## An overview of methods to deal with missing

 valuesJulie Josse
CUSO doctoral school in Statistics and Applied Probability
Les Diableret, February 8-9 2022


## Presentation Julie Josse: Stat and ML for bio-sciences

## Academic background:

- Engineer and Assistant Professor in Agronomy University (2007-2015)
- Visiting Researcher + Teaching at Stanford University (18 months)
- Professor at Ecole Polytechnique (IP Paris) (2016-2020). Still Teaching
- Visiting Researcher at Google Brain Paris (2019-2020). Still Collaborating
- Senior Researcher at Inria Montpellier (Sept. 2020-)


## Research topics:

- Dimensionality reduction to visualize high dimensional heterogeneous data
- Missing values: supervised learning, inference, matrix completion, MNAR
- Causal inference: estimating treatment effect, combining RCT and observational data, personalized recommendation
- Medical collaborations: Traumabase, IGR, CHU Nancy, Curie, etc.


## Implementations - transfert:

- $R$ community: book $R$ for Statistics, $R$ foundation, $R$ Forwards (widen the participation of minorities), R packages
- Rmisstastic https://rmisstastic.netlify.app/


## Outline

- Lecture 1: Introduction
- Single imputation, Multiple imputation
- Likelihood approaches
- Lecture 2: Low rank methods
- PCA with missing values - (Multiple) Imputation with PCA
- Practice
- MNAR data
- Heterogeneous data
- Lecture 3:
- Supervised learning with missing values
- Random Forest with missing values
- Linear regression with missing values
- Causal inference with missing values


## References



## Outline

## 1. Introduction

2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Missing values

are everywhere: unanswered questions in a survey, lost data, damaged plants, machines that fail...
"The best thing to do with missing values is not to have any"
$\Rightarrow$ Still an issue in the "big data" area


Data integration: data from different sources

## Traumabase

- 30000 patients
- 250 continuous and categorical variables: heterogeneous
- 20 hospitals
- 4000 new patients/ year

| Center | Accident | Age | Sex | Weight | Lactactes | BP | shock | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Beaujon | fall | 54 | m | 85 | NM | 180 | yes |  |
| Pitie | gun | 26 | m | NR | NA | 131 | no |  |
| Beaujon | moto | 63 | m | 80 | 3.9 | 145 | yes |  |
| Pitie | moto | 30 | w | NR | Imp | 107 | no |  |
| HEGP | knife | 16 | m | 98 | 2.5 | 118 | no |  |
| $\vdots$ |  |  |  |  |  |  |  | $\ddots$ |

[^0]
## Traumabase

- 30000 patients
- 250 continuous and categorical variables: heterogeneous
- 20 hospitals
- 4000 new patients/ year

| Center | Accident | Age | Sex | Weight | Lactactes | BP | shock | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Beaujon | fall | 54 | m | 85 | NM | 180 | yes |  |
| Pitie | gun | 26 | m | NR | NA | 131 | no |  |
| Beaujon | moto | 63 | m | 80 | 3.9 | 145 | yes |  |
| Pitie | moto | 30 | w | NR | Imp | 107 | no |  |
| HEGP | knife | 16 | m | 98 | 2.5 | 118 | no |  |
| $\vdots$ |  |  |  |  |  |  |  | $\ddots$ |

$\Rightarrow$ Estimate causal effect: Administration of the treatment
"tranexamic acid" on the outcome mortality for trauma brain patients.
Causal Inference (IPW) with covariates with missing values ${ }^{1}$

[^1]
## Traumabase

- 30000 patients
- 250 continuous and categorical variables: heterogeneous
- 20 hospitals
- 4000 new patients/ year

| Center | Accident | Age | Sex | Weight | Lactactes | BP | shock | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Beaujon | fall | 54 | m | 85 | NM | 180 | yes |  |
| Pitie | gun | 26 | m | NR | NA | 131 | no |  |
| Beaujon | moto | 63 | m | 80 | 3.9 | 145 | yes |  |
| Pitie | moto | 30 | w | NR | Imp | 107 | no |  |
| HEGP | knife | 16 | m | 98 | 2.5 | 118 | no |  |
| $\vdots$ |  |  |  |  |  |  |  | $\ddots$ |

$\Rightarrow$ Explain and Predict platelet levels, hemorrhagic shock given pre-hospital features

Ex linear, logistic regression/ random forests with covariates with missing values

## Missing values

Percentage of missing values


Different types of missing values
Multilevel data/ data integration: Systematic missing variable in one hospital

## Contingency tables with side information

- National agency for wildlife and hunting management (ONCFS) data
- Contingency tables: Water (722 wetland sites) - bird (species) count data, from 1990-2016 in 5 countries in North Africa
- Additional sites \& years info: meteo, geographical (altitude, etc.)

$\Rightarrow$ Aims: Assess the effect of time on species abundances;
Monitor the population and assess wetlands conservation policies.
$\Rightarrow 70 \%$ of missing values in contingency tables ${ }^{2} 3$

[^2]
## Multi-blocks data set



L'OREAL data: 100000 women in many countries - 300 questions in groups:

- Self-assessment questionnaire: life style, skin and hair characteristics, care and consumer habits
- Clinical assessments by a dermatologist: facial skin complexion, wrinkles, scalp dryness, greasiness
- Hair assessments by a hair dresser: abundance, volume, breakage, curly
- Skin and Hair photographs and measurements: sebum quantity, etc.
$\Rightarrow$ Aim: Clustering women for marketing targeting
$\Rightarrow$ Missing values structured by group ${ }^{45}$
${ }^{4}$ Handling missing values in exploratory multivariate data analysis. J., Husson. JSFDS 2012.
${ }^{5}$ Handling missing values in Multiple Factor Analysis. J., Husson. FQP 2013.


