ENG EC503: Introduction to Learning from Data

Course Information

Lectures: EPC 208, Tu/Th 1:30-3:15 PM

Discussion Sessions: CAS 208, Wed 6.30-8.15

Lecturer:
Venkatesh Saligrama, PHO Room 438
srv@bu.edu, `http://https://venkatesh-saligrama.github.io/`.
Office hours: Tue, 3:15-5:15pm

Graduate Student Teacher:
Mitchell Gilmore mgilm0re@bu.edu
Office hours: Fri, TBD, PHO 428

M.S. Teaching Fellow:
TBD,

We would recommend to try Piazza first (see below) for getting answers to your questions. The best way to reach us is via e-mail. When necessary, we can arrange a meeting time outside regular office hours.

Description: This is an introductory course in machine learning covering the basic theory, algorithms, and applications. The course surveys a variety of topics, covering both supervised and unsupervised learning problems, and including classification, regression, kernels, robustness and regularization, concepts from learning theory, clustering, dimensionality reduction, generative models, neural networks, and deep learning. A variety of contemporary applications will be explored through homework assignments and a term project.

Course Websites:

Website:	`http://learn.bu.edu/`
Discussion board:	`http://piazza.com/` Signup Link: `https://piazza.com/bu/spring2025/ec503` Class Link: `https://piazza.com/bu/spring2025/ec503` Access Code: rij44ad30wq
Gradescope:	https://www.gradescope.com/courses/963747 Access code: BKVP4W

Piazza: We will be using Piazza as a discussion board. You have all been registered and you should have received an invitation to join. The system is highly catered to getting you help quickly and efficiently from both the course staff and your fellow classmates. Rather than emailing questions, we encourage you to post your questions on Piazza. I am personally very active on Piazza.

Prerequisites: Probability, e.g., EC381 or EK500 or EC505, Linear Algebra, e.g., EK102 or MA142, Multivariate Calculus, e.g., MA225, and a good level of mathematical maturity. Prior experience with Matlab, e.g., EK 127 is important. Python will not be accepted for the assignments. Good computer programming skills, e.g., EC327, are desirable.

Textbooks: The required textbook for the class is:
[UML] Shalev-Shwartz and Ben-David. Understanding Machine Learning: From Theory to Algorithms (Cambridge University Press, 2014) The book is available online at
`http://www.cs.huji.ac.il/~shais/UnderstandingMachineLearning`.

Supplemental recommended books:
[CML] Daumé. A Course in Machine Learning.
`http://www.freetechbooks.com/a-course-in-machine-learning-t905.html`.
[CO] Boyd and Vandenberghe. Convex Optimization.
`http://web.stanford.edu/~boyd/cvxbook/`.

Grading: There will be regular homework assignments, two in-class exams, and a term project. Your grade will be formed as follows:

1. 20% Homework. We will count the best $n-1$ scores out of n assignments.
2. (22.5+22.5=) 45% In-class exams.
   Midterm 1: March 6
   Midterm 2: April 17
3. 30% Term project.
4. 5% Attendance and class participation **only if** your overall homework score exceeds 85%. (Participation can mean physical or online on Piazza, or office hours, etc.)

Attendance and class participation: You will find that active class attendance and compilation of class notes are essential in this course. We will be posting slides but we will also use the blackboard; it will be your responsibility to take notes. Because the topics we will cover build upon each other, if you fall behind you may find that you are lost and not able to follow the lectures.

Homework: Homeworks will be assigned regularly; a bit less frequently than on a weekly basis. They will be due one week after the date issued. We will use Gradescope. Deadlines will be strictly enforced. Although homeworks represent only 15% of the grade you will find that they are *essential* to the learning process. We strongly encourage you to work on them independently. Often, it is easy to follow another person's solution but much harder to come up with your own. Past experience has shown that the performance in the class is highly correlated with your ability to solve problem sets on your own! We will offer homework help at the discussion meeting times and office hours.

Rules of Conduct: You *may* collaborate in study groups on the solution of homeworks. An *acceptable* form of collaboration is to discuss with others possible approaches for solving the problems and then fill the details and write your solutions independently. Copying the solution that someone else has written is *unacceptable* and at times transparent. If you do collaborate you *should* acknowledge your collaborators in the write-up for each problem. We view this as essential!

