# Traitement statistique des données manquantes-Part I Introduction to modeling 

Christophe Biernacki

## To cite this version:

Christophe Biernacki. Traitement statistique des données manquantes-Part I Introduction to modeling. Doctorat. France. 2022. hal-03505648

HAL Id: hal-03505648

## https://hal.science/hal-03505648

Submitted on 31 Dec 2021

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.

# Traitement statistique des données manquantes 

Atelier Statistique de la SFdS - 10 et 11 mars 2021

# Part I <br> Introduction to modeling 

C. Biernacki

Laboratoire P. Painlevé, UMR CNRS 8524 \& Université de Lille \& Inria


Merci à : F. Antonazzo, C. Boyer, G. Celeux, Q. Grimonprez, J. Jacques, J. Josse, C. Keribin, V. Kubicki, F. Laporte, M. Marbac, A. Sportisse, J. Vandaele, V. Vandewalle

## Lectures

Books about missing data... but also about modeling and its applications


## Preamble

## What is this lesson?

- Be able to perform practically some (well-established or advanced) methods on missing data
■ Use them with discernment


## What is not this lesson?

- Not an exhaustive list of missing data methods (and related bibliography)
- Do not make specialists of missing data methods

This preamble is valid for all four parts:
1 Part I: Introduction to modeling
[2 Part II: Numerical and non-numerical data
3 Part III: Missing not at random data (MNAR) $\rightarrow$ ongoing research
4 Part IV: Binned data for big data analysis $\rightarrow$ ongoing research

## Outline

## 1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling

5 Frequentist estimation

6 Bayesian estimation

7 Model selection

## Everything begins from data!

## Data sets structure



Todays features: full mixed/missing


## Large data sets $(n)^{1}$



[^0]An opportunity for detecting weak signal


## High-dimensional data $(d)^{2}$



[^1]
## An opportunity for amplifying the signal

Each variable provides equal and own separation information



## Genesis of "Big Data"

The Big Data phenomenon mainly originates in the increase of computer and digital resources at an ever lower cost

■ Storage cost per MB: 700\$ in 1981, $1 \$$ in 1994, $0.01 \$$ in 2013
$\rightarrow$ price divided by 70,000 in thirty years
■ Storage capacity of HDDs: $\approx 1.02$ Go in $1982, \approx 8$ To today
$\rightarrow$ capacity multiplied by 8,000 over the same period

- Computeur processing speed: 1 gigaFLOPS ${ }^{3}$ in 1985, 33 petaFLOPS in 2013
$\rightarrow$ speed multiplied by 33 million

[^2]
## Digital flow

- Digital in 1986: $1 \%$ of the stored information, 0.02 Eo $^{4}$
- Digital in 2007: $94 \%$ of the stored information, 280 Eo (multiplied by 14,000)


[^3]
## Societal phenomenon

All human activities are impacted by data accumulation

■ Trade and business: corporate reporting system, banks, commercial transactions, reservation systems. .

- Governments and organizations: laws, regulations, standardizations, infrastructure...

■ Entertainment: music, video, games, social networks...

- Sciences: astronomy, physics and energy, genome,...
- Health: medical record databases in the social security system. .

■ Environment: climate, sustainable development, pollution, power...
■ Humanities and Social Sciences: digitization of knowledge, literature, history , art, architecture, archaeological data...

## Three challenges

- The storage challenge:
- Storage, transfer, preservation, availability
- Ex.: in Astrophysics, 50GB acquisition per day for the Euclid project (2021)

■ The data analysis challenge: $3 \mathrm{~V}, 4 \mathrm{~V}, 5 \mathrm{~V} \ldots$


- The societal and economic challenge:
- Protection of private life, right to be forgotten, property rights, operating rights, cost of energy storage or transfer
- Ex.: the PRISM project of the NSA


## Coding for data $\boldsymbol{x}$

- A set of $n$ individuals

$$
\boldsymbol{x}=\left\{\boldsymbol{x}_{1}, \ldots, \boldsymbol{x}_{n}\right\}
$$

with $\boldsymbol{x}_{i}$ a set of (possibly non-scalar) $d$ variables

$$
\boldsymbol{x}_{i}=\left\{\boldsymbol{x}_{i 1}, \ldots, \boldsymbol{x}_{i d}\right\}
$$

where $\boldsymbol{x}_{i j} \in \mathcal{X}_{j}$

- A n-uplet of individuals

$$
\boldsymbol{x}=\left(\mathbf{x}_{1}, \ldots, \boldsymbol{x}_{n}\right)
$$

with $\boldsymbol{x}_{i}$ a $d$-uplet of (possibly non-scalar) variables

$$
\boldsymbol{x}_{i}=\left(\boldsymbol{x}_{i 1}, \ldots, \boldsymbol{x}_{i d}\right) \in \mathcal{X}
$$

where $\mathcal{X}=\mathcal{X}_{1} \times \ldots \ldots \mathcal{X}_{d}$

We will pass from a coding to another, depending of the practical utility (useful for some calculus to have matrices or vectors for instance)

## Outline

1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling
[5] Frequentist estimation

6 Bayesian estimation

7 Model selection

More data. . . implies more missing data!

