USC University of Southern California

CSCI 566: Deep Learning and its Applications

Spring 2024—Fridays—1:00pm-4:20pm (4)

Location: THH 201

Course Description
This course offers a comprehensive introduction to the principles of machine learning (ML) and deep learning (DL), emphasizing both mathematical foundations and practical applications. You will gain insights into basic ML techniques, learn knowledge of advanced DL applications in fields like computer vision and natural language processing, and understand their transformative impact on areas such as image recognition and autonomous systems. The course includes hands-on assignments and a customizable final project, providing you with practical experience in implementing ML and DL solutions.

Course website: https://usc-asap.github.io/CICS566-S24/ (see up-to-date course schedule and TA office hours)
Course Piazza: https://piazza.com/usc/spring2024/csci566/info
Course Google Drive (mostly for slides; nothing there yet):
https://drive.google.com/drive/folders/1leqG80jRFPqsa52s758rETWbEg_3XRJa?usp=drive_link
(use your USC login – do not request access from personal Gmail)

Special Notice: This course is currently evolving. For Fall 2024, expect a refocused curriculum with an increased emphasis on deep learning in CSCI 567 and more ML foundation in CSCI 566, though the Spring 2024 semester will continue to cover a broader range of ML and DL topics.

Guest Lectures: Industry and academic professionals will join our lectures regularly, sharing their experiences in ML and data science, and providing career insights.

Prerequisite(s):

1. Proficiency in Python
2. College Calculus, Linear Algebra
3. Probability and Statistics
4. (Recommended) CSCI 567 (Machine Learning) or equivalent

Recommended Preparation: sufficient mathematical background; good programming skills; familiarity with concepts and methods in machine learning and AI.

Instructor: Yue Zhao
Office: PHE 432
Office Hours: TBA
Purpose: Office hours are intended for non-homework and project-related inquiries. Questions about course content and logistics are encouraged.

Teaching Assistants: Office hours and details are available on the course website: https://usc-asap.github.io/CICS566-S24/
Purpose: TA sessions focus on homework, project guidance, and administrative support

Contact: All course communications should be directed through Piazza for efficiency and inclusivity. We prioritize this platform over email due to the course's large size.

Required Readings and Supplementary Materials

- Deep Learning (MIT Press) by Ian Goodfellow, Yoshua Bengio, and Aaron Courville.
    - A free online version is available at http://www.deeplearningbook.org/
- Mathematics for Machine Learning (Cambridge University Press) by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong.
    - A free online version is available at https://mml-book.github.io/

Description and Assessment of Assignments
3 in-class quizzes, 2 coding assignments, 1 midterm exam, and 1-course project (a proposal, reports, presentations, etc.)

Grading Breakdown

Deliverable	**Points of the total grade**
In-class Quizzes	10 (5 each)
Assignments	30 (15 each)
Midterm	25
Course Project	40
TOTAL	**105**

In-class Quizzes
Quizzes will typically be graded on a completion basis, and the results will be reviewed together during the same class session to facilitate a collective discussion of the answers. Please note that quizzes must be completed within the designated class period, and they also function as a means of taking attendance. **No remote quizzes or makeup quizzes**.

Assignments
These assignments will be individual coding tasks. Your submissions should include both the written code and accompanying documentation. Typically, you will be given a starter code and a detailed description of the required implementation, along with some unit tests to check the compatibility of your code with the grading scripts.

Midterm Exam

- Format: Open book (no electronic devices allowed) with mostly multiple-choice questions and some computation-based questions (with unambiguous, single-number answers).
- Content: Covers material from Module 1.
- Makeup exams will not be provided.

Project (more information to come)
The project should be in groups of 4-6 students; the topic is open. This can be from applied ML projects to some rigorous ML research projects (which lead to publications; note this is an applied course which means this is rare).

Each group will have a dedicated TA, who will meet with you for two times to make sure the project is flowing smoothly.

- *Project Preference Form (5%).* This form is designed to assist with the team-matching process for the course research project. Students will indicate their project preferences and potential teammates. We will use Piazza to help this process
- *Pre-proposal (5%).* The pre-proposal is a short, 1-page document confirming your team and providing initial thoughts on your research project. This includes the dataset you plan to use.
- *Proposal and Literature Review (20%).* The proposal report should present initial research ideas and summarize them. It should provide an overview of the dataset and main research tasks. A detailed literature review, including recent papers and models related to the dataset, research tasks, and new ideas, is an essential part of the proposal.
- *Midterm Report (20%).* The midterm report should summarize current research progress and initial results. It should also present an updated list of research ideas that the team plans to explore.
- *Final Report and Presentation (50%).* The final report should resemble a research paper, motivating the problem, presenting novel approaches, describing the experiment, and discussing results.

Assignment Submission Policy

All assignments need to be submitted in an electronic form by **11:59pm PST** of the due date. There are NO late days for assignments.

