Unsupervised Learning

Partitioning Methods

Road Map

1. Basic Concepts
2. K-Means
3. K-Medoids
4. CLARA & CLARANS

Cluster Analysis

- Unsupervised learning (i.e., Class label is unknown)
- Group data to form new categories (i.e., clusters), e.g., cluster houses to find distribution patterns
- Principle: Maximizing intra-class similarity & minimizing interclass similarity

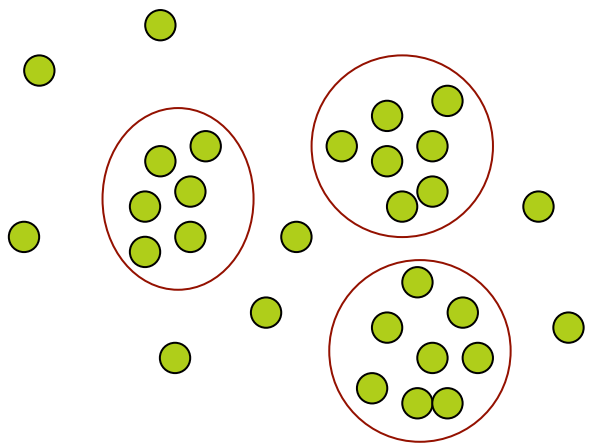

- Typical Applications
  - WWW, Social networks, Marketing, Biology, Library, etc.

Partitioning Methods

- Given
  - A data set of n objects
  - K the number of clusters to form
- Organize the objects into k partitions (k<=n) where each partition represents a cluster
- The clusters are formed to optimize an objective partitioning criterion
  - Objects within a cluster are similar
  - Objects of different clusters are dissimilar

Road Map

1. Basic Concepts
2. K-Means
3. K-Medoids
4. CLARA & CLARANS

K-Means

Choose 3 objects
(cluster centroids)

Assign each object
to the closest centroid
to form Clusters

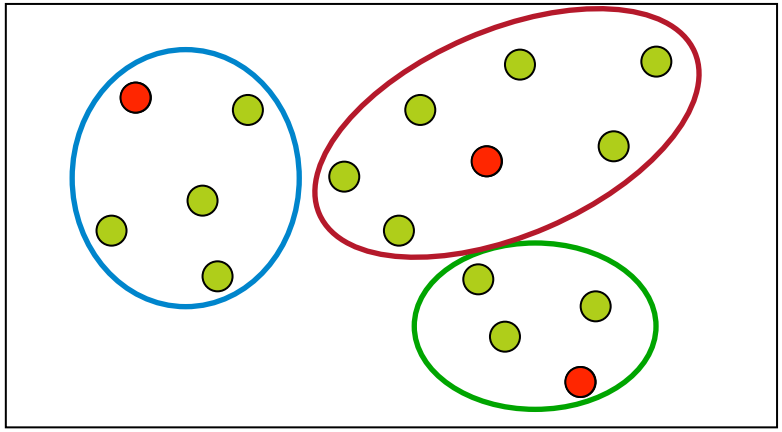

Goal:
create 3 clusters
(partitions)

Update cluster
centroids

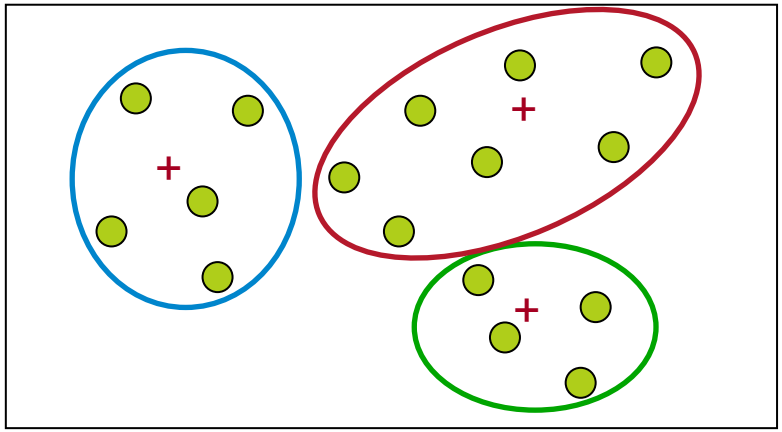

K-Means

Recompute Clusters

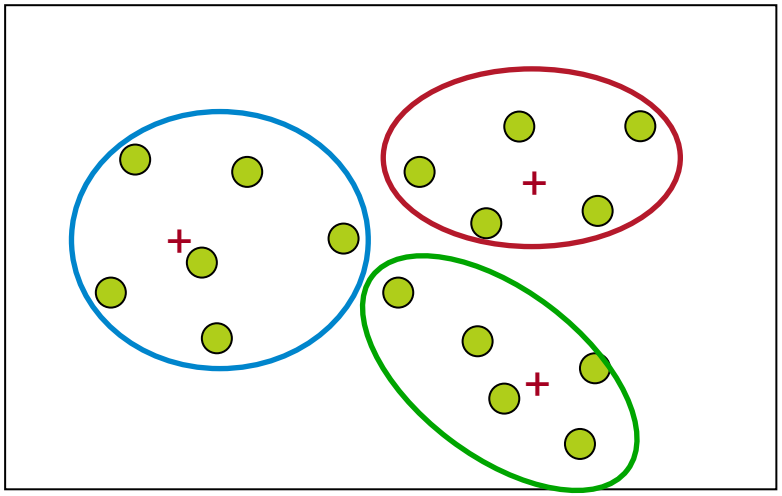

If Stable centroids, then stop

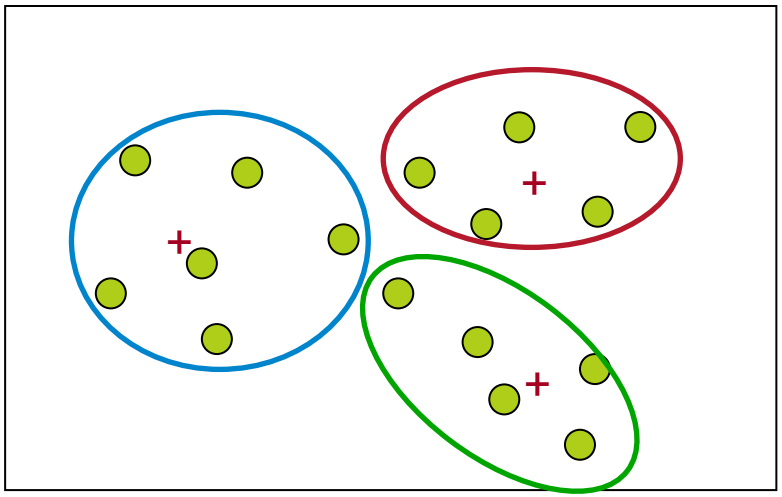

K-Means Algorithm

- Input
  - K: the number of clusters
  - D: a data set containing n objects
- Output: A set of k clusters
- Method:

(1) Arbitrary choose k objects from D as in initial cluster centers
(2) Repeat
(3) Reassign each object to the most similar cluster based on the mean value of the objects in the cluster
(4) Update the cluster means
(5) Until no change

K-Means Properties

- The algorithm attempts to determine k partitions that minimize the square-error function

$$E = \sum_{i-1}^{k} \sum_{p \in C_i} | p - m_i |^2$$

  - E: the sum of the squared error for all objects in the dataset
  - P: the data point in the space representing an object
  - m_i: is the mean of cluster C_i

- It works well when the clusters are compact clouds that are rather well separated from one another

K-Means: Advantages

- K-means is relatively scalable and efficient in processing large data sets
- The computational complexity of the algorithm is O(nkt)
  - n: the total number of objects
  - k: the number of clusters
  - t: the number of iterations
  - Normally: k<<n and t<<n

K-Means: Disadvantages

- Can be applied only when the mean of a cluster is defined
- Users need to specify k
- K-means is not suitable for discovering clusters with nonconvex shapes or clusters of very different size
- It is sensitive to noise and outlier data points (can influence the mean value)

K-Means demo

Demo

Variations of K-Means

- A few variants of the **k-means** which differ in
  - Selection of the initial k means
  - Dissimilarity calculations
  - Strategies to calculate cluster means
- How can we change K-Means to deal with categorical data?
- Handling categorical data: k-modes (Huang'98)
  - Replacing means of clusters with modes
  - Using new dissimilarity measures to deal with categorical objects
  - Using a frequency-based method to update modes of clusters
  - A mixture of categorical and numerical data

Road Map

1. Basic Concepts
2. K-Means
3. K-Medoids
4. CLARA & CLARANS

K-Medoids Method

- Minimize the sensitivity of k-means to outliers
- Pick actual objects to represent clusters instead of mean values
- Each remaining object is clustered with the representative object (**Medoid**) to which is the most similar
- The algorithm minimizes the sum of the dissimilarities between each object and its corresponding representative object

$$E = \sum_{i-1}^{k} \sum_{p \in C_i} | p - o_i |$$

  - E: the sum of absolute error for all objects in the data set
  - P: the data point in the space representing an object
  - Oi: is the representative object of cluster Ci

K-Medoids: The idea

- Initial representatives are chosen randomly
- The iterative process of replacing representative objects by no representative objects continues as long as the quality of the clustering is improved
- For each representative Object O
  - For each non-representative object R, swap O and R
- Choose the configuration with the lowest cost
- Cost function is the difference in absolute error-value if a current representative object is replaced by a non-representative object

