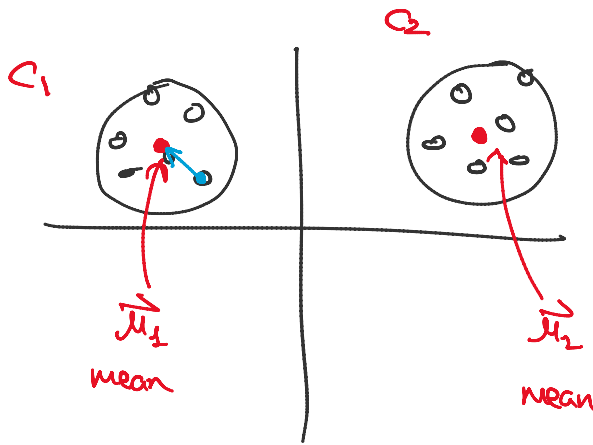


Lecture 20

Monday, November 13, 2023 10:04 AM

Representative Based Clustering



hard clustering

→ partition into k groups

$$\mathcal{C} = \{C_1, C_2, \dots, C_k\}$$

$$C_i \cap C_j = \emptyset \quad \forall i \neq j$$

$$\bigcup_{i=1}^k C_i = D$$

$$\{\vec{\mu}_1, \vec{\mu}_2, \dots, \vec{\mu}_k\}$$

SSE objective \equiv sum of squared errors

Informally: minimize the squared distance of the point from the corresponding mean

$$\min_{(\vec{\mu}_1, \vec{\mu}_2, \dots, \vec{\mu}_k)} = \sum_{i=1}^k \sum_{\vec{x}_j \in C_i} \|\vec{x}_j - \vec{\mu}_i\|^2$$

Search over all possible partitions

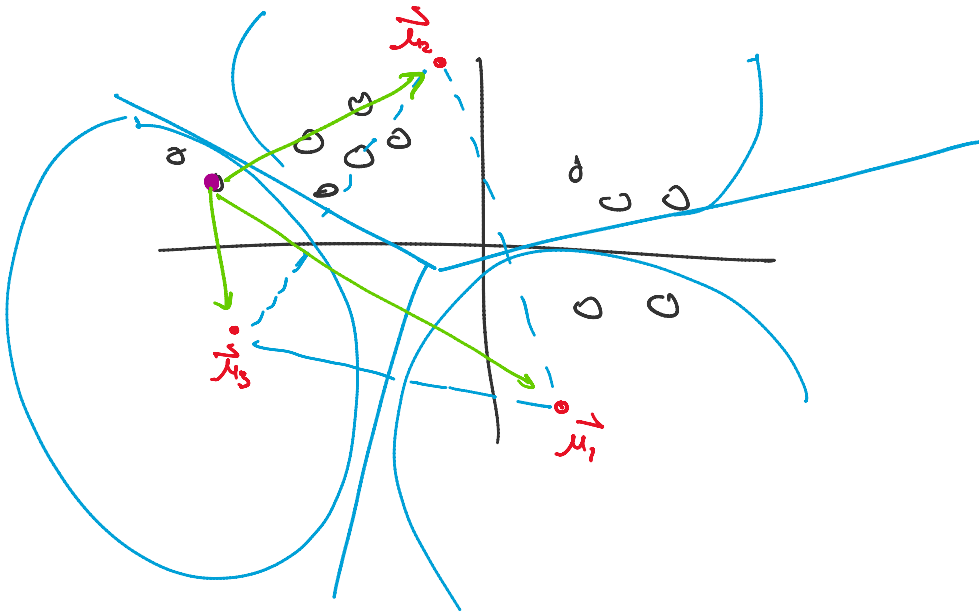
$k=1 \Rightarrow$ Obvious solution is $\vec{\mu}$ for the data

$k=2 \Rightarrow$ NP-hard!

