ENG EC500: Introduction to Online Learning

Course Information

Lecture Time and Place: PHO 201, Tue/Thu 3:30-5:15

Lecturer:

Francesco Orabona, 8 St. Mary's St., Room PHO 430
email: fo@bu.edu
website: `http://francesco.orabona.com/`
Office hours: Tue/Thu: 1:30-2:30

I would recommend to try Piazza first (see below) for getting answers to well-formulated questions. The best way to reach me is via e-mail; I can arrange a meeting time outside regular office hours.

Description: This course deals with the foundations and advances of online learning and online convex optimization. The main theme of the course is the design and theoretical understanding of algorithms that make sequential decisions in adversarial environments, striving to perform as close as possible to a fixed strategy that knows the future in advance. Special attention will be paid to parameter-free, efficient, and practical algorithms. The focus will be on theorems and proofs for the analysis of online learning algorithms. The class will also cover applications of online learning to stochastic optimization, boosting, portfolio selection, and statistical machine learning topics.

Course Websites:

Website:	`http://learn.bu.edu/`
Discussion board:	`http://piazza.com/` Signup Link: `https://piazza.com/bu/fall2019/engec500f5` Class Link: `https://piazza.com/bu/fall2019/engec500f5/home`

Piazza: We will be using Piazza as a discussion board. 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, I encourage you to post your questions on Piazza.

Prerequisites: A strong level of mathematical maturity is required. In particular, Linear Algebra, e.g., EK102 or MA142, Multivariate Calculus, e.g., MA225 are required. Prior knowledge of convex optimization, e.g. EC524 or ES524, is also suggested.

Textbooks: I will give lecture notes for each class, however there are few suggested books that might be useful as references:

- Cesa-Bianchi and Lugosi. Prediction, Learning, and Games.
- Hazan. Online Convex Optimization. Available at `https://ocobook.cs.princeton.edu/OCObook.pdf`.
- Lattimore and Szepesvári. Bandit Algorithms. Available at `https://tor-lattimore.com/downloads/book/book.pdf`.
- Shalev-Shwartz. Online Learning and Online Convex Optimization. Available at `https://www.cse.huji.ac.il/~shais/papers/OLsurvey.pdf`.

Grading: There will be regular homework assignments and a term project. Your grade will be formed as follows:

1. 45% Homework.
2. 45% Term project.
3. 10% Attendance and class participation **only if** your overall homework score exceeds 85%.

Attendance: You will find that active class attendance and compilation of class notes are essential in this course. I will heavily use the blackboard, so 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, problably 4-5 assignments, each of which consists of several theory questions on algorithm design and analysis. I will only accept hard copies of your homework during class, writted in Word or Latex, no handwritten notes. Homework submission by email will not be accepted. Deadlines will be strictly enforced.

Rules of Conduct: An *acceptable* form of collaboration is to discuss with others possible approaches for solving the problems. Yet, you'll have to 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.

Needless to say that I 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`.

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 (or poster) presentation and submit a written final report. The report should be typed and concise; you should use your judgment as to how much is too much or too little.

The project can be done in groups of maximum size of 3 people.

The project includes two submissions (both must be written in Latex or Word)

- proposal – 1 page, due 10/31
- final report – up to 8 pages, due 12/3

Finally, you will present your results with an oral presentation.

There are many alternatives for the project. I want you to take the responsibility and specify the topic; you should view this more as a research task rather than as a homework problem. I expect that on 10/31, you will formulate a concrete proposal for what you plan to do. You may get in touch with me if you want to discuss it.

Projects can be of different types. Here is a partial list:

- **Empirical evaluation.** Compare the performance (error, time, etc.) of different existing online learning algorithms on some datasets. Explore the effect of different tuning of parameters. Try to modify existing algorithms in a reasonable way and see how they perform.
- **Original theoretical research**. Pick an open problem in online learning and try to make progress on it. Complete solutions are not required or expected! Try to break down or simplify the problem and make progress.
- **A Survey paper.** Select a set of papers on some topic related to the course and write a critical and insightful survey report. You should demonstrate in-depth understanding of the chosen subject, providing a detailed and clear overview and comparison of existing results along with some key algorithmic and analysis techniques. Also, identify new research directions.

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.

Tentative Syllabus.

1. Introduction to online learning and guessing example
2. Notes on convex analysis and probability
3. Online convex optimization, Online Gradient Descent, and lower bound
4. Online Mirror Descent (OMD)
5. Exponentiated Gradient and p-norms
6. Learning with experts and lower bound
7. Follow-The-Regularized-Leader (FTRL) and Vovk-Azoury-Warmuth forecaster
8. Online classification, Perceptron, and Mistake Bounds
9. $L^\star$ bounds
10. Adaptivity to gradients: AdaGrad
11. Connection to stochastic optimization: online-to-batch, noise, and smoothness
12. Connection to statistical learning theory and boosting
13. Adaptation to competitor norm and KL bounds with reduction to coin-betting
14. Logarithmic regret for strongly convex losses and exp-concave losses
15. Online learning and concentration inequalities
16. Universal portfolio selection
17. Multi-armed Bandits (MAB), Exp3 algorithm, lower bounds
18. Optimal MAB algorithms, FTRL/OMD with Tsallis entropy
19. Stochastic MAB and Upper Confidence Bound (UCB) algorithm
20. Stochastic linear bandits and Linear UCB
21. Contextual bandits and Exp4