

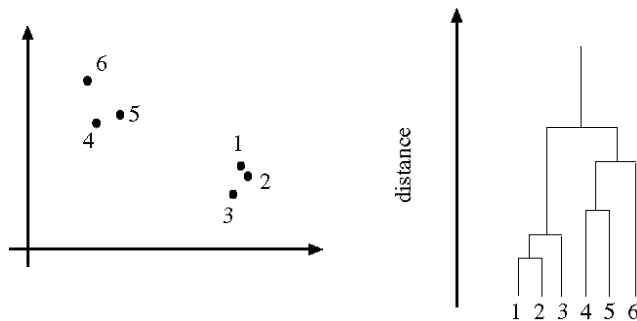
Out of order lecture given by Ranjini due to iPlant travel.

Clustering approaches

- Agglomerative clustering
 - initialize: every item is a cluster
 - attach item that is “closest” to a cluster to that cluster
 - repeat
- Divisive clustering
 - split cluster along best boundary
 - repeat
- Probabilistic clustering
 - Define a probabilistic grouping model

Simple clustering approaches

- Point-Cluster or Cluster-Cluster distance
 - single-link clustering (minimum distance from point to points in clusters or among pairs of points, one from each cluster)
 - complete-link clustering (maximum)
 - group-average clustering (average)
 - (terms are not important, but concepts are worth thinking about)
- Dendrograms
 - classic picture of output as clustering process continues



K-Means

- Choose a fixed number of clusters (“K”)
- Choose cluster centers (**means**) and point-cluster allocations (membership) to minimize the error

$$\sum_{i \in \text{clusters}} \left\{ \sum_{j \in \text{elements of } i\text{'th cluster}} \|x_j - \mu_i\|^2 \right\}$$

- \mathbf{x} 's could be any set of features for which we can compute a distance (careful with scaling)

K-Means

- Want to minimize

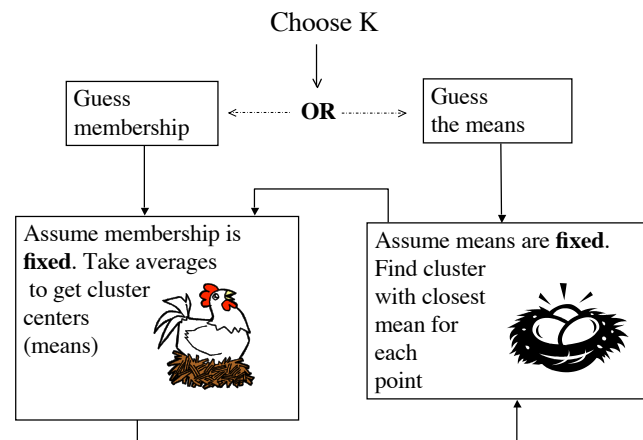
$$\sum_{i \in \text{clusters}} \left\{ \sum_{j \in \text{elements of } i\text{'th cluster}} \|x_j - \mu_i\|^2 \right\}$$

- **Cannot** do this optimization by search, because there are too many possible allocations.
- Standard difficulty which we handle with an iterative process (chicken and egg)

K-Means algorithm (intuition)

- If we know the cluster centers, the best cluster for each point is easy to compute
 - Just compute the distance to each to find the closest
- If we know the best cluster for each point, the cluster centers are also easy to compute
 - Just average the points in each cluster
- Algorithm
 - 1) Guess one of the two.
 - 2) Alternatively re-compute the values for each