Needless to say that we expect students to adhere to basic, common sense concepts of academic honesty; presenting another's work as your own or cheating on exams will not be tolerated. Knowingly allowing others to represent your work as their own is as serious an offense as submitting another's work as your own. BU takes academic integrity very seriously. More information on BU's Academic Conduct Code, with examples, may be found at `http://www.bu.edu/academics/policies/academic-conduct-code`.

Make-up Exams: There will be no make-up exams. If there is a legitimate reason for missing an exam, then the scores of other exams will be used appropriately to compensate for the missed exam. If there is no legitimate reason provided for missing an exam, a grade of zero will be assigned for the missed exam.

Term project: In lieu of a final, you will have to complete a project applying some of the knowledge you have acquired in this course. You will present your project in a brief oral presentation and submit a written final report. The report should be in NeurIPS format, minimum of 6 pages, no maximum number of pages.

You will have to work in groups of 3 people.

There are many alternatives for the project. We must take the responsibility and specify the topic, but I must approve each project proposal. It goes without saying that the project should be related to the topics of the class. For example, a pure deep learning project is not suited as a project for this class. We expect that before Spring break, you will formulate a concrete proposal for what you plan to do. You may get in touch with us if you want to discuss it.

Projects can be of three different types.

- **Implementing and studying a particular algorithm**. You can choose an application area that you are interested in, identify a specific problem, locate a relevant dataset, and solve the problem by implementing an algorithm that was not studied in class and comparing it to some baselines. You have to implement the algorithm from scratch, no libraries are allowed.
- **Original empirical research**. You may investigate any topic empirically and try to improve the state-of-the-art. Clearly, this can be risky as there are usually no guarantees that you will succeed. Still, we would like to see your new algorithm applied to at least two different datasets.
- **Original theoretical research**. You may investigate any topic theoretically and try to extend the available theory (or create your own theory). Clearly, this can be risky as there are usually no guarantees that you will get results. Still, we would like to see your new theory or method applied to at least two datasets. You may consider choosing something that has a case study as a fallback option.

Incomplete grades: Incomplete grades will not be given to students who wish to improve their grade by taking the course in a subsequent semester. An incomplete grade may be given for medical reasons if a doctor's note is provided. The purpose of an incomplete grade is to allow a student *who has essentially completed the course* and who has a legitimate interruption in the course, to complete the remaining material in another semester. Students will not be given an opportunity to improve their grades by doing extra work.

Drop dates: Students are responsible for being aware of the drop dates for the current semester. Drop forms will not be back-dated.

Inclusion: I consider this classroom to be a place where you will be treated with respect, and I welcome individuals of all ages, backgrounds, beliefs, ethnicities, genders, gender identities, gender expressions, national origins, religious affiliations, sexual orientations, ability and other visible and nonvisible differences. All members of this class are expected to contribute to a respectful, welcoming and inclusive environment for every other member of the class.

Accommodations for Students with Documented Disabilities: If you are a student with a disability or believe you might have a disability that requires accommodations, requests for accommodations must be made in a timely fashion to Disability & Access Services, 25 Buick St, Suite

300, Boston, MA 02215; 617-353-3658 (Voice/TTY). Students seeking academic accommodations must submit appropriate medical documentation and comply with the established policies and procedures `http://www.bu.edu/disability/accommodations/`.

Tentative Syllabus.

1. Introduction and logistics.
2. Review of Linear Algebra and Probability.
3. PAC learning.
4. Agnostic PAC learning.
5. Linear Regression.
6. Bias-Variance Trade-off and Model Selection.
7. Perceptron.
8. Support Vector Machines.
9. Duality and kernels.
10. Regularization.
11. Stochastic Gradient Descent.
12. Multiclass Classification.
13. Boosting.
14. Clustering.
15. Dimensionality Reduction.
16. Generative Models.
17. Nearest Neighborhood and Decision Trees.
18. Neural Networks and Back-propagation.
19. Deep Learning Overview.
20. Applications of Deep Learning.
21. Online Learning.