

Clustering Methods


Prof. Asim Tewari
IIT Bombay

What are Clustering Methods?

- Clustering refers to a very broad set of techniques for finding subgroups, or clusters, in a data set.
- Clustering is concerned with grouping together objects that are similar to each other and dissimilar to the objects belonging to other clusters.
- Clustering is unsupervised learning that assigns labels to objects in unlabeled data.






Example of clustering

- How many clusters do you see?



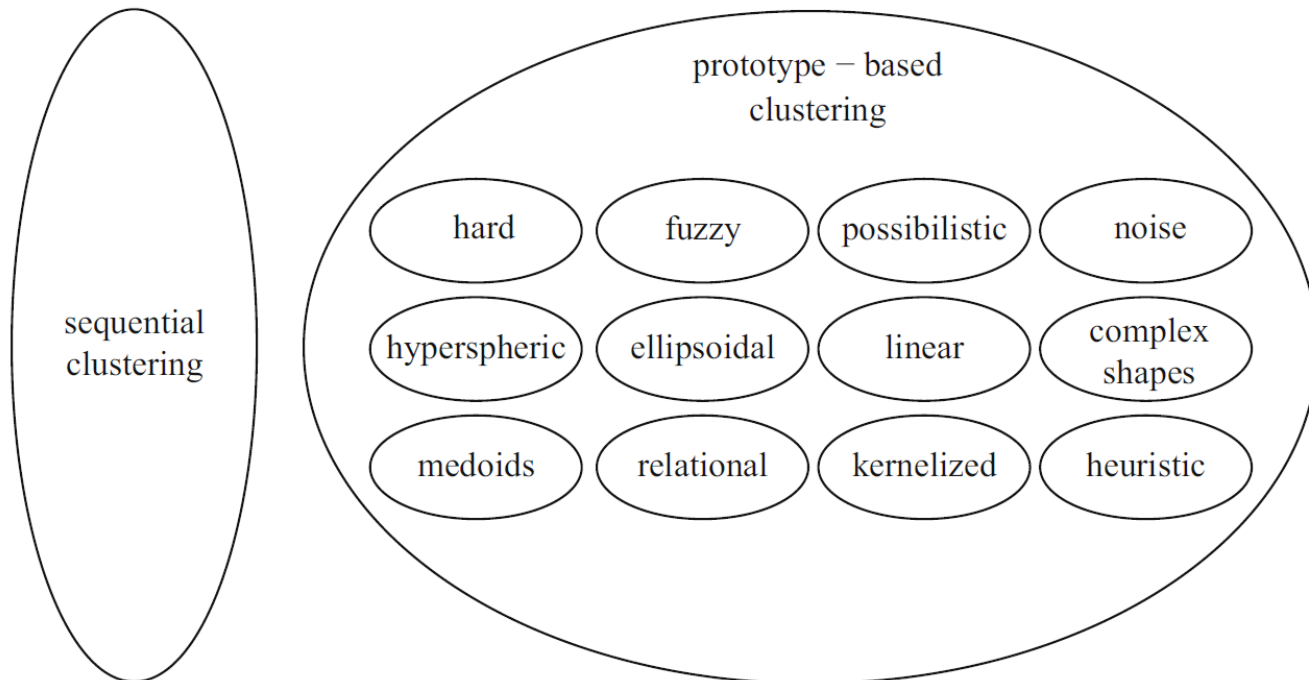
Example of clustering

- How many clusters do you see?



Types of Clustering Classification-1

- Sequential Clustering
- Prototype based clustering



Types of Clustering Classification-2

- K-means clustering
 - In this we seek to partition the observations into a pre-specified number of clusters
- Hierarchical clustering
 - In this we do not know in advance how many clusters we want; in fact, we end up with a tree-like visual representation of the observations, called a *dendrogram*

K-Means Clustering

Partitioning a data set into K distinct, non-overlapping clusters

- First specify the desired number of clusters K
- Then the K -means algorithm will assign each observation to exactly one of the K clusters
- The idea behind K -means clustering is that a good clustering is one for which the within-cluster variation is as small as possible

K-Means Clustering method

Let C_1, \dots, C_K denote sets containing the indices of the observations in each cluster. These sets satisfy two properties:

1. $C_1 \cup C_2 \cup \dots \cup C_K = \{1, \dots, n\}$. In other words, each observation belongs to at least one of the K clusters.
2. $C_{k_1} \cap C_{k_2} = \emptyset$ for all $k_1 \neq k_2$. In other words, the clusters are nonoverlapping: no observation belongs to more than one cluster.