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

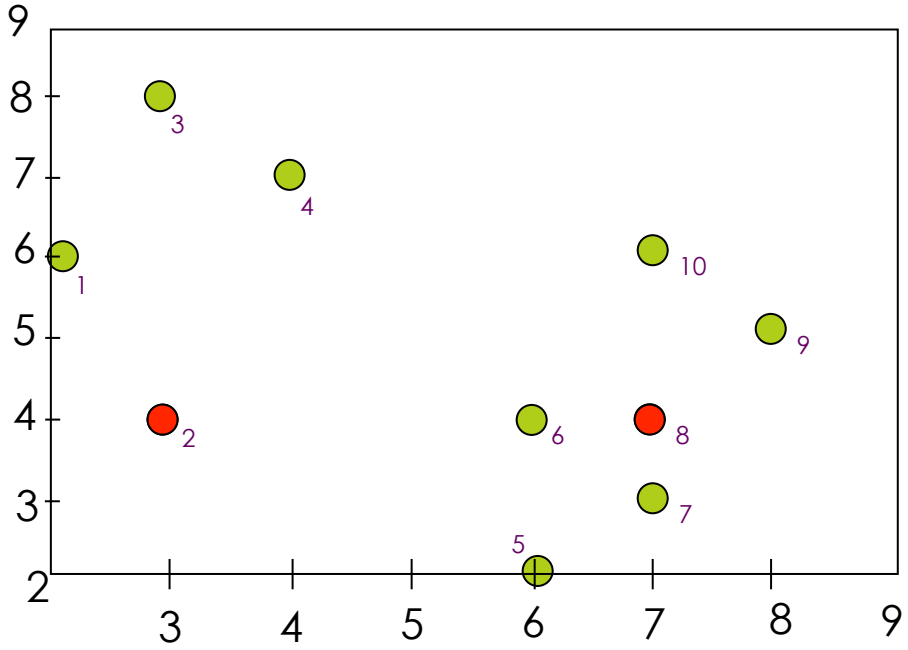

Goal: create two clusters

Choose randmly two medoids

$O_2 = (3,4)$
$O_8 = (7,4)$

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

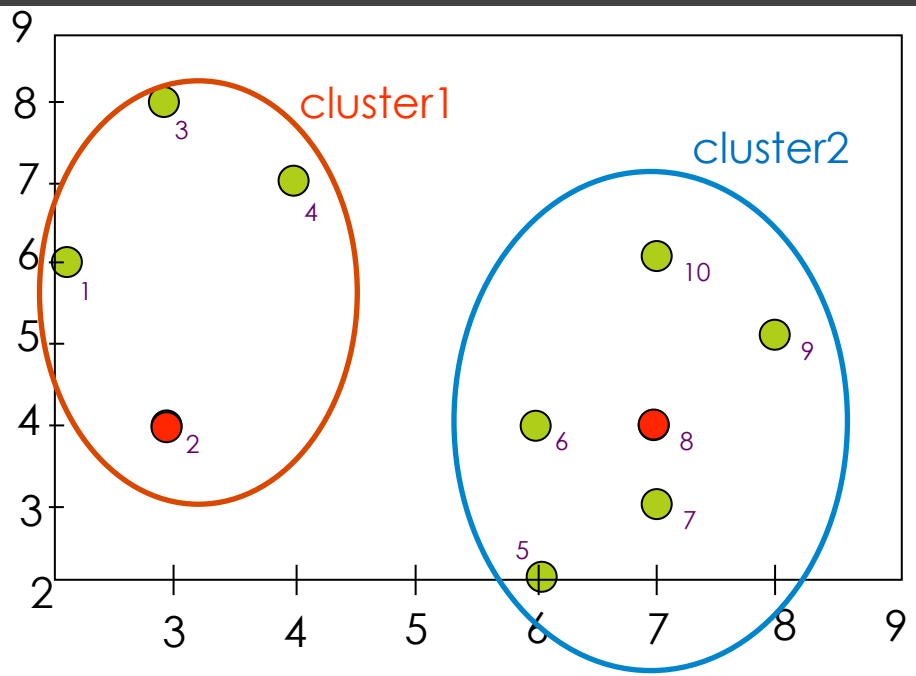

→Assign each object to the closest representative object

→Using L1 Metric (Manhattan), we form the following clusters

Cluster1 = $\{O_1, O_2, O_3, O_4\}$

Cluster2 = $\{O_5, O_6, O_7, O_8, O_9, O_{10}\}$

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

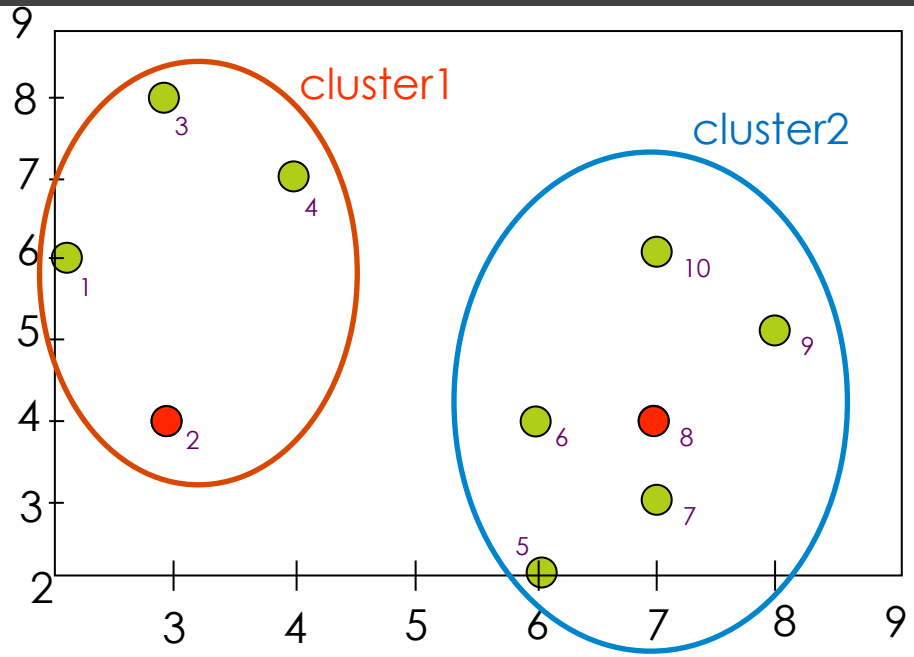

→Compute the absolute error criterion [for the set of Medoids (O2,O8)]

$$E = \sum_{i-1}^{k} \sum_{p \in C_i} |p - o_i| = |o_1 - o_2| + |o_3 - o_2| + |o_4 - o_2|$$

$$+ |o_5 - o_8| + |o_6 - o_8| + |o_7 - o_8| + |o_9 - o_8| + |o_{10} - o_8|$$

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

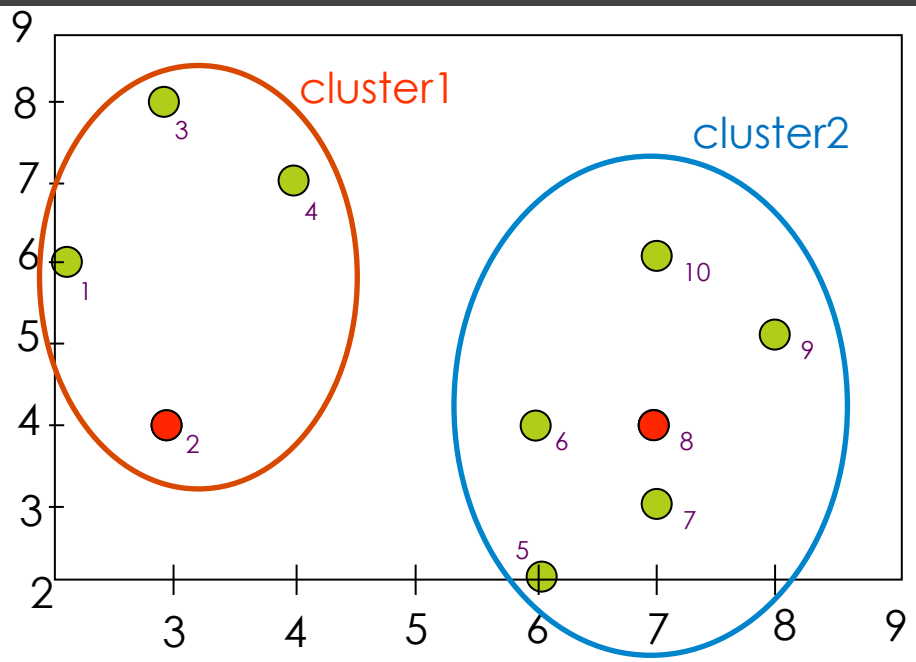

→The absolute error criterion [for the set of Medoids (O2,O8)]

$$E = (3+4+4)+(3+1+1+2+2) = 20$$

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

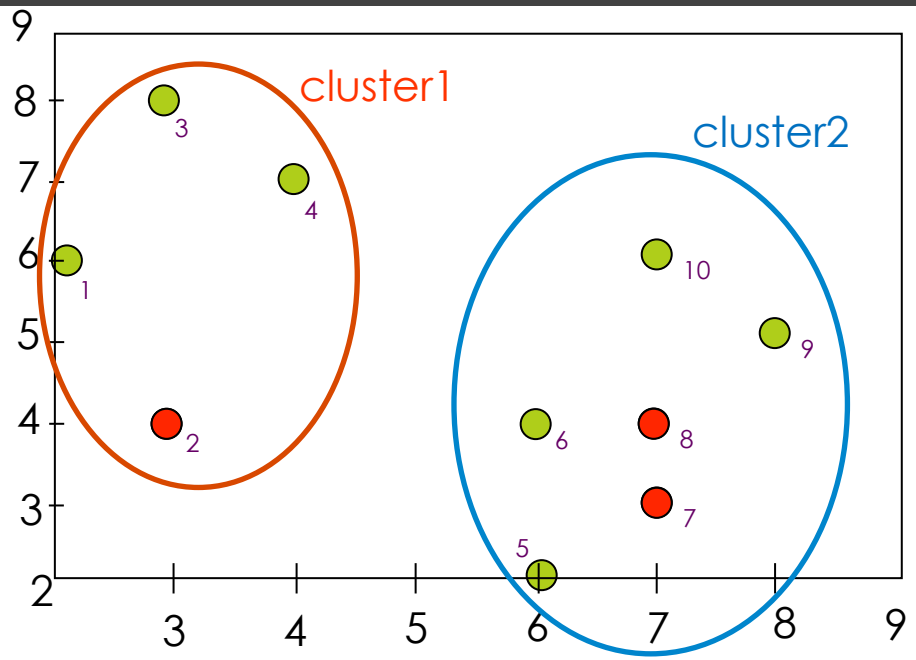

→Choose a random object O_7

→Swap O8 and O7

→Compute the absolute error criterion [for the set of Medoids (O2,O7)]

