CSE 575: Statistical Machine Learning

****Any information in this syllabus may be subject to change with reasonable advance notice.**

We will not be conducting any Zoom classes.

To ensure that students have access to the course materials, I will upload all lecture slides and recordings (where applicable) after each class on Canvas. students are responsible to check the records/slides if they are not in class.

Please note that the MT and final exam will only be conducted in person, and there will be no online option available. All students are required to be present in class for the MT and Final exam .

Note: The instructor runs a pre-packaged course developed by Dr. Baoxin Li , Hanghang Tong and Jingrui He with all lectures pre-recorded and delivered on-line.

About this Course

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
● K-means clustering

● Dimensionality reduction
● Principal component analysis
● Neural networks and deep learning
● Convolutional neural networks

Technologies covered include:

● Python
● Matlab
● Jupyter Notebooks
● Google Colab
● PyTorch

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

Course Map:

NO Make up for MT and final exam.

Task name	Tasks complete date	Comment
Module 1	July 1st-July 7th	Read the Syllabus Carefully
Module 2	July 8th- July 14th	
Module 3	July 15th - July 21st	
MT exam	July 23rd @11 am	In person exam
Module 4	July 22nd-July 28th	
Module 5	July 29th - Aug 4th	
Final exam	August 8th @11 am	In person exam
Project 1: Density Estimation and Classification	July 21st	
Project 2: Unsupervised Learning (K-means)	July 30	
Project 3: Classification Using Neural Networks and Deep Learning	August 6th	

Course policy:

Graded task date/time are fixed and we are not going to change the following their due date.

Makeup for graded tasks is not permitted, and we will not be extending the due date for individuals. Late penalties will not be accepted except for tasks specified in their description.

Instructor: Samira Ghayekhloo

Main class:	Tu/Thurs 11:00 AM - 12:15 PM @BYAC 110
Office hour:(**By appointment**)	Monday: 10-11 am In person at BYENG 514 OR Zoom: https://asu.zoom.us/j/7795505455

IA: Ali Sarabi

email:	alisarabi@asu.edu
Office hour:	

Grader : Koushik Sai Achyuth Ayila

email:	kayila@asu.edu

Grade Breakdown

Course Work	Quantity	Team or Individual	Percentage of Grade	weight of tasks
Auto-graded quizzes	5	Individual	15%	Equal weight for all 5 quizzes
Project #1	1	Individual	15%	
Project #2	1	Individual	15%	
Project #3	1	Individual	15%	
Midterm Exam	1	Individual	20%	
Final Exam	1	Individual	20%	
End-of-semester survey	1	Individual	1%	
Total			101%	

Grading

You must earn a cumulative grade of 70% to earn a “C” in this course. Grades in this course will include pluses (+) and minuses (-).

A+	≥ 97%	**C+**	≥ 77% and < 80%
A	≥ 93% and < 97%	**C**	≥ 70% and < 77%
A-	≥ 90% and < 93%	**D**	≥ 60% and < 70%
B+	≥ 87% and < 90%	**E**	< 60%
B	≥ 83% and < 87%		
B-	≥ 80% and < 83%		

* The instructor reserves the right to curve if necessary.

Required Prior Knowledge and Skills

This course will be very challenging, and learners are expected to learn the necessary technologies in their own time.

Proficient Mathematical Skills and Theoretical Understanding

- Basics of linear algebra
- Basics of probability and statistics
- Basics of calculus and set theory
- Basics of algorithm design and analysis.

Strong Application Skills

- Programming in Python
- Ability to effectively read Python code
- Confidence executing at least one programming language:
  - Python
  - MATLAB
  - R

Note: It is highly recommended that learners use Python to complete the coursework to get more support from the course team.

Proficient Experience

- High level programming language.
  - Python or MATLAB
- Ability to implement Machine Learning algorithms using Python
- Familiarity with any one of the following frameworks:
  - Jupiter Notebook
  - Google Collab
  - Pytorch
- Familiarity with the following tools/libraries:
  - NumPy
  - Pandas
  - TensorFlow
  - Keras
  - Matplotlib
  - Scikit Learn

Technology Requirements

Hardware

- Standard personal computer with major operating system
- Reliable, strong Internet connection
- Webcam
- Microphone

Software/Other

- GPU environment like Google Collab or personal setup on your own
- Jupiter Notebook
- Pytorch
- MATLAB
- Anaconda

Textbook and Readings:

At the graduate level, inquiry, research, and critical reading are part of the learning experience; however, this course does not have a required textbook. Any required readings are provided within or are accessible through the course of the ASU Library.