## Complete-case analysis



Deleting rows with missing values?
?lm, ?glm, na.action = na.omit
"One of the ironies of Big Data is that missing data play an ever more significant role" (R. Sameworth, 2019)

An $n \times p$ matrix, each entry is missing with probability 0.01
$p=5 \quad \Longrightarrow \approx 95 \%$ of rows kept
$p=300 \Longrightarrow \approx 5 \%$ of rows kept

## Distribution of missing values

## Missing values mechanisms taxonomy Rubin, 1976



MCAR


MNAR

Orange: missing values for Systolic Blood Pressure - Gravity index (GCS) is always observed

MCAR (completely at random): Proba to be missing does not depend on SBP neither on gravity
MAR: Proba depends on gravity (we do not measure for too severe patients) MNAR (not at random): Proba depends on SBP (low SBP not measured)

## Missing values mechanisms

- Random Variables:
- $X \in \mathbb{R}^{d}$ : the complete unvailable data
- $\widetilde{X} \in\{\mathbb{R} \cup\{\mathrm{NA}\}\}^{p}$ : incomplete data (available)
- $M \in\{0,1\}^{d}$ : the missing-data pattern obs $(M)$ (resp. mis $(M)$ ) indices of the observed (resp. missing) entries.
- Realizations:

$$
\begin{aligned}
& x=(1.1,2.3,3.1,8,5.27) \\
& \widetilde{x}=(1.1, \mathrm{NA},-3.1,8, \mathrm{NA}) \\
& m=(0,1,0,0,1) \\
& x_{\mathrm{obs}(\mathrm{~m})}=(1.1,3.1,8), \quad x_{\operatorname{mis}(m)}=(2.3,5.27)
\end{aligned}
$$

MCAR: For all $m \in\{0,1\}^{d}, P(M=m \mid X)=P(M=m)$
$\operatorname{MAR}^{6}$ : For all $m \in\{0,1\}^{d}, P(M=m \mid X)=P\left(M=m \mid X_{o b s(m)}\right)$

## Visualization

The first thing to do with missing values (as for any analysis) is descriptive statistics: Visualization of patterns to get hints on how and why they occur VIM (M. Templ), naniar (N. Tierney), FactoMineR (Husson et al.)


Right: PAS_m close to PAD_m: Often missing on both PAS \& PAD IOT: nested questions. Q1: yes/no, if yes Q2-Q4, if no Q2-Q4 "missing"

Note: Crucial before starting any treatment of missing values and after

## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Collaborators on inference/imputation with missing values

- W. Jiang, A. Sportisse, former PhD student at Polytechnique
- F. Husson, Professor Agronomy University (package missMDA, FactoMineR)
- G. Bogdan, Professor Wroclaw. C. Boyer, Associate Professor Sorbonne
- Traumabase project: J.P. Nadal, T. Gauss, S. Hamada


Logistic Regression with Missing Covariates - Parameter Estimation, Model Selection and Prediction within a Joint-Modeling Framework. (2019). CSDA Adaptive Bayesian SLOPE - High dimensional Model Selection with Missing Values. (2020). JCGS.

Estimation and Imputation in Probabilistic Principal Component Analysis with Missing Not At Random data. Neurips2020.

Missing Data Imputation using Optimal Transport. ICML2020.
Debiasing Stochastic Gradient Descent to handle missing values. Neurips2020.

## Solutions to handle missing values (M(C)AR)

Books: Schafer (2002), Little \& Rubin (2019), Kim \& Shao (2013), Carpenter \& Kenward (2013), van Buuren (2018), etc.

## Modify the estimation process to deal with missing values

Maximum likelihood: EM algorithm to obtain point estimates + Supplemented EM (Meng \& Rubin, 1991) / Louis formulae for their variability Ex logistic regression: EM to get $\hat{\beta}+$ Louis to get $\hat{V}(\hat{\beta})$

Aim: Estimate parameters \& their variance from an incomplete data
$\Rightarrow$ Inferential framework

## Solutions to handle missing values (M(C)AR)

Books: Schafer (2002), Little \& Rubin (2019), Kim \& Shao (2013), Carpenter \& Kenward (2013), van Buuren (2018), etc.

## Modify the estimation process to deal with missing values

Maximum likelihood: EM algorithm to obtain point estimates + Supplemented EM (Meng \& Rubin, 1991) / Louis formulae for their variability Ex logistic regression: EM to get $\hat{\beta}+$ Louis to get $\hat{V}(\hat{\beta})$

Cons: Difficult to establish - not many softwares even for simple models One specific algorithm for each statistical method...

Aim: Estimate parameters \& their variance from an incomplete data
$\Rightarrow$ Inferential framework

## Solutions to handle missing values (M(C)AR)

Books: Schafer (2002), Little \& Rubin (2019), Kim \& Shao (2013), Carpenter \& Kenward (2013), van Buuren (2018), etc.

## Modify the estimation process to deal with missing values

Maximum likelihood: EM algorithm to obtain point estimates +
Supplemented EM (Meng \& Rubin, 1991) / Louis formulae for their variability
Ex logistic regression: EM to get $\hat{\beta}+$ Louis to get $\hat{V}(\hat{\beta})$
Cons: Difficult to establish - not many softwares even for simple models One specific algorithm for each statistical method...

## Imputation (multiple) to get a complete data set

Any analysis can be performed
Ex logistic regression: Impute and apply logistic model to get $\hat{\beta}, \hat{V}(\hat{\beta})$

Aim: Estimate parameters \& their variance from an incomplete data
$\Rightarrow$ Inferential framework

## Mean imputation

- $\left(x_{i}, y_{i}\right) \underset{\text { i.i.d. }}{\sim} \mathcal{N}_{2}\left(\left(\mu_{x}, \mu_{y}\right), \Sigma_{x y}\right)$


$$
\begin{array}{c|c|}
\hline \mu_{y}=0 & \hat{\mu}_{y}=-0.01 \\
\cline { 2 - 3 } \sigma_{y}=1 & \hat{\sigma}_{y}=1.01 \\
\cline { 2 - 2 }=0.6 & \hat{\rho}=0.66 \\
\hline
\end{array}
$$

## Mean imputation

- $\left(x_{i}, y_{i}\right) \underset{\text { i.i.d. }}{\sim} \mathcal{N}_{2}\left(\left(\mu_{x}, \mu_{y}\right), \Sigma_{x y}\right)$
- $70 \%$ of missing entries completely at random on $Y$



## Mean imputation

- $\left(x_{i}, y_{i}\right)_{\text {i.i.d. }}^{\sim} \mathcal{N}_{2}\left(\left(\mu_{x}, \mu_{y}\right), \Sigma_{x y}\right)$
- $70 \%$ of missing entries completely at random on $Y$
- Estimate parameters on the mean imputed data


Mean imputation deforms joint and marginal distributions

## Mean imputation is bad for estimation



Ecological data: ${ }^{7} n=69000$ species -6 traits. Estimated correlation between
Pmass \& Rmass $\approx 0$ (mean imputation) or $\approx 1$ (EM PCA)
${ }^{7}$ Wright, I. et al. (2004). The worldwide leaf economics spectrum. Nature.