K-means flow chart



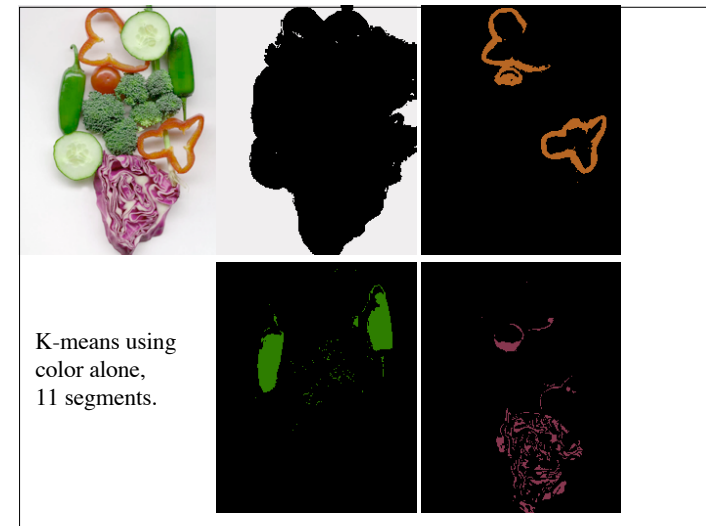
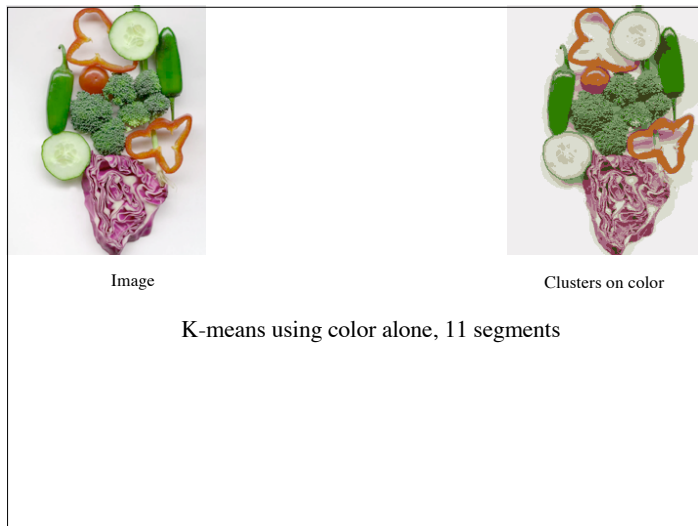
Image

Clusters on intensity

Clusters on color



K-means clustering using intensity alone and color alone
(Assuming 5 segments, i.e. k=5)

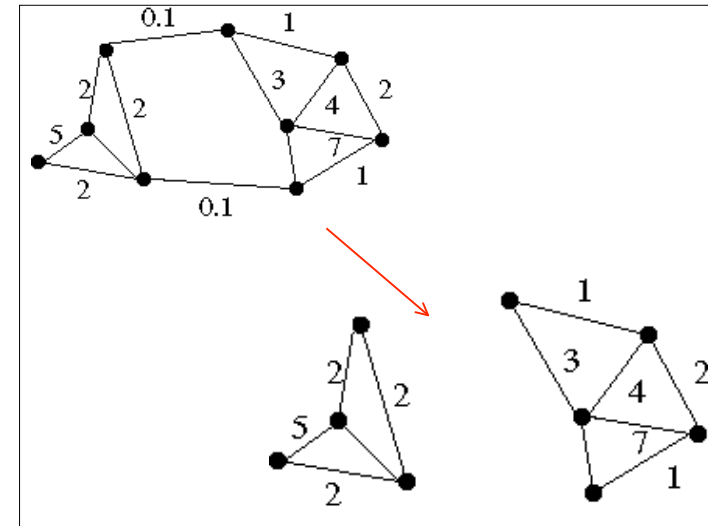
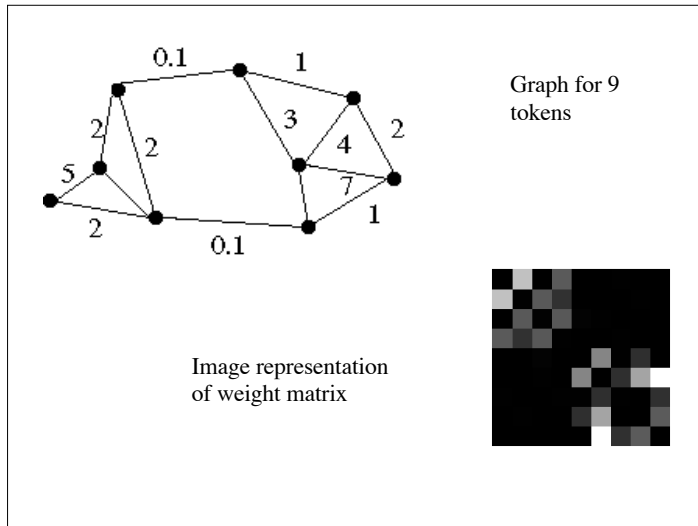


