STATS 161/261: Introduction to Pattern Recognition and Machine Learning

Course Description

This course provides an accessible introduction to machine learning aimed at advanced undergraduate and graduate students in statistics, computer science, electrical engineering or related disciplines. Topics covered include Bayes decision theory, parameter estimation, regression, PCA, K-means, SVMs, deep learning, neural networks, and hidden Markov models. Emphasis is on learning high-level concepts behind machine learning algorithms and gaining practical experience applying them to real data and problems.

Details

- Instructor: Allie K. Fletcher, akfletcher@ucla.edu
- Mondays and Wednesdays, 3:30pm – 4:45pm, Public Affairs Bldg 2238
- TA: Ruiqi Gao, ruiqigao@ucla.edu
- Office Hours:
  - Fletcher: Mondays and Wednesdays, 4:45pm – 5:45pm after class
  - Gao: Tuesdays 12:00pm – 1:00pm and Fridays 2:00pm – 3:00pm in BH 9406
- Pre-requisite:
  - UCLA STATS100B, Math 33A or equivalent
  - Students should be comfortable in programming in some language
- Tentative grading plan
  - C261: Midterm 20%, Final 35%, HW and labs 25%, Quizzes & Participation 10%, Project 10%,
  - C161: Midterm 20%, Final 35%, HW and labs 35%, Quizzes & Participation 10%
- Homework will include Python based assignments
- Midterm: Tentatively, May 8 in class
- Final Exam: Thursday, June 13, 3:00pm – 6:00pm (regular time as assigned by registrar)

Texts

- Required:
  - Alpaydin. Introduction to Machine Learning, 3rd edition
  - Bishop. Pattern Recognition and Machine Intelligence.
- Recommended:
  - Hastie, Tibshirani, Friedman. Elements of Statistical Learning.
  - Murphy. Machine Learning. "A Probabilistic Perspective.
  - Shai Shalev-Shwartz and Shai Ben-David, "Understanding Machine learning from theory to algorithms

Tentative Topics Covered

- Bayes classification, Bayes risk, ROC
- Multi-variate models and multiple linear regression
- Bias, variance and model selection
- Linear discrimination, logistic regression
- Nonlinear optimization
- Dimensionality reduction, PCA
- Clustering, k-means, EM
- SVM
- Multi-layer perceptrons
- Convolutional neural networks and deep learning

Tentative Lecture Schedule

Day	Date		Topic, notes, etc.	Readings	
				Alpaydin	Bishop
M	1-Apr	1	ML Intro: Supervised and Unsupervised Learning Overview		
			Note: good review/reference materials for probability	Ch 1	Ch 1
W	3-Apr	2	Probabilistic Classification, LRT, Max Likelihood, Bayes, High Dimensional Intuitions	Appendix A, 5.4	1.2, 2.1-2.3, App B
M	8-Apr	3	Supervised Learning: Probabilistic Methods versus Bayes Empirical Risk, PAC learning	2.1,2.6, 4.6, 5.8	3.1, 3.3.1
W	10-Apr	4	Supervised Learning: Multivariate Regression and Least Squares	3.1-3.4	1.3, 1.4
M	15-Apr	5	Bias-Variance, Model selection, Cross-Validation	2.3-2.8,4.1-4.2	1.5
W	17-Apr	6	Regularization and LASSO	4.3-4.8	1.3, 1.4, 3.2
M	22-Apr	7	Linear classifiers, Logistic regression, SVM	5.6	3.1.4
W	24-Apr	8	PCA, LDA and Factor Analysis	11.1-11.4	4.1.7
M	29-Apr	9	Clustering and K-Means	6.1 to 6.6	4.1.1-4.1.6, 12.1
W	1-May	10	Document clustering, EM and Latent Semantic Analysis	7.1 to 7.3	9.1-9.3
M	6-May	11	Midterm Review	7.1-7.4	9.2.2, 9.3-9.4
W	8-May		Midterm		
M	13-May	12	Nonparametric Methods--Density Estimation and Classification		
W	15-May	13	KNN, Nonparametric Regression, Perceptron and linear discrimination	8.1-8.5	2.5
M	20-May	14	Gradients, multi-layer perceptrons, and training (backpropagation)	Ch 10	2.5.2
W	22-May	15	Convolutional neural networks	not in text	5.1-5.3
M	27-May		Holiday - Memorial Day		
W	29-May	16	Neural network training examples, image classification with convolutional neural nets	not in text	5.5.6
M	3-Jun	17	Deep convolutional nets	not in text	5.3
W	5-Jun	18	Final review	not in text	not in text
M	10-Jun		Final exam		