SYLLABUS – EE 641: Fall 2026 (2 units)

A Computational Introduction to Deep Learning

Instructor:	Brandon Franzke	Office:	EEB 504B
Email:	franzke@usc.edu	Hours:	Wednesday: 13:00 - 14:00
			Thursday: 14:00 - 16:00

Machine learning using large datasets is the most transformative technology of the 21st century. Advances in generative ML promise solutions to almost any problem imaginable. 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 Wednesday (section: 30404) 15:00 – 16:50

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

Course resources

Course website https://ee641.usc-ece.com.

Gradescope https://gradescope.com. Electronically submit all homeworks. You will be automatically added to Gradescope during the first week of classes. Contact Dr. Franzke with technical issues.

Course materials

[1] *Neural Networks and Deep Learning*, Aggarwal, C., Springer International, 2018. online, via SpringerLink.

[2] *Deep Learning*, Goodfellow, I., Bengio, Y., Courville, A., The MIT Press, 2016. online, http://www.deeplearningbook.org.

[3] *Deep Learning Architectures, A Mathematical Approach*, Calin, O., Springer International, 2020. online, via SpringerLink.

NOTE: Texts are secondary to in-class lecture material and homework sets.

"AI" policy. You may use AI-powered tools in this course to enhance your learning and productivity. Use AI as a collaborative tool for understanding concepts, generating ideas, and troubleshooting. Approach AI-generated content critically and use it responsibly. Engage with AI as you would with a knowledgeable peer or tutor, using iterative conversations to deepen your understanding. You must attribute all AI-generated content in your work, including the prompts you used. You are fully accountable for the accuracy and appropriateness of any AI-assisted work. AI should supplement, not substitute, your own critical thinking and problem-solving. For assignments, you may use AI to clarify concepts or resolve issues, but submitted work must be your own. Submitting AI-generated work as your own without proper attribution or understanding is academic misconduct and will be treated as such.

You must develop complete mastery of all course material independent of AI assistance. Your knowledge and skills will be evaluated in contexts where AI tools are not accessible, mirroring real-world scenarios where you must rely solely on your own expertise. This ensures you can perform effectively in any situation, with or without AI support. Violations of this policy will result in severe academic penalties. The goal is to prepare you to use AI effectively in your future work while ensuring you develop a strong, self-reliant foundation in the course material.

"No technology" policy. Personal electronics are prohibited during lecture. Phones must be silenced and put away for the entire session. Laptops must remain closed. Slides are posted to the course website so you do not need a device to follow along. You will be asked to put away any device out in violation of this policy. Repeat offenders will be asked to leave for the day. **Exception:** a tablet lying flat on the desk, used with a stylus for handwritten notes or annotating course materials. A propped tablet counts as a laptop. Approved OSAS accommodations for assistive technology are always honored — email your letter as early as possible.

Course Outline

	Topics	Homework
Week 1 26 Aug	Introduction, PyTorch, and deep learning review. Advanced CNN architectures and applications.	**HW 1 assigned**.
Week 2 02 Sep	Object detection and segmentation.	**HW 2 assigned**.
Week 3 09 Sep	Multi-scale detection and spatial regression.	
Week 4 16 Sep	Generative Models: Energy based, Generative adversarial networks (GANs), and Variational autoencoders (VAEs).	**HW 3 assigned**.
Week 5 23 Sep	Generative models, continued. Recurrent Neural Networks (RNNs) and sequence-to-sequence models.	
Week 6 30 Sep	Attention mechanisms and self-attention.	**HW 4 assigned**.
Week 7 07 Oct	Transformers. Large language models (LLMs) and Mixture of experts (MoE).	
Week 8 14 Oct	Reinforcement learning principles. Project overview.	**HW 5 assigned**.
Week 9 21 Oct	Reinforcement learning applications.	
Week 10 28 Oct	**Exam (weeks 4–9)**.	**Draft project proposal due (02 Nov)**.
Week 11 04 Nov	Project proposal meetings.	**Revised project proposal due (06 Nov)**.
(11 Nov)	**No class, Veterans Day.**	
Week 12 18 Nov	Project meetings.	
(23 Nov)	**Status report due.**	
(25 Nov)	**No class, Thanksgiving Holiday.**	
Week 13 02 Dec	**Project presentations (mandatory) — EXTENDED TIME.**	
Sunday 06 Dec	**Project deliverables, due 23:59.**	

Grading Procedure

Homework (35%). Assignments include analytic and programming problems and encourage experimentation and curiosity. 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 substantively identical homework solutions counts as cheating.

Late homework is accepted with a 0.5% deduction per hour, up to 48-hours – **no exceptions**. Technical issues while submitting are not grounds for extension. No submissions will be accepted 48-hours after the due date. Graders score what is submitted and will not follow up if the file is incorrect, incomplete, or corrupt. It is your responsibility to ensure you submit the correct files and that they are accessible.

