New Jersey Institute of Technology
Ying Wu College of Computing
Computer Science Department

Advanced Machine Learning

Code: FA25-CS732-101
Mode: Synchronous online (Zoom)
Time: Wed 6:00pm–8:50pm

Instructor: Michael Houle
Webpage: https://people.njit.edu/faculty/meh43
Office: GITC 4317D (Newark)
Email: michael.houle@njit.edu (or directly on Canvas)

Note: Your messages will usually be answered by the end of the next day. Grades for all items will generally be posted within 7 to 14 days after their due date. For issues with your grades, contact the instructor directly.

Office Hours: Tue 4:00pm–6:00pm, online (Zoom).
Reserve an online appointment slot by following this calendar link. Please try to do so at least one day in advance. Other appointment times can also be arranged by email.

Teaching Assistant / Grader: None

Course Description

[From the NJIT catalog]: This course explores contemporary advanced topics in machine learning with emphasis on understanding neural networks. Students will examine foundations of generalization including uniform convergence bounds, complexity measures, and modern perspectives on why overparameterized models can still generalize well beyond their training data.

[Instructor's description]: This course develops a rigorous theoretical foundation for understanding how and why machine learning works. We will study the statistical properties of learning algorithms, focusing on the relationship between training performance and generalization to unseen data. Central questions include: under what conditions do algorithms succeed, what role does data quantity and model choice play, and how can mathematical tools guide the design of better methods? Topics include generalization guarantees and sample complexity, uniform convergence and concentration inequalities, classical learning theory (such as VC dimension and Rademacher complexity), optimization and implicit regularization, and recent advances in deep learning theory. Throughout, we will balance mathematical precision with intuition, showing how theory explains phenomena such as overfitting, underfitting, and the trade-offs between model capacity, data size, and algorithmic choice. By the end of the course, students will be able to reason about the performance of learning algorithms, interpret the limits of current theory, and connect foundational results to modern practice in machine learning and deep neural networks.

Prerequisites

CS732 Advanced Machine Learning has the following prerequisites: CS634 or CS670 or CS675 or CS677 or DS675 or DS677, or instructor permission.

This course emphasizes theoretical foundations and research-oriented learning. Students must have a solid undergraduate-level background in differential and integral calculus, linear algebra, and probability and statistics. Familiarity with concepts and proof techniques from real analysis, matrix analysis, and convex optimization is also highly desirable.

Course Textbooks

There is no required textbook for this class. However, the class material covers much of the content found in the following online resource:

- Tengyu Ma, Lecture Notes for Machine Learning Theory (https://github.com/tengyuma/cs229m_notes/blob/main/master.pdf)

Learning Outcomes

By the end of the course, you will be able to:

a. Explain the goals and challenges of machine learning, emphasizing the theoretical perspectives that underpin practice.
b. Evaluate the quality and reliability of online resources that discuss machine learning methods and theory.
c. Recognize machine learning problems that are well-suited to theoretical analysis.
d. Interpret and critically assess empirical results on learning algorithms using theoretical concepts.
e. Relate classical ideas in learning theory to modern approaches, including deep learning.
f. Apply theoretical insights to guide the design and use of learning algorithms in new contexts.

Coursework, Assessment, and Related Outcomes

Participation [10%]
Students are expected to participate in classroom activities throughout the term. They will receive credit based on their rate of live attendance, and their level of engagement in discussions in class and on Canvas forums. Special consideration will be given for cases in which the student is absent for health reasons, conference participation, or other valid, documented reasons.
[Outcomes: a, c, d, e]

Theory Assignments [50%]
There will be four assignments on the concepts, theory, and practice in advanced machine learning analysis. Some assignments may have a coding component, such as an empirical investigation of the gap between theoretical prediction and empirical performance (implemented in Python). Students will typically have two weeks to complete each assignment. Submission will be via Canvas.
[Outcomes: c, d, e, f]

Project [40%]
Students will conduct an individual project involving the theoretical concepts and techniques related to machine learning. This can be a new theoretical analysis of a machine learning problem or method, a survey of the existing research literature that utilizes ML theory in the analysis of a practical ML technique or application, and/or an empirical verification of ML theory in a practical setting (implemented in Python). Students are encouraged to choose a topic that relates to their own graduate research, wherever possible. The project will have three components: a detailed proposal (worth 10%), a final report (worth 25%), and a video presentation (worth 5%). Submission will be via Canvas. A selection of the project submissions will be invited for video presentation and synchronous Q&A in the Week 15 class.
[Outcomes: b, c, d, e, f]

Assignment Due Dates

All assignments and milestones are due during the weeks indicated in the table below.

Assignments [50%] Due at 23:59	Project [40%] Due at 23:59
Week 5: #1 [12.5%] Sunday Week 8: #2 [12.5%] Sunday Week 11: #3 [12.5%] Sunday Week 15: #4 [12.5%] Thursday	Week 10: Proposal [10%] Friday Week 14: Final Report [25%] Friday Week 14: Video Presentation [5%] Friday

Note that there is no class in Week 13 (Wednesday November 26) due to the Thanksgiving Recess. Assignment 4 is due on the Last Day of Classes (Thursday December 11).

