DAY 3

Intelligent Audio Systems:
A review of the foundations and applications of semantic audio analysis and music information retrieval

Jay LeBoeuf
Imagine Research
jay{at}imagine-research.com

Rebecca Fiebrink
Princeton University
fiebrink{at}princeton.edu

July 2010

These lecture notes contain hyperlinks to the CCRMA Wiki.

On these pages, you can find supplemental material for lectures - providing extra tutorials, support, references for further reading, or demonstration code snippets for those interested in a given topic .

Click on the  symbol on the lower-left corner of a slide to access additional resources.

WIKI REFERENCES...

Review from Day 2

- BBQ Tomorrow evening: 6-9PM
- Let's talk lab…
- AdaBoost? Cross Validation?
- AdaBoost with too easy of a problem!
- Bug with CV
- Your feature vectors (e.g., only first 100ms)

Did you get to the Weka lab? Optional lab?

- How did the lab go?
- Did you try other audio files – other instrument recognizers?

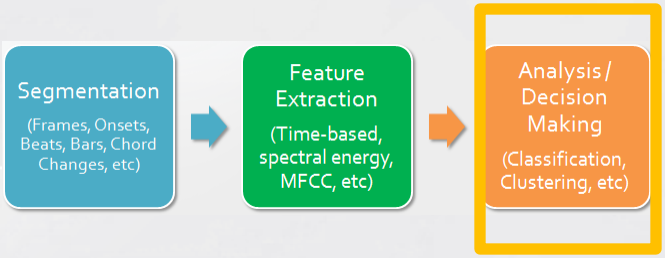

ANALYSIS AND DECISION MAKING

Supervised vs. Unsupervised

- Unsupervised - “clustering”
- Supervised – binary classifiers (2 classes)
- Multiclass is derived from binary

Clustering

- Unsupervised learning – find pockets of data to group together
- Statistical analysis techniques

Clustering

- K = # of clusters

- Choosing the number of clusters – note that choosing the "best" number of clusters according to minimizing total squared distance will always result in same # of clusters as data points.

Clustering

The basic goal of clustering is to divide the data into groups such that the points within a group are close to each other, but far from items in other groups.

Hard clustering – each point is assigned to one and only one cluster.

K-Means

The key points relating to *k-means clustering are:*

- k-means is an automatic procedure for clustering unlabelled data;
- it requires a pre-specified number of clusters;
- Clustering algorithm chooses a set of clusters with the minimum within-cluster variance
- Guaranteed to converge (eventually)
- Clustering solution is dependent on the initialization

(You get different results with each running)

Demo

- http://home.dei.polimi.it/matteucc/Clustering/tutorial_html/AppletKM.html

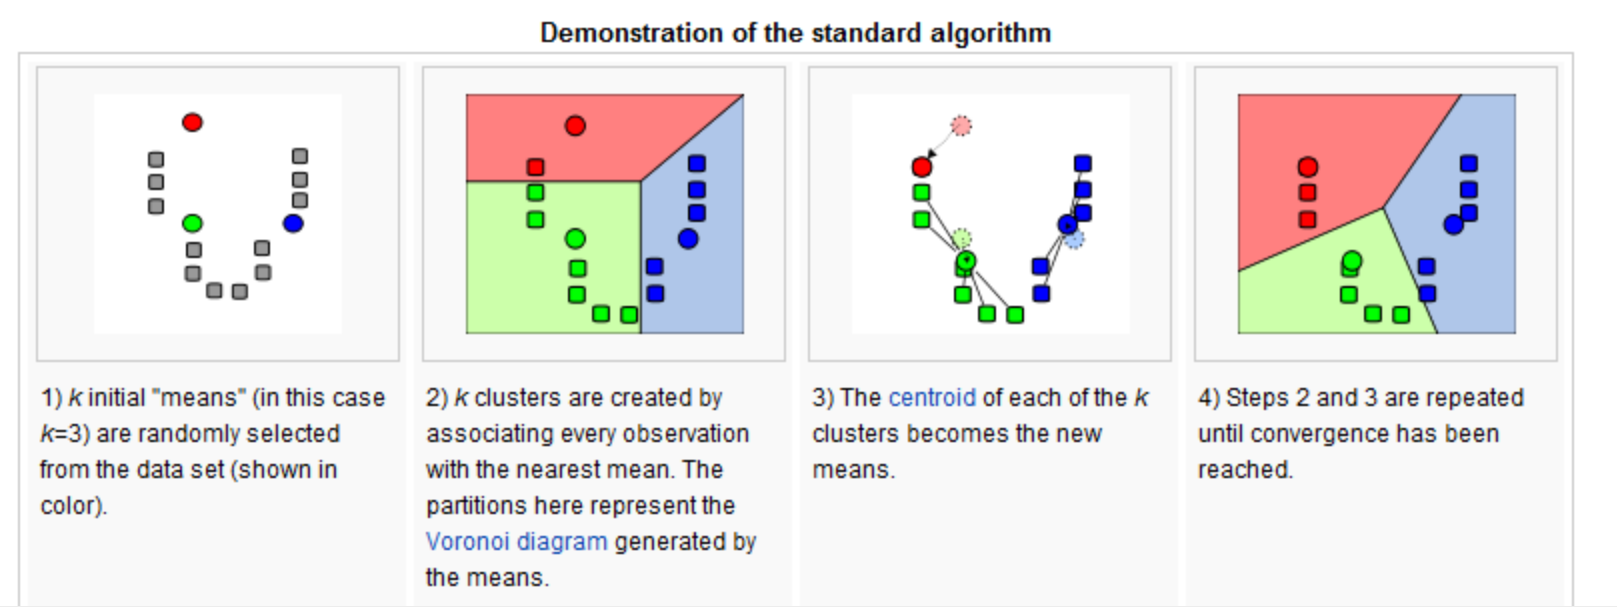

K-Means

The initialization method needs to be further specified. There are several possible ways to initialize the cluster centers:

- *Choose random data points as cluster centers*
- *Randomly assign data points to K clusters and compute means as initial centers*
- *Choose data points with extreme values*
- *Find the mean for the whole data set then perturb into k means*
- *Find ground-truth for data*