ESE 577 — Deep Learning Algorithms and Software

Instructor: Jorge Mendez-Mendez

Fall 2025

1 Lecture

Tuesdays and Thursdays 3:30 pm – 4:50 pm, Frey Hall 226. Attendance is mandatory. During lectures, we will do an overview of the technical contents of the course, tie together the high-level motivation for the ideas we'll cover, and go over concepts in detail. We will additionally do exercises on the whiteboard and some Python coding.

2 Course staff & office hours

Instructor: Jorge Mendez-Mendez Mondays and Wednesdays, 3:30 pm – 5:00 pm, Light Engineering Building, Room 145, or by appointment.

TA: Yuchen Tang Tuesdays and Fridays, 1:30 pm – 3:00 pm, location TBD

3 Official course description

This course is an introduction to deep learning which uses neural networks to extract layered high-level representations of data in a way that maximizes performance on a given task. Deep learning is behind many recent advances in AI, including Siri's speech recognition, Facebook's tag suggestions and self-driving cars. Topics covered include basic neural networks, convolutional and recurrent network structures, deep unsupervised and reinforcement learning, and applications to problem domains like speech recognition and computer vision. Classes will be a mix of short lectures and tutorials, hands-on problem solving, and project work in groups. Fall, 3 credits, grading ABCF.

4 Lecture schedule

The following schedule is tentative and subject to change.

	Tuesday	Thursday
Week 1	Aug 26 Introduction to machine learning	Aug 28 Regression, regularization
Week 2	Sep 02 Gradient descent on generic function	Sep 04 Gradient descent for ML
Week 3	Sep 09 Vanilla linear classification	Sep 11 Logistic regression for binary and multi-class classification
Week 4	Sep 16 Nonlinear models, systematic feature transformations	Sep 18 Domain-dependent feature transformations
Week 5	Sep 23 Fully-connected neural networks	Sep 25 Backpropagation: forward and backward passes
Week 6	Sep 30 Why convolutional networks, 1D and 2D convolution	Oct 02 2D convolution on 3D tensors, max-pooling, CNN architecture
Week 7	Oct 07 **MIDTERM REVIEW**	Oct 09 **MIDTERM**
Week 8	Oct 14 **FALL BREAK**	Oct 16 Modern neural nets: optimization, regularization
Week 9	Oct 21 (recorded due to travel) Modern neural nets: training bigger models, pipeline	Oct 23 (recorded due to travel) Transformers: tokenization, attention
Week 10	Oct 28 Transformers: self-attention, learning, applications	Oct 30 MDPs: formal definition, policy evaluation
Week 11	Nov 04 MDPs: optimal policy, value iteration	Nov 6 Reinforcement learning: model learning, Q-learning
Week 12	Nov 11 Reinforcement learning: what changed from supervised	Nov 13 Deep RL, deep Q-learning
Week 13	Nov 18 Challenges and improvements to deep Q-learning	Nov 20 Unsupervised learning: autoencoder
Week 14	Nov 25 Unsupervised learning: representation learning	Nov 27 **THANKSGIVING BREAK**
Week 15	Dec 02 Special lecture: Continual learning	Dec 04 Special lecture: The future of deep learning

5 Recommended books

- "MIT 6.390 lecture notes," posted each week on Brightspace
- "Deep learning with Python," (2nd edition) by François Chollet
- "Deep learning: A visual approach," by Andrew Glassner
- "Deep learning with Pytorch," by Eli Stevens, Luca Antiga, and Thomas Viehmann
- "Deep learning," by Ian Goodfellow, Yoshua Bengio, and Aaron Courville
- Deep learning has become so popular, that you can find blogs and tutorials for just about any topic. I encourage students to browse around to find additional material.

6 Grading

The following evaluations will take place throughout the semester:

- 10% quizzes: between 1 and 2 quizzes per week, at the end of class. I will keep only the top 75% of your grades (e.g., keep 15 out of 20 quizzes).
- 20% homework: approximately 10 homework assignments, roughly weekly. They will involve a mix of theory and practice, all autograded in Jupyter.
- 20% project: single larger project toward the end of the semester
- 20% midterm: date 10/09/2025
- 30% final: date determined by the registrar's office. (tentatively 12/16/2025 5:30 pm – 8:00 pm. Check `https://www.stonybrook.edu/commcms/registrar/registration/_exams/fall25-finals-2.php` for updates.)

Late days Assignments turned in late will receive a penalty of 20% per day. The full 20% penalty is applied at midnight immediately after the deadline for each assignment. Each student will be granted three automatic 1-day extensions on homework assignments.

Collaboration policy Students are responsible for writing their own quizzes, assignments, projects, and exams. For homework assignments, students are welcome (and encouraged) to discuss problems with one peer, **but each student must write their own assignment wrtieup and code individually**. The peer must be listed at the top of the writeup for each assignment. *Note: I will treat AI assistants as peers. That is, students are welcome to discuss problems with an AI assistant, but it is considered cheating to directly obtain an answer by querying the assistant. Please credit any AI assistant that you use.*

7 Student Accessibility Support Center Statement

If you have a physical, psychological, medical, or learning disability that may impact your course work, please contact the Student Accessibility Support Center, Stony Brook Union Suite 107, (631) 632-6748, or at `sasc@stonybrook.edu`. They will determine with you what accommodations are necessary and appropriate. All information and documentation is confidential.

Students who require assistance during emergency evacuation are encouraged to discuss their needs with their professors and the Student Accessibility Support Center. For procedures and information go to the following website: `https://ehs.stonybrook.edu/programs/fire-safety/emergency-evacuation/evacuation-guide-disabilities` and search Fire Safety and Evacuation and Disabilities.

8 Academic Integrity Statement

Each student must pursue his or her academic goals honestly and be personally accountable for all submitted work. Representing another person's work as your own is always wrong. Faculty is required to report any suspected instances of academic dishonesty to the Academic Judiciary. Faculty in the Health Sciences Center (School of Health Technology & Management, Nursing, Social Welfare, Dental Medicine) and School of Medicine are required to follow their school-specific procedures. For more comprehensive information on academic integrity, including categories of academic dishonesty please refer to the academic judiciary website at `http://www.stonybrook.edu/commcms/academic_integrity/index.html`.

9 Critical Incident Management

Stony Brook University expects students to respect the rights, privileges, and property of other people. Faculty are required to report to the Office of Student Conduct and Community Standards any disruptive behavior that interrupts their ability to teach, compromises the safety of the learning environment, or inhibits students' ability to learn. Faculty in the HSC Schools and the School of Medicine are required to follow their school-specific procedures. Further information about most academic matters can be found in the Undergraduate Bulletin, the Undergraduate Class Schedule, and the Faculty-Employee Handbook.