## "Classical" missing data

Today, it is easy to collect many features, so it favors

- data variety and/or mixed
- data missing
- data uncertainty (or interval data)

Mixed, missing, uncertain

| Observed individuals $\boldsymbol{x}^{\circ}$ |  |  |  |
| :---: | :---: | :---: | :---: |
| $?$ | 0.5 | $?$ | 5 |
| 0.3 | 0.1 | green | 3 |
| 0.3 | 0.6 | \{red, green | 3 |
| 0.9 | $[0.250 .45]$ | red | $?$ |
| $\downarrow$ | $\downarrow$ | $\downarrow$ | $\downarrow$ |
| continuous | continuous | categorical | integer |

## Missing data: notations

- $\boldsymbol{c}=\left\{\boldsymbol{c}_{1}, \ldots, \boldsymbol{c}_{n}\right\}$ : pattern of missing data for the full dataset

■ $\boldsymbol{c}_{i}=\left(c_{i 1}, \ldots, c_{i d}\right) \in\{0,1\}^{d}$ : pattern of missing data for individual $i \in\{1, \ldots, n\}$

$$
c_{i j}=1 \Leftrightarrow x_{i j} \text { is missing }
$$

- $\boldsymbol{x}^{\circ}$ : the observed values in $\boldsymbol{x}$
- $\boldsymbol{x}^{m}$ : the missing values in $\boldsymbol{x}$

$$
\boldsymbol{x}=\left\{\boldsymbol{x}^{\circ}, \boldsymbol{x}^{m}\right\} \text { is the full dataset with its observed and missing parts }
$$

Notation illustration

$$
\boldsymbol{x}^{\circ}=\left(\begin{array}{ccc}
? & 2.6 & 5 \\
\text { blue } & 1.9 & 4 \\
\text { red } & 2.3 & ?
\end{array}\right), \quad \boldsymbol{c}=\left(\begin{array}{ccc}
1 & 0 & 0 \\
0 & 0 & 0 \\
0 & 0 & 1
\end{array}\right)
$$

## More advanced missing data: latent class in clustering (1/3)

## Clustering everywhere ${ }^{5}$...

## Data Mining Algorithms

- Decision trees, regression, and cluster analysis continue to form a triad of core algorithms for most data miners. This has been very consistent over time.
- However, a wide variety of algorithms are being used.



## See Part II

${ }^{5}$ Rexer Analytics's Annual Data Miner Survey is the largest survey of data mining, data science, and analytics

More advanced missing data: latent class in clustering (2/3)

## Clustering

Detect hidden structures in data sets: opportunity to reveal new information...



## More advanced missing data: latent class in clustering (3/3)

- Aim: estimation of the partition $z$ and the number of clusters $K$

■ Notation : partition in $K$ clusters $G_{1}, \ldots, G_{K}: z=\left(z_{1}, \ldots, z_{n}\right)$, $\mathbf{z}_{i}=\left(z_{i 1}, \ldots, z_{i K}\right)^{\prime}$

$$
\boldsymbol{x}_{i} \in G_{k} \quad \Leftrightarrow \quad z_{i h}=\mathbb{I}_{\{h=k\}}
$$

Mixed, missing, uncertain

| Individuals $\boldsymbol{x}^{\circ}$ |  |  |  |  |  | Partition $\boldsymbol{z}$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |$\Leftrightarrow$

All the partition $z$ is unknown, thus it corresponds also to missing data...

## Other more advanced missing data: binned data

- Binned data are equivalent to a multivariate histogram
- Data are partially missing since precise value is unknown inside each bin

(a) Raw data

(b) Binned data


## Classical strategies for dealing with missing data

Two traditional solutions (for obtaining a filled dataset)

- Discard individuals with missing data: more variance or a biased subset
- Impute missing data ${ }^{6}$ : possible bias and underestimation of the variability


## General guidelines

- Obtaining a complete dataset is not the final goal
- Missing data management should take into account the initial analysis goal


## Purpose of Part I

Embed missing data management into the final goal paradigm...

[^4]Illustration of the risk of imputation

Data


## Illustration of the risk of imputation

Data after imputation of missing data by the mean


Illustration of the risk of imputation

Data after imputation of missing data by regression


## Outline

1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling

5 Frequentist estimation

6 Bayesian estimation

7 Model selection

## Prostate cancer data ${ }^{7}$

- 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 ( $\approx 1 \%$ )

[^5]
## Descriptors of the prostate cancer data set

| Courriate | Abreeriation | Number of Levelts (if categorical) |
| :---: | :---: | :---: |
| Age | Age |  |
| Weight | Wt |  |
| Peformance rating | PF | 4 |
| Cardionascular disease history | HX | 2 |
| Systolic Blood pressure | SBP |  |
| Diastlic blod pressure | DBP |  |
| Electrocardiogram code | EKG | 7 |
| Serum haemogolbin | HG |  |
| Sizze of primary tumour | 52 |  |
| Index of tumour stage and histolic grade | SG |  |
| Serum prostatic acid phosphatase | AP |  |
| Bone metastases | BMI | . |

## Aim

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

## Questions

- How many clusters?
- Which partition?

