10701: Intro to Machine Learning

Instructors:
Ziv Bar-Joseph, zivbj@cs.cmu.edu
Pradeep Ravikumar, pradeepr@cs.cmu.edu

Teaching Assistants:

Lectures: WEH 7500, Mondays and Wednesdays, 3:00 – 4:20 PM

Office Hours:
Ziv Bar-Joseph: Monday after class
Pradeep Ravikumar: GHC 8111, Tuesdays 3:00 – 4:00 PM

Course Description:
Machine learning is concerned with the study and development of automated systems that improve their performance through experience. Examples range from robots learning to better navigate based on experience gained by roaming their environments, medical decision aids that learn to predict which therapies work best for which diseases based on historical health records, and speech recognition systems that lean to better understand your speech based on experience listening to you.

Learning Objectives:
This course is designed to give a graduate-level student a thorough grounding in the methodologies, technologies, mathematics and algorithms currently needed by people who do research in machine learning, and related disciplines and applications.

Pre-requisites:
Students entering the class are expected to have a pre-existing working knowledge of probability, linear algebra, statistics and algorithms, though the class has been designed to allow students with a strong numerate background to catch up and fully participate. In addition, recitation sessions will be held to revise some basic concepts.

Outline of material:

- Foundations
    - Key Axes of ML: Data, Algorithms, Tasks
    - Data: Partially/Fully Observed, Interactive
    - Algorithms: Model-based, Model-Free
    - Tasks: Prediction, Description
    - Decision Theory, Generalization, Model Selection, Guarantees
- Regression
    - Linear, Polynomial

- Classification
    - Logistic Regression, Naïve Bayes, Support Vector Machines, Boosting, Surrogate Losses, Decision Trees
- Nonparametric Methods
    - K Nearest Neighbors, Kernel Regression and Density Estimation
    - Kernel Trick
- Unsupervised Learning
    - Graphical Models
    - Clustering
    - Latent Variables Models, Expectation Maximization
- Sequence Models
    - Hidden Markov Models
    - State Space Models
- Representation Learning
    - Random Features
    - Principal Component Analysis, Independent Component Analysis
    - Neural Networks, Deep Networks
- Reinforcement Learning
    - Markov Decision Processes
    - Value Iteration, Q Learning

Logistics:

Class Website:
http://www.cs.cmu.edu/~pradeepr/701
The class schedule, logistics, and lecture materials will be posted there.

Discussion, Announcements:
We will use Piazza for announcements, as well as the discussion board for the class. We will also be using piazza for in class quizzes and for determining class participation.

Textbooks:
Lectures are intended to be self-contained. For supplementary readings, with each lecture, we will have pointers to either online reference materials, or chapters from the following books:

- Pattern Recognition and Machine Learning, Christopher Bishop.
- Machine Learning: A probabilistic perspective, Kevin Murphy.
- The Elements of Statistical Learning: Data Mining, Inference and Prediction, Trevor Hastie, Robert Tibshirani, Jerome Friedman.
- Machine Learning, Tom Mitchell.

Homework:
There will be 4 homework assignments, approximately evenly spaced throughout the semester. The assignments will be posted on the course website, and on Piazza. We will use Gradescope for submitting, and grading assignments, while your code should be submitted to Autolab.

You will get a late day quota of 8 days, which you can distribute among the four homework as you wish, subject to a maximum of 3 days per homework. Homework submitted after your late day quota will lose all of points. The homework schedule is posted right at the beginning of the semester, so please plan in advance. We expect you to use the late day quota for conference deadlines and events of the like, so we cannot provide an additional extension for such cases. In the case of an emergency (sudden sickness, family problems, etc.), we can give you a reasonable extension. But we emphasize that this is reserved for true emergencies.

Collaboration Policy:
The homeworks are structured to give you experience in both written mathematical exercises and programming exercises. While it is completely acceptable for you to collaborate with other students in order to solve the problems, we assume that you will be taking full responsibility in terms of writing up your own solutions and implementing your own code, failing to do this will count as an academic integrity violation and will be reported to CMU authorities. You must indicate on each homework the students with whom you collaborated.

Exams:
There will be two exams, one scheduled halfway through, and the other at the last scheduled class for the semester. The precise dates are on the course website. Note that while exams will take place on dates in which a class is taught, on those days we will not have a regular class and instead will have the exam in the afternoon (from 5-7) given space constraints. The exam will consist of multiple choice and true/false questions, as well as short-answer questions and will be open book and open notes.

Class project:
There will be a class project. Each team would be able to choose a project from a small, pre-defined, set of topics that we will provide. You can form groups of up to 3 students. Further details on the project can be found on the website.

Grading:

Exam 1	15%
Exam 2	15%
Homeworks	40%
Project	25%

Class participation	5%

Accommodations for Students with Disabilities:
If you have a disability and are registered with the Office of Disability Resources, I encourage you to use their online system to notify me of your accommodations and discuss your needs with me as early in the semester as possible. I 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, I encourage you to contact them at access@andrew.cmu.edu.

Take care of yourself:
Take care of yourself. Do your best to maintain a healthy lifestyle this semester by eating well, exercising, avoiding drugs and alcohol, 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 and an important part of the college experience is learning how to ask for help. Asking for support sooner rather than later is almost always helpful.

If you or anyone you know experiences any academic stress, difficult life events, or feelings like anxiety or depression, we strongly encourage you to seek support. Counseling and Psychological Services (CaPS) is here to help: call 412-268-2922 and visit their website athttp://www.cmu.edu/counseling/. Consider reaching out to a friend, faculty or family member you trust for help getting connected to the support that can help.

If you or someone you know is feeling suicidal or in danger of self-harm, call someone immediately, day or night:
- CaPS: 412-268-2922
- Re:solve Crisis Network: 888-796-8226

If the situation is life threatening, call the police
- On campus: CMU Police: 412-268-2323
- Off campus: 911

If you have questions about this or your coursework, please let me know. Thank you, and have a great semester.