$$E = (3+4+4)+(2+2+1+3+3) = 22$$

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

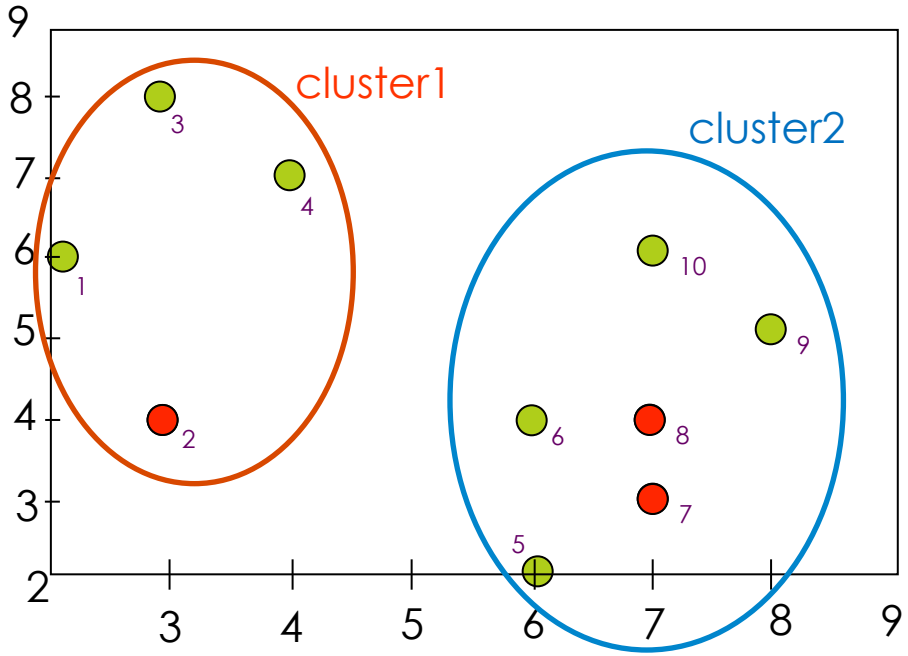

→Compute the cost function

Absolute error [for O_2,O_7] – Absolute error [O_2,O_8]

$$S = 22 - 20$$

S> 0 ⇒ it is a bad idea to replace O_8 by O_7

K-Medoids: Example

Data Objects

	A_1	A_2
O_1	2	6
O_2	3	4
O_3	3	8
O_4	4	7
O_5	6	2
O_6	6	4
O_7	7	3
O_8	7	4
O_9	8	5
O_{10}	7	6

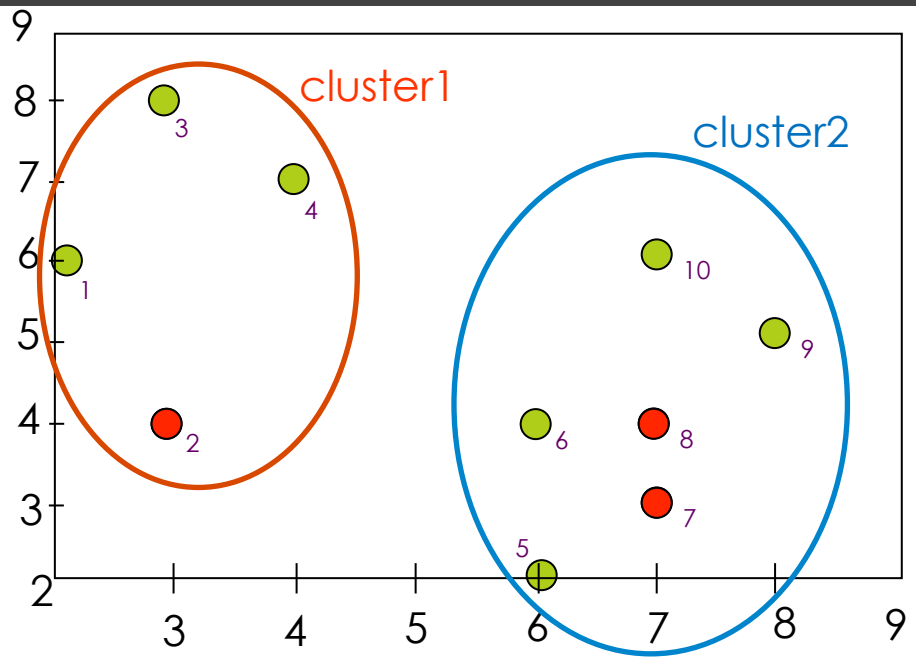

- In this example, changing the medoid of cluster 2 did not change the assignments of objects to clusters.
- What are the possible cases when we replace a medoid by another object?