Visually not so easy. . .



## Two strategies in competition

■ Strategy "mice ${ }^{8}+$ MixtComp $^{9}$ ": MixtComp on the dataset completed by mice

```
> data.imp=mice(data)
```

> data.comp.mice=complete(data.imp)

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

More information about Mixtcomp in Part II (be patient)..

[^6]
## Choosing $K$ with the ICL criterion ${ }^{10}$



$$
\begin{gathered}
\text { mice }+ \text { MixtComp } \\
\hat{K}=7
\end{gathered}
$$


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

## Partition quality with $K=2$

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

To be compared also to missing data removal:

- 475 patients with non-missing data
- MixtComp for clustering
- possibility to consider continuous, categorical or mixed data

| Strategy | continuous only | categorical only | mixed cont/cat |
| :---: | :---: | :---: | :---: |
| $\%$ misclassified | 9.46 | 47.16 | 8.63 |

- risk of information lost when removing missing data lines/columns
- avoid to complete missing data (imputation depends on the purpose)


## Outline

1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling
[5] Frequentist estimation

6 Bayesian estimation

7 Model selection

# Full modeling relies on full observed data ( $\boldsymbol{x}^{\circ}, \boldsymbol{c}$ ) 

## From "partial" to "full" modeling

## Partial modeling

- In statistics, most methods rely on

1 modeling $\mathrm{p}(\boldsymbol{x} ; \boldsymbol{\theta})$
2 estimating $\boldsymbol{\theta}$ from $\boldsymbol{x}$
■ Examples : density estimation, regression (just add covariates), clustering. . .

- However, it involves both observed and missing data since $\boldsymbol{x}=\left(\boldsymbol{x}^{\circ}, \boldsymbol{x}^{m}\right)$
- It is the reason why imputation/deletion of missing entries is popular!


## Full modeling

■ However, strictly speaking

- complete data are ( $\boldsymbol{x}, \boldsymbol{c}$ )
- observed data are $\left(\boldsymbol{x}^{\circ}, \boldsymbol{c}\right)$
- Consequently

1 modeling should concern $\mathrm{p}(\boldsymbol{x}, \boldsymbol{C} ; \boldsymbol{\theta}, \psi)=\mathrm{p}(\boldsymbol{c} \mid \boldsymbol{x} ; \psi) \mathrm{p}(\boldsymbol{x} ; \boldsymbol{\theta})$
2 estimating $\boldsymbol{\theta}$ should be performed only from ( $\boldsymbol{x}^{\circ}, \boldsymbol{c}$ )

## Missing data: typology of the missing mechanisms

- Missing completely at random (MCAR):

$$
\mathrm{p}(\boldsymbol{c} \mid \boldsymbol{x} ; \psi)=\mathrm{p}(\boldsymbol{c} ; \psi) \quad \forall \boldsymbol{x}
$$

- Missing at random (MAR):

$$
\mathrm{p}(\boldsymbol{c} \mid \boldsymbol{x} ; \psi)=\mathrm{p}\left(\boldsymbol{c} \mid \boldsymbol{x}^{o} ; \psi\right) \quad \forall \boldsymbol{x}^{m}
$$

■ Missing not at random (MNAR): the mechanism is not MCAR nor MAR

## Example of MNAR data

The probability to have a missing value on income depends on the value of income (rich people less inclined to reveal their income).

## 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.

[^7]
## Ignorable vs. non ignorable model

A missing mechanism is ignorable if likelihoods can be decomposed as

$$
L\left(\boldsymbol{\theta}, \psi ; \boldsymbol{x}^{\circ}, \boldsymbol{c}\right)=L\left(\psi ; \boldsymbol{c} \mid \boldsymbol{x}^{\circ}\right) \times L\left(\boldsymbol{\theta} ; \boldsymbol{x}^{\circ}\right)
$$

## Inference of $\theta$

"If the missing mechanism is ignorable then likelihood-based inferences for $\boldsymbol{\theta}$ from $L\left(\boldsymbol{\theta} ; \boldsymbol{x}^{\circ}\right)$ will be the same as likelihood based inference for $\boldsymbol{\theta}$ from $L\left(\boldsymbol{\theta}, \psi ; \boldsymbol{x}^{\circ}, \boldsymbol{c}\right) . "{ }^{\text {a }}$
${ }^{\text {a Statistical Analysis With Missing Data. Roderick J. A. Little and Donald B. Rubin. New York: John Wiley \& }}$ Sons, 1987, Section 6.2.

