

Clustering algorithms

Konstantinos Koutroumbas

Unit 3

- Fuzzy measures, Missing data
- Clustering: Sequential clustering

Proximity measures between vectors

Fuzzy measures – an alternative perspective

- Let $\mathbf{x} \in [0,1]^l$.
- In this context, x_k is **not** the **outcome** of a **measuring device**.
- Rather, it **indicates** the **degree** to which \mathbf{x} **possesses** the k -th characteristic.
- The closer the x_k to 1 (0), the **more likely** is that \mathbf{x} **possesses** (does not possess) the k -th characteristic.
- As x_k **approaches** 0.5, the **certainty** about the **possession or not** of the i -th feature from \mathbf{x} decreases.
- Let

$$\mathbf{x} = [x_1, x_2, \dots, x_k, \dots, x_l]^T \in [0,1]^l$$

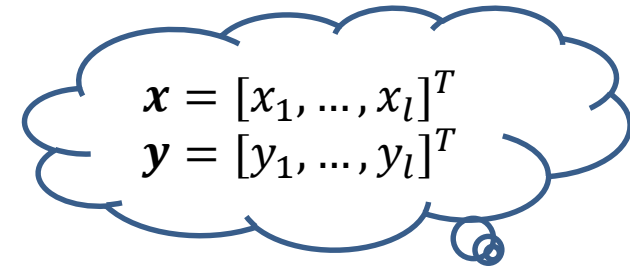
$$\mathbf{y} = [y_1, y_2, \dots, y_k, \dots, y_l]^T \in [0,1]^l$$

- A **measure of similarity between** x_k and y_k is the following

$$s(x_k, y_k) = \max(\min(1 - x_k, 1 - y_k), \min(x_k, y_k))$$

Then, as **measure of similarity between** \mathbf{x} and \mathbf{y} we can use the following

$$s^q(\mathbf{x}, \mathbf{y}) = \left(\sum_{k=1}^l s(x_k, y_k)^q \right)^{1/q}, q \in [1, +\infty)$$


$$\mathbf{x} = [x_1, \dots, x_l]^T$$
$$\mathbf{y} = [y_1, \dots, y_l]^T$$

Proximity measures between vectors

Fuzzy measures – an alternative perspective

Exercise 3: Let $l = 3$ and $q = 1$.

- (a) Consider the vectors $\mathbf{x}_1 = [1, 1, 1]^T$, $\mathbf{x}_2 = [0, 0, 1]^T$, $\mathbf{x}_3 = [\frac{1}{2}, \frac{1}{3}, \frac{1}{4}]^T$, $\mathbf{x}_4 = [\frac{1}{2}, \frac{1}{2}, \frac{1}{2}]^T$. Determine the similarities $s^1(\mathbf{x}_i, \mathbf{x}_i)$, $i = 1, 2, 3, 4$.
- (b) Consider the vectors $\mathbf{y}_1 = [\frac{3}{4}, \frac{3}{4}, \frac{3}{4}]^T$, $\mathbf{y}_2 = [1, 1, 1]^T$, $\mathbf{y}_3 = [\frac{1}{4}, \frac{1}{4}, \frac{1}{4}]^T$, $\mathbf{y}_4 = [\frac{1}{2}, \frac{1}{2}, \frac{1}{2}]^T$. Determine the similarities $s^1(\mathbf{y}_i, \mathbf{y}_j)$, $i, j = 1, 2, 3, 4$, $i \neq j$.
- (c) Draw your conclusions.

Proximity measures between vectors

Dynamic similarity measures

- These are useful for cases where **the two vectors** to be compared have **different lengths**.
- Such a situation may arise e.g., when **comparing two strings** stemming **from two different texts**.
- A simple example: The **Edit distance**.

Proximity measures between vectors – Missing data

Missing data

- For **some vectors** of the data set X , **some features values** are **unknown**.
- This issue arises **very often** in **practice**.
- It may be caused by a measurement device failure, inability to take measure due to specific physical conditions etc.
- Ways to deal with this situation:
 - ✓ **Discard** all **vectors** with **missing values** (not recommended for small data sets).
 - ✓ **Find** the mean value m_k of the **available k -th feature values** over that data set and **substitute** the **missing k -th feature values** with m_k .