Notes on K-Means

- K-means is “hard” clustering-each point is completely in exactly one cluster
- What you get is a function of starting “guess”
- The error goes down with every iteration
 - This means you get a local minimum
- Unfortunately, the dimension of the space is usually large, and high-dimensional space have lots of local maximum (standard problem!)
 - Dimensionality here is $K \cdot \dim(\mathbf{x})$
- Finding the global minimum for a real problem is very optimistic!

Graph theoretic clustering

- Represent distance between tokens using a weighted graph.
 - affinity matrix
- Cut up this graph to get subgraphs with strong interior links (and weak links between the subgraphs).



Measuring Affinity

Intensity

$$\text{aff}(x, y) = \exp\left\{-\left(\frac{1}{2\sigma_i^2}\right)\left(\|I(x) - I(y)\|^2\right)\right\}$$

Distance

$$\text{aff}(x, y) = \exp\left\{-\left(\frac{1}{2\sigma_d^2}\right)\left(\|x - y\|^2\right)\right\}$$

Texture

$$\text{aff}(x, y) = \exp\left\{-\left(\frac{1}{2\sigma_t^2}\right)\left(\|c(x) - c(y)\|^2\right)\right\}$$

Texture Descriptor

Eigenvectors and cuts

exam

- For some cluster*, consider a vector \mathbf{a} giving the association between each element and that cluster
- We want elements within this cluster to, on the whole, have strong affinity with one another
- If two elements, i and j , are part of the same cluster, then
 - \mathbf{a}_i and \mathbf{a}_j are both large
 - and the affinity A_{ij} is large
 - thus, $\mathbf{a}_i A_{ij} \mathbf{a}_j$ should be large
- Thus a good cluster is one where $\sum_i \sum_j \mathbf{a}_i A_{ij} \mathbf{a}_j$ is large.

* In the hand out, \mathbf{a} is sub-scripted by an "i" for cluster. But this is a distraction. I use "i" for something else below.



Eigenvectors and cuts

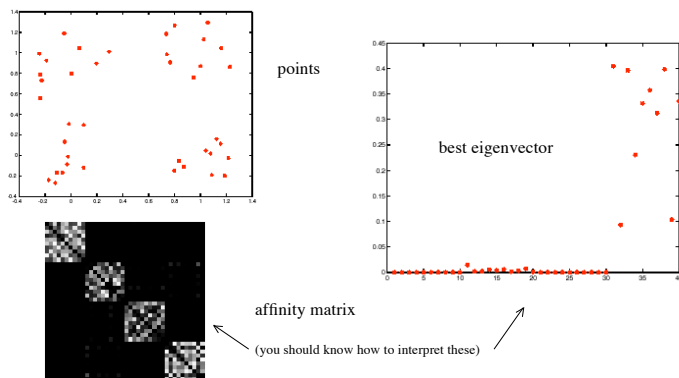
- $\sum_i \sum_j a_i a_j$ should be large for a coherent cluster represented by \mathbf{a} .
- This suggests maximizing $\mathbf{a}^T \mathbf{A} \mathbf{a}$
- But we need the constraint $\mathbf{a}^T \mathbf{a} = 1$ (why?)
 - Arguably it might be more logical to make the sum of the elements of \mathbf{a} to be one, but the standard (L_2) norm is easier to deal with.



