

# Unsupervised methods (clustering)

INFO 523 - Lecture 10

Dr. Greg Chism

One main question for this  
and the next lecture:

**What is modelling?**

# Mapping problems to machine learning tasks

- Predicting what customers might buy, based on past transactions
- Identifying fraudulent transactions
- Determining price elasticity (the rate at which a price increase will decrease sales, and vice versa) of various products or product classes
- Determining the best way to present product listings when a customer searches for an item
- Customer segmentation: grouping customers with similar purchasing behavior
- AdWord valuation: how much the company should spend to buy certain AdWords on search engines
- Determine which variable has larger effects on a particular outcome

# Solving classification problems



**Figure 5.2** Assigning products to product categories

# Solving classification problems


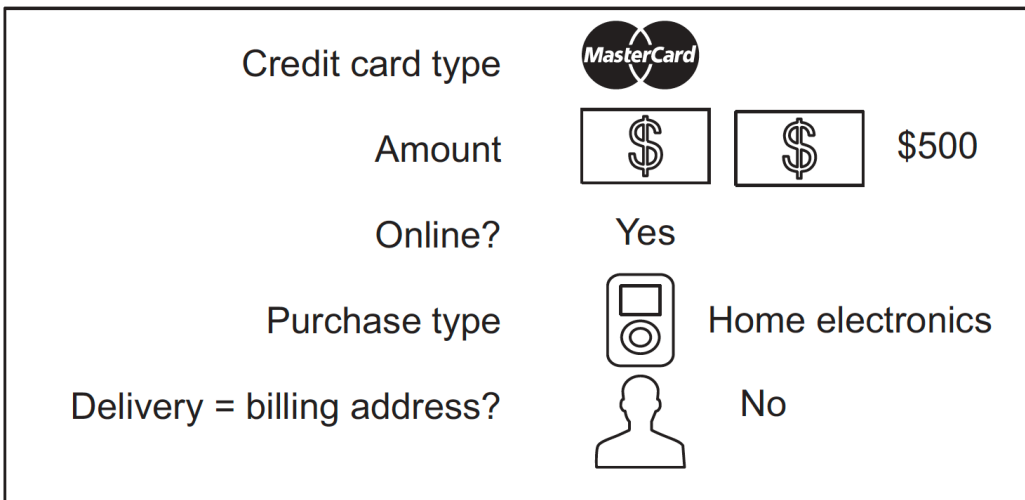

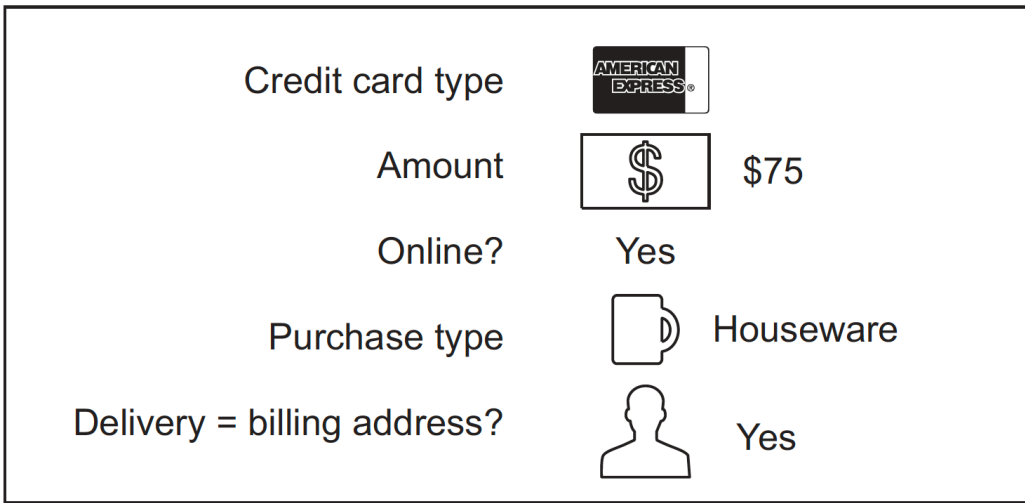
- Classification itself is an example of what is called *supervised learning*.
- Multicategory vs. two-category classification
  - using binary classifiers to solve multicategory problems : building one classifier for each category (a “one versus rest” classifier)
  - in most cases a suitable multiple-category implementation better than multiple binary classifiers: using the package *mlogit* instead of the base method *glm()* for logistic regression

# Solving classification problems

- Some common classification methods
  - Naive Bayes : a good first attempt at solving the product categorization problem.
  - Decision trees : an important extension of decision trees - random forests
  - Logistic regression : estimate class probabilities (the probability that an object is in a given class) in addition to class assignments
  - Support vector machines : SVMs make fewer assumptions about variable distribution than do many other methods

