EECS 445: Introduction to Machine Learning

Fall 2018

Course Staff

Professor
Dr. Wiens
Office: 3765 BBB
wiensj@umich.edu

Course Staff:
Thomas Huang (thomaseh)
Mark Jin (kinmark)
Anurag Koduri (kanuarg)
Vamsi Nimmagadda (vimmada)
Cristina Noujaim (cnjoujaim)
Shengpu Tang (tangsp)
Yi Wen (wennyi)

Course Description

This course is a programming-focused introduction to machine learning. Increasingly, extracting value from data is an important contributor to the global economy across a range of industries. The field of machine learning provides the theoretical underpinnings for data-analysis as well as, more broadly, for modern artificial intelligence; it has had a major impact on many real-world applications.

In this course, we will emphasize understanding of foundational algorithms and "tricks of the trade" through implementation and basic theoretical analysis. On the implementation side, the emphasis will be on practical applications of machine learning to computer vision, data mining, speech recognition, text processing, bioinformatics, and robot perception and control. Real datasets will be used whenever feasible to encourage understanding of practical issues. On the theoretical side, the course will give an undergraduate-level introduction to the foundations of machine learning topics including regression, classification, kernel methods, regularization, neural networks, graphical models, and unsupervised learning.

Prerequisites

EECS 281 In addition, we strongly suggest that students have familiarity with linear algebra (MATH 217, MATH 417) and probability (EECS 401).

Lectures

Monday & Wednesday, 6:00pm-7:30pm, AUD CHRYS

Discussion

In addition to attending lecture, you are encouraged to attend a discussion section once a week. During discussion the course staff will go over additional examples related to concepts introduced in lecture. You will also have the opportunity to ask additional questions and get clarification regarding concepts covered in lecture.

Canvas

Information about the course including assignments and supplementary readings will be posted on Canvas (https://canvas.umich.edu/gateway/). You are expected to check the site frequently.

Piazza

We will use Piazza for course-related questions & answers. You can access the course Piazza site from Canvas. You are encouraged to use Piazza to connect with other students in the class, answer each other's questions and notify course staff of typos. You can even ask questions privately to the staff. Course staff will monitor Piazza at regular intervals, but will not provide immediate responses. For more immediate feedback, students are encouraged to attend office hours. Please do not share answers to homework on Piazza.

Course Materials & Textbook

Course materials will be posted on Canvas.

Recommended Textbooks (optional):

Aurelien Geron, "Hands-On Machine Learning with Scikit-Learn and TensorFlow," O'Reilly, 2017.
Chris Bishop, "Pattern Recognition and Machine Learning," Springer, 2007.
Kevin Murphy, "Machine Learning: A Probabilistic Perspective," MIT Press, 2012.

i>clickers

We will use i>clickers in class. Go to Canvas... i>clicker ... Remote Registration to register your clicker. Bring the clicker, and extra AAA batteries, with you to each class.

Grading

Homework (20%): Homework assignments are due on the dates noted in the course schedule and should be submitted as a PDF via Gradescope. Homework must be individual work. Students may study in groups and may help each other understand the material, but must not do or copy each others' homework. Students are encouraged to learn the math formatting tool LaTeX. Hand-written homework, scanned to PDF, is acceptable, but must be clearly readable. You must ensure that each PDF file is relatively small (<50MB limit).

Exams (55%): There is a midterm exam (25%) and a final exam (30%). If you need special accommodation, please let the professor know within the first two weeks of the semester.

Mini-Projects (20%): There will be two mini-projects assigned over the course of the semester to strengthen understanding of fundamental concepts and provide an opportunity for hands-on learning using real datasets.

In-Class Problems (5%): In-class problems will be solved during lectures, graded for effort. No make-up quizzes will be allowed. The lowest 3 quiz grades will be dropped.

Course Evaluation (Bonus 0.5%): The final course evaluation provides useful feedback. Please forward the receipt you receive from the course evaluation system to course staff. If the course evaluation is completed, you will get 0.5% bonus points toward your final grade.

