Carnegie Mellon University

10-606: Mathematical Foundations for ML

Meeting Days, Times, Location:
Mondays and Wednesdays, 11:45 AM – 1:15 PM, Wean Hall 2302

Semester: Fall**, Year:** 2022

Units: 6**, Section(s):** A1

Instructor information

Name	Hoda Heidari
Contact Info	hheidari@andrew.cmu.edu
Office location	GHC 8229
Office hours	Mondays, 4:00 – 5:00 PM

TA Information [If applicable]

TA name	Michael Feffer Aishwarya Jadhav
TA Contact Info	<mfeffer@andrew.cmu.edu> <anjadhav@andrew.cmu.edu>
Office location	Porter A20 (To be confirmed)
Office hours	Aishwarya Jadhav: Tuesdays at 6 PM via Zoom Michael Feffer: Fridays at 10 AM in Porter A20 (to be confirmed)

Course Description

These mini courses provide a place for students to practice the necessary mathematical and computational background for further study in Machine Learning (ML), particularly for taking 10-601 and 10-701. Topics covered include probability (random variables, modeling with continuous and discrete distributions), linear algebra (inner product spaces, linear operators), and multivariate differential calculus (partial derivatives, matrix differentials).
The courses assume some background in each of the above, but will review and give practice in each. (They do not provide from-scratch coverage of all of the above, which would be impossible in the available time.) Some coding will be required: the courses will provide practice with translating the above mathematical and computational concepts into concrete programs.

Prerequisites: We assume you know Calculus I and basic linear algebra, including matrix multiplication and solving systems of linear equations. We'll assume that you've worked with

probabilities, including conditional probabilities and Bayes rule. We'll assume some programming experience, specifically experience in Python or enough experience to learn Python quickly, as some of the HW questions will be in Python.

Learning Objectives

- 10-607: familiarity with probability (random variables, modeling with continuous and discrete distributions), linear algebra (inner product spaces, linear operators), and multivariate differential calculus (partial derivatives, matrix differentials).

Learning Resources

There is no required textbook for this course. However, the following textbooks are recommended as a way to get an alternate presentation of some of the course material:

- Mathematics for Machine Learning. by Marc Peter Deisenroth, A. Aldo Faisal, and Cheng Soon Ong.
- Matrix Analysis (2nd ed.). Roger A. Horn, Charles R. Johnson. Cambridge University Press, 2013.
- Advanced Calculus (5th ed.). Wilfred Kaplan. Pearson, 2002.
- Introduction to Probability (2nd ed.). Dimitri P. Bertsekas, John N. Tsitsiklis. Athena Scientific, 2008.
- The Elements of Statistical Learning: Data Mining, Inference and Prediction (2nd ed.). Trevor Hastie, Robert Tibshirani, Jerome Friedman. Springer, 2008.

Assessments

The final course grade will be calculated using the following categories:

Assessment	Percentage of Final Grade
Homework assignments	60%
Two quizzes	30%
Participation	10%

- Homework assignments: Two problem sets to be completed individually; may contain coding exercise. Students will turn in their homework electronically using Gradescope.
- Quizzes: Two in-class written quizzes.
- Participation: attendance in lectures and participation in in-class polls

We convert final course scores to letter grades based on grade boundaries that are determined at the end of the semester. What follows is a **rough guide** to how course grades will be established, not a precise formula — we will fine-tune cutoffs and other details as we see fit after the end of the course.

Grade	Percentage Interval
A	90-100%
B	80-90%
C	70-80%
D	50-70%
R (F)	0-50%

Course Policies

- **Attendance & Participation**: Please sign the attendance sheet (by the TAs) when you come to the class, so we can keep track of your attendance and allocate participation points accurately.