K-Means Clustering method..

- Define a $W(C_k)$ as within-cluster variation measure.
- To obtain the k-means cluster solve the following problem

$$\text{minimize}_{C_1, \dots, C_K} \left\{ \sum_{k=1}^K W(C_k) \right\}$$

This formula says that we want to partition the observations into K clusters such that the total within-cluster variation, summed over all K clusters, is as small as possible

K-Means Clustering method..

- There are many possible ways to define the $W(C_k)$, within-cluster variation measure
- Most common is squared Euclidean distance

$$W(C_k) = \frac{1}{|C_k|} \sum_{i, i' \in C_k} \sum_{j=1}^p (x_{ij} - x_{i'j})^2$$




Where $|C_k|$ denotes the number of observations in the k th cluster

K-Means Clustering

K=2

K=3

K=4



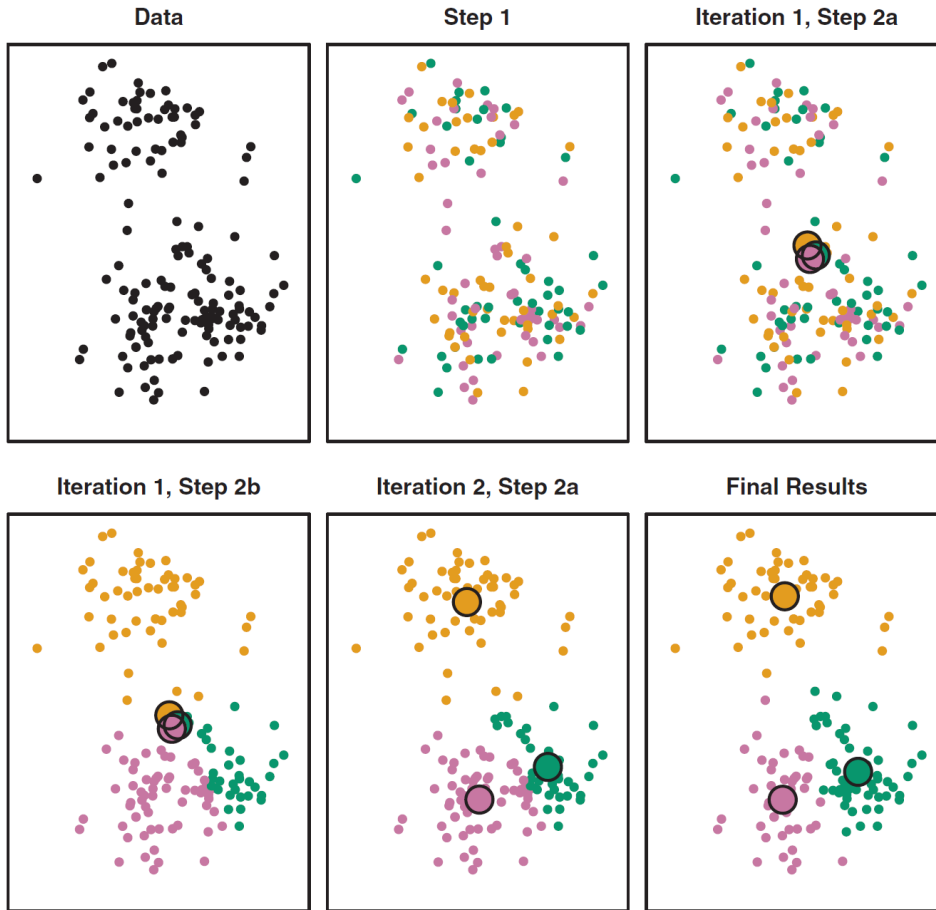
K-Means Clustering Algorithm

1. Randomly assign a number, from 1 to K , to each of the observations. These serve as initial cluster assignments for the observations.
2. Iterate until the cluster assignments stop changing:
 - a) For each of the K clusters, compute the cluster *centroid*. The k th cluster centroid is the vector of the p feature means for the observations in the k th cluster.
 - b) Assign each observation to the cluster whose centroid is closest (where *closest* is defined using Euclidean distance).

