CS688: Machine Learning

Instructor: Fang-Yi Yu

Fall 2025

Web: **https://cs.gmu.edu/ fangyiyu/**	E-mail: `fangyiyu@gmu.edu`
Class Hours: W 7:20-10:00pm	Class Room: Horizon Hall 2009
Office Hours: M 4:00-5:00pm	Office: Research Hall 350
GTA: TBD	E-mail:
Office Hours:	Location:

Course Description

Machine learning uses computational methods and information to improve performance or accuracy. In this course, we will explore various machine learning settings that can access different information, supervised, unsupervised, online, and reinforcement learning, and we will study the computational methods to process and utilize information efficiently. Finally we will study the possibilities and limitations of machine learning.

Resources

- Textbook: Bishop, Christopher M., and Hugh Bishop. Deep learning: Foundations and concepts. Springer Nature, 2023. `https://www.bishopbook.com/`
- Communication and class link: `https://piazza.com/gmu/fall2025/cs688004`
- Course website: Canvas `https://canvas.gmu.edu/courses/51367`

Prerequisites/Corequisites

CS 580 or CS 584 or permission of instructor. Programming experience is expected. Students must be familiar with basic probability and statistics concepts, linear algebra, optimization, and multivariate calculus.

Course Objectives

1. Develop the ability to apply core machine learning principles—grounded in probability, linear algebra, and optimization—to formulate and solve prediction, inference, and control problems.

2. Understand of implementation and usage algorithms across supervised, unsupervised, and reinforcement learning using modern software tools.
3. Learn to analyze ML problems end-to-end by identifying data requirements and preparation steps, selecting appropriate representations and model families, choosing training objectives and regularization, and performing model selection via sound validation procedures.
4. Understand and apply rigorous performance metrics and diagnostics, interpret results, compare methods, and reason about bias–variance and generalization to draw meaningful conclusions from data.
5. Build effective communication skills to clearly present assumptions, methods, results, and limitations of machine learning solutions to diverse audiences, emphasizing reproducibility and responsible practice.

Preliminary Topics

This preliminary list of topics may change based on time constraints, the interests of the class, or

- Supervised Learning: Linear Regression
- Classification & Logistic Regression
- Generalized Linear Models (GLMs)
- Generative Learning Algorithms
- Kernel Methods
- Support Vector Machines
- Regularization and Model Selection
- Privacy and Fairness
- Neural Networks & Backpropagation (Deep Learning)
- Calibration and Bayesian Decision Theory
- Unsupervised Learning: Clustering & EM
- Self-Supervised Learning & Foundation Models
- Online learning and Reinforcement Learning

Policies

Grading Policy

We'll calculate your final grade based on the following components. There will be no make-up or extra-credit assignments at the end of the semester; your grade should be a measure of your semester-long progress.

- Homework: 40% (best four out of five)
- Midterm 30%
- Final project 25%
- Participation 5%

Assessment

- A+ (rank >= 5%)
- A (score >= 95.0% or rank >= 10%)
- A– (score >= 90.0% or rank >= 20%)
- B+ (score >= 85.0% or rank >= 30%)
- B (score >= 80.0% or rank >= 40%)
- B– (>= 75.0%)
- C and F

Homework

There will be five homework assignments and the score of the lowest one will be dropped. Each homework assignment has a posted deadline, and late submission is not accepted unless a valid excuse is communicated to the instructor *before* the deadline. Assignments are considered individual efforts, and no sharing and discussion of problem solutions are allowed with anyone except the TAs or the instructor.

If you feel points have been incorrectly deducted, contact the grader: TA for homework and instructor for the midterm. Contesting of grades on any/all submissions must be requested within one week of receiving the grades. No grade changes will be considered after that deadline.

Midterm exam

Midterm will be on *Oct 22* during lecture. You are allowed one 8.5x11in sheet of notes, front and back. There will be no make up for the exam unless previously arranged (well in advance).

Final project

The final project will be graded in groups of size 2-3 people. The project will consist of the following:

1. Checkpoint, a written report that contains a research topic, a brief introduction, and a literature review of the topic,
2. Presentation in the class,
3. Final report that includes the research topic, introduction, literature, results, and discussion.

Honor Code

Please see the Office for Academic Integrity (https://oai.gmu.edu/) for a full description of the code and the honor committee process, and the Honor Code Policies of the Department of Computer Science (https://cs.gmu.edu/resources/honor-code/) regarding the course project. GMU is an Honor Code university. The principle of academic integrity is taken seriously and violations are treated gravely. If you rely on someone else's work in an aspect of the course project, you should give full credit in the proper, accepted form. Another aspect of academic integrity is the free play of ideas. Vigorous discussion and debate are encouraged in this course, with the firm expectation that all aspects of the class will be conducted with civility and respect for differing ideas, perspectives, and traditions. When in doubt (of any kind) please ask for guidance and clarification.

Please refer to the GMU Common Course Policies provided by the Stearns Center (https://stearnscenter.gmu.edu/home/gmu-common-course-policies/), which cover any policies not directly superceded in this syllabus.

Disabilities

If you have a documented learning disability or other condition which may affect academic performance, make sure this documentation is on file with the Office of Disability Services and talk to the instructor about accommodations.