# Department of Electrical and Computer Engineering
# University of Massachusetts - Amherst
# ECE 697SP - Statistical Image Processing
# Spring 2018 Syllabus

This information is available on the website

## INSTRUCTOR

Prof. Mario Parente, Knowles 113D, mparente@ecs.umass.edu
Office Hours: TBA or by appointment

## COURSE FORMAT

Lectures: 1:00 pm – 2:15 pm Tuesday and Thursday, Hasbrouck Laboratory room 242

## DESCRIPTION

This course, along with ECE 697SL Statistical Models for Learning, is designed to jointly provide an introduction to supervised and unsupervised machine learning with a focus on their application to statistical signal, image, and data processing. This course will mainly focus on topics in supervised learning - in short, topics in machine learning that involve prediction based on a labeled set of examples. The course content will span mathematical analysis, implementation, and example cases with real-world data, with some focus on image data.

## PREREQUISITES

This class assumes knowledge of signals and systems, linear algebra, and probability equivalent to that obtained in ECE 313, ECE 603, and Math 235. The course is only open to graduate students. If you are unsure whether you meet the requirements, please contact the instructor.

## TEXTBOOKS AND RESOURCES

There will be three textbooks that will be followed during portions of the course. The first one is available for purchase from the campus bookstore and online at the link given below:

- C. Bishop, "Machine Learning," Springer, 2006.

We will also follow the following two textbooks, available for purchase and as downloads from their authors' websites linked below. The second book is also available at the library as a course reserve.

- G. James, D. Witten, T. Hastie, R. Tibshirani, "An Introduction to Statistical Learning with Applications in R," Springer, 2013.
- T. Hastie, R. Tibshirani, and J. Friedman "The Elements of Statistical Learning: Data Mining, Inference, and Prediction," Springer, 2009.

Lecture notes for some of these topics are available at Andrew Ng's Machine Learning course website at Stanford University. We will also share slides from lectures in the course Moodle website. Note, however, that <u>none of these resources are meant to be replacements for lecture attendance</u>.

## GRADING

- Homework assignments : 30% of Grade
- Midterm: 30% of Grade
- Final Project: 30% of Grade
- Class Participation: 10% of Grade

There will be four to five *homework assignments*, due roughly every other week, that will be posted in the Moodle course website. Assignments will involve both analytical work and Python or Matlab coding and must be turned in at the beginning of the class period on the due date. One late homework (due the lecture following the original due date) is allowed for each student, and no further late submissions are accepted. It is encouraged to discuss the problem sets with others, but each student must turn in **their own unique** write-up or code implementation.

A *midterm exam* will take place (tentative date: March 29 in the evening) covering the content from the first half of the course. The exam will be open book/computer and open notes.

The course will also include an *original research project* that involves a project report (4 pages in IEEE conference proceeding style) and a 15-minute oral presentation in class. The goal of a project is to obtain original or replicated research results in machine learning involving the topics contained in the class that are of sufficient caliber for a conference publication. Project topics will be selected in consultation with the instructor during the second half of the semester. Suggested projects will be posted in the course Moodle website later in the semester. Students are also encouraged to discuss their project interests with the instructor and/or their advisor. Teams are allowed with individual grading, in which case a statement of contributions must be submitted with the report.

The instructor will assign a *class participation* grade to students based on their interaction and level of involvement with the instructor and the students during lectures, office hours, etc.

## ACCOMMODATION POLICY

The University of Massachusetts-Amherst is committed to providing an equal educational opportunity for all students. If you have a documented physical, psychological or learning disability on file with Disability Services (DS), Learning Disability Support Services (LDSS) or Psychological Disabilities Services (PDS), you may be eligible for reasonable academic accommodations to help you succeed in this course. If you have a documented disability that requires an accommodation, please notify me within the first two weeks of the semester so that we may make appropriate arrangements.

## ACADEMIC HONESTY POLICY

It is expected that all graduate students will abide by the UMass Academic Honesty Policy (available at the Graduate Dean's Office, the Ombuds Office, or online at https://www.umass.edu/honesty/). Acts of academic dishonesty will result in a grade of F in the course, and possibly additional sanctions including loss of funding, being placed on probation or suspension for a period of time, or being dismissed from the University. All students have the right of appeal through the Academic Honesty Board.

**LECTURE SCHEDULE (TENTATIVE)**

Each topic is annotated with corresponding book chapters for the textbooks listed earlier.

| Week | Topic | Bishop | JWHT | HTF |
|---|---|---|---|---|
| | Basic concepts | | | |
| 1 | Background: Vector Spaces and Linear Algebra. Multivariate Probability | Section 1 of CNX Notes for ECE 697CS, linear algebra and probability notes | | |
| 2 | Supervised learning setup. LMS. | Ch. 3.1.1 | | Ch.2, Ch.3 |
| 3 | Discriminative algorithms: Logistic Regression, Perceptron, Exponential Family | Ch.4.1, Ch. 4.2, Ch.4.3 | | Ch.4, Ch. 6.6.3 |
| 4 | Generative learning algorithms. Gaussian discriminant analysis. Naive Bayes. | Ch.4.1, Ch.8.2 | | Ch.4 |
| 5 | Support vector machines | Ch.7 | | Ch.12 |
| 6 | 1. Bias/variance tradeoff Model selection and feature selection | Ch. 3.2 | | Ch.7 |
| 7 | Evaluating and debugging learning algorithms, practical advice on structuring an ML project | | notes | |
| 8 | Classification and regression trees, boosting, bagging, random forests | | | Ch.10 Ch.15 |
| 9 | Neural Networks, Forward/Back propagation | Ch.5 | | Ch.11 |
| 10 | Vectorization and Other optimization tricks for NN | | notes | |
| 11 | Deep Learning Methods | | DL | |
| 12 | Deep Learning Platform | | DL | |
| 13 | Generative Adversarial Networks (GANs) | | paper | |