Proximity measures between vectors – Missing data

Missing data

- Ways to deal with this situation:

✓ Define $b_k = 0$, if **both** the k -th features x_k, y_k are **available** and **1 otherwise**. Then

$$\wp(x, y) = \frac{l}{l - \sum_{k=1}^l b_k} \sum_{\text{all } k: b_k=0} \phi(x_k, y_k)$$

where $\phi(x_k, y_k)$ denotes the **proximity measure** between two scalars x_k, y_k .

NOTE: The **proximity** is **based only on the features** for which both x_k, y_k are **available**.

- ✓ For the k -th feature, $k = 1, 2, \dots, l$, **find** the average proximity $\phi_{avg}(k)$ among all **available values** along the feature vectors in X . Then

$$\wp(x, y) = \sum_{k=1}^l \psi(x_k, y_k),$$

where $\psi(x_k, y_k) = \begin{cases} \phi(x_k, y_k), & \text{if both } x_k, y_k \text{ are available} \\ \phi_{avg}(k), & \text{otherwise} \end{cases}$

Proximity measures between vectors – Missing data

Missing data


Exercise 4: Consider the data set $X = \{\mathbf{x}_1, \mathbf{x}_2, \mathbf{x}_3, \mathbf{x}_4, \mathbf{x}_5\}$, with $\mathbf{x}_1 = [0,0]^T$, $\mathbf{x}_2 = [1,*]^T$, $\mathbf{x}_3 = [0,*]^T$, $\mathbf{x}_4 = [2,2]^T$, $\mathbf{x}_5 = [3,1]^T$ (“*” stands for missing values).

- (a) Compute the l_1 distances between all pairs of vectors, using all the four techniques for dealing with missing data.
- (b) In which of these techniques, the computed distances are dependent on the specific data set?

Clustering: A historic example(*)

Dr John Snow plotted the **location** of **cholera deaths** on a **map** during an outbreak at London in the 1850s

The **locations** were **clustered** around certain intersections where there were **polluted wells!!!**



Questions:

- Which are the **entities** and how they are **represented**?
- Which **dissimilarity measure** could be used?
- What is the **form** of the **resulted clusters**?
- What kind of clusters should be able to reveal the adopted **clustering criterion**?

Clustering algorithms

Number of possible clusterings

Let $X = \{x_1, x_2, \dots, x_N\}$ be a set of data points.

Question: In how many ways the N points of X can be assigned into m groups?

Answer:
$$S(N, m) = \frac{1}{m!} \sum_{i=0}^m (-1)^{m-i} \binom{m}{i} i^N$$

Examples:


- $S(15, 3) = 2,375,101$
- $S(20, 4) = 45,232,115,901$
- $S(25, 8) = 690,223,721,118,368,580$
- $S(100, 5) \approx 10^{68}!!$

NOTE: The above calculations are for fixed m . If this varies, then we have to enumerate **all clusterings**, for **all possible** values of m !!


⇒

Evaluating all possible clusterings is **impractical** even for **moderate values** of N .


Clustering algorithms



Data set



A “sensible” clustering



A “less sensible” clustering

- **Clustering algorithms** may be **viewed** as **schemes** that provide us with *sensible clusterings* by considering only *a small fraction* of all possible partitions of X .
- This *fraction* **depends on** the adopted **criteria**.
- Thus a **clustering algorithm** is a **learning procedure** that tries to **identify clusters** formed by the data vectors, **in accordance to the adopted criteria**.

Clustering algorithms

Major categories of clustering algorithms

A **vast amount** of **algorithms** **exists** based on **very diverse criteria**
⇒ **Strict categorization** is extremely **difficult** (rather **impossible**).

A rough categorization:

- **Sequential**: A **single clustering** is produced. **One** or **few sequential passes** on the data.
- **Hierarchical**: A **sequence** of (nested) **clusterings** is produced.
 - Agglomerative**
 - Matrix theory
 - Graph theory
 - Divisive**
 - Combinations** of the above (e.g., the Chameleon algorithm.)

