

SYLLABUS (18 Aug 2022)

# A Computational Introduction to Deep Learning

EE 541: Fall 2022 (2 units)

Machine learning using large datasets stands as one of the most transformative technologies of the 21st century. It enables reliable face and speech recognition, internet search and monetization, computer vision, and fully autonomous vehicles. Machine learning proficiency requires software skills as well as an understanding of the underlying mathematics and theoretical concepts. This class introduces important aspects of deep learning using a computation-first approach. It emphasizes using frameworks to solve reasonably well-defined machine learning problems. Two advanced courses provide a deeper study of mathematical concepts: EE 559 Machine Learning I: Supervised Methods and EE 641 Deep Learning Systems.

Instructor: Brandon Franzke  
Email: franzke@usc.edu  
Office: EEB 504B  
Zoom: meet: [998 5176 5591](https://meet.usc.edu/room/99851765591)  
code: 574987  
Hours: Monday: 17:00 - 18:00  
Wednesday: 12:00 - 14:00

## Lecture

Monday (section: 30799)  
15:00 - 16:50

## Attendance and Participation

This class is offered with **in-person enrollment ONLY**. Attendance is mandatory to all lectures. You are responsible for missed announcements or changes to the course schedule or assignments. Taping or recording lectures or discussions is strictly forbidden.

## Piazza

<https://piazza.com/usc/fall2022/ee541>

Piazza enables fast and efficient help from classmates and instructors. Use Piazza to post questions about course material, homeworks, and policies instead of emailing questions to the teaching staff.

## Canvas

<https://canvas.usc-ece.com>

Use Canvas to electronically submit your homework and view course grades. You will receive an email to register during the first week of classes. Contact Dr. Franzke with any technical issues.

## Autolab

<https://autolab.usc-ece.com>

Use Autolab to electronically submit programming portions of homework for “auto-grading”. You will receive an email to register during the first weeks of the course. Contact Dr. Franzke with technical issues.

### TAs and staff

TA:	Karkala Hegde (Shashank)	CP:	Aditya Anulekh Mantri
Email:	khegde@usc.edu	E-mail:	adityaan@usc.edu
Office:	PHE 320	Zoom:	meet: <a href="#">613 225 3283</a>
Zoom:	meet: <a href="#">724 621 7299</a>	Hours:	Tuesday: 11:00 - 12:00
Hours:	Thursday: 11:00 - 12:30		Friday: 11:00 - 12:00
	Friday: 15:00 - 16:30		
Grader:	Zeyu Liu		
E-mail:	liuzeyu@usc.edu		
Hours:	by appointment		

## Learning objectives

Upon completion of this course a student will be able to:

- Understand the fundamentals and implement linear regression, a linear classifier, and logistic regression.
- Apply common deep architectures such as multilayer perceptron and convolutional neural networks and understand where each is most applicable.
- Be proficient in Python programming, including loops, conditionals, lists, dictionaries, classes, and standard programmatic patterns.
- Be proficient in numerical Python using NumPy, SciPy, and matplotlib for design and analysis.
- Organize, store, and access datasets such as .npz files, .h5 files, pickle, and pandas.
- Understand the role of machine learning frameworks such as scikit-learn.
- Understand the role and use of deep-learning frameworks such as PyTorch.
- Use frameworks to train MLP, convolutional, and recurrent networks to solve machine learning problems.
- Be proficient with relevant computing and cloud computing resources such as:
  - linux command line interface and automation with shell scripts
  - version control software such as git
  - cloud GPUs to train deep networks and understand fundamentals of accelerated computing.

## Course materials

- “*Neural Networks and Deep Learning*”, Michael Nielson. (online: <http://neuralnetworksanddeeplearning.com>).
- “*Deep Learning with PyTorch*”, Eli Stevens, Luca Antiga, Thomas Viehmann, Manning, 2020. (online: <https://pytorch.org/assets/deep-learning/Deep-Learning-with-PyTorch.pdf>).
- “*Deep Learning*”, Ian Goodfellow, Yoshua Bengio, Aaron Courville, The MIT Press, 2016. (online: <http://www.deeplearningbook.org>).
- “*Python Programming And Numerical Methods: A Guide For Engineers And Scientists*”, Qingkai Kong, Timmy Siau, Alexandre Bayen, Elsevier, 2020. (online: <https://pythonnumericalmethods.berkeley.edu/notebooks/Index.html>).

