CS 536 Machine Learning II
Course for Graduate students, Fall 2026
Instructor: Lesia Semenova

Quick details

Monday/Thursday 12.10 – 1.30 pm. Livingston Campus, Tillet Hall 116

Overview and schedule.
Schedule and intro will be posted here: https://lesiasemenova.github.io/mlii-fall2026/

My office hours. After class. I'm usually happy to stay for ~20 minutes after class for questions or to discuss project ideas. For longer discussions, please follow up by email.
Lecturer (Rushi) Office hours – Thursday 11am-12pm, Tillet Hall

Contact.
For official or sensitive communication: lesia.semenova@rutgers.edu
For questions we will use Canvas.

Prerequisites. This is a graduate-level course. Students should be comfortable with basic linear algebra, probability, and core machine learning concepts. You should also be able to implement and debug ML experiments in Python (e.g., NumPy, scikit-learn, and PyTorch), including working with real datasets and writing reasonably organized, reproducible code.

In this course, we will cover the following topics in machine learning (This plan is subject to change as the course progresses):

- Machine learning basics
- Deep learning architectures:
  - Multi-Layer Perceptrons (MLP)
  - Activation Functions, Backpropagation, Vanishing/Exploding Gradients
  - Convolutional Neural Networks (CNN)
  - Recurrent Neural Networks (RNN)
- Attention Operations, Transformer, BERT
- Advanced topics on deep learning, including optimization and deep generative models.

Textbooks and Materials
There is no textbook for this course. However, the following books can be useful though as references on relevant topics:

- Deep Learning (DL), Goodfellow, Ian and Bengio, Yoshua and Courville, Aaron, MIT Press, 2016, ISBN: 9780262035613

- Dive into Deep Learning (https://d2l.ai/)
- Pattern Recognition and Machine Learning (PRML), Christopher C. Bishop, Springer, 2006, ISBN: 9780387310732
- Natural Language Processing with Distributed Representations (NLP), Kyunghyun Cho, https://arxiv.org/abs/1511.07916

Grading and Requirements

Your grade is based on two primary assessment types: a semester-long research project and in-class quizzes (which will closely mirror released practice materials).

Here is how you will be evaluated:
In class quizzes – 20%
Project – 65%
Engagement during the class and attendance – 15% (10% for interaction and 5% for attendance)

In class quizzes. We will have 3 total, 10 point each and the lowest score is dropped. Please also use this as an opportunity if you need to miss some quiz. Quizzes will be close to practice quizzes. Practice quizzes will be framed as homework for small extra credit.

Engagement includes asking questions, contributing to discussions, and providing thoughtful feedback during project presentations.

Attendance. I expect you to be present at lectures since engagement takes a portion of your grade. Attendance will be tracked via interactions during the class by our Graders/lecturers. We allow 4 dropped absences without changes to your score. First lecture attendance on September 3 does not count.

Class project. This is an almost semester-long project that aims to answer an interesting research question related to interpretability or explainability. Project directions may vary and can include, but are not limited to: designing a new method; evaluating existing methods by proposing a benchmark; applying algorithms to new problems or domains; proving interesting properties of models, so on.

There will be three milestones for the project:

1. **Project proposal (1 page, 5 points)** – describe your project idea, motivation, planned tools or methods, what you expect to learn, the research question you are addressing, and how your project sits within the existing literature.

2. **Midpoint check-in (up to 3 pages, 10 points)** – submit project updates, including initial results, lessons learned, successes and challenges. Indicate whether you changed direction and what you plan to do next. The report should include a link to your GitHub repository.
3. **Final presentation (presentation + write-up, 40 points)** – present your project, including clear motivation, main results, methods, observations, experiments, and lessons learned. The slides/report should include a link to your GitHub repository.
4. **Teamwork (10 points) –** how well you collaborated. We will have teamwork quizzes alongside midpoint and final presentations.

