CSCI 567 Summer 2018 Syllabus

Course meetings: You are responsible for everything covered in lectures and your enrolled discussion, including administrative announcements.

Strongly Recommended Textbooks: Tom E. Mitchell, *Machine Learning*
Pang-Ning Tan, Michael Steinbach, Anuj Karpatne, Vipin Kumar, *Introduction to Data Mining, Second Edition*
Also Recommended Textbooks: A "Recommended Textbooks" list is available at the end of this syllabus. These cover both mathematical preliminaries and machine learning topics.
Gradebook: blackboard.usc.edu
Forums:`piazza.com/usc/summer2018/csci567` -- your instructor will sign you up in the first week.

Blackboard will be used for posting of grades. Messages that do not need a particular instructor's direct attention should be posted to Piazza with the appropriate privacy setting. If you email your instructor, you *must* include the substring "CSCI 567" to begin a *meaningful* subject line and have tried to resolve the issue appropriately otherwise (e.g., questions about course material should be posted to Piazza first, using email only after an appropriate amount of time has passed without a response.) Such emails must be sent from your USC email account. Any emails not conforming to this might not be seen by your instructor and will not be considered to have been sent for administrative purposes.

Course Description
Statistical methods for building intelligent and adaptive systems that improve performance from experiences; Focus on theoretical understanding of these methods and their computational implications.
Recommended preparation: Undergraduate level training or coursework in linear algebra, multivariate calculus, basic probability and statistics; an undergraduate level course in Artificial Intelligence may be helpful but is not required.

Grade Calculations:

Artifact	Weight	Date	Time
Midterm Exam	40%	TBD	9:00 (arrival deadline)
Final Exam	60%	Tuesday, July 24	9:00 (arrival deadline)

Due to the condensed summer schedule, problem sets and programming assignments will not be collected for credit. You may work on these on your own schedule. TA(s) and other course staff will be available to discuss these with you.
When determining letter grades at the end of summer, I will recognize that homework / programming is usually ~40% of the grade and will determine letter grades accordingly.

Exams

You will be provided with paper on which to take the exam. Exams will be individual effort, closed-book and closed-notes. You will not be permitted any electronics on the exams, including digital watches. Any requirements for you to bring to the exam will be announced well in advance.

Exams will likely include a multiple-choice component, to be scored by Scantron. As such, you are required to have at least one #2 pencil with you during the exam. These will be graded *only* based on what is scanned by the Scantron machine, although you may choose to indicate on your packet what your response was for when the exam is returned. Free-response portions may be answered in pen or pencil, student's choice.

Students requiring alternate exam arrangements must make such requests within the first two weeks of the term, or as soon as possible after knowing of the conflict or requirement. Students who will be taking exams remotely must check the testing center availability within the first two weeks of the semester and email the instructor *within that time period* if the testing center cannot accommodate you on our test day at our test times (9:00am - 10:30am, Pacific). Deviations from the exam time schedule must be requested by June 1 or as soon as the conflict is known by the student. If the conflict could have been known by the student during the first two weeks, and is brought to instructor attention after that time period, the request may be declined for this reason. Deviations from the exam schedule not consistent with this paragraph will be declined and exams so taken will be given a zero.

Academic Conduct

Please keep in mind that CSCI 567 is a large course with staff that work very hard to run it. Abusive behavior towards course staff will not be tolerated and will result in a referral to the relevant office in cases that warrant it. We strive to run the class well, but recognize that we cannot do things the same way we could if there were 20 students total.

In addition, all students are expected to have read and understood the university's academic integrity guidelines for graduate students, available at https://sjacs.usc.edu/files/2015/03/GradIntegrity.pdf. Related items appearing in this syllabus are a supplement, not a replacement, for those policies.

The university's recommended sanction for plagiarism and similar offenses is an F in the course on first offense[1]. Every suspected incident of academic dishonesty will be reported to the relevant office and no suggestion for lighter punishment will be made and requests for me to do such will not be entertained. The time to consider the consequence of the potential fail in the course is **before** committing an infraction. In any course with a pending academic honesty violation report, you may not drop the course; if you drop the course and are later discovered to have violated the academic honesty policy, you will be re-enrolled.

[1] See https://studentaffairs.usc.edu/files/2015/09/appendix_a.pdf for recommended sanctions for *undergraduates*, and observe that graduate students may be sanctioned more severely.

CSCI 567 Projected Course Schedule: A Weekly Breakdown

Please note that this is a *projected* schedule and is subject to change. Any changes will be announced via the course Piazza page or in lecture. *May 28 (Mon), July 4 (Wed) are holidays*

Textbooks referenced:
[GT] refers to *Algorithm Design* by Michael Goodrich and Roberto Tamassia; this was likely a book used in your algorithms classes (e.g., CSCI 570).
Other references are either the top-listed books ([Mitchell] and [TSKK]) or are listed in recommended books below.
I will update this schedule with additional suggested readings later.