## Course Outline (tentative)

	Topics/Daily Activities	Reading & Homework	Deliverables
Week 1 (22 Aug)	Machine Learning inventory. Configuring your Python environment.	<b>HW 1 assigned.</b>	
Week 2 (29 Aug)	Getting started with Python. Estimation and MMSE.	Supp: [4] Ch. 1.	HW 1 due (02 Sep).
Week 3 (05 Sep)	Numerical Python. Regression and maximum likelihood.	lecture slides, supp: [4] Ch. 2-7. <b>HW 2 assigned.</b>	
Week 4 (12 Sep)	Logistic regression and multi-layer perceptrons (MLPs)	lecture slides, supp: [4] Ch. 14-16, 25. <b>HW 3 assigned.</b>	HW 2 due (16 Sep).
Week 5 (19 Sep)	<b>Quiz #1 (weeks 1-4).</b>		
Week 6 (26 Sep)	MLP backpropagation (scalar detail)	supp: [1] Ch. 1, [2] Ch. 6, [3] Ch. 6. <b>HW 4 assigned.</b>	HW 3 due (26 Sep).
Week 7 (03 Oct)	MLP backpropagation (scalar, vector/tensor overview).	lecture slides.	
Week 8 (10 Oct)	PyTorch: Introduction	supp: [3] Ch. 7-8, [1] Ch. 3.	HW 4 due (10 Oct).
Week 9 (17 Oct)	PyTorch: Building MLPs	supp: [3] Ch. 7-8, [1] Ch. 3. <b>HW 5 assigned.</b>	
Week 10 (24 Oct)	PyTorch: Optimizing training		HW 5 due (28 Oct).
Week 11 (31 Oct)	<b>Quiz #2 (weeks 6-10), Project overview.</b>		
Week 12 (07 Nov)	Introduction to linux CLI and AWS	lecture notes.	
Week 13 (14 Nov)	Advanced data engineering, embedding and auto-encoding	lecture notes. <b>HW 6 assigned.</b>	
Week 14 (21 Nov)	PyTorch: Convolutional Neural Networks (CNN)	Supp: [2] Ch. 8, [3] Ch. 9.	Project proposal due.
Week 15 (28 Nov)	<b>Quiz #3 (weeks 12-14).</b>		HW 6 due (02 Dec).
(12 Dec)	<b>Project report due.</b>		

## Grading Procedure

### Homework

Homework is assigned every 1-2 weeks. Assignments will include computational and numerical programming problems and will encourage experimentation and curiosity. No late submissions for credit. Your total homework score sums your best homework scores (as a percentage) after removing the one lowest score (of minimum 50%). You may discuss homework problems with classmates but each student must do their own original work. Cheating warrants an F in the course. Turning in identical homework establishes a rebuttable presumption of cheating.

## Quizzes

Quizzes are short (60-minute) non-cumulative assessments that cover the most recent material (approximately 5-weeks). Quizzes highlight important concepts and methods. They test ability to apply major principles, demonstrate conceptual understanding, and may require writing snippets of Python code. They occur during weeks 5, 9, and 13 (tentative). You may use a single 8.5"x11" reference sheet (front and back OK). You may not use any additional resources. You are expected to bring a scientific (non-graphing) calculator. Any cheating may result in an "F" in the course and will be referred to Student Affairs for other penalties. Alternate arrangements will be considered only for valid medical or family emergency excuses (proof required).

## Final Project

This course culminates with a final project in lieu of a final exam. Teams of two students apply deep-learning to problem selected from a set of instructor-defined options. Each option will include a *complete* starter dataset. Teams must experiment and document network architecture search, hyper-parameter optimization, and dataset augmentation. Students should treat the final project as a multi-week in-depth homework assignment and work to integrate concepts from the entire semester. All projects must use PyTorch as the primary deep learning framework unless approved explicitly in writing by the instructor. But teams may use or integrate additional languages for tooling and support. The instructor may include additional requirements.

The final project score consists of the following three elements:

**Project proposal (10%):** restate the selected problem and include a preliminary discussion about methods, techniques, or proposed variations/modifications/extensions to the problem handout. Proposals should be written in anticipation of fulfilling final report requirements.

**Written project report (50%):** a properly referenced written report that incorporates all source code as well as links to any external code sources. The project report should summarize the topic, provide relevant background (theoretical or applied), timeline and contributions, and document challenges and future extensions. It should provide discussion sufficient for an uninformed expert to understand the problem, model and training decisions, and implementation. Teams should provide quantifiable metrics to justify engineering tradeoffs.

The report must include the following additional elements

- (1) a mini "literature survey" of work already done in your problem domain. Understand network architectures used in similar models,
- (2) experimental results with various models (both ML and DL) and performance comparison. Use different architectures and hyperparameters,
- (3) evidence of good data handling and processing, *i.e.*, maintain separate test, train, and validation sets.

**Source code (40%):** submitted to instructor by providing link to pull from github.

## Course Grade

Homework 45%  
Quizzes 30%  
Final Project 25%

A if 90 – 100 points  
B if 80 – 89 points  
C if 70 – 79 points  
D if 60 – 69 points  
F if 0 – 59 points

(“+” and “-” within approx. 3% of grade boundary)

### **Cheating**

Cheating is not tolerated on homework or exams. Penalty ranges from F on assignment or exam to F in course to recommended expulsion.

# 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-andappropriate-sanctions>. Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus 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/contactus>. 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/studentaffairs/cwm/> provides 24/7 confidential support, and the sexual assault resource center webpage <http://sarc.usc.edu> describes reporting options and other resources.

## Academic Integrity

Academic integrity is critical the assessment and evaluation we perform which leads to your grade. In general, all work should be your own and any sources used should be cited. Gray-areas occur when working in groups. Telling someone how to do the problem or showing your solution is a VIOLATION. Reviewing examples from class or other sources to help a fellow classmate understand a principle is fine and encouraged. All students are expected to understand and abide by these principles. SCampus, the Student Guidebook, contains the University Student Conduct Code in Section 10, while the recommended sanctions are located in Appendix A. Students will be referred to the Office of Student Judicial Affairs and Community Standards for further review, should there be any suspicion of academic dishonesty.

## 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.

## Academic Accommodations

Any student requiring academic accommodations based on a disability is required to register with Disability Services and Programs (DSP) each semester. A letter of verification for approved accommodations can be obtained from DSP. Please be sure the letter is delivered to me as early in the semester as possible. DSP is located in GFS 120 and is open 08:30 – 17:00, Monday through Friday. The phone number for DSP is (213) 740-0776.