

QUEEN MARY SCHOOL HAINAN
QUEEN MARY UNIVERSITY OF LONDON

QHM5703 Principles of Machine Learning

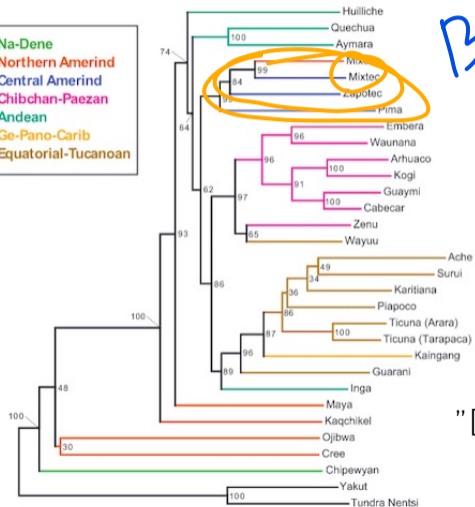
Unsupervised learning: Structure analysis

Dr Nikesh Bajaj
Dr Jesús Requena Carrión

Week 14: 11 Dec 2025


English spoken; American understood

- Na-Dene
- Northern Amerind
- Central Amerind
- Chibchan-Paezan
- Andean
- Ge-Pano-Carib
- Equatorial-Tucanoan



Betty Bought

a bit



Dr. Katherine Dalton

Better Butter

"[Languages] reflect how humans perceive and organise the world around them"

Know your population!

Agenda

Unsupervised learning

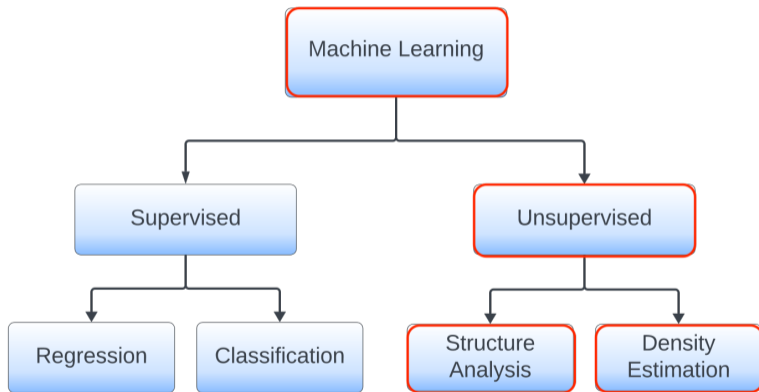
K-means clustering

Density-based clustering

Hierarchical clustering


Summary

Machine Learning taxonomy



Supervised learning

In supervised learning, we **designate** one attribute as a **label** and treat the rest as **predictors**. We set out to build a model that estimates the value of the label based on the value of the predictors.



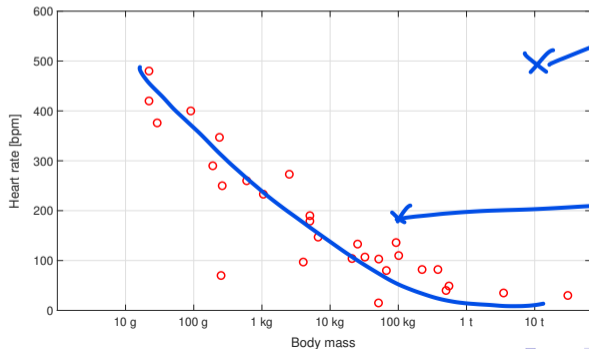
Note that:

- The model mirrors the underlying **structure** of the dataset.
- The notion of quality is defined as a function of the **discrepancy between the true and estimated** values of the label.

The space is empty

Unsupervised learning does not elevate any attribute to the category of label: all the attributes are treated equally. The essence of unsupervised learning is encapsulated in the simple question **where is my data?**

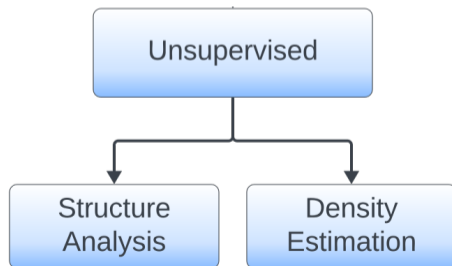
The attribute space is infinite and mostly empty, and the answer to this question will be a **model** that will allow us to identify the regions where we could expect to find samples.



