CS 7643 Deep Learning

Fall 2025, Tue/Thu 2:00 pm - 3:15 pm, classroom TBD, Georgia Tech-Europe

Course Information | Schedule | Grading | Late Policy | Prerequisites | Projects | External Resources

Course Information

This is an exciting time to be studying (Deep) Machine Learning, or Representation Learning, or for lack of a better term, simply Deep Learning!

Deep Learning is rapidly emerging as one of the most successful and widely applicable set of techniques across a range of domains (vision, language, speech, reasoning, robotics, AI in general), leading to significant commercial success and exciting new directions that may previously have seemed out of reach. With the recent advances in large language models (LLMs), successful commercial deployments of these techniques continue to grow.

This course will introduce students to the basics of Neural Networks (NNs) and expose them to some cutting-edge research. It is structured in modules (background, Convolutional NNs, NN Training, Sequence Modeling, Generative Modeling, Frontiers). Modules will be presented via instructor lectures and reinforced both homeworks that teach theoretical and practical aspects. The course will also include a project which will allow students to explore an area of Deep Learning that interests them in more depth.

Instructor

Benjamin Joffe

Class Info & Links

Lectures: T/Th: 2:00 pm - 3:15 pm, classroom TBD

Website: https://cc.gatech.edu/~bjoffe3/teaching/AY2026/cs7643_fall

Piazza: TBD

Canvas TBD

Gradescope TBD

Tentative Schedule (subject to changes)

Date	Topic	Optional Reading
W1: Aug 19	Intro lecture + class logistics Slides (pdf) PS0 is due 11:59pm 08/24 (NO grace period). See Piazza for instruction. IMPORTANT: All students MUST complete PS0!	• LeCun et al., Nature '15 • Shannon, 1956 • DL book: Linear Algebra background • DL book: Probability background • DL book: ML Background
W1: Aug 21	Machine learning intro, applications (CV, NLP, etc.), parametric models and their components PS0 due 08/24 11:59pm (no grace period) PS/HW1 out, due Sep. 11th 11:59pm	
W2: Aug 26	Supervised Learning, Linear Classification, Loss functions, Gradient Descent	• DL book: Linear Algebra background • DL book: Probability background • DL book: ML Background
W2: Aug 28	Gradient Descent, Backpropagation	
W3: Sep 2	Neural Networks, Backpropagation, Linear Algebra View	• DL book: Deep Feedforward Nets • Matrix calculus for deep learning
W3: Sep 4	Backpropagation, Jacobians, Autodiff	• Automatic Differentiation Survey, Baydin et al.
W4: Sep 9	Automatic Differentiation recap, Activation functions, Optimization	
W4: Sep 11	Optimization, Convolution Layers PS/HW1 due Sep. 11th 11:59pm, PS/HW2 out	• DL book: Optimization for Training Deep Models • DL book: Convolutional Networks
W5: Sep 16	Convolutions Cont., Convolution Gradients	
W5: Sep 18	Pooling, Convolutional Neural Networks (CNNs)	
W6: Sep 23	CNNs continued, Regularization, Augmentation, Transfer Learning Project Proposal Due Sep 23th 11:59pm	• DL book: Regularization for DL
W6: Sep 25	Recurrent Neural Networks, Long Short-Term Memory	• DL Book: Sequential Modeling and Recurrent Neural Networks (RNNs)
W7: Sep 30	Attentions and Transformers PS/HW2 due Sep 30th 11:59pm, PS/HW3 out	• Attention is all you need • BERT Paper • The Illustrated Transformer • Formal Algorithms for Transformers
W7: Oct 2	Transformers Cont. and Vision Transformers	
W8: Oct 7	Object Detection and Segmentation	• Fully Convolutional Networks for Semantic Segmentation • An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale
W8: Oct 9	Project Planning Session	
W9: Oct 14	Generative Models (Part I): Generative Adversarial Networks	• Generative Adversarial Networks
W9: Oct 16	Generative Models (Part II): Denoising Diffusion Probabilistic Models	
W10: Oct 21	Variational Autoencoders (VAEs) PS/HW3 due Oct. 21st 11:59pm, HW4 out	• Tutorial on Variational Autoencoders
W10: Oct 23	RL background Project Check-in due Oct 23rd 11:59pm	• Sutton & Bartow Chapter 1 • Survey paper on Deep RL • MDP Notes (courtesy Byron Boots)
W11: Oct 28	No class: GT-E RECESS	
W11: Oct 30	No class: GT-E RECESS	
W12: Nov 4	RL: RL Part 2 - Q-Learning, DQN, Policy Gradient	• Notes on Q-learning (courtesy Byron Boots)
W12: Nov 6	RL: Policy Gradients, REINFORCE, Actor-Critic, Large Language Model Reinforcement Learning from Human Feedback.	• Policy iteration notes (courtesy Byron Boots) • Policy gradient notes (courtesy Byron Boots)
W13: Nov 11	No class: Bank Holiday	
W13: Nov 13	Large Language Models (LLMs) PS/HW4 due Nov 13th 11:59pm	
W14: Nov 18	Vision-Language Models	
W14: Nov 20	3D Vision	
W15: Nov 25	Deep Learning in Physical World and Robot Learning	
W15: Nov 27	Robot Learning Cont. and Class Conclusion	
W16: Dec 2	Project presentations Final Project report due Dec 7th 11:59pm	

