# Introduction to cluster analysis and classification: Formalizing clustering 

Christophe Biernacki

## - To cite this version:

Christophe Biernacki. Introduction to cluster analysis and classification: Formalizing clustering. Summer School on Clustering, Data Analysis and Visualization of Complex Data, May 2018, Catania, Italy. hal-01810379

HAL Id: hal-01810379<br>https://inria.hal.science/hal-01810379

Submitted on 7 Jun 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

# Introduction to cluster analysis and classification: Formalizing clustering 

C. Biernacki<br>Summer School on Clustering, Data Analysis and Visualization of Complex Data May 21-25 2018, University of Catania, Italy

de Lille


Cnran

## Outline

1 Need to formalize

2 Formalizing estimation

3 Formalizing selection

4 More advanced formalizing

5 MixtComp in MASSICCC

6 To go further

## Clustering: an ill-posed problem

■ We have seen how to perform and evaluate clustering. . .
■ . . . but we do not know what is a cluster
■ Thus we have built something without defining it!

It is a serious problem
Impossible to provide guarantees on by-products of clustering (ex.: some user decisions) since no guarantees on clustering itself is really available

$$
\boldsymbol{x} \longrightarrow \hat{\boldsymbol{z}}=f(\boldsymbol{x}) \longrightarrow \underbrace{\text { decision }=g(\hat{\boldsymbol{z}})}_{\text {need guarantees }}
$$

## Expected guarantees on clustering

$■ \hat{\boldsymbol{z}}, \hat{\boldsymbol{\mu}}, \hat{K}$, etc. are estimates of theoretical quantities $\boldsymbol{z}, \boldsymbol{\mu}, K$, etc.
■ It can be thus expected classical guarantees provided for any estimate in statistics
■ Typically: consistency, bias, variance

- Examples:

$$
\begin{gathered}
\mathrm{p}(\hat{K}=K)=1 \text { as } n \rightarrow \infty \\
\mathrm{p}(\hat{z}=\boldsymbol{z}) \text { for finite } n
\end{gathered}
$$

- In the previous lessons we were very far from such a requirement. . .


## Key idea

Formalize the rigorous definition of a cluster

## The model-based clustering paradigm

$$
\text { a cluster } \Longleftrightarrow \text { a distribution }
$$

- It recasts all previous/next questions into model design/estimation/selection
- It takes benefits from all theoretical statistics environment
- How to choose the best metric $\mathbf{M}_{(k)}$ ?

■ How to choose the number $K$ of clusters?
■ Clusters of different sizes are they well estimated?

- How to choose the data unit?

■ How to select features?

- How to deal with mixed data?

■ How to deal with missing data?
■ How to deal with outliers?
■...

## What about empirical clustering?

Somewhere it works pretty well even if it has the previous mentioned limits
11 Interesting to understand why
[2 Interesting to overcome their limits then

In fact many empirical methods are hidden model-based clustering ones!

## Reformulate K-means: the hidden Gaussian assumption

$$
W_{\mathbf{1}}(\mathbf{z})=\sum_{i=1}^{n} \sum_{k=1}^{K} z_{i k}\left\|\boldsymbol{x}_{i}-\boldsymbol{\mu}_{k}\right\|_{1}^{2}
$$

$=-2 \sum_{i=1}^{n} \sum_{k=1}^{K} z_{i k} \ln [\underbrace{\frac{1}{K}}_{!} \underbrace{\frac{1}{(2 \pi)^{d / 2}|\mathbf{I}|^{1 / 2}} \exp \left(-\frac{1}{2}\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{k}\right)^{\prime} \mathbf{l}\left(\boldsymbol{x}_{i}-\boldsymbol{\mu}_{k}\right)\right)}_{N_{d}\left(\boldsymbol{\mu}_{k}, \mathbf{l}\right)}]+\mathrm{cst}$

## Model

$d$-variate Gaussian with variance matrix I and same cluster sample size (see later)

## Reformulate $K$-means: the hidden estimate choice

$$
\begin{gathered}
W_{\mathbf{l}}(\boldsymbol{z})=-2 \ell_{c}(\boldsymbol{\mu} ; \boldsymbol{x}, \boldsymbol{z})+\mathrm{cst} \\
\Downarrow \\
\left.\left(\hat{\mathbf{z}}^{\ell_{c}}, \hat{\boldsymbol{\mu}}^{\ell_{c}}\right)=\begin{array}{c}
\hat{\mathbf{z}}^{W}=\arg \max _{(\boldsymbol{z}, \boldsymbol{\mu})} \ell_{c}(\boldsymbol{\mu} ; \boldsymbol{x}, \boldsymbol{z})
\end{array}\right\} \Rightarrow \hat{\mathbf{z}}^{W} \equiv \hat{\mathbf{z}}^{\ell_{c}}
\end{gathered}
$$

## Estimate

Maximum of the so-called complete-likelihood (see later for its statistical properties)

## Outline

1 Need to formalize

2 Formalizing estimation

3 Formalizing selection

4 More advanced formalizing

5 MixtComp in MASSICCC

6 To go further

## Parametric mixture model

- Parametric assumption:

$$
\mathrm{p}_{k}\left(\mathbf{x}_{1}\right)=\mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\alpha}_{k}\right)
$$

thus

$$
\mathrm{p}\left(\mathbf{x}_{1}\right)=\mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\theta}\right)=\sum_{k=1}^{K} \pi_{k} \mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\alpha}_{k}\right)
$$

- Mixture parameter:

$$
\boldsymbol{\theta}=(\boldsymbol{\pi}, \boldsymbol{\alpha}) \text { with } \boldsymbol{\alpha}=\left(\boldsymbol{\alpha}_{1}, \ldots, \boldsymbol{\alpha}_{K}\right)
$$

- Model: it includes both the family $\mathrm{p}\left(\cdot ; \boldsymbol{\alpha}_{k}\right)$ and the number of groups $K$

$$
\mathbf{m}=\left\{p\left(\mathbf{x}_{1} ; \boldsymbol{\theta}\right): \boldsymbol{\theta} \in \Theta\right\}
$$

The number of free continuous parameters is given by

$$
\nu=\operatorname{dim}(\Theta)
$$

## The clustering process in mixtures

1 Estimation of $\boldsymbol{\theta}$ by $\hat{\boldsymbol{\theta}}$
$\boxed{2}$ Estimation of the conditional probability that $\mathbf{x}_{i} \in G_{k}$

$$
t_{i k}(\hat{\boldsymbol{\theta}})=p\left(Z_{i k}=1 \mid \mathbf{X}_{i}=\mathbf{x}_{i} ; \hat{\boldsymbol{\theta}}\right)=\frac{\hat{\pi}_{k} p\left(\mathbf{x}_{i} ; \hat{\boldsymbol{\alpha}}_{k}\right)}{p\left(\mathbf{x}_{i} ; \hat{\boldsymbol{\theta}}\right)}
$$

3 Estimation of $\mathbf{z}_{i}$ by maximum a posteriori (MAP)

$$
\hat{z}_{i k}=\mathbb{I}_{\left\{k=\arg \max _{h=1, \ldots, K} t_{i h}(\hat{\boldsymbol{\theta}})\right\}}
$$





## Geometric interpretation of $\boldsymbol{\Sigma}_{k}$

$$
\boldsymbol{\Sigma}_{k}=\underbrace{\lambda_{k}}_{\text {volume }} \cdot \underbrace{\mathbf{D}_{k}}_{\text {orientation }} \cdot \underbrace{\mathbf{A}_{k}}_{\text {shape }} \cdot \mathbf{D}_{k}^{\prime}
$$



## Estimation of $\boldsymbol{\theta}$ by complete-likelihood

Maximize the complete-likelihood over $(\boldsymbol{\theta}, \mathbf{z})$

$$
\ell_{c}(\boldsymbol{\theta} ; \mathbf{x}, \mathbf{z})=\sum_{i=1}^{n} \sum_{k=1}^{k} z_{i k} \ln \left\{\pi_{k} p\left(\mathbf{x}_{i} ; \boldsymbol{\alpha}_{k}\right)\right\}
$$