Professor Christopher Bishop, PhD has given ASU permission to provide his textbook Pattern Recognition and Machine Learning to you in this course.

You may print portions of the book or the entire book, but you may not share the PDF, in whole or part, with any parties outside of this course.

Please note:
1. Material in Pattern Recognition and Machine Learning that corresponds with topics covered in the lecture videos is highlighted in the course's Recommended Readings, which can be found in the overview section at the beginning of each week.

2. PRMLT | Pattern Recognition and Machine Learning Toolbox is a companion to Professor Bishop's textbook. On it you will find a package that is a MATLAB implementation of the algorithms described in the book, some of which are covered in this course.

For interested learners, Drs. He, Tong, and Li recommend:

- The Elements of Statistical Learning: Data Mining, Inference, and Prediction, Second Edition. Trevor Hastie, Robert Tibshirani, and Jerome Friedman. Springer, 2009.

- Semi-Supervised Learning. Olivier Chapelle, Bernhard Sch¨olkopf, and Alexander Zien. The MIT Press, 2006.

- Kernel Methods for Pattern Analysis. John Shawe-Taylor and Nello Cristianini. Cambridge University Press, 2004.

- Pattern Classification, Second Edition. Richard Duda, Peter Hart, and David Stork. Wiley, 2000.

- Machine Learning. Tom Mitchell. McGraw Hill, 1997.

- Introduction to Data Mining. Pang-Ning Tan, Michael Steinbach, and Vipin Kumar. Addison Wesley, 2005.

- Data Mining: Theories, Algorithms, and Examples. Nong Ye. CRC Press, 2013.

Other Policies:

Remarks on Electronic Communication:

Questions about class content, homework, projects, etc., must be posted on ***Discussion page on Canvas***; otherwise, your emails will be ignored. For other class-related questions, email the instructor or TAs. Before sending an email please follow the excellent advice http://www.wikihow.com/Email-a-Professor.

The instructor and TAs will try to respond to any question as soon as possible and ***within 48 hours (weekdays). Please check existing questions before posting a new one since it may have already been asked and answered***. Send emails of questions directly to the instructor or TAs only for other class-related questions.

We are using Ed discussion page for each module. Please write down your questions on specified module number. IA will reply to your question less than 24 hours.

Attendance policy:

According to the university rules, attendance and participation in class activities is an essential part of the learning process, and students are expected to attend class regularly.

Excused absences do not relieve students from responsibility for any part of the course work required during the period of absence. Faculty will provide accommodations that may include access to recordings of class activities.

Excused absences for classes will be given without penalty to the grade in the case of (1) a university-sanctioned
event [ACD 304-02]; (2) religious holidays [ACD 304-04]; a list can be found here https://eoss.asu.edu/cora/holidays]; (3) work performed in the line-of-duty according [SSM 201-18]; Excused absences do not relieve students from responsibility for any part of the course work required during the period of absence.

Although attendance is not mandatory, We don't have any zoom class, you can join in person if you like.

I will upload all the lectures slides / records* after each class on Canvas.

***Recordings may be used to accommodate student absences in some sessions.* No make-up for the MT and Final exam.**

End-of-semester survey:

It is extremely important that you respond to the final anonymous survey solicited by the university at the end of the school year. The overall feedback helps me make changes for the next year. The survey is often released 1-2 weeks before the final at: https://fultonapps.asu.edu/eval/

Policy regarding expected classroom behavior (e.g., use of pagers, recording devices):

Any violent or threatening conduct by an ASU student in this class will be reported to the ASU Police Department and the Office of the Dean of Students.

Academic Integrity:

Mandatory Statement
Students in this class must adhere to ASU's academic integrity pol icy, 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.
All engineering students are expected to adhere to the ASU Academic Integrity Honor Code.

All work submitted for the course cannot have been submitted for any other course or any previous section of this same course. Student academic integrity violations are reported to the Fulton Schools of Engineering
Academic Integrity Office (AIO). Withdrawing from this course will not absolve you of responsibility for an academic integrity violation and any sanctions that are applied. The AIO maintains a record of all violations and has access to academic integrity violations committed in all other ASU college/schools.

