Note This is a redesigned course from the previous offerings of CSCI567. While the topic coverage in this course is similar to the previous ones, there will be differences in focus and pedagogy.

Introduction The chief objective of this course is to teach methods in pattern classification and machine learning. Key components include statistical learning approaches, including but not limited to various parametric and nonparametric methods for supervised and unsupervised learning problems. Particular focuses on the conceptual understanding of these methods, their applications and hands-on experience.

Preparation Required: (1) undergraduate level training or coursework in linear algebra, calculus and multivariate calculus, basic probability and statistics; (2) Skills in programming with Python 3.0 (self-studying `scikit-learn` and related packages is expected.) are required. (3) Basic skills in using `git` for maintaining code development are required. An undergraduate level course in Artificial Intelligence may be helpful but is not required.

Teaching assistants Soravit Changpinyo, Weilun Chao, Guangyu Li, Zhiyun Lu, Yiming Yang, and Ke Zhang.

Homework typesetting For some homework assignments, using LaTeX(various TeXeditors and compiling environment on Windows, Mac OS X and Unix/Linux are available, including WYSIWYG ones) to typeset is required.

Format classroom lectures, discussions, homework, in-class three quizzes.

Outside-classroom learning environments We will use Blackboard (blackboard.usc.edu) to post grades. We will use Piazza forums for online discussions (link address TBA). We will use Github repositories for homework assignments and other matters. Messages that do not need a particular instructors attention should be posted to Piazza with the appropriate privacy setting.

Preparation If you would like to prepare or refresh your skills in relevant maths, the followings would be good starting points (there are plenty of equally good online resources – the following is just a sample)

- For calculus, please check Prof. Strang's free online textbook

  http://ocw.mit.edu/resources/res-18-001-calculus-online-textbook-spring-2005/

- For linear algebra, please check (again) Prof. Strang's OpenCourseWare site

  http://ocw.mit.edu/courses/mathematics/18-06-linear-algebra-spring-2010/index.htm

- Probability and statistics, please check

  MIT course

Grading 5 homework assignments (40%, roughly 8% each), 3 quizzes (total: 60%, roughly 20% each)

Grading Reconsideration Reconsideration requests for any graded artifact must be made within one week of our release of grades for the item. A form will be available on the course Piazza page for you to use to explain your reconsideration request. You must fill out the form within that week. Once the reconsideration period has passed, grades for each artifact are considered final.

For all the quizzes, grading reconsideration will require filling out the same form and handing your graded quizzes to the instructor only within that one week period.

Requests for grading reconsideration submitted in any other fashion will not be considered.

Academic standard You are responsible for everything covered in lectures and discussions, including administrative announcements. Attending the discussion is required. The discussion provides more detailed and in-depth exposition of the lectured material.

Policy on homework assignments

- Extension and late turn-in: You have a total of 5 “grace days”. In order to use a grace day, you must submit a form (which will be provided) within 24 hours of the due time of the assignment. If you use up “grace days”, then your submission will be considered as “being late” and this will cost you 25%, 50% and 100% of the maximum points of the assignments (within 24 hours, 48 hours, or 72 hours of the due times). **Note** that some assignments will have a limit to the number of grace days that may be used for that assignment.

- Working in group: permitted but each member needs to write up and submit solutions separately. Standards on academic integrity are strictly enforced. **Note** that please see the academic honesty supplement later in the syllabus.

- Turn-in: You will need to get a Github repo for turning homework assignments for the class and other material. If you do not have one, please obtain a Github handle before 8/25. The homework submissions by any means other than the Github repo are **not** accepted. In particular, either email or physical submissions are **not** accepted.

Note: There is no grace period. Even if you submit a few minutes after the deadline, you will need to use a grace day (even if the wireless network in your dorm room is down or you have a github issue, etc.). It is your job to be on time and not cut it too close. Remember Murphy’s Law and leave time for things to go wrong.

