USC Viterbi School of Engineering

CSCI 467: Introduction to Machine Learning (Fall 2019)

Units: 4

Instructor: Mohammad Reza Rajati, PhD
PHE 412
rajati@usc.edu – Include CSCI 467 in subject

Office Hours: Wednesday 3:30 pm –5:00 pm

TA(s): Soumyaroop Nandi
soumyarn@usc.edu – Include CSCI 467 in subject

Office Hours: TBD

Grader: Roger Lin
ruizhel@usc.edu – Include CSCI 467 in subject

Lecture: Monday, Wednesday, 2:00 pm –3:20 pm in MHP B7B

Discussion 1: Monday 4:00 pm –5:50 pm in TBD

Discussion 2: Tuesday 4:00 pm –5:50 pm in TBD

Webpages: Piazza Class Page for everything except grades
and USC Blackboard Class Page for grades
and GitHub for code submission

– All HWs, handouts, solutions will be posted in PDF format

Prerequisite: (CSCI 270 and MATH 225) and 1 from (EE 364 or MATH 407 or BUAD 310).

Other Requirements: Computer programming skills.
Using Python is mandatory.
Students must know Python or must be willing to learn it.

Tentative Grading: Programming Assignments 35%
Problem Sets 15%
Midterm Exam 20%
Final Exam 30%
Participation on Piazza* 5%

Letter Grade Distribution:

$\geq$ 93.00	A	73.00 - 76.99	C
90.00 - 92.99	A-	70.00 - 72.99	C-
87.00 - 89.99	B+	67.00 - 69.99	D+
83.00 - 86.99	B	63.00 - 66.99	D
80.00 - 82.99	B-	60.00 - 62.99	D-
77.00 - 79.99	C+	$\leq$ 59.99	F

Disclaimer: Although the instructor does not expect this syllabus to drastically change, he reserves every right to change this syllabus any time in the semester.

Note on e-mail vs. Piazza: If you have a question about the material or logistics of the class and wish to ask it electronically, please post it on the piazza page (not e-mail). You may post it anonymously if you wish. Often times, if one student has a question/comment, other also have a similar question/comment. Use e-mail with the professor, TA, graders only for issues that are specific to you individually (e.g., a scheduling issue or grade issue).

Catalogue Description: Methods for building intelligent and adaptive systems from statistical analyses; theoretical understanding of such methods and the computational implications. .

Course Description: This is an introductory undergraduate course on Machine Learning with a focus on applications. The primary approach of instruction in this course is *Learing by Doing*. The focus of the course is to provide the students with basic understanding of Machine Learning algorithms and to make them use the algorithms to analyze data and convert them into information for decision-making.

Course Objectives: Upon successful completion of this course a student will

- Broadly understand major algorithms used in machine learning.
- Understand supervised and unsupervised learning techniques.
- Understand regression methods.
- Understand resampling methods, including cross-validation and bootstrap.
- Understand decision trees, dimensionality reduction, regularization, clustering, and kernel methods.
- Understand feedforward neural networks and deep learning.

Exam Dates:

- **Midterm Exam:** Wednesday October 23, 2:00-3:20 PM.
- **Final Exam:** Friday, Dec 13, 2:00 PM- 4:00 PM as set by the university

Textbooks:

- **Required Textbook:**

  1. Gareth James, Daniela Witten, Trevor Hastie, and Robert Tibshirani, *An Introduction to Statistical Learning with Applications in R*, Springer, 2013. (ISLR)
     Available at `http://faculty.marshall.usc.edu/gareth-james/ISL/ISLR%20Seventh%20Printing.pdf`

- **Recommended Textbooks:**

1. *Applied Predictive Modeling*, 1st Edition
   Authors: Max Kuhn and Kjell Johnson; Springer; 2016. **ISBN-13:** 978-1-4614-6848-6
2. *Machine Learning: A Concise Introduction*, 1st Edition
   Author: Steven W. Knox; Wiley; 2018. **ISBN-13:** 978-1-119-43919-6
3. *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd Edition
   Authors: Trevor Hastie, Robert Tibshirani, and Jerome Friedman; Springer; 2008. (ESL) **ISBN-13:** 978-0387848570
4. *Machine Learning: An Algorithmic Perspective*, 2nd Edition
   Author: Stephen Marsland; CRC Press; 2014. **ISBN-13:** 978-1-4614-7137-0
5. *Deep Learning*, 1st Edition
   Authors: Ian Goodfellow, Yoshua Bengio, and Aaron Courville; MIT Press; 2016. (DL) **ISBN-13:** 978-0262035613
6. *Neural Networks and Learning Machines*, 3rd Edition
   Author: Simon Haykin; Pearson; 2008. **ISBN-13:** 978-0131471399
7. *Neural Networks and Deep Learning: A Textbook*, 1st Edition
   Authors: Charu Aggrawal; Springer; 2018. **ISBN-13:** 978-3319944623
8. *Introduction to Machine Learning*, 2nd Edition
   Author: Ethem Alpaydine; MIT Press; 2010. (AL) **ISBN-13:** 978-8120350786
9. *Machine Learning*, 1st Edition
   Author: Tom M. Mitchell; McGraw-Hill Education; 1997. **ISBN-13:** 978-0070428072