Grading

- 64% Homework (4 homeworks)
- 36% Final Project
- 1% (potential bonus) Class Participation: top endorsed answers/questions/comments on Piazza

Late policy for deliverables

- There will be no make-up work provided for missed assignments. Of course, emergencies (illness, family emergencies) will happen. In those instances, please submit an Class Absence Verification Form to Dean of Students office (see here for rules). The Dean of Students is equipped to verify emergencies and pass confirmation on to all your classes. For consistency, we ask all students to do this in the event of an emergency. Do not send any personal/medical information to the instructor; all such information should go through the Dean of Students.
- Every homework deliverable and project deliverable will have a 48-hour grace period during which no penalty will apply. This is intended to allow you time to verify that your submission has been submitted (we recommend you re-download it and look it over to make sure all questions/deliverables have been answered). Canvas will show your submission as late, but you do not have to ask for this grace period. Deliverables after the grace period will receive a grade of 0.

Prerequisites

CS 7643 should NOT be your first exposure to machine learning. Ideally, you need:

- Intro-level Machine Learning
  - CS 4641 for the undergraduate section and CS 7641/ISYE 6740/CSE 6740 or equivalent for the graduate section.
- Algorithms
  - Dynamic programming, basic data structures, complexity (NP-hardness)
- Calculus and Linear Algebra
  - Positive semi-definiteness, multivariate derivates (be prepared for lots and lots of gradients!)
- Programming
  - This is a demanding class in terms of programming skills.
  - HWs will involve Python and PyTorch.
  - Your library of choice for project.
- Ability to deal with abstract mathematical concepts

Plagiarism & Academic Integrity

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. All students enrolled at Georgia Tech, and all its campuses, are to perform their academic work according to standards set by faculty members, departments, schools and colleges of the university; and cheating and plagiarism constitute fraudulent misrepresentation for which no credit can be given and for which appropriate sanctions are warranted and will be applied. For information on Georgia Tech's Academic Honor Code, please visit http://www.catalog.gatech.edu/policies/honor-code/ or http://www.catalog.gatech.edu/rules/18/.