Grading Scale

A	93-100
A-	90-92
B+	87-89
B	83-86
B-	80-82
C+	77-79
C	73-76
C-	70-72
D+	67-69
D	63-66
D-	60-62
F	59 and below

Letter grades decided by rounding floating point grades up to the nearest whole number (e.g., 92.2 -> A; 59.8 -> D-).

Placeholder Schedule (this will change, and I am still revising it)

	Topics/Daily Activities	Readings and Homework	Deliverable/ Due Dates
Week 1 Jan 12	1. Course Introduction 2. Framing ML Problems 3. Shape the Course Together		
Module 1: ML and DL Foundations			
Week 2 Jan 19	1. ML as Function Approximation 2. Linear Models ***Guest Discussion**: Life in Data Science Startup		
Week 3 Jan 26	1. Mathematical Foundations Review (TA) 2. Classical ML & Overfitting Revisited • Decision Trees • k-Nearest Neighbors • Clustering • Anomaly Detection ***Guest Discussion**: Life in Risk modeling		
Week 4 Feb 2	Neural Network Basics • Perceptron Revisited • Gradient Descent • Forward Propagation	**Quiz 1**	**Course Project Teams Formed** **Pre-prosal DUE**
Week 5 Feb 9	1. Neural Network Continued: • Backpropagations 2. Cloud computing service tutorial (TA) 3. Deep Learning Software Tutorial (TA)	**Assignment 1 OUT**	
Week 6 Feb 16	**MIDTERM EXAM**		**Proposal and Literature Review DUE**
Week 7 Feb 23	Different Types of Neural Networks: 1. Convolutional Neural Networks 2. Recurrent Neural Networks 3. Graph Neural Networks		

Module 2 Deep Learning Applications			
Week 8 Mar 1	1. Transformers and AutoDiff 2. Deep Generative Models 3. Deep Reinforcement Learning		**Assignment 1 DUE**
Week 9 Mar 8	1. Deep Learning for Natural Language Processing *** Guest discussion**: life in the NLP industry	**Assignment 2 OUT**	
Week 10 Mar 15	**NO CLASS; Spring Recess**		
Week 11 Mar 22	1. Deep Learning for Computer Vision + midterm debrief *** Guest discussion**: life in the CV industry	**Quiz 2**	**Course Project Mid-report DUE**
Module 3: Advanced Topics in Machine Learning and Deep Learning **Advanced topics are all subject to change**			
Week 12 Mar 29	1. Advanced Graph Neural Networks 2. ML Scalability and Systems 3. ML Advice?		**Assignment 2 DUE**
Week 13 Apr 5	1. Automated ML and Model Selection 2. Foundational Models, Pretraining, and In-context Learning *** Guest discussion**: content to be decided		
Week 14 Apr 12	**Team Project Presentations**		
Week 15 Apr 19	**Team Project Presentations**		
Week 16 Apr 26	**Team Project Presentations**		

FINAL	**Final Report Due on University Final Exam Day, Wed May 3**	**(No in-class Exam)**	**Final Project Report due**

Statement on Academic Conduct and Support Systems

Academic Conduct

Plagiarism – presenting someone else's ideas as your own, either verbatim or recast in your own words – is a serious academic offense with serious consequences. Please familiarize yourself with the discussion of plagiarism in *SCampus* in Section 11, *Behavior Violating University Standards* https://scampus.usc.edu/1100-behavior-violating-university-standards-and-appropriate-sanctions. Other forms of academic dishonesty are equally unacceptable. See additional information in *Campus* and university policies on scientific misconduct, http://policy.usc.edu/scientific-misconduct.

Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the *Office of Equity and Diversity* http://equity.usc.edu or to the *Department of Public Safety* http://capsnet.usc.edu/department/department-public-safety/online-forms/contact-us. This is important for the safety of the whole USC community. Another member of the university community – such as a friend, classmate, advisor, or faculty member – can help initiate the report, or can initiate the report on behalf of another person. *The Center for Women and Men* http://www.usc.edu/student-affairs/cwm/ provides 24/7 confidential support, and the sexual assault resource center webpage http://sarc.usc.edu describes reporting options and other resources.

Support Systems

A number of USC's schools provide support for students who need help with scholarly writing. Check with your advisor or program staff to find out more. Students whose primary language is not English should check with the *American Language Institute* http://dornsife.usc.edu/ali, which sponsors courses and workshops specifically for international graduate students. *The Office of Disability Services and Programs* http://sait.usc.edu/academicsupport/centerprograms/dsp/home_index.html provides certification for students with disabilities and helps arrange the relevant accommodations. If an officially declared emergency makes travel to campus infeasible, *USC Emergency Information*

http://emergency.usc.edu will provide safety and other updates, including ways in which instruction will be continued by means of blackboard, teleconferencing, and other technology.