# Solving scoring problems

Probability of fraud



**Figure 5.3** Notional example of determining the probability that a transaction is fraudulent

# Solving scoring problems

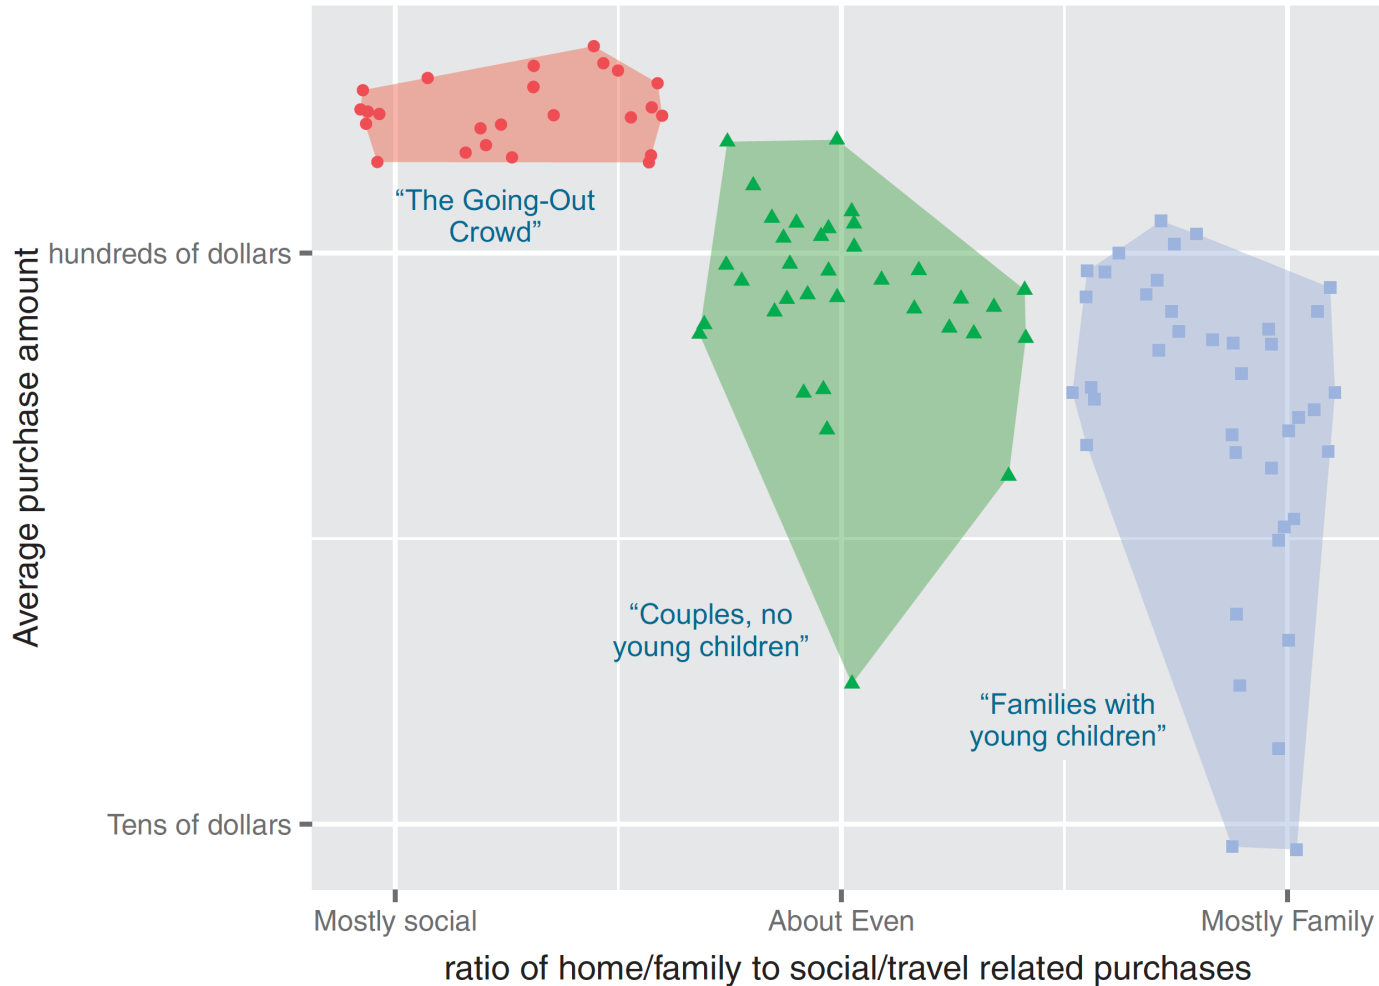
- Linear regression
  - Linear regression builds a model such that the predicted numerical output is a linear additive function of the inputs.
  - a good first model to try when trying to predict a numeric value
- Logistic regression
  - Logistic regression always predicts a value between 0 and 1
  - I.e., if what you want to estimate is the probability that a given transaction is fraudulent or legitimate.



# Working without known targets


- *unsupervised learning*: there's not (yet) a specific outcome that you want to predict
  - discover similarities and relationships in the data. rather than predicting outputs based on inputs
- Common clustering methods
  - **K-means** clustering
  - **A priori** algorithm for finding association rules
  - Nearest neighbor

# WHEN TO USE BASIC CLUSTERING



**Figure 5.4** Notional example of clustering your customers by purchase pattern and purchase amount

# WHEN TO USE BASIC CLUSTERING



Sometimes basic clustering algorithms (k-means) fail...

# WHEN TO USE ASSOCIATION RULES



Bikini, sunglasses, sunblock, flip-flops



Swim trunks, sunblock



Tankini, sunblock, sandals



Bikini, sunglasses, sunblock



One-piece, beach towel

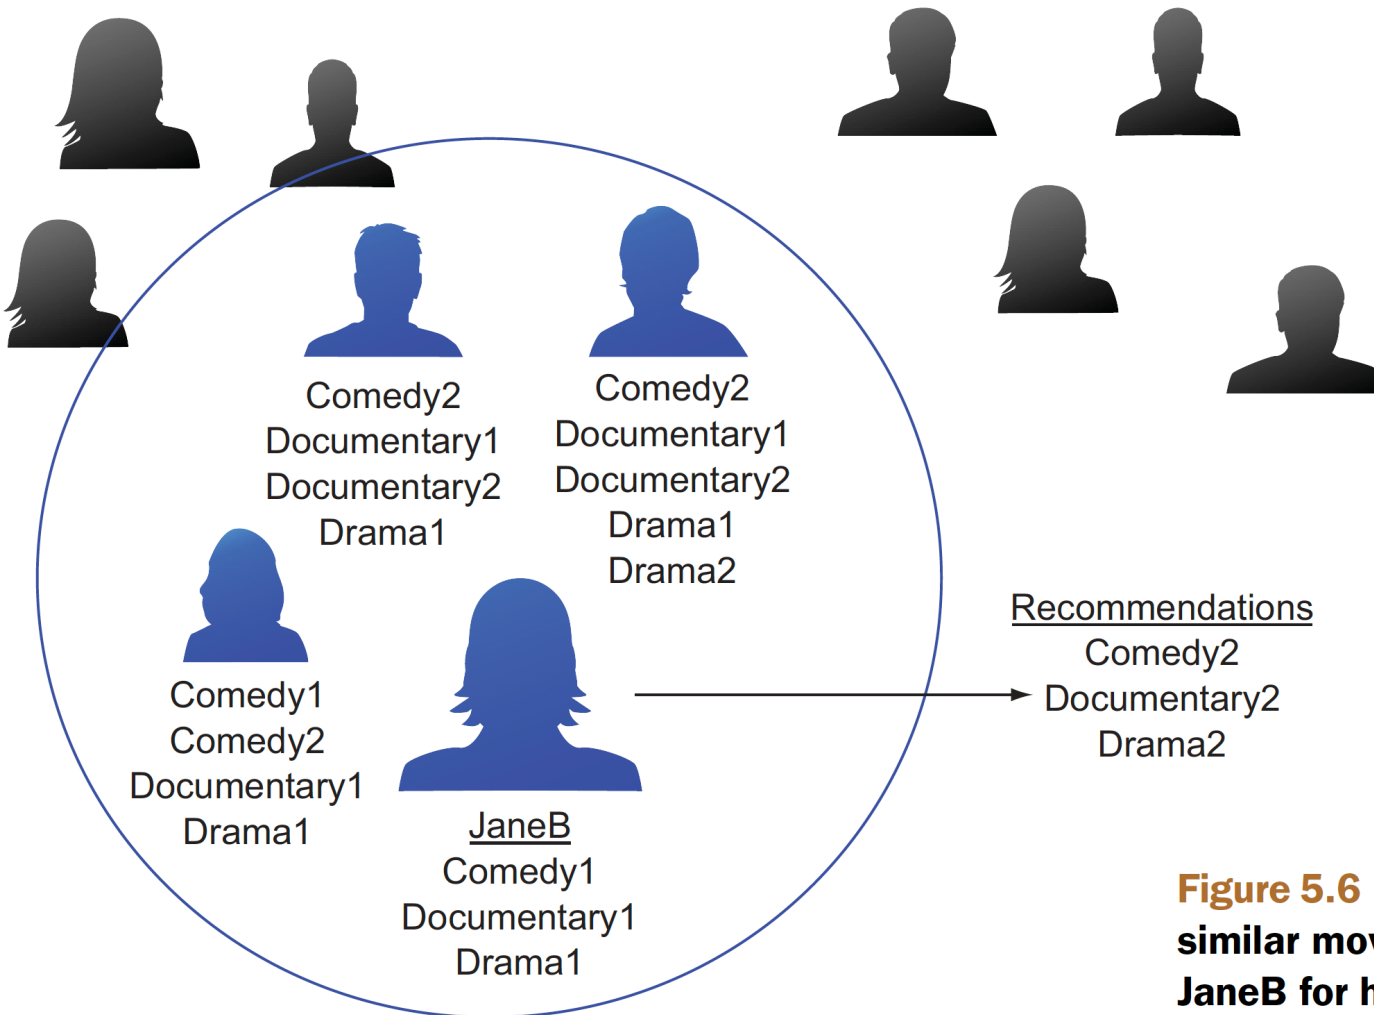
80% of purchases include both a bathing suit and sunblock.

