Fall 2014
ECE 539: STATISTICAL LEARNING AND OPTIMIZATION

Logistics:

Instructor:	Prof. Anand D. Sarwate Office: CoRE 517 Email: `asarwate@ece.rutgers.edu` Web: `http://www.ece.rutgers.edu/~asarwate/`
Office hours:	TBA
Lectures:	Monday and Thursday 1:40 AM - 3:00 PM in BME 126
Website:	`http://sakai.rutgers.edu/`

Topics: This course will cover the foundations of statistical learning from a theoretical perspective, as well as the computational approaches. Topics will include:

- parametric estimation, nonparametric estimation, and statistical decision theory
- classification, regression, and other examples of statistical learning problems
- the empirical risk minimization framework
- computational issues arising in optimization
- minimax theory and lower bounds
- applications and modifications to account for privacy, missing data, or other issues (time permitting)

More on prerequisites: Students should have a strong background in probability beyond the undergraduate level. The course will make extensive use of concepts such as density functions, expectation, conditional probabilities, independence, and Markov's inequality. Advanced topics such as martingales and measure concentration will be introduced as needed. Students are *not* expected to have taken a previous course in machine learning, but some experience with machine learning concepts and algorithms may make the material easier. Students will be expected to have the mathematical maturity to read supplementary material, including the proofs. If you have any concerns about whether you should take the course, please contact the instructor before classes begin.

Additional references: Students may find some the following as useful supplements to the material in the course. Additional references will be added throughout the course.

- Vladimir N. Vapnik, *The Nature of Statistical Learning Theory*, Springer, 2000.

- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar, *Foundations of Machine Learning*, MIT Press, 2012.
- Luc Devroye, László Györfi, and Gábor Lugosi, *A Probabilistic Theory of Pattern Recognition*, Springer, 1996.
- Felipe Cucker and Ding Xuan Zhou, *Learning Theory: An Approximation Theory Viewpoint*, Cambridge, 2007.
- Alexandre B. Tsybakov, *Introduction to Nonparametric Estimation*, Springer, 2010.

Topics and Schedule: This is subject to change. Days that are marked as cancelled may be replaced by guest lectures on related topics or video lectures.

Lecture	Topics
9/4	Introduction to the course + some history
9/8	What do we mean by learning?
9/11	Statistical notions of learning and parametric estimation
9/15	Nonparametric estimation
9/16	Statistical decision theory
9/18	Computational learning theory
9/22	Probably Approximately Correct (PAC) learning
9/25	Finishing up the theory of learning
9/29	Empirical risk minimization (ERM)
10/2	CANCELLED
10/6	Vapnik-Chervonenkis (VC) dimension
10/9	Analyzing ERM
10/13	Examples of ERM: classification
10/16	Examples of ERM: regression
10/20	Introduction to minimax theory
10/23	Minimax bounds for estimation
10/27	Minimax bounds for learning
10/30	Segue: computation concerns
11/3	Computational approaches to ERM
11/6	Convex optimization for ERM
11/10	Speeding up the optimization
11/13	CANCELLED
11/17	Stochastic optimization
11/20	Stochastic gradient descent
11/24	TBD
11/25	TBD (Tuesday with a Thursday schedule?)
11/27	NO CLASS: Thanksgiving
12/1	TBD
12/4	Project Presentations
12/8	Project Presentations

Grading and assessment:

Homeworks	20% (4 assignments)
Short Quizzes (3)	15% (5% each)
Scribing lectures:	10%
Project proposal/meetings	5%
Project report draft	10%
Final project presentation	20%
Final project report	20%

Scribing Lectures: Every student will be expected to scribe either two lectures (working in pairs) or one lecture (working alone) using a LaTeX template provided by the instructor. The writeup of the lecture should be written in full sentences, fill in details and ideas that were only mentioned in lecture, and provide references that may have been missing. This is not simply "typing up your notes." but preparing a readable document that can be used by other students to remind and enhance their own knowledge of the material. Signups for scribing will happen during the first week of class.

