centergy 5227
Office phone: 404-894-3930
Office hours: I will typically hold scheduled office hours the day before homework is due. More details will be provided soon. I am also available to meet in Centergy 5227 by appointment.

Teaching Assistant

TBA

Grading

30% Homework, 11 (± 1) assignments
20% Quiz #1 (preliminary date: September 28)
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 this class, **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 and Wednesday from 3:05-4:25p in Klaus 2443.
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**.

T-square, and Piazza

Notes for the lectures, homework assignment, and supplemental materials will be posted to the class T-square site. Student can also review their grades on this site.

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. Below is a list of books that I have found helpful over the years for learning (and teaching) the material in this class.

Linear algebra and function spaces:

- Strang: *Linear Algebra and its Applications*
  http://goo.gl/6fW3Hq
- Strang: *Computational Science and Engineering*
  http://goo.gl/ycbDil
- Horn and Johnson: *Matrix Analysis*
  http://goo.gl/9a313C
- Young: *An Introduction to Hilbert Space*
  http://goo.gl/WbPVu2

Mathematics of signal processing:

- Moon and Stirling: *Mathematical Methods and Algorithms for Signal Processing*
  http://goo.gl/sRvZsM
- Vetterli et al: *Foundations of Signal Processing*
  http://goo.gl/r1mtcJ
- Scharf: *Statistical Signal Processing*
  http://goo.gl/NvwJLd

Outline

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

1. Signal representations in vector spaces
   (a) Introduction to discretizing signals using a basis: the Shannon-Nyquist sampling theorem
   (b) Linear vector spaces, linear independence, and basis expansions
   (c) Norms and inner products
   (d) Orthobases and the reproducing formula
   (e) Parseval's theorem and the general discretization principle
   (f) Important bases: Fourier, discrete cosine, lapped orthogonal, splines, wavelets
   (g) Signal approximation in an inner product space
   (h) Gram-Schmidt and the QR decomposition
2. Linear inverse problems
   (a) Introduction to linear inverse problems, examples
   (b) The singular value decomposition (SVD)
   (c) Least-squares solutions to inverse problems and the pseudo-inverse
   (d) Stable inversion and regularization
   (e) Weighted least-squares and linear estimation
   (f) Least-squares with linear constraints
3. Computing the solutions to least-squares problems
   (a) Cholesky and LU decompositions
   (b) Structured matrices: Toeplitz, diagonal+low rank, banded systems
   (c) Large-scale systems: Steepest descent
   (d) Large-scale systems: The conjugate gradient method
4. Low-rank updates for streaming solutions to least-squares problems
   (a) Recursive least-squares
   (b) The Kalman filter
   (c) Adaptive filtering using LMS
5. Matrix approximation using least-squares
   (a) Low-rank approximation of matrices using the SVD
   (b) Total least-squares
   (c) Principal components analysis
   (d) Signal and noise subspaces in array processing
6. Beyond least-squares (topics as time permits)
   (a) Approximation in non-Euclidean norms
   (b) Regularization using non-Euclidean norms
   (c) Recovering vectors from incomplete information (compressed sensing)
   (d) Recovering matrices from incomplete information (matrix completion)