

CS 782: Advanced Machine Learning

George Mason University

Department of Computer Science

Course title and number	CS 782: Advanced Machine Learning
Term	Fall 2023
Class times and location	Tue 19:20 - 22:00, Exploratory Hall, Room L102
Piazza	https://piazza.com/gmu/fall2023/cs782

1. Course Description and Prerequisites

This course, structured as a seminar-style engagement, is designed to facilitate an in-depth exploration into the evolving domain of machine learning. Students will be expected to undertake critical assessments of research papers, formulate articulate reviews, and present research papers and form discussions within the classroom setting. Additionally, they will collaborate in teams to conduct a semester-long research project in machine learning.

Students will be provided with a selection of research papers by the course supervisor, encompassing ten emerging and imperative themes within the machine learning research sphere. These include, but are not limited to, recommender systems, fairness and bias in ranking, fairness in machine learning, bias in natural language processing, security in ML, privacy in ML, explainability in ML, graph neural networks, neuro-symbolic learning, and large language models. Furthermore, the course supervisor will provide guidance regarding potential project themes and methodologies.

The objectives of this course are twofold. First, to expose students to a range of state-of-the-art advanced machine learning research topics, and secondly, to enhance students' research skills through paper reading, presentation, review writing, and conducting a semester-long research project.

Formally, you must have received a grade of B- or better in CS 681 or CS 687 or CS 688. Students should be familiar with machine learning and deep learning basics. Programming experience in Python with widely-used libraries, such as numpy, PyTorch, Tensorflow, is required (project will use PyTorch to build machine learning algorithms). Please expect lots of programming in the project.

2. Instructor Information

Instructor: Ziwei Zhu

Email: zzhu20 at gmu dot edu

Office: Engineering 4609

Zoom: <https://gmu.zoom.us/j/2481531650>

Office hours: 4 pm to 5 pm every Tuesday (both in-person or virtually)

Attention: write '[CS 782]' in the subject if the email is related to this course

3. Grading Policies

Your final letter grade will be given based on:

Letter grade	Points (out of 100)
A+	98-100
A	93-97
A-	90-92
B+	86-89
B	83-85
B-	80-82
C+	76-79
C	73-75
C-	70-72
D	60-69
F	0-59

The overall course points will be determined using the following weights.

- 1) In-class paper presentation: 24%
- 2) Paper review: 15%
- 3) In-class report: 18%
- 4) Project: 43%

Details are as follows:

In-class paper presentation: There will be 12 paper presentation sessions, each of them is with 4 papers. Every student needs to **pick 3 papers** (in 3 different presentation sessions with different topics) to present in the classroom. Each presentation is **25 min with 10 min Q&A**. Do the best to keep the length of the presentation **25+/-5 min**. Students need to prepare the presentation slides (can reuse slides from the authors). Slides need to be submitted to blackboard. Each presentation accounts for **8%** of the final grade. Grading will be based on slides preparation, presentation skills, presentation fluency, Q&A quality, etc. After the presentation, the student needs to send professor slides through email for the purpose of grading.

Paper review: During the semester, every student needs to **choose 3 papers** to write paper reviews. These 3 papers should be from 3 different presentation sessions and should not be from the 3 sessions that the student pick to give presentations. After the students finish the paper choosing for in-class presentation, the professor will **assign 3 sessions** to each student, and the student can choose one paper from each of the 3 assigned sessions. Each review accounts for **5%** of the final grade. A review template will be provided. Grading will be based on writing quality, understanding of the paper, and critical analysis. A pdf review needs to be submitted to blackboard. The deadline for one review is the end of the Monday right before the assigned session.

In-class report: Besides the sessions the student gives presentations, students need to attend at least **6 paper presentation sessions**. And during the presentation sessions, students need to write a brief in-class report to summarize key points of the papers presented during the presentation session (report form will be provided). Each report accounts for **3%** of the final grade. A student cannot submit the report for the presentation session that she/he gives paper presentation. In other words, students are expected to **attend at least 9 presentation sessions**.

Project: will be a team project with **1-3 people**. 6 project topic options will be provided, and teams can choose the project topics from provided ones, or teams can propose new topics (but need to discuss with the professor). A team needs to submit **weekly report** about the progress of the project (13 report * 1% = 13%); a **final report** (12%); and a github repo with all project **code and data** (10%). In addition, there will be **two meetings** with professor to confirm project progress (2 meetings * 4% = 8%).

4. Research Papers for Presentation

Sign up sheet:

https://docs.google.com/spreadsheets/d/1UaaUCUHZDKtyhhu9L0sH7oOitbgVbDs7_Nm311JzcPc/edit?usp=sharing

5. Research Project Topic Options

Sign up sheet:

https://docs.google.com/spreadsheets/d/1JJweuGjte9IYjy_PGppNm3hasnYw9oA8jWVWUR8AIU8/edit?usp=sharing

6. Course Topics, Calendar of Activities, Project Milestones (subject to change)

Week	Topic	Project Milestones
1 (08/22)	Introduction	
2 (08/29)	Advances in Recommendation	Weekly Report 1; Discuss with professor
3 (09/05)	Fairness and bias in ranking 1	Weekly Report 2;
4 (09/12)	Fairness and bias in ranking 2	Weekly Report 3;
5 (09/19)	Fairness in ML	Weekly Report 4;
6 (09/26)	Bias in NLP 1	Weekly Report 5; Discuss with professor
7 (10/03)	Bias in NLP 2	Weekly Report 6;
8 (10/10)	No class	Weekly Report 7;
9 (10/17)	Security in ML	Weekly Report 8;
10 (10/24)	Privacy in ML	Weekly Report 9;
11 (10/31)	Explainability in ML	Weekly Report 10; Discuss with professor
12 (11/07)	Graph Neural Networks	Weekly Report 11;

13 (11/14)	Neuro-Symbolic Learning	Weekly Report 12;
14 (11/21)	Large Language Models	Weekly Report 13;
15 (11/28)	TBA	Weekly Report 14; Report, data, code submission

7. Academic Integrity and GMU Honor Code

Collaboration in thinking through problems can be highly beneficial and is allowed in this class. However, you may not share or look at any written material (code, answers to problems) that will be part of your or another student's submission. Please make sure you are cognizant of the GMU Honor Code: <https://oai.gmu.edu/mason-honor-code/full-honor-code-document/>. In addition, the CS department has its own Honor Code policies (<https://cs.gmu.edu/resources/honor-code/>) regarding programming assignments. Any deviation from the GMU or the CS department Honor Code is considered a Honor Code violation.

8. Accommodations and Resources for Disabilities

If you have a documented learning disability or other condition that may affect academic performance you should: 1) make sure this documentation is on file with the Office of Disability Services (SUB I, Rm. 222; 993-2474; <http://www.gmu.edu/student/drc> to determine the accommodations you need; and 2) talk with the instructor to discuss your accommodation needs.

9. Campus Closure or Emergency Class Cancellation/Adjustment Policy

If the campus closes, or if a class meeting needs to be canceled or adjusted due to weather or other concern, students should check Piazza for updates on how to continue learning and for information about any changes to events or assignments.