

ORI 397 Introduction to Deep Learning Methodology and Applications

Lectures

Lecture Time: 3:30 PM – 4:45 PM, Tuesday & Thursday.

Lecture Location: ETC 4.150

Instructor

Hairong Wang: <https://www.me.utexas.edu/people/faculty-directory/wang>

Email: hairong.wang@austin.utexas.edu

Office: ETC 5.120

Office Hour: Tuesday 5:00 PM - 6:30 PM. You can also email me for an appointment for a Zoom or in-person meeting.

GroupMe

We highly encourage you to ask questions and help answer others' questions by joining the GroupMe chat https://groupme.com/join_group/109603991/2U6zFFda. We will check messages on a daily basis. Please post any technical questions in the group chat instead of directly messaging me.

Course Overview

Deep learning has transformed a wide range of fields, with applications in life sciences, finance, manufacturing, automation systems, etc. The advancements in neural networks (aka “deep learning”) have significantly improved the performance of many real-world applications, from data analysis to prediction systems. This course introduces the fundamental concepts and methodologies behind deep learning architectures, focusing on developing end-to-end models that can be applied to a variety of tasks.

Learning Objectives

After taking this course, students will learn to implement, train, and fine-tune their own deep neural networks, gaining a solid understanding of key techniques in model design, optimization, and regularization. Students will also explore cutting-edge research in the field of deep learning and apply their knowledge to real-world problems. Through multiple hands-on assignments and the final course project, students will acquire the toolset for setting up deep learning tasks and practical engineering tricks for training and fine-tuning deep neural networks.

Prerequisites

- Basic understanding of Python programming
- Familiarity with linear algebra, calculus, and probability theory
- No prior deep learning knowledge required, though familiarity with machine learning is beneficial

Course Tools and Frameworks

Python will be our main programming language. Jupyter Notebooks for interactive coding and visualization.

Textbook & Other Materials

- Goodfellow, Ian, et al. Deep learning. Vol. 1. No. 2. Cambridge: MIT press, 2016.

Class notes: Electronic lecture slides will be provided before each lecture. Optional reading materials will be shared along with the slides.

Grading

Grade Constitution

- 3 Assignments: 60 %
- Final Project: 40 % (Project Proposal: 5 %; Presentation: 15 %; Final Report: 20 %)

Tentative date

- Assignment0: released 08/26; due 09/02 11:59pm CT
- Assignment1: released 09/12; due 10/07 11:59pm CT
- Assignment2: released 10/03; due 10/28 11:59pm CT
- Assignment3: released 10/24; due 11/18 11:59pm CT
- Project report: due 12/09 11:59pm CT

Project Overview

- The project is a group assignment comprising 2-4 members.
- Students will present their final projects, showcasing the application of deep learning techniques to a specific problem or dataset. This project will demonstrate the practical implementation of the methodologies and concepts covered throughout the course.
- A dataset will be provided for students to use, but they also have the option to select their own dataset for the project.

Letter Grade Minimum Averages Required (Tentative): A: 90; B: 80; C: 70; D: 60; F: <60. I would like to reserve the right to adjust the Minimum Average to avoid certain extreme cases.

Homework Policy: Collaboration on homework is allowed, but submitted solutions must be your own and reflect individual effort (not identical to others' work).

Late Work Policy: To provide flexibility, each homework and project has a 24-hour grace period with no penalty. Canvas may mark these submissions as late, but no action is required on your

part. Use this time to confirm your files are uploaded correctly. Work submitted after the grace period will not be accepted and will receive a grade of zero.

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).

Grade Check & Regrade Request: You may request a re-grade within one week of receiving your graded homework or project. Please note that we reserve the right to re-grade the entire assignment or project, which means you could lose points as well as gain them.

Accessible, Inclusive, and Compliant Statement

The university is committed to creating an accessible and inclusive learning environment for students with disabilities consistent with university policy and federal and state law. Please let us know if you experience any barriers to learning so we can work with you to ensure you have equal opportunity to participate fully in this course. If you are a student with a disability, or think you may have a disability, and need accommodations please contact Disability and Access (D&A). Please refer to D&A's website for contact and more information: <http://disability.utexas.edu/>. If you are already registered with D&A, please deliver your accommodation letter to us as early as possible in the semester so we can discuss your approved accommodations and needs in this course.

University Policies & Student Resources

In addition to the requirements outlined in the General Information Catalog, this course follows all university policies regarding academic integrity, accessibility, safety, and student conduct. Please review the [University Policies and Resources for Students Canvas](#) page, which compiles key services and guidance to help you navigate the course and the university.

Course Outline

Class Web Page: All class materials will be posted on Canvas <https://utexas.instructure.com/courses/1426145>.

Chapter 1: Introduction & Course Logistics (1 week)

- Course logistics, policies, and tools
- Overview of deep learning and applications

Chapter 2: Fundamentals of Deep Learning (~ 2 weeks)

- Supervised learning
- Gradient descent and optimization basics
- Backpropagation and computation graphs
- Neural networks

- Guidance on how to pick a project

Chapter 3: Convolutional Networks and Training Deep Models (~ 3 weeks)

- Convolutional networks (convolution, pooling, CNN architectures)
- Strategies for effectively training deep networks (preprocessing, initialization, normalization, regularization, optimization, hyperparameter search)
- Advanced training practices including transfer learning and ensembles
- Project helping session

Chapter 4: Sequence Models and Transformers (~ 2 weeks)

- Recurrent Neural Networks (RNN, LSTM, GRU)
- Attention mechanism and Transformers (BERT, GPT, etc.)

Chapter 5: Applications in Computer Vision (~ 1 weeks)

- DL software and hardware foundations
- Object detection and semantic segmentation
- Neural rendering (e.g., NeRF)

Chapter 6: Generative Models and Representation Learning (~ 2 weeks)

- Autoregressive models & VAEs
- Diffusion models
- Generative Adversarial Networks (GANs)
- Self-supervised learning methods

Chapter 7: Natural Language Processing (~ 1 week)

- Large Language Models

Chapter 8: Reinforcement Learning (~ 1 week)

- RL basics: MDP, value iteration, deep Q-learning

Chapter 9: Frontier Topics & Guest Lectures (1 week)

Chapter 10: Final Project Presentation (1 week)