- Equivalent to traditional methods

| Metric | $\mathbf{M}=\mathbf{I}$ | $\mathbf{M}$ free | $\mathbf{M}_{k}$ free |
| :--- | :--- | :--- | :--- |
| Gaussian model | $[\pi \lambda I]$ | $[\pi \lambda C]$ | $\left[\pi \lambda_{k} C_{k}\right]$ |

- Bias of $\hat{\boldsymbol{\theta}}$ : heavy if poor separated clusters
- Associated optimization algorithm: CEM (see later)
- CEM with $[\pi \lambda I]$ is strictly equivalent to $K$-means

■ CEM is simple et fast (convergence with few iterations)

## Estimation of $\boldsymbol{\theta}$ by observe-likelihood

Maximize the observe-likelihood on $\theta$

$$
\ell(\boldsymbol{\theta} ; \mathbf{x})=\sum_{i=1}^{n} \ln p\left(\mathbf{x}_{i} ; \boldsymbol{\theta}\right)
$$

- Convergence of $\hat{\boldsymbol{\theta}}$, asymptotic efficiency, asymptotically unbiased
- General algorithm for missing data: EM
- EM is simple but slower than CEM
- Interpretation: it is a kind of fuzzy clustering


## Principle of EM and CEM

- Initialization: $\boldsymbol{\theta}^{0}$
- Iteration $\mathrm{n}^{\circ} q$ :

■ Step E: estimate probabilities $\mathbf{t}^{q}=\left\{t_{i k}\left(\boldsymbol{\theta}^{q}\right)\right\}$

- Step C: classify by setting $\mathbf{t}^{q}=\operatorname{MAP}\left(\left\{t_{i k}\left(\boldsymbol{\theta}^{q}\right)\right\}\right)$

■ Step M: maximize $\boldsymbol{\theta}^{q+1}=\arg \max _{\boldsymbol{\theta}} \ell_{c}\left(\boldsymbol{\theta} ; \mathbf{x}, \mathbf{t}^{q}\right)$
■ Stopping rule: iteration number or criterion stability

## Properties

- $\oplus$ : simplicity, monotony, low memory requirement
- $\ominus$ : local maxima (depends on $\boldsymbol{\theta}^{0}$ ), linear convergence (EM)


## Gaussian M-step

$$
\begin{aligned}
\pi_{k}^{(q+1)} & =\frac{n_{k}^{(q)}}{n} \\
\boldsymbol{\mu}_{k}^{(q+1)} & =\frac{1}{n_{k}^{(q)}}\left(\sum_{i=1}^{n} t_{i k}\left(\boldsymbol{\theta}^{(q)}\right) \mathbf{x}_{i}\right) \\
\boldsymbol{\Sigma}_{k}^{(q+1)} & =\frac{1}{n_{k}^{(q)}}\left(\sum_{i=1}^{n} t_{i k}\left(\boldsymbol{\theta}^{(q)}\right)\left(\mathbf{x}_{i}-\boldsymbol{\mu}_{k}^{(q+1)}\right)\left(\mathbf{x}_{i}-\boldsymbol{\mu}_{k}^{(q+1)}\right)^{\prime}\right)
\end{aligned}
$$

## Example of an EM in the univariate case











Note : low at the beginning but increase of the log-likelihood

## Local maxima




## Comparison EM-CEM in practice

| $n$ | overlap $^{1}$ |  | KL $(\boldsymbol{\theta}, \hat{\boldsymbol{\theta}})$ |  | $\operatorname{err}(\boldsymbol{z}, \hat{\mathbf{z}})$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  |  | EM | CEM | EM | CEM |  |
| 20 | low | 0.2770 | 0.2771 | 0.3383 | 0.3217 |  |
|  | middle | 0.4916 | 0.3699 | 0.2050 | 0.1700 |  |
|  | high | 0.4108 | 0.3132 | 0.0983 | 0.0667 |  |
| 200 | low | 0.0209 | 0.0822 | 0.3342 | 0.3188 |  |
|  | middle | 0.0187 | 0.0425 | 0.1638 | 0.1587 |  |
|  | high | 0.0172 | 0.0209 | 0.0530 | 0.0500 |  |
| 2000 | low | 0.0014 | 0.0454 | 0.3112 | 0.3113 |  |
|  | middle | 0.0017 | 0.0246 | 0.1620 | 0.1619 |  |
|  | high | 0.0017 | 0.0059 | 0.0509 | 0.0510 |  |

[^0]
## Categorical variables: latent class model

- Categorical variables: $d$ variables with $m_{j}$ modalities each, $x_{i}^{j} \in\{0,1\}^{m_{j}}$ and

$$
\mathbf{x}_{i}^{j h}=1 \quad \Leftrightarrow \quad \text { variable } j \text { of } \mathbf{x}_{i} \text { takes level } h
$$

■ Conditional independence:

$$
\mathrm{p}\left(\boldsymbol{x}_{i} ; \boldsymbol{\alpha}_{k}\right)=\prod_{j=1}^{d} \prod_{h=1}^{m_{j}}\left(\alpha_{k}^{j h}\right)^{x_{i}^{j h}}
$$

and

$$
\alpha_{k}^{j h}=\mathrm{p}\left(x_{i}^{j h}=1 \mid z_{i k}=1\right)
$$

with $\boldsymbol{\alpha}_{k}=\left(\alpha_{k}^{j h} ; j=1, \ldots, d ; h=1, \ldots, m_{j}\right)$

## Integer: Poisson mixture model

- integer variables: $d$ variables $\mathbf{x}_{i}^{j} \in \mathbb{N}$
- Intra conditional independence:

$$
\mathrm{p}\left(\boldsymbol{x}_{i}^{i n t} ; \boldsymbol{\alpha}_{k}^{i n t}\right)=\prod_{j=1}^{d} \frac{\left(\alpha_{k}^{j}\right)^{x_{i}^{j}}}{\alpha_{k}^{j}!} e^{-\alpha_{k}^{j}}
$$

## SPAM E-mail Database ${ }^{3}$

■ $n=4601$ e-mails composed by 1813 "spams" and 2788 "good e-mails"
$\square d=48+6=54$ continuous descriptors ${ }^{2}$
■ 48 percentages that a given word appears in an e-mail ("make", "you'...)
■ 6 percentages that a given char appears in an e-mail (";", "\$"...)

- Transformation of continuous descriptors into binary descriptors

$$
x_{i}^{j}= \begin{cases}1 & \text { if word/char } j \text { appears in e-mail } i \\ 0 & \text { otherwise }\end{cases}
$$

[^1]
## An EM run with a binary data set

## Initial binary data



## An EM run with a binary data set



## An EM run with a binary data set

## Iteration 2



## An EM run with a binary data set



## An EM run with a binary data set

Iteration 4


## An EM run with a binary data set



## An EM run with a binary data set



## An EM run with a binary data set



## An EM run with a binary data set



## An EM run with a binary data set



## An EM run with a binary data set

## Iteration 10



## An EM run with a binary data set

## Iteration 11



## An EM run with a binary data set

## Iteration 12



## An EM run with a binary data set

Iteration 13


## An EM run with a binary data set

## Iteration 14



## An EM run with a binary data set

## Iteration 15



## An EM run with a binary data set

Iteration 16


## An EM run with a binary data set

Iteration 17


## An EM run with a binary data set

Iteration 18


## An EM run with a binary data set

Iteration 19


## An EM run with a binary data set

Iteration 20


## An EM run with a binary data set

Final summary


## Mixed data: classical approaches

Usually, unify data type by transformation :
■ Quantify continuous variables: loose some information
■ MCA dof categorical variable: loose the meaning
■...