80% of purchases that include a bathing suit also include sunblock.

So customers who buy a bathing suit might also appreciate a recommendation for sunblock.

**Figure 5.5** Notional example of finding purchase patterns in your data

# WHEN TO USE NEAREST NEIGHBOR METHODS



**Figure 5.6** Look to the customers with similar movie-watching patterns as JaneB for her movie recommendations.

# Unsupervised methods

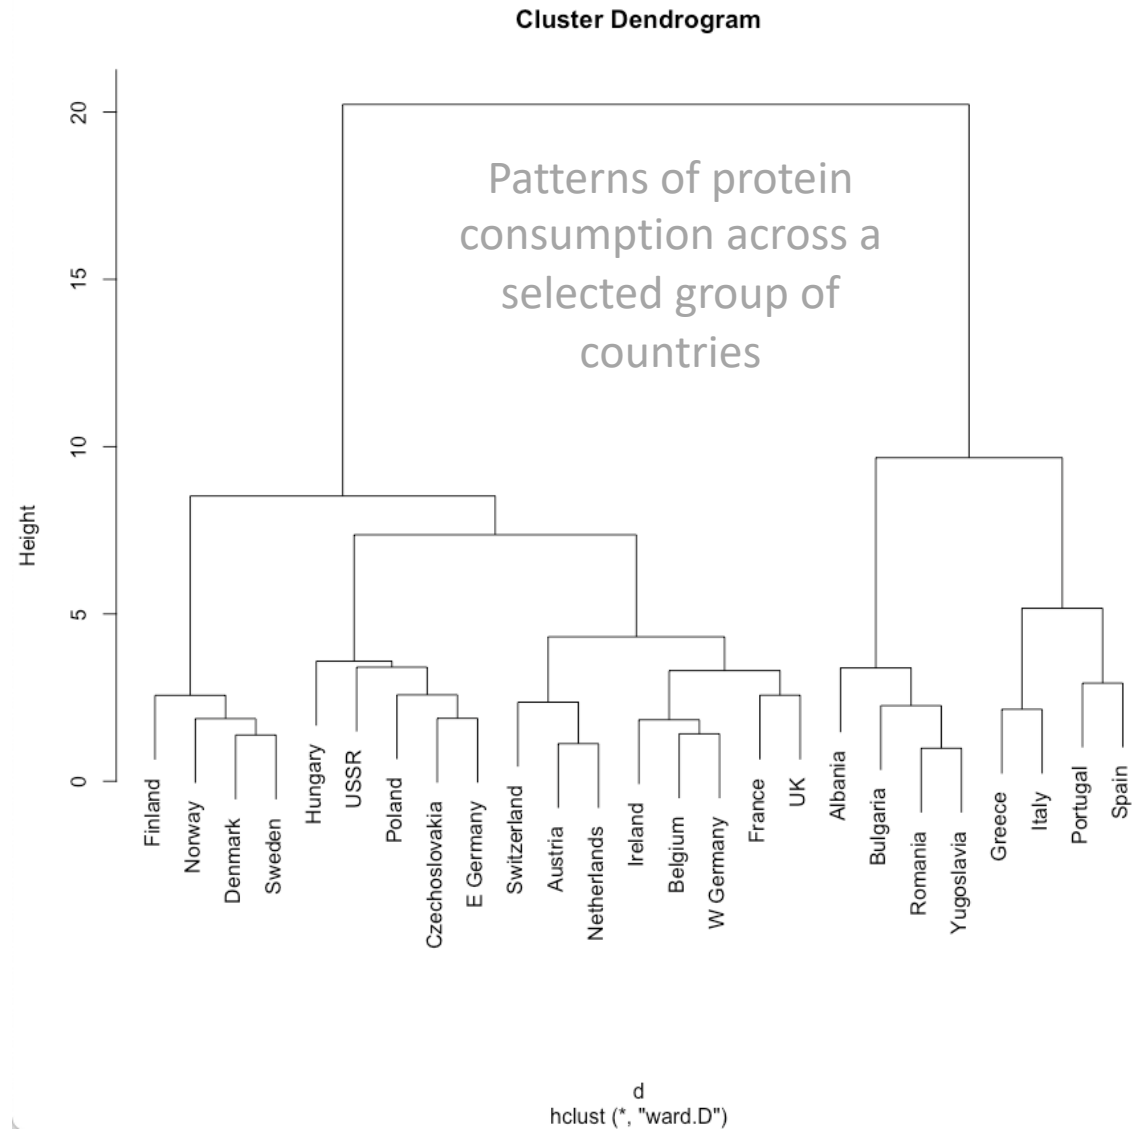
- Methods to discover unknown relationships in data...Central to data mining!
- Two main classes of unsupervised methods
  - *Cluster analysis*: finds groups in your data with similar characteristics
    - Hierarchical clustering
    - k-means
  - *Association rule*: mining finds elements or properties in the data that tend to occur together.