■ M(C)AR is ignorable $\longrightarrow$ is the case in Part I

- MNAR is not ignorable $\longrightarrow$ will be the case in Part III

Likelihood-based inference concerns both frequentist and Bayesian paradigms

## Outline

1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling

5 Frequentist estimation

6 Bayesian estimation

7 Model selection

## Observed-data log-likelihood estimation of $\boldsymbol{\theta}$

■ Principle: MLE (Maximum Likelihood Estimate)

$$
\hat{\boldsymbol{\theta}}=\arg \max _{\boldsymbol{\theta} \in \Theta} \ell\left(\boldsymbol{\theta} ; \boldsymbol{x}^{\circ}\right)
$$

where the observed-data log-likelihood is defined by

$$
\ell\left(\boldsymbol{\theta} ; \boldsymbol{x}^{\circ}\right)=\ln \mathrm{p}\left(\boldsymbol{x}^{\circ} ; \boldsymbol{\theta}\right)
$$

■ Properties: we have, with $\boldsymbol{\theta}^{*}=\arg \min _{\boldsymbol{\theta} \in \Theta} \mathrm{KL}\left(\mathrm{p}, \mathrm{p}_{\boldsymbol{\theta}}\right)^{11}$

$$
\hat{\boldsymbol{\theta}} \xrightarrow{\text { a.s. }} \boldsymbol{\theta}^{*} \quad \text { and } \quad \sqrt{n}\left(\hat{\boldsymbol{\theta}}-\boldsymbol{\theta}^{*}\right) \xrightarrow{d} \mathrm{~N}_{\nu}\left(\mathbf{0}, \mathbf{J}^{-1} \mathbf{K J}^{-1}\right)
$$

where $\nu$ is the number of free continuous parameters in $\Theta$, and

$$
\begin{aligned}
& \mathbf{J}=-\mathrm{E}_{X_{1}} \nabla^{2} \ln \mathrm{p}\left(X_{1} ; \boldsymbol{\theta}^{*}\right) \\
& \mathbf{K}=\operatorname{Var}_{X_{1}} \nabla \ln \mathrm{p}\left(X_{1} ; \boldsymbol{\theta}^{*}\right)
\end{aligned}
$$

- Algorithm: EM and variants (see later)

[^8]
## Expectation-Maximization (EM) algorithm

We need to defined first the complete-data log-likelihood as

$$
\ell_{c}(\boldsymbol{\theta} ; \boldsymbol{x})=\mathrm{p}(\boldsymbol{x} ; \boldsymbol{\theta})
$$

- Initialisation: $\boldsymbol{\theta}^{(0)}$
- Iteration (q):

■ E-step: expectation of the complete-data log-likelihood

$$
Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(q)}\right)=\mathrm{E}\left[\ell_{c}(\boldsymbol{\theta} ; \boldsymbol{x}) \mid \boldsymbol{x}^{o} ; \boldsymbol{\theta}^{(q)}\right]
$$

- M-step: maximization of $Q$ over $\boldsymbol{\theta}$

$$
\boldsymbol{\theta}^{(q+1)}=\arg \max _{\boldsymbol{\theta} \in \Theta} Q\left(\boldsymbol{\theta}, \boldsymbol{\theta}^{(q)}\right)
$$

- Stopping rule: iteration number $Q$ or criterion stationarity


## Properties

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


## Tutorial and practical work about EM (1/5)

- Theoretical part ( $1 / 2$ ): express the MLE of $\mu$ in the case where
- $\boldsymbol{x}=\left(x_{1}, \ldots, x_{n}\right) \sim \mathcal{N}(\mu, 1)$ with $\mu=1$ to be estimated (variance is known)
- $x^{\circ}=\left(x_{1}, \ldots, x_{n}\right.$ o $)$ is the known sample of size $n^{\circ}$
- $x^{m}=\left(x_{n}{ }^{\circ}+1, \ldots, x_{n}{ }^{0}+n^{m}\right)$ is the known sample of size $n^{m}$
- Theoretical part (2/2): express EM steps in the previous case
- Practical part: write an R script
- Implementing the previous EM algorithm
- Check on the same figure both the evolution of $\mu^{(q)}$ values over the iteration number and the MLE value $\hat{\mu}$
- Check on another figure the evolution of $\ell\left(\mu^{(q)} ; \boldsymbol{x}^{\circ}\right)$ over $q$


## Tutorial and practical work about EM $(2 / 5)$

- MLE:

$$
\begin{aligned}
\ell\left(\hat{\mu} ; \boldsymbol{x}^{\circ}\right) & =-\frac{n^{\circ}}{2} \ln (2 \pi)-\frac{1}{2} \sum_{i=1}^{n^{\circ}}\left(x_{i}-\hat{\mu}\right)^{2} \\
\frac{\partial \ell\left(\hat{\mu}^{-} ; \boldsymbol{x}^{\circ}\right)}{\partial \hat{\mu}^{-}}=0 & \Leftrightarrow \hat{\mu}=\overline{\boldsymbol{x}}^{\circ}
\end{aligned}
$$

Remark
Since MLE has a closed-form, the EM algorithm is dummy!

