

SYLLABUS
Deep Learning Systems

EE 641: Spring 2023 (2 units)

Machine learning using large datasets is the most transformative technology of the 21st century. Recent advances in generative ML promise a applications to almost any imaginable problem. This course provides in-depth knowledge of deep learning systems theory and practice. It builds on prerequisite software skills from EE541: A Computational Introduction to Deep Learning and analytical skills from EE 559: Machine Learning I: Supervised Methods.

Lecture

Monday (section: 31225)
12:00 – 13:50

Enrollment is in-person ONLY. Attendance is mandatory to all lectures and all discussions. Taping or recording lectures or discussions is strictly forbidden without the instructor’s explicit written permission.

Instructor: Brandon Franzke
Email: franzke@usc.edu
Office: EEB 504B
Hours: Wednesday: 14:00 – 15:30 (in-person, remote)
Thursday: 12:00 – 14:00 (remote only)

Course materials

- “*Neural Networks and Deep Learning*”, Charu Aggarwal, Springer International, 2018. (online via SpringerLink).
- “*Deep Learning*”, Ian Goodfellow, Yoshua Bengio, Aaron Courville, The MIT Press, 2016. (online: <http://www.deeplearningbook.org>).
- “*Deep Learning Architectures, A Mathematical Approach*”, Ovidiu Calin, Springer International, 2020. (optional, online via SpringerLink).

Teaching assistants

| | | | |
|---------|--------------------------|------------|------------------|
| TA: | Karkala Hegde (Shashank) | Grader/CP: | Shoumik Nandi |
| Email: | khegde@usc.edu | Hours: | by appointment |
| Office: | PHE 320 | E-mail: | shoumikn@usc.edu |
| Hours: | Thursday 14:00 - 15:30 | | |
| | Friday 11:00 - 12:30 | | |

Piazza <https://piazza.com/usc/spring2023/ee641>

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

Electronically submit homework and view grades. You will receive a registration email during the first week of classes. Contact Dr. Franzke with technical issues.

Course Outline (tentative)

| | Topics | Homework | Deliverables |
|-------------------|--|----------------|---|
| Week 1 09 Jan | Introduction, PyTorch and deep learning review | HW 0 assigned. | |
| (16 Jan) | No class, Martin Luther King Day. | | |
| Week 2 23 Jan | CNN architectures, image segmentation | HW 1 assigned. | HW 0 due. |
| Week 3 30 Jan | Generative Adversarial Networks (GAN) | | |
| Week 4 06 Feb | Recurrent Neural Networks (RNN) | | HW 1 due. |
| Week 5 13 Feb | Sequence to sequence models (seq2seq), attention mechanisms | HW 2 assigned. | |
| (20 Feb) | No class, Presidents' Day. | | |
| Week 6 27 Feb | Transformers, self-attention, bidirectional encoders | | |
| Week 7 06 Mar | Exam #1 (weeks 1–6). AutoML, hyper-parameter tuning, and training optimization | | HW 2 due. |
| (13 Mar) | No class, Spring Break. | | |
| Week 8 20 Mar | Reinforcement learning (RL) | HW 3 assigned. | |
| Week 9 27 Mar | Reinforcement learning, deep Q-learning | | Preliminary proposal due (27 Mar). Proposal return (31 Mar). |
| Week 10 03 Apr | Project meetings | | HW 3 due. Revised proposal due (07 Apr). |
| Week 11 10 Apr | ML-Ops, pipelines, production deployments | | |
| Week 12 17 Apr | Diffusion and modern generative models | | Project status report due (21 Apr). |
| Week 13 24 Apr | Exam #2 (weeks 7–12). Explainability, bias and ethical AI | | |
| 05 May | Project presentations, 11:00 - 13:00 (mandatory). | | |
| 08 May | Project deliverables, due 12:00. | | |

Grading Procedure

Homework

Homework is assigned every 2-3 weeks. Assignments include analytic and programming problems and encourage experimentation and curiosity. No late submissions for credit. You may discuss homework problems with classmates but each student must submit their own original work. Cheating warrants an “F” on the assignment. Turning in identical homework establishes a rebuttable presumption of cheating.

Exams

Exams are non-cumulative and cover the most recent material (approximately 6-weeks). They test your ability to apply major principles, demonstrate conceptual understanding, and may require writing Python code. They occur during weeks 7 and 13 (tentative). You are expected to bring a scientific (non-graphing) calculator. You may use a single 8.5”x11” reference sheet (front and back OK). You may not

use any additional resources. 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 three students (teams of two with instructor approval) design and implement a deep-learning system to a self-identified problem. Students should treat the project as a platform to demonstrate mastery of problem specification, model selection, data analysis, testing, debugging, and results validation and analysis. The instructor will guide teams with difficulty identifying a suitable problem. Teams will prepare and present their approved project and show how it applies course concepts and deep learning best-practices. Attendance and participation during the project presentation session is mandatory.

Course Grade

| | | | |
|---------------|-----|---|--------------------|
| Homework | 35% | A | if 90 – 100 points |
| Exams | 30% | B | if 80 – 89 points |
| Final Project | 35% | C | if 70 – 79 points |
| | | D | if 60 – 69 points |
| | | F | if 0 – 59 points |

(“+” and “-” at $\approx 2.5\%$ 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.

Final Project

Requirements

Groups are encouraged to devise solutions to novel problems of personal interest to their background or research. But teams may select a problem with prior solutions provided their effort demonstrates novelty in addition to mastery of the course material. Groups may abstract problems from original context to fit within the project timeline and simplify constraints and scale. Projects must include sufficient mathematical and hypothetical complexity and include or extend substantive material from the course. All projects must obtain the instructor's written approval. All projects must use PyTorch as the primary deep learning framework unless approved explicitly in writing by the instructor. But projects may use additional languages for tooling and support. The instructor may provide additional requirements when introducing the final project assignment.

Scoring and Milestones

| | | |
|---|---------|-----|
| Topic Proposal (initial and revised) | week 9 | 14% |
| Status Report - data, training, integration | week 12 | 8% |
| Presentation and demo | final | 20% |
| Project report | | 25% |
| Model and source code | | 25% |
| Model Card | | 4% |
| Video | | 4% |

Project Deliverables

Topic proposal: describe the problem, prior related work, candidate data sets/sources, proposed technical approach, and expected outcomes. It should communicate that your topic is adequately prepared and it should outline immediate next steps. But the proposal is merely a guidepost and reasonable deviations in method, approach, and scope are expected.

Written report: summarize the topic, provide relevant background (theoretical or applied), timeline and contributions, and document challenges and extensions. It should provide discussion sufficient that an uninformed expert can understand the models, analytic decisions, outcomes, and implementation. Teams should provide quantifiable metrics to justify engineering tradeoffs.

Presentation: approximately 20 minute presentation to describe the project problem, your approach, and results. It should provide only what is necessary to understand the *what* and *why* and include minimal theoretical background.

Source code: submitted as a GitHub repository archive file (zip). It must include README file(s) that describe the repository structure, execution instructions, and special technical requirements. It should not include any training data or model files.

Video: a 4-minute video that describes the topic, your implementation, and your results. You may choose to upload this to a video sharing site such as YouTube but that is not required.

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.