# Distances – Many indexes...

- Euclidean distance (length of the line segment between points, L2 distance)
  - $\text{edist}(x, y) \leftarrow \sqrt{(x[1]-y[1])^2 + (x[2]-y[2])^2 + \dots}$
  - When all the data is real-valued (quantitative)
- Hamming distance
  - $\text{hdist}(x, y) \leftarrow \text{sum}((x[1] \neq y[1]) + (x[2] \neq y[2]) + \dots)$
  - For categorical and binary variables (male /female , or small /medium /large )
- Manhattan (city block) distance : L1 distance
  - $\text{mdist}(x, y) \leftarrow \text{sum}(\text{abs}(x[1]-y[1]) + \text{abs}(x[2]-y[2]) + \dots)$
  - Number of horizontal and vertical units from one point to the other.
- Cosine similarity
  - $\text{dot}(x, y) \leftarrow \text{sum}(x[1]*y[1] + x[2]*y[2] + \dots)$
  - $\text{cossim}(x, y) \leftarrow \text{dot}(x, y) / (\sqrt{\text{dot}(x,x)*\text{dot}(y,y)})$
  - $1 - 2*\text{acos}(\text{cossim}(x,y))/\text{pi}$

# Hierarchical clustering

Distance metric: Euclidean  
Cluster algorithm: Ward's





# Hierarchical clustering

- The dendrogram suggests five clusters (or more? Or less? Let's talk about that later!) = > draw the rectangles on the dendrogram

- `rect.hclust(pfit, k=5)` 

- Extracting the clusters found by `hclust()`

```
groups <- cutree(pfit, k=5)
```

```
print_clusters <- function(labels, k) {
```

```
  for(i in 1:k) {
```


```
    print(paste("cluster", i))
```

```
    print(protein[labels==i,c("Country","RedMeat","Fish","Fr.Veg
```

```
  ]
```

```
}
```

```
print_clusters(groups, 5)
```



# VISUALIZING CLUSTERS...again, many options!

- Let's check out a PCA: `prcomp()`


- PCA

```
library(ggplot2)
princ <- prcomp(pmatrix)
nComp <- 2
project <- predict(princ,
  newdata=pmatrix)[,1:nComp]
```

- Plot

```
project.plus <-
cbind(as.data.frame(project),
  cluster=as.factor(groups),
  country=protein$Country)
ggplot(project.plus, aes(x=PC1,
  y=PC2)) +
  geom_point(aes(shape=cluster)) +
  geom_text(aes(label=country),
  hjust=0, vjust=1)
```

Why do we need to use the `predict` function here?