Course Topic Schedule

Week 1	Course Introduction
Week 2	Empirical Risk and Asymptotic Analysis
Week 3	Concentration Inequalities and Sub-Gaussian Distribution
Week 4	Uniform Convergence for Finite Hypotheses
Week 5	Uniform Convergence for Infinite Hypotheses
Week 6	Rademacher Complexity Basics
Week 7	Advanced Rademacher Bounds
Week 8	Generalization for Specific Models
Week 9	Neural Networks and Generalization
Week 10	Deep Neural Networks: Data Dependent Bounds
Week 11	Nonconvex Optimization and Neural Tangent Kernel
Week 12	Implicit Regularization
Week 13	NO CLASS due to Thanksgiving Recess
Week 14	Unsupervised, Self-Supervised, and Online Learning
Week 15	Selected Student Project Presentations

Grading Policies

Letter Grades
In accordance with the graduate grade legend, the raw total percentage assessment score will be converted to a final letter grade that will appear on your transcript. The conversion table for this course is:

Letter Grade	Percentage Range
A	90 — 100
B+	80 — <90
B	70 — <80
C+	65 — <70
C	60 — <65
F	<60

In cases where the project outcomes are judged to be original contributions meeting the publication standard of reputable international research conferences or journals, the letter grade may be raised one level higher than what is suggested by the above conversion table.

Incomplete
A grade of **I** (incomplete) is given in rare cases where work cannot be completed during the semester due to documented long-term illness or unexpected absence for other serious reasons. A student requesting special consideration should be in good standing (i.e. with a passing grade on

coursework submitted before the absence). When special consideration is granted, the student receives a provisional **I** if there is no other way to make up for the documented lost time; in such cases, an email with a timeline for makeup work will be sent to the student. Note that according to NJIT regulations, an **I** must always be resolved by the end of the following semester.

Late Submission Policy
Generally speaking, assignments, project milestones, and participation exercises will be accepted late without penalty, but only up until the time grading has begun, or solutions or other feedback have been released to members of the class. At that time, the class will be informed that submissions have closed, that no further submissions will be accepted, and that any missing work will be given a mark of zero. Students should be aware that grades, solutions, and/or feedback may be released at any time after the deadline, at the sole discretion of the lecturer and without prior warning. *The only way to ensure that work will be accepted without penalty is to submit it before the deadline*. No individual extensions will be granted. However, special consideration may be given in rare cases when a student is unable to complete an assignment for serious, unavoidable reasons — these must be communicated and documented promptly.

Grading Feedback
Assignment marks will in most cases be accompanied with class discussion of the solutions. Individual grading feedback will be given where appropriate. Further clarifications can be provided by contacting the instructor.

Grade Corrections
Check the grades in course work and report errors promptly. Please try and resolve any issue within one week of the grade notification.

Other Course Policies

Email
Use of your NJIT email or Canvas inbox is strongly encouraged.

Requesting Accommodations
If you need an accommodation due to a disability, please contact Marsha Williams-Nicholas, Associate Director of the Office of Accessibility Resources and Services, Kupfrian Hall 201 to discuss your specific needs. A Letter of Accommodation Eligibility from the office authorizing student accommodations is required.

NJIT Services for Students, Including Technical Support
Please follow this link.

Canvas Accessibility Statement
Please follow this link.

Collaboration on Assignments

You are expected to tackle all the problems **on your own**. However, some of the assignment problems may be quite challenging! For difficulties that persist, you are welcome to raise questions in the Canvas Discussion Forum, or talk to the instructor during office hours. In consulting with others, you are allowed to exchange general ideas and approaches only: unless you are given explicit permission to do so in the assignment statement, the full solutions themselves must be worked out by you alone.

Generative AI Tools and Other External Resources
Sometimes you may come across code, text, derivations, or other helpful information online, or you may be able to generate it using AI tools such as ChatGPT or other Large Language Models (LLMs). In most cases, you will be allowed to integrate this information into your solution. However, you must always give the appropriate credit and citations (e.g. links, LLM prompts) for the material you use (especially when you use the code and text you found online). In particular, in the assignments for this course, you are encouraged to use LLMs for help with proofs and derivations. Assignment submissions will not be accepted unless you supply the transcript of your 'conversation' with it (either printouts or working links to sessions), which should show all the prompts used, and the responses that guided you to the delivered solution. Your 'conversation' with it must be entirely yours, with all prompts written by you in your own words. If you do not make use of LLMs or external sources, you must declare it. Failure to give appropriate credit when using the work of others (whether human or AI) is considered plagiarism, and may lead to disciplinary action under NJIT's Academic Integrity policy (see below).

Statement on Academic Integrity

"Academic Integrity is the cornerstone of higher education and is central to the ideals of this course and the university. Cheating is strictly prohibited and devalues the degree that you are working on. As a member of the NJIT community, it is your responsibility to protect your educational investment by knowing and following the academic code of integrity policy that is found at: http://www5.njit.edu/policies/sites/policies/files/academic-integrity-code.pdf .

Please note that it is my professional obligation and responsibility to report any academic misconduct to the Dean of Students Office. ***Any student found in violation of the code by cheating, plagiarizing or using any online software inappropriately will result in disciplinary action. This may include a failing grade of F, and/or suspension or dismissal from the university.*** *If you have any questions about the code of Academic Integrity, please contact the Dean of Students Office at dos@njit.edu ."*