K-Means Algorithm

→ greedy, randomized, iterative algorithm





$$\text{Cluster}(\vec{x}_j) = \arg \min_{i=1 \dots k} \{ \|\vec{x}_j - \vec{\mu}_i\|^2 \}$$

$k\text{-means}(D, k)$
↖ new many clusters?

1) Initialization:

randomly select any k points from D

$$\vec{\mu}_i = \vec{x}_j \text{ for some random } j \in [1, 2, \dots, n]$$

(try to select points that are far apart)

2) Given $\vec{\mu}_1, \vec{\mu}_2, \dots, \vec{\mu}_k$

for each $\vec{x}_j \in D$

assign to closest mean $\vec{\mu}_i$

$$\text{cluster}(\vec{x}_j) = \arg \min_{i=1 \dots k} \{ \|\vec{x}_j - \vec{\mu}_i\|^2 \}$$

$O(n \cdot k \cdot d)$
↖

3) given the clusters: (C_1, C_2, \dots, C_k)

3)

given the clusters: C_1, C_2, \dots, C_k

Update the means:

for $i = 1 \dots k$

$$\vec{\mu}_i = \frac{1}{|C_i|} \sum_{\vec{x}_j \in C_i} \vec{x}_j$$

$O(n)$

$$P(C_i | x_j) = \begin{cases} 0 \\ 1 \end{cases}$$

4) Convergence: repeat until

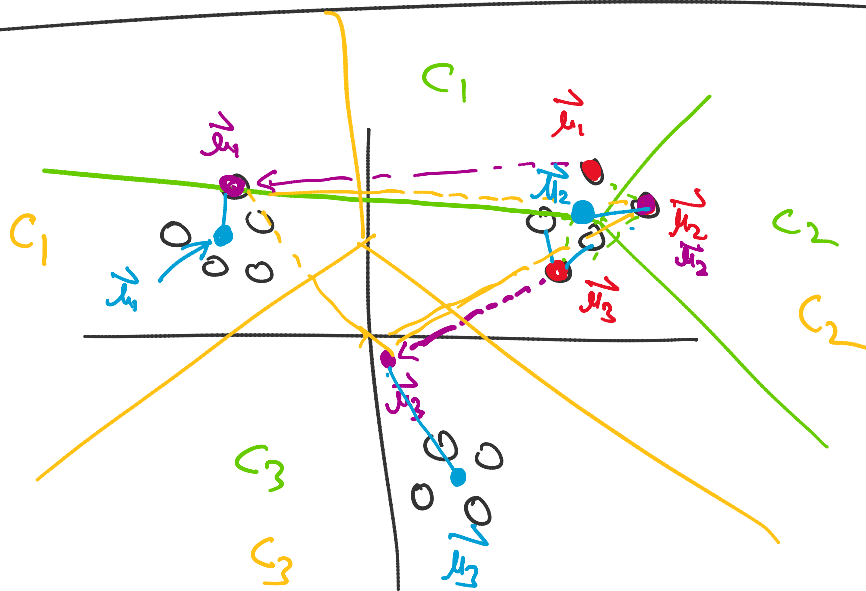
$$\sum_{i=1}^k \left\| \vec{\mu}_i^{(t+1)} - \vec{\mu}_i^{(t)} \right\|^2 \leq \epsilon$$

