# CS 1678/2078 Introduction to Deep Learning

Instructor: Scott Jordan

Email: scott.jordan@pitt.edu

Class is held MW 11:00-12:15 in SENSQ 5505.

Office hours: M 12:15-1:15 in SENSQ 6105

**Course Description:**

This course will cover the basics of modern deep neural networks. The first part of the course will introduce function approximation, neural network architectures, activation functions, and operations. It will present different loss functions and describe how training is performed via backpropagation. In the second part, the course will describe specific types of neural networks, e.g. convolutional, recurrent, and transformers, as well as their applications in computer vision and natural language processing. The course will also briefly discuss foundation models, self-supervised and deep reinforcement learning, and generation approaches.

**Prerequisites:** Math 220 (Calculus I), Math 280 or 1180 (Linear Algebra), CS 1501 (Algorithm Implementation).

**Programming language/framework:** We will use Python, NumPy/SciPy, and Jax.

**Textbooks:** The following textbook is recommend but not required to be purchased. There will be suggested readings from the book.

- Ian Goodfellow, Yoshua Bengio, Aaron Courville. *Deep Learning*. online version

**Communication:** There is a slack channel used for this course. I will post announcements to slack. You may also ask and answer questions there. It is expected that you will check slack for announcements regularly. Slack is asynchronous communication, do not expect

immediate responses or outside normal working hours. It takes a while to answer technical questions; these questions are best asked in office hours. Office hours should be your first instinct for getting help. They are a dedicated time to assist in your learning.

## Grading

Grading for CS 1678 will be based on the following components:

- Homework (5 assignments x 11% each = 55%)
- First exam (20%)
- Second exam (20%)
- Participation (5%)

Grading for CS 2078 will be based on the following components:

- Homework (5 assignments x 9% each = 45%)
- First exam (15%)
- Second exam (15%)
- Participation (5%)
- Project (20%)
    - Proposal (5%)
    - Status report (5%)
    - Presentation (5%)
    - Final report (5%)

### Assignment Submission Mechanics

Homework, project reports and presentations slides are due at 11:59pm on the due date. You will submit your homework using Canvas, under "Assignments" and the corresponding homework ID. You should submit a single zip file with source/results/report/slides, as requested. Name the file `YourFirstName_YourLastName.zip`. *Please comment your code!* Grades will be posted on Canvas.

Note that Canvas will also contain an automatically computed running average column that you can use to gauge how you're doing in the class based on grades that are already available. Generally, Overall scores over 90% map to some type of A, over 80% to B, and over 70% to C.

## Exams and Quizzes

There will be two in-class exams. The second exam is **not** cumulative and will not cover material from the first exam. *There will be no make-up exams unless you or a very close friend/relative is seriously ill!*

## Participation

Students are expected to regularly attend the class lectures, and should actively engage in in-class discussions. Attendance will not be taken, but keep in mind that if you don't attend, you cannot participate. You can actively participate by, for example, responding to the instructor's or others' questions, asking questions or making meaningful remarks and comments about the lecture, or bringing in relevant articles you saw in the news. The grading rubric will be as follows: 1 = you attended infrequently, 2 = you attended frequently but did not speak in class, 3 = you attended frequently and spoke a few times, 4 = you participated frequently, 5 = you participated every other week or more.

## Homework Late Policy

For homework assignments only (not for project components), you get 5 "free" late days counted in *minutes*, i.e., you can submit a total of 120 hours late. For example, you can submit one homework 24 hours late, and another 96 hours late. The 120-hour "budget" is total for all assignments, NOT per assignment. Once you've used up your free late days, you will incur a penalty of 25% from the total assignment credit possible for each late *day*. A late *day* is anything from 1 minute to 24 hours. Note that because of the generous late days policy, the instructor will be reluctant to give extra days to individual students due to illness that is not severe/prolonged, due to travel, etc. However, if you are facing an especially difficult situation or exceptional circumstances, please speak with the instructor (ahead of the assignment deadline, if possible) and provide documentation.

## Collaboration Policy and Academic Honesty

You will do your work (homework and quizzes) individually. The work you turn in must be your own work. You are allowed to discuss the assignments with your classmates, but do not look at code they might have written for the assignments, or at their written answers. You are also **not** allowed to search for code on the internet, use solutions posted online unless you are explicitly allowed to look at those. When in doubt about what you can or

cannot use, ask the instructor! Plagiarism will cause you to fail the class and receive disciplinary penalty. Please consult the University Guidelines on Academic Integrity. All project components involve group work.

