Ira A. Fulton Schools of Engineering
Arizona State University

Course Syllabus

Statistical Machine Learning (CSE 575)

Course Description

Deriving generalizable models from some given training data is central to statistical machine learning. Statistical machine learning has found wide applications in many fields including artificial intelligence, computer vision, natural language processing, finance, bioinformatics, and etc. This course provides a systematic introduction to common learning paradigms in statistical machine learning, accompanied by an exploration of a set of foundational algorithms. Main topics covered include supervised learning, unsupervised learning, and deep learning.

Specific topics covered include:

- Mathematical foundations for machine learning
- Maximum likelihood estimation
- Naive Bayes classification
- Logistic regression
- Support vector machines
- Probabilistic graphical models
- Mixture models
- K-means clustering
- Spectral clustering
- Dimensionality reduction
- Principal component analysis
- Neural networks and deep learning
- Convolutional neural networks

Learning Outcomes

Learners completing this course will be able to:

- Distinguish between supervised learning and unsupervised learning
- Apply common probability distributions in machine learning applications
- Use cross validation to select parameters
- Use maximum likelihood estimate (MLE) for parameter estimation
- Implement fundamental learning algorithms such as logistic regression and k-means clustering
- Implement more advanced learning algorithms such as support vector machines and convolutional neural networks
- Design a deep network using an exemplar application to solve a specific problem
- Apply key techniques employed in building deep learning architectures

Estimated Workload/ Time Commitment Per Week

Average of 15-20 hours per week

Required Prior Knowledge and Skills

- Basics of linear algebra, statistics, calculus, and algorithm design and analysis
- Programming in Python

Technology Requirements

Hardware
Standard with major OS

Software and Other
Standard - technology integrations will be provided through Coursera

Textbook and Readings

CSE 575 does not have a required textbook; however, a reference textbook is provided. Information about it and several other textbooks the faculty recommends can be found on the course site.

Course Content

Instruction

Video Lectures
Other Videos
Readings
Live Sessions (office hours, webinars, etc.)

Assessments

Knowledge check quizzes (ungraded)
Graded quizzes (auto-graded
Individual Project
Midterm and Final Exam (proctored, auto-graded and/or instructor-graded)

Details of the main instructional and assessment elements this course comprises follow:

Lecture videos. In each module the concepts you need to know will be presented through a collection of short video lectures. You may stream these videos for playback within the browser by clicking on their titles or download the videos. You may also download the slides that are used in the videos. The lecture slides, where available, are provided with the video.

Knowledge Checks. Each module includes at least one knowledge check quiz to help you assess your understanding of the topics. You will be allowed unlimited attempts for each one. There is no time limit on how long you take to complete each attempt at the quiz, and these quizzes are not counted toward your final grade in the class.

Graded Quizzes. Each module includes one graded quiz, comprised of 10 questions. You will be allowed one attempt for each of these quizzes. There is no time limit on how long you take to complete each attempt at the quiz. *There is a 10% grade penalty for each day late past the deadline.*

Graded Discussion Prompts. Each module includes one graded discussion prompt. Each prompt provides a space for you to respond. After responding, you can see your peers' responses, and are required to respond to at least one classmate to receive full credit. *There is a 10% grade penalty for each day late past the deadline.*

Proctored Exams: You will have two (2) proctored exams, a midterm and a final. ProctorU is an online proctoring service that allows students to take exams online while ensuring the integrity of the exam for the institution. Additional information and instructions are provided in Module 1 of the course.

Projects: This course includes one individual project in three parts. The first part is introduced during Week 1 and is due at the end of Week 3. The second part, will be due at the end of Week 5, and the final part will be due in the final week of the term. The project is instructor graded. *There is a 10% grade penalty for each day late past the deadline.*

Course Grade Breakdown

Course Work	**Quantity**	**Percentage of Grade**
Graded discussion	8	5%
Module quizzes	8	5%
Homework assignments	6	20%
Project	1	30%
Mid-term exam	1	15%

Final exam	1	25%

NOTES: 1) Not all course work types are required. 2) Homework assignments may utilize Coursera programming, Jupyter Notebook, and code block assignments with custom auto-graders.

Grade Scale

NOTE: You must earn a cumulative grade of 70% to earn a "C" in this course. Grades in this course will *not* include pluses or minuses.

A	88% - 100%
B	75% - 87%
C	60% - 74%
D	50% - 59%
E	<50%

Course Schedule

Live Events - Weekly (meet with the course instructor and your classmates to learn more about course topics and discuss assignments):

Exact day and time TBD

Please Note:

- *These events will be recorded and uploaded to the course.*

Virtual Office Hours - Weekly (another chance to get your questions answered from the course instructor and/or teaching assistants):

Exact day and time TBD (Check the Live Events page in the course for your local time and access details.)

Week/Module
Week 1: Orientation & Introduction
Week 2: Supervised learning- Naive Bayes & Logistic Regression
Week 3: Linear Machines and SVM
Week 4: Graphical Models
Midterm Exam
Week 5: Unsupervised Learning & Data Clustering
Week 6: Dimensionality Reduction
Week 7: Neural Networks & Deep Learning
Week 8: Exemplar Deep Learning Application
Final Exam

Assignment Deadlines

Unless otherwise noted, all graded work is due on Sunday 11:59 pm Arizona time for the week it is assigned. A late penalty of 10% for each day late will be applied for work submitted after the scheduled due date and time.