of iterations?
 T

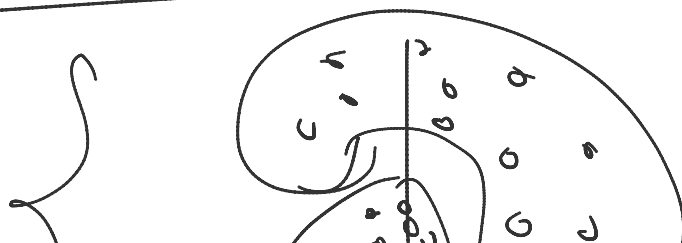
t : iteration

$O(n \cdot k \cdot d \cdot T)$

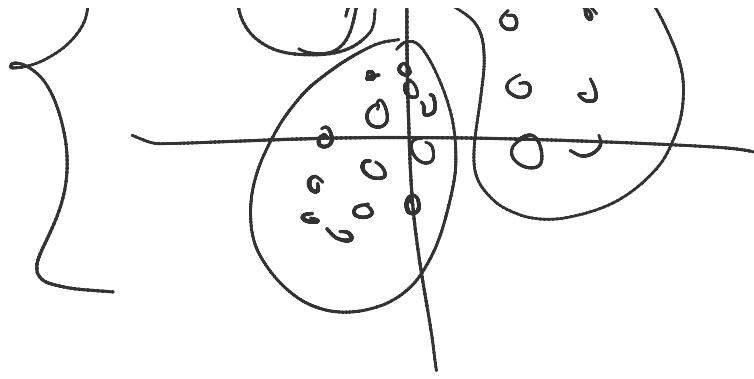
meta: run multiple times & report the \mathcal{C} with smallest SSE



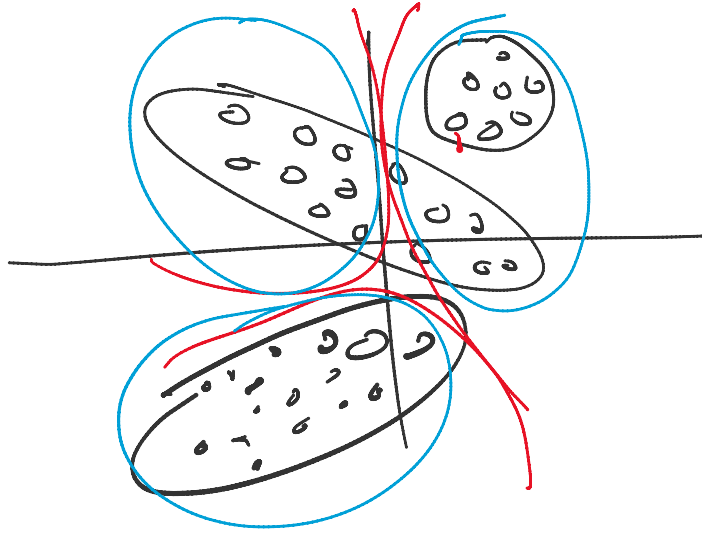
Works well for "Convex" and well separated clusters



will not work here!



will not work here!



Soft-clustering \leftarrow probability $P(C_i | \vec{x}_j)$

\hookrightarrow Expectation Maximization Algorithm

\hookrightarrow maximise the likelihood

Clusters are from a mixture of Gaussians

$$\text{each cluster } C_i \sim N(\vec{x}_j | \underset{\substack{\uparrow \\ \text{mean}}}{\vec{\mu}_i}, \underset{\substack{\uparrow \\ \text{cov}}}{\Sigma_i})$$

$P(C_i)$ = cluster prior prob

$$f(\vec{x}_j) = \sum_{i=1}^k N(\vec{x}_j | \vec{\mu}_i, \Sigma_i) \cdot P(C_i)$$

$$P(x_j) = \sum_{i=1}^k \underline{P(x_j | C_i)} \cdot P(C_i)$$

mixture model

$$L = \text{likelihood} \equiv P(D|\theta)$$

$$\theta \equiv \text{parameters} \rightarrow \underbrace{\mu_i, \Sigma_i, P(C_i)}_{\text{Unknown}} : \forall i=1 \dots k$$

$$\max_{\theta} P(D|\theta) \equiv \max_{\theta} \ln P(D|\theta)$$

$$= \max_{\theta} \left\{ \ln \left(\sum_{i=1}^k N(x_j | \mu_i, \Sigma_i) P(C_i) \right) \right\}$$

$$\frac{\partial L}{\partial \mu_i}, \frac{\partial L}{\partial \Sigma_i}, \frac{\partial L}{\partial P(C_i)}$$

direct solution is "recursive"