Unless explicitly allowed by your instructor, the use of generative AI tools on any course assignment or exam will be considered academic dishonesty and a violation of the ASU Academic Integrity Policy. Students confirmed to be engaging in non- allowable use of generative AI will be sanctioned according to the academic integrity policy and FSE sanctioning guidelines.

Student Copyright Responsibilities:

Mandatory Statements

You must refrain from uploading to any course shell, discussion board, or website used by the course instructor or other course forum, material that is not the student's original work, unless the student first complies with all applicable copyright laws; faculty members reserve the right to delete materials on the grounds of suspected copyright infringement.

The contents of this course, including lectures and other instructional materials, are copyrighted materials. Students may not share outside the class, including uploading, selling or distributing course content or notes taken during the conduct of the course. Any recording of class sessions is authorized only for the use of students enrolled in this course during their enrollment in this course. Recordings and excerpts of recordings may not be distributed to others. (see ACD 304–06, "Commercial Note Taking Services" and ABOR Policy 5-308 F.14for more information).

Policy against threatening behavior, per the Student Services Manual, SSM 104–02:

Students, faculty, staff, and other individuals do not have an unqualified right of access to university grounds, property, or services (see SSM 104-02). Interfering with the peaceful conduct of university-related business or activities or remaining on campus grounds after a request to leave may be considered a crime. All incidents and allegations of violent or threatening conduct by an ASU student (whether on- or off-campus) must be reported to the ASU Police Department (ASU PD) and the Office of the Dean of Students.

Disability Accommodations:

Suitable accommodations are made for students having disabilities. Students needing accommodations must register with the ASU Disabilities Resource Center and provide documentation of that registration to the instructor. Students should communicate the need for an accommodation in enough time for it to be properly arranged. See ACD 304-08 Classroom and Testing Accommodations for Students with Disabilities.

Harassment and Sexual Discrimination:

Arizona State University is committed to providing an environment free of discrimination, harassment, or retaliation for the entire university community, including all students, faculty members, staff employees, and guests. ASU expressly prohibits discrimination, harassment, and retaliation by employees, students, contractors, or agents of the university based on any protected status: race, color, religion, sex, national origin, age, disability, veteran status, sexual orientation, gender identity, and genetic information.

Title IX is a federal law that provides that no person be excluded on the basis of sex from participation in, be denied benefits of, or be subjected to discrimination under any education program or activity. Both Title IX and university policy make clear that sexual violence and harassment based on sex is prohibited. An individual who believes they have been subjected to sexual violence or harassed on the basis of sex can seek support, including counseling and academic support, from the university. If you or someone you know has been harassed on the basis of sex or sexually assaulted, you can find information and resources at https://sexualviolenceprevention.asu.edu/faqs.

As a mandated reporter, I am obligated to report any information I become aware of regarding alleged acts of sexual discrimination, including sexual violence and dating violence. ASU Counseling Services, https://eoss.asu.edu/counseling is available if you wish to discuss any concerns confidentially and privately. ASU online students may access 360 Life Services, https://goto.asuonline.asu.edu/success/online-resources.html.

Creators:

Jingrui He

Jingrui He, PhD is an associate professor in the School of Computing and Augmented Intelligence (SCAI) within the Fulton Schools of Engineering (FSE) at Arizona State University (ASU). She received her PhD from Carnegie Mellon University. She joined ASU in 2014 and directs the Statistical Learning Lab (STAR Lab). Her research focuses on rare category analysis, heterogeneous machine learning, active learning and semi-supervised learning, with applications in social media analysis, healthcare, manufacturing process, etc.

Baoxin Li

Baoxin Li, PhD 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.

Hanghang Tong

Hanghang Tong, PhD is an associate professor at School of Computing and Augmented Intelligence (SCAI) within the Fulton Schools of Engineering (FSE) at Arizona State University (ASU) since August 2014. Before that, he was an assistant professor at the Computer Science Department, City College, City University of New York, a research staff member at IBM T.J. Watson Research Center and a Post-doctoral fellow in Carnegie Mellon University. His research interest is in large scale data mining for graphs and multimedia.