SYLLABUS

Deep Learning Systems

EE 599: Fall 2020 (4 units)

Neural networks for nonlinear regression, classification, reinforcement learning. Back-propagation learning for multilayer perceptrons, convolutional and recurrent networks. Applications in audio processing, vision, and autonomy.

Machine learning from large datasets is arguably the most transformative technology of the century. It enables reliable speech recognition, face recognition, internet search, computer vision, self-driving cars, and countless other technical marvels. In the last decade, deep learning has surpassed traditional EE methods for inference in many applications in which accurate models prove evasive, but data is plentiful.

The computer science community has advanced deep learning and created tools that bring the most powerful deep learning techniques to general engineering applications. This class embraces these tools but it frames deep learning practice in the context of other EE graduate classes and will develop skills at training neural networks using python-based packages and cloud computing resources.

Instructor:	Brandon Franzke	
Email:	franzke@usc.edu	
Office:	EEB 504B (remote)	
Zoom:	meet:	998 5176 5591
	code:	574987
Hours:	Mo	14:00 – 15:30
	Th	10:00 – 11:30

Lecture

Tuesday and Thursday (section: 31290)
15:00 – 16:50

Discussion

This course has an additional discussion section that provides deeper insight and facility with coding deep neural networks and presents applications of concepts from lecture. You may attend either session each week.

Section #1: Friday @ 11:00 – 11:50

Section #2: Friday @ 13:00 – 13:50

Piazza

Piazza gets you help fast and efficiently from classmates, the TAs, and me. I encourage you to post questions on Piazza rather than emailing questions to the teaching staff.

https://piazza.com/usc/fall2020/ee599/home

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 issues.

https://canvas.usc-ece.com

TAs and graders

TA:	Olaoluwa (Oliver) Adigun	TA:	Digbalay Bose
Zoom:	meet: 951 2677 9992 code: seek_help	Zoom:	meet: 932 3409 0397 code: 695410
Office hours:	Mo 15:00 – 17:00	Office hours:	We 15:00 – 17:00
Email:	adigun@usc.edu	Email:	dbose@usc.edu
Grader:	Karkala Hegde	Grader:	Ashwin Telagimathada Ravi
Office hours:	by appointment	Office hours:	by appointment
E-mail:	khegde@usc.edu	E-mail:	telagima@usc.edu

Course materials

"Deep Learning", Ian Goodfellow, Yoshua Bengio, Aaron Courville, The MIT Press, 2016. (online: http://www.deeplearningbook.org)

"Neural Networks and Deep Learning", Michael Nielson. (online: http://neuralnetworksanddeeplearning.com)

"Deep Learning with Python", Francois Chollet, Manning, 2018.

"Neural Networks and Deep Learning, A Textbook", Charu C. Aggarwal, Springer International Publishing, 2018. (usc access: https://rd.springer.com/book/10.1007/978-3-319-94463-0)

"Learning from Data, A Short Course", Yaser S. Abu-Mostafa, Malik Magdon-Ismail, Hsuan-Yien Lin, AMLbook.com.

Prerequisites

EE 503 and EE 510.

Learning Objectives

Upon completion of this course students will be able to:

1. Understand the relationship between inference methods based on statistical models such as regression and data-driven methods like neural networks.
2. Understand the relationship between common loss and regularization functions and maximum likelihood parameter estimation, maximum a posteriori probability parameter estimation, Kullback–Leibler divergence, and cross-entropy.
3. Understand stochastic gradient descent learning in neural networks and its relation to other methods such as least-mean squares (LMS) adaptive linear filtering and the EM algorithm.
4. Be able to derive back-propagation equations for dense, convolutional, and recurrent layers (back-propagation in time).
5. Understand the commonly used methods in deep learning: activations, optimizers, loss functions, regularization, drop-out, batch normalization, and down/up sampling layers.
6. Become proficient with good data systems engineering practices including a pipeline development for data acquisition, cleaning, augmentation, training, testing, and closed-loop refinement.
7. Understand the role of deep learning networks in modern speech, computer vision, and signal processing systems.
8. Appreciation of the importance of data, its quality, and its limitations.
9. Propose and complete a novel team-based project in deep learning.