K-Medoids

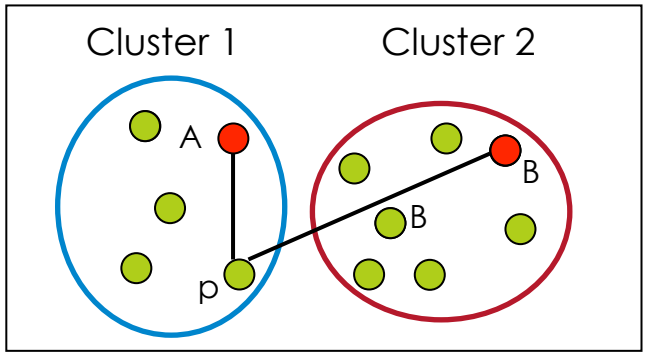

- Representative object
- Random Object

Currently P assigned to A

First case

The assignment of P to A does not change

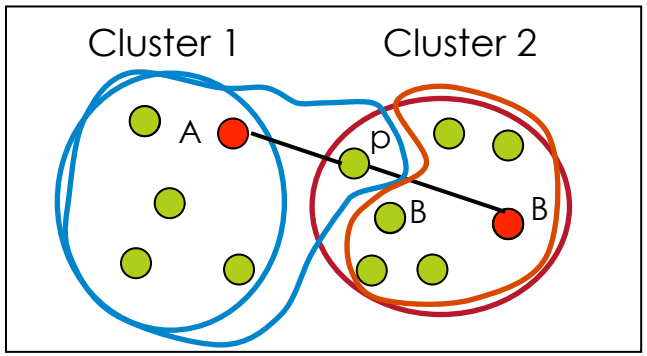

- Representative object
- Random Object

Currently P assigned to B

Second case

P is reassigned to A

K-Medoids

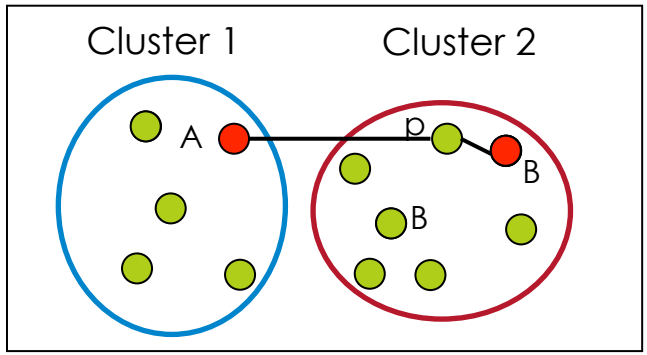

Representative object

Random Object

Currently P assigned to B

Third case

P is reassigned to the new B

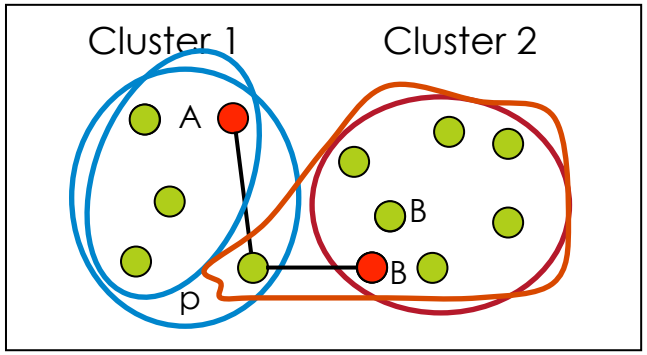

Representative object

Random Object

Currently P assigned to A

Fourth case

P is reassigned to B

PAM: Partitioning Around Medoids

- Input
  - K: the number of clusters
  - D: a data set containing n objects
- Output: A set of k clusters
- Method:

(1) Arbitrary choose k objects from D as representative objects (seeds)
(2) Repeat
(3) Assign each remaining object to the cluster with the nearest representative object
(4) For each representative object O_j
(5) Randomly select a non representative object O_{random}
(6) Compute the total cost S of swapping representative object Oj with O_{random}
(7) if S<0 then replace O_j with O_{random}
(8) Until no change

PAM Properties

- The complexity of each iteration is $O(k(n-k)^2)$
- For large values of n and k, such computation becomes very costly
- Advantages
  - K-Medoids method is more robust than k-Means in the presence of noise and outliers
- Disadvantages
  - K-Medoids is more costly than k-Means
  - Like k-means, k-medoids requires the user to specify k
  - It does not scale well for large data sets

Road Map

1. Basic Concepts
2. K-Means
3. K-Medoids
4. CLARA & CLARANS

CLARA

- CLARA (Clustering Large Applications) uses a sampling-based method to deal with large data sets

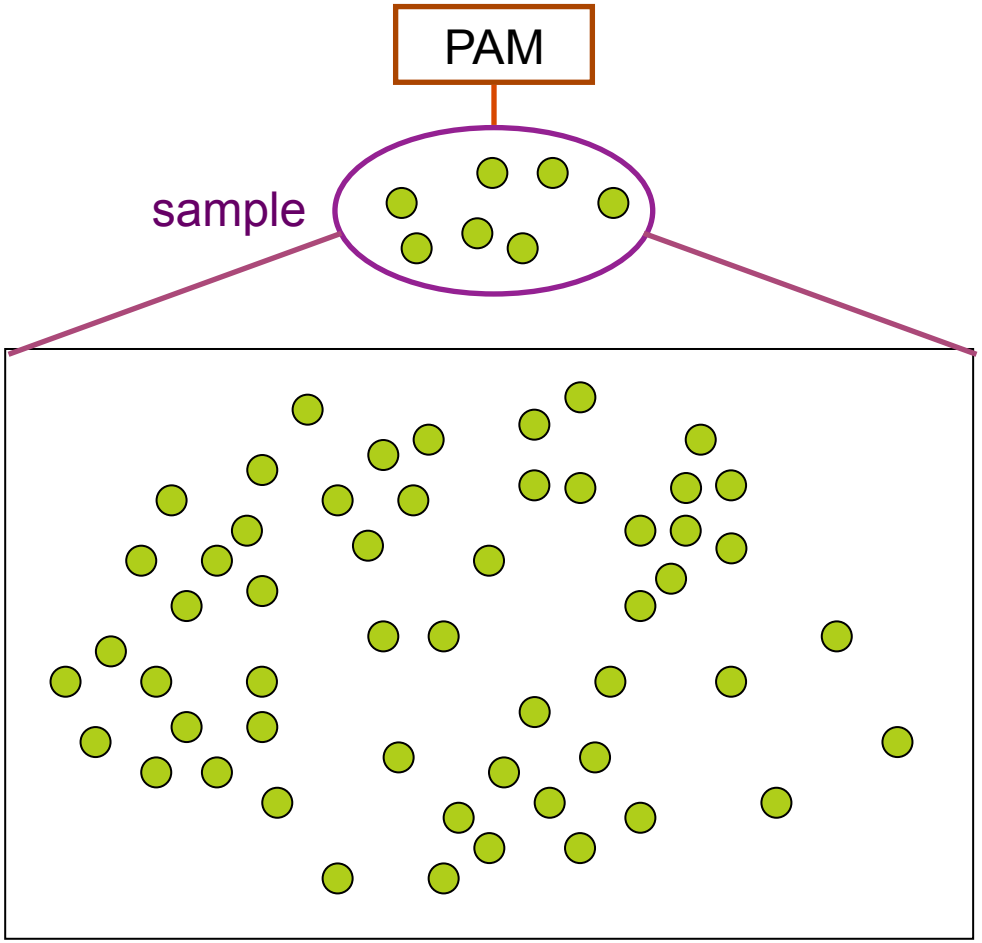

- A random sample should closely represent the original data

- The chosen medoids will likely be similar to what would have been chosen from the whole data set

CLARA

- Draw multiple samples of the data set
- Apply PAM to each sample
- Return the best clustering