Grading Policies:

- The letter grade distribution table guarantees the *minimum* grade each student will receive based on their final score. When appropriate, relative performance measures will be used to assign the final grade, at the discretion of the instructor.
  - Final grades are non-negotiable and are assigned at the discretion of the instructor. If you cannot accept this condition, you should not enroll in this course.
- Your lowest grade in problem sets and your lowest grade in programming assignments will be dropped from the final grade.
- *Participation on Piazza has up to 5% extra credit, which is granted on a competetive basis *at the discretion of the instructor.*
- **Homework Policy**
  - Homework is assigned on an approximately weekly basis. A one-day grace period can be used for each homework with 10% penalty. *Absolutely no late homework will be accepted after the grace period. A late assignment results in a zero grade.*
  - Homework solutions should be typed or *scanned* using scanners or mobile scanner applications like CamScan and uploaded (photos taken by cell-phone cameras and in formats other than pdf will NOT be accepted). Programs and simulation results have to be uploaded on github as well.

  - Students are encouraged to discuss homework problems with one another, but each student must do their own work and submit individual solutions written/ coded in their own hand. Copying the solutions or submitting identical homework sets is written evidence of cheating. The penalty ranges from F on the homework or exam, to an F in the course, to recommended expulsion.
  - Posting the homework assignments and their solutions to online forums or sharing them with other students is strictly prohibited and infringes the copyright of the instructor. Instances will be reported to USC officials as academic dishonesty for disciplinary action.

- **Exam Policy**
  - **Make-up Exams:** No make-up exams will be given. If you cannot make the above dates due to a class schedule conflict or personal matter, you must drop the class. In the case of a required business trip or a medical or family emergency, a signed letter from your manager or counselor or physician has to be submitted. This letter must include the contact of your physician or counselor or manager.
  - Midterm and final exams will be closed book and notes. Calculators are allowed but computers and cell-phones or any devices that have internet capability are not allowed. One letter size cheat sheet (back and front) is allowed for the midterm. Two letter size cheat sheets (back and front) are allowed for the final.
  - All exams are cumulative, with considerable emphasis on material presented since the last exam.

- **Attendance:**
  - Students are required to attend all the lectures and discussion sessions and actively participate in class discussions. Use of cellphones and laptops is prohibited in the classroom. If you need your electronic devices to take notes, you should discuss with the instructor at the beginning of the semester.

Important Notes:

- Textbooks are secondary to the lecture notes and homework assignments.
- Handouts and course material will be distributed.
- Please use your USC email to register on Piazza and to contact the instructor and TAs.

Tentative Course Outline

Monday	Wednesday
Aug 26th **1** **Introduction to Statistical Learning** (ISLR Chs.1,2, ESL Chs.1,2) Supervised vs. Unsupervised Learning	28th **2** **Introduction to Statistical Learning** (ISLR Chs.1,2, ESL Chs.1,2) Model Assessment
Sep 2nd Labor Day	4th **3** **Linear Regression** (ISLR Ch.3, ESL Ch. 3) Estimating Coefficients Estimating the Accuracy of Coefficients
9th **4** **Linear Regression** (ISLR Ch.3, ESL Ch. 3) Qualitative Variables Comparison with K-Nearest Neighbors	11th **5** **Linear Regression** (ISLR Ch.3, ESL Ch. 3) Qualitative Variables Comparison with K-Nearest Neighbors
16th **6** **Classification** (ISLR Ch. 4, ESL Ch. 4) Logistic Regression Class Imbalance Multinomial Regression	18th **7** **Classification** (ISLR Ch. 4, ESL Ch. 4) Subsampling and Upsampling SMOTE Bayesian Linear Discriminant Analysis
23rd **8** **Classification** (ISLR Ch. 4, ESL Ch. 4) Multinomial Regression Linear Discriminant Analysis	25th **9** **Classification** (ISLR Ch. 4, ESL Ch. 4) Measures for Evaluating Classifiers Quadratic Discriminant Analysis* Comparison with K-Nearest Neighbors
30th **10** **Resampling Methods** (ISLR Ch. 5, ESL Ch. 7) Cross-Validation The Bootstrap	Oct 2nd **11** **Linear Model Selection and Regularization** (ISLR Ch.6, ESL Ch. 3) Subset Selection Shrinkage Methods
7th **12** **Linear Model Selection and Regularization** (ISLR Ch.6, ESL Ch. 3) Dimension Reduction Methods* Considerations in High Dimensions*	9th **13** **Tree-based Methods** (ISLR Ch. 8, ESL Chs. 9, 10) Regression and Classification Trees
14th **14** **Tree-based Methods** (ISLR Ch. 8, ESL Chs. 9, 10) Bagging, Boosting, and Random Forests	16th **15** **Tree-based Methods** (ISLR Ch. 8, ESL Chs. 9, 10) Bagging, Boosting, and Random Forests

