Noiselab@UCSD Home News Publications Gerstoft Teaching Group Group-pictures Computing@Noiselab

ECE228 and SIO209 Machine learning for physical applications, Spring 2019

Professor Peter Gerstoft, Spiess Hall 462, Gerstoft@ucsd.edu
TA Siva Prasad Varma Chiluvuri
TA Harshuk gupta
Location: CENTR 113
Time: Monday and Wednesday 5-6:20pm

Sylabus: Machine learning has received enormous interest. To learn from data we use probability theory, which has been the mainstay of statistics and engineering for centuries. The class will focus on implementations for physical problems. Topics: Gaussian probabilities, linear models for regression, linear models for classification, neural networks, kernel methods, support vector machines, graphical models, mixture models, sampling methods, sequential estimation.
Prerequisites: graduate standing.

It is not a computer science class, so we go slowly through the fundamentals to appreciate the methods, implement these. We will discuss their use in Physical Sciences. In the first part of the class we focus on theory and implementations. We will then transition to focus on machine learning final projects.

Books: Main book: Chris M Bishop Pattern Recognition and Machine Learning . A third party Matlab implementation of many of the algorithms in the book.
Other good books:
Hastie and Tibshirani The Elements of Statistical Learning
Kevin P. Murphy: Machine Learning: A Probabilistic Perspective. UCSD license . Matlab codes used in Murphy's book.

Online resources: While not required, I recommend taking these classes. Both are online classes are excellent.
Statistical Learning by Hastie and Tibshirani. My favorite class.
Andrew Ng's Coursera class, Machine learning. This was the first class offered by Coursera.

Grading: Full scale of the letter grade. Grade consists of About 50 % homework and 50% final-project (10% poster, 10 % code and 30 % Report). Your and my purpose is to lean, so a good effort is sufficient.

Homework Cody homework will be graded. I'm a strong believer in leaning by doing. Thus we will have computer-based homework each week.

Final project: Propose a topic before May 1. Otherwise it will be based on my paper: Niu et al, 2017 on arXiv. We will make teams on April 26. Report due June 16.
Suggested topics:

1. Tracking ships using acoustics. Based on my paper: Niu et al, 2017 on arXiv. This can be solved in TensorFlow or SciKit-learn (in Python) or matlab . Data and SVM example
2. Graph signal processing for localizing small events. Based on my paper Riahi 2017.
3. Classifying plankton. This would be based on Jules Jaffe's underwater microscope. This might require convolutional networks, but random forest and support vector machines would also work.
4. Identifying earthquakes from data on a mobile phone (example is in ipython). Extensive example with 3 Gb of data. See background paper in Science

Schedule

Lecture	Date	Topic	Required reading	Assignments
1	April 1	**Introduction, probability theory** [Slides]	Bishop CH 1.0-1.2	
2	April 3	**Gaussian probability theory** [Slides, Annotated Slides]	Bishop CH 1.5, 1.6, 2.3	
3	April 8	**Non Parametric methods, linear models for Regression** [Slides Annotated Slides]	Bishop CH 2.5, 3.1, 3.2	Homework 1 due before class
4	April 10	**Linear models for Regression** [Slides Annotated Slides]	Bishop CH 3.0-3.3	Homework 2 due before class
5	April 15	**Sparse models I, Lecture by Santosh Nannuru** [Slides Annotated Slides] After class we install Tensorflow, Python, Jupyter notebook	Read Chapter 2 for an introduction to sparse problems Murphy 13.1, 13.3, 13.6.1	Homework 3 due before class
6	April 17	**Sparse models II,** [Slides Annotated Slides] Emma Ozanich: Tracking ships using acoustics. Based on our paper: Niu et al, 2017	Murphy 13.7, 13.8	Homework 1 and 2 due 4/20
7	April 22	**Linear models for Classification** [Slides Annotated Slides]	Bishop CH 4.0-4.3.2	Homework 4 due before class
8	April 24	**Backpropagation** [Slides Annotated Slides] Eric Orenstein: Automatic Analysis of Planktonic Image Data Sets	Bishop CH 4.3.4, 5.0-5.2	
9	April 29	**Kernel methods** [Slides Annotated Slides]	Bishop CH 6 and from last lecture 5.3 to 5.3.2	Homework 5 due before class
10	May 1	**Support vector machines** [Slides Annotated Slides]	Bishop CH 7	Homework 6 due before class solution
11	May 6	**K-means, EM and Mixture models**[Slides Annotated Slides]	Bishop CH 9	**Project proposals due**
12	May 8	Mike Bianco: Dictionary leaning **Trees**	Read Bianco's Dictionary learning paper	Homework 7 due before class Solution
13	May 13	**ICASSP Trees, Random Forrest, Gradient boosting** [Slides Annotated Slides] Bayes net	Hastie CH 9 Bishop CH 8.1, 8.2	
14	May 15	**ICASSP Graphs, maybe PCA** [Slides Annotated Slides]	Bishop CH 8.2 12	Homework 8 due before class
15	May 20	**mid-project presentation of final-projects**		
16	May 22	**Sequential estimation**[Slides Annotated Slides]	Bishop CH 13	Homework 9 due before class
17	May 29			
18	June 3			
19	June 5	**Final project poster presentation, Jacobs Hall Lobby, 5 pm**		Final report due Saturday 15 June.

Spring 2018 class

Spring 2017 class