Required textbooks There will be no required textbooks. However, we suggest one of the following to help you to study:

- Kevin Murphy's *Machine Learning: A Probabilistic Perspective*
- *Elements of Statistical Learning* by Hastie, Tibshirani and Friedman http://web.stanford.edu/~hastie/ElemStatLearn/ We will mark suggested readings from these two books.

Other useful reading materials are:

- Max Welling's quick intro: https://www.ics.uci.edu/~welling/teaching/ICS273Afall11/IntroMLBook.pdf
- Alex Smolas book http://alex.smola.org/drafts/thebook.pdf
- Gareth James book 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
- *Pattern Recognition and Machine Learning* by C Bishop (available from online and campus bookstores)
- Andrew Moore's Tutorial http://www.autonlab.org/tutorials/
- *Pattern Classification* by Duda, Hart and Stork
- *All of Statistics* by L. Wasserman

References for frequently used maths

- *The Matrix cookbook* http://orion.uwaterloo.ca/~hwolkowi/matrixcookbook.pdf
- The Wisconsin collection http://pages.cs.wisc.edu/~andrzeje/lmml.html
- Khan Academy: http://www.khanacademy.org/

Tentative Schedule Please see the last page of this document.

Statement for Students with Disabilities Any student requesting academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me (or to TA) as early in the semester as possible. DSP is located in STU 301 and is open 8:30 a.m.5:00 p.m., Monday through Friday. The phone number for DSP is (213) 740-0776.

Statement on Academic Integrity USC seeks to maintain an optimal learning environment. General principles of academic honesty include the concept of respect for the intellectual property of others, the expectation that individual work will be submitted unless otherwise allowed by an instructor, and the obligations both to protect ones own academic work from misuse by others as well as to avoid using anothers work as ones own. All students are expected to understand and abide by these principles. Scampus, the Student Guidebook, contains the Student Conduct Code in Section 11.00, while the recommended sanctions are located in Appendix A: http://www.usc.edu/dept/publications/SCAMPUS/gov/. Students will be referred to the Office of Student Judicial Affairs and Community Standards for further review, should there be any suspicion of academic dishonesty. The Review process can be found at: http://www.usc.edu/student-affairs/SJACS/.

Academic Honesty Supplement The following are a supplement for this course to the university academic honesty guidelines.

Any code obtained with help of a course producer, TA, or instructor must be clearly marked in comments with who and when. This may not constitute a meaningful portion of your project. The substring “assistance from” (without the quotation marks) must appear in the comment acknowledgement near the code on which you received assistance, followed by the name of who you got assistance from.

The help you receive from classmates should be limited to conceptual help – how an algorithm you need to implement works as a big picture rather than at the level of code.

If you receive help from others, whether an instructor, course producer, or fellow student, follow the Kenny Loggins Rule: You may discuss high-level ideas and receive hints regarding how to solve portions of the assignments. However, neither party should keep any written record from this discussion. Afterwards, take a 30-minute break and do something unrelated to the course (watching an episode of your favorite cartoon show, for example). You may now return to your assignment. When you write a section of code based on help received, add a comment acknowledging the help, including the substring assistance from as part.

You are explicitly prohibited from seeking help outside of course resources for the programming projects. The following is an exhaustive list of course resources: The instructor, TA(s), and course producers. Your fellow students. Remember that this means discussing concepts, not sharing code. The public portion of the mentioned references (textbooks, tutorials, manuscripts)' websites. The course Piazza page for this semester. Course lectures and discussions, along with any notes provided by instructors.

The previous bullet point means that if you seek information towards solving part of a programming assignment online, find it, and use it, you risk this as an academic honesty violation. If you use code fragments from one of the textbooks mentioned above, you must include a comment acknowledging the source, including the substring assistance from and the textbook citation (textbook name, authors, and page number is sufficient). If you are using git, use your private repository so your code is not available to your fellow students. Course staff will help set you up with a private repository early in the semester.