## Tutorial and practical work about EM (3/5)

- E-step of EM:

$$
\begin{aligned}
Q\left(\mu^{+}, \mu^{-}\right) & =\mathrm{E}\left[\left.-\frac{n}{2} \ln (2 \pi)-\frac{1}{2} \sum_{i=1}^{n}\left(x_{i}-\mu^{+}\right)^{2} \right\rvert\, \boldsymbol{x}^{\circ} ; \mu^{-}\right] \\
& =-\frac{n}{2} \ln (2 \pi)-\frac{1}{2} \sum_{i=1}^{n^{\circ}}\left(x_{i}-\mu^{+}\right)^{2}-\frac{1}{2} \mathrm{E}\left[\sum_{i=n^{\circ}+1}^{n}\left(x_{i}-\mu^{+}\right)^{2} \mid \boldsymbol{x}^{\circ} ; \mu^{-}\right]
\end{aligned}
$$

- M-step of EM:

$$
\begin{aligned}
\frac{\partial Q\left(\mu^{+}, \mu^{-}\right)}{\partial \mu^{+}}=0 & \Leftrightarrow-\frac{1}{2} \mathrm{E}\left[\left.\partial \frac{\sum_{i=1}^{n}\left(x_{i}-\mu^{+}\right)^{2}}{\partial \mu^{+}} \right\rvert\, \boldsymbol{x}^{\circ} ; \mu^{-}\right]=0 \\
& \Leftrightarrow \sum_{i=1}^{n^{\circ}}\left(x_{i}-\mu^{+}\right)+n^{m} \mathrm{E}\left[\left(x-\mu^{+}\right) ; \mu^{-}\right]=0 \\
& \Leftrightarrow \mu^{+}=\frac{n^{\circ} \overline{\boldsymbol{x}}^{\circ}+n^{m} \mu^{-}}{n}
\end{aligned}
$$

## Tutorial and practical work about EM (4/5)

```
# data drawing
n.obs = 50
x.obs = rnorm(n.obs)
n.mis = 50
# closed-form ML
mu.ML = mean(x.obs)
# EM
n.it = 20
mu.EM = vector(mode = "numeric", length = n.it)
mu.EM[1] = rnorm(1) # starting value
L.EM = vector(mode = "numeric", length = n.it)
L.EM[1] = sum(dnorm(x.obs, mean = mu.EM[1], sd = 1, log = TRUE))
for (i.it in 2:n.it){
    mu.EM[i.it] = (sum(x.obs) + n.mis*mu.EM[i.it-1]) / (n.obs+n.mis)
    L.EM[i.it] = sum(dnorm(x.obs, mean = mu.EM[i.it-1], sd = 1, log = TRUE))
}
plot(mu.EM)
lines(1:n.it,rep(mu.ML,n.it))
# draw L
plot(L.EM)
```


## Tutorial and practical work about EM (5/5)



## Stochastic EM (SEM) algorithm

The purpose is still to maximize $\ell$ over $\boldsymbol{\theta}$

- Initialisation: $\boldsymbol{\theta}^{(0)}$
- Iteration (q):
- SE-step: draw missing values

$$
\boldsymbol{x}^{m(q)} \sim \mathrm{p}\left(\cdot \mid \boldsymbol{x}^{\circ} ; \boldsymbol{\theta}^{(q)}\right)
$$

- M-step: maximize $\ell_{c}$

$$
\boldsymbol{\theta}^{(q+1)}=\arg \max _{\boldsymbol{\theta} \in \Theta} \ell_{c}\left(\boldsymbol{x}^{\circ}, \boldsymbol{x}^{m(q)} ; \boldsymbol{\theta}\right)
$$

- Stopping rule: iteration number $Q$


## Properties

- The mean of the sequence $\left(\boldsymbol{\theta}^{(q)}\right)$ approximates $\hat{\boldsymbol{\theta}}$ :

$$
\hat{\boldsymbol{\theta}} \simeq \frac{1}{Q-Q^{\text {burn-in }}+1} \sum_{q=Q^{\text {burn-in }}}^{Q} \boldsymbol{\theta}^{(q)}
$$

where $Q^{\text {burn-in }}$ is a so-called burn-in period to reach a stationary regime

- $\oplus$ : often easier to express than EM, avoids local maxima, the standard deviation of the sequence ( $\boldsymbol{\theta}^{(q)}$ ) produces a confidence interval
- $\ominus$ : no punctual convergence (does not improve $\ell$ at each iteration), can be slower than EM


## Tutorial and practical work about SEM (1/4)

Same exercise as EM but now with SEM.