Course Outline (tentative)

	Topics/Daily Activities	Readings and Homework	Deliverable/ Due Dates
Week 1 (24 Aug)	**Introduction and Motivation. Machine learning: approaches and EE topics**	Lecture slides	Tools setup
Week 2 (31 Aug)	**Estimation and Detection**	Lecture slides	HW 1 assigned
Week 3 (07 Sep)	**Regression and Classification**	Lecture slides	
Week 4 (14 Sep)	**Multi-layer perceptrons (MLP) and back-propagation (BP)**	Lecture slides	HW 1 due, HW 2 assigned
Week 5 (21 Sep)	**Loss functions, activations, parameters**	Lecture slides	
Week 6 (28 Sep)	**Application of MLP**	Lecture slides	HW 2 due, HW 3 assigned
Week 7 (05 Oct)	**Batch normalization, drop out, initialization, stochastic gradient descent (SGD)**	Lecture slides	
Week 8 (12 Oct)	**Review/Midterm**		HW 3 due
Week 9 (19 Oct)	**Working with Data**	Lecture slides	HW 4 assigned
Week 10 (26 Oct)	**Convolutional neural networks (CNN)**	Lecture Slides	
Week 11 (02 Nov)	**Recurrent neural networks (RNN)**	Lecture Slides	HW 4 due, HW 5 assigned
Week 12 (09 Nov)	**Deep reinforcement learning**	Lecture Slides	Initial Project Proposal due
Week 13 (16 Nov)	**Generative adversarial networks (GAN)**	Lecture Slides	Revised Project Proposal due HW 5 due
Week 14 (24 Nov)	**Application: Natural language processing (NLP)**	Lecture Slides Communication tips	Final project work.
FINAL	**Deep Learning Symposium**		Final project due

Grading Procedure

Homework

Homework will be assigned every 1-2 weeks. Problems will be a mix of analytical and computational problems. No late assignments for credit.

You may discuss homework problems with classmates but each student must do his or her own work. Cheating warrants an F in the course. Turning in identical homework establishes a rebuttable presumption of cheating.

Exam

Ensure that you can stream live video (mute or audio off) during the entirety of the work time for proctoring. You must make prior arrangements with me if that is not possible. 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 non-graphing scientific calculator. You must show how you arrived at your answers to receive full credit. Any cheating may result in an "F" in the course and will be referred to Student Affairs for other penalties. Make up exams will only be given for valid medical or family emergency excuses (proof required).

Final project

This course has a team-based final project In lieu of a final exam.

- Final project teams will be 2-3 students each. Ph.D. students may complete individual projects with approval.
- You will develop a novel topic with the help of your team members and guidance from instructors.
- Each team will be assigned a mentor. Mentors will be typically be the instructor or a TA, but in some cases may be Ph.D. students or faculty associated with the topic. Ph.D. students enrolled in this class may have their research advisors serve as their mentor in coordination with the instructor.
- Each team will produce the following with the percentage of the overall project grade shown:
    - Preliminary proposal: 5%, due week 12
    - Revised proposal (after mentor feedback): 15%, due week 13
    - Project presentation: 30%, during schedule final exam period
    - Project report: 40%, due approximately 1 week after presentations (depending on final exam schedules).
    - Project video (approximately 5-minute video to be posted to YouTube): 10% due one day after projects due.
- Example project proposals and reports will be provided.
- Final presentations will take place at an end-of-semester "Deep Learning Symposium" that will have parallel sessions chaired by the TAs. For schedules of symposiums from previous years, see: http://deeplearning.usc-ece.com.
- Your videos will be posted in a playlist on the DeepLearning USC YouTube channel.

Course Grade

Homework	35%
Exam	25%
Final Project	40%

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. 2% of grade boundary)

Attendance and Participation

Attendance is mandatory to all lectures and discussions. You are responsible for missed announcements and changes to the course schedule and assignments. Your attendance may be synchronous or asynchronous. Make synchronous attendance a priority. Per university guidance: you should plan to attend every synchronous session of this class regardless of when it occurs in your time zone. Some unreasonable hours (sessions outside 07:00 – 22:00) may preclude this general rule.

Synchronous class dynamics are improved substantially with visible participants in the class. Arrange to have your cameras on (default: "camera on, audio off") during synchronous online sessions. You must make prior arrangements with me if that is not possible.

USC policy requires that all classes conducted online be recorded for asynchronous viewing with transcriptions made available. These recordings are considered "educational records" subject to federal privacy laws (FERPA) as students may be personally identifiable in class recordings by voice, name, or image. Students are not permitted to create their own class recordings without prior written permission. Violations of these policies will be met with the appropriate disciplinary sanction.

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.