

# Deep Learning

## Course Information

- **Instructor:** Danfei Xu (danfei@gatech.edu)
- **Course Prefix and Number:** CS 4644 A
- **Term:** Fall 2026
- **Lecture:** Tuesday/Thursday, 5:00 PM – 6:15 PM
- **Course Website:** [https://sites.cc.gatech.edu/classes/AY2025/cs7643\\_fall/](https://sites.cc.gatech.edu/classes/AY2025/cs7643_fall/)

## Course Description

Deep Learning introduces neural networks and representation learning across multiple domains, including computer vision, natural language processing, speech, and robotics. The course is organized into modules covering machine learning fundamentals, convolutional neural networks, training techniques, sequence modeling and transformers, generative models, and frontier topics. Instruction combines lectures with substantial programming assignments and a culminating group project.

This course is cross-listed as CS 4644 (undergraduate) and CS 7643 (graduate).

## Prerequisites

- Introductory machine learning: CS 4641 (undergraduates) or CS 7641 / ISYE 6740 / CSE 6740 (graduates)
- Algorithms, data structures, and computational complexity
- Calculus, linear algebra, and probability
- Programming experience in Python; familiarity with PyTorch is helpful but not required

## Course Topics

The course is organized into the following modules:

1. **Background and Fundamentals:** Machine learning review; loss functions; optimization; linear classifiers.
2. **Neural Networks:** Neural network basics; computation graphs; backpropagation; optimization algorithms for deep learning.

3. **Convolutional Neural Networks:** Convolution operations; pooling; CNN architectures; image classification; transfer learning.
4. **Training Deep Networks:** Activation functions; batch normalization; dropout; regularization; data augmentation; learning rate scheduling.
5. **Sequence Modeling and Transformers:** Recurrent neural networks (RNNs); LSTMs; sequence-to-sequence models; attention mechanisms; the transformer architecture.
6. **Computer Vision Applications:** Object detection; instance and semantic segmentation; 3D vision; neural rendering.
7. **Generative Models:** Autoencoders; variational autoencoders (VAEs); generative adversarial networks (GANs); diffusion models.
8. **Language Models:** Large language models (LLMs); pre-training and fine-tuning; reinforcement learning from human feedback (RLHF); instruction tuning.
9. **Frontier Topics:** Vision-language models; robot learning; selected advanced topics.

## Required Course Materials

The recommended textbook is *Deep Learning* by Ian Goodfellow, Yoshua Bengio, and Aaron Courville (MIT Press, 2016), freely available at <https://www.deeplearningbook.org/>. Lecture slides and supplementary readings will be provided on the course website.

Programming assignments use Python and PyTorch.

## Grading Policy

- **Homework (64%):** Four programming assignments, equally weighted at 16% each.
- **Final Project (36%):**
  - Project Proposal: 1%
  - Milestone Report: 10%
  - Final Written Report: 20%
  - Poster Presentation: 5%
- **Class Participation:** Up to 1% bonus credit for engagement on the course discussion forum.

## Late Policy

Every homework assignment and project deliverable includes a 48-hour grace period with no penalty. Submissions received after the grace period will receive zero credit. Exceptions for documented emergencies may be requested through the Dean of Students office.

## **Final Project**

The final project is a substantial group effort constituting 36% of the course grade. Projects may involve applying deep learning to a new problem domain, reproducing and extending a recent research paper, or conducting original research. Deliverables include a project proposal, a milestone progress report, a final written report (6–8 pages in a standard conference format), and a poster presentation. Detailed guidelines will be provided on the course website.

## **Attendance Policy**

Students are expected to attend all scheduled lectures. Lecture slides will be posted on the course website after each class. Attendance will not be formally taken, but active engagement with course material—including class discussions—is expected.

## **Academic and Research Honesty/Integrity Statement**

Georgia Tech aims to cultivate a community based on trust, academic integrity, and honor. Students are expected to act according to the highest ethical standards. Review the [Student Code of Conduct](#) and the [Academic Honor Code](#), especially [Appendix A: Graduate Addendum to the Academic Honor Code](#).

Discussion of assignment concepts is encouraged; however, submitted work must be entirely the student's own. Code solutions may not be copied from or directly adapted from external sources, including online repositories, prior semesters' solutions, or other students' work. All public resources used must be explicitly cited. Any student suspected of cheating or plagiarism will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.

## **Core IMPACTS**

Not applicable.

## **Accommodations for Students with Disabilities**

If you are a student with learning needs that require special accommodation, [contact the Office of Disability Services](#) as soon as possible to make an appointment to discuss your special needs and to obtain an accommodations letter. Please also e-mail me as soon as possible in order to set up a time to discuss your learning needs.

## **Student-Faculty Expectations**

At Georgia Tech, we believe that it is important to strive for an atmosphere of mutual respect, acknowledgement, and responsibility between faculty members and the student body. [The Student-Faculty Expectations](#) articulates some basic expectations that you can have of me and that I have of you. In the end, simple respect for knowledge, hard work, and cordial interactions will help build the environment we seek. Therefore, I encourage you to remain committed to the ideals of Georgia Tech while in this class.