



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

<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

Summer School on Clustering, Data Analysis and Visualization of Complex Data
May 21-25 2018, University of Catania, Italy



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

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

Expected guarantees on clustering

- \hat{z} , $\hat{\mu}$, \hat{K} , *etc.* are estimates of theoretical quantities z , μ , K , *etc.*
- It can be thus expected classical guarantees provided for any estimate in statistics
- Typically: [consistency](#), [bias](#), [variance](#)
- Examples:

$$p(\hat{K} = K) = 1 \quad \text{as } n \rightarrow \infty$$

$$p(\hat{z} = z) \text{ for finite } n$$

- 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

a cluster \iff 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

- 1 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

$$\begin{aligned}
 W_{\mathbf{I}}(\mathbf{z}) &= \sum_{i=1}^n \sum_{k=1}^K z_{ik} \|\mathbf{x}_i - \boldsymbol{\mu}_k\|_{\mathbf{I}}^2 \\
 &= -2 \sum_{i=1}^n \sum_{k=1}^K z_{ik} \ln \left[\underbrace{\frac{1}{K}}_{!} \underbrace{\frac{1}{(2\pi)^{d/2} |\mathbf{I}|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x}_i - \boldsymbol{\mu}_k)' \mathbf{I} (\mathbf{x}_i - \boldsymbol{\mu}_k)\right)}_{N_d(\boldsymbol{\mu}_k, \mathbf{I})} \right] + \text{cst}
 \end{aligned}$$

Model

d -variate Gaussian with variance matrix \mathbf{I} and same cluster sample size (see later)

Reformulate K -means: the hidden estimate choice

$$W_1(\mathbf{z}) = -2\ell_c(\boldsymbol{\mu}; \mathbf{x}, \mathbf{z}) + \text{cst}$$

$$\Downarrow$$

$$\left. \begin{array}{l} \hat{\mathbf{z}}^W = \arg \min_{\mathbf{z}} W_1(\mathbf{z}) \\ (\hat{\mathbf{z}}^{\ell_c}, \hat{\boldsymbol{\mu}}^{\ell_c}) = \arg \max_{(\mathbf{z}, \boldsymbol{\mu})} \ell_c(\boldsymbol{\mu}; \mathbf{x}, \mathbf{z}) \end{array} \right\} \Rightarrow \hat{\mathbf{z}}^W \equiv \hat{\mathbf{z}}^{\ell_c}$$

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:**

$$p_k(\mathbf{x}_1) = p(\mathbf{x}_1; \alpha_k)$$

thus

$$p(\mathbf{x}_1) = p(\mathbf{x}_1; \theta) = \sum_{k=1}^K \pi_k p(\mathbf{x}_1; \alpha_k)$$

- **Mixture parameter:**

$$\theta = (\pi, \alpha) \text{ with } \alpha = (\alpha_1, \dots, \alpha_K)$$

- **Model:** it includes both the family $p(\cdot; \alpha_k)$ and the number of groups K

$$\mathbf{m} = \{p(\mathbf{x}_1; \theta) : \theta \in \Theta\}$$

The number of free *continuous* parameters is given by

$$\nu = \dim(\Theta)$$




The clustering process in mixtures

- 1 Estimation of θ by $\hat{\theta}$
- 2 Estimation of the **conditional probability** that $\mathbf{x}_i \in G_k$

$$t_{ik}(\hat{\theta}) = p(Z_{ik} = 1 | \mathbf{X}_i = \mathbf{x}_i; \hat{\theta}) = \frac{\hat{\pi}_k p(\mathbf{x}_i; \hat{\alpha}_k)}{p(\mathbf{x}_i; \hat{\theta})}$$



- 3 Estimation of \mathbf{z}_i by *maximum a posteriori* (**MAP**)

$$\hat{\mathbf{z}}_{ik} = \mathbb{I}_{\{k = \arg \max_{h=1, \dots, K} t_{ih}(\hat{\theta})\}}$$



Geometric interpretation of Σ_k

$$\Sigma_k = \underbrace{\lambda_k}_{\text{volume}} \cdot \underbrace{\mathbf{D}_k}_{\text{orientation}} \cdot \underbrace{\mathbf{A}_k}_{\text{shape}} \cdot \mathbf{D}_k'$$



Estimation of θ by *complete*-likelihood

Maximize the *complete-likelihood* over (θ, \mathbf{z})

$$\ell_c(\theta; \mathbf{x}, \mathbf{z}) = \sum_{i=1}^n \sum_{k=1}^K z_{ik} \ln \{ \pi_k p(\mathbf{x}_i; \boldsymbol{\alpha}_k) \}$$