## Proposal

Model-based directly on raw data

## Mixed data: conditional independence everywhere

The aim is to combine continuous, categorical and integer data

$$
\mathbf{x}_{1}=\left(\mathbf{x}_{1}^{\text {cont }}, \mathbf{x}_{1}^{\text {cat }}, \mathbf{x}_{1}^{\text {int }}\right)
$$

The proposed solution is to mixed all types by inter-type conditional independence

$$
\mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\alpha}_{k}\right)=\mathrm{p}\left(\mathrm{x}_{1}^{\text {cont }} ; \boldsymbol{\alpha}_{k}^{\text {cont }}\right) \times \mathrm{p}\left(\mathrm{x}_{1}^{\text {cat }} ; \boldsymbol{\alpha}_{k}^{\text {cat }}\right) \times \mathrm{p}\left(\mathbf{x}_{1}^{\text {int }} ; \boldsymbol{\alpha}_{k}^{\text {int }}\right)
$$

In addition, for symmetry between types, intra-type conditional independence

Only need to define the univariate pdf for each variable type!

■ Continuous: Gaussian

- Categorical: multinomial

■ Integer: Poisson

## Missing data: a seminal paper

Biometrika (1976), 63, 3, pp. 581-92

# Inference and missing data 

By DONALD B. RUBIN

Educational Testing Service, Princeton, New Jersey

## Summary

When making sampling distribution inferences about the parameter of the data, $\theta$, it is appropriate to ignore the process that causes missing data if the missing data are 'missing at random' and the observed data are 'observed at random', but these inferences are generally conditional on the observed pattern of missing data. When making directlikelihood or Bayesian inferences about $\theta$, it is appropriate to ignore the process that causes missing data if the missing data are missing at random and the parameter of the missing data process is 'distinct' from $\theta$. These conditions are the weakest general conditions under which ignoring the process that causes missing data always leads to correct inferences.

Some key words: Bayesian inference; Incomplete data; Likelihood inference; Missing at random; Missing data; Missing values; Observed at random; Sampling distribution inference.

## Missing data: current solutions

| $X_{1}$ | $X_{2}$ | $X_{3}$ | Cluster |
| :---: | :---: | :---: | :---: |
| 1.23 | $?$ | 3.42 | $?$ |
| $?$ | $?$ | 4.10 | $?$ |
| 4.53 | 1.50 | 5.35 | $?$ |
| $?$ | 5.67 | $?$ | $?$ |

## Discarded solutions

- Suppress units and/or variables with missing data $\Rightarrow$ loss of information

■ Imputation of the missing data by the mean or more evolved methods $\Rightarrow$ uncertainty of the prediction not taken into account

## Retained solution

Use an integrated approach which allows to take into account all the available information to perform clustering

## Missing data: notations and MNAR assumption

■ $O_{i} \subseteq\{1, \ldots, d\}$ the set of the observed variables from sample $i$

- $\boldsymbol{x}_{i}^{O}$ the observed data from sample $i$
- $M_{i}$ the set of the missing variables for sample $i$
- $\boldsymbol{\mu}_{i k}^{O}$ the sub-vector of $\boldsymbol{\mu}_{k}$ associated to index $O_{i}$ (the same for $M_{i}$ )

■ $\boldsymbol{\Sigma}_{i k}^{\mathrm{OM}}$ the sub-matrix of $\boldsymbol{\Sigma}_{k}$ associated to row $O_{i}$ and columns $M_{i}$ (the same for any other combination)

## Assumption on the missingness mecanism

Missing At Randon (MAR): the probability that a variable is missing does not depend on its own value given the observed variables.

## Missing data: maximum likelihood estimator

Observed log-likelihood...

$$
\ell\left(\boldsymbol{\theta} ; \mathbf{x}^{\mathrm{O}}\right)=\sum_{i=1}^{n} \log \left(\sum_{k=1}^{K} \pi_{k} \mathrm{p}\left(\boldsymbol{x}_{i}^{\mathrm{O}} ; \boldsymbol{\alpha}_{k}\right)\right)=\ln [\sum_{k=1}^{K} \pi_{k} \underbrace{\int_{\mathbf{x}_{i}^{M}} \mathrm{p}\left(\mathbf{x}_{i}^{O}, \mathbf{x}_{i}^{M} ; \boldsymbol{\alpha}_{k}\right) d \mathbf{x}_{i}^{M}}_{\text {MAR assumption }}]
$$

## EM and Gaussian case: E step

$\boldsymbol{\theta}$ and $\boldsymbol{\theta}^{+}$the parameters for two successive steps (idem for missing data)

$$
\begin{aligned}
z_{i k}^{+} & =P\left(Z_{i k}=1 \mid \boldsymbol{x}_{i}^{\mathrm{O}} ; \boldsymbol{\theta}\right)=\frac{\pi_{k} \phi\left(\boldsymbol{x}_{i}^{\mathrm{O}} ; \boldsymbol{\Sigma}_{k}\right)}{\sum_{\ell=1}^{K} \pi_{\ell} \phi\left(\boldsymbol{x}_{i}^{\mathrm{O}} ; \boldsymbol{\Sigma}_{\ell}\right)} \\
\boldsymbol{x}_{i k}^{\mathrm{M}^{+}} & =E\left[\boldsymbol{X}_{i}^{M} \mid \boldsymbol{x}_{i}^{\mathrm{O}}, Z_{i k}=1 ; \boldsymbol{\theta}\right]=\boldsymbol{\mu}_{i k}^{\mathrm{M}}+\boldsymbol{\Sigma}_{i k}^{\mathrm{MO}}\left(\boldsymbol{\Sigma}_{i k}^{\mathrm{OO}}\right)^{-1}\left(\boldsymbol{x}_{i}^{\mathrm{O}}-\boldsymbol{\mu}_{i k}^{\mathrm{O}}\right) .
\end{aligned}
$$

## Interpretation

- $z_{i k}^{+}$: class posterior probability membership given the available information $\boldsymbol{x}_{i}^{\mathrm{O}}$.
- $x_{i k}^{\mathrm{M}^{+}}$: conditional imputation of the missing data given the cluster.


## EM and Gaussian case: M step

$$
\begin{aligned}
\pi_{k}^{+} & =\frac{1}{n_{k}^{+}} \sum_{i=1}^{n} z_{i k}^{+}, \boldsymbol{\mu}_{k}^{+}=\frac{1}{n_{k}^{+}} \sum_{i=1}^{n} z_{i k}^{+} \boldsymbol{x}_{i k}^{+} \\
\boldsymbol{\Sigma}_{k}^{+} & =\frac{1}{n_{k}^{+}} \sum_{i=1}^{n} z_{i k}^{+}\left[\left(\boldsymbol{x}_{i k}^{+}-\boldsymbol{\mu}_{k}^{+}\right)\left(\boldsymbol{x}_{i k}^{+}-\boldsymbol{\mu}_{k}^{+}\right)^{\prime}+\boldsymbol{\Sigma}_{i k}^{+}\right]
\end{aligned}
$$

where $n_{k}^{+}=\sum_{i=1}^{n} z_{i k}^{+}, \boldsymbol{x}_{i k}^{+}=\binom{\boldsymbol{x}_{i}^{O}}{\boldsymbol{x}_{i k}^{\mathrm{M}}}, \boldsymbol{\Sigma}_{i k}^{+}=\left(\begin{array}{cc}0_{i}^{O} & 0_{i}^{\mathrm{OM}} \\ 0_{i}^{\mathrm{MO}} & \boldsymbol{\Sigma}_{i k}^{\mathrm{M}^{+}}\end{array}\right)$with 0 the $d \times d$ null matrix, and $\boldsymbol{\Sigma}_{i k}^{\mathrm{M}^{+}}=\boldsymbol{\Sigma}_{i k}^{\mathrm{MO}}\left(\boldsymbol{\Sigma}_{i k}^{\mathrm{O}}\right)^{-1} \boldsymbol{\Sigma}_{i k}^{\mathrm{OM}}$.

Interpretation of $\boldsymbol{\Sigma}_{i k}^{\mathrm{M}^{+}}$
Variance correction due to the under-estimation of variability caused by the imputation of missing data.