## Note on Disabilities

If you have a disability for which you are or may be requesting an accommodation, you are encouraged to contact both your instructor and Disability Resources and Services (DRS), 140 William Pitt Union, (412) 648-7890, drsrecep@pitt.edu, (412) 228-5347 for P3 ASL users, as early as possible in the term. DRS will verify your disability and determine reasonable accommodations for this course.

## Note on Medical Conditions

If you have a medical condition which will prevent you from doing a certain assignment, you must inform the instructor of this **before** the deadline. You must then submit documentation of your condition within a week of the assignment deadline.

## Statement on Classroom Recording

To ensure the free and open discussion of ideas, students may not record classroom lectures, discussion and/or activities without the advance written permission of the instructor, and any such recording properly approved in advance can be used solely for the student's own private use.

## Project (CS 2078 only)

Students will complete a project which studies in more depth one of the topics we cover in class. Students should work in groups of two or three. These projects should focus on one of the following:

- a novel approach which addresses one of the problems covered in class, properly evaluated
- a definition of a new problem, along with detailed argumentation of why this problem is important and challenging, an approach to solve this problem, and an evaluation of this approach

- an extensive analysis and experimental evaluation of one or more of the approaches covered in class
- A related research project suggested or approved by the instructor

In the **project proposal**, students should include the following: clear problem statement, extensive literature review, detailed outline of the approach, and planned experimental setup. A good source for learning about what work has been done in your domain of interest are search engines, Google Scholar, and arxiv.org. Students are encouraged to discuss a draft of the proposal with the instructor before the proposal is due. Proposals should be 5-7 pages long, and are expected to involve substantial effort, and be well thought out. Note: a successful project does not have to show improved performance but can be an investigation that can clearly identify important properties about deep learning models or training.

The mid-semester **project status report** will describe the students' progress on the project, and any problems encountered along the way. The status report should be well formatted and possibly use a known conference format (e.g. ICLR, AAAI, NeurIPS, CVPR, ACL, RLC), but can be more informal than a conference paper. The progress report should include the following sections: Introduction, Related Work, Approach, and Results. In Results, include your experimental setup (this can change later). If you have results but they do not yet look great, include them anyway. Comment on any challenges encountered as well.

The **project presentation**, scheduled for the last day of class, will describe the students' approach and their experimental findings in a clear and engaging fashion. This will be a chance to get feedback from the class before final submission of your report. Presentations will be about 10-15 minutes long. Please submit a copy of your slides to Canvas on the same day as your presentation.

The **project final report** should be formatted and should read like a conference paper, with clear problem definition and argumentation of why this problem is important, overview of related work, detailed explanation of the approach, and well-motivated experimental evaluation. Each student should document what part of the project they did, and how duties and tasks were divided.

Tentative Course Schedule:

| Date | Topic | Other info | Due |
|---|---|---|---|
| 8/26 | Introduction – Prereq Review | | |
| 8/28 | Function approximation | | |
| 9/2 | No Class (Labour Day) | | |
| 9/4 | Optimization basics | | |
| 9/9 | ML Basics | HW 1 Released | |
| 9/11 | | | |
| 9/16 | Backprop and deep learning | | |
| 9/18 | models | | |
| 9/23 | Training neural networks | HW 2 Released | HW 1 |
| 9/25 | | | |
| 9/30 | | | |
| 10/2 | Computer Vision with Deep | | |
| 10/7 | Learning | HW 3 Released | HW 2 |
| 10/9 | | | |
| 10/14 | No Class (Fall break) | | |
| 10/16 | Recurrent Neural Networks | | |
| 10/21 | | | HW 3 |
| 10/23 | | HW 4 Released | |
| 10/28 | Transformers | | |
| 10/30 | | | |
| 11/4 | Reinforcement Learning | | HW 4 |
| 11/6 | | HW 5 Released | |
| 11/11 | No Class (Veterans Day) | | |
| 11/13 | Reinforcement Learning | | |
| 11/18 | Topics in Deep Learning | | |
| 11/20 | | | HW 5 |
| 11/25 | No class (Thanksgiving) | | |
| 11/27 | | | |
| 10/2 | Review | | |
| 10/4 | Exam | | |
| 10/9 | Project Presentations | | |