- **Equivalent** to traditional methods

Metric	$\mathbf{M} = \mathbf{I}$	\mathbf{M} free	\mathbf{M}_k free
Gaussian model	$[\pi \lambda I]$	$[\pi \lambda C]$	$[\pi \lambda_k C_k]$

- **Bias** of $\hat{\theta}$: 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 θ by *observe*-likelihood

Maximize the *observe*-likelihood on θ

$$\ell(\theta; \mathbf{x}) = \sum_{i=1}^n \ln p(\mathbf{x}_i; \theta)$$

- **Convergence** of $\hat{\theta}$, 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: θ^0
- Iteration $n^o q$:
 - Step E: estimate probabilities $\mathbf{t}^q = \{t_{ik}(\theta^q)\}$
 - Step C: classify by setting $\mathbf{t}^q = \text{MAP}(\{t_{ik}(\theta^q)\})$
 - Step M: maximize $\theta^{q+1} = \arg \max_{\theta} \ell_c(\theta; \mathbf{x}, \mathbf{t}^q)$
- Stopping rule: iteration number or criterion stability

Properties

- \oplus : simplicity, monotony, low memory requirement
- \ominus : local maxima (depends on θ^0), linear convergence (EM)

Gaussian M-step


$$n_k^{(q)} = \sum_{i=n'+1}^n t_{ik}(\boldsymbol{\theta}^{(q)})$$

$$\pi_k^{(q+1)} = \frac{n_k^{(q)}}{n}$$

$$\boldsymbol{\mu}_k^{(q+1)} = \frac{1}{n_k^{(q)}} \left(\sum_{i=1}^n t_{ik}(\boldsymbol{\theta}^{(q)}) \mathbf{x}_i \right)$$



$$\boldsymbol{\Sigma}_k^{(q+1)} = \frac{1}{n_k^{(q)}} \left(\sum_{i=1}^n t_{ik}(\boldsymbol{\theta}^{(q)}) (\mathbf{x}_i - \boldsymbol{\mu}_k^{(q+1)}) (\mathbf{x}_i - \boldsymbol{\mu}_k^{(q+1)})' \right)$$

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	$\hat{KL}(\theta, \hat{\theta})$		$\text{err}(z, \hat{z})$	
		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

¹high: 30%, middle: 15%, low:5%

Categorical variables: latent class model

- **Categorical variables:** d variables with m_j modalities each, $\mathbf{x}_i^j \in \{0, 1\}^{m_j}$ and

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

- **Conditional independence:**

$$p(\mathbf{x}_i; \alpha_k) = \prod_{j=1}^d \prod_{h=1}^{m_j} (\alpha_k^{jh})^{\mathbf{x}_i^{jh}}$$

and

$$\alpha_k^{jh} = p(\mathbf{x}_i^{jh} = 1 | z_{ik} = 1)$$

with $\alpha_k = (\alpha_k^{jh}; j = 1, \dots, d; h = 1, \dots, m_j)$

Integer: Poisson mixture model

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

$$p(\mathbf{x}_i^{int}; \boldsymbol{\alpha}_k^{int}) = \prod_{j=1}^d \frac{(\alpha_k^j)^{x_i^j}}{\alpha_k^j!} e^{-\alpha_k^j}$$

SPAM E-mail Database³

- $n = 4601$ e-mails composed by 1813 “spams” and 2788 “good e-mails”
- $d = 48 + 6 = 54$ continuous descriptors²
 - 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**


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

²There are 3 other continuous descriptors we do not use

³<https://archive.ics.uci.edu/ml/machine-learning-databases/spambase/>


An EM run with a binary data set

Initial binary data




An EM run with a binary data set

Iteration 1




An EM run with a binary data set

Iteration 2




An EM run with a binary data set

Iteration 3




An EM run with a binary data set

Iteration 4




An EM run with a binary data set

Iteration 5




An EM run with a binary data set

Iteration 6




An EM run with a binary data set

Iteration 7




An EM run with a binary data set

Iteration 8




An EM run with a binary data set

Iteration 9




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



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

$$\mathbf{x}_1 = (\mathbf{x}_1^{cont}, \mathbf{x}_1^{cat}, \mathbf{x}_1^{int})$$

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

$$p(\mathbf{x}_1; \alpha_k) = p(\mathbf{x}_1^{cont}; \alpha_k^{cont}) \times p(\mathbf{x}_1^{cat}; \alpha_k^{cat}) \times p(\mathbf{x}_1^{int}; \alpha_k^{int})$$