## Missing data: SEM algorithm

A SEM algorithm to estimate $\boldsymbol{\theta}$ by maximizing the observed-data log-likelihood

- Initialisation: $\boldsymbol{\theta}^{(0)}$
- Iteration nb $q$ :
- E-step: compute conditional probabilities $\mathrm{p}\left(\boldsymbol{x}^{M}, \boldsymbol{z} \mid \mathcal{D} ; \boldsymbol{\theta}^{(q)}\right)$
- S-step: draw ( $\left.\boldsymbol{x}^{M(q)}, \boldsymbol{z}^{(q)}\right)$ from $\mathrm{p}\left(\boldsymbol{x}^{M}, \boldsymbol{z} \mid \boldsymbol{x}^{0} ; \boldsymbol{\theta}^{(q)}\right)$
- M-step: maximize $\boldsymbol{\theta}^{(q+1)}=\arg \max _{\boldsymbol{\theta}} \ln \mathrm{p}\left(\mathrm{x}^{\circ}, \boldsymbol{x}^{M(q)}, \boldsymbol{z}^{(q)} ; \boldsymbol{\theta}\right)$

■ Stopping rule: iteration number

## Properties: simpler than EM and interesting properties!

- Avoid possibly difficult E-step in an EM

■ Classical M steps

- Avoids local maxima
- The mean of the sequence $\left(\boldsymbol{\theta}^{(q)}\right)$ approximates $\hat{\boldsymbol{\theta}}$
- The variance of the sequence $\left(\boldsymbol{\theta}^{(q)}\right)$ gives confidence intervals


## Missing data: SE algorithm

A SE algorithm estimates then $\left(\boldsymbol{x}^{M}, \boldsymbol{z}^{M}\right)$

- Iteration nb $q$ :
- E-step: compute conditional probabilities $\mathrm{p}\left(x^{M}, \boldsymbol{z}^{M} \mid \mathcal{D} ; \hat{\boldsymbol{\theta}}\right)$
- S-step: draw ( $\left.x^{M(q)}, \boldsymbol{z}^{M(q)}\right)$ from $\mathrm{p}\left(x^{M}, \boldsymbol{z}^{M} \mid \mathcal{D} ; \hat{\boldsymbol{\theta}}\right)$

■ Stopping rule: iteration number

## Properties

■ simplicity because of conditional independence
$\square$ the mean/mode of the sequence $\left(\boldsymbol{x}^{M(q)}, \boldsymbol{z}^{M(q)}\right)$ estimates $\left(\boldsymbol{x}^{M}, \boldsymbol{z}^{M}\right)$

- confidence intervals are also derived


## Missing data: illustration with the cancer data set (1/2)

■ Strategy "mice ${ }^{4}+$ mixture": mixture on the dataset completed by mice
> data.imp=mice(data)
> data. comp.mice=complete(data.imp)

■ Strategy "full mixture": mixture on the observed (no completed) dataset

[^2]MMissing data: illustration with the cancer data set (2/2)

| Strategy | mice + mixture | full mixture |
| :---: | :---: | :---: |
| \% misclassified | 12.8 | 8.1 |

Avoid to complete missing data (imputation depends on the purpose)

## Outline

1. Need to formalize

2 Formalizing estimation

3 Formalizing selection

4 More advanced formalizing

5 MixtComp in MASSICCC

6 To go further

## Keep in mind

## George E.P. Box (1987)

"Essentially, all models are wrong, but some are useful"

■ $\mathcal{M}=\{\mathbf{m}\}$ will denote the set of competing models

- The true distribution p is not necessarily in $\mathcal{M}$
- Density estimation: AIC, BIC

■ Clustering: ICL, CL, NEC

Importance of model selection: example

Model $=$ number of clusters + parametric structure of clusters

Too simple model: bias


## Too complex model: variance


_ true borderline

-     - = borderline with [ $\pi \lambda /$ ] borderline with $\left[\pi \lambda_{k} C_{k}\right]$


## Importance of model selection: bias/variance trade-off

■ Partition error rate: $\operatorname{err}\left(\mathbf{z}_{1}, \mathbf{z}_{2}\right) \geq 0$ a distance-like between two partitions $\mathbf{z}_{1}, \mathbf{z}_{2}$
■ Gap between true and model partition:

$$
\boldsymbol{\theta}_{\mathbf{m}}^{*}=\arg \min _{\boldsymbol{\theta} \in \Theta_{\mathbf{m}}} \operatorname{err}(\mathbf{z}, \mathbf{z}(\boldsymbol{\theta}))
$$

- MLE:

$$
\hat{\boldsymbol{\theta}}_{\mathbf{m}}=\arg \max _{\boldsymbol{\theta} \in \Theta} \ell(\boldsymbol{\theta} ; \mathcal{D})
$$

- Fundamental decomposition of $\operatorname{err}\left(\mathbf{z}, \mathbf{z}\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}}\right)\right)$ :

$$
\begin{aligned}
& \operatorname{err}\left(\mathbf{z}, \mathbf{z}\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}}\right)\right) \\
& \quad=\left\{\operatorname{err}\left(\mathbf{z}, \mathbf{z}\left(\boldsymbol{\theta}_{\mathbf{m}}^{*}\right)\right)-\operatorname{err}(\mathbf{z}, \mathbf{z})\right\}+\left\{\operatorname{err}\left(\mathbf{z}, \mathbf{z}\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}}\right)\right)-\operatorname{err}\left(\mathbf{z}, \mathbf{z}\left(\boldsymbol{\theta}_{\mathbf{m}}^{*}\right)\right)\right\} \\
& \quad=\left\{\operatorname{bias}_{\mathbf{m}}\right\}+\left\{\text { variance }_{\mathbf{m}}\right\}
\end{aligned}
$$

Importance of model selection: illustration of the variance effect

30 samples from a bivariate mixture with two components

$$
\begin{gathered}
\pi_{1}=\pi_{2}=0.5, \quad \boldsymbol{\mu}_{1}=(0,0)^{\prime}, \quad \boldsymbol{\mu}_{2}=(2,2)^{\prime}, \quad \boldsymbol{\Sigma}_{1}=\boldsymbol{\Sigma}_{2}=\mathbf{I} \\
\mathcal{M}=\{\text { spherical, general }\} \\
\begin{array}{ccc}
\hline n & \mathbf{m} & \operatorname{err}\left(\mathbf{z}, \hat{\mathbf{z}}_{\mathbf{m}}\right) \\
\hline 40 & \text { spherical } & 0.0967 \\
& \text { general } & 0.1100 \\
\hline 200 & \text { spherical } & 0.0840 \\
& \text { general } & 0.0872 \\
\hline
\end{array}
\end{gathered}
$$

## Some heuristics entropy-based criteria: examples

■ A fundamental decomposition of $\ell(\boldsymbol{\theta} ; \mathbf{x})$ : for any "fuzzy partition" $\mathbf{c}=\left\{c_{i k}\right\}$

$$
\begin{aligned}
\ell(\boldsymbol{\theta} ; \mathbf{x}) & =\sum_{i=1}^{n} \sum_{k=1}^{K} c_{i k} \ln \left\{\pi_{k} \mathrm{p}\left(\mathbf{x}_{i} ; \boldsymbol{\alpha}_{k}\right)\right\}-\sum_{i=1}^{n} \sum_{k=1}^{K} c_{i k} \ln t_{i k}(\boldsymbol{\theta}) \\
& =\ell_{c}(\boldsymbol{\theta} ; \mathbf{x}, \mathbf{c})+\xi(\boldsymbol{\theta} ; \mathbf{c}) \\
& =\text { complete-data log-likelihood + entropy }
\end{aligned}
$$

- NEC criterion (Normalized Entropy Criterion): retain m minimizing

$$
\mathrm{NEC}_{K}= \begin{cases}\frac{\xi\left(\hat{\boldsymbol{\theta}}_{K} ; \mathbf{t}\left(\hat{\boldsymbol{\theta}}_{K}\right)\right)}{\ell\left(\hat{\boldsymbol{\theta}}_{K} ; \mathbf{x}\right)-\ell\left(\hat{\boldsymbol{\theta}}_{1} ; \mathbf{x}\right)} & \text { if } K>1 \\ 1 & \text { if } K=1\end{cases}
$$

■ CL criterion (Completed Likelihood): retain maximizing

$$
\mathrm{CL}=\ell_{c}(\hat{\boldsymbol{\theta}} ; \mathbf{x}, \hat{\mathbf{z}})=\underbrace{\ell(\hat{\boldsymbol{\theta}} ; \mathbf{x})}_{\text {model adequacy }}-\underbrace{\xi(\hat{\boldsymbol{\theta}} ; \hat{\mathbf{z}})}_{\text {partition evidence }}
$$

■ Behaviour: not completely satisfactory but something happens...

