SYLLABUS (v2)

A Computational Introduction to Deep Learning

EE 541: Fall 2021 (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 self-driving 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	
	code: 574987	
Hours:	Monday	12:00 – 13:30
	Wednesday	14:00 – 16:00 (remote)

Lecture

Monday (section: 30799)	15:00 – 16:50
Thursday (section: 30754)	14:00 – 15:50

Piazza

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.

https://piazza.com/usc/fall2021/ee541

Canvas *(replaces Blackboard)*

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 the instructor with any technical issues.

https://canvas.usc-ece.com

Autolab

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.

https://autolab.usc-ece.com

TAs and graders

TA:	Ruochen Yang		TA:	Panagiotis (Panos) Kyriakis
Zoom:	meet: 989 2791 1975		Zoom:	
	code: 589717		Office hours:	TBA
Office hours:	Tuesday	09:30 – 11:30	Email:	pkyriaki@usc.edu
Email:	ruocheny@usc.edu			

Grader: Himali Paresh Munshi
Office hours: by appointment
E-mail: hmunshi@usc.edu

Course materials

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

Corequisites

EE 503 and EE 510.

Recommended Preparation

Graduate standing with a typical undergraduate ECE background including the equivalent of EE 105 and EE 155. Some exposure to a high-level scripting language such as Python or Matlab.

Learning Objectives

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

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

Course Outline (tentative)

	Topics/Daily Activities	Readings and Homework	Deliverables / Due Dates

Week 1 (23 Aug)	**Setting up your Python environment**	Lecture slides Supp: [4] Ch. 1 HW 1 assigned	
Week 2 (30 Aug)	**Getting started with Python**	Lecture slides Supp: [4] Ch. 2-7 HW 2 assigned	HW 1 due
(06 Sep)	**No class: Labor Day, University holiday**		
Week 3 (13 Sep)	**Numerical Python**	Lecture slides Supp: [4] Ch. 14-16, 25 HW 3 assigned	HW 2 due
Week 4 (20 Sep)	**Regression and Classification**	Lecture slides	
Week 5 (27 Sep)	**PyTorch: Multilayer Perceptrons (MLP)**	Lecture slides Supp: [1], Ch. 1, [2] Ch. 6, [3] Ch. 6 HW 4 assigned	**QUIZ 1** HW 3 due
Week 6 (04 Oct)	**PyTorch: Optimizing training**	Lecture slides Supp: [3] Ch. 7-8, [1] Ch. 3 HW 5 assigned	HW 4 due
Week 7 (11 Oct)	**Introduction to Data Engineering**	Lecture slides HW 6 assigned	HW 5 due
Week 8 (18 Oct)	**PyTorch: Convolutional Neural Networks (CNN)**	Lecture slides Supp: [2] Ch. 8, [3] Ch. 9 HW 7 assigned	HW 6 due
Week 9 (25 Oct)	**Introduction to linux CLI and AWS**	Lecture slides HW 8 assigned	**QUIZ 2** HW 7 due
Week 10 (01 Nov)	**Introduction to linux CLI and AWS**	Lecture slides	
Week 11 (08 Nov)	**Generative adversarial networks (GAN)**	Lecture slides Supp: [3], Ch. 20	HW 8 due
Week 12 (15 Nov)	**Deep reinforcement learning (RL)**	Lecture slides HW 9 assigned	
Week 13 (22 Nov)	**PyTorch: Recurrent Neural Networks (RNN)**	Lecture slides Supp: [3], Ch. 10	
Week 14 (29 Nov)	**Final projects**		**QUIZ 3** HW 9 due
FINAL (13 Dec)	**Project reports and videos due**		

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.

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 (45 minute) non-cumulative tests 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 14 (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. Treat the final project as a multi-week in-depth homework assignment that integrates concepts from the entire semester. The teaching team will prepare three problem statements. Each problem will include a representative dataset to get you started. You must experiment and document network architecture search, hyper-parameter optimization, and dataset augmentation.

Example projects:

- Semantic classification of images of faces
    - Classify emotion from facial expressions – *e.g.*, sad, happy, bored, neutral.
- Clothing coordination
    - Design a CNN to match shoes, pants, shirts, accessories, etc.
- Language classification
    - Design an RNN to perform real-time audio language classification – *e.g.*, Hindi, Mandarin, and English.

You will submit two deliverables: (1) a properly referenced written report that includes all source code as well as links to any external code sources and (2) a short video describing your findings, methods, and analyses.

Course Grade

Homework	50% (drop lowest one)	A	if 90 – 100 points
Quizzes	25%	B	if 80 – 89 points
Final Project	25%	C	if 70 – 79 points
		D	if 60 – 69 points
		F	if 0 – 59 points
		("+" and "-" within approx. 2% of grade boundary)	

Attendance and Participation

Attendance is mandatory to all lectures and discussions. You are responsible for missed announcements or changes to the course schedule or assignments. Taping or recording lectures or discussions is strictly forbidden.

Cheating

Cheating is not tolerated on homework or exams. Penalty ranges from F on 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.