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

581

Printed in Great Britain

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, θ , 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 direct-likelihood or Bayesian inferences about θ , 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 θ . 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

- $O_i \subseteq \{1, \dots, d\}$ the set of the observed variables from sample i
- \mathbf{x}_i^O the observed data from sample i
- M_i the set of the missing variables for sample i
- $\boldsymbol{\mu}_{ik}^O$ the sub-vector of $\boldsymbol{\mu}_k$ associated to index O_i (the same for M_i)
- $\boldsymbol{\Sigma}_{ik}^{OM}$ 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 mechanism

Missing At Random (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...

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

EM and Gaussian case: E step

θ and θ^+ the parameters for two successive steps (*idem* for missing data)

$$z_{ik}^+ = P(Z_{ik} = 1 | \mathbf{x}_i^O; \theta) = \frac{\pi_k \phi(\mathbf{x}_i^O; \Sigma_k)}{\sum_{\ell=1}^K \pi_\ell \phi(\mathbf{x}_i^O; \Sigma_\ell)}$$

$$\mathbf{x}_{ik}^{M+} = E \left[\mathbf{x}_i^M \mid \mathbf{x}_i^O, Z_{ik} = 1; \theta \right] = \boldsymbol{\mu}_{ik}^M + \Sigma_{ik}^{MO} \left(\Sigma_{ik}^{OO} \right)^{-1} (\mathbf{x}_i^O - \boldsymbol{\mu}_{ik}^O).$$

Interpretation

- z_{ik}^+ : class posterior probability membership given the available information \mathbf{x}_i^O .
- \mathbf{x}_{ik}^{M+} : conditional imputation of the missing data given the cluster.

EM and Gaussian case: M step

$$\begin{aligned}\pi_k^+ &= \frac{1}{n_k^+} \sum_{i=1}^n z_{ik}^+, \quad \mu_k^+ = \frac{1}{n_k^+} \sum_{i=1}^n z_{ik}^+ \mathbf{x}_{ik}^+ \\ \Sigma_k^+ &= \frac{1}{n_k^+} \sum_{i=1}^n z_{ik}^+ \left[(\mathbf{x}_{ik}^+ - \mu_k^+)(\mathbf{x}_{ik}^+ - \mu_k^+)' + \Sigma_{ik}^+ \right]\end{aligned}$$

where $n_k^+ = \sum_{i=1}^n z_{ik}^+$, $\mathbf{x}_{ik}^+ = \begin{pmatrix} \mathbf{x}_i^O \\ \mathbf{x}_{ik}^{M^+} \end{pmatrix}$, $\Sigma_{ik}^+ = \begin{pmatrix} 0_i^O & 0_i^{OM} \\ 0_i^{MO} & \Sigma_{ik}^{M^+} \end{pmatrix}$ with 0 the $d \times d$ null matrix, and $\Sigma_{ik}^{M^+} = \Sigma_{ik}^{MO} (\Sigma_{ik}^O)^{-1} \Sigma_{ik}^{OM}$.

Interpretation of $\Sigma_{ik}^{M^+}$

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 θ by maximizing the **observed**-data log-likelihood

- Initialisation: $\theta^{(0)}$
- Iteration nb q :
 - **E-step**: compute conditional probabilities $p(x^M, z | \mathcal{D}; \theta^{(q)})$
 - **S-step**: draw $(x^{M(q)}, z^{(q)})$ from $p(x^M, z | x^O; \theta^{(q)})$
 - **M-step**: maximize $\theta^{(q+1)} = \arg \max_{\theta} \ln p(x^O, x^{M(q)}, z^{(q)}; \theta)$
- 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 $(\theta^{(q)})$ approximates $\hat{\theta}$
- The variance of the sequence $(\theta^{(q)})$ gives confidence intervals

Missing data: SE algorithm

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

- Iteration nb q :
 - **E-step**: compute conditional probabilities $p(\mathbf{x}^M, \mathbf{z}^M | \mathcal{D}; \hat{\theta})$
 - **S-step**: draw $(\mathbf{x}^{M(q)}, \mathbf{z}^{M(q)})$ from $p(\mathbf{x}^M, \mathbf{z}^M | \mathcal{D}; \hat{\theta})$
- Stopping rule: iteration number