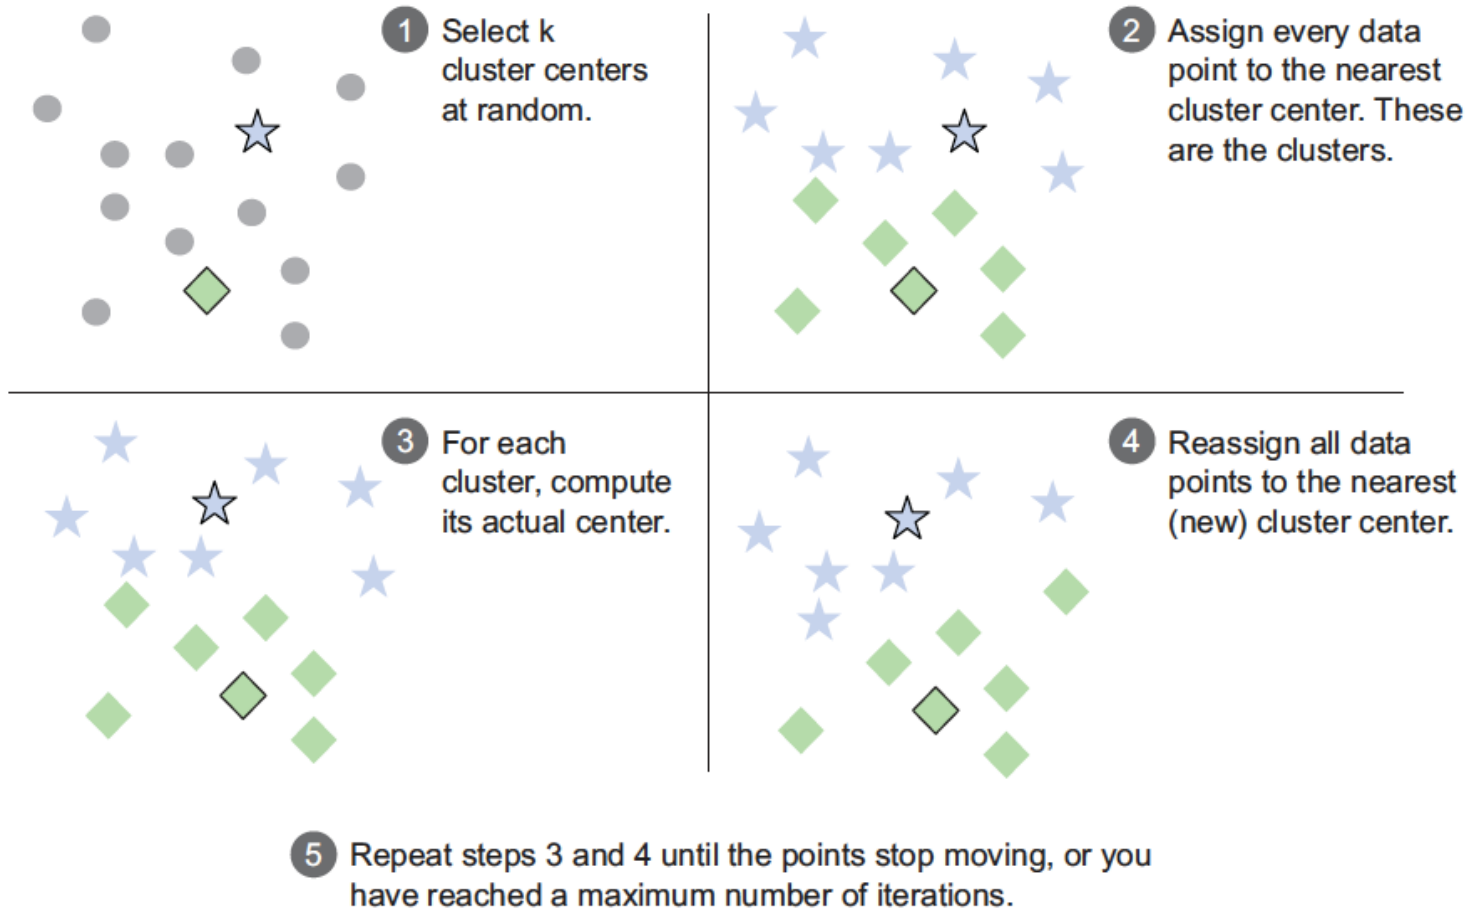


# The k-means algorithm

- The data is all numeric and the distance metric is squared Euclidean (in theory, you could also run it with other distance metrics)
- Plus side...
  - easy to implement
  - can be faster than hierarchical clustering on large datasets
  - works best on data that looks like a mixture of Gaussians
- Negative side
  - must pick  $k$  in advance
  - fairly unstable : the final clusters depend on the initial cluster centers.  
=> This algorithm isn't guaranteed to have a unique "final solution".
  - Not a probabilistic algorithm...

# The k-means procedure

Cluster Dendrogram



# Basic general strategy

1. Cluster the data as usual.
  2. Draw a new dataset (of the same size as the original) by resampling the original dataset with replacement (meaning that some of the data points may show up more than once, and others not at all). Cluster the new dataset.
  3. For every cluster in the original clustering, find the most similar cluster in the new clustering (the one that gives the maximum Jaccard coefficient) and record that value.
    - If this maximum Jaccard coefficient is less than 0.5, the original cluster is considered to be dissolved.
- Repeat steps 2–3 \*many\* times.
  - **Do NOT forget about the difference between training and testing datasets!**

# Cluster stability


- The cluster stability of each cluster in the original clustering is the mean value of its Jaccard coefficient over all the bootstrap iterations.
- clusters with a stability value
  - $0.6 \leq$  : unstable.
  - $0.6 \sim 0.75$  : measure a pattern in the data, but there isn't high certainty about which points should be clustered together.
  - $0.85 \geq$  : highly stable (they're likely to be real clusters).

# PICKING THE NUMBER OF CLUSTERS

- Cross validation (CV) – extremely common!
- Internal indexes (to minimize or maximize).
- For instance:
  - Variances of within cluster and between clusters
  - Rate-distortion method
  - F-ratio
  - Davies-Bouldin index (DBI)
  - Bayesian Information Criterion (BIC)
  - Silhouette Coefficient

# The general idea (e.g. sum of squared errors)

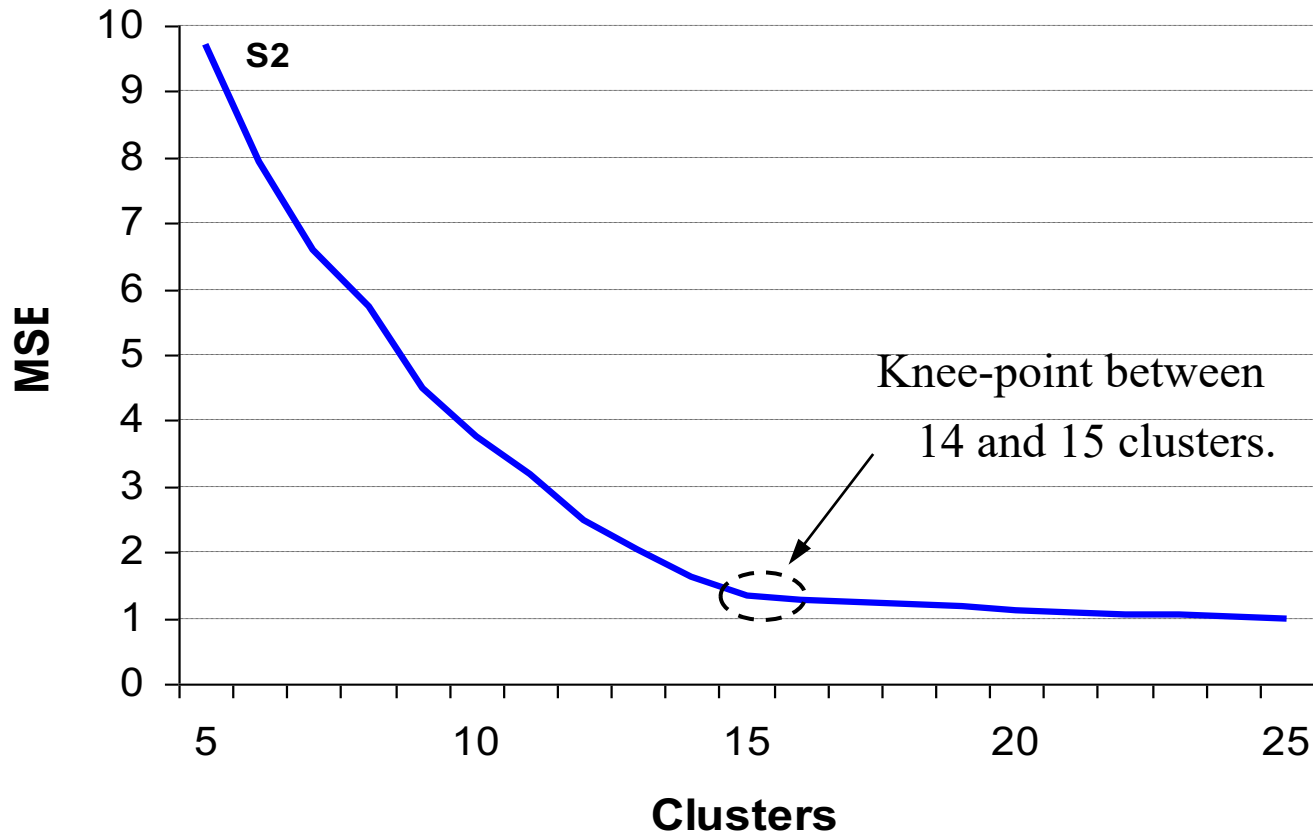
- Minimize within cluster variance
- Maximize between cluster variance





# Sum of squared errors

- The more clusters the smaller the value.
- Small knee-point near the correct value.
- But how to detect?