Eigenvectors and cuts

- We want to maximize $\mathbf{a}^T \mathbf{A} \mathbf{a}$ subject to $\mathbf{a}^T \mathbf{a} = 1$
- This is an eigenvalue problem - choose the eigenvector of \mathbf{A} with largest eigenvalue
- This gives the cluster with greatest internal affinity
 - Ideally, most elements of the eigenvalue are near zero, and the others tell us which tokens are in the cluster

Example eigenvector



Normalized cuts

- Previous criterion evaluates **within** cluster similarity, but does not promote large differences **between** clusters across cluster difference
- N-cuts proposes maximizing the within cluster similarity **compared** to the across cluster difference
- Write graph as \mathbf{V} , one cluster as \mathbf{A} and the other as \mathbf{B} . ($\mathbf{V} = \mathbf{A} \cup \mathbf{B}$).
- Maximize

$$\left(\frac{\text{assoc}(\mathbf{A}, \mathbf{A})}{\text{assoc}(\mathbf{A}, \mathbf{V})} \right) + \left(\frac{\text{assoc}(\mathbf{B}, \mathbf{B})}{\text{assoc}(\mathbf{B}, \mathbf{V})} \right)$$

- (Solution follows to keep notes self-contained).

Normalized cuts

- Previous criterion evaluates **within** cluster similarity, but does not promote large differences **between** clusters across cluster difference.
- For simplicity, consider the task of splitting the tokens into two groups A and B. The union of the two groups is V.
- N-cuts proposes maximizing the within cluster similarity **compared** to the across cluster difference.
- Define $\text{cut}(A,B)$ to be the sum of the weights of the edges that you remove to split up the image.
- Define $\text{assoc}(A,V)$ to be the sum of all the weights between elements in A and elements in V.

Normalized cuts

- Two equivalent formulations

- Minimize

$$\left(\frac{\text{cut}(A,B)}{\text{assoc}(A,V)} \right) + \left(\frac{\text{cut}(A,B)}{\text{assoc}(B,V)} \right)$$

- Maximize

$$\left(\frac{\text{assoc}(A,A)}{\text{assoc}(A,V)} \right) + \left(\frac{\text{assoc}(B,B)}{\text{assoc}(B,V)} \right)$$

Normalized cuts

Optional

- Let \mathbf{y} be a vector whose elements are (ideally) 1 if the element is in A, and -b if it's in B.
 - b is theoretically defined for the derivation, but \mathbf{y} is going to be estimated.
- Write the matrix of the graph as W, and the matrix which has the row sums of W on its diagonal as D. Let $\mathbf{1}$ be a vector with all ones.
- With some algebra, the criterion becomes $\min_{\mathbf{y}} \left(\frac{\mathbf{y}^T (D - W) \mathbf{y}}{\mathbf{y}^T D \mathbf{y}} \right)$
- And we have a constraint $\mathbf{y}^T D \mathbf{1} = 0$
- This is hard to do, because \mathbf{y} 's values are quantized

Normalized cuts

Optional

- Instead, solve the generalized eigenvalue problem

$$\max_{\mathbf{y}} (\mathbf{y}^T (D - W) \mathbf{y}) \text{ subject to } (\mathbf{y}^T D \mathbf{y} = 1)$$

- which gives

$$(D - W) \mathbf{y} = \lambda D \mathbf{y}$$

- Now look for a quantization threshold that maximizes the criterion --- i.e all components of \mathbf{y} above that threshold go to one, all below go to -b

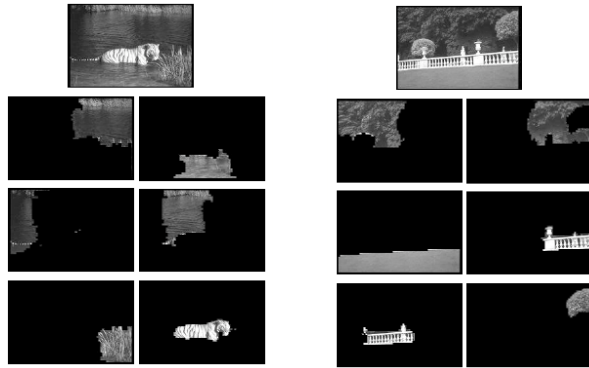


Figure from "Image and video segmentation: the normalised cut framework",
by Shi and Malik, copyright IEEE, 1998