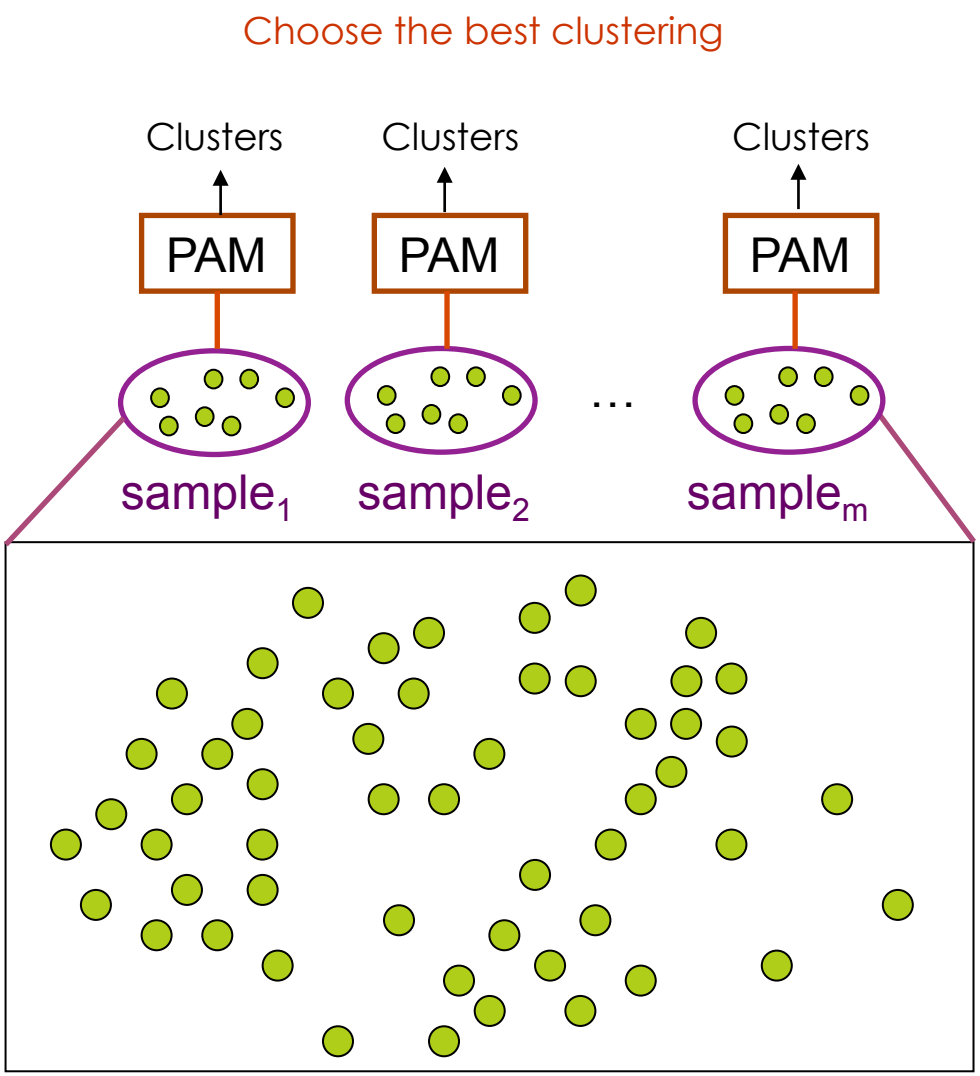

CLARA Properties

- Complexity of each Iteration is: $O(ks^2 + k(n-k))$
  - s: the size of the sample
  - k: number of clusters
  - n: number of objects

- PAM finds the best k medoids among a given data, and CLARA finds the best k medoids among the selected samples

- Problems
  - The best k medoids may not be selected during the sampling process, in this case, CLARA will never find the best clustering
  - If the sampling is biased we cannot have a good clustering
  - Trade off-of efficiency

CLARANS

- CLARANS (Clustering Large Applications based upon RANdomized Search) was proposed to improve the quality and the scalability of CLARA
- It combines sampling techniques with PAM
- It does not confine itself to any sample at a given time
- It draws a sample with some randomness in each step of the search

CLARANS: The idea

Clustering view

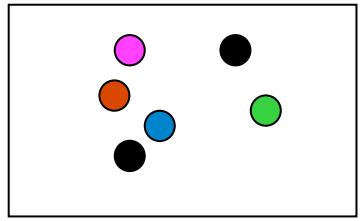

● medoids

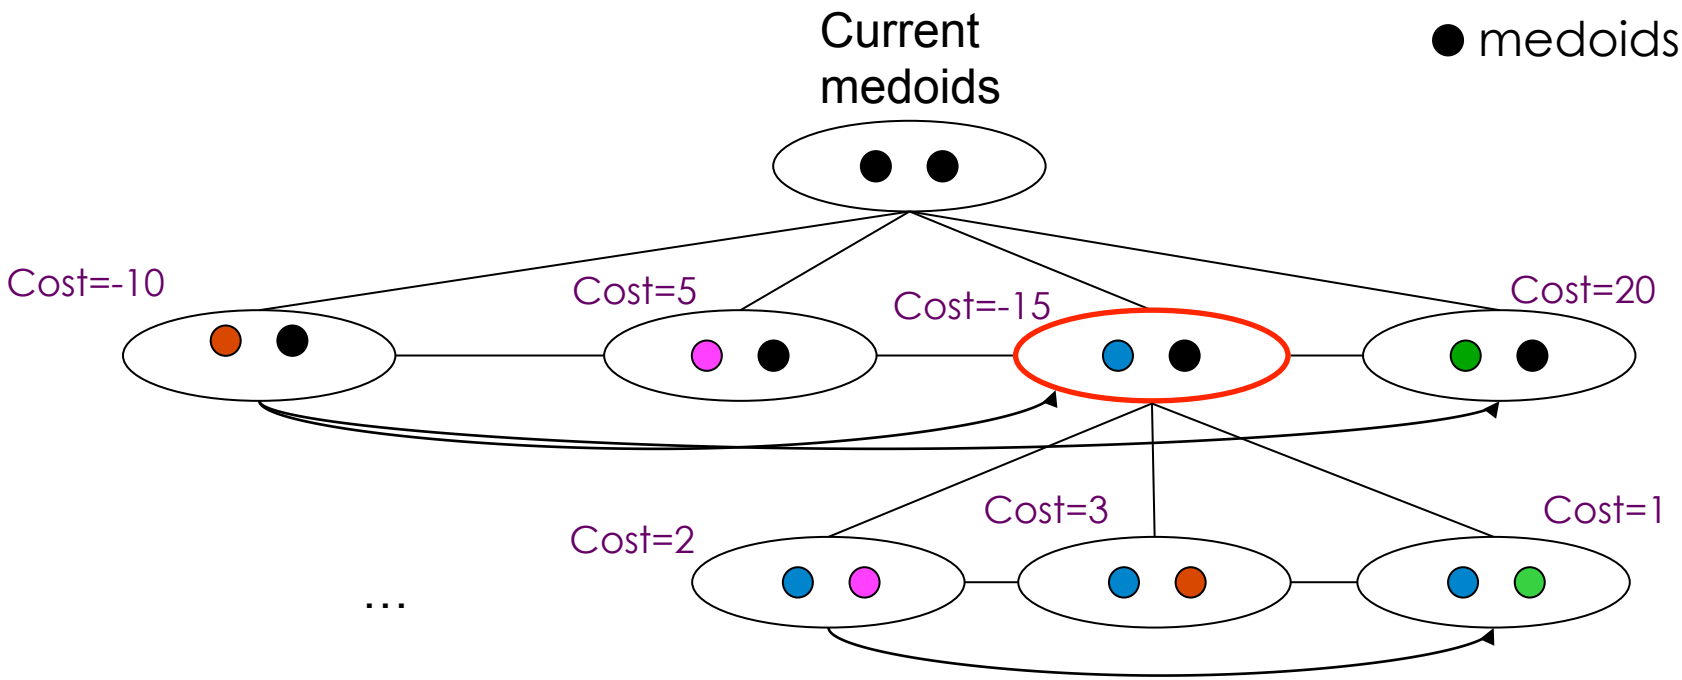

Keep the current medoids

CLARANS: The idea

CLARA

- Draws a sample of nodes at the beginning of the search
- Neighbors are from the chosen sample
- Restricts the search to a specific area of the original data