## Tutorial and practical work about SEM (2/4)

■ SE-step:

$$
\begin{aligned}
\boldsymbol{x}^{m+} & \sim \mathrm{p}\left(\cdot \mid \boldsymbol{x}^{o} ; \boldsymbol{\theta}^{-}\right) \\
& =\mathrm{p}\left(\mid ; \boldsymbol{\theta}^{-}\right) \\
& =\mathcal{N}_{n^{m}}\left(\mu^{-} \mathbf{1}, \mathbf{I}\right)
\end{aligned}
$$

■ M-step: we have just to maximise the complete-data log-likeliood

$$
\ell_{c}\left(\mu^{+} ; \boldsymbol{x}^{o}, \boldsymbol{x}^{m+}\right)=-\frac{n}{2} \ln (2 \pi)-\frac{1}{2} \sum_{i=1}^{n^{o}}\left(x_{i}^{o}-\hat{\mu}^{+}\right)^{2}-\frac{1}{2} \sum_{i=1}^{n^{m}}\left(x_{i}^{m+}-\mu^{+}\right)^{2}
$$

thus leading to the standard result

$$
\mu^{+}=\frac{n^{o} \overline{\boldsymbol{x}}^{o}+n^{m} \overline{\boldsymbol{x}}^{m+}}{n}
$$

## Tutorial and practical work about SEM (3/4)

```
# data drawing
n.obs = 50
x.obs = rnorm(n.obs)
n.mis = 50
# closed-form ML
mu.ML = mean(x.obs)
# SEM
n.it = 500
mu.SEM = vector(mode = "numeric", length = n.it)
mu.SEM[1] = rnorm(1) # starting value
L.SEM = vector(mode = "numeric", length = n.it)
L.SEM[1] = sum(dnorm(x.obs, mean = mu.EM[1], sd = 1, log = TRUE))
for (i.it in 2:n.it){
    x.mis = rnorm(n.obs,mu.SEM[i.it-1],1)
    mu.SEM[i.it] = (sum(x.obs) + sum(x.mis)) / (n.obs+n.mis)
    L.SEM[i.it] = sum(dnorm(x.obs, mean = mu.SEM[i.it-1], sd = 1, log = TRUE))
}
plot(mu.SEM)
lines(1:n.it,rep(mu.ML,n.it))
# mu.SEM.final
n.burn = 10
mu.SEM.final = mean(mu.SEM[(n.burn+1):n.it])
# draw L
plot(L.SEM)
```


## Tutorial and practical work about SEM (4/4)



## Outline

1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling
[5 Frequentist estimation

6 Bayesian estimation

7 Model selection

## Bayesian estimation of $\boldsymbol{\theta}$

■ Preliminary remark: within the Bayesian paradigm, the parameter $\boldsymbol{\theta}$ is considered as a random vector thus it is classical to note $\mathrm{p}(\boldsymbol{x} ; \boldsymbol{\theta})=\mathrm{p}(\boldsymbol{x} \mid \boldsymbol{\theta})$

- Principle ${ }^{12}$ :

$$
\hat{\boldsymbol{\theta}}^{\mathrm{Bayes}}=\mathrm{E}\left[\boldsymbol{\theta} \mid \boldsymbol{x}^{\circ}\right]=\int_{\theta} \boldsymbol{\theta} \mathrm{p}\left(\boldsymbol{\theta} \mid \boldsymbol{x}^{\circ}\right) d \boldsymbol{\theta}
$$

where

- $\mathrm{p}\left(\boldsymbol{\theta} \mid \boldsymbol{x}^{\circ}\right)=\ell\left(\boldsymbol{\theta} ; \boldsymbol{x}^{\circ}\right) \mathrm{p}(\boldsymbol{\theta}) / \mathrm{p}\left(\boldsymbol{x}^{\circ}\right)$ is the posterior distribution of $\boldsymbol{\theta}$
- $\mathrm{p}(\boldsymbol{\theta})$ is the prior distribution of $\boldsymbol{\theta}$
- Properties:
- as for the MLE, asymptotic consistency and normality of $\hat{\boldsymbol{\theta}}^{\text {Bayes }}$
- need to define a prior distribution...
- posterior often not so easy to compute thus requires specific algorithms
- Algorithms: Gibbs, Metropolis-Hastings...

[^9]
## Gibbs algorithm

The purpose is to generate a sequence $\boldsymbol{\theta}^{\left(Q^{\text {burn-in }}\right)}, \ldots, \boldsymbol{\theta}^{(Q)}$ drawn from $\mathrm{p}\left(\boldsymbol{\theta} \mid \boldsymbol{x}^{\circ}\right)$

- Initialisation: $\boldsymbol{\theta}^{(0)}$
- Iteration (q):
- Draw $\boldsymbol{x}^{m(q)}$ :

$$
\boldsymbol{x}^{m(q)} \sim \mathrm{p}\left(\cdot \mid \boldsymbol{x}^{0}, \boldsymbol{\theta}^{(q)}\right)
$$

- Draw $\theta^{m(q)}$

$$
\boldsymbol{\theta}^{(q+1)} \sim \mathrm{p}\left(\cdot \mid \boldsymbol{x}^{0}, \boldsymbol{x}^{m(q)}\right)
$$

■ Stopping rule: iteration number $Q$

## Properties

- The mean of the sequence $\left(\boldsymbol{\theta}^{(q)}\right)$ approximates $\hat{\boldsymbol{\theta}}^{\text {Bayes }}$ :

$$
\hat{\boldsymbol{\theta}}^{\text {Bayes }} \simeq \frac{1}{Q-Q^{\text {burn-in }}+1} \sum_{q=Q^{\text {burn-in }}}^{Q} \boldsymbol{\theta}^{(q)}
$$

- Very similar to the SEM algorithm


## Tutorial and practical work about Gibbs (1/4)

- Same exercise as EM and SEM but now with Gibbs.
- Take the prior $\mu \sim \mathcal{N}(0,1)$.