## Imputation methods

- by regression takes into account the relationship: Estimate $\beta$-impute $\hat{y}_{i}=\hat{\beta}_{0}+\hat{\beta}_{1} x_{i} \Rightarrow$ variance underestimated and correlation overestimated
- by stochastic reg: Estimate $\beta$ and $\sigma$ - impute from the predictive $\hat{y}_{i} \sim \mathcal{N}\left(x_{i} \hat{\beta}, \hat{\sigma}^{2}\right) \Rightarrow$ preserve distributions

Here $\hat{\beta}, \hat{\sigma}^{2}$ estimated with complete data, but MLE can be obtained with EM

Stochastic regression imputation


$$
\begin{array}{c|c|}
\mu_{y}=0 \\
\sigma_{y}=1 \\
\rho_{x y}=0.6
\end{array} \quad 0.01
$$

| 0.01 |
| :--- |
| 0.72 |
| 0.78 |


| 0.01 |
| :--- |
| 0.99 |
| 0.59 |

## Imputation with joint model with gaussian distribution

$\Rightarrow$ Assumption joint gaussian model $z_{i}=\left(x_{i}, y_{i}\right), z_{i} \sim \mathcal{N}(\mu, \Sigma)$

- Bivariate case with missing values on $y$ (stochastic regression):
- estimate $\beta$ and $\sigma$
- impute from the predictive $\hat{y}_{i} \sim \mathcal{N}\left(x_{i} \hat{\beta}, \hat{\sigma}^{2}\right)$
- Extension to the multivariate case:
- Estimate $\mu$ and $\Sigma$ from an incomplete data with EM
- Impute by drawing from the conditional distribution

$$
\begin{aligned}
& Z_{\text {mis }} \mid Z_{\text {obs }} \sim \mathcal{N}\left(\mu_{\text {mis } \mid \text { obs }}, \Sigma_{\text {mis } \mid \text { obs }}\right) \\
& \quad \mu_{\text {mis } \mid \text { obs }}=\mathbb{E}\left[X_{\text {mis }}\right]+\Sigma_{\text {mis }, \text { obs }} \Sigma_{\text {obs,obs }}^{-1}\left(X_{\text {obs }}-\mathbb{E}\left[X_{\text {obs }}\right]\right) \\
& \quad \Sigma_{\text {mis } \mid \text { obs }}=\Sigma_{\text {mis }}-\Sigma_{\text {mis }, \text { obs }} \Sigma_{\text {obs }, \text { obs }}^{-1} \Sigma_{\text {obs }, \text { mis }} . \text { Schur complements. }
\end{aligned}
$$

> library (norm)
> pre <- prelim.norm(as.matrix(don))
> thetahat <- em.norm(pre)
> imp <- imp.norm(pre, thetahat, don)

## Imputation methods for multivariate data

## Assuming a joint model

- Gaussian distribution: $z_{i} \sim \mathcal{N}(\mu, \Sigma)$ (Amelia Honaker, King, Blackwell)
- low rank: $Z_{n \times d}=\mu_{n \times d}+\varepsilon \varepsilon_{i j} \stackrel{\text { iid }}{\sim} \mathcal{N}\left(0, \sigma^{2}\right)$ with $\mu$ of low rank $k$ (softimpute Hastie \& Mazuder; missMDA J. \& Husson, mimi ${ }^{8}$ )
- latent class - nonparametric Bayesian (dpmpm Reiter)
- deep learning using variational autoencoders (miwae, Mattei, 2018, VAEAC Ivanov et al., 2019), using GAN (GAIN, Yoon et al. 2018)


## Using conditional models (joint implicitly defined)

- with logistic, multinomial, poisson regressions (mice van Buuren)
- iterative impute each variable by random forests (missForest Stekhoven)

Imputation for categorical, mixed, blocks/multilevel data ${ }^{9}$, etc.
$\Rightarrow$ Rmistatic platform, more than 150 packages ${ }^{10}$

[^3]
## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Single imputation methods: Danger!



| $\mu_{y}=0$ | 0.01 |
| :---: | :---: |
| $\sigma_{y}=1$ | 0.5 |
| $=0.6$ | 0.30 |
| $\mu_{y} 95 \%$ |  |

## Confidence interval for a mean

Let $Y=\left(Y_{1}, \ldots, Y_{n}\right)^{\prime}$ be i.i.d. independent Gaussian random with expectation $\mu_{y}$ and variance $\sigma_{y}^{2}>0$.

- The empirical mean $\bar{Y}=n^{-1} \sum_{i=1}^{n} Y_{i}$
- $\bar{Y} \sim \mathcal{N}\left(\mu_{y}, \sigma_{y}^{2} / n\right)$
- A confidence interval for $\mu$

$$
\mathbb{P}\left(\bar{Y}-\frac{\sigma_{y}}{\sqrt{n}} z_{1-\alpha / 2} \leq \mu \leq \bar{Y}+\frac{\sigma_{y}}{\sqrt{n}} z_{1-\alpha / 2}\right)=1-\alpha
$$

## Confidence interval for a mean

Let $Y=\left(Y_{1}, \ldots, Y_{n}\right)^{\prime}$ be i.i.d. independent Gaussian random with expectation $\mu_{y}$ and variance $\sigma_{y}^{2}>0$.

- The empirical mean $\bar{Y}=n^{-1} \sum_{i=1}^{n} Y_{i}$
- $\bar{Y} \sim \mathcal{N}\left(\mu_{y}, \sigma_{y}^{2} / n\right)$
- A confidence interval for $\mu$

$$
\mathbb{P}\left(\bar{Y}-\frac{\sigma_{y}}{\sqrt{n}} z_{1-\alpha / 2} \leq \mu \leq \bar{Y}+\frac{\sigma_{y}}{\sqrt{n}} z_{1-\alpha / 2}\right)=1-\alpha
$$

Variance unknown:

$$
\begin{gathered}
\frac{\sqrt{n}}{\widehat{\sigma}_{y}}\left(\bar{Y}-\mu_{y}\right) \sim T(n-1) \\
{\left[\bar{y}-\frac{\widehat{\sigma}_{y}}{\sqrt{n}} q t_{1-\alpha / 2}(n-1), \bar{y}+\frac{\hat{\sigma}_{y}}{\sqrt{n}} q t_{1-\alpha / 2}(n-1)\right]}
\end{gathered}
$$

## Simulation

(1) Generate bivariate Gaussian data ( $\mu_{y}=0, \sigma_{y}=1, \rho=0.6$ )
(2) Put missing values on $y$
(3) Imput missing entries
(4) Compute the confidence interval of $\mu_{y}$ - count if the true value $\mu_{y}=0$ is in the confidence interval
(5) Repeat the steps 1-4, 10000 times
$\Rightarrow$ Give the coverage

## Single imputation methods: Danger!

$$
\left[\bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}} ; \bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}}\right]
$$



$$
\begin{array}{|c|c|}
\hline \mu_{y}=0 & 0.01 \\
\sigma_{y}=1 & 0.5 \\
\hline \rho=0.6 & 0.30 \\
\cline { 2 - 3 } & 39 . \mu_{y} 95 \%
\end{array}
$$