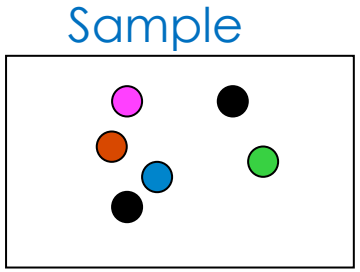

● medoids

First step of the search
Neighbors are from the chosen sample

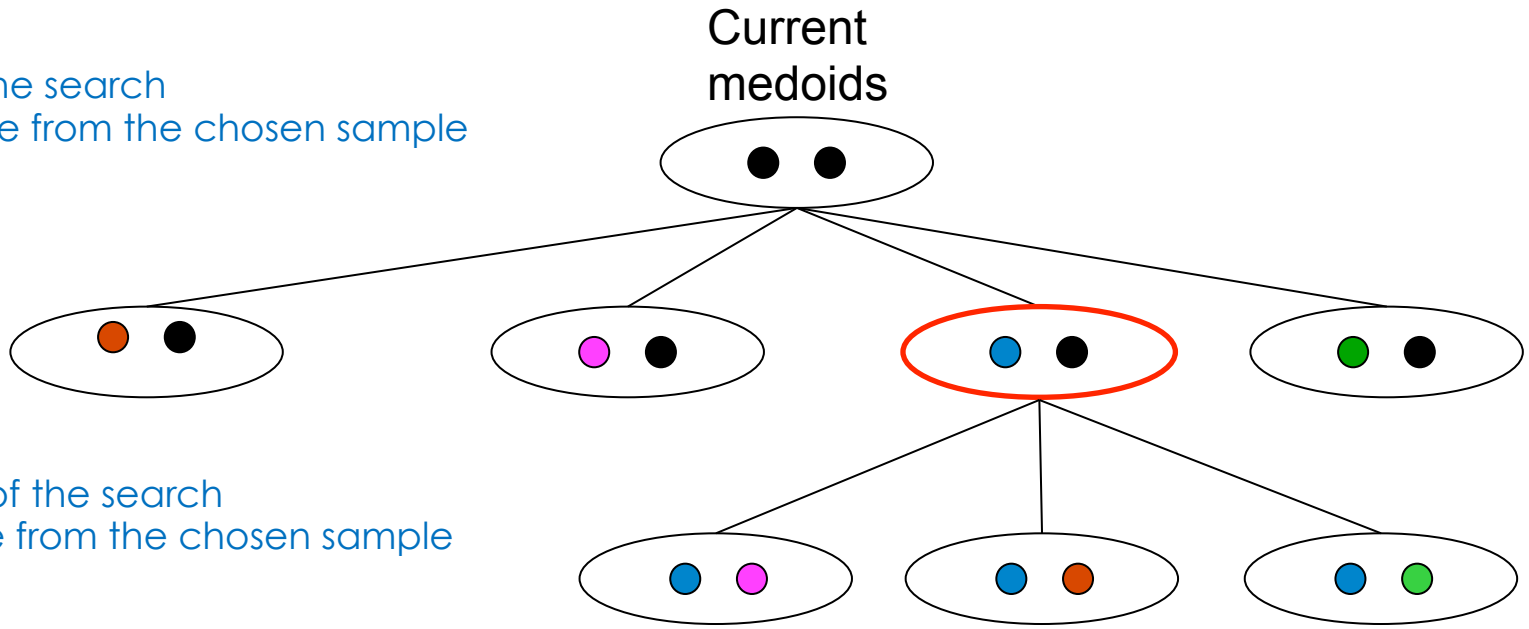

second step of the search
Neighbors are from the chosen sample

…

CLARANS: The idea

CLARANS

- Does not confine the search to a localized area
- Stops the search when a local minimum is found
- Finds several local optimums and output the clustering with the best local optimum