Some heuristics entropy-based criteria: NEC illustration



## Some heuristics entropy-based criteria: CL illustration

Interpretation as a penalized within-cluster inertia criterion:

$$
\mathrm{CL}_{([p \lambda l], K)}=-\frac{n d}{2} \ln \left(W_{K}\right)-n \ln (K)+\mathrm{cst}
$$





## Theoretical model selection criteria

## The most widespread principle



| criterion | penalty | interpretation | user purpose |
| :---: | :---: | :---: | :---: |


| general criteria in statistics |  |  |  |
| :---: | :---: | :---: | :---: |
| AIC | $\nu$ | model complexity | prediction |
| BIC | $0.5 \nu \ln (n)$ | model complexity | identification |

specific criterion for the clustering aim

| ICL | $0.5 \nu \ln (n)$ | model complexity | well-separated |
| :---: | :---: | :---: | :---: |
|  | $-\sum_{i, k} \hat{z}_{i k} \ln t_{i k}(\hat{\boldsymbol{\theta}})$ | + partition entropy | clusters |

## BIC criterion: integrated likelihood

■ Posterior likelihood of $\mathbf{m}$ :

$$
\mathrm{p}(\mathbf{m} \mid \mathcal{D}) \propto \mathrm{p}(\mathcal{D} \mid \mathbf{m}) \underbrace{\mathrm{p}(\mathbf{m})}_{\text {prior on } \mathbf{m}}
$$

■ Ideal model in a Bayesian context:

$$
\hat{\mathbf{m}}^{*} \in \arg \max _{\mathbf{m} \in \mathcal{M}} \mathrm{p}(\mathbf{m} \mid \mathcal{D})
$$

■ Integrated likelihood: if $\mathbf{p ( m )}=\mathrm{cst}$, it is equivalent to maximize

$$
\mathrm{p}(\mathcal{D} \mid \mathbf{m})=\int_{\Theta} \mathrm{p}(\mathcal{D} ; \boldsymbol{\theta}, \mathbf{m}) \underbrace{\mathrm{p}(\boldsymbol{\theta} \mid \mathbf{m})}_{\text {prior on } \theta} d \boldsymbol{\theta}
$$

- Difficulties:

■ Choose the prior $\mathrm{p}(\boldsymbol{\theta} \mid \mathbf{m})$

- Evaluate the integral


## BIC criterion: genesis

■ Laplace-Metropolis approximation: under standard regularity conditions, we have

$$
\ln \mathrm{p}(\mathcal{D} \mid \mathbf{m})=\ell(\hat{\boldsymbol{\theta}} ; \mathcal{D})-\frac{\nu}{2} \ln (n)+O_{p}(1)
$$

■ BIC criterion (Bayesian Information Criterion): retain maximizing

$$
\mathrm{BIC}_{\mathrm{m}}=\ell\left(\hat{\theta_{\mathrm{m}}} ; \mathcal{D}\right)-\frac{\nu_{\mathrm{m}}}{2} \ln (n)
$$

## BIC criterion: consistency

- Consistency: BIC asymptotically selects the best

$$
\mathbf{m}^{*}=\arg \inf _{\mathbf{m} \in \mathcal{M}} \mathrm{KL}\left(\mathbf{p}, \mathrm{p}_{\boldsymbol{\theta}_{\mathbf{m}}^{*}}\right)
$$

- Theoretical illustration of consistency: $\mathbf{m}_{1} \subseteq \mathbf{m}_{2}, \mathbf{m}_{1}$ being the true model, $\Delta \nu=\nu_{2}-\nu_{1}, \Delta \ell=\ell\left(\hat{\boldsymbol{\theta}}_{2} ; \mathcal{D}\right)-\ell\left(\hat{\boldsymbol{\theta}}_{1} ; \mathcal{D}\right)$, we have

$$
2\left(\mathrm{BIC}_{2}-\mathrm{BIC}_{1}\right)+\Delta \nu \ln (n)=2 \Delta \ell \xrightarrow{d} \chi_{\Delta \nu}^{2}
$$

With $\mu=\Delta \nu$ and $\sigma^{2}=2 \Delta \nu$ the mean and the variance of $\chi_{\Delta \nu}^{2}$

$$
\mathrm{p}\left(\chi_{\Delta \nu}^{2}>\Delta \nu \ln (n)\right) \leq \mathrm{p}\left(\left|\chi_{\Delta \nu}^{2}-\mu\right|>\Delta \nu \ln (n)-\mu\right) \leq \frac{\sigma^{2}}{(\Delta \nu \ln (n)-\mu)^{2}} \xrightarrow{n \rightarrow \infty} 0
$$

by using the Chebyschev inequality. Thus, asymptotically, BIC will select $\mathbf{m}_{1}$

- Special case of $K$ : be careful on the $\chi^{2}$ approximation validity...


## ICL criterion: genesis

■ Revisiting the fundamental decomposition: if $\boldsymbol{z}$ known, retain $\mathbf{m}$ maximizing

$$
\underbrace{\ln \mathrm{p}(\mathbf{x}, \mathbf{z} \mid \mathbf{m})}_{\text {all data evidence }}=\underbrace{\ln \mathrm{p}(\mathbf{x} \mid \mathbf{m})}_{\text {data } \mathbf{x} \text { evidence }}+\underbrace{\ln \mathrm{p}(\mathbf{z} \mid \mathbf{x}, \mathbf{m})}_{\text {partition } \mathbf{z} \text { evidence }}
$$

Thus models leading to overlapping groups are more penalized (low z evidence)
■ ICL criterion (Integrated Classification Likelihood): replace z by $\hat{\mathbf{z}}$

$$
\mathrm{ICL}=\ln \mathrm{p}(\mathrm{x}, \hat{\mathbf{z}} \mid \mathbf{m})
$$

- BIC-like approximation of ICL:

$$
\ln \mathrm{p}(\mathrm{x}, \mathbf{z} \mid \mathbf{m})=\ln \mathrm{p}\left(\mathbf{x}, \mathbf{z} ; \hat{\boldsymbol{\theta}}_{\mathbf{x}, \mathbf{z}} \mid \mathbf{m}\right)-\frac{\nu}{2} \ln n+O_{\mathrm{p}}(1)
$$

In case of the right model $\mathbf{m}: \hat{\boldsymbol{\theta}}_{\mathbf{x}, \mathbf{z}} \xrightarrow{\text { a.s. }} \boldsymbol{\theta}^{*}$ and $\hat{\boldsymbol{\theta}}_{\mathbf{x}} \xrightarrow{\text { a.s. }} \boldsymbol{\theta}^{*}$. Thus, for $n$ large enough, $\hat{\boldsymbol{\theta}}_{\mathbf{x}, \mathbf{z}} \approx \hat{\boldsymbol{\theta}}_{\mathbf{x}}$. Then, we take $\hat{\mathbf{z}}=\operatorname{MAP}\left(\hat{\boldsymbol{\theta}}_{\mathbf{x}}\right)$ (or also $\hat{\mathbf{z}}=\mathbf{t}\left(\hat{\boldsymbol{\theta}}_{\mathbf{x}}\right)$ ). It gives

$$
\begin{aligned}
\text { ICLbic } & =\ln \mathrm{p}\left(\mathbf{x}, \hat{\mathbf{z}} ; \hat{\boldsymbol{\theta}}_{\mathbf{x}}\right)-\frac{\nu}{2} \ln n \\
& =\mathrm{BIC}-\xi\left(\hat{\boldsymbol{\theta}}_{\mathbf{x}} ; \hat{\mathbf{z}}\right) \\
& =\mathrm{CL}-\frac{\nu}{2} \ln n
\end{aligned}
$$

