Introduction to Machine Learning

CptS 437

Spring 2020

Monday / Wednesday / Friday 10:10 – 11:00, Sloan 175

Course Overview

Machine learning is the study of computer algorithms and models that learn automatically from data. It is a key area of artificial intelligence and has applications in many domains, including biology, social science, statistics, and image processing. This introductory course covers key topics in machine learning, including linear models for regression and classification, decision trees, support vector machines and kernel methods, neural networks and deep learning, ensemble methods, unsupervised learning and dimension reduction.

Course Instructor

Instructor:	Teaching assistant:
Diane Cook	Mahdi Pedram
EME 121	Dana 114
335-4985	
MW 11:00 – 11:30	Tu 1:30 – 3:30
djcook@wsu.edu	mahdi.pedram@wsu.edu

Prerequisites

Required: CptS 223 or CptS 233 or CptS 215 (or equivalent). In addition, students are expected to have some familiarity with basic linear algebra (vectors, matrices, matrix-vector computations, vector and matrix norms, linear independence), multivariate calculus (derivatives of univariate functions, derivatives of multivariate functions, chain rule), and basic probability and statistics (discrete and continuous probability distributions, sum rule, product rule, marginal probability distributions, conditional probability distributions, joint probability distributions, independence and conditional independence, Bayes Theorem, variance and covariance, expectation).

Required Instructional Material

Required textbook: Hal Daumé, A Course in Machine Learning, 2017. Available for download from http://ciml.info/. Additional online materials may also be recommended for individual lectures, see course schedule.

of the class materials are available on Blackboard, including the syllabus, homework assignments, papers, and lecture materials. Homework assignment descriptions, submissions, and grades are all handled via Blackboard as well. Log in to Blackboard with your WSU ID and password at https://learn.wsu.edu/webapps/login/.

Specific Course Learning Outcomes and Assessments

Following completion of this course, students should (1) have an understanding of major supervised, unsupervised and reinforcement learning techniques, (2) have a basic understanding of evaluation methodologies, (3) have a working knowledge of how to apply machine learning technologies to real-world datasets, and (4) have gained experience designing and applying machine learning techniques in team settings..

This class provides a unique opportunity to strengthen skills in each of the WSU Seven Learning Goals and Outcomes: 1) Critical and Creative Thinking, 2) Quantitative Reasoning, 3) Scientific Literacy, 4) Information Literacy, 5) Communication, 6) Diversity, and 7) Depth, Breadth, and Integration of Learning. The methods and measures for each goal is summarized in the table.

WSU Learning Outcome	Goal (by end of course)	Course topics that address the learning outcome	Evaluation
Critical and Creative Thinking	Understand the method and applicability of alternative machine learning strategies	• Decision trees, nearest neighbors, k-means cluster, neural network, linear regression, logistic regression, SVMs	• Homework assignments • Exams • Project
Quantitative Reasoning	Grasp properties involved in algorithm assessment	• Decision boundaries, margin, performance measures, validation	• Homework assignments • Exams • Semester project
Scientific Literacy	Be aware of and understand state-of-the-art research in machine learning	• Guest lectures on dep learning, generative adversarial networks, tensor flow	• Exams
Information Literacy	Be able to access and utilize literary resources to understand a machine learning challenge	• Research projects	• Semester project
Communication	Present the results of a research project and service learning orally and in writing	• Research project	• Project poster presentation • Project demonstration
Diversity	Be aware of ethical issues related to machine learning	• Lectures on supervised and unsupervised learning	• Exams • Semester project
Depth, Breadth, and Integration of Learning	Understand issues related to practical application of machine learning technologies	• Multi-disciplinary research project	• Semester project

Course Requirements

(1) *Homework Assignments (30%).* You will be assigned six homework assignments to complete. All assignments will have written components and programming components. The homework assignments will expose you to the machine learning methods we discuss in class and data from a diversity of applications that illustrate how the methods can be used. All programs will be written in Python, They will be assigned and submitted using Google's Collaboratory online Python programming environment. Completed homework assignments are due by 11:59pm on the due date.

(2) *Three Midterm Exams (45%).* Three in-class exams will be given during the semester. The exams will cover all material discussed in class up to the lecture prior to the exam date.

(3) *Semester Project (25%).* To obtain experience designing, enhancing, and applying machine learning techniques in a team-based setting, you will complete a semester project. This project will take the place of a final exam. See the Semester Project section below for additional details.

Semester Grades

There are sites that post common mappings from scores to grades, such as found at https://en.wikipedia.org/wiki/Academic_grading_in_the_United_States. Semester grades will be no lower than what is posted on these sites.

Semester Project

A requirement for this class is that you design and complete a machine learning project (graded out of 100 points). Each project will include implementation of a machine learning technique not described in class or enhancement of a described technique, with application to a real-world dataset or problem. Students are encouraged to work in teams consisting of 2-3 students. Due dates related to the project are listed below.

- February 26: Project ideas and requirements will be summarized in class.
- March 23: Project proposals are due as part of Homework #4 (worth 10 points of the project grade). The proposal should include a brief problem statement, proposed methods, novelty of the technique beyond what was discussed in class, application, and evaluation. The proposal should include a list of team members with assigned roles as well.
- April 27, April 29, May 1: Project teams will present a poster in class describing the project (worth 30 points of the project grade). Sample posters are available on Blackboard. The poster presentation should be 7 minutes in length, allowing 1-2 minutes for follow-up questions and discussion. Send a Powerpoint or PDF files with your one-page poster by April 25 at 5:00pm to djcook@wsu.edu.
- May 7, 10:00am: Project due date (also the final exam date/time). Provide a link to working code with instructions on running it and a video demonstrating how to run the code or highlighting project results (worth 40 points of the project grade – the remaining 20 points are assigned based on project scope and completeness).