The idea of imputation is both seductive and dangerous (Dempster and Rubin, 1983)

## Single imputation methods: Danger!

$$
\left[\bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}} ; \bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}}\right]
$$




| $\mu_{y}=0$ | 0.01 |
| :---: | :---: |
| $\sigma_{y}=1$ | 0.5 |
| $\rho=0.6$ | 0.30 |
| Cl $\mu_{y} 95 \%$ | 39.4 |


| 0.01 |
| :--- |
| 0.72 |
| 0.78 |
| 61.6 |

The idea of imputation is both seductive and dangerous (Dempster and Rubin, 1983)

## Single imputation methods: Danger!

$$
\left[\bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}} ; \bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}}\right]
$$





| $\mu_{y}=0$ | 0.01 |
| :---: | :---: |
| $\sigma_{y}=1$ | 0.5 |
| $=0.6$ | 0.30 |
| Cl $_{y} 95 \%$ | 39.4 |


| 0.01 |
| :--- |
| 0.72 |
| 0.78 |
| 61.6 |


| 0.01 |
| :--- |
| 0.99 |
| 0.59 |
| 70.8 |

The idea of imputation is both seductive and dangerous (Dempster and Rubin, 1983)

## Single imputation methods: Danger!

$$
\left[\bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}} ; \bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}}\right]
$$





| $\mu_{y}=0$ | 0.01 |
| :---: | :---: |
| $\sigma_{y}=1$ | 0.5 |
| $\rho=0.6$ |  |
| Cl $\mu_{y} 95 \%$ | 0.30 |
|  | 39.4 |


| 0.01 |
| :--- |
| 0.72 |
| 0.78 |
| 61.6 |


| 0.01 |
| :--- |
| 0.99 |
| 0.59 |
| 70.8 |

The idea of imputation is both seductive and dangerous (Dempster and Rubin, 1983)
$\Rightarrow$ Standard errors of the parameters $\left(\hat{\sigma}_{\hat{\mu}_{y}}\right)$ calculated from the imputed data

## Underestimation of variance

Classical confidence interval for $\mu_{y}\left[\bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}} ; \bar{y}-q t_{n-1} \frac{\hat{\sigma}_{y}}{\sqrt{n}}\right]$

Asymptotic variance with MCAR values (Little \& Rubin, 2019. p158):

$$
\frac{\hat{\sigma}_{y}^{2}}{n_{o b s}}\left(1-\hat{\rho}^{2} \frac{n-n_{o b s}}{n_{o b s}}\right)=\frac{\hat{\sigma}_{y}^{2}}{n}\left(1+\frac{n-n_{o b s}}{n_{o b s}}\left(1-\hat{\rho}^{2}\right)\right)
$$

$\Rightarrow$ When the $\rho=1$, we trust the prediction and the coverage given by stochastic regression is OK.
$\Rightarrow$ Coverage of single imputation is too low: need to take into account the uncertainty associated to the predictions.

## Single imputation: Underestimation of the variability

$\Rightarrow$ Incomplete Traumabase

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\cdots$ | Y |
| :---: | :---: | :---: | :---: | :---: |
| NA | 20 | 10 | $\cdots$ | shock |
| -6 | 45 | NA | $\cdots$ | shock |
| 0 | NA | 30 | $\cdots$ | no shock |
| NA | 32 | 35 | $\cdots$ | shock |
| -2 | NA | 12 | $\cdots$ | no shock |
| 1 | 63 | 40 | $\cdots$ | shock |

## Single imputation: Underestimation of the variability

$\Rightarrow$ Incomplete Traumabase

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\cdots$ | Y |
| :---: | :---: | :---: | :---: | :---: |
| NA | 20 | 10 | $\cdots$ | shock |
| -6 | 45 | NA | $\cdots$ | shock |
| 0 | NA | 30 | $\cdots$ | no shock |
| NA | 32 | 35 | $\cdots$ | shock |
| -2 | NA | 12 | $\cdots$ | no shock |
| 1 | 63 | 40 | $\cdots$ | shock |

$\Rightarrow$ Completed Traumabase

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | Y |
| :---: | :---: | :---: | :--- | :---: |
| 3 | 20 | 10 | $\ldots$ | shock |
| -6 | 45 | 6 | $\ldots$ | shock |
| 0 | 4 | 30 | $\ldots$ | no shock |
| -4 | 32 | 35 | $\ldots$ | shock |
| -2 | 75 | 12 | $\ldots$ | no shock |
| 1 | 63 | 40 | $\ldots$ | shock |

## Single imputation: Underestimation of the variability

$\Rightarrow$ Incomplete Traumabase

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\cdots$ | Y |
| :---: | :---: | :---: | :---: | :---: |
| NA | 20 | 10 | $\cdots$ | shock |
| -6 | 45 | NA | $\cdots$ | shock |
| 0 | NA | 30 | $\cdots$ | no shock |
| NA | 32 | 35 | $\cdots$ | shock |
| -2 | NA | 12 | $\cdots$ | no shock |
| 1 | 63 | 40 | $\cdots$ | shock |

$\Rightarrow$ Completed Traumabase

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | Y |
| :---: | :---: | :---: | :---: | :---: |
| 3 | 20 | 10 | $\ldots$ | shock |
| -6 | 45 | 6 | $\ldots$ | shock |
| 0 | 4 | 30 | $\ldots$ | no shock |
| -4 | 32 | 35 | $\ldots$ | shock |
| -2 | 75 | 12 | $\ldots$ | no shock |
| 1 | 63 | 40 | $\ldots$ | shock |

A single value can't reflect the uncertainty of prediction
Multiple impute 1) Generate $M$ plausible values for each missing value

| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 3 | 20 | 10 | s |
| -6 | 45 | 6 | s |
| 0 | 4 | 30 | no s |
| -4 | 32 | 35 | s |
| -2 | 75 | 12 | no s |
| 1 | 63 | 40 | s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| -7 | 20 | 10 | s |
| -6 | 45 | 9 | s |
| 0 | 12 | 30 | no s |
| 13 | 32 | 35 | s |
| -2 | 10 | 12 | no s |
| 1 | 63 | 40 | s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 7 | 20 | 10 | s |
| -6 | 45 | 12 | s |
| 0 | -5 | 30 | no s |
| 2 | 32 | 35 | s |
| -2 | 20 | 12 | no s |
| 1 | 63 | 40 | s |

library (mice); mice(traumadata)
library (missMDA); MIPCA(traumadata)