EM: Expectation Maximization

"guess" the labels

$$P(C_i | x_j)$$

Expectation

Prob of cluster i given point x_j

very easy to compute $\mu_i, \Sigma_i, P(C_i)$

very easy to compute $\mu_i, \Sigma_i, P(c_i)$
maximization

EM(D, k)

1) Initialization: guess $\vec{\mu}_i, \Sigma_i, P(c_i)$
mean cov size

Choose k points from D at random as $\vec{\mu}_i$
 $\Sigma_i = I$
 $P(c_i) = 1/k$

2) Expectation step

$E[\ln P(D|\theta)]$

Given $\vec{\mu}_i, \Sigma_i, P(c_i) \forall i=1 \dots k$

Compute the posterior prob $P(c_i | x_j) \forall x_j \forall c_i$

$$P(c_i | \vec{x}_j) = \frac{P(\vec{x}_j | c_i) \cdot P(c_i)}{P(\vec{x}_j)} \equiv \frac{\text{likelihood} \cdot \text{prior}}{\text{evidence}}$$

$$P(c_i | x_j) = \frac{P(\vec{x}_j | c_i) \cdot P(c_i)}{\sum_{i=1}^k P(\vec{x}_j | c_i) \cdot P(c_i)}$$

mixture model

$O(n \cdot k \cdot d^2)$
↑
 Σ_i^{-1}

$w_{ij} \equiv P(c_i | x_j) = \frac{N(\vec{x}_j | \vec{\mu}_i, \Sigma_i) \cdot P(c_i)}{\sum_{i=1}^k N(\vec{x}_j | \vec{\mu}_i, \Sigma_i) \cdot P(c_i)}$

↑
weight for cluster c_i

weight for cluster C_i
from point \vec{x}_j

Probabilistic 'labels'

$$\sum_{i=1}^K N(\vec{x}_j | \mu_i, \Sigma_i) P(C_i)$$

weight of \vec{x}_j for cluster i

3) Maximization Step: Given $(w_{ij}) \equiv P(C_i | \vec{x}_j) \forall i, j$
update $\mu_i, \Sigma_i, P(C_i)$

$$\underline{\underline{\mu_i}} = \text{weighted mean} = \frac{\sum_{j=1}^n w_{ij} \cdot \vec{x}_j}{\sum_{j=1}^n w_{ij}} \quad O(n)$$

$$\Sigma_i = \text{weighted cov} = \frac{\sum_{j=1}^n w_{ij} (\vec{x}_j - \vec{\mu}_i) (\vec{x}_j - \vec{\mu}_i)^T}{\sum_{j=1}^n w_{ij}}$$

$$P(C_i) = \frac{\text{fraction of weight in } C_i}{n} = \frac{\sum_{j=1}^n w_{ij}}{n}$$

4) Repeat until convergence

$$\sum_{i=1}^K \left\| \mu_i^{(t+1)} - \mu_i^{(t)} \right\|^2 \leq \epsilon$$

5) Once we have $\Theta = \{ \mu_i, \Sigma_i, P(C_i) \} \forall i$

↳ we "can" assign each point to the most probable cluster

probable cluster

$$\text{cluster}(\vec{x}_j) = \arg \max_{i=1 \dots k} \left\{ \underbrace{P(c_i | \vec{x}_j)}_{w_{ij}} \right\}$$

