COS 302 / SML 305 / ECE 305
Mathematics for Numerical Computing and Machine Learning

Princeton University

Instructor:	Prof. Ryan P. Adams (OH: Weds 3-4pm in COS 411)
Teaching Assistants:	Boya Zeng (OH: Tues 12-1pm & 4:30-5:30pm; Thurs 4:30-5:30pm)
Lectures:	Mon/Weds 10:40am-12:00pm in Commons D121
Precepts:	P01 – Thu 9:35-10:25am in Friend Center 008
	P03 – Thu 12:15-1:05pm in Friend Center 111
URL:	https://www.cs.princeton.edu/courses/archive/fall26/cos302/

Course Description

This course provides a comprehensive and practical background for students interested in continuous mathematics for computer science. The goal is to prepare students for higher-level subjects in artificial intelligence, machine learning, computer vision, natural language processing, graphics, and other topics that require numerical computation. This course is intended for students who wish to pursue these more advanced topics, but who have not taken (or do not feel comfortable) with university-level multivariable calculus (e.g., MAT 201/203) and probability (e.g., ORF 245 or ORF 309).

Topics will include vectors, matrices, norms, orthogonality, projection, eigenvalues, singular value decomposition, basic vector calculus, introductory probability, Monte Carlo, information theory, convex optimization, Lagrange multipliers, and gradient descent. Assignments will have both conceptual and coding components. Students will complete the coding portions in Python. Familiarity with programming will be assumed, but expertise in Python is not required.

Course Website

The course URL is https://www.cs.princeton.edu/courses/archive/fall26/cos302/. We'll also use Ed at https://edstem.org/. Most questions about the course, lecture/precept material, or the assignments should be addressed via Ed. The course instructors will regularly check this discussion board with the goal of posting responses within 24 hours. Students taking the class are also encouraged to post responses. Code examples can be posted, but don't post anything you wouldn't be expected to share with other students in the class as per the collaboration policy. Long, detailed questions are probably best answered during office hours. Use your judgement.

Precepts

There will be 50-minute weekly precepts led by the teaching assistants. New material may be presented in the precepts (see course schedule), so attendance is strongly encouraged.

Textbook and Course Materials

This course will use a freely-available textbook: Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong. *Mathematics for Machine Learning*. Cambridge University Press. 2020. Get the PDF at https://mml-book.github.io/. If you notice errors in the book, please let me know and I will pass them on to the authors personally.

Requirements and Grading

The assignments together represent 36% of the final grade, with the lowest one being dropped. The in-class quizzes together represent 24% of the final grade, again with the lowest one being dropped. There is a midterm and a final, each representing a further 20% of the final grade.

Assignments and Quizzes

This semester the course alternates between homework weeks and quiz weeks: an assignment is due (almost always) on a Wednesday at 6:00pm ET, and in the intervening weeks there is a short quiz at the start of lecture.

There are seven homework assignments, due roughly every other week. After dropping the lowest one, each of the remaining six represents 6% of your total grade. These homework assignments will involve components that are theoretical or conceptual, as well as some practical implementation using Colab notebooks. You will turn in your assignment as a PDF (compiled via LaTeX) to Gradescope. If you are not already familiar with LaTeX, it is recommended that you use Overleaf. Templates of the assignments will be posted there to help get you started. Each assignment also includes a video problem: you will record a video of at most two minutes in which you explain your solution to one problem of your choice from that assignment, uploaded to Canvas. Homework assignments should be done individually.

There are five in-class quizzes. After dropping the lowest one, each of the remaining four represents 6% of your total grade. The quizzes are short and multiple choice, and they are meant to be easy if you did the homework yourself and are keeping up with the material. There is no quiz in a week with an exam. There are no makeup quizzes; if you must miss one, that becomes your dropped quiz.

Key Dates

- First lecture: Wednesday 2 September
- In-class quizzes: 16 Sept, 30 Sept, 4 Nov, 18 Nov, and Monday 7 Dec

- Midterm exam (in class): Wednesday 14 October
- Last lecture: Monday 7 December
- Final assignment due: Tuesday 15 December (Dean's Date)
- Final exam: scheduled by the registrar during the final exam period (16–22 December)

Collaboration Policy

We want you to be able to discuss the class material with each other, but we want the homework you submit to be your own work. More specifically:

- You may never:
  - Share code.
  - Share writeups.
- You may always:
  - Discuss the related concepts and the high-level approach.
  - Discuss the results of your experiments at a high level, e.g., "I got 90% test accuracy."
- You should be wary of discussing details of proofs, your code, or results at an implementation level, rather than at the "big idea" level.
- In your assignment writeup, state who you discussed the problems with.
- It is prohibited to search the internet for assignment solutions.
- You are welcome to use AI language models (e.g., ChatGPT, Claude, etc.) in the following two ways. First, you may use it to analyze past assignments (i.e., assignments you have already submitted for grading). For example, you could use them to explore different solutions to problems on assignments that have been submitted, or debug previously submitted code. Second, you may also use language models to explain concepts and ask questions as you might do in office hours, without reference to the specific assignment. Just as you would not ask a TA to do a problem for you, do not ask an AI to do a problem for you. Nevertheless, these are powerful interactive learning tools and it is fine to use them in a good-faith manner. Any other use of language models beyond these two uses will not be allowed for this course.
- As always, remember that you are bound by the Princeton honor code, and violations can have serious consequences.

Remember that the quizzes and exams will check that you are keeping up with the material; outsourcing the homework will catch up with you quickly.

Late Policy

Homework assignments may be turned in up to a week late for a 50% penalty. There will be no exceptions and no further extensions. This class does not use "late days". Plan ahead. Note that the final assignment is due on Dean's Date and university rules mean it cannot be accepted late at all.

Changelog

- 1 September 2026 – Initial F26 version.