## ICL criterion: robustness to model misspecification

- A bivariate mixture of a uniform and a Gaussian cluster:

■ non-Gaussian component: $\pi_{1}=0.5, \mathrm{p}_{1}\left(\mathbf{x}_{1}\right)=0.25 \mathbf{I}_{[-1,1]}\left(x^{1}\right) \mathbf{I}_{[-1,1]}\left(x^{2}\right)$
■ Gaussian component: $\pi_{2}=0.5, \boldsymbol{\mu}_{2}=(3.3,0)^{\prime}, \boldsymbol{\Sigma}_{2}=\mathbf{I}$
■ 50 simulated data sets of size $n=200$


| $K$ | 1 | 2 | 3 | 4 | 5 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| BIC | $\cdot$ | $\mathbf{6 0}$ | $\cdot$ | 32 | 8 |
| ICLbic | $\cdot$ | $\mathbf{1 0 0}$ | . | . | . |

## ICL criterion: consistency?

- Assumption: true model with two groups and parameter $\boldsymbol{\theta}_{2}^{*}$

■ Theoretical result:

- Preliminaries: $\delta_{n}=n\left(\boldsymbol{\theta}_{2}^{*}-\boldsymbol{\theta}_{2}^{* p}\right)^{\prime} \mathbf{J}\left(\boldsymbol{\theta}_{2}^{*}\right)\left(\boldsymbol{\theta}_{2}^{*}-\boldsymbol{\theta}_{\boldsymbol{2}}^{* p}\right), \mathbf{J}\left(\boldsymbol{\theta}_{2}^{*}\right)$ the Fisher matrix for a data unit calculated with the true parameter $\boldsymbol{\theta}_{2}$ and $\boldsymbol{\theta}_{2}^{* \rho}$ its projected value on the parameter subspace associated to the one component case, $\mu_{n}=\mathrm{E}\left[\chi_{\Delta \nu}^{2}\left(\delta_{n}\right)\right]=\Delta \nu+\delta_{n}$, $\sigma_{n}^{2}=\operatorname{Var}\left[\chi_{\Delta \nu}^{2}\left(\delta_{n}\right)\right]=2\left(\Delta \nu+\delta_{n}\right)$
- Asymptotically: by Chebishev inequality, with $\mu_{n}-\Delta \nu \ln n-2 n \ln 2>0$

$$
\mathrm{p}(\text { choose wrong model })=\mathrm{p}\left(\text { ICLbic }_{2}<\mathrm{ICLbic}_{1}\right) \leq \frac{\sigma_{n}^{2}}{\left(\mu_{n}-\Delta \nu \ln n-2 n \ln 2\right)^{2}}
$$

Thus it goes towards 0 for well-separated groups

- Experimental result: 100 samples from a univariate Gaussian mixture

$$
\pi_{1}=\pi_{2}, \quad \mu_{1}=0, \quad \mu_{2}=\Delta \mu, \quad \sigma_{1}^{2}=\sigma_{2}^{2}=1
$$

| $\Delta \mu$ | 2.9 |  | 3.0 |  | 3.1 |  | 3.2 |  | 3.3 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $n$ | BIC | ICL | BIC | ICL | BIC | ICL | BIC | ICL | BIC | ICL |
| 100 | 94 | 23 | 96 | 31 | 97 | 44 | 95 | 45 | 97 | 60 |
| 400 | 100 | 9 | 100 | 21 | 100 | 48 | 100 | 70 | 100 | 85 |
| 700 | 100 | 8 | 100 | 15 | 100 | 39 | 100 | 72 | 100 | 96 |
| 1000 | 100 | 6 | 100 | 16 | 100 | 56 | 100 | 75 | 100 | 91 |

## Large $n$ : BIC behaviour (1/2)

- The mixture density is wrong (as all models)

■ Mixtures allow to estimate any distribution by increasing the number of components (high flexibility)




## Large $n$ : BIC behaviour (2/2)

Since BIC is consistent, as $n$ grows, it adds components for improving the true density estimation

## Real example

Marketing context :
Looking for classes of customers



It was only $\mathrm{d}=2$ in this example
It was only $\mathrm{n}<10000$ in this example
Reality is even worst: $(\mathrm{n}=10$ customers, $d=77$ mixed, 1 day computer for 20 classes, more than 40 classes!

## Missing data: illustration with the cancer data set (1/2)

■ Strategy "mice ${ }^{5}+$ mixture": mixture on the dataset completed by mice
> data.imp=mice(data)
> data.comp.mice=complete(data.imp)

■ Strategy "full mixture": mixture on the observed (no completed) dataset

[^3]Missing data: illustration with the cancer data set (2/2)

mice + mixture
$\hat{K}-7$

$$
\hat{K}=7
$$


full mixture
$\hat{K}=2$
... may lose some cluster information when imputation before clustering

## CAUTION

Impossible to use BIC/ICL for kernel/spectral clustering (data set has changed. . . )

## Reformulate $K$-means: elbow as a slope heuristics $(1 / 3)$

■ SH (Slope Heuristics) criterion: retain m maximizing

$$
\mathrm{SH}_{\mathbf{m}}=\ell\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}} ; \mathcal{D}\right)-2 \text { variance }_{\mathbf{m}}
$$

■ Estimating the penalty: optimal penalty is linear in $\nu_{\mathbf{m}}$

$$
\text { 2variance }{ }_{\mathbf{m}}=\kappa \nu_{\mathbf{m}} .
$$

and also

$$
2 \text { variance }_{\mathbf{m}}=\underbrace{2\left\{\ell\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}} ; \mathcal{D}\right)-\mathrm{p}(\mathcal{D})\right\}}_{\approx \kappa \nu_{\mathbf{m}}}+\underbrace{2\left\{\mathrm{p}(\mathcal{D})-\ell\left(\boldsymbol{\theta}_{\mathbf{m}} ; \mathcal{D}\right)\right\}}_{\text {bias } \approx \text { cst for too complex models }}
$$

thus, for complex enough models, $\ell\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}} ; \mathcal{D}\right)$ behaves linearly with $\nu_{\mathbf{m}}$ and the corresponding slope is $\kappa / 2$

- Capushe ${ }^{6}$ (CAlibrated Penalty Using Slope HEuristics): $\kappa / 2$ can be estimated by a linear regression of $\ell\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}} ; \mathcal{D}\right)$ on $\frac{\kappa}{2} \nu_{\mathbf{m}}$

[^4]Reformulate $K$-means: elbow as a slope heuristics $(2 / 3)$



Reformulate $K$-means: elbow as a slope heuristics $(2 / 3)$


# Outline 

1. Need to formalize

2 Formalizing estimation

3 Formalizing selection

4 More advanced formalizing

5 MixtComp in MASSICCC

6 To go further

## Outliers: Two possibilities

■ "After": exclude data outside the confidence area of clusters


95\% confidence area

■ "During": model outliers as a particular cluster in the mixture

$$
\mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\theta}\right)=\sum_{k=1}^{K} \pi_{k} \mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\alpha}_{k}\right)+\pi_{0} \mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\alpha}_{0}\right)
$$

## Outliers: "during" example with acoustic emission control

■ Data: $n=2061$ event locations in a rectangle of $\mathbb{R}^{2}$ representing the vessel

- Model: Diagonal Gaussian mixture + uniform (noise)
- Groups: sound locations $=$ vessel defects