Date	Topic	Related Reading	Notes
May 16 (Wed)	Course Overview ML overview Nearest Neighbor classification		
May 21 (Mon)	Core ML concepts; typical steps to developing a ML system		Lecture 5/16 at noon;
	Fundamentals: Linear Algebra	[LADW] Ch. 1-6	
	Fundamentals: Linear Algebra	[LADW] Ch. 1-6	
May 28	**No class: Memorial Day holiday**		
	Fundamentals: Probability and Statistics		
	Linear Regression	[MLaPP] 1.4.5, 7.1-7.3, 7.5.1, 7.5.2, 7.5.4, 7.6	
June 4	Regression with nonlinear basis Regularized regression Maximum Likelihood Estimation	[MLaPP] 1.4.7, 1.4.8 [ESL] 7.1, 7.2, 7.3, 7.10	
	Concept Learning	[Mitchell] Chapter 2	
	Decision Trees	[Mitchell] Chapter 3	
June 11	Boosting/ensemble	[MLaPP] 16.4.1-16.4.5, 16.4.8, 16.4.9 [ESL] 16.3	
	Generative Models, Naive Bayes	[Mitchell] Chapter 6 [MLaPP] 3.5 [ESL] 6.6.3	

	Review / Catch-Up		
Monday June 18: Midterm Exam: 9:00 - 10:30am; Room TBD			
	Hidden Markov Models	[MLaPP] 17.1-17.4, 17.5.1-17.5.2	
	Linear discriminant analysis Perceptron	[MLaPP] 4.2.1 - 4.2.5, 8.5.1-8.5.4	
June 25	Logistic regression	[MLaPP] 1.4.6, 8.1-8.3 [ESL] 4.1-4.2, 4.4	
	softmax; multi-way classification		
	neural networks / mlp	[Mitchell] Chapter 4 [MLaPP] 16.5.1-16.5.6, 28 [ESL] 11.3-11.7	
July 2	DNN, CNN, RNN, and LSTM		
	Kernel Methods	[MLaPP] 14.1, 14.2.1-14.2.4, 14.4.1, 14.4.3 [ESL] 5.8, 6.3, 6.7	
	No class: July 4 holiday		
July 9	Review: Linear Programming	[GT] Chapter 26	
	SVM	[MLaPP] 14.5.2-14.5.4 [ESL] 12.1-12.3	
	clustering, mixture models	[MLaPP] 11.1-11.3, 11.4.1-11.4.4, 11.5 [ESL] 14.3.1-14.3.9, 8.5	
July 16	mixture models / density estimation		
	“Advanced Clustering”	To be posted	
	Basic learning theory	[Mitchell] Chapter 7	
July 23	Review / Catch up / Enrichment		
Tuesday July 24: Final Exam, 9:00am - 10:30am; Room TBD			

Recommended Textbooks

Abbreviations in brackets is used to reference in the projected course schedule above.

Math Preparation:

- Calculus: Prof. Strang's free online textbook. https://ocw.mit.edu/resources/res-18-001-calculus-online-textbook-spring-2005/
- Linear Algebra: Also from Prof. Strang: https://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/index.htm
- [LADW] Linear Algebra Done Wrong by Sergei Treil: https://www.math.brown.edu/~treil/papers/LADW/LADW.html
- Probability and Statistics: Orloff and Bloom: https://ocw.mit.edu/courses/mathematics/18-05-introduction-to-probability-and-statistics-spring-2014/
- The Matrix Cookbook: http://www.math.uwaterloo.ca/~hwolkowi//matrixcookbook.pdf

Additional Machine Learning Resources:

- Max Welling's quick intro: https://www.ics.uci.edu/~welling/teaching/ICS273Afall11/IntroMLBook.pdf
- Alex Smolas's book: http://alex.smola.org/drafts/thebook.pdf
- Gareth James et al's *An Introduction to Statistical Learning* : http://www-bcf.usc.edu/~gareth/ISL/ISLR%20First%20Printing.pdf
- A Course in Machine Learning by Hal Daumé III: http://ciml.info
- *Bayesian Reasoning and Machine Learning* by David Barber: http://web4.cs.ucl.ac.uk/staff/D.Barber/pmwiki/pmwiki.php?n=Brml.HomePage
- [Bishop] *Pattern Recognition and Machine Learning* by C Bishop
- Andrew Moore's tutorials: https://www.autonlab.org/tutorials
- *Pattern Classification* by Duda, Hart and Stork
- *All of Statistics* by L. Wasserman
- [ESL] Trevor Hastie, Robert Tibshirani, and Jerome Friedman, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction. Second Edition*
- [MLaPP] Kevin Murphy, *Machine Learning: A Probabilistic Perspective*