EE 599 Deep Learning Systems (31290)

Units: 4
Fall 2020, TTh 3:00 – 4:50

Location: TBD

Instructor: Brandon Franzke
Office: EEB 504B
Office Hours: TBA
Contact Info:
franzke@usc.edu

Teaching Assistant: TBD
Office: TBD
Office Hours: TBD
Contact Info: TBD

Course Catalogue Description

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.

Course Description

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.

Learning Objectives and Outcomes

- Understand the relationship between inference methods based on statistical models such as regression and data-driven methods like neural networks.
- 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.
- 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.
- Be able to derive back-propagation equations for dense, convolutional, and recurrent layers (back-propagation in time).
- Understand the commonly used methods in deep learning: activations, optimizers, loss functions, regularization, drop-out, batch normalization, and down/up sampling layers.
- Become proficient with good data systems engineering practices including a pipeline development for data acquisition, cleaning, augmentation, training, testing, and closed-loop refinement.
- Understand the role of deep learning networks in modern speech, computer vision, and signal processing systems.
- Appreciation of the importance of data, its quality, and its limitations.
- Propose and complete a novel team-based project in deep learning.

Prerequisite(s): : EE 503 and EE 510
Co-Requisite(s): None
Concurrent Enrollment: None

Course Notes

The course Piazza page will be the primary source of communication (link TBD). All homework assignments, lecture notes, and handouts will be distributed via Piazza and/or Canvas. GitHub will be used to distribute some code. Large datasets will be distributed via an ECE server.

Technological Proficiency and Hardware/Software Required

Access to a computer and familiarity with Python. Course will also utilize commercial cloud storage and computing platforms.

Required Readings and Supplementary Materials

Optional Textbooks and Resources:

- Ian Goodfellow, Yoshua Bengio, Aaron Courville. *Deep Learning*, The MIT Press, 2016.
  - available online: http://www.deeplearningbook.org
- Michael Nielson. *Neural Networks and Deep Learning*.
  - available online: http://neuralnetworksanddeeplearning.com
- Francois Chollet. *Deep Learning with Python*. Manning, 2018.
- Charu C. Aggarwal. *Neural Networks and Deep Learning, A Textbook*. Springer International Publishing, 2018.
  - available online (usc domain): https://rd.springer.com/book/10.1007/978-3-319-94463-0
- Yaser S. Abu-Mostafa, Malik Magdon-Ismail, Hsuan-Yien Lin. *Learning from Data, A Short Course*. AMLbook.com

In addition to this, helpful webpages, and code examples will be distributed. Concepts will be introduced through concrete example scripts which students will build upon. *Lecture slides include numerous links to research papers and on-line examples.*

Description and Assessment of Assignments

Homework will be assigned and collected every 1-2 weeks. Problems will be a mix of analytical and computational problems.

The course has one midterm and a team-based final project.

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.[1]
- Your videos will be posted in a playlist on the DeepLearning USC YouTube channel.

[1] This is the symposium from 2019 and 2020 offerings of a 4-unit EE599 special topics class. That 4-unit class is being split into two 2-unit classes. This is the second in that series and the projects are expected to be similar in scope and style to those shown from the EE599.

Grading Breakdown

Assignment	% of Grade
Homeworks	30
Midterm	30
Final Project	40
TOTAL	100

Assignment Submission Policy

Assignments will provide submission guidelines by the canvas webpage and must be received by 23:59 of the due date.

Grading Timeline

Homeworks will usually be returned by the time the next assignment is made. And partial solutions will be posted within 4 days of the due date.

Additional Policies

No late assignments for credit. Lecture attendance is not be recorded. If you do not attend lecture and do the homework, you will struggle very soon. It is imperative not fall behind. Respectful and professional conduct is expected at all times and in relation to your classmates, the instructor, and TAs.

Course Schedule: A Weekly Breakdown

	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

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 Part B, Section 11, "Behavior Violating University Standards" policy.usc.edu/scampus-part-b. Other forms of academic dishonesty are equally unacceptable. See additional information in SCampus and university policies on scientific misconduct, policy.usc.edu/scientific-misconduct.

Support Systems:

Student Health Counseling Services - (213) 740-7711 – 24/7 on call
engemannshc.usc.edu/counseling
Free and confidential mental health treatment for students, including short-term psychotherapy, group counseling, stress fitness workshops, and crisis intervention.

National Suicide Prevention Lifeline - 1 (800) 273-8255 – 24/7 on call
suicidepreventionlifeline.org
Free and confidential emotional support to people in suicidal crisis or emotional distress 24 hours a day, 7 days a week.

Relationship and Sexual Violence Prevention Services (RSVP) - (213) 740-4900 – 24/7 on call
engemannshc.usc.edu/rsvp
Free and confidential therapy services, workshops, and training for situations related to gender-based harm.

Office of Equity and Diversity (OED) | Title IX - (213) 740-5086
equity.usc.edu, titleix.usc.edu
Information about how to get help or help a survivor of harassment or discrimination, rights of protected classes, reporting options, and additional resources for students, faculty, staff, visitors, and applicants. The university prohibits discrimination or harassment based on the following protected characteristics: race, color, national origin, ancestry, religion, sex, gender, gender identity, gender expression, sexual orientation, age, physical disability, medical condition, mental disability, marital status, pregnancy, veteran status, genetic information, and any other characteristic which may be specified in applicable laws and governmental regulations.

Bias Assessment Response and Support - (213) 740-2421
studentaffairs.usc.edu/bias-assessment-response-support
Avenue to report incidents of bias, hate crimes, and microaggressions for appropriate investigation and response.

The Office of Disability Services and Programs - (213) 740-0776
dsp.usc.edu
Support and accommodations for students with disabilities. Services include assistance in providing readers/notetakers/interpreters, special accommodations for test taking needs, assistance with architectural barriers, assistive technology, and support for individual needs.

USC Support and Advocacy - (213) 821-4710
studentaffairs.usc.edu/ssa
Assists students and families in resolving complex personal, financial, and academic issues adversely affecting their success as a student.
Diversity at USC - (213) 740-2101
diversity.usc.edu
Information on events, programs and training, the Provost's Diversity and Inclusion Council, Diversity Liaisons for each academic school, chronology, participation, and various resources for students.

USC Emergency - UPC: (213) 740-4321, HSC: (323) 442-1000 – 24/7 on call
dps.usc.edu, emergency.usc.edu
Emergency assistance and avenue to report a crime. Latest updates regarding safety, including ways in which instruction will be continued if an officially declared emergency makes travel to campus infeasible.

USC Department of Public Safety - UPC: (213) 740-6000, HSC: (323) 442-120 – 24/7 on call
dps.usc.edu
Non-emergency assistance or information.