

Course Title: Computational Mathematics for Machine Learning

Course Number: BU MET CS 550 A1 (Fall 2024)

Course Format: On Campus

Instructor Name: Eugene Pinsky epinsky@bu.edu

Computer Science Department, Metropolitan College, Boston University

1010 Commonwealth Avenue, Room 327, Boston, MA 02215

Course Times: Monday 6:00-8:45 p.m.

Office Hours: Fridays 11-1 or by appointment

Location: 871 Commonwealth Ave CGS 527

Teaching Assistants:

Tharunya Katikireddy email: tkatiki@bu.edu

Siddhant Shah email: siddshah@bu.edu

Course Objectives: mathematics is fundamental to data science and machine learning. This course reviews essential mathematical concepts and procedures which are fundamental. These concepts are illustrated by Python and/or R code and by many visualizations.

Teaching Approach and Goals: This course discusses mathematical concepts and computational methods for data science using simple self-contained examples, intuition, and visualization. These examples will help develop intuitive explanations behind mathematical concepts. Extensive visualizations will be used to illustrate core mathematical concepts. The

emphasis is on mathematics and computational algorithms at the heart of many algorithms for data analysis and machine learning. This course will advance students' mathematical skills that can be used effectively in data analytics and machine learning.

Books (Required):

Mathematics for Machine Learning by Marc Deisenroth, Cambridge University Press, 2020

Prerequisites: basic knowledge of Python or R

Courseware: Blackboard, Course Notes

Additional materials will be added to "From Your Professor" section under group discussion section.

BU Community Public Health Policies

All students returning to campus must follow all the compliance requirements as specified in the student link.

Class Policies

Weekly programming assignments submitted through blackboard on-line. Late homework is accepted with 50% penalty. Final projects are submitted through blackboard on-line. Students will present their projects on the last day of class. Both quizzes and final are closed-book.

This is an on-campus class. Class attendance is **mandatory** and will be recorded. Unexcused absences could result in lower grades.

Academic Conduct Code – “Cheating and plagiarism will not be tolerated in any Metropolitan College course. They will result in no credit for the assignment or examination and may lead to disciplinary actions. Please take the time to review the Student Academic Conduct Code:

Academic conduct code as specified below:

http://www.bu.edu/met/metropolitan_college_people/student/resources/conduct/code.html.

NOTE: [This should not be understood as a discouragement for discussing the material or your particular approach to a problem with other students in the class. On the contrary – you should share your thoughts, questions and solutions]

Grading Criteria:

Final 30%, Project 20%, Homework 35%, Quizzes 15%

Homework: This is a “computational” class and students must have practice. Most homework assignments will consist of both programming problems and pencil-and-paper problems. Programming can be done either in Python or in R.

Project: The project is open ended and the topics can be chosen by students (subject to instructor’s approval). In the project, students have to illustrate the usage of Python or R to implement computational algorithms to real problems in data analysis. Students will present their projects on the last day of the course.

Course Overview:

1. Numeric Python and Pandas review, arrays, vectorizations, time-series with Pandas, plotting with Matplotlib.
2. Review of calculus. Basics of derivatives and integration. Chain rule for multiple variable. Multi-variate Taylor series. Computing gradients.
3. Analytic geometry: vectors, norms, distances, lengths. Inner products. Orthonormal basis. Rotation and orthogonal projections.
4. Linear algebra: scalars, vectors, matrices, tensors. Matrix multiplication and inversion. Norms and Scalar Gram matrices. Matrix decomposition.
5. Eigenvalues and eigenvectors. Eigenvector computation. Singular vector decomposition. Dimensionality reduction via principal component analysis.
6. Random variables. Discrete and continuous probability distributions (Bernoulli, uniform, Normal, binomial, Poisson). Marginal and conditional probability. Maximum likelihood estimation.
7. Expectation, variance, and covariance. Histograms and empirical distributions. Correlations and covariance matrices, Bayes rule. Parametric and non-parametric significance tests.
8. Density estimation and Gaussian mixture models. Parameter learning by Maximum Likelihood. Expectation maximization algorithm.
9. Mathematical tools (approximation, convex optimization, quadratic programming, integration), gradient-based optimization, Newton-Raphson method. Lagrange multipliers.
10. Linear models and ordinary least square regression (OLS) analysis. Parameter estimation. Generalized linear models.