- **Academic Integrity & Collaboration**: We encourage you to discuss course content and assignments with your classmates. However, these discussions must be kept at a conceptual level only.
    - You **may NOT** view, share, or communicate about any artifact that will be submitted as part of an assignment. Example artifacts include, but are not limited to: code, pseudocode, diagrams, and text.
    - You **may** look at another student's *code output* and discuss it at a conceptual level, as long as it is not output that appears directly in the homework submission.
    - You **may** look at another student's code error messages and discuss what the error means at a conceptual level. However, you **may NOT** give specific instructions to fix the error.
    - All work that you present must be your own.
    - Using any external sources of code or algorithms in any way must have approval from the instructor *before* submitting the work. For example, you must get instructor approval before using an algorithm you found online for implementing a heuristic function in a programming assignment.

  Violations of these policies will be reported as an academic integrity violation. Information about academic integrity at CMU may be found at https://www.cmu.edu/academic-integrity. Please contact the instructor if you ever have any questions regarding academic integrity or these collaboration policies.

- **Late-work/Make-up work policy**: Unless otherwise stated, assignments are due on the days indicated in the class schedule. However, I recognize that sometimes “life happens.” In these instances, you may use your allotted four flex days for homework assignments. These days allow you to submit the assignment up to four days late without penalty. You can use these days for any assignment and for any reason. You do not need to provide any reason: simply email the teaching staff and tell us how many of your flex days you would like to use. Once you’ve exhausted your flex days, then point deductions will

occur for any assignment submitted after the deadline. An assignment submitted 24 hours after the due date will only be eligible for 80% of the maximum number of points allotted. Assignments submitted more than 24 hours after the due date will not be accepted.
Note that you *cannot* use your flex days for *in-class* quizzes. If you experience extenuating circumstances that prohibit you from attending in-class exams, please let us know as soon as possible and produce the necessary documentation (e.g., doctor's note). We will evaluate these instances on a case-by-case basis.

- **Accommodations for students with disabilities**: If you have a disability and have an accommodations letter from the Disability Resources office, we encourage you to discuss your accommodations and needs with us as early in the semester as possible. We will work with you to ensure that accommodations are provided as appropriate. If you suspect that you may have a disability and would benefit from accommodations but are not yet registered with the Office of Disability Resources, we encourage you to visit their website.

- **Statement on student wellness**: Please take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, getting enough sleep, and taking some time to relax. This will help you achieve your goals and cope with stress. All of us benefit from support during times of struggle. There are many helpful resources available on campus. Please take advantage of them to support your overall health and wellness during challenging and stressful times (see, for example, Counseling and Psychological Services).

- **Mobile Devices**: The use of mobile devices is only allowed for participating in in-class polls. Please step out of the classroom if you need to respond to an urgent text or call.

Course Schedule

Date	Theme/Topic	Learning Outcomes Addressed	Assignments Due
Aug 29	Preliminaries	Overview; sets	
Aug 31	Preliminaries	Datatypes; functions	
Sep 2	Recitation	Python	HW 1 release date.
Sep 7	Linear Algebra	Vectors, matrices, inner product spaces	
Sep 9	Recitation	Preliminaries	
Sep 12	Linear Algebra	Linear systems and operators	
Sep 14	Linear Algebra	Regression	

Sep 16	Recitation	Linear Algebra	
Sep 19	Matrix Calculus	Multivariate calculus; derivatives	HW 1 due date.
Sep 21	Matrix Calculus	In-class quiz; lecture overflow	Quiz 1
Sep 23	Recitation	Linear Algebra	HW 2 release date.
Sep 26	Matrix Calculus	Lagrange multipliers	
Sep 28	Matrix Calculus	Nonlinear systems	
Sep 30	Recitation	Matrix Calculus	
Oct 3	Probability	Events, Probabilities, and Random Variables	
Oct 5	Probability	Mean and variance; Bayes rule	
Oct 7	Recitation	Probability	HW 2 due date.
Oct 10	Probability	Multivariate Distributions; Sampling	
Oct 12	Probability	In-class quiz; lecture overflow	Quiz 2