Course Participation

Students are highly encouraged to attend all classes and actively participate in discussions. While participation is not part of the course grade, lectures will provide information not always available in the text and in slides that will be valuable for homework assignments and the exams.

Policy Regarding Late Work: Assignments are expected to be emailed by 11:59pm on the listed due date. After that time, 15% will be deducted per day for the first two days. Assignments turned in more than two days late will not be accepted.

Students with Disabilities: Reasonable accommodations are available for students with a documented disability. If you have a disability and may need accommodations to fully participate in this class, please either visit the Access Center (Washington Building 217) or call 509-335-3417 to make an appointment with an Access Advisor. All accommodations MUST be approved through the Access Center.

Academic Integrity Policy: Academic integrity is the cornerstone of higher education. As such, all members of the university community share responsibility for maintaining and promoting the principles of integrity in all activities, including academic integrity and honest scholarship. Academic integrity will be strongly enforced in this course. Students who violate WSU's Academic Integrity Policy (identified in Washington Administrative Code (WAC) 504-26-010(3) and -404) will fail the assignment, will not have the option to withdraw from the course pending an appeal, and will be reported to the Office of Student Conduct.

Cheating includes, but is not limited to, plagiarism and unauthorized collaboration as defined in the Standards of Conduct for Students, WAC 504-26-010(3). You need to read and understand all of the definitions of cheating: http://app.leg.wa.gov/WAC/default.aspx?cite=504-26-010. If you have any questions about what is and is not allowed in this course, you should ask course instructors before proceeding. If you wish to appeal a faculty member's decision relating to academic integrity, please use the form available at conduct.wsu.edu.

Safety Information: Washington State University is committed to maintaining a safe environment for its faculty, staff, and students. Safety is the responsibility of every member of the campus community and individuals should know the appropriate actions to take when an emergency arises. In support of our commitment to the safety of the campus community the University has developed a Campus Safety Plan, http://safetyplan.wsu.edu . It is highly recommended that you visit this web site as well as the University emergency management web site at http://oem.wsu.edu/ to become familiar with the information.

Course Calendar (Tentative)

Date	Topic	Read before class	Due by 9am
1/13	Syllabus		HW #1 assigned
1/15	Introduction to machine learning	Daumé Chapter 1	
1/17	Python / Colab overview		
1/20	Machine Luther King Day		
1/22	Decision tree	Mitchell Chapter 3 [1]	
1/24	Decision tree		
1/27	Limits of learning, inductive bias, underfitting, overfitting	Daumé Chapter 2	HW #1 due HW #2 assigned
1/29	Nearest neighbors	Daumé Chapter 3	
1/31	Decision boundaries		
2/3	K-means clustering		
2/5	Sklearn		
2/7	Perceptron	Daumé Chapter 4	
2/10	**Exam 1**		
2/12	Perceptron		HW #2 due HW #3 assigned
2/14	Linear separability, margin		
2/17	**President's Day**		
2/19	Ranking	Daumé Chapter 6	
2/21	Practical issues, normalization, hyperparameters	Daumé Chapter 5 Supplemental material [2]	
2/24	Evaluating model performance		
2/26	Significance testing, confidence intervals		
2/28	Multi-class classification, imbalanced class distributions		HW #3 due HW #4 assigned
3/2	Linear regression	Supplemental material [3]	
3/4	Loss functions, regularization	Daumé Chapter 7	
3/6	Bias and fairness	Daumé Chapter 8	
3/9	Naïve Bayes classifier	Daumé Chapter 9	
3/11	Logistic regression		
3/13	**Exam 2**		
3/16, 3/18, 3/20	**Spring Break**		
3/23	Logistic regression	Supplemental material [4]	HW #4 due HW #5 assigned
3/25	Interpretability		
3/27	Neural networks, backpropagation	Daumé Chapter 10	
3/30	Neural networks		
4/1	Deep neural networks, tensor flow	Supplemental material [7,8]	

4/3	Support vector machines	Daumé Chapter 11	
4/6	Ensemble methods	Daumé Chapter 13	
4/8	Ensemble methods		
4/10	Unsupervised learning	Supplemental material [5]	HW #5 due HW #6 assigned
4/13	Dimensionality reduction, PCA		
4/15	**Exam 3**		
4/17	Reinforcement learning	Supplemental material [6]	
4/20	Reinforcement learning		
4/22	Generative adversarial networks	Supplemental material [9]	
4/24	Autoencoders	Supplemental material [10]	HW #6 due
4/27	**Poster presentations**		Poster due 4/25 5:00pm
4/29	**Poster presentations**		
5/1	**Poster presentations**		Project due 5/7 10:00am

[1] http://www.cs.princeton.edu/courses/archive/spr07/cos424/papers/mitchell-dectrees.pdf
[2] http://cs229.stanford.edu/materials/ML-advice.pdf
[3] http://cs229.stanford.edu/notes/cs229-notes1.pdf (Part I, Section 1)
[4] http://cs229.stanford.edu/notes/cs229-notes1.pdf (Part II, Section 5)
[5] http://www.cs.otago.ac.nz/cosc453/student_tutorials/principal_components.pdf
[6] http://incompleteideas.net/book/bookdraft2017nov5.pdf
[7]https://www.researchgate.net/publication/285164623_An_Introduction_to_Convolutional_Neural_Networks
[8] https://ip.cadence.com/uploads/901/cnn_wp-pdf
[9] https://arxiv.org/abs/1812.02849
[10] https://cs.stanford.edu/~quocle/tutorial2.pdf