Unsupervised learning

There are two main approaches to answer the question *where is my data*:

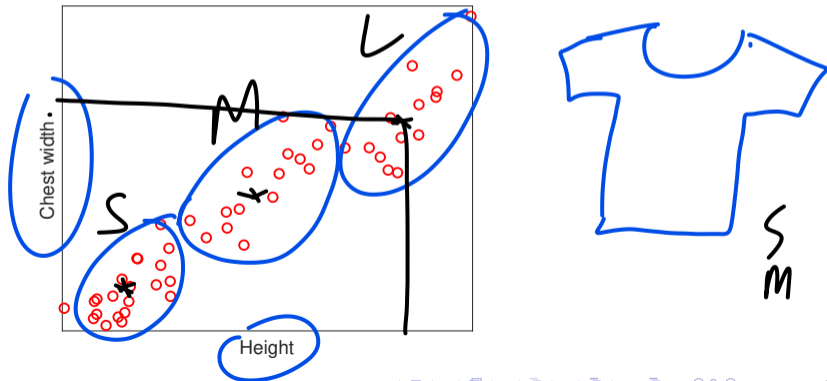
- **Density estimation**: Creates models that allow us to quantify the probability of finding a sample within a region of the attribute space (**probability density**).
- **Structure analysis**: Creates models that identify regions within the attribute space (**cluster analysis**) or directions (**component analysis**) with a high density of samples.



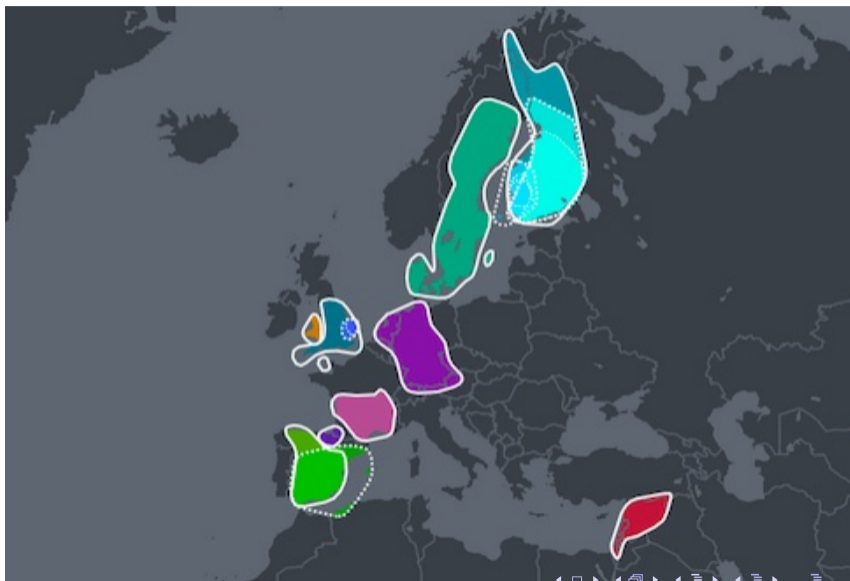
Unsupervised learning: Summaries

Unsupervised learning can be used to provide summaries of a population in the form of **prototypes** samples.

Look at the dataset below. If you had to produce 3 different t-shirt sizes, which sizes would you choose?



Example: DNA analysis




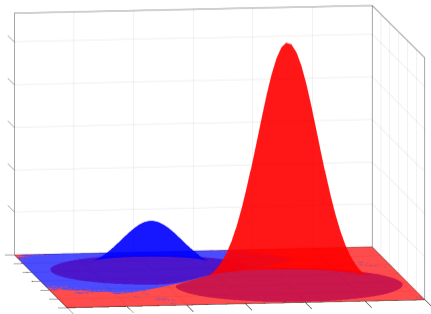
Unsupervised learning: Quantitative applications

Unsupervised learning can be used to build **class densities** that describe the probability of finding a sample from a given class in a region.

A probability density is also useful to identify **anomalies**, i.e. samples that are likely to belong to a different population.

