Electrical & Computer ENGINEERING

Carnegie Mellon University Course Syllabus

18-698 / 42-632 Neural Signal Processing

Spring Semester, 2014

Course Personnel:

Instructor:

Byron Yu
Office: B204 Hamerschlag Hall
Email: byronyu@cmu.edu
Phone: (412) 268-9658
Office Hours: Mon 4-5pm, Tue 3-4pm

Teaching Assistants:

João Semedo
Email: jsemedo@andrew.cmu.edu
Office Hours: TBD

Course Management Assistant:

Jen Loughran
Hamerschlag Hall 1112
jackson1@andrew.cmu.edu
Office Hours: Monday-Friday: 8:30 a.m.-5:00p.m.

Course Description:

The brain is among the most complex systems ever studied. Underlying the brain's ability to process sensory information and drive motor actions is a network of roughly 10^11 neurons, each making 10^3 connections with other neurons. Modern statistical and machine learning tools are needed to interpret the plethora of neural data being collected, both for (1) furthering our understanding of how the brain works, and (2) designing biomedical devices that interface with the brain. This course will cover a range of statistical methods and their application to neural data analysis. The statistical topics include latent variable models, dynamical systems, point processes, dimensionality reduction, Bayesian inference, and spectral analysis. The neuroscience applications include neural decoding, firing rate estimation, neural system characterization, sensorimotor control, spike sorting, and field potential analysis.

Course goals:

There are two primary goals for the course: (1) to introduce the statistical tools used to study large-scale neural activity, and (2) to bring out the real-world challenges of working with experimental data. By the end of the course, students should be able to ask research-level questions in neural signal processing, as well as develop new statistical tools for problems in their own research. In short, this course serves as a stepping stone to research in neural signal processing.

Pre-requisites:
This course is ideally suited for students with a solid background in basic probability and linear algebra. Prior knowledge of neuroscience is welcome, but not required. Students with experience in neuroscience should be aware that the first 3 weeks will cover basic neuroscience.

Students should already be familiar with concepts such as:

Probability -- independence, conditional probability, Bayes rule, multivariate Gaussian distribution, Poisson distribution, Poisson process

Linear algebra -- basic matrix operations (sums and products), matrix inversion, eigenvectors and eigenvalues, singular value decomposition

For those unfamiliar with the concepts above, I would recommend *Probability Theory and Random Processes* (36-217).

If you are unsure whether this class is for you, please talk with me.

Graduate Course Area: Signal Processing and Communications

Class Schedule:
Lecture:
Tuesdays & Thursdays, 1:30 - 2:50p.m., Doherty 1212

Required Textbook:
Pattern Recognition and Machine Learning
Christopher Bishop. Springer, 2007

Optional textbooks:
Principles of Neural Science
Eric Kandel, James Schwartz, Thomas Jessell. McGraw-Hill Medical, 2000.

Theoretical Neuroscience
Peter Dayan and L.F. Abbott. MIT Press, 2001.

Information Theory, Inference, and Learning Algorithms
David J.C. MacKay. Cambridge University Press, 2003.

Matlab for Neuroscientists
Pascal Wallisch, Michael Lusignan, Marc Benayoun, Tanya I. Baker, Adam S. Dickey, and Nicholas G. Hatsopoulos. Academic Press, 2009.

Course Blackboard:
In order to access the course blackboard from an Andrew Machine, go to the login page at: http://www.cmu.edu/blackboard. You should check the course blackboard regularly for announcements.

Assignments and exams:
There will be approximately 9 problem sets during the semester and regular reading assignments. There will be a midterm exam in class on **Thursday, March 6** and a final exam during the week May 5-13, date TBD.

Most problem sets will have a Matlab component, in which students will implement various algorithms and apply them to neural data. This link has information about how to obtain Matlab software.

Students may discuss problem sets, but each student must turn in his/her own work. *You may not simply copy another student's work.* All students are bound by the CMU Academic Integrity Code.

Late policy for problem sets: Each student is allowed two late problem sets during the semester (up to 24 hours after the deadline). Problem sets that are turned in outside of this grace period will receive zero credit.