## Tutorial and practical work about Gibbs (2/4)

- Draw $\boldsymbol{x}^{m+}$ : identical to the SE-step of SEM
- Draw $\mu^{+}$:

$$
\begin{aligned}
\mu^{+} & \sim \mathrm{p}\left(\cdot \mid x^{\circ}, \boldsymbol{x}^{m+}\right) \\
& =\mathcal{N}\left(\frac{n^{\circ} \overline{\boldsymbol{x}}^{\circ}+n^{m} \overline{\boldsymbol{x}}^{m+}}{n+1}, 1\right)
\end{aligned}
$$

after using standard Gaussian results applied to the following expression

$$
\mathrm{p}\left(\mu^{+} \mid x^{o}, \boldsymbol{x}^{m+}\right)=\frac{\mathrm{p}\left(x^{o}, \boldsymbol{x}^{m+} \mid \mu^{+}\right) \mathrm{p}\left(\mu^{+}\right)}{\mathrm{p}\left(x^{o}, \boldsymbol{x}^{m+}\right)}
$$

## Tutorial and practical work about Gibbs (3/4)

```
# data drawing
n.obs = 50
x.obs = rnorm(n.obs)
n.mis = 50
# closed-form ML
mu.ML = mean(x.obs)
# Gibbs
n.it = 500
mu.gibbs = vector(mode = "numeric", length = n.it)
mu.gibbs[1] = rnorm(1) # starting value
for (i.it in 2:n.it){
    x.mis = rnorm(n.obs,mu.gibbs[i.it-1],1)
    mu.gibbs[i.it] = (sum(x.obs) + sum(x.mis)) / (n.obs+n.mis+1)
}
plot(mu.gibbs)
lines(1:n.it,rep(mu.ML,n.it))
# mu.gibbs.final
n.burn = 10
mu.gibbs.final = mean(mu.gibbs[(n.burn+1):n.it])
# draw L
plot(L.SEM)
```

Tutorial and practical work about Gibbs (4/4)



## Outline

1 Data

2 Missing data

3 Embedding interest through an example

4 "Full" modeling
[5] Frequentist estimation

6 Bayesian estimation

7 Model selection

## Model definition

- Model $\mathbf{m}$ : it corresponds to a family of distributions $\mathrm{p}(\cdot ; \boldsymbol{\theta})$

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

■ Dimension $\nu$ : it corresponds to the number of free continuous parameters

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

Interest of choosing a model

- A too simple model involves bias modeling
- A too complex model involves variance of estimation


## Importance of model selection: example in supervised learning

Model $\mathbf{m}=$ parametric structure of the borderline shape between too groups

## Too simple model: bias


true model: estimated spherical borderline too simple model: estimated linear borderline

## Too complex model: variance


$\square$ true borderline
_ - - estimated linear borderline
estimated quadratic borderline

## Integrated likelihood

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

$$
\mathrm{p}\left(\mathbf{m} \mid \boldsymbol{x}^{\circ}\right) \propto \mathrm{p}\left(\boldsymbol{x}^{\circ} \mid \mathbf{m}\right) \underbrace{\mathrm{p}(\mathbf{m})}_{\text {prior on } \mathbf{m}}
$$

- Ideal model in a Bayesian context: with $\mathcal{M}$ a family of competing models

$$
\hat{\mathbf{m}}^{*} \in \arg \max _{\mathbf{m} \in \mathcal{M}} \mathrm{p}\left(\mathbf{m} \mid \boldsymbol{x}^{0}\right)
$$

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

$$
\mathrm{p}\left(\boldsymbol{x}^{o} \mid \mathbf{m}\right)=\int_{\Theta} \mathrm{p}\left(\boldsymbol{x}^{o} ; \boldsymbol{\theta}, \mathbf{m}\right) \underbrace{\mathrm{p}(\boldsymbol{\theta} \mid \mathbf{m})}_{\text {prior on } \boldsymbol{\theta}} d \boldsymbol{\theta}
$$

- 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

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

with $\hat{\boldsymbol{\theta}}_{\mathbf{m}}$ the MLE and $\nu_{\mathbf{m}}$ the number of parameters associated to $\mathbf{m}$