Course Outline with Assignments

Module 1: Orientation & Introduction

Lesson 1: Introduction to Machine Learning

Lesson 2: Machine Learning Examples

Assignments

- ❑ Week 1 Graded Discussion
- ❑ Mini-assignment: Introduction to Machine Learning
- ❑ Week 1 Graded Quiz

Module 2: Supervised learning- Naive Bayes & Logistic Regression

Lesson 1: Review key concepts in calculus, set theory, and linear algebra
Lesson 2: Review of probability theory
Lesson 3: Review of random variables and their distributions
Lesson 4: Common densities in machine learning
Lesson 5: Set-up of supervised learning; Regression
Lesson 6: Classification; Density estimation
Lesson 7: Generative & discriminative models

Assignments

- ❏ Mini-assignment: Estimators
- ❏ Week 2 Graded Discussion
- ❏ Week 2 Graded Quiz

Module 3: Linear Machines and SVM

Lesson 1: Implement the fundamental learning algorithm Naive Bayes
Lesson 2: Implement the fundamental learning algorithm Logistic Regression
Lesson 3: Linear Machines: Basics
Lesson 4: The Concept of Margin

Assignments

- ❏ Week 3 Graded Discussion
- ❏ Week 3 Graded Quiz
- ❏ Project part 1

Module 4: Graphical Models

Lesson 1: SVM: Linearly-separable Case
Lesson 2: SVM: Non-linearly-separable Case
Lesson 3: Bayesian Networks
Lesson 4: Hidden Markov Models: Basics
Lesson 5: Hidden Markov Models: Learning & Inference

Midterm Exam

Assignments

- ❏ Midterm Exam - Proctored
- ❏ Week 4 Graded Discussion
- ❏ Assignment #1: Activity Recognition Submission

- ❑ Week 4 Graded Quiz

Module 5: Unsupervised Learning & Data Clustering

Lesson 1: Set-up of the unsupervised learning problem
Lesson 2: Gaussian Mixture Models & the EM Algorithm
Lesson 3: The k-means Algorithm
Lesson 4: Analyzing the k-means Algorithm
Lesson 5: The Basics of Spectral Clustering

Assignments

- ❑ Mini-assignment: K-Means
- ❑ Project part 2
- ❑ Week 5 Graded Discussion
- ❑ Week 5 Graded Quiz

Module 6: Dimensionality Reduction

Lesson 1: Graph Cut Formulation
Lesson 2: Going Beyond MinCut
Lesson 3: Practical Considerations in Implementation
Lesson 4: Introduction to the Problem of Dimensionality Reduction

Assignments

- ❑ Mini-assignment, part 1&2
- ❑ Week 6 Graded Discussion
- ❑ Week 6 Graded Quiz

Module 7: Neural Networks & Deep Learning

Lesson 1: Principal Component Analysis: the Basic Idea
Lesson 2: PCA: the Algorithm & Further Discussion
Lesson 3: Introduction to neural networks & deep learning
Lesson 4: Key enabling techniques for deep learning

Assignments

- ❑ Mini-assignment
- ❑ Week 7 Graded Discussion
- ❑ Week 7 Graded Quiz

Module 8: Exemplar Deep Learning Applications

Lesson 1: Basic Deep Architectures
Lesson 2: Deep learning for image-based recognition

Lesson 3: Deep learning for video-based recognition

Lesson 4: GAN and its applications

Final Exam

Assignments

- ❏ Final Exam - Proctored

Assignments

- ❏ Optional: Portfolio Inclusion Report for ASU MCS Degree
- ❏ Course Survey

Policies

All ASU and Coursera policies will be enforced during this course. For policy details, please consult the MCS Graduate Handbook 2018 -- 2019 and/or the MCS Onboarding Course.

Academic Integrity

Students in this class must adhere to ASU's academic integrity policy, which can be found at https://provost.asu.edu/academic-integrity/policy). Students are responsible for reviewing this policy and understanding each of the areas in which academic dishonesty can occur. In addition, all engineering students are expected to adhere to both the ASU Academic Integrity Honor Code and the Fulton Schools of Engineering Honor Code. All academic integrity violations will be reported to the Fulton Schools of Engineering Academic Integrity Office (AIO). The AIO maintains records of all violations and has access to academic integrity violations committed in all other ASU college/schools.

Course Faculty

The following faculty member created this course.

Baoxin Li

Baoxin Li is currently a professor and the chair of the Computer Science & Engineering Program and a Graduate Faculty Endorsed to Chair in the Electrical Engineering and Computer Engineering programs. From 2000 to 2004, he was a Senior Researcher with SHARP Laboratories of America, where he was the technical lead in developing SHARP's HiIMPACT Sports™ technologies. He was also an Adjunct Professor with the Portland State University from 2003 to 2004. His general research interests are on visual computing and machine learning, especially their application in the context of human-centered computing.