## Multiple imputation

1) Generate $M$ plausible values for each missing value

| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 3 | 20 | 10 | s |
| -6 | 45 | 6 | s |
| 0 | 4 | 30 | no s |
| -4 | 32 | 35 | s |
| 1 | 63 | 40 | s |
| -2 | 15 | 12 | no s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| -7 | 20 | 10 | s |
| -6 | 45 | 9 | s |
| 0 | 12 | 30 | no s |
| 13 | 32 | 35 | s |
| 1 | 63 | 40 | s |
| -2 | 10 | 12 | no s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 7 | 20 | 10 | s |
| -6 | 45 | 12 | s |
| 0 | -5 | 30 | no s |
| 2 | 32 | 35 | s |
| 1 | 63 | 40 | s |
| -2 | 20 | 12 | no s |

2) Perform the analysis on each imputed data set: $\hat{\beta}_{m}, \widehat{\operatorname{Var}}\left(\hat{\beta}_{m}\right)$
3) Combine the results (Rubin's rules):

$$
\begin{aligned}
& \hat{\beta}=\frac{1}{M} \sum_{m=1}^{M} \hat{\beta}_{m} \\
& T=\frac{1}{M} \sum_{m=1}^{M} \widehat{\operatorname{Var}}\left(\hat{\beta}_{m}\right)+\left(1+\frac{1}{M}\right) \frac{1}{M-1} \sum_{m=1}^{M}\left(\hat{\beta}_{m}-\hat{\beta}\right)^{2}
\end{aligned}
$$

imp.mice <- mice(traumadata)
lm.mice.out <- with(imp.mice, glm(Y ~ ., family = "binomial"))
$\Rightarrow$ Variability of missing values taken into account

## Multiple imputation

(1) Generating $M$ imputed data sets

First idea: several stochastic regression
for $m=1, \ldots, M$, draw $\hat{y}_{i}$ from the predictive $\mathcal{N}\left(x_{i} \hat{\beta}, \hat{\sigma}^{2}\right)$
(2) Performing the analysis on each imputed data set
(3) Combining: variance $=$ within + between imputation variance

|  | $M=1$ | $M=50$ |
| :---: | :---: | :---: |
| $\mu_{y}=0$ | 0.01 | 0.01 |
| $\sigma_{y}=1$ | 0.99 | 0.99 |
| $\rho=0.6$ | 0.59 | 0.59 |
| $C l \mu_{y} 95 \%$ | 70.8 | 81.8 |

## Multiple imputation

(1) Generating $M$ imputed data sets

First idea: several stochastic regression
for $m=1, \ldots, M$, draw $\hat{y}_{i}$ from the predictive $\mathcal{N}\left(x_{i} \hat{\beta}, \hat{\sigma}^{2}\right)$
(2) Performing the analysis on each imputed data set
(3) Combining: variance $=$ within + between imputation variance

|  | $M=1$ | $M=50$ |
| :---: | :---: | :---: |
| $\mu_{y}=0$ | 0.01 | 0.01 |
| $\sigma_{y}=1$ | 0.99 | 0.99 |
| $\rho=0.6$ | 0.59 | 0.59 |
| $C l \mu_{y} 95 \%$ | 70.8 | 81.8 |

$\Rightarrow$ Variability of the parameters is missing: "improper" imputation

## Multiple imputation

(1) Generating $M$ imputed data sets

First idea: several stochastic regression
for $m=1, \ldots, M$, draw $\hat{y}_{i}$ from the predictive $\mathcal{N}\left(x_{i} \hat{\beta}, \hat{\sigma}^{2}\right)$
(2) Performing the analysis on each imputed data set
(3) Combining: variance $=$ within + between imputation variance

|  | $M=1$ | $M=50$ |
| :---: | :---: | :---: |
| $\mu_{y}=0$ | 0.01 | 0.01 |
| $\sigma_{y}=1$ | 0.99 | 0.99 |
| $\rho=0.6$ | 0.59 | 0.59 |
| $C l \mu_{y} 95 \%$ | 70.8 | 81.8 |

$\Rightarrow$ Variability of the parameters is missing: "improper" imputation
$\Rightarrow$ Prediction variance $=$ estimation variance plus noise

## Regression: variance of prediction

$$
\begin{aligned}
& y_{n+1}=x_{n+1}^{\prime} \beta+\varepsilon_{n+1} \\
& \hat{y}_{n+1}=x_{n+1}^{\prime} \hat{\beta} \\
& \hat{\beta}=\left(X^{\prime} X\right)^{-1} X^{\prime} Y
\end{aligned}
$$

$$
\begin{aligned}
V\left[\hat{y}_{n+1}-y_{n+1}\right] & =V\left[x_{n+1}^{\prime}(\hat{\beta}-\beta)-\varepsilon_{n+1}\right] \\
& \left.=x_{n+1}^{\prime} V(\hat{\beta}-\beta) x_{n+1}+\sigma^{2}\right] \\
& =\hat{\sigma}^{2}\left(x_{n+1}^{\prime}\left(X^{\prime} X\right)^{-1} x_{n+1}+1\right)
\end{aligned}
$$

Cl for the prediction

$$
\left[x_{n+1}^{\prime} \hat{\beta}+-t_{n-p}(1-\alpha / 2) \hat{\sigma} \sqrt{\left(x_{n+1}^{\prime}\left(X^{\prime} X\right)^{-1} x_{n+1}+1\right)}\right]
$$

## Multiple imputation continuous data: bivariate case

$\Rightarrow$ Proper multiple imputation with $y_{i}=x_{i} \beta+\varepsilon_{i}$
(1) Variability of the parameters, $M$ plausible: $(\hat{\beta})^{1}, \ldots,(\hat{\beta})^{M}$
$\Rightarrow$ Bootstrap
$\Rightarrow$ Posterior distribution: Data Augmentation (Tanner \& Wong, 1987)

2 Noise: for $m=1, \ldots, M$, missing values $\hat{y}_{i}^{m}$ are imputed by drawing from the predictive distribution $\mathcal{N}\left(x_{i} \hat{\beta}^{m},\left(\hat{\sigma}^{2}\right)^{m}\right)$

$$
\begin{array}{ccc} 
& \text { Improper } & \text { Proper } \\
\mathrm{Cl} \mu_{y} 95 \% & 0.818 & 0.935
\end{array}
$$

## Multiple imputation

$\Rightarrow$ Aim: provide estimation of the parameters and of their variability (taken into account the variability due to missing values)

Single imputation: a single value can't reflect the uncertainty of prediction $\Rightarrow$ underestimate the standard errors
(1) Generating $M$ imputed data sets: variance of prediction