Monday	Wednesday
21st **16** **Ensemble Learning*** (ESL Ch. 16) Combination Methods Boosting and Regularization Paths Learning Ensembles Mixture of Experts and Stacking	23rd **17** **Midterm**
28th **18** **Support Vector Machines** (ISLR Ch. 9, ESL Ch. 12) Maximal Margin Classifier Support Vector Classifiers Support Vector Machines The Kernel Trick L1 Regularized SVMs Multi-class and Multilabel Classification The Vapnik-Chervonenkis Dimension* Support Vector Regression*	30th **19** **Unsupervised Learning** (ISLR Ch. 10, ESL Ch. 14) K-Means Clustering Hierarchical Clustering Competetive Learning and Self-Organizing Maps*
Nov 4th **20** **Unsupervised Learning** (ISLR Ch. 10, ESL Ch. 14) Practical Issues in Clustering Association Rules* Gaussian Mixtures and Soft K-Means*	6th **21** **Unsupervised Learning** (ISLR Ch. 10, ESL Ch. 14) Principal Component Analysis Anomaly Detection* Association Rules* Gaussian Mixtures and Soft K-Means*
11th **22** **Active and Semi-Supervised Learning** Semi-Supervised Learning Self-Training Co-Training Yarowsky Algorithm Refinements Active vs. Passive Learning Stream-Based vs. Pool-Based Active Learning Query Selection Strategies	13th **23** **Neural Networks and Deep Learning** (ESL Ch. 11, DL Ch. 6) The Perceptron Feedforward Neural Networks
18th **24** **Neural Networks and Deep Learning** (ESL Ch. 11, DL Ch. 6) Feedforward Neural Networks Backpropagation and Gradient Descent Overfitting	20th **25** **Neural Networks and Deep Learning** (DL Chs. 6, 7) Autoencoders and Deep Feedforward Neural Networks* Regularization Early Stopping and Dropout Adversarial Training*

Monday	Wednesday
25th **26** **Neural Networks and Deep Learning** (DL Chs. 9, 10) Convolutional Neural Networks	27th Thanksgiving Recess
Dec 2nd **27** **Neural Networks and Deep Learning** (DL Chs. 9, 10) Sequence Modeling* Recurrent Neural Networks* Sequence-to-Sequence Modeling* Long Short Term Memory (LSTM) Neural Networks*	4th **28** **Reinforcement Learning*** Definitions Task-Reward-Policy Formulation Total Discounted Future Reward Optimal Policy Value Function Q-Function The Bellman Equation Q-Learning Exploration- Exploitation Temporal Difference Learning Extensions to Stochastic Environments and Rewards Deep Reinforcement Learning

Notes:

- Items marked by * will be covered only if time permits.

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 one's own academic work from misuse by others as well as to avoid using another's work as one's own. All students are expected to understand and abide by these principles. SCampus, the Student Guidebook, contains the University Student Conduct Code (see University Governance, Section 11.00), while the recommended sanctions are located in Appendix A. See: http://scampus.usc.edu.

Emergency Preparedness/Course Continuity in a Crisis In case of a declared emergency if travel to campus is not feasible, USC executive leadership will announce an electronic way for instructors to teach students in their residence halls or homes using a combination of Blackboard, teleconferencing, and other technologies. See the university's site on Campus Safety and Emergency Preparedness: http://preparedness.usc.edu

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. Website: http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html

(213) 740-0776 (Phone), (213) 740-6948 (TDD only), (213) 740-8216 (FAX) ability@usc.edu.