

# Stat 231 / CS 276A

## Pattern Recognition and Machine Learning

Fall 2018, MW 3:30-4:45 PM, Franz Hall 1260  
[www.stat.ucla.edu/~sczhu/Courses/UCLA/Stat\\_231/Stat\\_231.html](http://www.stat.ucla.edu/~sczhu/Courses/UCLA/Stat_231/Stat_231.html)

### Course Description

This course introduces fundamental concepts, theories, and algorithms for pattern recognition and machine learning, which are used in computer vision, speech recognition, data mining, statistics, information retrieval, and bioinformatics. Topics include: Bayesian decision theory, parametric and non-parametric learning, data clustering, component analysis, boosting techniques, support vector machine, and deep learning with neural networks.

### Prerequisites

You don't have to take exactly these courses as long as you know the materials.

- Math 33A Linear Algebra and Its Applications, Matrix Analysis.
- Stat 100B Intro to Mathematical Statistics.
- CS 180 Intro to Algorithms and Complexity.
- Programming skills in Matlab or Python.

### Textbook

Textbook is not mandatory if you can understand the lecture notes and handouts.

- R. Duda, et al., *Pattern Classification*, John Wiley & Sons, 2001. [Good for CS students]
- T. Hastie, et al., *The Elements of Statistical Learning*, Springer, 2009. [Good for Stat students]
- C. Bishop, *Pattern Recognition and Machine Learning*, Springer, 2006. [with advanced materials]

### Instructors

- Prof. [Song-Chun Zhu](mailto:sczhu@stat.ucla.edu), [sczhu@stat.ucla.edu](mailto:sczhu@stat.ucla.edu), Office: Boelter Hall 9404  
Office Hours: Tuesday 2-3pm
- TAs for STAT 231A: [Ruiqi Gao](mailto:ruiqigao@ucla.edu), [ruiqigao@ucla.edu](mailto:ruiqigao@ucla.edu), [Erik Nijkamp](mailto:enijkamp@ucla.edu), [enijkamp@ucla.edu](mailto:enijkamp@ucla.edu).  
TAs for CS276A: Luyao Yuan, [yuanluxu@cs.ucla.edu](mailto:yuanluxu@cs.ucla.edu)  
Office Hours: Session 1: Tuesday 6-7pm and Session 2: 7-8pm for project overview and assistant.  
Individual questions will be answered by the readers through CCLE online.

### Grading Plan

<b>Five projects:</b>	
• <a href="#">0. Exercise on Classification using Deep Learning and CNN</a> Experience the training and testing of problems in machine learning	10%
• <a href="#">1. Face modeling by AAM Model, Auto-encoder and FLD</a> How many bits do you need to represent a face? Or telling apart male from female faces?	15%
• <a href="#">2. Face detection by AdaBoost and RealBoost</a> How many features do you need to detect a face in a crowd?	15%
• <a href="#">3. Object detection by Fast-RCNN</a> How to detect the class of an object and its location?	10%
• <a href="#">4. Face social attributes and sentiment analysis by SVM</a> How do we measure the social dimensions of faces in political elections and social network?	10%

## Grading Policy

- You are encouraged to work and discuss in a group, but each person must finish his/her own project. Submit your report of the project, and your code through the CCLE website.
- You have a total of 5 late days (not including weekends) for the entire class (5 projects) to cover your various reasons, but after using the four late days, no credit will be given for late homework/project.

## Schedule

Tentative Schedule for 2018 (Course materials will be posted on the CCLE site.)

Lecture	Topics	Handouts
1	<b>Introduction to Pattern Recognition</b> [Problems, applications, examples]	
2	<b>Bayesian Decision Theory I</b> [Bayes rule, discriminant functions]	
3	<b>Bayesian Decision Theory II</b> [loss functions and Bayesian error analysis]	
4	<b>Component Analysis and Dimension Reduction I:</b> [PCA, face modeling by Active Appearance Model (AAM), Auto-encoder]	
5	<b>Component Analysis and Dimension Reduction II</b> [Fisher Linear Discriminant ] [Multi-dimensional scaling (MDS)]	
6	<b>Component Analysis and Dimension Reduction III</b> [Local Linear Embedding (LLE), Intrinsic dimension]	
7	<b>Boosting Techniques I</b> [perceptron, backpropagation and Adaboost]	
8	<b>Boosting Techniques II</b> [RealBoost and Example on face detection]	
9	<b>Boosting Techniques III</b> [analysis, logit boost, cascade and decision policy]	
10	<b>Boosting Techniques III</b> [analysis, logit boost, cascade and decision policy]	
11	<b>Non-metric method I</b> [Decision tree and random forest]	
12	<b>Non-metric method II</b> [Syntactic pattern recognition and example on human parsing]	
13	<b>Support vector machine I</b> [Kernel-induced feature space]	
14	<b>Support vector machine II</b> [Support vector classifier]	
15	<b>Support vector machine III</b> [Loss functions, Latent SVM]	
16	<b>Parametric Learning</b> [ Maximum Likelihood Estimation ] [ Sufficient Statistics and Maximum entropy]	
17	<b>Non-parametric Learning I</b> [ Parzen window and K-nn classifier]	
18	<b>Non-parametric Learning II</b> [K-nn classifier and Error analysis]	
19	<b>Deep Learning I</b>	
20	<b>Deep Learning II</b>	