(2) Performing the analysis on each imputed data set ${ }^{11}$, ${ }^{12}$
(3) Combining: variance $=$ within + between imputation variance

$$
\hat{\beta}=\frac{1}{M} \sum_{m=1}^{M} \hat{\beta}_{m} T=\frac{1}{M} \sum \widehat{\operatorname{Var}}\left(\hat{\beta}_{m}\right)+\left(1+\frac{1}{M}\right) \frac{1}{M-1} \sum\left(\hat{\beta}_{m}-\hat{\beta}\right)^{2}
$$

[^4]
## Multiple imputation

$\Rightarrow$ Aim: provide estimation of the parameters and of their variability (taken into account the variability due to missing values)

Single imputation: a single value can't reflect the uncertainty of prediction $\Rightarrow$ underestimate the standard errors
(1) Generating $M$ imputed data sets: variance of prediction

"1) Variance of estimation of the parameters +2 ) Noise"
(2) Performing the analysis on each imputed data set ${ }^{11},{ }^{12}$
(3) Combining: variance $=$ within + between imputation variance

$$
\hat{\beta}=\frac{1}{M} \sum_{m=1}^{M} \hat{\beta}_{m} T=\frac{1}{M} \sum \widehat{\operatorname{Var}}\left(\hat{\beta}_{m}\right)+\left(1+\frac{1}{M}\right) \frac{1}{M-1} \sum\left(\hat{\beta}_{m}-\hat{\beta}\right)^{2}
$$

[^5]
## Joint modeling

$\Rightarrow$ Hypothesis $z_{i} \sim \mathcal{N}(\mu, \Sigma)$
Algorithm Expectation Maximization Bootstrap:
(1) Bootstrap rows: $Z^{1}, \ldots, Z^{M}$

EM algorithm: $\left(\hat{\mu}^{1}, \hat{\Sigma}^{1}\right), \ldots,\left(\hat{\mu}^{M}, \hat{\Sigma}^{M}\right)$
(2 Imputation: $\hat{z}_{\text {imiss }}^{m}$ drawn from $\mathcal{N}\left(\hat{\mu}^{m}, \hat{\Sigma}^{m}\right)$
Easy to parallelized. Implemented in Amelia (website)


Amelia Earhart

## Fully conditional modeling

## Multiple Imputation by Chained Equations (MICE) - Single Iteration

| Age | Income | Gender |
| :---: | :---: | :---: |
| 33 | N.A. | F |
| 18 | 12,000 | N.A. |
| N.A. | 13,542 | M |



| Age | Income | Gender |
| :---: | :---: | :---: |
| 33 | 12,771 | F |
| 18 | 12,000 | F |
| N.A. | 13,542 | M |

(3) Bayesian Linear Regression $\begin{aligned} & \text { Age } \sim \text { Income, Gender }\end{aligned}$

| Age | Income | Gender |
| :---: | :---: | :---: |
| 33 | N.A | F |
| 18 | 12,000 | F |
| 35.3 | 13,542 | M |


| Age | Income | Gender |
| :---: | :---: | :---: |
| 33 | 12,771 | F |
| 18 | 12,000 | F |
| 35.3 | 13,542 | M |




[^6]
## Fully conditional modeling: one model/variable

(1) Initial imputation: mean imputation
(2) For a variable $j$
$2.1\left(\hat{\beta}^{-j}, \hat{\sigma}^{-j}\right)$ drawn from a Bootstrap: $\left(\hat{\beta}^{-j}, \hat{\sigma}^{-j}\right)^{1}, \ldots,\left(\hat{\beta}^{-j}, \hat{\sigma}^{-j}\right)^{M}$
2.2 Imputation of the missing values in variable $j$ with a model of $X_{j}$ on the other $X_{-j}$ : stochastic regression imputation from

$$
\mathcal{N}\left(\left(x_{i,-j}\right)^{\prime} \hat{\beta}^{-j}, \hat{\sigma}^{-j}\right)
$$

(3) Cycling through variables
$\Rightarrow$ Iteratively refine the imputation.
$\Rightarrow$ With continuous variables \& regression/variable: gibbs $\mathcal{N}(\mu, \Sigma)^{14} 15$
Implemented in mice (website) and Python*
"There is no clear-cut method for determining whether the MICE algorithm has converged"


Stef van Buuren

[^7]
## Single Iterative Random Forests Imputation

(1) Initial imputation: mean imputation - random category

Sort the variables according to the amount of missing values
(2) Fit a RF $X_{j}^{\text {obs }}$ on variables $X_{-j}^{o b s}$ and then predict $X_{j}^{\text {miss }}$
(3) Cycling through variables
(4) Repeat step 2.2 and 3 until convergence

- number of trees: 100
- number of variables randomly selected at each node $\sqrt{d}$
- number of iterations: 4-5

Implemented in the R package missForest

[^8]
## Joint versus Conditional modeling

$\Rightarrow$ Imputed values are both seen as draws from a Joint distribution

## Conditional modeling takes the lead?

- Flexible: one model/variable. Easy to deal with interactions and variables of different nature (binary, ordinal, categorical...)
- Many statistical models are conditional models
- Tailor to your data
- Super powerful in practice
$\Rightarrow$ Drawbacks: one model/variable... tedious? Computational costly.


## What to do with high correlation or when $n<p$

- JM shrinks the covariance $\Sigma+k \mathbb{I}$ (selection of $k$ ?)
- CM: ridge regression or predictors selection/variable


## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Ignorable missing values mechanism

- The full joint data distribution of $(Z, M)$ with density $p(z, m \mid \theta, \phi)^{17}$
- The (full) observed distribution ${ }^{18}$ :

$$
\begin{aligned}
p\left(z_{\mathrm{obs}}, m ; \theta, \phi\right) & =\int p(z, m ; \theta, \phi) d z_{\mathrm{mis}} \\
& =\int p(z ; \theta) p(m \mid z ; \phi) d z_{\mathrm{mis}}
\end{aligned}
$$

- With M(C)AR data:

$$
\begin{aligned}
p\left(z_{\mathrm{obs}}, m ; \theta, \phi\right) & =\int p(z ; \theta) p\left(m \mid z_{\mathrm{obs}} ; \phi\right) d z_{\mathrm{mis}} \\
& =p\left(m \mid z_{\mathrm{obs}} ; \phi\right) \int p(z ; \theta) d z_{\mathrm{miss}} \\
& =p\left(m \mid z_{\mathrm{obs}} ; \phi\right) p\left(z_{\mathrm{obs}} ; \theta\right)
\end{aligned}
$$

$\Rightarrow$ Likelihood inference can be based on $p\left(z_{\mathrm{obs}} ; \theta\right)$

[^9]
## Expectation - Maximization (Dempster et al., 1977)

Rationale to get ML estimates: max the observed data likelihood $L_{o b s}(\theta)$ through max of $L_{\text {comp }}(\theta)$. Augment the data to simplify the problem.