Scribe notes for the lecture are due **2 weeks** after the lecture. Students who want feedback on the draft of their notes may submit them within a week of the lecture. So for example

Homework: Homeworks will be handed out and due according to the following approximate schedule:

Homework	Assigned	Due
Homework 0	9/4	never
Homework 1	9/8	9/25
Homework 2	9/25	10/13
Homework 3	10/13	11/3
Homework 4	11/4	11/24

Homework is due at the beginning of class on the due date. There are no extensions or late submissions without the prior consent of the instructor. Such consent will only be given in exceptional circumstances. Each student must prepare their own original solutions to the homework.

- *Collaboration policy:* I expect students to independently solve the problems on the homework and write up the solutions individually. Discussing with other students is permissible, but these are not meant to be "group assignments." Treating the homework as such is cheating yourself. Copying the solution from another student is a form of *academic misconduct* (see below) and there may be severe penalties. If you are concerned about whether your collaboration is appropriate, please contact the instructor. If you have difficulties with the homework, arrange to come for office hours.

- *Some tips:* Since this is a more mathematical course, doing the homework is very important! Developing proficiency in applying the concepts will give you a depth of understanding that will help you understand the papers for your project. If you have problems, come to office hours, and if you miss a question, make sure you understand the solution when it is posted.

Final Project: More than half of the grade in the class will be based on a final project. More details on projects will be provided in class; this is a brief summary. The projects are individual and students should explore some topic related to the class through reading research papers. Example projects could include comparing algorithms to solve a particular statistical learning problem, performing a literature review on a class of statistical optimization methods, or testing statistical learning methods in a particular application domain. Students are particularly encouraged to find ways in which the ideas from the class may intersect with their own research.

Evaluation of the final projects will be based on four components:

- **Project proposal:** a 1-2 page description of what the goal of the project will be. In particular, students should select the set of papers on which they will base their project and clearly delineate the scope of the work. The proposal will be due in early October.
- **Project report draft:** this is a mid-project review to evaluate student's progress towards the project goals. The draft will be due in November.
- **Project presentation:** Students will give a short presentation to the class on their project, describing the problem clearly and what the outcomes were. Presentations will be during the last two classes of the term.
- **Project final report:** A final report describing the project – its goals, methods, and outcomes. Reports will be due the last day of instruction.

General policies

Academic Integrity. Students should be familiar with the Academic Integrity policy at Rutgers University:

`http://academicintegrity.rutgers.edu`

If you have any questions about whether your actions may compromise academic integrity in some way, please contact the instructor as soon as possible.

Misrepresenting your work or contributions hurts not only yourself, but others. As the website states, "every member of the University community therefore bears a responsibility for ensuring that the highest standards of academic integrity are upheld." Among other things, this means *copying homework solutions* and *letting others copy your solutions* are both serious offenses. Violations of the academic integrity code may result in disciplinary proceedings.

In your project reports, if you are found to have plagiarized any material from other sources it will be grounds for receiving a failing grade on the project, in the class overall, and you may be referred for disciplinary action to the Graduate Division. **Do not plagiarize.** If you have any questions are concerns about your writing, ask the instructor. Ignorance of the rules ("but I didn't think this was plagiarizing") will not be grounds for leniency. It is *your responsibility* to read and understand proper academic conduct.

Responsibilities. As a student, it is your responsibility to manage your schedule such that you can come to lectures prepared and on-time. Homework is due at the beginning of class on the due date indicated – late homework will not be accepted.

Incomplete grades and dropping the course. Incomplete grades will not be given to students who wish to improve their grade by taking the course in a subsequent semester. An incomplete grade may be given for medical reasons if a doctor's note is provided. The purpose of an incomplete grade is to allow a student *who has essentially completed the course* and who has a legitimate interruption in the course, to complete the remaining material in another semester. Students will not be given an opportunity to improve their grade by doing "extra work".

Students are responsible for being aware of the drop dates for the current semester. Drop forms will not be back-dated.