Grading policy:
Homework solutions will be posted three days after the due date. Students will be responsible for grading their own homework. We will provide students with a clear grading rubric. Students are encouraged to identify where they went wrong in their solution and correct their mistakes. For questions in which students made a clear effort, students can receive 50% of the original points for these corrections. Scanned copies of the graded/corrected assignments are due at dates specified in the course schedule. If graded homework is not submitted by the deadline, students cannot retrieve points for their corrections. If no graded version of the homework is submitted, then student will receive a grade of zero. The goal behind this policy is to ensure that students to take an active role in addressing their misconceptions and in evaluating their performance. Mini-projects and exams will be graded by the course staff.

Late submission policy for homework and projects:
You can be up to 3 days late, automatically losing 10% of the assignment total for each 24 hour period starting immediately. For example, suppose that you score 80/100 on an assignment, but you submit 1 hour late. This will result in a deduction of 10 pts, leading to a total of 70/100. >24 and <=48 hours late results in a total of 60/100 (i.e., -20pts). No submissions will be accepted after the 3 days unless accompanied with a note from the Dean.

Commitment to Equal Opportunity

As indicated in the General Standards of Conduct for Engineering Students, we are committed to a policy of equal opportunity for all persons and do not discriminate on the basis of age, marital status, sexual orientation, gender identity/expression, ethnicity, religion, ability - and other visible and non-visible differences. Please feel free to contact us with any problem, concern, or suggestion. All members of this class are expected to contribute to a respective, welcoming and including environment for every other member of the class.

Disability Policy

If you think you need an accommodation for a disability, please let the instructor know. As soon as you make us aware of your needs, we can work with the Services for Students with Disabilities (SSD) office to help us determine appropriate academic accommodations. SSD (734-763-3000; http://ssd.umich.edu) typically recommends accommodations through a Verified Individualized Services and Accommodations (VISA) form. Any information you provide is private and confidential and will be treated as such. For special accommodations during any exam, we will need to receive the necessary paperwork issued by the SSD office by September 30th, 2017.

Mental Health and Wellbeing

University of Michigan is committed to advancing the mental health and wellbeing of its students. If you or someone you know is feeling overwhelmed, depressed, and/or in need of support, services are available. For help, contact Counseling and Psychological Services (CAPS) at (734) 764-8312 and https://caps.umich.edu/ during and after hours, on weekends and holidays, or through its counselors physically located in schools on both North and Central Campus. You may also consult University Health Service (UHS) at (734) 764-8320 and https://www.uhs.umich.edu/mentalhealthsvcs, or for alcohol or drug concerns, see www.uhs.umich.edu/aodresources For a listing of other mental health resources available on and off campus, visit: http://umich.edu/~mhealth/

Honor Code

The Honor Code outlines certain standards of ethical conduct for persons associated with the College of Engineering at the University of Michigan. The policies of the Honor Code apply to graduate and undergraduate students, faculty members, and administrators. Read about the UM Honor Code here: (http://www.crlt.umich.edu/faculty/honor.html). There is also an Engineering Honor Code: (http://www.engin.umich.edu/students/honorcode/code/). In this class, as in many others at the University, you will be expected to include and sign the Honor Pledge on each assignment you submit. The Honor Pledge is as follows:

> I have neither given nor received unauthorized aid on this assignment, nor have I concealed any violations of the Honor Code.

The Honor code is based on these tenets:

- Engineers must possess personal integrity both as students and as professionals. They must be honorable people to ensure safety, health, fairness, and the proper use of available resources in their undertakings.
- Students in the College of Engineering community are honorable and trustworthy persons.
- The students, faculty members, and administrators of the College of Engineering trust each other to uphold the principles of the Honor Code. They are jointly responsible for precautions against violations of its policies.
- It is dishonorable for students to receive credit for work that is not the result of their own efforts.

Among other things, the Honor Code forbids plagiarism. To plagiarize is to use another person's ideas, writings, etc. as one's own, without crediting the other person. Thus, you must credit information obtained from other sources, including web sites, e-mail or other written communications, conversations, articles, books, etc.

When answering in-class quizzes you should answer only with your registered i>clicker.