K-Means Clustering Algorithm..

- This algorithm is guaranteed to decrease the value of the objective at each step
- This means that as the algorithm is run, the clustering obtained will continually improve until the result no longer changes (the objective function will never increase)
- When the result no longer changes, a **local optimum** has been reached
- The resulting clustering will depend on the initial (random) cluster assignment

K-Means Clustering Algorithm..



Top left: the observations are shown.
Top center: in Step 1 of the algorithm, each observation is randomly assigned to a cluster.
Top right: in Step 2(a), the cluster centroids are computed. These are shown as large colored disks. Initially the centroids are almost completely overlapping because the initial cluster assignments were chosen at random.
Bottom left: in Step 2(b), each observation is assigned to the nearest centroid.
Bottom center: Step 2(a) is once again performed, leading to new cluster centroids.
Bottom right: the results obtained after ten iterations.

K-Means Clustering Algorithm..



- *K-means clustering with a different initial random assignment. Above each plot is the value of the objective function.*
- *Three different local optima were obtained*

Hierarchical Clustering

- In agglomerative Hierarchical Clustering we start with each object in a cluster of its own and then repeatedly merge the closest pair of clusters until we end up with just one cluster containing everything
- This results in a tree-based representation of the observations, called a *dendrogram*

Hierarchical Clustering


- *Dissimilarity measure between a pair of groups is defined by developing the notion of Linkage*
- *The four most common types of linkage are complete, average, single and centroid*

<i>Linkage</i>	<i>Description</i>
Complete	Maximal intercluster dissimilarity. Compute all pairwise dissimilarities between the observations in cluster A and the observations in cluster B, and record the <i>largest</i> of these dissimilarities.
Single	Minimal intercluster dissimilarity. Compute all pairwise dissimilarities between the observations in cluster A and the observations in cluster B, and record the <i>smallest</i> of these dissimilarities. Single linkage can result in extended, trailing clusters in which single observations are fused one-at-a-time.
Average	Mean intercluster dissimilarity. Compute all pairwise dissimilarities between the observations in cluster A and the observations in cluster B, and record the <i>average</i> of these dissimilarities.
Centroid	Dissimilarity between the centroid for cluster A (a mean vector of length p) and the centroid for cluster B. Centroid linkage can result in undesirable <i>inversions</i> .

Hierarchical Clustering Algorithm

1. Assign each object to its own single-object cluster. Calculate the dissimilarity between each pair of clusters.
2. Choose the closest pair of clusters and merge them into a single cluster (so reducing the total number of clusters by one).
3. Calculate the dissimilarity between the new cluster and each of the old clusters.
4. Repeat steps 2 and 3 until all the objects are in a single cluster.


Hierarchical Clustering Algorithm...




An illustration of the first few steps of the hierarchical clustering algorithm,

Effect of linkage on Hierarchical Clustering


Average Linkage



Complete Linkage



Single Linkage



- *Average, complete, and single linkage applied to an example data set.*
- *Average and complete linkage tend to yield more balanced clusters.*

Practical consideration in Clustering

- Should the observations or features first be standardized in some way? For instance, maybe the variables should be centered to have mean zero and scaled to have standard deviation one
- In the case of K -means clustering, how many clusters should we look for in the data?
- In the case of hierarchical clustering,
 - What dissimilarity measure should be used?
 - What type of linkage should be used?
 - Where should we cut the dendrogram in order to obtain clusters?
- How to validate the clusters obtained