ENEE759G: Advanced Topics in Computer Engineering - Unsupervised Learning

Spring 2017 (TuTh 2-3:15; CSS 1114)
Instructor: Joseph JaJa
Course Syllabus

Course Objectives: The course will cover core statistical machine learning techniques for unsupervised learning. Topics covered will include: density estimation, latent variable models, mixture models, clustering, directed and undirected graphical models and inference, learning graphical models, and generative deep models.

Relationship with ENEE 633: The focus of this course is not intended for those interested in taking ENEE633. It is strongly advised that students enroll in **only one** of ENEE633 and ENEE759G.

Course prerequisites: Graduate standing.

Prerequisite topics: A strong undergraduate background in probability and statistics, linear Algebra, advanced calculus, algorithms, and nonlinear optimization is required for this course.

Textbooks: No textbook is required for this course but the following can be used as references:

T. Hastie, R. Tibshirani, and J. Friedman, The Elements of Statistical Learning, second edition, Springer, 2009.

K. Murphy, Machine Learning: A Probabilistic Perspective, The MIT Press, 2012.

Core Topics:

1. **Introduction**
    - Basic framework and concepts
    - Probability density estimates for low-dimensional data
    - Maximum Likelihood Estimation and Bayesian Estimation.
    - Matrix Factorization and Nonlinear Optimization
2. **Latent Variable Models**
    - Principal Component Analysis
    - Introduction to Factor Analysis
    - Independent Component Analysis
    - The EM algorithm with application to Gaussian Mixture Models
3. **Clustering**
    - Proximity measures and evaluation methodologies
    - The k-means algorithm and its variant the k-medoid algorithm
    - Hierarchical clustering
    - Spectral clustering
    - Overview of other types of clustering algorithms such as Self Organizing Maps, density based clustering, or affinity propagation.

4. **Directed Graphical Models**
    - Basic definitions and properties
    - Naïve Bayesian networks
    - Inference: Exact and Approximate
    - Learning Bayesian networks

5. **Undirected Graphical Models**
    - Basic definitions and concepts
    - Markov properties
    - Factor graphs
    - Inference algorithms
6. **Deep Generative Models**
    - RBMs and deep belief networks
    - Deep auto-encoders
    - Recurrent neural networks

Midterm: March 16 – 40%; Final: Date TBA – 60%

Homeworks: Assignments will be given out throughout the semester and will be corrected but not graded. The assignments will provide a good technical background to the material covered in class, and some of the assignment problems may appear on the midterm or final exams.

Contact Information

Instructor:	Joseph JaJa
Office:	3433 A. V. Williams
Office Hours:	Tu and Th 4-5 or by appointment
Email:	Joseph@umiacs.umd.edu
Phone:	405-1925