Properties

- simplicity because of conditional independence
- the mean/mode of the sequence $(\mathbf{x}^{M(q)}, \mathbf{z}^{M(q)})$ estimates $(\mathbf{x}^M, \mathbf{z}^M)$
- confidence intervals are also derived

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

- **Strategy “mice⁴ + 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

⁴<http://cran.r-project.org/web/packages/mice/mice.pdf>

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**



true modèle: $[\pi \lambda_k I]$

too simple model: $[\pi \lambda I]$

Too complex model: **variance**



— true borderline

- - - borderline with $[\pi \lambda I]$

. . . borderline with $[\pi \lambda_k C_k]$

Importance of model selection: bias/variance trade-off

- **Partition error rate:** $\text{err}(\mathbf{z}_1, \mathbf{z}_2) \geq 0$ a distance-like between two partitions $\mathbf{z}_1, \mathbf{z}_2$
- **Gap between true and model partition:**

$$\theta_m^* = \arg \min_{\theta \in \Theta_m} \text{err}(\mathbf{z}, \mathbf{z}(\theta))$$

- **MLE:**

$$\hat{\theta}_m = \arg \max_{\theta \in \Theta} \ell(\theta; \mathcal{D})$$

- **Fundamental decomposition of $\text{err}(\mathbf{z}, \mathbf{z}(\hat{\theta}_m))$:**

$$\begin{aligned} \text{err}(\mathbf{z}, \mathbf{z}(\hat{\theta}_m)) &= \left\{ \text{err}(\mathbf{z}, \mathbf{z}(\theta_m^*)) - \text{err}(\mathbf{z}, \mathbf{z}) \right\} + \left\{ \text{err}(\mathbf{z}, \mathbf{z}(\hat{\theta}_m)) - \text{err}(\mathbf{z}, \mathbf{z}(\theta_m^*)) \right\} \\ &= \left\{ \text{bias}_m \right\} + \left\{ \text{variance}_m \right\} \end{aligned}$$

Importance of model selection: illustration of the variance effect

30 samples from a bivariate mixture with two components

$$\pi_1 = \pi_2 = 0.5, \quad \mu_1 = (0, 0)', \quad \mu_2 = (2, 2)', \quad \Sigma_1 = \Sigma_2 = \mathbf{I}$$

$$\mathcal{M} = \{\text{spherical}, \text{general}\}$$

n	\mathbf{m}	$\text{err}(\mathbf{z}, \hat{\mathbf{z}}_{\mathbf{m}})$
40	spherical	0.0967
	general	0.1100
200	spherical	0.0840
	general	0.0872

Some heuristics entropy-based criteria: examples

- A fundamental decomposition of $\ell(\boldsymbol{\theta}; \mathbf{x})$: for any “fuzzy partition” $\mathbf{c} = \{c_{ik}\}$

$$\begin{aligned}\ell(\boldsymbol{\theta}; \mathbf{x}) &= \sum_{i=1}^n \sum_{k=1}^K c_{ik} \ln \{\pi_k p(\mathbf{x}_i; \boldsymbol{\alpha}_k)\} - \sum_{i=1}^n \sum_{k=1}^K c_{ik} \ln t_{ik}(\boldsymbol{\theta}) \\ &= \ell_c(\boldsymbol{\theta}; \mathbf{x}, \mathbf{c}) + \xi(\boldsymbol{\theta}; \mathbf{c}) \\ &= \text{complete-data log-likelihood} + \text{entropy}\end{aligned}$$

- NEC criterion (*Normalized Entropy Criterion*): retain \mathbf{m} minimizing



$$\text{NEC}_K = \begin{cases} \frac{\xi(\hat{\boldsymbol{\theta}}_K; \mathbf{t}(\hat{\boldsymbol{\theta}}_K))}{\ell(\hat{\boldsymbol{\theta}}_K; \mathbf{x}) - \ell(\hat{\boldsymbol{\theta}}_1; \mathbf{x})} & \text{if } K > 1 \\ 1 & \text{if } K = 1 \end{cases}$$

