



USC Viterbi School of Engineering

CSCI 566: Deep Learning and its Applications

Spring 2023 — Fridays — 1:00pm-4:20pm (Units: 4)

Location: [THH](#) 201

For the most up-to-date version of *this document*, go here:

https://docs.google.com/document/d/1iFUMaC8f_8Thbcj9mjoQvebpDDYTjxj-zAWU0C7YoGc/edit?usp=sharing

Course Description

Recently, deep learning has advanced many AI-related problems: image retrieval, video analysis, natural language processing, self-driving, medical applications, and more. 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.

Course website: <https://csci566-spring2023.github.io/>

Course Piazza: <https://piazza.com/usc/spring2023/csci566/info>

Course gDrive:

https://drive.google.com/drive/folders/11YcbSZcJUdRSbrNBWiOs_HM2VkxdNKWS?usp=sharing

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: Jesse Thomason

Office: SAL 244

Office Hours: Mondays, 10am-11am

Purpose: Non-homework questions (e.g., about projects, research ideas, etc.); questions about course material; general advising, etc.

Teaching Assistants and Office Hours: Consult the course website for the most up-to-date OH for TAs: <https://csci566-spring2023.github.io>

Purpose: Homework and project questions; project advising meetings; course admin

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.

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) Neural Network Methods in Natural Language Processing (Goldberg)
 - <https://www.amazon.com/Language-Processing-Synthesis-Lectures-Technologies/dp/1627052984>

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)

Deliverable	Points of the total grade
Pop-up Quizzes	5
Assignment #1	10
Assignment #2	10

Paper Roleplaying Breakout	10
Midterm	20
Course Project	45
TOTAL	100

Pop-up Quizzes

These quizzes will “pop up” in lecture. They will, generally, be graded just for completion, with results gone over in the same class period to discuss the answers collectively. They must be completed *during the class period*.

Assignments #1 and #2

These will be coding assignments you will complete *individually*. Your submissions will be a combination of written code and documentation. Generally, you can expect to be provided with starter code and a description of what’s to be implemented, as well as some unit tests to ensure components of your code are compatible with the scripts to be used for grading.

Midterm Exam

The midterm will be an open-book, multiple choice exam covering the material in Module 1. Note that there is no corresponding final exam; evaluation in the second half of the course will center around your team research projects.

Paper Roleplaying Breakout

Students will take on an active Role in a Paper Roleplaying session lasting about 90 minutes (held at the end of classes in Module 3). The general idea of these sessions is that groups of 10-15 students will gather together with an administering TA to discuss a single paper in detail. Students can attend as observers or take on an active Role [again, each student will take on one Role during one breakout session; that’s the graded part]. Roles include things like: creating a 2-page summary of the paper highlighting the main points and presenting that summary verbally to your team; acting as a “reviewer” of the paper as though it were a conference submission; an industry practitioner aiming to monetize aspects of the paper; a hacker trying to reimplement the methods in the paper; being a teacher who scopes out interesting discussion questions to ask the group about the paper and leads that discussion; or a societal impact assessment officer who studies how the results in the paper could lead to help or harm for different groups of people. For your breakout and chosen Role, you will turn in a 2 page document based on your Role and will spend some time during the breakout session presenting the arguments or case of your Role to the group and TA. If you’re curious to get an idea in advance, the original ideas for these seminar breakouts are detailed [here](#). Grading will be broken down as:

- [4pts] Oral presentation of your role writeup/findings with respect to the paper

- [3pts] Engagement with other folks as they play their roles and discuss their writeups/findings
- [3pts] 2 page writeup/findings turned in on gDrive; grading depends on satisfaction of the description of the role

Project grading breakdown:

Project Deliverable	Points of the project 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	15
TOTAL	45

Project Requirement

Students are expected to form 4-person 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) introduce 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) tackle 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 methods and results during the class, and write a final project report based on the progress.

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

Project proposal. Your project proposals should be prepared in the form 10 or fewer slides. Your slides should answer the following questions.

Q1: What will your project aim to do? Articulate your objectives using absolutely no jargon.

Q2: What is new in your approach and why do you think it will be successful? You will need to highlight briefly the existing work in the space to make your innovation clear.

Q3: Who cares? If you are successful, what difference will it make? Think about both positive (e.g., helping others, advancing research) and negative (e.g., harming others, advancing nefarious actors) outcomes.

Q4: What are the risks? What could go wrong, and how will you pivot? You should identify one or two places your expectations for the project might not be met in the first few weeks that you would have time to recover from, and what your recover strategy will be.

Q5: How much will it cost? Estimate the CPU and GPU computation time (e.g., in CARC units and GCP credit).

Q6: Who will do what? Outline your expectations for your team to hold yourselves accountable to one another and to us.

Q7: Thinking ahead, what is your expected timeline? What do you hope to turn in by the midterm report? The final report?

We will share a gForm to turn in slides.

You will condense these longer, more detailed slides into a single lightning pitch slide for the Project Pitch assignment.

Project pitch (in class). Each project team will spend 60 seconds doing a *lightning pitch* (*elevator pitch*) of two slides. The first slide must contain only the team name and working project title [team name is fixed; project title can continue to evolve]. The second slide will contain the project content. The content slide should state:

- What your project aims to do, articulating your objectives using absolutely no jargon.
- What is done currently, and how does that feed into what is new in your approach?
- What is your key insight to move past what is done currently / has been done before?