# Cluster validation

## Supervised classification:

- Ground truth class labels known
- Accuracy, precision, recall

Precision =  $5/5 = 100\%$   
Recall =  $5/7 = 71\%$


## Cluster analysis:

- No class labels



## Validation need to:

- Compare clustering algorithms
- **Solve the number of clusters**
- Avoid finding patterns in noise

Oranges:



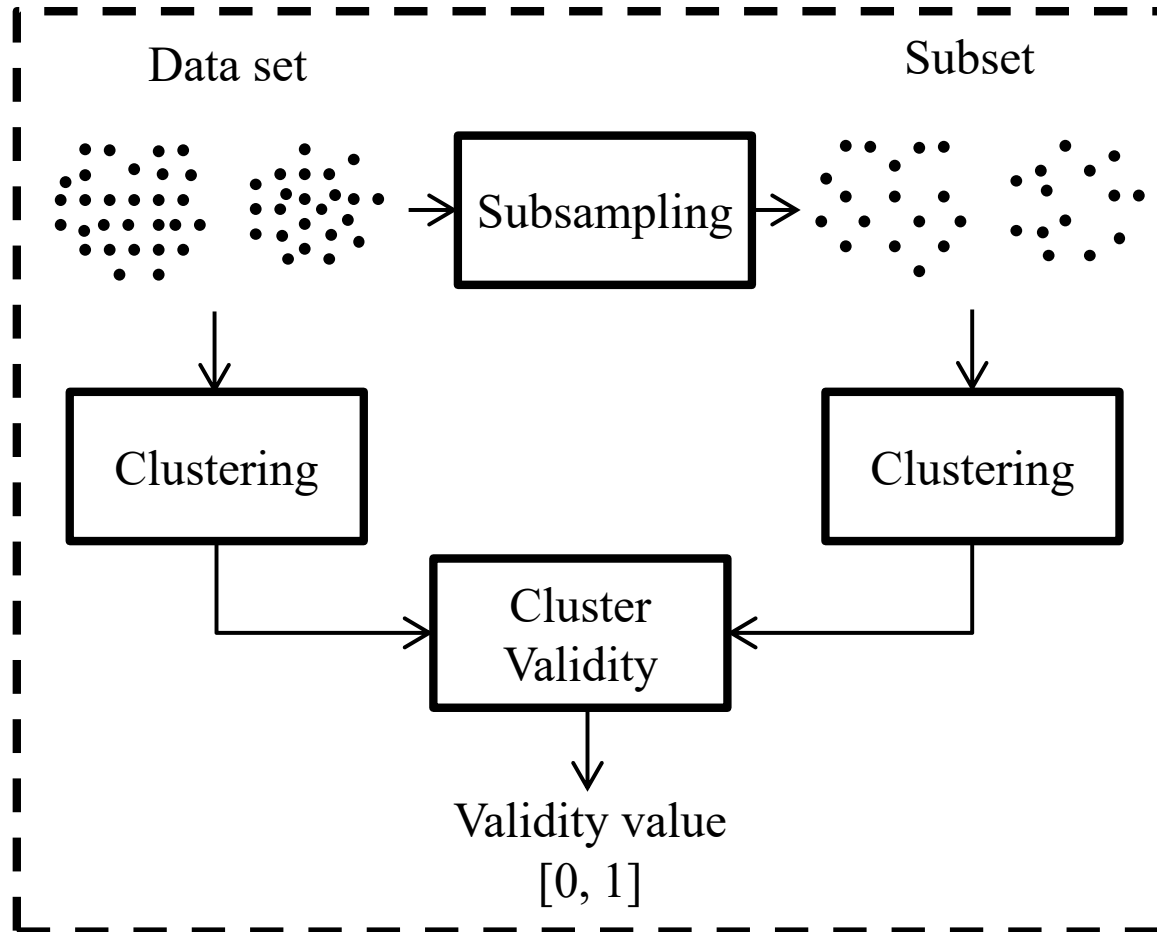
Apples:



Precision =  $5/5 = 100\%$   
Recall =  $3/5 = 60\%$


# Cross-validation

Compare clustering of full data against sub-sample



# Cross-validation: Correct

Correct number of clusters:  $k=5$

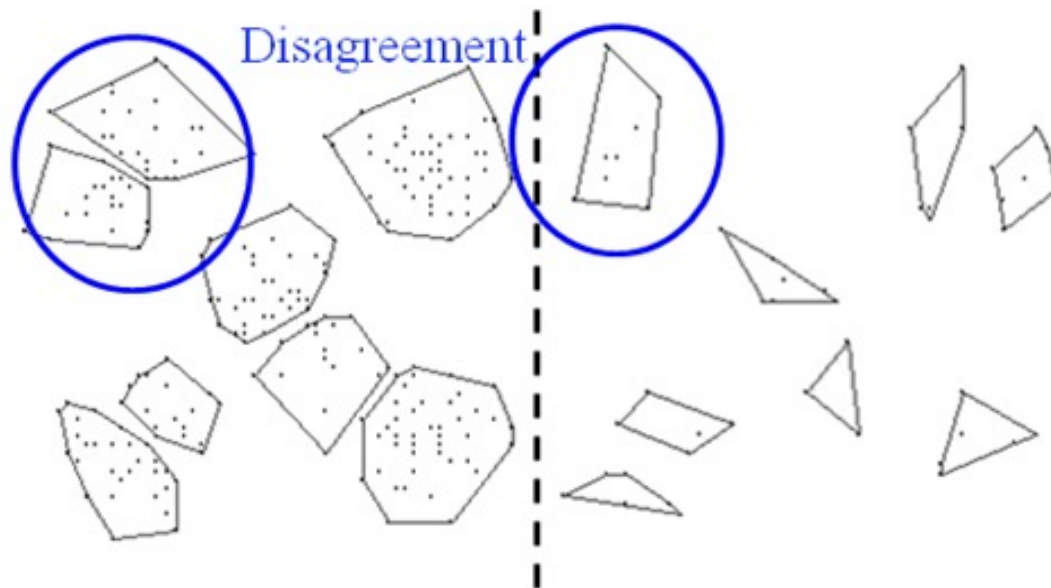


Same results

# Cross-validation

Incorrect

**Incorrect number of clusters:  $k=8$**



Different results

# Stability approach in general

1. Add randomness
2. Cross-validation strategy
3. Solve the clustering
4. Compare clustering


# Adding randomness

- Three choices:
  1. Subsample
  2. Add noise
  3. Randomize the algorithm
- What subsample size?
- How to model noise and how much?
- Use k-means?

