TEXAS A&M UNIVERSITY

SYLLABUS

Course title and number	CSCE 636: Deep Learning
Term	Spring 2024
Meeting time and location	MWF 11:30 am-12:20 pm, HRBB 124
Credit hours	3

Course Description and Prerequisites

An introduction to the field of deep learning, including basic machine learning, supervised learning, logistic regression, loss functions, neural networks, optimization, error back-propagation, regularization and generalization, unsupervised learning and expectation and maximization, convolutional neural networks, graph neural networks, attention and transformers, large language models, advanced topics of deep learning, applications to natural language processing and computer vision.

Students are expected to have some level of familiarity with basic linear algebra (including vectors, matrices, matrix-vector computations, vector and matrix norms, linear independence, matrix rank, singularity, positive definiteness, eigenvalues/eigenvectors, matrix decomposition, orthogonality), multivariate calculus (including derivatives of univariate functions, derivatives of multivariate functions, chain rule, Taylor expansion), and basic probability and statistics (including discrete and continuous probability distributions, sum rule, product rule, marginal probability distributions, conditional probability distributions, joint probability distributions, independence and conditional independence, Bayes Theorem, variance and covariance, expectation). Students need to have access to a GPU and be proficiency in Python programming.

Learning Outcomes or Course Objectives

Student learning outcomes include (1) understanding the foundation, major techniques, applications, and challenges of deep learning; (2) the ability to apply basic deep learning algorithms for solving real-world problems. The learning outcomes will be assessed based on a combination of homework assignments, exams, projects, and presentations.

Instructor Information

Name	Shuiwang Ji
Telephone number	(979) 458-1547
Email address	sji@tamu.edu
Office hours	Tuesday and Thursday: 2:00 pm – 3:00 pm
Office location	Zoom (Link on Canvas)

TA Information

Name	Xiner Li
Email address:	lxe@tamu.edu
Office Hours:	Tuesday & Wednesday: 3pm - 4pm
Office Location	Zoom (Link on Canvas)

Name Shurui Gui
Email address: shurui.gui@tamu.edu
Office Hours: Thursday & Friday: 3:00 pm – 4:00 pm
Office Location Zoom (Link on Canvas)

Grader Information

Name: Shubham Parashar
Email address: shubhamprshr@tamu.edu

Textbooks (Required)

YS Abu-Mostafa, M Magdon-Ismail, HT Lin: Learning from Data **(LFD)** https://amlbook.com/ e-Chapters at https://amlbook.com/eChapters.html (Need a copy of the book to access e-Chapters)
C. M. Bishop and H. Bishop: Deep Learning: Foundations and Concepts **(DL)** https://www.bishopbook.com/

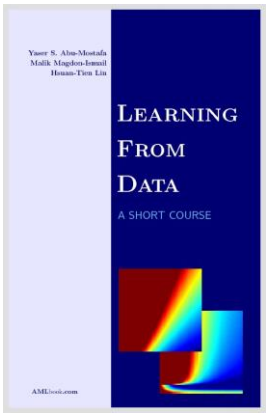

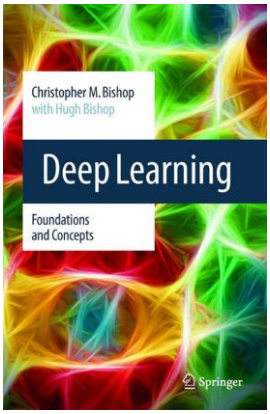

Course Delivery and Management

This course will be managed via Canvas (https://canvas.tamu.edu/). All hand-outs will be distributed via Canvas, including assignments, lecture materials, etc.

Grading Policies

Homework (5): 40%: There will be five homework assignments containing both written and programming components. The total number of points for each assignment may be different.

Exam (2): 35%: There will be two exams covering the foundations of neural networks and deep learning. Exam 1 will be in class. Exam 1 will be 15% and final exam will be 20%.
Final exam schedule: 10:30 a.m. – 12:30 p.m. on Tuesday, May 7, 2024

Project (1): 25%: There will be one semester-long project. The project is for individual student, not for group. Details of project will be presented during the first class. This will be a structured project in which each student is given the same task, but you can explore different solutions to this task. In the end, each student needs to submit their results (code and prediction results) and a report summarizing the methods and results. Students are required to use LaTex for typesetting the report and the NeurIPS LaTex template (https://nips.cc/) is recommended.

All homework assignments are individual and collaboration among students is strictly prohibited. Project reports should be treated as scientific publications, and all rules governing paper-writing apply.

Exam Policy

Exam 1 will be in class, and final exam schedule is determined by the University. All exams are closed-book, closed-notes, closed-Internet. You are allowed to use a calculator. You are also allowed to bring ONE sheet of letter-sized (8.5 x 11 inches) paper ("cheat sheet") on which you can write anything you wish to. Note that you can choose to use two-sided, typed or hand-written with any font size, and the only requirement on the cheat sheet is that it has to be one sheet of letter-sized paper. All exams are strictly closed to neighbors and other students, and any violations will be reported to Aggie Honor System Office.