The content slide can contain animations. For the project pitch, teams will elect one speaker each, but the whole team will receive the same grade for the pitch. Pitches will be done via a shared Google Slides deck; a link will be shared via Piazza to collect pitch slides. The 60 second time limit will be *strictly enforced* via the embarrassing “clap-off-the-stage method.” It is okay to finish earlier than your 60 seconds.

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. If you find your proposed problem has been well addressed already, you should tweak your project to claim something novel. Make it clear in your survey: (1) what has been done related to your proposal (an organized view); and (2) what are the limitations or challenges remaining to be solved. **Format:** 1.5-2 content pages (double-column; unlimited space for references) literature review write-up. ([latex template](#)) After these 2 pages and any number of pages containing only references, you should include a single page summary of contributions from each team member on your project; team members are expected to shoulder roughly equal amounts of work for projects. Spell out your individual contributions clearly on that final page.

Project mid-term report. The project midterm report should provide details about the set of initial experiments you have done for the project. This report 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](#)) After these 3 pages, you should include a single page summary of contributions from each team member on your project; team members are expected to shoulder roughly equal amounts of work for projects. Spell out your individual contributions clearly on that final page. *Suggested Outline* [you may deviate from this]:

- Sec 1. Introduction (0.5 page)
- Sec 2. Problem statement/definition, notations, etc. (~0.5 page)
- Sec 3. Description of (tentative) solution (~0.7 page)
- Sec 4. Experiment settings: dataset, evaluation procedure/metrics, baselines, other implementation details (0.5 page)
- Sec 5. Results and discussion (0.8 page)

Grading Breakdown:

- Substance of work (35%)
- Novelty & contribution (20%)
- Experiments (30%)
- Presentation & writing (15%)

Project final presentation. A 5-min slide-based project presentation to briefly overview your proposed approach, the results obtained so far, and analysis as well as interesting findings. You should cover the following material:

- What your project aimed to do, articulating your objectives using absolutely no jargon.
- What was done before and what your key insight was to go beyond those approaches in your course project.
- Your findings, each presented as: a hypothesis, the methodology you used to evaluate the hypothesis, your experiments and results, and a brief discussion of whether your findings support the hypothesis. [You may have one central such finding or multiple; it will depend on the project. You should cover at least one in the 5min prez.]
 - The bulk of your presentation time should be here.
- What next steps you'd like to accomplish immediately, if any, before the final report deadline.
- What next steps your work leaves open for future research.

Format: A 5-min long slide-based presentation in class. The project presentation will be graded by the instructor. Additionally, you will submit one multiple choice question about your project a week before your presentation date.

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) proposes a new task formulation in the scope of deep learning and create resources for studying the new task; or (4) conducts 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](#)). After these 8 pages, you should include a single page summary of contributions from each team member on your project; team members are expected to shoulder roughly equal amounts of work for projects. Spell out your individual contributions clearly on that final page.

Grading Breakdown:

- Substance of work (35%)
- Novelty & contribution (20%)
- Experiments (30%)
- Presentation & writing (15%)

Assignment Submission Policy

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

Grading Scale

A	93-100
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

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

Grading Timeline

Assignments will be graded within two weeks after the due date, mostly.

Additional Policies

Late assignments will have their total grade reduced by 5% for every day late they are turned in. Each student will have **5 late day tokens** 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). Late day tokens *cannot* be redeemed for in-class presentations. Late day tokens *cannot* be redeemed for any other purpose than removing 5% per-day penalties on late assignments. Late day tokens may be used by a subset of group members for project deliverables that are turned in late (e.g., members who do not utilize tokens would receive late day penalties, while those that redeem tokens would not; that is to say, late day tokens are all *personal*, not group-level).

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

	Topics/Daily Activities	Readings and Homework	Deliverable/ Due Dates
Week 1 Jan 13	Course Introduction / Applications of Deep Learning		
Module 1: Neural Network Basics			
Week 2 Jan 20	Framing Problems for ML + Loss functions and Optimization		
Week 3 Jan 27	NO CLASS		
Week 4 Feb 3	Training Neural Networks + Convolutional Neural Networks		Course Project Teams Formed
Week 5 Feb 10	Activations, Initializations, Optimization, and Regularization + Cloud computing service tutorial (Gautam Salhotra & Bingjie Tang) + Deep Learning Software Tutorial (Deqing Fu & Tejas Srinivasan)	Assignment 1 OUT	
Week 6 Feb 17	Recurrent Neural Networks		Course Project Proposal
Week 7 Feb 24	MIDTERM EXAM		Assignment 1 DUE Monday Feb 27th, 11:59pm

Module 2 Deep Learning Applications			
Week 8 Mar 3	Deep Learning for Natural Language Processing + project pitches		
Week 9 Mar 10	Transformer Networks and Deep Learning for Computer Vision + midterm debrief	Assignment 2 OUT	Project survey
Week 10 Mar 17	NO CLASS		
Week 11 Mar 24	Multimodal Deep Learning		
Module 3: Advanced Topics in Deep Learning Advanced topics are all subject to change			
Week 12 Mar 31	Deep Learning for Agents + paper breakouts		Course Project Mid-report DUE Apr 3, 11:59pm Assignment 2 DUE Apr 3, 11:59pm
Week 13 Apr 7	Deep Reinforcement Learning + paper breakouts		
Week 14 Apr 14	Deep Learning and the World + paper breakouts		
Week 15 Apr 21	Team Project Presentations		
Week 16 Apr 28	Team Project Presentations		

FINAL	Final Report Due on University Final Exam Day, Wed May 3	(No in-class Exam)	Final Project Report due Wed May 3 11:59pm
--------------	---	-------------------------------	---

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.