Introduction to Machine Learning

10-315, Fall 2019

Carnegie Mellon University

Teaching team:	Aarti Singh, Instructor
	Mary Stech, Class Assistant
	Fabricio Flores, Education Associate
	Siddharth Ancha, TA
	TBA, TA
Lecture:	**Day and Time:** Monday and Wednesday, 10:30 - 11:50 am **Location:** HOA 160
Recitation:	**Day and Time:** Friday, 10:30-11:50 am **Location:** HOA 160
Office Hours:	M, T, W, R - time and location TBA

Course Description:
Machine Learning is concerned with computer programs that automatically improve their performance through experience (e.g., programs that learn to recognize human faces, recommend music and movies, and drive autonomous robots). This course covers the core concepts, theory, algorithms and applications of machine learning. We cover supervised learning topics such as classification (Naïve Bayes, Logistic regression, Support Vector Machines, neural networks, k-NN, decision trees, boosting) and regression (linear, nonlinear, kernel, nonparametric) as well as unsupervised learning (density estimation, MLE, MAP, clustering, PCA, dimensionality reduction).

Pre-requisites: (15122) and (21127 or 21128 or 15151) and (21325 or 36217 or 36218 or 36225 or 15359)

Learning Outcomes: After completing the course, students will be able to:
*select and apply an appropriate supervised learning algorithm for classification problems (e.g., naive Bayes, support vector machine, logistic regression, neural networks).
*select and apply an appropriate supervised learning algorithm for regression problems (e.g., linear regression, ridge regression, nonparametric kernel regression).

*recognize different types of unsupervised learning problems, and select and apply appropriate algorithms (e.g., clustering, linear and nonlinear dimensionality reduction).
*work with probabilities (Bayes rule, conditioning, expectations, independence), linear algebra (vector and matrix operations, eigenvectors, SVD), and calculus (gradients, Jacobians) to derive machine learning methods such as linear regression, naive Bayes, and principal components analysis.
*understand machine learning principles such as model selection, overfitting, and underfitting, and techniques such as cross-validation and regularization.
*implement machine learning algorithms such as logistic regression via stochastic gradient descent, linear regression, or k-means clustering.
*run appropriate supervised and unsupervised learning algorithms on real and synthetic data sets and interpret the results.

Tentative lecture schedule (subject to change):

Date	**Topic**	**HW**
Aug 26 M	Introduction	
Aug 28 W	Bayes Optimal Classifier, Decision boundary	QNA1 out
Sept 2 M	**Labor Day – No Class**	
Sept 4 W	Naive Bayes, MLE	QNA 1 due, HW 1 out (Naïve Bayes and LR)
Sept 9 M	Naive Bayes, MAP, continuous features	
Sept 11 W	Logistic Regression	
Sept 16 M	Neural Networks	
Sept 18 W	Convolutional Neural Networks	HW 1 due, QNA 2 out
Sept 23 M	Case study NN implementation details	
Sept 25 W	Support Vector Machines I	QNA 2 due, HW 2 out (NNs and SVMs)
Sept 30 M	Support Vector Machines II	
Oct 2 W	Kernel Trick	
Oct 7 M	Decision Trees	
Oct 9 W	Boosting	HW2 due
Oct 14 M	kNN and model selection	
Oct 16 W	**Midterm Exam**	QNA 3 out

Oct 21 M	Linear Regression	
Oct 23 W	**Student presentations 1**	QNA3 due, HW3 out (Regression)
Oct 28 M	Regularized Least Squares regression	
Oct 30 W	Nonlinear and Kernel regression	
Nov 4 M	PCA	
Nov 6 W	Nonlinear dimensionality reduction	HW3 due, HW4 out (Unsupervised)
Nov 11 M	Clustering	
Nov 13 W	Mixture models and EM	
Nov 18 M	Nonparametric density estimation	
Nov 20 W	Learning theory	HW4 due, QNA4 out
Nov 25 M	Learning theory	
Nov 27 W	**Thanksgiving – No Class**	QNA 4 due
Dec 2 M	**Student presentations 2**	
Dec 4 W	**Final exam**	

Recommended Textbooks:

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

Grading:

- 4 Homeworks (40%)
- 4 QnAs (15%)
- Midterm and final exam (10+15=25%)
- Project (20%)

Communication: All class discussions, announcements and other communication will take place via Piazza.

Policies:

- **Collaboration**
  - You may discuss the questions
  - Each student writes their own answers
  - Each student must write their own code for the programming part
  - Please don't search for answers on the web, Google, previous years' homeworks, etc.

- please ask us if you are not sure if you can use a particular reference
- list resources used (references, discussants) on top of submitted homework

- **Waitlist**

  we'll let everyone in as long as there is space in room
  wait to see how many students drop
  keep attending lectures and doing HW

- **Audits and Pass/Fail**

  Audits NOT allowed
  Pass/Fail allowed

- **Academic integrity**

  Any violations of academic integrity will always be reported to the university authorities (your Department Head, Associate Dean, Dean of Student Affairs, etc.) as an official Academic Integrity Violation, in compliance with CMU's Policy on Academic Integrity (https://www.cmu.edu/policies/student-and-student-life/academic-integrity.html), and will carry severe penalties.

Support:

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. You are not alone. 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 often 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 at http://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 the instructors know.