


DM534: Introduction to Computer Science
 Autumn term 2018

Exercise Clustering: Clustering, Color Histograms

Exercise Clustering-1 *k*-means, choice of *k*, and compactness

Given the following data set with 8 objects (in \mathbb{R}^2) as in the lecture:



Compute a complete partitioning of the data set into $k = 3$ clusters using the basic k-means algorithm (due to Forgy and Lloyd). The initial assignment of objects to clusters is given using the triangle, square, and circle markers.

Objects x are assigned to the cluster with the least increase in squared deviations $SSQ(x, c)$ where c is the cluster center.










$$SSQ(x, c) = \sum_{i=1}^d |x_i - c_i|^2$$

Start with computing the initial centroids, and draw the cluster assignments after each step and explain the step. Remember to use the least squares assignment!

You can use the data set sketches on the next page.


Give the final quality of the clustering (TD^2). How does it compare with the solutions for $k = 2$ discussed in the lecture? Can we conclude on $k = 3$ or $k = 2$ being the better parameter choice on this data set?

Also compute solutions with $k = 4$, $k = 5$, starting from some random initial assignments of objects to clusters. What do you observe in terms of the TD^2 measure?



Exercise Clustering-2 Furthest First Initialization

Given the following data set with 11 objects (in \mathbb{R}^2):



Aim is now to perform a furthest-first initialization as seen in the lecture.

You should use the following distance measures in order to measure the distance between two points $p = (p_1, p_2)$ and $q = (q_1, q_2)$.

$$\begin{aligned} \text{dist}_2(p, q) &= \left(|p_1 - q_1|^2 + |p_2 - q_2|^2\right)^{\frac{1}{2}} \\ \text{dist}_1(p, q) &= |p_1 - q_1| + |p_2 - q_2| \\ \text{dist}_\infty(p, q) &= \max(|p_1 - q_1|, |p_2 - q_2|) \end{aligned}$$

It might help to fill out the similarity matrix noting all pair-wise distances between all points (note: only the upper triangle is required since the distance functions are symmetric). You find table sketches on the next page.

Let us choose point 3 as our first center. Define the next 3 centers according to the three different norms. (In case two or more points have the same distance, choose the point with the lower point number). Does the sequence of points differ between the norms?

L_2 norm:

	1	2	3	4	5	6	7	8	9	10	11
1	█										
2	█	█									
3	█	█	█								
4	█	█	█	█							
5	█	█	█	█	█						
6	█	█	█	█	█	█					
7	█	█	█	█	█	█	█				
8	█	█	█	█	█	█	█	█			
9	█	█	█	█	█	█	█	█	█		
10	█	█	█	█	█	█	█	█	█	█	
11	█	█	█	█	█	█	█	█	█	█	█

L_1 norm:

	1	2	3	4	5	6	7	8	9	10	11
1	█										
2	█	█									
3	█	█	█								
4	█	█	█	█							
5	█	█	█	█	█						
6	█	█	█	█	█	█					
7	█	█	█	█	█	█	█				
8	█	█	█	█	█	█	█	█			
9	█	█	█	█	█	█	█	█	█		
10	█	█	█	█	█	█	█	█	█	█	
11	█	█	█	█	█	█	█	█	█	█	█

L_∞ norm:

	1	2	3	4	5	6	7	8	9	10	11
1	█										
2	█	█									
3	█	█	█								
4	█	█	█	█							
5	█	█	█	█	█						
6	█	█	█	█	█	█					
7	█	█	█	█	█	█	█				
8	█	█	█	█	█	█	█	█			
9	█	█	█	█	█	█	█	█	█		
10	█	█	█	█	█	█	█	█	█	█	
11	█	█	█	█	█	█	█	█	█	█	█

Exercise Clustering-3 Color-Histograms and Distancefunctions

As a warm-up on distance measures: For each of the following distance measures (Euclidean, Manhattan, maximum, weighted Euclidean, quadratic form)

$$\begin{aligned} \text{dist}_2(p, q) &= \left(|p_1 - q_1|^2 + |p_2 - q_2|^2 + |p_3 - q_3|^2 \right)^{\frac{1}{2}} \\ \text{dist}_1(p, q) &= |p_1 - q_1| + |p_2 - q_2| + |p_3 - q_3| \\ \text{dist}_\infty(p, q) &= \max(|p_1 - q_1|, |p_2 - q_2|, |p_3 - q_3|) \\ \text{dist}_w(p, q) &= \left(w_1|p_1 - q_1|^2 + w_2|p_2 - q_2|^2 + w_3|p_3 - q_3|^2 \right)^{\frac{1}{2}} \\ \text{dist}_M(p, q) &= \left((p - q)M(p - q)^T \right)^{\frac{1}{2}} \end{aligned}$$

calculate the distance between $p = (2, 3, 5)$ and $q = (4, 7, 8)$. As w use $(1, 1.5, 2.5)$ and as M use both of the following:

$$M_1 = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \quad M_2 = \begin{pmatrix} 1 & 0.9 & 0.7 \\ 0.9 & 1 & 0.8 \\ 0.7 & 0.8 & 1 \end{pmatrix}$$

Given 5 pictures as in Figure 1 with 36 pixels each.




Figure 1: 6×6 pixel pictures

- Extract from each picture a color histogram with the bins *red*, *orange*, and *blue* (the white pixels are ignored).
- Which pictures are most similar to the query q , using Euclidean distance? Give a ranking according to similarity to q .
- The results are not entirely satisfactory. What could you change in the feature extraction or in the distance function to get better results? Report the improved feature extraction and features or the improved distance function.