## Units: changing the data units

- Principle of data units transformation $\mathbf{u}$ :

$$
\begin{array}{lll}
\mathbf{u}: & \mathcal{X}=\mathcal{X}^{\text {id }} & \longrightarrow \mathcal{X}^{\mathbf{u}} \\
& \mathbf{x}=\mathbf{x}^{\text {id }}=\mathbf{i d}(\mathbf{x}) & \longmapsto
\end{array}
$$

- u is a bijective mapping to preserve the whole data set information quantity
- We denote by $\mathbf{u}^{-1}$ the reciprocal of $\mathbf{u}$, so $\mathbf{u}^{-1} \circ \mathbf{u}=\mathbf{i d}$

■ Thus, id is only a particular unit u

- Often a meaningful restriction ${ }^{7}$ on $\mathbf{u}$ : it proceeds lines by lines and rows by rows

$$
\mathbf{u}(\mathbf{x})=\left(\mathbf{u}\left(\mathbf{x}_{1}\right), \ldots, \mathbf{u}\left(\mathbf{x}_{n}\right)\right) \quad \text { with } \quad \mathbf{u}\left(\mathbf{x}_{i}\right)=\left(\mathbf{u}_{1}\left(x_{i 1}\right), \ldots, \mathbf{u}_{d}\left(x_{i d}\right)\right)
$$

■ Advantage to respect the variable definition, transforming only its unit
$■ \mathbf{u}\left(\mathbf{x}_{i}\right)$ means that $\mathbf{u}$ applied to the data set $\mathbf{x}_{i}$, restricted to the single individual $i$
$■ \mathbf{u}_{j}$ corresponds to the specific (bijective) transformation unit associated to variable $j$

[^5]
## Units: revisiting units as a modelling component

■ Explicitly exhibiting the "canonical" unit id in the model

$$
\mathrm{p}_{\mathbf{m}}=\left\{\cdot \in \mathcal{X} \mapsto \mathrm{p}(\cdot ; \boldsymbol{\theta}): \boldsymbol{\theta} \in \Theta_{\mathbf{m}}\right\}=\left\{\cdot \in \mathcal{X}^{\text {id }} \mapsto \mathrm{p}(\cdot ; \boldsymbol{\theta}): \boldsymbol{\theta} \in \Theta_{\mathbf{m}}\right\}=\mathrm{p}_{\mathbf{m}}^{\text {id }}
$$

- Thus the variable space and the probability measure are embedded
- As the standard probability theory: a couple (variable space, probability measure)!

■ Changing id into $\mathbf{u}$, while preserving $\mathbf{m}$, is expected to produce a new modelling

$$
\mathrm{p}_{\mathbf{m}}^{\mathbf{u}}=\left\{\cdot \in \mathcal{X}^{\mathbf{u}} \mapsto \mathrm{p}(\cdot ; \boldsymbol{\theta}): \boldsymbol{\theta} \in \Theta_{\mathbf{m}}\right\} .
$$

A model should be systematically defined by a couple ( $\mathbf{u}, \mathbf{m}$ ), denoted by $\mathrm{p}_{\mathbf{m}}^{u}$

## Units: interpretation and identifiability of $p_{m}^{u}$

■ Standard probability theory (again): there exists a measure $\mathbf{u}^{-1}(\mathbf{m})$ s.t. ${ }^{8}$

$$
\mathbf{u}^{-1}(\mathbf{m}) \in\left\{\mathbf{m}^{\prime} \in \mathbb{M}: p_{\mathbf{m}^{\prime}}^{\mathbf{i d}}=\mathrm{p}_{\mathbf{m}}^{\mathbf{u}}\right\}
$$

■ There exists two alternative interpretations of strictly the same model:

- $\mathrm{p}_{\mathbf{m}}^{\mathbf{u}}$ : data measured with unit $\mathbf{u}$ arise from measure $\mathbf{m}$;
- $p_{\mathbf{u}^{-1}(\mathbf{m})}^{\text {id }}$ : data measured with unit id arise from measure $\mathbf{u}^{-1}(\mathbf{m})$
- Two points of view:


## Statistician

The model $p_{\mathbf{m}}^{\mathbf{u}}$ is not identifiable over the couple ( $\mathbf{m}, \mathbf{u}$ )

## Practitioner

Freedom to choose the interpretation which is the most meaningful for him

[^6]
## Units: opportunity for designing new models

Great opportunity to build easily numerous new meaningful models $p_{m}^{u}$ !

- Just combine a standard model family $\{\mathbf{m}\}$ with a standard unit family $\{\mathbf{u}\}$

■ New family can be huge! Combinatorial problems can occur...
■ Some model stability can exist in some (specific) cases: $\mathbf{m}=\mathbf{u}^{-1}(\mathbf{m})$

## Units: model selection

As any model, possible to choose between $p_{\mathbf{m}_{1}}^{\mathbf{u}_{1}}$ and $p_{\mathbf{m}_{2}}^{\mathbf{u}_{2}}$

However, caution when using likelihood-based model selection criteria (as BIC)

■ Prohibited to compare $\mathbf{m}_{1}$ in unit $\mathbf{u}_{1}$ and $\mathbf{m}_{2}$ in unit $\mathbf{u}_{2}$

- But allowed after transforming in identical unit id
- Thus compare their equivalent expression: $p_{\mathbf{u}_{1}^{-1}\left(\boldsymbol{m}_{1}\right)}^{\text {id }}$ and $p_{\mathbf{u}_{2}^{-1}\left(\boldsymbol{m}_{2}\right)}^{\text {id }}$

■ Example for abs. continuous $\mathbf{x}$ and differentiable $\mathbf{u}$, the density transform in id is:

$$
\mathrm{p}_{\mathbf{u}^{-1}(\mathbf{m})}^{\mathrm{id}}=\left\{\cdot \in \mathcal{X}^{\mathbf{i d}} \mapsto \mathrm{p}(\mathbf{u}(\cdot) ; \boldsymbol{\theta}) \times\left|\mathbf{J}^{\mathbf{u}}(\cdot)\right|: \boldsymbol{\theta} \in \Theta_{\mathbf{m}}\right\}
$$

with $\mathbf{J u}^{\mathbf{u}}(\cdot)$ the Jacobian associated to the transformation $\mathbf{u}$

## Units: prostate cancer data (1/2)

- Individuals: 506 patients with prostatic cancer grouped on clinical criteria into two Stages 3 and 4 of the disease
■ Variables: $d=12$ pre-trial variates were measured on each patient, composed by
■ Eight continuous variables (age, weight, systolic blood pressure, diastolic blood pressure, serum haemoglobin, size of primary tumour "SZ", index of tumour stage and histolic grade, serum prostatic acid phosphatase "AP")
- Two ordinal variables (performance rating, cardiovascular disease history)
- Two categorical variables with various numbers of levels (electrocardiogram code, bone metastases)

■ Some missing data: 62 missing values ( $\approx 1 \%$ )

- Two historical units for performing the clustering task:

■ Raw units id: [McParland \& Gormley, 2015] ${ }^{9}$
■ Transformed data u: since SZ and AP are skewed, [Jorgensen \& Hunt, 1996] ${ }^{10}$ propose

$$
\mathbf{u}_{S Z}=\sqrt{\cdot} \text { and } \mathbf{u}_{A P}=\ln (\cdot)
$$

[^7]
## Units: prostate cancer data (2/2)

■ Model $\mathbf{m}$ : full mixed data $\mathbf{x}=\left(\mathbf{x}^{\text {cont }}, \mathbf{x}^{\text {cat }}, \mathbf{x}^{\text {ordi }}, \mathbf{x}^{\text {int }}, \mathbf{x}^{\text {rank }}\right)$ (missing data are allowed also) are simply modeled by inter conditional independence

$$
\mathrm{p}\left(\mathbf{x} ; \boldsymbol{\alpha}_{k}\right)=\mathrm{p}\left(\mathrm{x}^{\text {cont }} ; \boldsymbol{\alpha}_{k}^{\text {cont }}\right) \times \mathrm{p}\left(\mathrm{x}^{\text {cat }} ; \boldsymbol{\alpha}_{k}^{\text {cat }}\right) \times \mathrm{p}\left(\mathrm{x}^{\text {ordi }} ; \boldsymbol{\alpha}_{k}^{\text {ordi }}\right) \times \ldots
$$

