ENEE759G: Advanced Topics in Computer Engineering - Unsupervised Learning

Fall 2019 (TuTh 2-3:15 AJC2119)
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, nonlinear dimensionality reduction strategies, mixture models, clustering, directed and undirected graphical models and inference, and generative deep models.

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 reference materials will be posted on the ELMS web site of the class. The following textbooks can also serve 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.

S. Theodoridis and K. Koutroumbas, Patter Recognition, fourth edition, Elsevier, 2009.

Core Topics:

1. **Introduction and Background**
    - Basic framework and concepts
    - Nonparametric Probability density estimation
    - Maximum Likelihood Estimation and Bayesian Estimation.
    - Matrix Factorization (including SVD) and Nonlinear Optimization
2. **Latent Variable Models**
    - Principal Component Analysis
    - Introduction to Factor Analysis
    - Independent Component Analysis
    - Nonnegative Matrix Factorization
    - The EM algorithm with application to Gaussian Mixture Models
3. **Nonlinear Dimensionality Reduction**
    - Kernel PCA
    - Laplacian Eigenmaps
    - Local Linear Embedding
    - Isometric Mapping
4. **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, and affinity propagation.

5. **Directed Graphical Models**
    - Basic definitions and properties
    - Naïve Bayesian networks
    - Inference: Exact and Approximate
6. **Undirected Graphical Models**
    - Basic definitions and concepts
    - Markov properties
    - Factor graphs
    - Inference algorithms
7. **Deep Generative Models**
    - RBMs and deep belief networks
    - Deep auto-encoders

Midterm: Date TBA 40%; Final: Date TBA – 50%

Homeworks: Assignments will be given out throughout the semester, will be corrected, and will contribute **10%** toward the final score. The assignments will provide a complementary 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:	3214 Iribe
Office Hours:	Tu and Th 4-5 or by appointment
Email:	josephj@umd.edu
Phone:	405-1925