Discriminant Analysis


LDA
QDA



Embedded deployment

In some scenarios, we can distinguish between a main problem and a nested problem, for instance:

- Build a recommendation system (main problem) by...
- Creating categories of customers (nested problem)

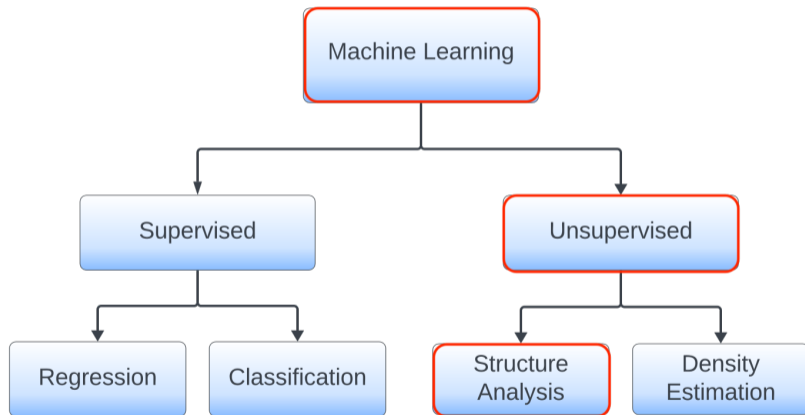
The image shows the word "Netflix" written in a handwritten, cursive style. The letters are black and slanted to the right. A single horizontal line is drawn underneath the entire word, serving as a baseline.

A solution to the nested problem is embedded within the solution of the main problem. New and strange idea:

- We can only assess the quality of a nested solution embedding it into the main solution.

Unsupervised learning models are usually solutions to a nested problem (e.g. customer segmentation). Designing test tasks for unsupervised models is not straightforward. Always think about the task first, then how to use data.

Machine Learning taxonomy



• Agenda

Unsupervised learning

K-means clustering

Density-based clustering

Hierarchical clustering


Summary

Cluster analysis

Clustering is a family of unsupervised learning algorithms that describe the structure of a dataset as groups, or clusters, of **similar samples**.

A notion of **similarity** is therefore needed in order for us to partition a dataset into clusters.

distance




Similarity as proximity

$$x_1 = (0, 0, 1, 2)$$
$$x_2 = (3, 1, 2, 2)$$

Clusters can be defined as groups of samples that are **close** to one another. In this case, we use proximity as our notion of similarity.

Mathematically, there are different ways of defining a distance. Given two samples x_i and x_j consisting of P attributes, $x_{i,1}, \dots, x_{i,P}$ and $x_{j,1}, \dots, x_{j,P}$, the **squared distance** $d_{i,j}$ is defined as

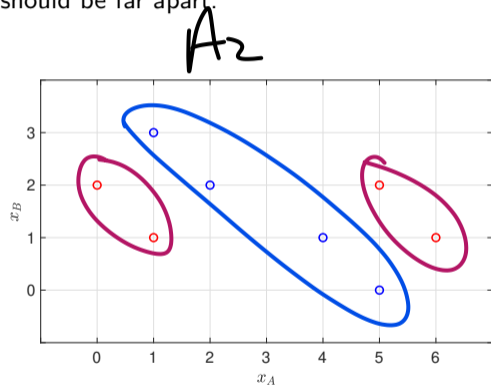

$$d_{i,j}^2 = (x_{i,1} - x_{j,1})^2 + \dots + (x_{i,P} - x_{j,P})^2$$



$$d_{x_1, x_2}^2 = (0-3)^2 + (0-1)^2 + (1-2)^2 + (2-2)^2 = 11$$


A proximity-based quality metric

Using distance as our notion of similarity, samples within the same cluster should be close to one another and samples from different clusters should be far apart.



Our next step will be to create a **cost function** for a clustering arrangement based on the notion of distance between samples.

