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

Princeton University

Instructor:	Prof. Ryan P. Adams (OH Mon and Wed 1:30-3pm on Zoom)
Teaching Assistants:	Joshua Aduol (OH: Tue 12:30-1:30pm, Thu 11am-12pm)
	Jad Rahme (OH: Tue 3-4pm, Thu 10-11am)
	Geoffrey Roeder (OH: Mon 4-5pm, Fri 1:30-2:30pm)
Lectures:	Pre-recorded videos
Precepts:	P01 – Thu 9:00-9:50am
	P02 – Thu 10:00-10:50am
	P03 – Fri 12:30-1:20pm
URL:	https://www.cs.princeton.edu/courses/archive/spring21/cos302/
Contact:	cos302-s21@lists.cs.princeton.edu

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/spring21/cos302/. We'll also use Ed at https://edstem.org/us/courses/3069/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. Long, detailed questions are probably best answered during office hours. Questions that are not appropriate for the discussion board may be sent to the staff via email. Use your judgement.

Precepts

There will be a synchronous 50-minute weekly precept 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 brand new 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 60% 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 ten homework assignments. After dropping the lowest one, each of the remaining ten represents approximately 7% 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.

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.

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.

Changelog

- 22 January 2021 – Initial S21 version.