Fall 2015 ECE 532 Theory and Applications of Pattern Recognition

ECE 532 is a introduction to machine learning and pattern recognition that focuses on matrix methods and features real-world applications, ranging from classification and clustering to prediction and data analytics.

Overview: Students will be exposed to real-world applications in machine learning and pattern recognition, learn to translate physical problems into mathematical models, and develop basic computational skills for problem-solving. The proposed course tackles this aim by studying the basics of linear algebra in the context of applications such as handwritten character recognition, neural networks, and collaborative filtering.

Lectures: Monday and Wednesday, 11:00am-12:15pm, 2317 Engineering Hall
Instructors: Robert Nowak and Laurent Lessard
Office Hours: TBA
Assistants: TBA

Textbook: **Matrix Methods in Data Mining and Pattern Recognition**
by Lars Elden

http://epubs.siam.org/doi/book/10.1137/1.9780898718867

Grading and Exams: The course grade will be determined by a combination of exams, projects and homework, and course participation according to the distribution:

Exam 1, Thursday October 8, 7:15-9:15 PM (**15%**)
Exam 2, Thursday November 19, 7:15-9:15 PM (**15%**)
Projects, (**45%**)
Homework Assignments, usually 1 per week (**15%**)
Course Participation (**10%**)

Programming Assignments and Projects: Homework and the project component of the course will investigate theory, methods and applications using Matlab or Python. Students familiar with Matlab are encouraged to try-out Python (and vice-versa for Python users). Matlab widely used in industry and academia for scientific data analysis and engineering development. Python is a free, open-source, general-purpose, high-level programming language that is also widely used in industry.

Course Syllabus (tentative)

Week 1	Vectors and Matrices in Machine Learning
Week 2	Vectors and Matrices (Bioinformatics examples)
Week 3	Linear Systems and Least Squares (Face recognition)
Week 4	Linear Independence and Orthogonality (Classification)
Week 5	Singular Value Decomposition (Principal Component Analysis)
Week 6	Midterm
Week 7	Singular Value Decomposition (Dimensional Reduction Methods)
Week 8	Ridge Regression and Kernel Methods
Week 9	The Lasso and Proximal Point Algorithms
Week 10	Beyond Least Squares: Alternate Loss Functions
Week 11	Support Vector Machines and Hinge Loss
Week 13	Stochastic Gradient Descent Algorithms
Week 14	Unsupervised Learning and Clustering
Week 15	Neural Networks and Deep Learning

Possible Project Topics

Recommender Systems and Collaborative Filtering
http://www.slideshare.net/erikbern/collaborative-filtering-at-spotify-16182818

Matrix Completion
http://statweb.stanford.edu/~candes/papers/SVT.pdf

Nonlinear Dimensionality Reduction
http://www.beermapperapp.com

Support Vector Machines
http://svms.org/tutorials/Moore2001.pdf

Deep Learning
http://arxiv.org/abs/1406.3332

Sparse Coding and Dictionary Learning
http://www.di.ens.fr/willow/pdfs/icml09.pdf

Active Learning
http://nowak.ece.wisc.edu/next_template.html

Neuronal Spike Sorting
http://www.scholarpedia.org/article/Spike_sorting

Sparse Methods for Machine Learning
http://www.di.ens.fr/~fbach/nips2009tutorial/nips_tutorial_2009_sparse_methods.pdf

Topic Modeling
http://www.cl.uni-heidelberg.de/courses/ss12/topicmodels/intro.pdf

Independent Component Analysis
http://www.stat.ucla.edu/~yuille/courses/Stat161-261-Spring14/HyvO00-icatut.pdf

Spectral Clustering
http://cs.nyu.edu/~dsontag/courses/ml14/notes/Luxburg07_tutorial_spectral_clustering.pdf

Climate Data Analysis
http://www.princeton.edu/~rvdb/tex/LocalWarming/LocalWarming.pdf

Image Segmentation
http://www.cis.upenn.edu/~jshi/GraphTutorial/

Anomaly Detection
http://www.cs.bu.edu/faculty/crovella/paper-archive/sigc04-network-wide-anomalies.pdf

Deconvolution and Deblurring
http://www.mathcs.emory.edu/~nagy/courses/fall06/ID_lecture1.pdf

Genomic Data Analysis and Classification
http://public.lanl.gov/mewall/kluwer2002.html

Spectral Learning Algorithms for Natural Language Processing
http://www.cs.columbia.edu/~scohen/naacl13tutorial/