



USC Viterbi School of Engineering

Note: Due to the high demand, the course will have an **entrance exam**.

1. D-clearance will be cleared **strictly** based on students' exam results (and major for tie-break).

CSCI 566: Deep Learning and its Applications

Spring 2022—Monday—6:00pm-9:20pm (Units: 4)

Location: **GFS** 106

Course Description

Recently, deep learning has advanced many AI-related problems from image retrieval, video analysis, to natural language processing, to self-driving, and to medical applications. Our goal is to guide students to get familiar with these recent cutting-edge deep learning (DL) advances in computer vision and natural language processing. Through this course, students will gain a basic understanding of DL algorithms, and how to set up and solve problems involving deep learning techniques. The course will include a couple of practical assignments and a final course project. For the final course project, students will be encouraged to pick their own topics, but can also select from a provided list of projects. Check out our course website for the lecture materials: <https://csci566-spring2022.github.io/> Students will be invited to the course Piazza once they received the clearance: <https://piazza.com/usc/spring2022/csci566>

Prerequisite(s):

1. Proficiency in Python
2. College Calculus, Linear Algebra
3. Probability and Statistics
4. Equivalent knowledge of CSCI 567 (Machine Learning)

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

Instructor: Xiang Ren

Office: Zoom (<https://usc.zoom.us/j/91954035563>)

Office Hours: Wednesday 11:00am-12:00pm (for assignment related meetings, please go to TA OHs)

Teaching Assistants:

- Aaron Chan
- Bingjie Tang
- Nathan Dennler
- Shariq Iqbal

TA Office Hours (@ Zoom):

Tuesday 4-5pm

- Shariq (Odd Weeks)
 - <https://usc.zoom.us/j/98618243144?pwd=SkFzTjd2KzZmNFhCVGM2eEt0U3doQT09>
- Aaron (Even Weeks)
 - <https://usc.zoom.us/j/93216623193?pwd=cmk3TEZDb3lK1lvUHViejhrN1pJUT09>

Thursday 4-5pm

- Nathan (Weeks 1-9)
 - <https://usc.zoom.us/my/dennler>
- Bingjie (Weeks 10-17)
 - <https://usc.zoom.us/j/2396895777?pwd=SnVpcUt1NWVQdzNSYzMrWWRsaytNZz09>

Contact: Please **ONLY** use Piazza for any communication. This term we will be using Piazza for class discussion and questions. The system is highly catered to getting you to learn fast and efficiently from classmates, the TA, and myself. Rather than emailing questions to the teaching staff, I encourage you to post your questions on Piazza. Find our class signup link at: <https://piazza.com/usc/spring2022/csci566>

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/>
- (Supplementary) Deep Neural Methods for NLP (Goldberg)

Description and Assessment of Assignments

Pop-up Quizzes, 2 Coding Assignments, 1 Midterm exam, 1 research paper presentation, and 1 Course project (a proposal, reports, presentations, etc.)

Grading Breakdown

Extra credit: participation (5 points)

Entry	Points of the total grade
Participation (Pop-up Quiz)	5
Assignment #1	10
Assignment #2	10
Paper presentation	10
Midterm	20
Course Project	50
TOTAL	105

Project grading breakdown:

Entry	Points of the grade
Project Proposal & Teaming	2
Project Pitch	1
Project survey	5
2 meetings with TA	2
Project Mid-report	12
Final Presentation	8
Final Report	20
TOTAL	50

Project Requirement

Students are expected to form 4-people teams and each team is asked to come up with a project idea, related to deep learning for AI applications. The project is expected to either (1) introduces new ideas to improve deep learning techniques for existing AI applications (e.g., computer vision, NLP, or Robotics), with a clear piece of technical contribution/innovation, or (2) tackles a new and interesting problem setting (in a new domain) that is motivated by the unique problems posed by the application domain, using deep learning techniques, or (3) propose a new task formulation in the scope of deep learning and create resources for studying the new task, or (4) conduct a systematic analysis on issues of existing deep learning models for a task. Students will be encouraged to pick their own topics, but can also select from a provided list of project topics. They will conduct a literature review on the topic and generate a survey paper, implement baseline models, present project method and results during the class, and generate a final project report based on the progress.

Project teaming. Finalize the 4-people project teams and inform the TAs via filling up the online form (one submission for each team).

Project proposal. 5-6 slides on problem (scope & definition) + "today" (status of literature) + challenges + directions to innovate. We will share a google slide deck to collect the slides.

Project pitch (in class). Each project team will spend 4 mins going through 2-3 slides to introduce about their project proposal, focusing on the problem statement, main challenges to address, and directions to tackle the problem.

