K-means algorithm

GI07/M012

Mark Herbster

University College London
Department of Computer Science

September 30, 2014

Objective

- The aim is to partition unlabelled data into k classes.
- This is “achieved” by optimising the following (nonconvex) objective

$$\operatorname*{arg\,min}_{C_1,\ldots,C_k;\; c_1,\ldots,c_k} \sum_{i=1}^{k} \sum_{x \in C_i} \|x - c_i\|^2 \tag{1}$$

- Thus we find a disjoint partition

$$C_1, \ldots, C_k$$

- As well as centers (or prototypes)

$$c_1, \ldots, c_k$$

k-means

Inputs:	Data: $X := (x_1, x_2, \ldots, x_\ell) \subset \mathbb{R}^n$ Number of classes: k
Initialization:	Choose random centers $c_1 \ldots, c_k$
Algorithm:	1. **for** $i = 1, \ldots, k$ **do** $C_i = \{x \in X \mid i = \arg\min_{1 \le j \le k} \|c_j - x\|^2\}$ 2. **for** $i = 1, \ldots, k$ **do** $c_i = \arg\min_{z \in R^n} \sum_{x \in C_i} \|z - x\|^2$ 3. **While** not converged **go to** step 1.

Convergence

Theorem: k-means converges

Proof.

1. The objective decreases in step 1
2. The objective decreases in step 2
3. The objective is bounded below
4. Hence we converge in “value”

Questions

1. Does gradient descent converge finitely or infinitely?
2. Is convergence of k-means finite or infinite?

K-means clustering - an example

- Suppose we have 4 boxes of different sizes and we want to divide them into 2 classes.
- Each box represents one point with two attributes (X,Y):

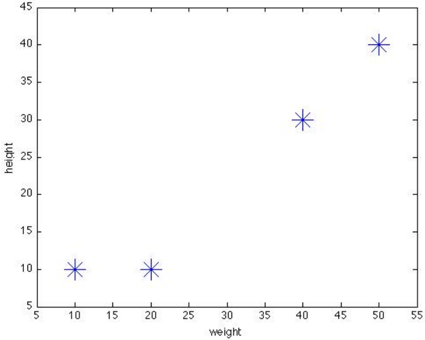

K-means clustering - an example

- *Initial centers:* suppose we choose points A and B as the initial centers, so $c_1 = (10, 10)$ and $c_2 = (20, 10)$
- *Object - centre distance*: calculate the Euclidean distance between cluster centres and the objects. For example, the distance of object C from the first center is:

$$\sqrt{(40-10)^2 + (30-10)^2} = 36.06$$

We obtain the following distance matrix:

0	10	36.06	50
10	0	28.28	42.43

K-means - an example

- *Object clustering*: We assign each object to one of the clusters based on the minimum distance from the centre:

1	0	0	0
0	1	1	1

- *Determine centres*: Based on the group membership, we compute the new centers: $c_1 = (10, 10)$, $c_2 = (\frac{20+40+50}{3}, \frac{10+30+40}{3}) = (36.7, 26.7)$

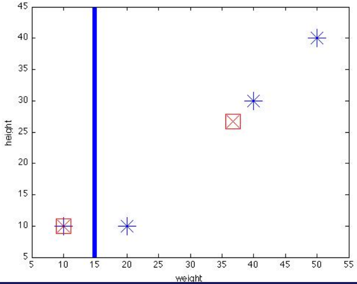

K-means clustering - an example

- *Recompute the object-centre distances*: We compute the distances of each data point from the new centres:

0	10	36.06	50
31.4	23.6	4.7	18.9

- *Object clustering*: We reassign the objects to the clusters based on the minimum distance from the centre:

1	1	0	0
0	0	1	1

- *Determine the new centres*: $c_1 = (\frac{10+20}{2}, \frac{10+10}{2}) = (15, 10)$, $c_2 = (\frac{40+50}{2}, \frac{30+40}{2}) = (45, 35)$

K-means clustering - an example

- *Recompute the object-centres distances*:

5	5	32	46.1
43	35.4	7.1	7.1

- *Object clustering*:

1	1	0	0
0	0	1	1

- The cluster membership did not change from one iteration to another and so the k-means computation terminates.

K-means clustering - an example

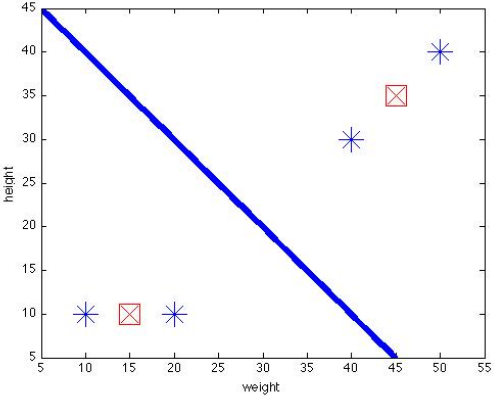

The end