# Sub-sample size

- Too large (80%): same clustering always
- Too small (5%): may break cluster structure
- Recommended 20-40%

**Spiral dataset**



**60% subsample**

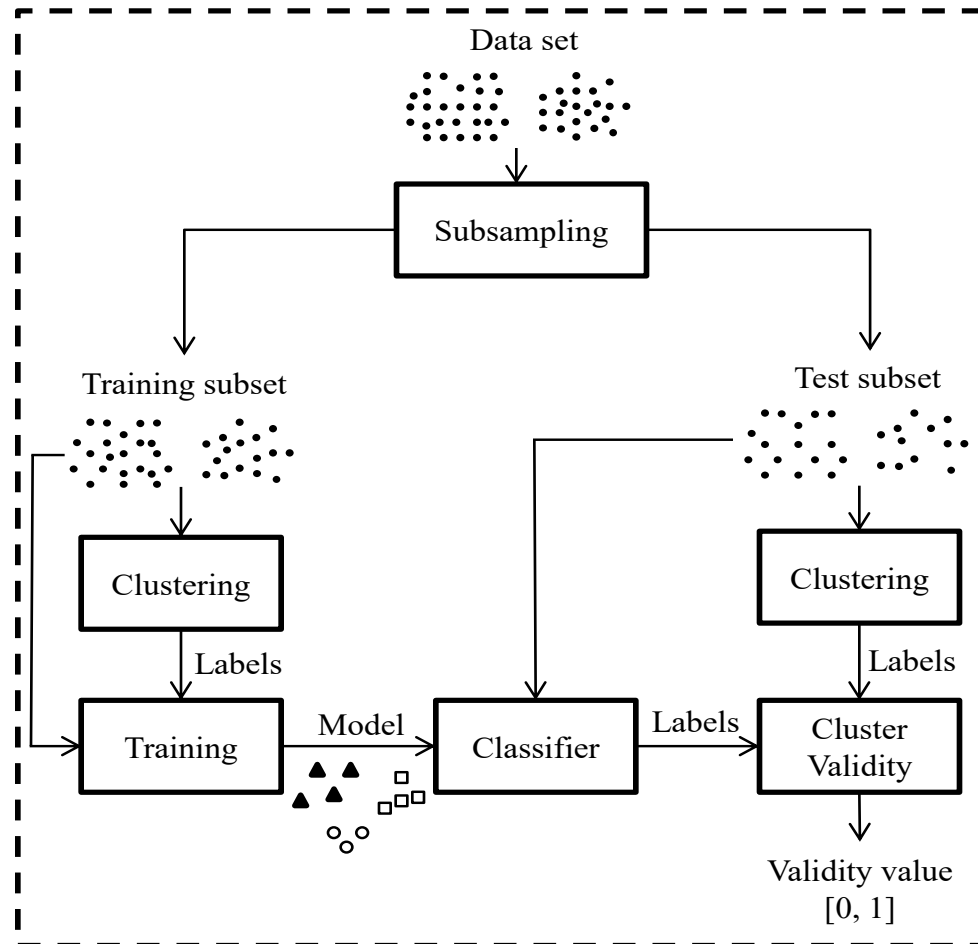


**20% subsample**





# Classification approach



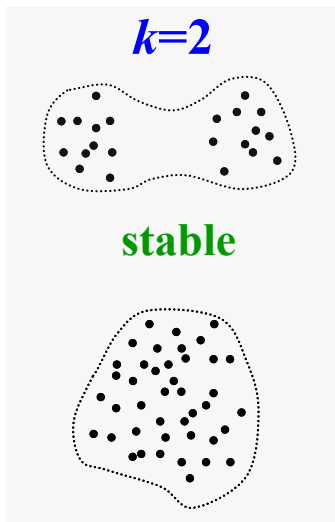
# Problem

Stability can also come from other reasons:

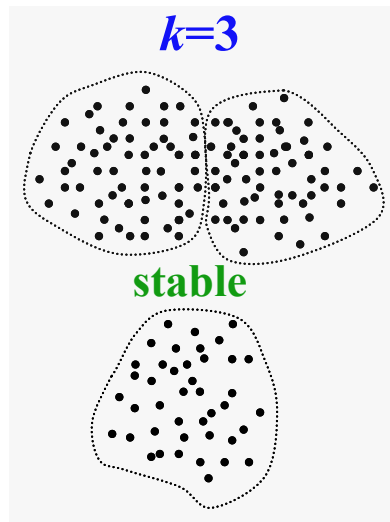
- Different cluster sizes
- Wrong cluster model