Original data

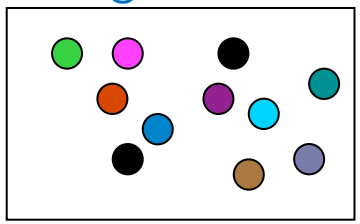

● medoids

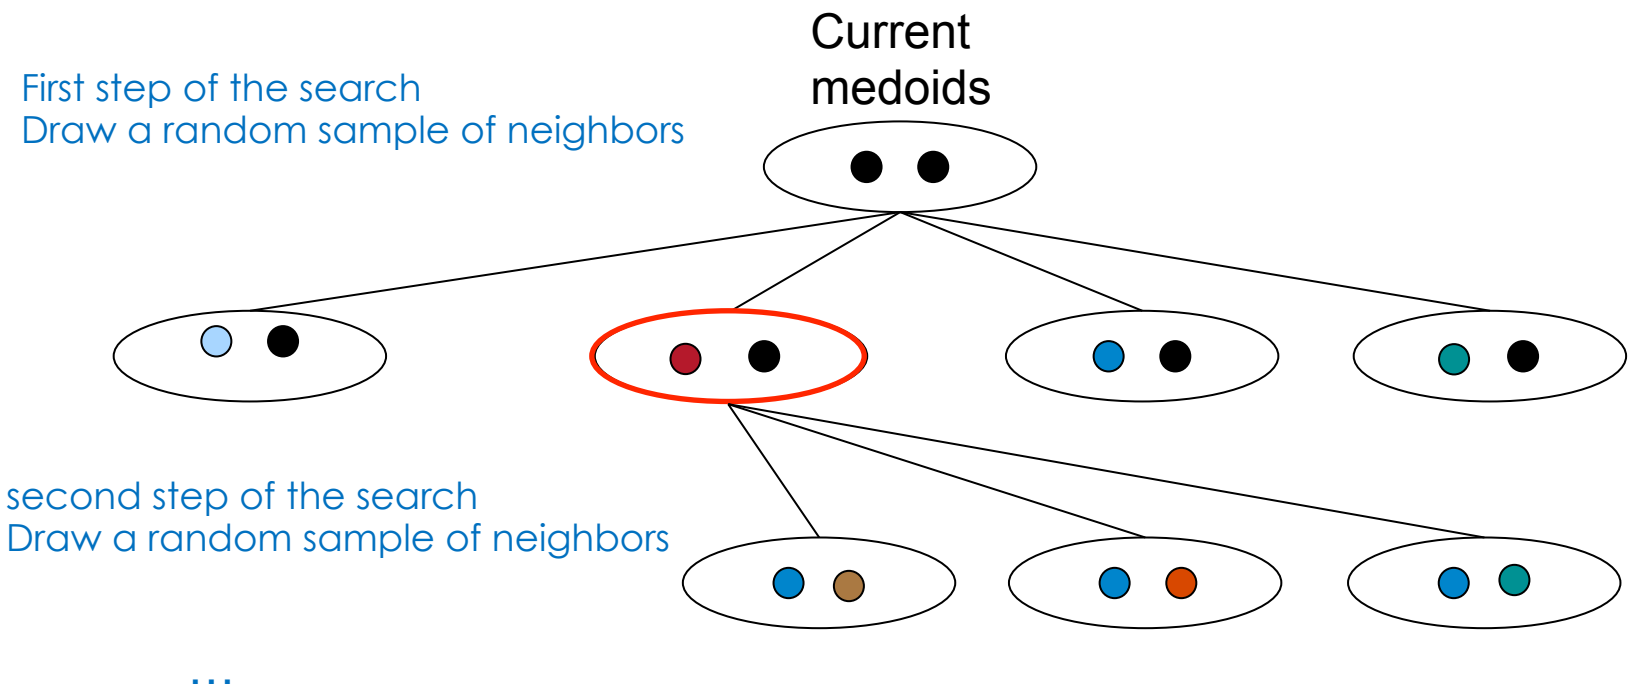

The number of neighbors sampled from the original data is specified by the user

CLARANS Properties

- Advantages
  - Experiments show that CLARANS is more effective than both PAM and CLARA
  - Handles outliers
- Disadvantages
  - The computational complexity of CLARANS is $O(n^2)$, where n is the number of objects
  - The clustering quality depends on the sampling method

Summary

- Partitioning methods find sphere-shaped clusters
- K- mean is efficient for large data sets but sensitive to outliers
- PAM used centers of the clusters instead of means
- CLARA and CLARANS are used for clustering large databases

Task1: Image Compression

- Assume you have a 538-pixel by 538-pixel image
- In a straightforward 24-bit color representation of this image, each pixel is represented as three 8-bit numbers (ranging from 0 to 255) that specify red, green and blue intensity values
- Our bird photo contains thousands of colors, but we would like to reduce that number to 16
- By making this reduction, it would be possible to represent the photo in a more efficient way by storing only the RGB values of the 16 colors present in the image
- How do solve this problem?

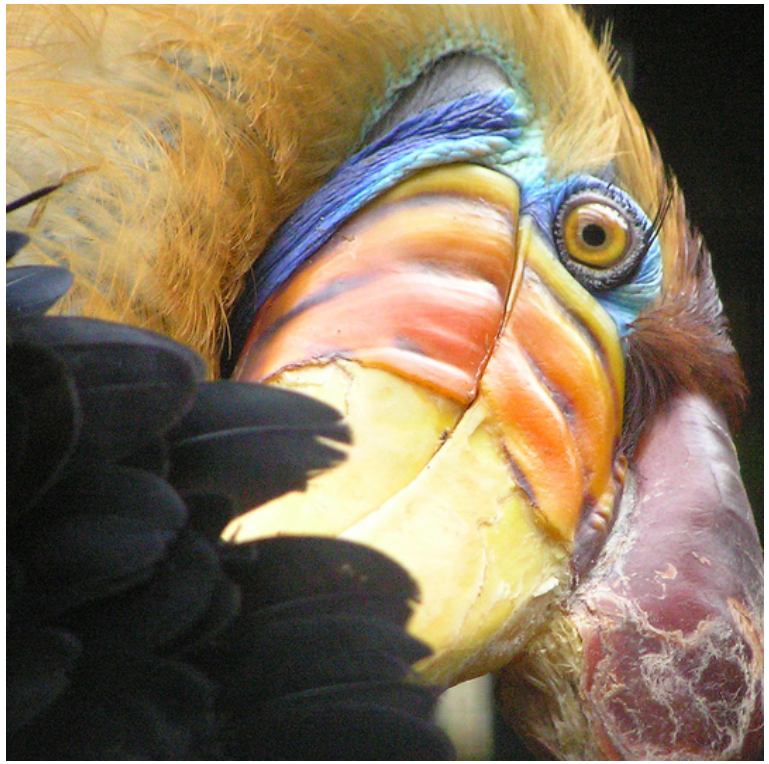

Task2: Documents Categorization

- We want cluster a set of documents based on their structural hierarchy. Documents that belong to the same topic should belong to the same cluster, documents that belong to the same sub-topic should belong to the same cluster, etc.

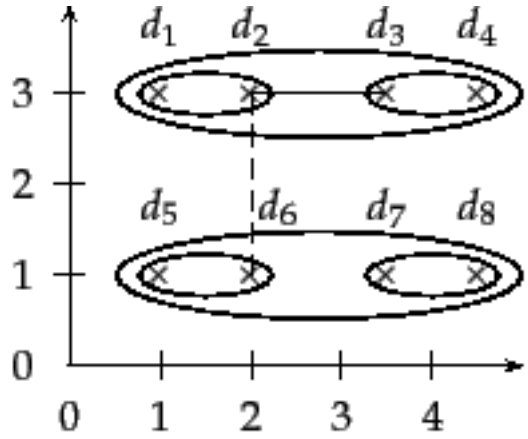

- Documents are described by the set of their contained terms
- Assume that similar documents belong to the same topic
- Propose a clustering Algorithm based on K-means that can achieve this task
- Imagine you have a knowledge base that gives information about topic hierarchies, where each topic is described by a set of terms. How would you cluster documents in this case?

Task3: Search Result Diversification

Query Google for Images of Barbara Liskov:

Results

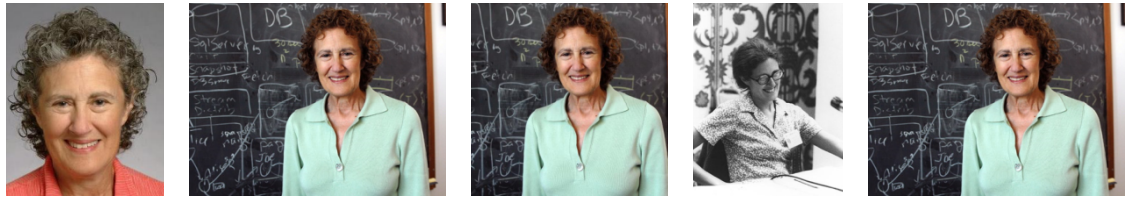

- We need to diversify search results
- How could you approach this problem using Clustering?

Task4: Fraudulent credit-card transactions

- We aim at identifying frauds in transactions.

- Assume for each transaction we have the following information

  - Costumer Id
  - Amount
  - List and types of bought items
  - Location

- How do you achieve this task?