UCSD COGS 118A: Spring 2017

Introduction to machine learning I: Syllabus

Zhuowen Tu

Lecture Time:

MWF 11:00a-11:50a CENTR 212

Lab hour: Wednesday, 2:00p-5:00pm, CSB 115, computer lab

Text Books:

1. Kevin P. Murphy, "Machine Learning: a Probabilistic Perspective", 2013. here

2. R. Duda, P. Hart, D. Stork, "Pattern Classification", second edition, 2000.

This course is self-contained; having the textbook is helpful but not absolutely necessary.

Course Description:

This course is one part of a two-course foundation that forms a rigorous introduction to machine learning. COGS 118A and COGS 118B are independent courses that may be taken in either order.

Introduction to machine learning (I) will prepare the students in basics of the statistical classification methods which will likely serve the foundation for data analysis and inference in a variety of applications. It will also be helpful in learning more advanced statistical machine learning algorithms, which have been applied in a wide range of scenarios for studying and predicting cognitive models, financial models, social behaviors, brain growth patterns, and visual inference.

You will need to use either Matlab or Python to do your assignments and final project. You can pick either one (Matlab or Python) or get consent from the instructor if you would pick a programming language of your choice.

Prerequisites:

Mathematics 20F (Linear Algebra) or Mathematics 31AH (Honors Linear Algebra), and Mathematics 180A (Introduction to Probability) or ECE 109 (Engineering Probability & Statistics), and COGS 109 (Modeling and Data Analysis) or CSE 11 (Introduction to Computer Science & Object-Oriented Programming: Java), or consent of instructor.

Grading policy:

Assignments: 35%

Midterms: 40%

Final project: 20%

Classroom participation: 5%

Bonus points: 5%

Late policy: every 5% of the total points will be deducted for every extra day past due for the homework assignments and the final project.

Policy on Integrity of Scholarship

Calendar and Class Notes

Date and Lecture #	Topic
Week 1: 04/01/17 (Monday)	Course overview, introduction to machine learning, real-world applications and impacts, cognitive science applications (pdf)
04/05/17 (Wednesday)	Review of linear algebra Probability theory (Events, Fields, Independence)
04/07/17 (Friday)	Joint probability, Conditional probability, Entropy, mutual information, KL divergence
Week 2: 04/10/17 (Monday)	Function estimation, maximum likelihood estimation
04/12/17 (Wednesday)	Error metrics,
04/14/27 (Friday)	Optimization, convexity
Week 3: 04/17/17 (Monday)	Gradient descent,
04/19/17 (Wednesday)	Least square estimation
04/21/17 (Friday)	Linear regression
Week 4: 04/24/17 (Monday)	Logistic regression
04/26/17 (Wednesday)	Midterm 1
04/28/17 (Friday)	Linear Discriminant Analysis,
Week 5: 05/01/17 (Monday)	Complexity, VC-dimension, Structural Risk Minimization
05/03/17 (Wednesday)	Support Vector Machine

05/05/17 (Friday)	Support Vector Machine
Week 6: 05/08/17 (Monday)	Support Vector Machine
05/10/17 (Wednesday)	Cross-validation, Lagrange multiplier
05/12/17 (Friday)	Kernels
Week 7: 05/15/17 (Monday)	K-nearest neighborhood classifier
05/17/17 (Wednesday)	Decision Tree
05/19/17 (Friday)	K-d tree
Week 8: 05/22/17 (Monday)	Bagging
05/24/17 (Wednesday)	Boosting
05/26/17 (Friday)	Midterm 2
Week 9: 05/29/17 (Monday)	Memorial Day Observance
05/31/17 (Wednesday)	Random forests
06/02/17 (Friday)	PAC theory
Week 10: 06/05/17 (Monday)	Perceptron, Neural networks
06/07/17 (Wednesday)	Convolutional neural networks
06/09/17 (Friday)	Modern systems