Handling Data with Integrity & Unfair Advantage

You may not falsify or misrepresent methods, data, results, or conclusions, regardless of their source. You may not possess, look at, use, or in any way derive advantage from the solutions of homework, exams or papers prepared in prior years (or from other courses), whether these solutions were former students' work products or solutions that have been made available by University of Michigan faculty or on the Internet.

Day	Date	Topic	Released	Due @ 4pm	Readings
Wed	**5-Sep-18**	Lec 1: Introduction; Linear Classification			
Fri	**7-Sep-18**		HW1		
Mon	**10-Sep-18**	Lec 2: Learning Linear Classifiers, Perceptron Algorithm			
Wed	**12-Sep-18**	Lec 3: Linear Classifiers Non-Separable Case, Gradient Descent			
Mon	**17-Sep-18**	Lec 4: Linear Regression - Empirical Risk & Least Squares, Regularization			
Wed	**19-Sep-18**	Lec 5: Support Vector Machines; Primal Formulation; Geometric Margin			Bishop pg. 325-331, Appx.E
Fri	**21-Sep-18**		Project 1	HW1	Murphy pg. 496-505
Mon	**24-Sep-18**	Lec 6: Dual Formulation; Kernels			
Wed	**26-Sep-18**	Lec 7: Performance Measures; Dimensionality Reduction			
Fri	**28-Sep-18**			Graded HW1	
Mon	**1-Oct-18**	Lec 8: Decision Trees; Entropy			
Wed	**3-Oct-18**	Lec 9: Bagging; Random Forest (Guest Lecturer)			
Fri	**5-Oct-18**		HW2	**Project 1**	
Mon	**8-Oct-18**	Lec 10: Boosting; Adaboost			
Wed	**10-Oct-18**	Lec 11: Introduction to Neural Networks			Bishop pg. 225-249
Fri	**12-Oct-18**			HW2	Murphy pg 562-587
Mon	**15-Oct-18**	Fall Break; no lecture			
Wed	**17-Oct-18**	**MIDTERM (6pm-8pm); no lecture**			
Fri	**19-Oct-18**		Project 2	Graded HW2	
Mon	**22-Oct-18**	Lec 12: Training Deep Neural Networks			Bishop pg. 267-272
Wed	**24-Oct-18**	Lec 13: Convolutional Neural Networks			Murphy pg. 563-572
Fri	**26-Oct-18**				
Mon	**29-Oct-18**	Lec 14: Recurrent Neural Networks; Autoencoders			
Wed	**31-Oct-18**	Lec 15: Introduction to Clustering; K-means, Hierarchical Clustering			Bishop pg. 424-430
Fri	**2-Nov-18**		HW3	**Project 2**	Murphy pg 352, 873-878
Mon	**5-Nov-18**	Lec 16: Spectral Clustering			
Wed	**7-Nov-18**	Lec 17: Collaborative Filtering			
Mon	**12-Nov-18**	Lec 18: Introduction to Generative Models; Gaussian Mixture Models			
Wed	**14-Nov-18**	Lec 19: Expectation Maximization			
Fri	**16-Nov-18**		HW4	HW3	
Mon	**19-Nov-18**	Lec 20: Bayesian Networks; d-separation			Bishop pg. 359-362, 372-383; Murphy 307-311, 324-327
Wed	**21-Nov-18**	**no lecture (THANKSGIVING)**			
Mon	**26-Nov-18**	Lec 21: Learning Bayesian Networks		Graded HW3	Murphy pg. 320-323
Wed	**28-Nov-18**	Lec 22: Hidden Markov Models; modeling			Murphy pg. 589, 604, 612-616; Bishop pg. 610-615, 629-631
Mon	**3-Dec-18**	Lec 23: Hidden Markov Models; alogrithms			Murphy pg. 609-612; Bishop pg. 618-625
Wed	**5-Dec-18**	Lec 24: Introduction to Reinforcement Learning			Géron Chapter 16
Fri	**7-Dec-18**			HW4	
Mon	**10-Dec-18**	Lec 25: Reinforcement Learning; Course Summary			
Fri	**14-Dec-18**			Graded HW4	
Thu	**20-Dec-18**	**FINAL EXAM (8am-10am)**			