Introduction to Machine Learning

Brown University CSCI 1950-F, Spring 2011 Prof. Erik Sudderth

Lecture 18: Clustering, Mixture Models, Expectation-Maximization (EM) Algorithm

> Many figures courtesy Kevin Murphy's textbook, Machine Learning: A Probabilistic Perspective



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning

Reconstruction Error versus Iteration



C. Bishop, Pattern Recognition & Machine Learning



For compressing new data, more codewords is always better. Cross-validation fails for unsupervised learning!

Gaussian Mixture Models



Mixture of 3 Gaussian Distributions in 2D Contour Plot of Joint Density, Marginalizing Cluster Assignments

Gaussian Mixture Models

Surface Plot of Joint Density, Marginalizing Cluster Assignments

Fitting Gaussian Mixtures



Posterior Assignment Probabilities









C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



C. Bishop, Pattern Recognition & Machine Learning



We are hoping EM will find a good local optimum...

Density Shape Matters



Real Clusters may not be Round

