SYLLABUS – EE 641: Fall 2025 (2 units)

A Computational Introduction to Deep Learning

Instructor:	Brandon Franzke	Office:	EEB 504B
Email:	franzke@usc.edu	Hours:	Monday: 12:00 – 13:25
			Thursday: 15:10 – 17:00

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: 3404) 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.

Teaching assistants

TA:	Alexios Rustom
Email:	arustom@usc.edu
Office:	(see Course Website)

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 the core work must be your own. Submitting AI-generated work as your own without proper understanding or attribution is academic misconduct and will be treated as such.

You must develop complete mastery of all course material independently 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.

Course Outline

	Topics	Homework
Week 1 27 Aug	Introduction, PyTorch, and deep learning review. Advanced CNN architectures and applications.	**HW 0 assigned**.
Week 2 03 Sep	Object detection and segmentation.	**HW 1 assigned**.
Week 3 10 Sep	Generative Models: Energy based, Generative adversarial networks (GANs), and Variational autoencoders (VAEs).	
Week 4 17 Sep	Recurrent Neural Networks (RNNs) and Long Short-Term Memory (LSTM).	**HW 2 assigned**.
Week 5 24 Sep	Sequence-to-sequence models and attention mechanisms.	
Week 6 01 Oct	Transformers and self-attention.	**HW 3 assigned**.
Week 7 08 Oct	**Exam #1 (weeks 1–6)**.	
Week 8 15 Oct	Large language models (LLMs) and Mixture of experts (MoE).	
Week 9 22 Oct	Reinforcement learning principles. Project Overview	**HW 4 assigned**.
Week 10 29 Oct	Reinforcement learning applications.	**Draft project proposal due (31 Oct)**.
Week 11 05 Nov	Project proposal meetings.	**Revised project proposal due (07 Nov)**.
Week 12 12 Nov	Semi-supervised training, similarity learning, and contrastive learning.	
Week 13 19 Nov	**Exam #2 (weeks 8–12)**.	
(22 Nov)	**Status report due.**	
(26 Nov)	**No class, Thanksgiving Holiday.**	
Week 14 03 Dec	**Project presentations (mandatory) — EXTENDED TIME.**	
Sunday 07 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.

Exams (20%). 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. 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.

Exams include multiple-choice and short answer questions. They 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 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

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. Groups should identify and solve novel problems from their personal interests or research areas. Groups may select problems with existing solutions provided their implementation demonstrates innovation beyond reproducing prior work. All projects must demonstrate mastery of course concepts through their technical approach and analysis. Groups may abstract problems from their original contexts to meet project timeline constraints. Projects must incorporate substantial mathematical complexity and extend core course material. Groups 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. The instructor may provide additional requirements when introducing the final project assignment.

Scoring and Milestones.	Topic Proposal (initial and revised)	week 10	4% + 10%
	Status Report - data, training, integration	week 14	8%
	Presentation and demo	final	20%
	Project report		20%
	Model Card		4%
	Model and source code		30%
	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 15 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 by adding the instructors as a collaborator to your GitHub repository. 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. Details will be provided in the project handout.

Model card: A model card is a short document that provides context and transparency about your trained model. It summarizes the intended use, performance, and limitations of your model. It must include intended use, factors, metrics, evaluation data, training data, and ethical considerations.

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 Accommodations

Any student requesting academic accommodations based on a disability is required to register with to 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 DSP 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 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 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/.