Exam (20%). The course has one exam, during week 10 (tentative). It covers weeks 4–9 — generative models through reinforcement learning; earlier material is exercised through homework. The exam tests your ability to apply major principles, demonstrate conceptual understanding, and may require writing Python code. 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.

The exam includes multiple-choice and short answer questions. It may also include free-response or open-ended questions to demonstrate conceptual understanding. You are expected to write reasonably correct code as well as determine expected behavior of novel computer code. Grading primarily follows correct reasoning but may include deductions for major syntax errors, algorithmic inefficiency, or poor implementation.

Final Project (45%). This course culminates with a final project in lieu of a final exam. Teams of two students (teams of three 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

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 "–" at $\approx$ 1.5% of grade boundary).

Cheating

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

Final Project

Requirements. Teams of two students design and implement a deep learning system that addresses a substantial problem in machine learning. Teams should identify problems from personal interests or research areas. Teams may select problems with existing solutions provided their implementation demonstrates innovation beyond reproducing prior work. Projects must incorporate substantial depth and demonstrate mastery of problem specification, model selection, experimental design, and rigorous evaluation. Teams must obtain written approval from the instructor before beginning their projects.

All projects must use PyTorch as the primary deep learning framework unless approved explicitly in writing by the instructor. Projects may use additional languages and frameworks for tooling and support.

Scoring and Milestones.

Initial Proposal	week 10	4%
Revised Proposal	week 11	8%
Status Report	week 12	8%
Presentation	week 13	20%
Final Report		25%
Model Card		3%
Video		2%
Source Code		30%

Project Deliverables.

Proposals: The initial proposal describes the problem, prior related work, candidate datasets, proposed technical approach, and expected outcomes. The revised proposal incorporates instructor feedback and reflects any changes based on early implementation work. Proposals communicate that the topic is adequately prepared and outline immediate next steps. The proposal is a guidepost — reasonable deviations in method, approach, and scope are expected as understanding evolves.

Status report: Documents implementation progress, preliminary results, technical challenges encountered, and remaining work. This checkpoint demonstrates substantial progress toward a working system.

Presentation: Approximately 15-minute presentation describing the project problem, approach, and results. Focus on what was accomplished and why, with minimal theoretical background. Attendance and participation during all presentations is mandatory.

Final report: A comprehensive technical document that captures the complete project including problem specification, methodology, experimental results, and critical analysis. The report must provide sufficient detail for an uninformed expert to understand models, analytic decisions, outcomes, and implementation. Teams provide quantifiable metrics justifying engineering tradeoffs. The report must explicitly address extensions, what question was answered, and what questions remain.

Model card: A document providing transparency about the trained model. It summarizes intended use, performance metrics, limitations, training data characteristics, fairness considerations, and ethical implications.

Video: A 3–4 minute video describing the topic, implementation, and results for a broader technical audience. The video should be engaging and teach viewers something interesting about the work.

Source code: Submitted as GitHub repository to Gradescope. Code must include README files describing repository structure, execution instructions, and technical requirements. Training data

and model files are not included in the repository. Code should be well-organized, documented, and reproducible.

Academic Accommodations

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

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 Student Accessibility Services* https://osas.usc.edu 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 Brightspace, teleconferencing, and other technology.

Discrimination, sexual assault, and harassment are not tolerated by the university. You are encouraged to report any incidents to the *Office of Civil Rights Compliance* http://ocrc.usc.edu or to the *Department of Public Safety* https://dps.usc.edu/contact/feedback/. 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.

Academic Conduct

The University of Southern California is foremost a learning community committed to fostering successful scholars and researchers dedicated to the pursuit of knowledge and the transmission of ideas. Academic misconduct is in contrast to the university's mission to educate students through a broad array of first-rank academic, professional, and extracurricular programs and includes any act of dishonesty in the submission of academic work (either in draft or final form).

This course will follow the expectations for academic integrity as stated in the *USC Student Handbook*. All students are expected to submit assignments that are original work and prepared specifically for the course/section in this academic term. You may not submit work written by others or "recycle" work prepared for other courses without obtaining written permission from the instructor(s). Students suspected of engaging in academic misconduct will be reported to the *Office of Academic Integrity*.

Other violations of academic misconduct include, but are not limited to, cheating, plagiarism, fabrication (*e.g.*, falsifying data), knowingly assisting others in acts of academic dishonesty, and any act that gains or is intended to gain an unfair academic advantage.

Academic dishonesty has a far-reaching impact and is considered a serious offense against the university. Violations will result in a grade penalty, such as a failing grade on the assignment or in the course, and disciplinary action from the university itself, such as suspension or even expulsion.

For more information about academic integrity see the student handbook https://policy.usc.edu/studenthandbook/, or the Office of Academic Integrity's website https://academicintegrity.usc.edu/, and university policies on Research and Scholarship Misconduct https://policy.usc.edu/research-and-scholarship-misconduct/.