E step (conditional expectation):

$$
Q\left(\theta, \theta^{\ell}\right)=\int \log (p(z ; \theta)) p\left(z_{m i s s} \mid z_{o b s} ; \theta^{\ell}\right) d z_{m i s s}
$$

M step (maximization):

$$
\theta^{\ell+1}=\operatorname{argmax}_{\theta} Q\left(\theta, \theta^{\ell}\right)
$$

Result: when $\theta^{\ell+1} \max Q\left(\theta, \theta^{\ell}\right)$ then $L_{o b s}\left(\theta^{\ell+1}\right) \geq L_{o b s}\left(\theta^{\ell}\right)$.


## Estimation of the mean and covariance matrix

```
Ex: Hypothesis zi.~\mathcal{N}(\mu,\Sigma)
=> Point estimates with EM:
> library(norm)
> pre <- prelim.norm(as.matrix(don))
> thetahat <- em.norm(pre)
> getparam.norm(pre,thetahat)
```

Exercice: EM with bivariate data

## Estimation of the mean and covariance matrix

```
Ex: Hypothesis zi.~\mathcal{N}(\mu,\Sigma)
# Point estimates with EM:
> library(norm)
> pre <- prelim.norm(as.matrix(don))
> thetahat <- em.norm(pre)
> getparam.norm(pre,thetahat)
```

Exercice: EM with bivariate data
$\Rightarrow$ Variances:

- Supplemented EM (Meng, 1991), Louis formulae
- Bootstrap approach:
- Bootstrap rows: $Z^{1}, \ldots, Z^{B}$
- EM algorithm: $\left(\hat{\mu}^{1}, \hat{\Sigma}^{1}\right), \ldots,\left(\hat{\mu}^{B}, \hat{\Sigma}^{B}\right)$


## Logistic regression with missing covariates: Parameter estima-

 tion, model selection and prediction (Jiang, J., et al, CSDA, 2018)$x=\left(x_{i j}\right)$ a $n \times d$ matrix of quantitative covariates
$y=\left(y_{i}\right)$ an $n$-vector of binary responses $\{0,1\}$
Logistic regression model: $\mathbb{P}\left(y_{i}=1 \mid x_{i} ; \beta\right)=\frac{\exp \left(\beta_{0}+\sum_{j=1}^{d} \beta_{j} x_{i j}\right)}{1+\exp \left(\beta_{0}+\sum_{j=1}^{d} \beta_{j} x_{i j}\right)}$
Covariables: $x_{i} \underset{\text { i.i.d. }}{\sim} \mathcal{N}(\mu, \Sigma)$
Log-likelihood with $\theta=(\mu, \Sigma, \beta)$ :
$\mathcal{L L}(\theta ; x, y)=\sum_{i=1}^{n}\left(\log \left(\mathrm{p}\left(y_{i} \mid x_{i} ; \beta\right)\right)+\log \left(\mathrm{p}\left(x_{i} ; \mu, \Sigma\right)\right)\right)$.

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | Y |
| :---: | :---: | :---: | :--- | :---: |
| NA | 20 | 10 | $\ldots$ | shock |
| -6 | 45 | NA | $\ldots$ | shock |
| 0 | NA | 30 | $\ldots$ | no shock |
| NA | 32 | 35 | $\ldots$ | shock |
| 1 | 63 | 40 | $\ldots$ | shock |
| -2 | NA | 12 | $\ldots$ | no shock |

## Logistic regression with missing covariates: Parameter estima-

 tion, model selection and prediction (Jiang, J., et al, CSDA, 2018)$x=\left(x_{i j}\right)$ a $n \times d$ matrix of quantitative covariates
$y=\left(y_{i}\right)$ an $n$-vector of binary responses $\{0,1\}$
Logistic regression model: $\mathbb{P}\left(y_{i}=1 \mid x_{i} ; \beta\right)=\frac{\exp \left(\beta_{0}+\sum_{j=1}^{d} \beta_{j} x_{j}\right)}{1+\exp \left(\beta_{0}+\sum_{j=1}^{d} \beta_{j} x_{j}\right)}$
Covariables: $x_{i} \underset{\text { i.i.d. }}{\sim} \mathcal{N} \mathcal{N}(\mu, \Sigma)$
Log-likelihood with $\theta=(\mu, \Sigma, \beta)$ :
$\mathcal{L L}(\theta ; x, y)=\sum_{i=1}^{n}\left(\log \left(\mathrm{p}\left(y_{i} \mid x_{i} ; \beta\right)\right)+\log \left(\mathrm{p}\left(x_{i} ; \mu, \Sigma\right)\right)\right)$.

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | $M_{1}$ | $M_{2}$ | $M_{3}$ | $\ldots$ | $Y$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| NA | 20 | 10 | $\ldots$ | 1 | 0 | 0 | $\ldots$ | shock |
| -6 | 45 | NA | $\ldots$ | 0 | 0 | 1 | $\ldots$ | shock |
| 0 | NA | 30 | $\ldots$ | 0 | 1 | 0 | $\ldots$ | no shock |
| NA | 32 | 35 | $\ldots$ | 1 | 0 | 0 | $\ldots$ | shock |

## Stochastic Approximation EM - package misaem

$\operatorname{argmax} \mathcal{L} \mathcal{L}\left(\theta ; x_{\text {obs }}, y\right)=\int \mathcal{L} \mathcal{L}(\theta ; x, y) d x_{\text {mis }}$

- E-step: Evaluate the quantity

$$
\begin{aligned}
Q\left(\theta, \theta^{\ell}\right) & =\mathbb{E}\left[\mathcal{L} \mathcal{L}(\theta ; x, y) \mid x_{\mathrm{obs}}, y ; \theta^{\ell}\right] \\
& =\int \mathcal{L} \mathcal{L}(\theta ; x, y) \mathrm{p}\left(x_{\mathrm{mis}} \mid x_{\mathrm{obs}}, y ; \theta^{\ell}\right) d x_{\mathrm{mis}}
\end{aligned}
$$

- M-step: $\theta^{\ell+1}=\operatorname{argmax}_{\theta} Q\left(\theta, \theta^{\ell}\right)$
$\Rightarrow$ Unfeasible computation of expectation
MCEM (Wei \& Tanner, 1990): Generate samples of missing data from $\mathrm{p}\left(x_{\text {mis }} \mid x_{\text {obs }}, y ; \theta^{\ell}\right)$ and replace the expectation by an empirical mean $\Rightarrow$ Require a huge number of samples

SAEM (Lavielle, 2014) almost sure convergence to MLE (Metropolis Hasting - Variance estimation with Louis formulae).