Grading breakdown:

Problem sets	30%
Midterm exam	30%
Final exam	40%

Course Outline:

1. What is neural signal processing?
   (1 lecture)
2. Neuroscience basics. Membrane potential. Action potential. Synaptic transmission.
   (5 lectures)
   PNS Ch 1, 2, 7
   Excerpts from PNS Ch 9, 10, 12
3. Spike train analysis. Spike histogram. Tuning curve. Poisson process.
   (4 lectures)
   TN Ch 1
4. Classification. Naive Bayes.
   Neuroscience application: discrete neural decoding
   (3 lectures)
   PRML Ch 4
5. Graphical models.
   (1 lecture)
   PRML Ch 8.1-8.2
6. Mixture models. Expectation-maximization.
   Neuroscience application: spike sorting
   (4 lectures)
   PRML Ch 9

7. Model selection. Cross-validation.
   Neuroscience applications: spike sorting, dimensionality reduction
   (2 lectures)
   PRML Ch 1.3, 3.4
8. Principal components analysis. Factor analysis.
   Neuroscience applications: spike sorting, dimensionality reduction
   (4 lectures)
   PRML Ch 12
9. Kalman filter.
   Neuroscience application: continuous neural decoding
   (4 lectures)
   PRML Ch 13

Education Objectives (Relationship of Course to Program Outcomes)

(a) an ability to apply knowledge of mathematics, science, and engineering:

In this course, students will learn to apply signal processing and machine learning methods to problems in neuroscience.

(b) an ability to design and conduct experiments, as well as to analyze and interpret data:

Students will analyze and interpret neural data (collected during neurophysiological experiments) in the problem sets.

(c) an ability to identify, formulate, and solve engineering problems:

Students will be asked to formulate models and design algorithms to be used in neural engineered systems.

Academic Integrity:
Students at Carnegie Mellon are engaged in preparation for professional activity of the highest standards. Each profession constrains its members with both ethical responsibilities and disciplinary limits. To assure the validity of the learning experience a university establishes clear standards for student work.
In any presentation, creative, artistic, or research, it is the ethical responsibility of each student to identify the conceptual sources of the work submitted. Failure to do so is dishonest and is the basis for a charge of cheating or plagiarism, which is subject to disciplinary action.
Cheating includes but is not necessarily limited to:
1. Plagiarism, explained below.
2. Submission of work that is not the student's own for papers, assignments or exams.
3. Submission or use of falsified data.
4. Theft of or unauthorized access to an exam.

5. Use of an alternate, stand-in or proxy during an examination.
6. Use of unauthorized material including textbooks, notes or computer programs in the preparation of an assignment or during an examination.
7. Supplying or communicating in any way unauthorized information to another student for the preparation of an assignment or during an examination.
8. Collaboration in the preparation of an assignment. Unless specifically permitted or required by the instructor, collaboration will usually be viewed by the university as cheating. Each student, therefore, is responsible for understanding the policies of the department offering any course as they refer to the amount of help and collaboration permitted in preparation of assignments.
9. Submission of the same work for credit in two courses without obtaining the permission of the instructors beforehand.

Plagiarism includes, but is not limited to, failure to indicate the source with quotation marks or footnotes where appropriate if any of the following are reproduced in the work submitted by a student:

1. A phrase, written or musical.
2. A graphic element.
3. A proof.
4. Specific language.
5. An idea derived from the work, published or unpublished, of another person.

In addition to the above, please also review fully and carefully Carnegie Mellon University's policies regarding *Undergraduate Academic Discipline* (http://www.cmu.edu/policies/documents/AcadRegs.html); and *Graduate Academic Discipline* (http://www.cmu.edu/policies/documents/GradDisc.html). In addition to the terms of the *Graduate Academic Discipline* policy, it is ECE's policy that an ECE graduate student may not drop a course in which a disciplinary action is assessed or pending without the course instructor's explicit approval. Further, an ECE course instructor may set his/her own course-specific academic integrity policies that do not conflict with Carnegie Mellon University's policies; course-specific policies should be made available to the students in writing.

This policy applies, in all respects, to 18-698/42-632.