Happens when  $k < k^*$

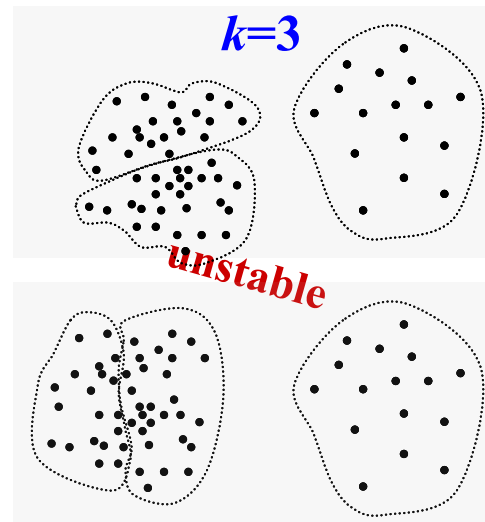
Too few clusters  
different size



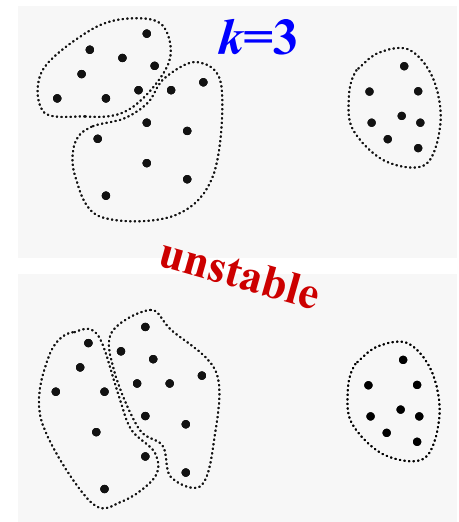
Too many clusters  
wrong model



Too many clusters  
different density

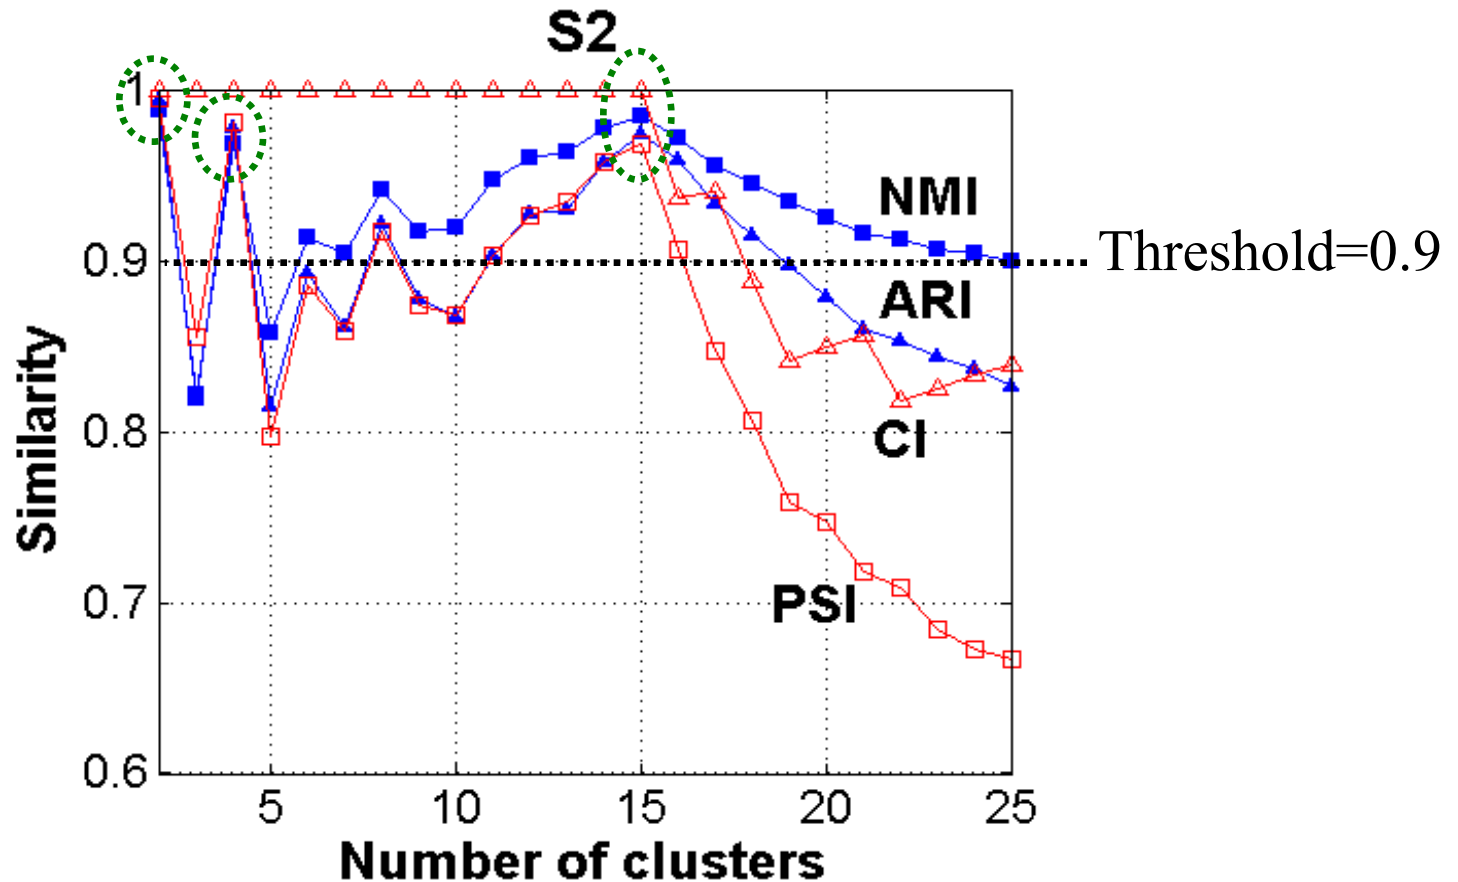


Too many clusters  
different size



# Solution

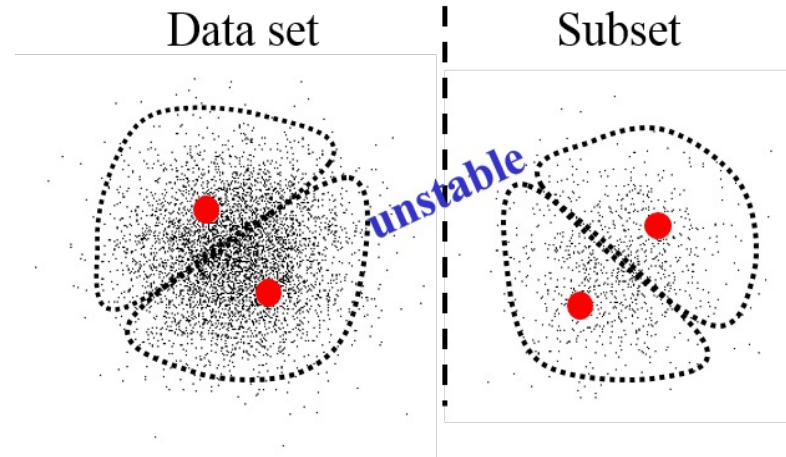
Instead of selecting  $k$  with maximum stability, select last  $k$  with stable result.



# Effect of cluster shapes

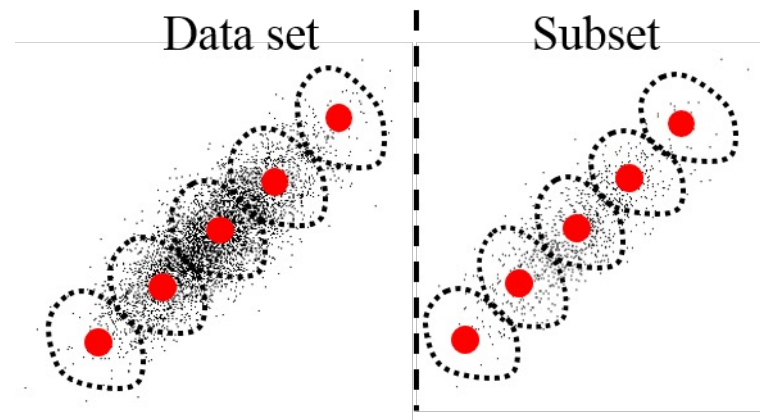
## Correct model:

- Works ok.



## Wrong model:

- Elliptical cluster  
Minimizing TSE  
would find 5  
spherical clusters



# Clustering takeaways

- The goal of clustering is to discover or draw out **similarities among subsets of your data**.
- In a good clustering, points in the **same cluster** should be more similar (nearer) to each other than they are to points in **other clusters**.
- When clustering, the units that each variable is **measured in matter**. Different units cause different distances and potentially different clusterings.
- Clustering is often used for data exploration or as a precursor to **supervised** learning methods.
- Different clustering algorithms might give different results. You should consider different approaches, with different numbers of clusters.
- NOT to forget probabilistic vs non-probabilistic algorithms.