Stats 453 Course Policy

Spring 2025

Last updated: Feb 5, 7pm, 2025

Basic info

- Course description

STAT 453 is a intro-level deep learning course that covers basic concepts in deep learning. This course has an emphasis on PyTorch programming, with a focus on course projects.

Knowledge about calculus, inear algebrea, and prior programming experience are required.

Instructor: Yiqiao (Joe) Zhong, (yiqiao.zhong@wisc.edu)
TA: Zhexuan Liu (zliu999@wisc.edu)

- Lecture info

Location: Bardeen 140
Time: Tuesdays, Thursdays 2:30 pm — 3:45 pm

- Office hour

Yiqiao— Tuesdays 1:30 pm—2:30 pm (Medical Science Center 1122)
Zhexuan— Thursdays 1:00 pm—2:00 pm (Medical Science Center 1475)

- Textbook

Disclaimer: Textbook is used as a reference. Exams will be based on lectures.
Machine Learning with PyTorch and Scikit-Learn: Develop machine learning and deep learning models with Python
Sebastian Raschka (Author), Yuxi (Hayden) Liu (Author), Vahid Mirjalili (Author)
Packt Publishing; 1st edition (February 25, 2022)

Grading policy

Breakdown (100 points in total):

- Assignment 20%,
- Midterm 20% & Final exam 30%
- Course projects 30%
  A. Project proposal (by group)
  B. Project presentation (by group)
  C. Project report (by group)

- Assignment (20%)

- **3—4 Exercises:** jupyter notebooks with missing code blocks or calculations to be filled
- **Submission format:** your completed jupyter notebooks; must submit via Canvas
- **- Frequency:** Three to four assignments in total
- **- Due:** roughly one week after assignment is posted
- **- Late submission:** A deduction will be applied before solutions are posted, zero grade after solutions are posted 😱

- Exams (50%)

- **Midterm:** in-class, date ~~Tuesday~~ Thursday Mar 6, same location (75 mins), 2 pages of cheatsheets, no need for calculators
- **Final exam:** May 6th, 2:45pm—4:45pm (2 hrs), 2 pages of cheatsheets, no need for calculators
- **Scope:** All questions are based on lecture slides. Final exam will be based primarily on the lectures after the midterm exam.

- **Format:** There will be multiple-choice questions, simple calculations, code explanations, etc. See sample exams on Canvas.

- Course projects (30%)

Course projects are graded by groups. Each group consists of up to 4 students (no exception allowd).

Policy about course projects will be updated soon.

Tips:

1. Please find teammates as soon as possible. Past experience shows that last-minute teams often perform poorly.
2. Start brainstorming project ideas early. Describe ideas and estimated resources clearly in the project proposal.
3. Try to get better at communication and leadership.

Communications

- Discussion on Canvas

Feel free to ask questions, or give feedback **(no piazza this year, all communications are on Canvas)**

- Email

Contact the instructor or the TA in emergent situations, e.g., getting sick, time conflict with exams, and other accommodations

Syllabus

Part 1: Introduction

- L01: Course overview, introduction to deep learning
- L02: The brief history of deep learning
- L03: Single-layer neural networks: The perceptron algorithm

Part 2: Mathematical and computational foundations

- L04: Linear algebra and calculus for deep learning
- L05: Parameter optimization with gradient descent
- L06: Automatic differentiation with PyTorch
- L07: Cluster and cloud computing resources*

Part 3: Introduction to neural networks

- L08: Multinomial logistic regression
- L09: Multilayer perceptrons and backpropagation
- L10: Regularization to avoid overfitting
- L11: Input normalization and weight initialization
- L12: Learning rates and advanced optimization algorithms

Part 4: Basics for computer vision and language modeling

- L13: Introduction to convolutional neural networks
- L14: Convolutional neural networks architectures
- L15: Introduction to recurrent neural networks

Part 5: Deep generative models

- L16: Factor analysis, autoencoders, variational autoencoders
- L17: Introduction to generative adversarial networks

- L18: Diffusion models

Part 6: Large language models

- L19: Recurrent neural networks for sequence data
- L20: Self-Attention and Transformers
- L21: Prompts and in-context learning
- L22: Foundation models, model safety, explainability
-

Part 7: class projects and final exam

- Course summary
- Student project presentations
- Final Exam
- Final report (online submissions)