- CL criterion (*Completed Likelihood*): retain \mathbf{m} maximizing

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

- 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:

$$CL_{([p\lambda I], K)} = -\frac{nd}{2} \ln(W_K) - n \ln(K) + \text{cst}$$



Theoretical model selection criteria

The most widespread principle

$$\underbrace{\text{Criterion}}_{\text{to be maximized}} = \underbrace{\text{maximum log-likelihood}}_{\text{model-data adequacy}} - \underbrace{\text{penalty}}_{\text{"cost" of the model}}$$

criterion	penalty	interpretation	user purpose
-----------	---------	----------------	--------------

general criteria in statistics

AIC	ν	model complexity	prediction
BIC	$0.5\nu \ln(n)$	model complexity	identification

specific criterion for the clustering aim

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

BIC criterion: integrated likelihood

- Posterior likelihood of \mathbf{m} :

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

- Ideal model in a Bayesian context:

$$\hat{\mathbf{m}}^* \in \arg \max_{\mathbf{m} \in \mathcal{M}} p(\mathbf{m}|\mathcal{D})$$

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

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

- Difficulties:

- Choose the prior $p(\theta|\mathbf{m})$
- Evaluate the integral

BIC criterion: genesis

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

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

- **BIC criterion (*Bayesian Information Criterion*)**: retain \mathbf{m} maximizing

$$\text{BIC}_{\mathbf{m}} = \ell(\hat{\boldsymbol{\theta}}_{\mathbf{m}}; \mathcal{D}) - \frac{\nu_{\mathbf{m}}}{2} \ln(n)$$

BIC criterion: consistency

- **Consistency**: BIC asymptotically selects the best

$$\mathbf{m}^* = \arg \inf_{\mathbf{m} \in \mathcal{M}} \text{KL}(\mathbf{p}, \mathbf{p}_{\theta_{\mathbf{m}}^*})$$

- **Theoretical illustration of consistency**: $\mathbf{m}_1 \subseteq \mathbf{m}_2$, \mathbf{m}_1 being the true model, $\Delta\nu = \nu_2 - \nu_1$, $\Delta\ell = \ell(\hat{\theta}_2; \mathcal{D}) - \ell(\hat{\theta}_1; \mathcal{D})$, we have

$$2(\text{BIC}_2 - \text{BIC}_1) + \Delta\nu \ln(n) = 2\Delta\ell \xrightarrow{d} \chi_{\Delta\nu}^2$$

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

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

by using the Chebyshev inequality. Thus, asymptotically, BIC will select \mathbf{m}_1

- **Special case of K**: be careful on the χ^2 approximation validity...

ICL criterion: genesis

- **Revisiting the fundamental decomposition:** if \mathbf{z} known, retain \mathbf{m} maximizing

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

Thus models leading to overlapping groups are more penalized (low \mathbf{z} evidence)

- **ICL criterion (*Integrated Classification Likelihood*):** replace \mathbf{z} by $\hat{\mathbf{z}}$

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

- **BIC-like approximation of ICL:**


$$\ln p(\mathbf{x}, \mathbf{z} | \mathbf{m}) = \ln p(\mathbf{x}, \mathbf{z}; \hat{\boldsymbol{\theta}}_{\mathbf{x}, \mathbf{z}} | \mathbf{m}) - \frac{\nu}{2} \ln n + O_p(1)$$

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

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

ICL criterion: robustness to model misspecification

- A bivariate mixture of a uniform and a Gaussian cluster:
 - non-Gaussian component: $\pi_1 = 0.5$, $p_1(\mathbf{x}_1) = 0.25 \mathbf{1}_{[-1,1]}(x^1) \mathbf{1}_{[-1,1]}(x^2)$
 - Gaussian component: $\pi_2 = 0.5$, $\boldsymbol{\mu}_2 = (3.3, 0)'$, $\boldsymbol{\Sigma}_2 = \mathbf{I}$
- 50 simulated data sets of size $n = 200$



K	1	2	3	4	5
BIC	.	60	.	32	8
ICLbic	.	100	.	.	.

ICL criterion: consistency?

- **Assumption:** true model with two groups and parameter θ_2^*
- **Theoretical result:**
 - Preliminaries: $\delta_n = n(\theta_2^* - \theta_2^{*P})' \mathbf{J}(\theta_2^*)(\theta_2^* - \theta_2^{*P})$, $\mathbf{J}(\theta_2^*)$ the Fisher matrix for a data unit calculated with the true parameter θ_2 and θ_2^{*P} its projected value on the parameter subspace associated to the one component case, $\mu_n = \mathbb{E}[\chi_{\Delta\nu}^2(\delta_n)] = \Delta\nu + \delta_n$, $\sigma_n^2 = \text{Var}[\chi_{\Delta\nu}^2(\delta_n)] = 2(\Delta\nu + \delta_n)$
 - Asymptotically: by Chebishev inequality, with $\mu_n - \Delta\nu \ln n - 2n \ln 2 > 0$

$$p(\text{choose wrong model}) = p(\text{ICLbic}_2 < \text{ICLbic}_1) \leq \frac{\sigma_n^2}{(\mu_n - \Delta\nu \ln n - 2n \ln 2)^2}$$

