MATH156: MACHINE LEARNING

Spring 2021

GENERAL INFORMATION

Instructor	Hanbaek Lyu	(Email: `hlyu@math.ucla.edu`, Office: MS 6156)
Lectures	MWF 3:00PM - 3:50PM using zoom (Link posted on CCLE)	Course webpage
Office hours	(tentative)T 2:00PM - 4:00PM	
Textbook	Bishop, Christopher M, *Pattern Recognition and Machine Learning*, Springer, 2006 Lecture notes will be provided on CCLE Supplementary Python codes and Juypiter notebooks will be provided in the course repository	
Prerequisites	115A, 164, 170A or 170E or Statistics 100A, and CS 31 or Program in Computing 10A. Strongly recommended requisite: Program in Computing 16A or Statistics 21.	
TA	Yushan Han	(Email: `yushanh1@ucla.edu`)

COURSE DESCRIPTION

Introductory course on mathematical models for pattern recognition and machine learning. Topics include parametric and nonparametric probability distributions, curse of dimensionality, correlation analysis and dimensionality reduction, and concepts of decision theory. Advanced machine learning and pattern recognition problems, including data classification and clustering, regression, kernel methods, artificial neural networks, hidden Markov models, and Markov random fields. Projects in Python to be part of final project presented in class. P/NP or letter grading.

GRADING

- Final course score will be computed by following scheme:

  Scheme: Homeworks (30%) + Final (40%) + Final project (30%)

- All grades will be posted on Gradescope. The final course grade will be posed on MyUCLA.

HOMEWORK

- Homeworks will be assigned weekly on every Wednesdays on Gradescope, and are due at the beginning of the class on the following Wednesday. Submit as a PDF file on Gradescope.
- No late homework will be accepted.
- Two lowest homework scores will be dropped.
- A random sample of problems will be graded by the TA.
- Solutions on some selected problems will be posted on the course website.
- Discussing homework problems with the instructor, TA, or classmates are encouraged. But you need to write your own solution with your own understanding.
- Some homeworks may contain Python implementations in Jupyter notebook. To get started with Python and Jupyter notebook, see:
  1. Jeff Heaton, "2020, Installing TensorFlow 2.0, Keras, & Python 3.7 in Mac OSX" (link)
  2. Jeff Heaton, "2019, Installing TensorFlow, Keras, & Python 3.7 in Windows" (link)
  3. Additional resources: [1] [2]

EXAMS

- There is no midterm and one final exam (open book, take-home tests)

  Final: Friday, 6/11 9 AM - Saturday, 6/12 9 AM

  (Exam will be released on Gradescope at 9AM. The exam will expire after 24 hours once you open it. Submit a PDF scan of your solution on Gradescope. Deadeline for submission is Saturday, 6/12 AM.)

- There is no make-up exam. You should attend the final exam to pass the course.

FINAL PROJECT

- By the end of the second week of the class, students will form groups with no more than 4 students. Email the instructor to notify if you form a group. By the end of the second week, remaining students will be randomly assigned to groups of 3-4 students.
- Groups will choose dataset of interest, and objective of data analysis using machine learning (e.g., classification of MNIST data of hand-written digits via naive Bayes classifier or RBM; Image classification via PCA or NMF; Natural language processing via RNN).
- Example sourses of data: (1) UCI Machine Learning Repository; (2) kaggle
- Structure of the final paper:

  Introduction: Summary of the problem, methods, related results with references.

  Problem statement: Precise description of the problem you are trying to address. Also discuss why addressing this problem is important.

  Methods: Detailed description of methods used or developed.

  Theory: A clear and convincing discussion of the theoretical properties of the proposed method. It should also discuss under what assumptions the methods should work and under what conditions they may fail.

  Dataset description: Detailed description of the dataset being analyzed.

  Resuls: The results of applying the methods to the data set. Also discuss why the results makes sense, possible implications, and how it compares to the literature.

  Simulation studies*(optional): Simulation studies. Results of applying the method to simulated data sets.

  Conclusion: What is the answer to the question? What did you learn about the methods? What is the contribution of the paper and what are the open problems?

- Submit a PDF scan of the final paper as well as the corresponding Pychon notebook on Gradescope by Saturday, 6/12 5 PM. The notebook should reproduce the results in the final paper. No late submission will be accepted.

TENTATIVE COURSE SCHEDULE

Below is a tentative course schedule based on the departmental guideline. There could be a slight change depending on our progress.

Week	Date	Section	Topics
1	M 3/29	1.2.1-1.2.4	Review of probability
	W 3/31	1.1, 1.2.5, 1.2.6	Polynomial curve fitting
	F 4/2	1.1, 1.2.5, 1.2.6	Polynomial curve fitting
2	M 4/5	1.5	Decision theory
	W 4/7	1.6	Information theory
	F 4/9	2.3	The Gaussian distribution
3	M 4/12	2.4, 2.5	The exponential family, Nonparametric methods
	W 4/14	3.1	Linear Basis Function Models
	F 4/16	3.2	The Bias-Variance Decomposition
4	M 4/19	3.3	Bayesian Linear Regression
	W 4/21	3.5	The Evidence Approximation
	F 4/23	4.1	Discriminant Functions
5	M 4/26	4.2	Probabilistic Generative Models
	W 4/28	4.3	Probabilistic Discriminative Models
	F 4/30	4.5	Bayesian Logistic Regression
6	M 5/3	5.1	Feed-forward Network Functions
	W 5/5	5.2	Network Training
	F 5/7	5.3	Error Backpropagation
7	M 5/10	9.1	K-means Clustering
	W 5/12	9.2	Mixtures of Gaussians
	F 5/14	9.3	The EM algorithm
8	M 5/17	12.1	Principal Component Analysis
	W 5/19	12.2	Probabilistic PCA
	F 5/21	12.4	Nonlinear Latent Variable Models
9	M 5/24		Nonnegative Matrix Factorization
	W 5/26	8.1	Bayesian Networks
	F 5/28	8.3	Markov Random Fields
10	M 5/31		No class (Memorial Day)
	W 6/2	8.4	Inference in Graphical Models
	F 6/4		Review
11	F 6/11		**Final**
	Sat 6/12		**Final paper due** (5PM)