A proximity-based cost function



Assume we have two clusters C_0 and C_1 . The intra-cluster sample scatter $I(C_0)$ and $I(C_1)$ is the average of square distances between samples in the same cluster.

$$I(C_0) = \frac{1}{2N_0} \sum_{x_i, x_j \text{ in } C_0} d_{i,j}^2$$

$$I(C_1) = \frac{1}{2N_1} \sum_{x_i, x_j \text{ in } C_1} d_{i,j}^2$$

min ↓

and the inter-cluster sample scatter $O(C_0, C_1)$ is defined as the average squared distances between samples in different clusters:

between

$$O(C_0, C_1) = \frac{1}{2N_0, 1} \sum_{x_i \text{ in } C_0, x_j \text{ in } C_1} d_{i,j}^2$$

max ↑

$$\frac{d_{12}^2 + d_{13}^2 + d_{14}^2}{3} + \frac{d_{21}^2 + d_{23}^2 + d_{24}^2}{3}$$

The **best** clustering arrangement has the **lowest** intra-cluster sample scatter and **highest** inter-cluster sample scatter. We can show that reducing the intra-cluster scatter increases the inter-cluster scatter!

~~big~~ ○

$N_{0,1}$

A proximity-based cost function: Problem formulation

$$I(C_1) + I(C_2)$$

Since minimising intra-cluster sample scatter result in maximising inter-cluster sample scatter, we can choose one metric as cost function.

Assume we have a clustering arrangement $\mathcal{A} = \{C_1, C_2, \dots, C_K\}$. We can define an overall cost function for given clustering arrangement as:


$$E\{\mathcal{A}\} = \sum_{C_k \in \mathcal{A}} I(C_k)$$

The **objective of a clustering algorithm** is to find an clustering arrangement \mathcal{A}_{opt} such that it has the lowest $E\{\mathcal{A}_{opt}\}$

A proximity-based quality metric


The intra-cluster and inter-cluster sample scatters allow us to compare and rank different clustering arrangements.

$E \{A_1\}$ ↓



$$I(1) + I(2)$$

$E \{A_2\}$ ↑



$$I(1)$$


The question is, can we create an algorithm capable of automatically identifying the **best clustering** arrangement? This is an **optimisation** question.

K-means clustering: Prototypes

A simple way to describe a cluster is by using **cluster prototypes**, such as the centre of a cluster.

$$p_0 = \left(\frac{15}{4}, \frac{9}{4} \right)$$

$$p_1 = \left(\frac{12}{3}, \frac{15}{3} \right)$$



	x_1	x_2	
}	1	1	0
	1	3	0
}	2	3	0
	1	5	0
	3	5	0
	4	5	0
	5	5	0


K-means clustering: Intra-cluster sample scatter

$$N_0 = 4$$

$$I(C_0) = \frac{1}{2} \sum d_i^2$$

Given a cluster C_0 consisting of N_0 samples, its centre (or mean) μ_0 can be calculated as:


$$\mu_0 = \frac{1}{N_0} \sum_{x_i \text{ in } C_0} x_i$$



Interestingly, the intra-cluster sample scatter can be calculated using the **distance between each sample and the cluster prototype** d_i^2 :

$$I(C_0) = \sum_{x_i \text{ in } C_0} d_i^2$$

Therefore, our notion of clustering quality can be expressed as follows: in a good clustering arrangement, samples are **close to their prototype**.



Other formulas with different scaling:


$$I(C_0) = N_0 \sum_{x_i \text{ in } C_0} d_i^2$$

$$I(C_0) = \frac{1}{N_0} \sum_{x_i \text{ in } C_0} d_i^2$$

$$I = 4 \times \left(\frac{1}{4} (d_1^2 + d_2^2 + d_3^2 + d_4^2) \right)$$

K-means clustering: Intra-cluster sample scatter: Question

Consider two clusters C_o and C_o arrangement \mathcal{A} and prototypes $P_o = (1, 2)$ and $P_o = (5, 1)$



Compute

- $I(C_o) = 1^2 + 1^2 + 1^2 + 1^2 = 4$ (crossed out) $= 1$
- $I(C_o) = 1^2 + 1^2 + 1^2 + 1^2 = 4$ (crossed out) $= 1$
- $E\{\mathcal{A}\} = I(C_o) + I(C_o) = 4 + 4 = 8$
 $1 + 1 = 2$

K-means clustering

1.

K-means partitions a dataset into K clusters represented by their **mean** and proceeds iteratively as follows:


- Prototypes are obtained as the centre (or mean) of each cluster.
- Samples are re-assigned to the cluster with the closest prototype.

As the K-means algorithm proceeds, we will see samples been reassigned to different clusters until at some point we reach a stable solution, where no sample is reassigned.