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
1 000	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



Reality is even worse: $n=10^6$ 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 “mice⁵ + 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

⁵<http://cran.r-project.org/web/packages/mice/mice.pdf>

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



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



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 \mathbf{m} maximizing

$$SH_{\mathbf{m}} = \ell(\hat{\boldsymbol{\theta}}_{\mathbf{m}}; \mathcal{D}) - 2\widehat{\text{variance}}_{\mathbf{m}}$$

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

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

and also



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

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



- **CAPUSHE⁶ (CALibrated Penalty Using Slope HEuristics)**: $\kappa/2$ can be estimated by a linear regression of $\ell(\hat{\boldsymbol{\theta}}_{\mathbf{m}}; \mathcal{D})$ on $\frac{\kappa}{2} \nu_{\mathbf{m}}$

⁶<http://cran.r-project.org/web/packages/capushe/>

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






- “During”: model outliers as a particular cluster in the mixture

$$p(\mathbf{x}_1; \boldsymbol{\theta}) = \sum_{k=1}^K \pi_k p(\mathbf{x}_1; \boldsymbol{\alpha}_k) + \pi_0 p(\mathbf{x}_1; \boldsymbol{\alpha}_0)$$

Outliers: “during” example with acoustic emission control

- **Data:** $n = 2\,061$ 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} :

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

- \mathbf{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 $\mathbf{u}^{-1} \circ \mathbf{u} = \text{id}$
- Thus, id is only a particular unit \mathbf{u}
- Often a **meaningful** restriction⁷ on \mathbf{u} : it proceeds lines by lines and rows by rows

$$\mathbf{u}(\mathbf{x}) = (\mathbf{u}(\mathbf{x}_1), \dots, \mathbf{u}(\mathbf{x}_n)) \quad \text{with} \quad \mathbf{u}(\mathbf{x}_i) = (\mathbf{u}_1(x_{i1}), \dots, \mathbf{u}_d(x_{id}))$$

- Advantage to respect the variable definition, transforming only its unit
- $\mathbf{u}(\mathbf{x}_i)$ 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

⁷Possibility to relax this restriction, including for instance linear transformations involved in PCA (principal component analysis). But the variable definition is no longer respected.

Units: revisiting units as a modelling component

- Explicitly exhibiting the “canonical” unit **id** in the model

$$p_m = \{\cdot \in \mathcal{X} \mapsto p(\cdot; \theta) : \theta \in \Theta_m\} = \{\cdot \in \mathcal{X}^{\text{id}} \mapsto p(\cdot; \theta) : \theta \in \Theta_m\} = p_m^{\text{id}}$$

- Thus the variable space and the probability measure are **embedded**
- As the **standard probability theory**: a couple (variable space, probability measure)!
- Changing **id** into **u**, while preserving **m**, is expected to produce a new modelling

$$p_m^u = \{\cdot \in \mathcal{X}^u \mapsto p(\cdot; \theta) : \theta \in \Theta_m\}.$$

A model should be systematically defined by a couple **(u,m)**, denoted by p_m^u

Units: interpretation and identifiability of p_m^u

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

$$u^{-1}(m) \in \{m' \in \mathbb{M} : p_{m'}^{\text{id}} = p_m^u\}$$

- There exists **two alternative interpretations** of strictly the same model:
 - p_m^u : data measured with **unit u** arise from **measure m**;
 - $p_{u^{-1}(m)}^{\text{id}}$: data measured with **unit id** arise from **measure $u^{-1}(m)$**
- Two points of view:

Statistician

The model p_m^u is not identifiable over the couple (m, u)

Practitioner

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

⁸This set is usually restricted to a single element

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 $p_{m_1}^{u_1}$ and $p_{m_2}^{u_2}$

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

- **Prohibited** to compare m_1 in unit u_1 and m_2 in unit u_2
- But **allowed** after transforming in **identical unit id**
- Thus compare their equivalent expression: $p_{u_1^{-1}(m_1)}^{id}$ and $p_{u_2^{-1}(m_2)}^{id}$
- Example for abs. continuous x and differentiable u , the **density transform** in **id** is:

$$p_{u^{-1}(m)}^{id} = \{\cdot \in \mathcal{X}^{id} \mapsto p(u(\cdot); \theta) \times |J^u(\cdot)| : \theta \in \Theta_m\}$$

with $J^u(\cdot)$ the **Jacobian** associated to the transformation 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 histologic 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

$$\mathbf{u}_{SZ} = \sqrt{\cdot} \text{ and } \mathbf{u}_{AP} = \ln(\cdot)$$

⁹McParland, D. and Gormley, I. C. (2015). Model based clustering for mixed data: clustmd. arXiv preprint arXiv:1511.01720.