For project materials, we allow 2 days of late submissions of the materials. After it every new day will cost you 5% of the allocated grade.

What makes a strong project?
Strong projects ask a clear research question, engage meaningfully with prior work, and demonstrate careful reasoning and experimentation. Novelty is not required; depth, clarity, and critical thinking are valued. Negative results and well-executed evaluations are welcome.

You may work individually or in a team of two or three. Teams of four or more are discouraged.
If your work results in an interesting finding, I will help with framing the results for a potential publication.

Your final project is a group grade, but we reserve the right to adjust individual grades based on peer evaluations. If a student's peer evaluations unanimously indicate that they did not contribute meaningfully to the project, that student may receive a zero for the teamwork points, and their final project grade (the 40 points) may be penalized.

Grades will be assigned toward the end of the semester based on quizzes, project progress and quality, and engagement. If at any point you are concerned about how you're doing, please talk to me. I do not curve. Grades will be assigned as follows:

- **A**: 89.5 to 100
- **B+:** 84.5 to 89.49
- **B:** 79.5 – 84.49
- **C+:** 74.5 – 79.49
- **C:** 69.5 – 74.49

My goal is for everyone to get to a good grade. We will evaluate holistically and work with you through the semester to achieve that.

Logistics

Syllabus Changes. The schedule, readings, and assignments may be adjusted during the semester to better match class interests, pacing, or guest availability.

Attendance. Regular attendance is expected.

Collaboration. Collaboration is encouraged within project teams. Any external assistance, tools, or prior work that meaningfully influenced your project should be clearly acknowledged.

Absence. If you plan to be absent for medical or other reason, please fill in a google form provided on Canvas to let our Lecturers/Graders know.

Computing Resources: Google Colab, iLab at Rutgers.

Academic Integrity. Students are expected to follow Rutgers' Academic Integrity Policy. All submitted work must be your own (or your team's) and properly acknowledge sources. Violations will be handled according to university policy.

Accessibility and Accommodations. Students with disabilities or other circumstances requiring accommodations should contact the Office of Disability Services and notify me as early as possible so appropriate arrangements can be made.

Student Well-Being. Graduate school can be demanding. If you are experiencing stress or personal difficulties that affect your participation, please reach out early. Rutgers offers counseling and mental health resources for students.

Mandatory Reporting and Support Resources. As an instructor, I am a mandatory reporter under Rutgers policy. This means that if you disclose to me information about sexual misconduct, sexual assault, dating or domestic violence, or stalking, I am required to share that information with the appropriate university office so that support resources can be offered.

If you would like to speak with someone confidentially, Rutgers provides confidential resources such as Counseling Services, CAPS, and other designated confidential advocates. If you have questions about reporting options or available resources, I can help direct you to the appropriate offices.

Recording and Privacy. Students may not record or redistribute class sessions, presentations, or discussion materials without permission.

LLM-use

LLMs may be used in this course as **assistive tools**, but not as substitutes for understanding, experimentation, or original reasoning. LLMs are known to "hallucinate" or generate technically incorrect information, including fabricated citations, flawed mathematical proofs, and non-functional code. You are responsible for the correctness, clarity, and substance of all submitted and presented work.

Allowed uses include:

- clarifying concepts or terminology
- debugging code you are running yourself
- improving writing clarity and organization
- helping with quick experiment implementation (which you then run and verify)
- checking for obvious errors in math or notation
- exploring alternative formulations
- brainstorming project ideas (final choice, framing, and execution must be yours)

Not allowed:

- using LLMs to generate core content that you do not understand or could not reproduce independently
- submitting images, plots, diagrams, or figures generated directly by an LLM (or other generative tools) instead of results produced by your own code or analysis
- submitting results, explanations, or analyses that you cannot explain or defend if asked

If you are unsure whether a particular use of an LLM is appropriate, ask.

The standard we will use is simple: **you should understand and be able to reproduce everything you submit and present.**