K-means 'is' just a special case of EM

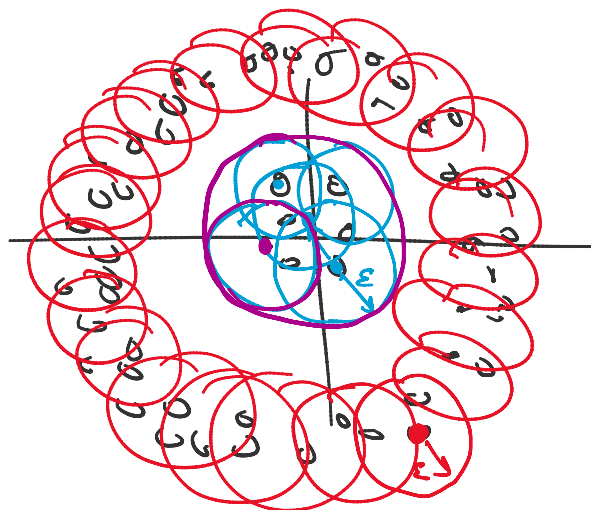
where $P(c_i | x_j) = \begin{cases} 1 \\ 0 \end{cases} \equiv w_{ij}$

further

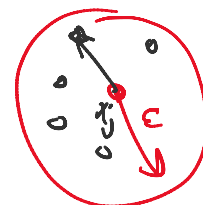
$$\sum_{i=1}^k w_{ij} = 1$$

\Rightarrow one-hot prob assignment

Non-Convex clusters?



density-based
clustering



$\epsilon \leftarrow$ radius
minpts \leftarrow minimum #
of points

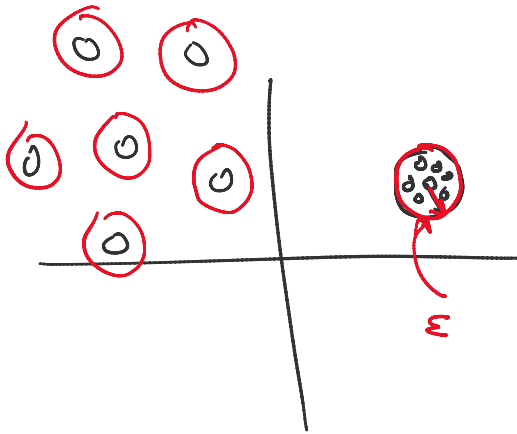
ϵ -ball or

ϵ -neighborhood

$\text{minpts} \leftarrow$ minimum #
of points
"density
constraint"

ϵ - ball or
 ϵ - Neighborhood

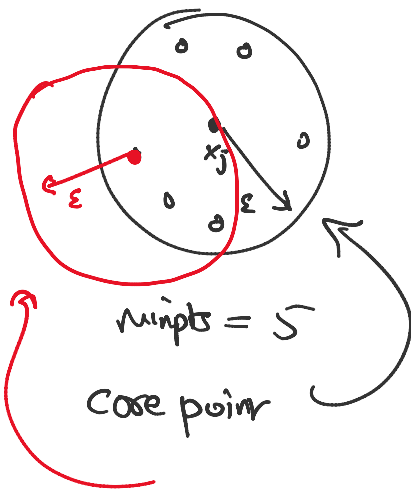
$$|N_{\epsilon}(\vec{x}_j)| \geq \text{minpts}$$



1) Compute "core" points

\vec{x}_j such that

$$|N_{\epsilon}(\vec{x}_j)| \geq \text{minpts}$$



border point
 $|N_{\epsilon}| \leq \text{minpts}$

find connected components of core points

2) $\forall x_i \text{ \& } x_j$

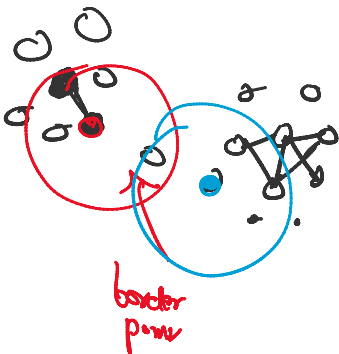
If $x_i \in N_{\epsilon}(x_j)$

and $x_j \in N_{\epsilon}(x_i)$

add an edge between them

3) find connected components over
the graph from step 2

4) for each component & each
core point in there,
include all border points
in the cluster.



DBSCAN = density based clustering