You are encouraged to discuss problems and papers with others as long as this does not involve the copying of code or solutions. After discussions, all materials that are part of a submission should be wholly your own. Do NOT search for code directly implementing the assignment and submit snippets or variations of them. You can search for conceptual information but NOT code solutions. Any public material that you use (open-source software, help from a textbook, or substantial help from a friend, etc.) should be acknowledged explicitly in anything you submit to us. If you have any doubts about whether something is legal or not, please do check with the class Instructor. We will actively check for cheating, and any act of dishonesty will result in a Fail grade. Any student suspected of cheating or plagiarizing on any deliverable including assignments will be reported to the Office of Student Integrity, who will investigate the incident and identify the appropriate penalty for violations.

AI-Based Assistance

We will use the AI-based assistance policy developed by David Joyner. We treat AI-based assistance, such as ChatGPT and Copilot, the same way we treat collaboration with other people: you are welcome to talk about your ideas and work with other people, both inside and outside the class, as well as with AI-based assistants.

However, all work you submit must be your own. You should never include in your assignment anything that was not written directly by you without proper citation (including quotation marks and in-line citation for direct quotes).

Including anything you did not write in your assignment without proper citation will be treated as an academic misconduct case.

If you are unsure where the line is between collaborating with AI and copying from AI, we recommend the following heuristics:

Heuristic 1: Never hit "Copy" within your conversation with an AI assistant. You can copy your own work into your conversation, but do not copy anything from the conversation back into your assignment.

Instead, use your interaction with the AI assistant as a learning experience, then let your assignment reflect your improved understanding.

Heuristic 2: Do not have your assignment and the AI agent open at the same time. Similar to above, use your conversation with the AI as a learning experience, then close the interaction down, open your assignment, and let your assignment reflect your revised knowledge.

This heuristic includes avoiding using AI directly integrated into your composition environment: just as you should not let a classmate write content or code directly into your submission, so also you should avoid using tools that directly add content to your submission.

Deviating from these heuristics does not automatically qualify as academic misconduct; however, following these heuristics essentially guarantees your collaboration will not cross the line into misconduct.

Students with Disabilities

If you are a student with learning needs that require special accommodation, contact the Office of Disability Services at http://disabilityservices.gatech.edu/, 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.

Subject to Change Statement

The syllabus and course schedule may be subject to change. Changes will be communicated via the Canvas announcement tool. It is the responsibility of students to check Piazza, email messages, and course announcements to stay current in their courses.

Project Details (36% of course grade)

The class project is meant for students to (1) gain experience implementing deep models and (2) try Deep Learning on problems that interest them. The amount of effort should be at least the level of 1.5 homework assignment per group member (2-4 people per group). The deliverables are

```
- Project Proposal (5%): Due Sep 23
- Milestone Written Report (10%): Due Nov 6
- Presentation (5%): Dec 2
- Final Written Report (20%): Due Dec 7
```

The final report is a PDF write-up describing the project in a self-contained manner will be the sole deliverable. Your final write-up is required to be between 6 - 8 pages using a standard Computer Science conference paper template such as CVPR and NeurIPS (we will release the LaTeX template). Please use this template so we can fairly judge all student projects without worrying about altered font sizes, margins, etc. After the class, we will post all the final reports online so that you can read about each others' work. Additionally, we will allow people to upload additional code, videos, and other supplementary material as zip file similar to code upload for assignments. While the PDF may link to supplementary material, external documents, and code, such resources may or may not be used to evaluate the project. The final PDF should completely address all of the points in the rubric described below.

Rubric

We will release a detailed project rubric on Piazza.

Related Classes / Online Resources

- CS231n Convolutional Neural Networks for Visual Recognition, Stanford
- Machine Learning, Oxford
- Deep Learning, New York University
- Deep Learning, CMU
- Deep Learning, University of Maryland
- Hugo Larochelle's Neural Networks class

Book

- Deep Learning, Ian Goodfellow, Aaron Courville, and Yoshua Bengio, MIT Press

Overviews

- Deep Learning, Yann LeCun, Yoshua Bengio, and Geoffrey Hinton, Nature
- Representation Learning: A Review and New Perspectives, Yoshua Bengio, Aaron Courville, and Pascal Vincent

© 2025 Georgia Tech