The final solution is a local optimum, not necessarily the global one.

.


K-means clustering $K=2$



$I(c) \uparrow$



$I(c) \downarrow$



$I(c) \downarrow$

0

0

How many clusters?

K-means requires that we specify the number of clusters K that we want to partition our dataset into.

A natural question is, what's the right number of clusters? The answer to this question depends on the application:

- In some cases we are given the number of clusters (e.g. t-shirt sizes).
- In a discovery scenario we want to find the underlying structure and the number of clusters is unknown.

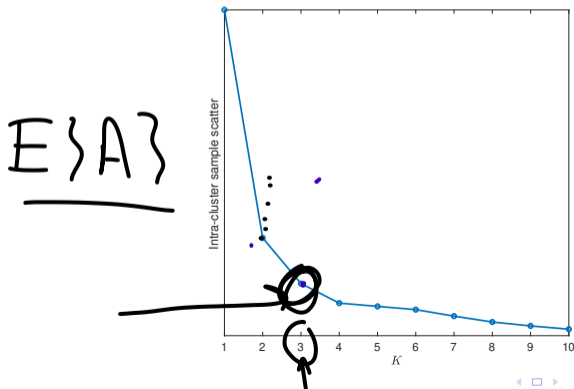
Validation strategies can suggest a suitable value for the hyperparameter K . Choosing the value of K producing the lowest $I(C_0)$ would not work however, as $I(C_0)$ always decreases as the number of clusters increase.

* Many of Attributes: it very diff gives you bad clusters

The elbow method

Assume the true number of clusters is K_T . For $K > K_T$, we should expect the increase in quality to be slower than for $K < K_T$, as we will be splitting true clusters.



The true number of clusters can be identified by observing the value of K beyond which the improvement slows down.



The elbow method: Example

Increasing value of K , split the clusters into smaller ones.

$k=1$
 $k=2$
 $k=3$
 $k=4$



Agenda

Unsupervised learning

K-means clustering


Density-based clustering

Hierarchical clustering

Summary


Non-convex clusters

Similarity
Connected



Non-convex clusters: K-means

K-means produces spherical clusters, as samples are arranged around a prototype.



Density-based clustering: DBSCAN




In a **non-convex** cluster, we can reach any sample by taking small jumps from sample to sample.

Non-convex scenarios suggest a different notion of cluster as group of samples that are **connected**, rather than simply close: *if I am similar to you, and you are similar to them, I am similar to them too.*

This notion of cluster as a group of connected samples is behind many clustering algorithms, such as DBSCAN (*density-based spatial clustering of applications with noise*).

DBSCAN belongs to the family of **density-based** algorithms, where an estimation of the density of samples around each sample is used to partition the dataset into clusters.

DBSCAN





DBSCAN defines two quantities, a radius r and a threshold t . A density is first calculated as the number of samples in a neighbourhood of radius r around each sample (excluding itself). Then, three types of samples are identified:

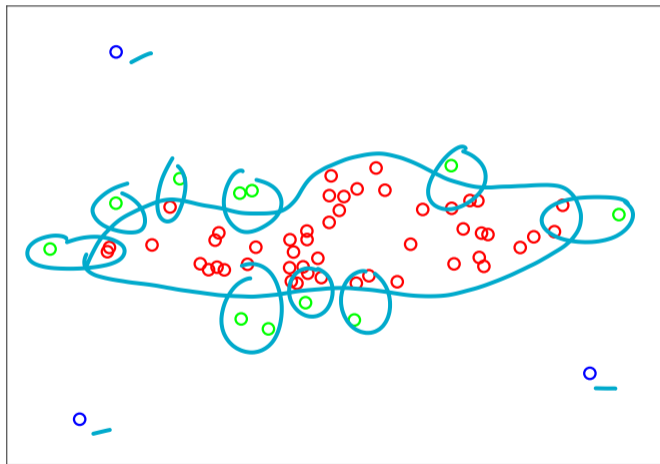
- Core: its density is equal or higher than the threshold t .
- Border: its density is lower than the threshold t , but contains a core sample within its neighbourhood.
- Outlier: Any other sample.

