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

Princeton University

Instructor:	Prof. Ellen Zhong (OH Wed 2-3pm in COS 314)
Teaching Assistants:	Ziyu Xiong (OH: TBD)
	Lacey Delucia (OH: TBD)
Lectures:	Mon/Weds 10:40am-12:00pm in Engineering Quad D221
Precepts:	P01 – Thu 9:35-10:25am in Friend Center 109
	P02 – Thu 12:15-1:05pm in Friend Center 108
URL:	https://www.cs.princeton.edu/courses/archive/fall25/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 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/fall25/cos302/. We'll also use Ed at https://edstem.org/us/courses/83454/discussion. 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. Use your judgement.

Precepts

There will be two 50-minute weekly precepts led by the teaching assistants. Precepts will review and complement the material presented in lecture and aim to bridge the gap between lectures and assignments. 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/.

Requirements and Grading

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

Assignments

There are eleven homework assignments. Seven are graded and four are optional. After dropping the lowest graded homework, 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. Homework assignments should be done individually.

When a quiz is scheduled during a given week, the homework will be optional and ungraded.

Quizzes

There will be 5 short in-class quizzes spaced throughout the semester. The lowest quiz score will be dropped. After dropping the lowest quiz score, each of the remaining four represents 6% of your total grade.

Collaboration Policy

Problem sets should be written up individually and should reflect your own individual work. However, you may discuss with your peers, TAs, and instructors.

You may never copy or share complete solutions or ask others if your answer is correct, whether in person or via Ed or Canvas.

If you work on the problem set with anyone other than TAs and instructors, list their names at the top of the problem set.

AI Assistants Policy

Our policy for using ChatGPT and other AI assistants is identical to our policy for using human assistants:

- Just like you can come to office hours and ask a human questions (about the lecture material, clarifications about problem set questions, tips for getting started, etc.), you are very welcome to do the same with AI assistants.
- But: just like you are not allowed to ask an expert friend to do your homework for you, you also should not ask an expert AI.
- If it is ever unclear, just imagine the AI as a human and apply the same norm as you would with a human.

If you work with any AI on a problem set, briefly describe which AI and the prompts you used at the top of each problem.

Late Policy

Homeworks will not be accepted more than 5 days after the deadline.

Each student has 5 late days to use over the course of the semester. A late day extends the due date of a homework assignment by 24 hours. You may use multiple late days on a single assignment, but once all 5 are used, no further late days are available.

After your late days are exhausted, homework may still be submitted up to five days late for a 50% penalty.

There will be no exceptions and no further extensions. Plan ahead!