Grading Disputes

Questions on assignment grading should be discussed with the TA. Questions on exam grading should be discussed with the instructor. Grading dispute period for the final exam will be determined and announced at the time of grade posting based on university deadline for submitting letter grades. Grading disputes for all other graded materials (other than the final exam) must be presented to the respective party within ONE week upon receiving grading results (timed as email notification for electronic ones or returning of graded materials for hard-copies). All grades after the dispute time windows are considered final.

Late Policies

For homework assignment, 20% is deducted for each late day for up to three days (including weekends) after which submissions are not accepted. Late project submissions will not be accepted.

Excused absence will not be counted towards late days. If an exam overlaps with an excused absence, the instructor will provide the student an opportunity to make up an exam by a date agreed upon by the student and instructor. If the instructor has a regularly scheduled make up exam, students are expected to attend unless they have a university approved excuse.

Unexpected excused absences: In cases where prior notification of excused absence is not feasible (e.g., accident or emergency) the student must provide notification by the end of the second working day after the absence, including an explanation of why notice could not be sent prior to the class.

Student Rule 7 explains attendance policies and excused absences. https://student-rules.tamu.edu/rule07/

Grading Scale

Final letter grades will be based on absolute percentage as follows*:*

A = [90, 100]
B = [80, 90)
C = [70, 80)
D = [60, 70)
F = <60

[] denotes inclusive; () denotes exclusive;

Course Outline

Week	Topic	References	Key dates
1	Introduction, linear regression	LFD: Chapter 3	
2	Logistic regression	LFD: Chapter 3	
3	Softmax regression	PDF handout	HW1
4	Multi-layer networks	LFD: e-Chapter 7; DL: 6.4	
5	Multi-layer networks	LFD: e-Chapter 7; DL: 6.4	
6	Convolutional neural networks	DL: Chapters 7, 9, and 10.1-10.2	HW2
7	Convolutional neural networks	DL: Chapters 7, 9, and 10.1-10.2	
8	A unified view of loss functions	PDF handout	Exam 1
9	Neural networks and kernel methods	PDF handout	HW3
10	Attention and transformers	DL: Chapter 12	
11	Large language models	DL: Chapter 12	
12	Graph neural networks	DL: Chapter 13	HW4
13	Graph neural networks	DL: Chapter 13	
14	Principal component analysis	PDF handout; DL: 19.1.1	Project due
15	Latent variable models and EM	DL: Chapter 15; LFD: e-Appendix C	HW5

* LFD and DL refer to the two textbooks, respectively.

Americans with Disabilities Act (ADA)

Texas A&M University is committed to providing equitable access to learning opportunities for all students. If you experience barriers to your education due to a disability or think you may have a disability, please contact Disability Resources in the Student Services Building or at (979) 845-1637 or visit disability.tamu.edu. Disabilities may include, but are not limited to attentional, learning, mental health, sensory, physical, or chronic health conditions. All students are encouraged to discuss their disability related needs with Disability Resources and their instructors as soon as possible.

Academic Integrity

An Aggie does not lie, cheat or steal, or tolerate those who do.

Texas A&M University students are responsible for authenticating all work submitted to an instructor. If asked, students must be able to produce proof that the item submitted is indeed the work of that student. Students must keep appropriate records at all times. The inability to authenticate one's work, should the instructor request it, may be sufficient grounds to initiate an academic misconduct case" (Section 20.1.2.3, Student Rule 20).

You can learn more about the Aggie Honor System Office Rules and Procedures, academic integrity, and your rights and responsibilities at aggiehonor.tamu.edu.

Statement on Mental Health and Wellness

Texas A&M University recognizes that mental health and wellness are critical factors that influence a student's academic success and overall wellbeing. Students are encouraged to engage in proper self-care by utilizing the resources and services available from Counseling & Psychological Services (CAPS). Students who need someone to talk to can call the TAMU Helpline (979-845-2700) from 4:00 p.m. to 8:00 a.m. weekdays and 24 hours on weekends. 24-hour emergency help is also available through the National Suicide Prevention Hotline (800-273-8255) or at suicidepreventionlifeline.org.

Statement on COVID-19

To help protect Aggieland and stop the spread of COVID-19, Texas A&M University urges students to be vaccinated and to wear masks in classrooms and all other academic facilities on campus, including labs. Doing so exemplifies the Aggie Core Values of respect, leadership, integrity, and selfless service by putting community concerns above individual preferences. COVID-19 vaccines and masking — regardless of vaccination status — have been shown to be safe and effective at reducing spread to others, infection, hospitalization, and death.