Clustering algorithms

Major categories of clustering algorithms

A rough categorization:

Cost function optimization.

- For most of the cases a *single clustering* is obtained.
- They can be further **categorized** through the notion of “**belongness**”.
Hard clustering (each **point belongs** exclusively to **a** single **cluster**):

- Basic hard clustering algorithms (e.g., *k*-means)
- *k*-medoids algorithms
- Mixture decomposition
- Branch and bound
- Simulated annealing
- Deterministic annealing
- Boundary detection
- Mode seeking
- Genetic clustering algorithms

Probabilistic clustering (a hard clustering case where probabilistic framework is utilized)

Fuzzy clustering (each **point belongs** to **more** than one **clusters** simultaneously).

Possibilistic clustering (it is based on the notion of the “*degree of compatibility*” of a point with a cluster).

Clustering algorithms

Major categories of clustering algorithms

A rough categorization:

Other.

- Algorithms based on **graph theory** (e.g., Spectral clustering, Minimum Spanning Tree, regions of influence, directed trees).
- **Density-based** algorithms.
- **Competitive learning** algorithms (basic competitive learning scheme, Kohonen self organizing maps).
- **Subspace clustering** algorithms.
- **Ensemble of clusterings**
- **Kernel-based** methods.

Sequential clustering algorithms

The common traits shared by the sequential clustering algorithms are:

- One or very **few passes** on the data are **required**.
- The number of clusters m is **not known a-priori**, except (possibly) an **upper bound**, q .
- The **clusters** are **defined** with the **aid** of
 - ✓ An appropriately defined distance $d(x, C)$ of a point from a cluster.
 - ✓ A threshold θ associated with the distance.

Sequential clustering algorithms

Basic Sequential Clustering Algorithm Scheme (BSAS)

- $m = 1 \setminus \{\text{number of clusters}\}$

- $C_m = \{x_1\}$

- **For** $i = 2$ to N

- **Find** C_k : $d(x_i, C_k) = \min_{1 \leq j \leq m} d(x_i, C_j)$

- **If** $(d(x_i, C_k) > \theta)$ AND $(m < q)$ then

- $m = m + 1$

- $C_m = \{x_i\}$


- **Else**

- $C_k = C_k \cup \{x_i\}$

- Where necessary, update representatives (*)

- **End** {if}

- **End** {for}



(*) When the mean vector m_C is used as representative of the cluster C with n_C elements, the updating in the light of a new vector x becomes

$$m_C^{new} = (n_C m_C^{old} + x) / (n_C + 1)$$

Sequential clustering algorithms

Basic Sequential Clustering Algorithm Scheme (BSAS)

Remarks:


- The **order of presentation of the data** in the algorithm plays important role in the clustering results. **Different order of presentation may lead to totally different clustering results**, in terms of the **number of clusters** as well as the **clusters themselves**.
- The **clustering results** depend on the choice of the value of θ .
- In BSAS the **decision** for a vector \mathbf{x} is **reached prior** to the **final cluster formation**.
- **BSAS** perform a **single pass** on the data. Its complexity is $O(N)$ (when point representatives are used).
- If clusters are represented by **point representatives**, **compact clusters** are favored.

Sequential clustering algorithms





Example in MATLAB 1:

Consider the data vectors depicted in the figure below and perform a “visual clustering” on it.





1. Apply the BSAS algorithm on X , presenting its elements in the order $x_8, x_6, x_{11}, x_1, x_5, x_2, x_3, x_4, x_7, x_{10}, x_9, x_{12}, x_{13}, x_{14}, x_{15}$, for $\theta = 2.5$ and $q = 15$.
2. Repeat step 1, now with the order of presentation to the algorithm as $x_7, x_3, x_1, x_5, x_9, x_6, x_8, x_4, x_2, x_{10}, x_{15}, x_{13}, x_{14}, x_{11}, x_{12}$.
3. Repeat step 1, now with $\theta = 1.4$.
4. Repeat step 1, now with $q = 2$.