Reading Material Please note the following suggested readings

- MLaPP refers to Murphy's textbook
- ESL refers to Elements of Statistical Learning

Schedule

- Schedule is subject to finer adjustment, depending on the progress of the class.
- **Quiz dates are finalized and will not be changed**.
- **Discussion topics** in the table in most cases cover the main lectures on Tuesdays and Thursdays. For example, The discussions for the lectures on 8/22 and 8/24 occur on the Friday 8/25 and the Monday 8/28 and Tuesday 8/29.
- **Homework relase and due dates**. Note that "(D)" means due and "R" means "Release". Due dates are always on the immediate Sundays 11:59pm following the Thursday's lecture dates. For example HW1 is due on 9/24 11:59pm. Release dates are on the immediate Friday 8am following the Thursday's lecture dates. For example HW1 is to be released on 9/8 8am.

Date	Topics	Discussion	Notes
8/22	Course overview; ML overview; Nearest neighbor classification	Install VM, Git primer	
8/24	Core ML concepts; typical steps of developing a ML system		
8/29	Linear regression [MLaPP] 1.4.5, 7.1-7.3, 7.5.1, 7.5.2, 7.5.4, 7.6	Linear algebra, Nearest neighbor, Linear Regression (No discussion on 9/4)	
8/31	Regression with nonlinear basis; regularized regression [MLaPP] 1.4.7, 1.4.8 [ESL] 7.1, 7.2, 7.3, 7.10		
9/5	Linear discriminant analysis, Perceptron [MLaPP] 4.2.1-4.2.5, 8.5.1-8.5.4	Perceptron, Numerical optimization	
9/7	Logistic regression [MLaPP] 1.4.6, 8.1-8.3 [ESL] 4.1-4.2, 4.4		HW1(R)
9/12	Softmax; multi-way classification	Classification, Neural networks, Backpropagation	
9/14	Neural networks/MLP [MLaPP] 16.5.1-16.5.6, 28 [ESL] 11.3-11.7		
9/19	DNN, CNN, RNN and LSTM	Deep learning architectures, Kernels	
9/21	Kernel methods [MLaPP] 14.1, 14.2.1-14.2.4, 14.4.1, 14.4.3 [ESL] 5.8, 6.3, 6.7		HW1(D), HW2(R)
9/26	SVM [MLaPP] 14.5.2-14.5.4 [ESL] 12.1-12.3	SVM, Convex duality, Entropy and Mutual Information, Tree	
9/28	Decision trees		
10/3	Boosting/ensemble [MLaPP] 16.4.1-16.4.5, 16.4.8, 16.4.9 [ESL] 16.3	Boosted trees, Review of Quiz 1	
10/5	Quiz 1 on Classification and Regression		HW2(D), HW3(R)

Date	Topics to be covered	Discussion	Notes
10/10	Basic Learning Theory	Probability, K-means	
10/12	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		
10/17	mixture models/ density estimation	EM algorithm, NB for document classification, Gaussian NB	
10/19	Generative models: naive Bayes [MLaPP] 3.5 [ESL] 6.6.3		HW3(D), HW4(R)
10/24	Hidden Markov models (HMMs) [MLaPP] 17.1-17.4, 17.5.1-17.5.2	Details of Baum-Welsh (Forward-Backward algorithms, Viterbi algorithms)	
10/26	Float day for review/Q&A		
10/31	Quiz 2	Review of Quiz 2, Variational inference, text modeling	
11/2	Hierarchical models; topic models [MLaPP] 10.1, 10.2.1-10.2.3, 10.3-10.5		HW4(D), HW5(R)
11/7	Dimensionality reduction and visualization [MLaPP] 12.2 [ESL] 14.5.1	Dimensionality reduction, autoencoder, manifold learning, collaborative filtering	
11/9	Recommender systems and other applications		
11/14	Large-scale ML systems	Markov decision process	
11/16	Intro to reinforcement learning		HW5(D)
11/21	Review/Tie-in topic	No discussion on 11/24	
11/23	no class (Thanksgiving)		
11/28	Quiz 3	Review Quiz (Friday discussion sections)	
11/30	Emergent Topics		