DBSCAN

$t=3$ $\epsilon=0.5$



DBSCAN




DBSCAN

The DBSCAN algorithm proceeds as follows:

- Identify core, border and outlier samples.
- Pair of core samples that are within each other's neighbourhood are connected. Connected core samples form the **backbone** of a cluster.
- Border samples are assigned to the cluster that has more core samples in the neighbourhood of the border sample.
- Outlier samples are not assigned to any cluster.


DBSCAN

Solutions for a threshold $t = 3$ and different $r = 0.7$.



DBSCAN

Solutions for a threshold $t = 3$ and different $r=0.5$.



DBSCAN

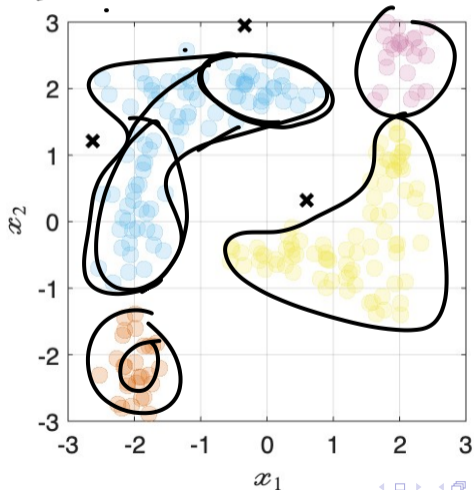
Solutions for a threshold $t = 3$ and different $r = 0.4$.

$t = 5$

$r = 0.1$

$\gamma = 0.1$

$t = 10$



Agenda

Unsupervised learning

K-means clustering

Density-based clustering

Hierarchical clustering

Summary

Hierarchical clustering

Given a dataset consisting of N samples, there exist two **trivial** clustering solutions: **one single cluster** that includes all the samples, and the solution where **each sample is a cluster** on its own.

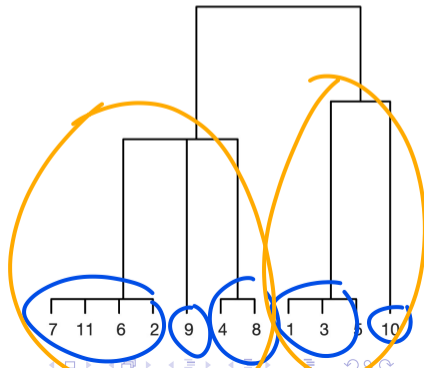

K-means produces K clusters, but we need to choose K within $1 \leq K \leq N$. In DBSCAN clusters are discovered automatically, but the final number of clusters depends on the values of the radius r and the threshold value t .

This ambiguity ultimately reveals that the structure of a dataset can be explored at different levels that expose different properties.

Hierarchical clustering

Hierarchical clustering is a family of clustering approaches that proceed by progressively building clustering arrangements at **different levels**.

The resulting collection of clustering arrangements is hierarchical in the sense that a cluster in one level contains all the samples from one or more clusters in the level below.



Hierarchical clustering

The representation of the relationship between clusters at different levels is called a **dendrogram**. At the bottom we find the arrangement where each sample is one cluster and at the top, the whole dataset.


There exist two basic strategies to build a **dendrogram**:

- The **divisive** or **top-down** approach splits clusters starting from the top of the dendrogram and stops at the bottom level.
- The **agglomerative** or **bottom-up** merges two clusters, starting from the bottom until we reach the top level.





$K=1$ \longrightarrow $K=N$

Hierarchical clustering




There are different options to decide which clusters to merge or split at each level. Common strategies in agglomerative clustering include:

- **Single linkage:** uses the distance between the two closest samples from two clusters. This option results in clusters of arbitrary shapes.
- **Complete linkage:** uses the distance between the two further samples from each pair of clusters. This choice produces clusters that tend to have a spherical shape.
- **Group average:** uses the average distance between samples in two cluster and also produces spherical shapes, although they are more robust to outliers.




Hierarchical clustering



$k=1$

Hierarchical clustering



Agenda

Unsupervised learning

K-means clustering

Density-based clustering

Hierarchical clustering

Summary

Unsupervised learning

- Unsupervised learning provides with answers for the basic question *where is my data? (in the attribute space)*
- Our answer is a mathematical/computer model. This model can tell us where in the space we have samples (clustering) or the probability to find a sample in a region within the space (density estimation).
- Sometimes we say that our data is unlabelled. What we **really mean** is that we don't treat any attribute as a label that we want to predict. Datasets are neither labelled nor unlabelled.
- Lacking such a target as a label means that **our quality metric is not obvious**.