Sequential clustering algorithms







Sequential clustering algorithms



Sequential clustering algorithms



Sequential clustering algorithms



Sequential clustering algorithms

Basic Sequential Clustering Algorithm Scheme (BSAS)

Estimating the number of clusters in the data set:

Let $BSAS(\theta)$ denote the $BSAS$ algorithm when the dissimilarity threshold is θ .



- For $\theta = a$ to b step c

- Run s times $BSAS(\theta)$, each time presenting the data in a different order.

- Estimate the number of clusters m_θ , as the most frequent number resulting from the s runs of $BSAS(\theta)$.

- Next θ

- Plot m_θ versus θ and identify the number of clusters m as the one corresponding to the widest flat region in the above graph.



- Consider as final clustering, the clustering that results for the θ in the middle of the widest flat region.

Sequential clustering algorithms

MBSAS, a Modification of BSAS

- In **BSAS** a **decision** for a data vector x is **reached prior** to the **final cluster formation**, which is determined after all vectors have been presented to the algorithm.
- MBSAS deals with this issue, at the cost of processing the data twice.
- **MBSAS** consists of:
 - A **cluster determination phase** (first pass on the data), which is the **same as BSAS** with the **exception** that **no vector is assigned to an already formed cluster**. At the end of this phase, **each cluster consists** of a **single element**.
 - A **pattern classification phase** (second pass on the data), where **each** one of the **unassigned vectors** is **assigned** to its **closest cluster**.

Exercise: Write the pseudocode for MBSAS (in the spirit of the BSAS pseudocode).

Remarks:

- In MBSAS, a decision for a vector x during the pattern classification phase is reached taking into account all clusters.
- MBSAS is **sensitive** to the **order of presentation** of the vectors.
- MBSAS requires **two passes** on the **data**. Its complexity is $O(N)$.

Sequential clustering algorithms

Refinement stages

The problem of **closeness of clusters**: “In all the above algorithms it may happen that two formed clusters lie very close to each other”.

(they may be **parts** of the **same physical cluster**)

A simple merging procedure

(A) **Find** C_i, C_j ($i < j$) such that $d(C_i, C_j) = \min_{k,r=1,\dots,m,k \neq r} d(C_k, C_r)$

If $d(C_i, C_j) \leq M_1$ then $\{ M_1 \text{ is a user-defined threshold} \}$

- **Merge** C_i, C_j to C_i and eliminate C_j .
- If necessary, update the cluster representative of C_i .
- Rename the clusters C_{j+1}, \dots, C_m to C_j, \dots, C_{m-1} , respectively.
- $m = m - 1$
- Go to (A)

Else

- Stop

End {if}

Sequential clustering algorithms

Refinement stages

The problem of **sensitivity to the order of data presentation**:

“A vector \underline{x} may have been assigned to a cluster C_i at the current stage but another cluster C_j may be formed at a later stage that lies closer to \underline{x} ”

A simple reassignment procedure

- **For** $i = 1$ to N
 - **Find** C_j such that $d(\underline{x}_i, C_j) = \min_{k=1, \dots, m} d(\underline{x}_i, C_k)$
 - **Set** $b(i) = j$ \{ $b(i)$ is the index of the cluster that lies closest to \underline{x}_i \}
- **End** {for}

- **For** $j = 1$ to m
 - **Set** $C_j = \{\underline{x}_i \in X: b(i) = j\}$
 - If necessary, update representatives
- **End** {for}

Sequential clustering algorithms

Example in MATLAB 2:

Generate and plot a data set X_1 , that consists of $N = 400$ 2-dim. data vectors. These vectors form **four groups**, each one of which contains vectors that stem from Gaussian distributions with **means** $\mathbf{m}_1 = [0, 0]^T$, $\mathbf{m}_2 = [4, 0]^T$, $\mathbf{m}_3 = [0, 4]^T$, $\mathbf{m}_4 = [5, 4]^T$, respectively, and respective **covariance matrices** $S_1 = I$, $S_2 = \begin{bmatrix} 1 & 0.2 \\ 0.2 & 1.5 \end{bmatrix}$, $S_3 = \begin{bmatrix} 1 & 0.4 \\ 0.4 & 1.1 \end{bmatrix}$, $S_4 = \begin{bmatrix} 0.3 & 0.2 \\ 0.2 & 0.5 \end{bmatrix}$. Then do the following:

1. Determine the number of clusters formed in X_1 by doing the following:
 - a. Determine the maximum, d_{max} , and the minimum, d_{min} , distances between any two points in the data set.
 - b. Determine the values of Θ for which the BSAS will run. These may be defined as $\Theta_{min}, \Theta_{min} + s, \Theta_{min} + 2s, \dots, \Theta_{max}$, where $\Theta_{min} = 0.25 \frac{d_{min} + d_{max}}{2}$, $\Theta_{max} = 1.75 \frac{d_{min} + d_{max}}{2}$ and $s = \frac{\Theta_{min} + \Theta_{max}}{n_\Theta}$, n_Θ is the number of successive values of Θ that will be considered. Use $n_\Theta = 50$.





Sequential clustering algorithms

Example in MATLAB 2 (cont.):

- c. For each of the previously defined values of Θ , run the BSAS algorithm $n_{times} = 10$, so that the data vectors are presented with different ordering to BSAS in each run. From the n_{times} estimates of the number of clusters, select the most frequently met value, m_{θ} , as the most accurate. Let \mathbf{m}_{tot} be the n_{θ} -dimensional vector, which contains the m_{θ} values.
 - d. Plot m_{θ} versus Θ . Determine the widest flat region, r , of Θ 's (excluding the one that corresponds to the single-cluster case) and let n_r be the number of Θ 's in $\{\Theta_{min}, \Theta_{min} + s, \dots, \Theta_{max}\}$ that also lie in r . If n_r is "significant" (e.g., greater than 10% of n_{θ}), the corresponding number of clusters, m_{best} , is selected as the best estimate and the mean of the values of Θ in r is chosen as the corresponding best value for Θ (Θ_{best}). Otherwise, the single-cluster clustering is adopted.
2. Run the BSAS algorithm for $\Theta = \Theta_{best}$ and plot the data set using different colors and symbols for points from different clusters.
 3. Apply the reassignment procedure on the clustering results obtained in the previous step and plot the new clustering.

Sequential clustering algorithms

Example in MATLAB 2 (cont.):



Sequential clustering algorithms

A two-threshold sequential scheme (TTSAS)

- The formation of the clusters, as well as the assignment of vectors to clusters, is carried out concurrently (like BSAS and unlike MBSAS)
- **Two thresholds** θ_1 and θ_2 ($\theta_1 < \theta_2$) are **employed**.
- The **general idea** is the following:

If the distance $d(x, C)$ of x from its closest cluster, C , is **greater** than θ_2 then:

–A **new cluster** represented by x is created.

Else if $d(x, C) < \theta_1$ then

– x is **assigned** to C .

Else

–The **decision** is **postponed** to a **later stage**.

End {if}

- The unassigned vectors are presented iteratively to the algorithm until all of them are classified.

Remarks:

- In practice, a few passes (≥ 2) of the data set are required.
- TTSAS is less sensitive to the order of data presentation, compared to BSAS.

Sequential clustering algorithms

The maxmin algorithm

W may be initialized by (a) the two most distant points or (b) the mean of the data set.

Let W be the set of all points that have been chosen to define clusters up to the current iteration step. The definition of clusters is carried out as follows:

- For each $x \in X - W$ determine $d_x = \min_{z \in W} d(x, z)$
- **Determine y :** $d_y = \max_{x \in X - W} d_x$
- **If d_y is greater** than a prespecified **threshold (θ)** then
 - **y defines a new cluster**
- **else**
 - the **cluster determination phase** of the algorithm **terminates**.
- **End {if}**

After the definition of the clusters, each unassigned vector is assigned to its closest cluster.

Remarks:

- The maxmin algorithm is more computationally demanding than MBSAS.
- Its result is independent of the order of data presentation to the algorithm.
- It is expected to produce better clustering results than MBSAS.
- Its performance may be degraded in the presence of noise.