In addition, for symmetry between types, intra conditional independence for each
■ Results:
■ New units $\mathbf{u}_{S Z}$ and $\mathbf{u}_{A P}$ are selected by ICL
■ New units allow to select two groups and provides a lower error rate


| clusters |  |
| :---: | :---: |
| 1 | 2 |
| 287 | 5 |
| 52 | 162 |

Table : Raw units: $11 \%$ misclassified

| clusters |  |
| :---: | :---: |
| 1 | 2 |
| 270 | 22 |
| 23 | 191 |

Table : New units: $9 \%$ misclassified

## Variable selection for Gaussians ${ }^{11}$

Definition

$$
\mathrm{p}\left(\mathbf{x}_{1} ; \boldsymbol{\theta}\right)=\underbrace{\left\{\sum_{k=1}^{K} \pi_{k} \mathrm{p}\left(\mathbf{x}_{1}^{S} ; \boldsymbol{\mu}_{k}, \mathbf{\Sigma}_{k}\right)\right\}}_{\text {clustering variables }} \times \underbrace{\left\{\mathrm{p}\left(\mathbf{x}_{1}^{U} ; \mathbf{a}+\mathbf{x}_{1}^{R} \mathbf{b}, \mathbf{C}\right)\right\}}_{\text {redundant variables }} \times \underbrace{\left\{\mathrm{p}\left(\mathbf{x}_{1}^{W} ; \mathbf{u}, \mathbf{V}\right)\right\}}_{\text {independent variables }}
$$

where
■ all parts are Gaussians
■ S: set of variables useful for clustering
■ U: set of redondant clustering variables, expressed with $R \subseteq S$
■ W: set of variables independent of clustering

Trick
Variable selection is recasted as a particular model selected by BIC

[^8]
## Outline

1 Need to formalize

2 Formalizing estimation

3 Formalizing selection

4 More advanced formalizing

5 MixtComp in MASSICCC

6 To go further

## MASSICCC?



A high quality and easy to use web platform where are transfered mature research clustering (and more) software towards (non academic) professionals

Here is the computer you need!


## Prostate cancer data ${ }^{12}$

- Individuals: 506 patients with prostatic cancer grouped on clinical criteria into two Stages 3 and 4 of the disease
■ Variables: $d=12$ pre-trial variates were measured on each patient, composed by eight continuous variables (age, weight, systolic blood pressure, diastolic blood pressure, serum haemoglobin, size of primary tumour, index of tumour stage and histolic grade, serum prostatic acid phosphatase) and four categorical variables with various numbers of levels (performance rating, cardiovascular disease history, electrocardiogram code, bone metastases)
- Some missing data: 62 missing values ( $\approx 1 \%$ )

We forget the classes (Stages of the desease) for performing clustering

## Questions

- How many clusters?

■ Which partition?

[^9]
## Enter into MASSICCC

## massiccc.lille.inria.fr



IRY IT I

## Data upload without preprocessing



## Run clustering analysis



## It is running on the (Inria) cloud. . .

MASSICCC Dashboard Help Profile Logout

| OVERVIEW |
| :--- |
| FILES |
| INPUTS |
| RESUITS |

RESULTS

Select a job execution from the list below

| 03 | (10) | Run Demo On Cancer Data Set MixtComp-Examplecsv | 44\% | 5 Feb 16:59 | \% |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 02 | (12) | MixtComp Cluster Functional-Example.csy |  | 3 Feb 19:15 |  |
| 01 | (10) | Essai Prostate Vendredi Soir MixtComp-Example.csy |  | 3 Feb 19:03 |  |

## Several quick result overviews. . . without post-processing



## Variable significance on global partition

## .ll Variable <br> Importance

This chart represents the discriminating level of each variable. A high value (close to one) means that the variable is highly discriminating. A low value (close to zero) means that the variable is poorly discriminating. Click on one of the bars to display the distribution of this variable and, to also display the similarities between this variable and all the others. The color of the bars reflects the similarities between all the variables and the selected variable.

Sort Variables: $\downarrow \uparrow$


## + similarity between variables

## Variable "Age" difference between clusters

This chart summarizes the distribution of the selected
variable.

Age

Boxplot of the distribution per class for Age


Age (Gaussian)
$\checkmark$ Hide model parameters
mean: 71.534 , sigma: 6.760 mean: 71.313 , sigma: 7.463

## Variable "SG" difference between clusters

$\mathcal{f}$ Variable Parameters

This chart summarizes the distribution of the selected
variable.

SG

Boxplot of the distribution per class for SG


## Variable "BM" difference between clusters



This chart summarizes the distribution of the selected variable.

BM

Mode and 95\% CI per class for BM


BM (Multinomial)
$\checkmark$ Hide model parameters
Class 1
Class 2
scatter: $[0.993,0.007]$ scatter: $[0.633,0.367]$

## Individual cluster separation (with the cluster weight)



## Scoring cancer data following the clustering task



## Curve "cookies" data set

The Kneading dataset comes from Danone Vitapole Paris Research Center and concerns the quality of cookies and the relationship with the flour kneading process ${ }^{13}$. There are 115 different flours for which the dough resistance is measured during the kneading process for 480 seconds. One obtains 115 kneading curves observed at 241 equispaced instants of time in the interval [ $0 ; 480$ ]. The 115 flours produce cookies of different quality: 50 of them have produced cookies of good quality, 25 produced medium quality and 40 low quality.


[^10]
## Upload curves data



## Run a clustering task with three clusters



## Overview of the three clusters of cookies



## Curve "cookies" result

Using a basis functional model-based design for functional data ${ }^{14}$

${ }^{14}$ Jacques and Preda (2013)

## Outline

1 Need to formalize

2 Formalizing estimation

3 Formalizing selection

4 More advanced formalizing

5 MixtComp in MASSICCC

6 To go further

## Some remaining questions

- More on dependent data (like times series)

■ High-dimensional data

- Missing not at random data (MNAR)


## Next lesson

Introduction to cluster analysis and classification: Bi-clustering and co-clustering


[^0]:    ${ }^{1}$ high: $30 \%$, middle: $15 \%$, low:5\%

[^1]:    ${ }^{2}$ There are 3 other continuous descriptors we do not use
    $3^{3}$ https://archive.ics.uci.edu/ml/machine-learning-databases/spambase/

[^2]:    ${ }^{4}$ http://cran.r-project.org/web/packages/mice/mice.pdf

[^3]:    ${ }^{5}$ http://cran.r-project.org/web/packages/mice/mice.pdf

[^4]:    ${ }^{6}$ http://cran.r-project.org/web/packages/capushe/

[^5]:    ${ }^{7}$ Possibility to relax this restriction, including for instance linear transformations involved in PCA (principal component analysis). But the variable definition is no longer respected.

[^6]:    ${ }^{8}$ This set is usually restricted to a single element

[^7]:    ${ }^{9}$ McParland, D. and Gormley, I. C. (2015). Model based clustering for mixed data: clustmd. arXiv preprint arXiv:1511.01720.
    ${ }^{10}$ Jorgensen, M. and Hunt, L. (1996). Mixture model clustering of data sets with categorical and continuous variables. In Proceedings of the Conference ISIS, volume 96, pages 375-384.

[^8]:    ${ }^{11}$ Raftery and Dean (2006), Maugis et al. (09a), Maugis et al. (09b)

[^9]:    ${ }^{12}$ Byar DP, Green SB (1980): Bulletin Cancer, Paris 67:477-488

[^10]:    ${ }^{13}$ Lévéder et al, 04