Clustering

- **K-means** is a prototype-based clustering that produces spherical clusters where K is a hyperparameter that has to be set.
- **DBSCAN** is a density-based option suitable for non-convex scenarios and does not require specifying the number of clusters. We need to set r and t and they determine the final number of clusters.
- **Hyierarchical clustering** allows to explore the structure of a dataset at multiple levels.

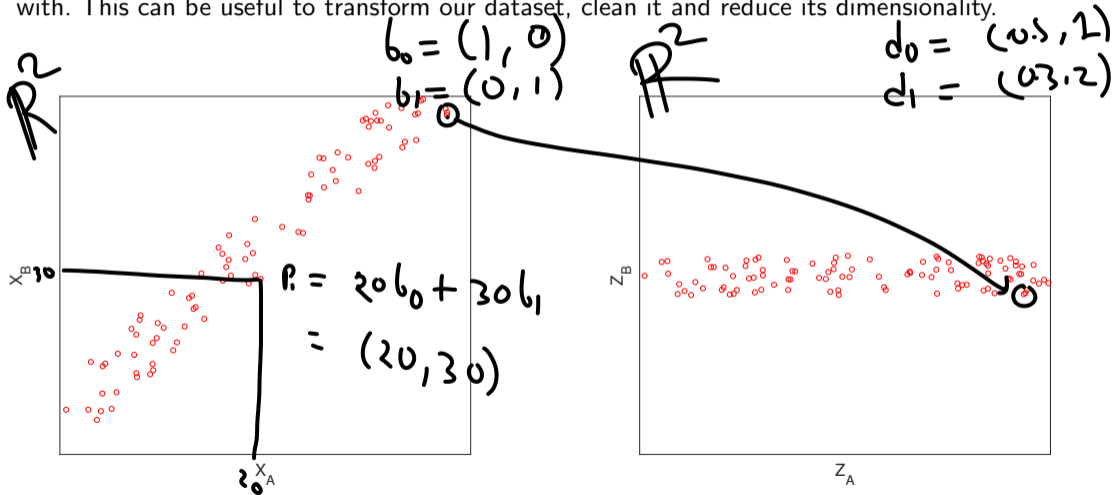
Comparing clustering solutions

- We could consider **comparing** the solutions from two different algorithms. However, if they use different definitions of clustering quality, this comparisons will make little sense.
- Clustering is ultimately implemented with an **application** in mind so we should create a final notion of clustering quality based on the specific goals of the application.

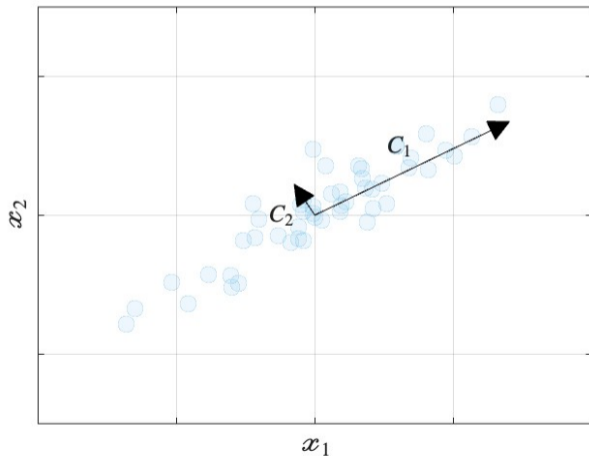
What about component analysis?

basis discovery

Component analysis allows us to identify the **directions in the space** that our data are aligned with. This can be useful to transform our dataset, clean it and reduce its dimensionality.




Principle Component Analysis (PCA)



What about component analysis?

Principal Component Analysis : Independent Study

Reading:

- Chapter 5: Principles of Machine Learning.
https://PMLBook.github.io/Ch_discovery.html#basis-discovery 
- A Tutorial on Principal Component Analysis by Jonathon Shlens.
<https://arxiv.org/abs/1404.1100>



Queen Mary
University of London