- BIC criterion (Bayesian Information Criterion): retain maximizing

$$
\mathrm{BIC}_{\mathbf{m}}=\ell\left(\hat{\boldsymbol{\theta}}_{\mathbf{m}} ; \boldsymbol{x}^{o}\right)-\frac{\nu_{\mathbf{m}}}{2} \ln (n)
$$

## BIC criterion: consistency

■ Consistency: BIC asymptotically selects

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

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

## Tutorial and practical work about BIC (1/4)

- Let the sample $\boldsymbol{x}=\left(x_{1}, \ldots, x_{n}\right)$ where $x_{1}, \ldots, x_{n} \stackrel{i i d}{\sim} \mathcal{N}(0,1)$
- Let the uncertain sample $\boldsymbol{x}^{\delta}=\left(\left[x_{1}-\delta, x_{1}+\delta\right], \ldots,\left[x_{n}-\delta, x_{n}+\delta\right]\right)$

■ Let two competing models: $\mathbf{m}_{1}=\{\mathcal{N}(\mu, 1): \mu \in \mathbb{R}\}$ and $\mathbf{m}_{2}=\mathcal{N}(1,1)$

- Compute BIC with $\mathbf{m}_{1}$ and with $\mathbf{m}_{2}$ associated to the sample $\boldsymbol{x}^{\delta}$

■ Vary $\delta$ and see its effect on the model selection by BIC. Conclusion?

## Tutorial and practical work about BIC (1/4)

Denoting by $\Phi(\cdot ; \mu)$ the cumulative distribution of $\mathcal{N}(\mu, 1)$, we have

$$
\begin{aligned}
\mathrm{BIC}^{\delta} & =\ell\left(\hat{\mu}^{\mathrm{MLE}} ; \boldsymbol{x}^{\delta}\right)-\frac{\nu}{2} \ln n \\
& =\sum_{i=1}^{n} \ln \left(\Phi\left(x_{i}+\delta ; \hat{\mu}^{\mathrm{MLE}}\right)-\Phi\left(x_{i}-\delta ; \hat{\mu}^{\mathrm{MLE}}\right)\right)-\frac{\nu}{2} \ln n
\end{aligned}
$$

Then, use
■ $\nu=1$ and $\hat{\mu}=\hat{\mu}_{1}^{\mathrm{MLE}}$ for $\mathbf{m}_{1}$
■ $\nu=0$ and $\mu=1$ for $\mathbf{m}_{2}$
Since the MLE $\hat{\mu}_{1}^{\text {MLE }}$ is not so easy to obtain, just graphically display the $\mathrm{BIC}^{\delta}$ value for different values of $\mu$ on the same figure

## Tutorial and practical work about BIC (3/4)

```
# data drawing
n = 10
x = rnorm(n)
# paramerization
delta = 0.1
mu = seq(-2,2,0.1)
n.mu = length(mu)
# M1 : N(mu,1)
L1 = vector(mode = "numeric", length = n.mu)
for (i.mu in 1:n.mu){
    L1[i.mu] = sum(log(pnorm(x+delta,mu[i.mu],1)- pnorm(x-delta,mu[i.mu],1)))
}
BIC1 = L1 - 0.5*log(n)
# M2 : N(1,1)
L2 = vector(mode = "numeric", length = n.mu)
for (i.mu in 1:n.mu){
    L2[i.mu] = sum(log(pnorm(x+delta,1,1)- pnorm(x-delta,1,1)))
}
BIC2 = L2
# figure
plot(mu,BIC1)
lines(mu,BIC2)
```


## Tutorial and practical work about BIC (4/4)



## End of Part I


[^0]:    ${ }^{1}$ S. Alelyani, J. Tang and H. Liu (2013). Feature Selection for Clustering: A Review. Data Clustering: Algorithms and Applications, 29

[^1]:    ${ }^{2}$ S. Alelyani, J. Tang and H. Liu (2013). Feature Selection for Clustering: A Review. Data Clustering: Algorithms and Applications, 29

[^2]:    ${ }^{3}$ FLOP $=$ FLoating-point Operations Per Second

[^3]:    ${ }^{4}$ Exabyte

[^4]:    ${ }^{6}$ See others lessons for classical methods of imputation

[^5]:    ${ }^{7}$ Byar DP, Green SB (1980): Bulletin Cancer, Paris 67:477-488

[^6]:    ${ }^{8}$ See other lessons: http://cran.r-project.org/web/packages/mice/mice.pdf
    ${ }^{9}$ See Part II: https://cran.r-project.org/web/packages/RMixtComp/index.html

[^7]:    Some key words: Bayesian inference; Incomplete data; Likelihood inference; Missing at random; Missing data; Missing values; Observed at random; Sampling distribution inference.

[^8]:    ${ }^{11} \mathrm{p}$ is the (unknown) true distribution, $\mathrm{p}_{\boldsymbol{\theta}}=\mathrm{p}(\cdot ; \boldsymbol{\theta})$ and KL is the Kullback-Leibler divergence.

[^9]:    ${ }^{12}$ Other definitions are possible.

