

CS 675 final exam review sheet

1. Feature selection: chi-square, F-score, Pearson, multivariate
 - a. Given the formula rank features by the statistics
2. Kernel methods – kernel k-means
 - a. Distance between two points in kernel space
 - b. Kernel version of nearest means and nearest neighbors
3. Dimensionality reduction: Principal component analysis, maximum margin criterion
 - a. Variance and means of projected data
 - b. Laplacian of graph – how does Laplacian connect to dimensionality reduction and clustering?
 - c. Relaxing an NP-hard problem into a smooth objective
4. Clustering: k-means, graph methods
 - a. Basic min-cut approach to clustering
 - b. Spectral clustering (Shi and Malik)
 - c. K-means algorithm – run on a simple example
5. Empirical and regularized risk
 - a. Loss functions
 - b. Effect of L1 vs L2 norm regularization
6. Decision trees
 - a. Simple example on how to split a feature to create a new node (greedy algorithm)
 - b. CART algorithm
7. Ensemble methods:
 - a. Bagging, Random forest
 - b. Boosting – Basic Adaboost algorithm
8. Regression
 - a. Linear regression, support vector regression, regression with decision trees
 - b. Regularized regression
9. HMMs: four problems
 - a. Calculate probability of hidden and observed state sequence
 - b. Find the most probable hidden sequence (Viterbi)
 - c. Calculate probabilities from hidden sequences (maximum likelihood)
 - d. Baum-Welch (forward probabilities, total probability, EM algorithm)
10. Runtimes:
 - a. One iteration of gradient descent, one iteration of K-means
 - b. Viterbi algorithm for HMMs
 - c. CART (one recursion)
 - d. Naïve-Bayes
11. Neural networks
 - a. Optimization objective for a simple perceptron (least squares)
 - b. Objective for a single layer neural network
12. Convolutional neural networks
 - a. Convolution kernel
 - b. Pooling
13. Multiple choice
 - a. Comparison of classifiers
 - b. Big data methods
14. Advanced topics
 - a. Random projections for dimensionality reduction and classification
 - b. What are the properties of random projections and signed random projections?
 - c. Method for learning features based on k-means for image data