¹⁰Jorgensen, 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.

Units: prostate cancer data (2/2)


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

$$p(\mathbf{x}; \alpha_k) = p(\mathbf{x}^{cont}; \alpha_k^{cont}) \times p(\mathbf{x}^{cat}; \alpha_k^{cat}) \times p(\mathbf{x}^{ordi}; \alpha_k^{ordi}) \times \dots$$

In addition, for symmetry between types, **intra conditional independence** for each

- **Results:**

- New units \mathbf{u}_{SZ} and \mathbf{u}_{AP} 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 Gaussians¹¹

Definition

$$p(\mathbf{x}_1; \boldsymbol{\theta}) = \underbrace{\left\{ \sum_{k=1}^K \pi_k p(\mathbf{x}_1^S; \boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k) \right\}}_{\text{clustering variables}} \times \underbrace{\left\{ p(\mathbf{x}_1^U; \mathbf{a} + \mathbf{x}_1^R \mathbf{b}, \mathbf{C}) \right\}}_{\text{redundant variables}} \times \underbrace{\left\{ p(\mathbf{x}_1^W; \mathbf{u}, \mathbf{V}) \right\}}_{\text{independent variables}}$$

where

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

Trick

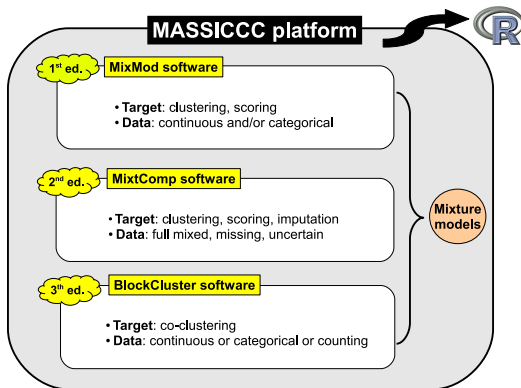
Variable selection is recasted as a particular model selected by BIC

¹¹Raftery and Dean (2006), Maugis *et al.* (09a), Maugis *et al.* (09b)

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 transferred mature research clustering (and more) software
towards (non academic) professionals

Here is the computer you need!



Prostate cancer data¹²

- **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 histologic 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 disease) for performing **clustering**

Questions

- How many clusters?
- Which partition?

¹²Byar DP, Green SB (1980): Bulletin Cancer, Paris 67:477-488

Enter into MASSICCC

massiccc.lille.inria.fr

Massive Clustering with Cloud Computing

Clustering of heterogeneous data with missing values.

Hosted in the cloud. No installation or configuration required.

Upload your data, and get results straight away.

Developed by *inria*

TRY IT !

Data upload without preprocessing

MASSICCC Dashboard Help Profile Logout

OVERVIEW
FILES
INPUTS
RESULTS

Age Wt PF HX SBP DBP EKG HG SZ SG A

Contli Contli Categ Categ Contli Contli Categ Contli Contli Contli

Save

Preview

	Age	Wt	PF	HX	SBP	DBP	EKG	HG	SZ	SG	AP	BM
0	75	76	1	1	15	9	5	138	14142	8	1.0986	1
1	76	?	?	?	?	?	?	?	5.3852	9	2.4849	?
2	54	116	1	1	13	7	4	146	6.4807	?	1.9459	1
3	69	102	1	2	14	8	5	134	1.7321	9	1.0986	1
4	66	?	?	?	?	?	?	?	1.0000	9	2.3979	?