Project survey report. A literature survey (e.g., "Related Work" section in a conference paper) of the relevant studies for your proposed problem/topic. At this point, you should have a clear picture on whether people have done the problem you proposed and to what stage. If you find your proposed problem has been well addressed already, you should tweak your project to claim something novel; Otherwise, make it clear in your survey that: (1) what have been done related to your proposal (an organized view); (2) what are the limitations or challenges remained to be solved. **Format:** a 1.5-2 page (double-column) literature review write-up. ([latex template](#))

Project mid-term report. The project midterm report should provide details about the set of initial experiments you have done for the project. This should involve reproducing the results of a state-of-the-art baseline model for the task of interest with code that you have implemented, or a pilot version of your proposed approach. The report should also talk about the specific dataset(s) you choose to use, the evaluation protocol and metric you decide to have, and the experiment settings. In your report, also perform an analysis of what remaining errors this model makes and describe how you plan to create a new model for the final project that will address these errors (i.e., improve over this version). **Format:** a 3-page (double-column, excluding references) report; will be graded by instructor and TAs. ([latex template](#))

Project final presentation. A 7-min slide-based project presentation to briefly overview your proposed approach, the results obtained so far, and analysis as well as interesting findings. **Format:** A 7-min long PPT presentation in class. We will share a google slide deck to collect the slides. The project presentation will be graded by all TAs and the instructor.

Project final report. The final report will be expected to present a novel contribution that either (1) introduces new ideas to improve deep learning techniques for existing AI applications (e.g., computer vision, NLP, or Robotics), with a clear piece of technical contribution/innovation, or (2) tackles a new and interesting problem setting (in a new domain) that is motivated by the unique problems posed by the application domain, using deep learning techniques, or (3) propose a new task formulation in the scope of deep learning and create resources for studying the new task, or (4) conduct a systematic analysis on issues of existing deep learning models for a task. **Format:** an 8-page (double-column, excluding references) final report for the project. ([latex template](#))

Paper Presentation Requirement

Each project team will spend **two** 10-min slide presentations, each covering one research paper that is most related/significant for your project. The presentation is expected to introduce the paper in an accessible way to the whole class, which should include: problem/background, brief summary of related work for it, its contributions, method/proposed work, and experiments/results. Each presentation can be done by 1 or 2 students together (by splitting the presentation content between the two).

Slides of the presentations will be uploaded to a google slide deck in advance. The presentation will be graded by the instructor in terms of its content quality, the accessibility of the slides, and presentation quality.

Assignment Submission Policy

All assignments need to be submitted in an electronic form by **11:59pm PST** of the due date.

Grading Scale

A	95-105
A-	90-94
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

Grading Timeline

Assignments will be graded within one week after the due date.

Additional Policies

There will be a total of **5 late days** for the entire class, to be used in **integer amounts** and distributed as the student sees fit. Any exception needs to be discussed within the first 2 weeks of the semester (no

exception otherwise). Note that there is no late day for the project final report (or any project-related in-class presentation).

Course Schedule: Weekly Breakdown (note: schedule may be up to adjustment)

	Topics/Daily Activities	Readings and Homework	Deliverable/ Due Dates
Week 1 1/10	Course Introduction / Applications of Deep Learning / Entrance exam		
Module 1: Neural Network Basics			
Week 2 1/17	No class - MLK Day		
Week 3 1/24	Machine Learning 101 + Loss functions and Optimization		
Week 4 1/31	Neural Network Overview + Convolutional Neural Networks		Course Project Team
Week 5 2/7	Training Neural Networks + Deep Learning Software Tutorial (by TA) + Cloud computing service tutorial (by TA)	Assignment 1 OUT	
Week 6 2/14	Recurrent Neural Networks		Course Project Proposal
Module 2 Deep Learning Applications			
Week 7 2/21	No class - President's Day	Online midterm exam in this week, time TBD	Assignment 1
Week 8 2/28	Deep Learning for Computer Vision + project pitch		
Week 9 3/7	Deep Learning for Natural Language Processing + midterm re-cap	Assignment 2 OUT	Project survey

Week 10 3/14	No class - spring recess		
Week 11 3/21	Reinforcement Learning 101 + Deep Learning for Graphs		Assignment 2
Module 3: Advanced Topics in Deep Learning Advanced topics are all subject to change			
Week 12 3/28	Attention, Memory, Relation Networks + 5 paper presentations		Course Project Mid-report
Week 13 4/4	Generative models & language models + 5 paper presentations		
Week 14 4/11	Self-supervised learning + 5 paper presentations		
Week 15 4/18	Transfer learning / Multi-task learning + 5 paper presentations		
Week 16 4/25	Multi-modal learning + 5 paper presentations		
Week 17 5/2	Team Project Presentation (3.5-4 hours)		Final report due May 8 11:59 pm PST
FINAL	No Final		

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/academicssupport/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.