Run clustering analysis

MASSICCC [Dashboard](#) [Help](#) [Profile](#) [Logout](#)

OVERVIEW

FILES

INPUTS

RESULTS

INPUTS

Parameters

Title

Run demo on cancer data set

Data File

MixtComp-Example.csv

Package

MxMod

MixtComp

BlockCluster

Function

Cluster

Labels Column

Cluster Groups

1-7

Create







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

MASSICCC Dashboard Help Profile Logout


OVERVIEW
FILES
INPUTS
RESULTS

RESULTS

Select a job execution from the list below

03		Run Demo On Cancer Data Set MixtComp-Example.csv	<div><div>44%</div></div>	5 Feb 16:59 
02		MixtComp Cluster Functional-Example.csv		3 Feb 19:15 
01		Essai Prostate Vendredi Soir MixtComp-Example.csv		3 Feb 19:03 

Several quick result overviews. . . without post-processing




Variable significance on global partition

Variable 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.

[Read more](#)



Sort Variables :

+ similarity between variables

Variable “Age” difference between clusters




Variable Parameters

This chart summarizes the distribution of the selected variable.

Age

Boxplot of the distribution per class for Age



Age (Gaussian)

▼ Hide model parameters

Class 1

Class 2

mean: 71.534, sigma: 6.760 mean: 71.313, sigma: 7.463

Variable “SG” difference between clusters




Variable Parameters

This chart summarizes the distribution of the selected variable.

SG

Boxplot of the distribution per class for SG



SG (Gaussian)

▼ Hide model parameters

Class 1

mean: 8.940, sigma: 1.154

Class 2

mean: 12.087, sigma: 1.405


Variable “BM” difference between clusters

Variable Parameters

This chart summarizes the distribution of the selected variable.

BM

Mode and 95% CI per class for BM




BM (Multinomial)

▼ 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

MASSICCC Dashboard Help Profile Logout

OVERVIEW
FILES
INPUTS
RESULTS

INPUTS

Parameters

Title

Scoring following the clustering task

Data File

MixtComp-Example.csv

Package

MixMod

MixtComp

BlockCluster

Function

Predict

Cluster

Predict


Classification Model

03	Avril Demo On Cancer Data Set MixtComp-Example.csv	5 Feb 16:59
02	MixtComp Cluster Functional-Example.csv	3 Feb 19:15
01	Essai Prostate Vendredi Soir MixtComp-Example.csv	3 Feb 19:03

Create

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 process¹³. 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.



¹³Lévêder *et al*, 04

Upload curves data

MASSICCC Dashboard Help Profile Logout

OVERVIEW

FILES

INPUTS

RESULTS

Set: ☐ All ☐ As: Categorical

Upload a file with a list of datatypes for each column.

Aucun fichier choisi

Function

Preview

	Function
0	0.251.226202169594.2.257.61097125343.4.263.758...
1	0.241.129520478231.2.245.716088727869.4.250.18...
2	0.194.07006418218.2.196.013131806268.4.197.956...
3	0.137.021447956417.2.154.635389904923.4.170.65...
4	0.244.120130204111.2.245.627062897663.4.247.13...

Run a clustering task with three clusters

MASSICCC [Dashboard](#) [Help](#) [Profile](#) [Logout](#)

OVERVIEW
FILES
INPUTS
RESULTS

INPUTS

Parameters

Title

Clustering of cookies into three clusters

Data File

Functional-Example.csv

Package

MixMod

MixtComp

BlockCluster

Function

Cluster

Labels Column

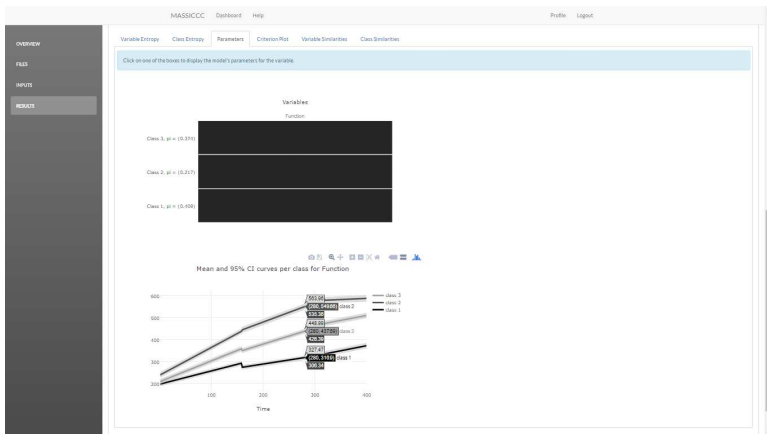
Cluster Groups

3

Variable Params


Create

Overview of the three clusters of cookies



Curve “cookies” result

Using a basis functional model-based design for functional data¹⁴



¹⁴Jacques 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