Unbiased estimates: $\hat{\beta}_{1}, \ldots, \hat{\beta}_{d}-\hat{V}\left(\hat{\beta}_{1}\right), \ldots, \hat{V}\left(\hat{\beta}_{d}\right)$-good coverage

## Stochastic Approximation EM

Starting from an initial guess $\theta_{0}$, the $k$ th iteration consists of three steps:

- Simulation: For $i=1,2, \cdots, n$, draw one sample $x_{i, \text { mis }}^{(k)}$ from

$$
\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }}, y_{i} ; \theta_{k-1}\right)
$$

- Stochastic approximation: Update the function $Q$

$$
Q_{k}(\theta)=Q_{k-1}(\theta)+\gamma_{k}\left(\mathcal{L} \mathcal{L}\left(\theta ; x_{\mathrm{obs}}, x_{\mathrm{mis}}^{(k)}, y\right)-Q_{k-1}(\theta)\right)
$$

where $\left(\gamma_{k}\right)$ is a decreasing sequence of positive numbers.

- Maximization: $\theta_{k}=\operatorname{argmax}_{\theta} Q_{k}(\theta)$.


## Stochastic Approximation EM

Starting from an initial guess $\theta_{0}$, the $k$ th iteration consists of three steps:

- Simulation: For $i=1,2, \cdots, n$, draw one sample $x_{i, \text { mis }}^{(k)}$ from

$$
\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }}, y_{i} ; \theta_{k-1}\right)
$$

- Stochastic approximation: Update the function $Q$

$$
Q_{k}(\theta)=Q_{k-1}(\theta)+\gamma_{k}\left(\mathcal{L} \mathcal{L}\left(\theta ; x_{\mathrm{obs}}, x_{\text {mis }}^{(k)}, y\right)-Q_{k-1}(\theta)\right)
$$

where $\left(\gamma_{k}\right)$ is a decreasing sequence of positive numbers.

- Maximization: $\theta_{k}=\operatorname{argmax}_{\theta} Q_{k}(\theta)$.

Convergence: Allassonniere et al. 2010)
The choice of the sequence $\left(\gamma_{k}\right)$ is important for ensuring the almost sure convergence of SAEM to a MLE.

## Metropolis-Hastings algorithm

Target distribution

$$
\begin{aligned}
f_{i}\left(x_{i, \text { mis }}\right) & =\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }}, y_{i} ; \theta\right) \\
& \propto \mathrm{p}\left(y_{i} \mid x_{i} ; \beta\right) \mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }} ; \mu, \Sigma\right)
\end{aligned}
$$

## Metropolis-Hastings algorithm

Target distribution

$$
\begin{aligned}
f_{i}\left(x_{i, \text { mis }}\right) & =\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }}, y_{i} ; \theta\right) \\
& \propto \mathrm{p}\left(y_{i} \mid x_{i} ; \beta\right) \mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }} ; \mu, \Sigma\right)
\end{aligned}
$$

Proposal distribution $g_{i}\left(x_{i, \text { mis }}\right)=\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }} ; \mu, \Sigma\right)$

$$
\begin{gathered}
x_{i, \text { mis }} \mid x_{i, \mathrm{obs}} \sim \mathcal{N}_{p}\left(\mu_{i}, \Sigma_{i}\right) \\
\mu_{i}=\mu_{i, \mathrm{mis}}+\Sigma_{i, \mathrm{mis}, \mathrm{obs}} \Sigma_{i, \mathrm{obs}, \mathrm{obs}}^{-1}\left(x_{i, \mathrm{obs}}-\mu_{i, \mathrm{obs}}\right) \\
\Sigma_{i}=\Sigma_{i, \mathrm{mis}, \mathrm{mis}}-\Sigma_{i, \mathrm{mis}, \mathrm{obs}} \Sigma_{i, \mathrm{obs}, \mathrm{obs}}^{-1} \Sigma_{i, \mathrm{obs}, \mathrm{mis}}
\end{gathered}
$$

## Metropolis-Hastings algorithm

Target distribution

$$
\begin{aligned}
f_{i}\left(x_{i, \text { mis }}\right) & =\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }}, y_{i} ; \theta\right) \\
& \propto \mathrm{p}\left(y_{i} \mid x_{i} ; \beta\right) \mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }} ; \mu, \Sigma\right)
\end{aligned}
$$

Proposal distribution $g_{i}\left(x_{i, \text { mis }}\right)=\mathrm{p}\left(x_{i, \text { mis }} \mid x_{i, \text { obs }} ; \mu, \Sigma\right)$

$$
\begin{gathered}
x_{i, \mathrm{mis}} \mid x_{i, \mathrm{obs}} \sim \mathcal{N}_{p}\left(\mu_{i}, \Sigma_{i}\right) \\
\mu_{i}=\mu_{i, \mathrm{mis}}+\Sigma_{i, \mathrm{mis}, \mathrm{obs}} \Sigma_{i, \mathrm{obs}, \mathrm{obs}}^{-1}\left(x_{i, \mathrm{obs}}-\mu_{i, \mathrm{obs}}\right) \\
\Sigma_{i}=\Sigma_{i, \mathrm{mis}, \mathrm{mis}}-\Sigma_{i, \mathrm{mis}, \mathrm{obs}} \Sigma_{i, \mathrm{obs}, \mathrm{obs}}^{-1} \Sigma_{i, \mathrm{obs}, \mathrm{mis}}
\end{gathered}
$$

Metropolis

- $z_{i m}^{(k)} \sim g_{i}\left(x_{i, m i s}\right), u \sim \mathcal{U}[0,1]$
- $r=\frac{f_{i}\left(z_{i m}^{(k)}\right) / g_{i}\left(z_{i m}^{(k)}\right)}{f_{i}\left(z_{i, m-1}^{(k)}\right) / g_{i}\left(z_{i, m-1}^{(k)}\right)}$
- If $u<r$, accept $z_{i m}^{(k)}$

Only need a few steps of Markov chains in each iteration of SAEM!

## Variance estimation

Observed Fisher information matrix (FIM) wrt $\beta$

$$
\mathcal{I}(\theta)=-\frac{\partial^{2} \mathcal{L} \mathcal{L}\left(\theta ; x_{\text {obs }}, y\right)}{\partial \theta \partial \theta^{T}}
$$

## Variance estimation

Observed Fisher information matrix (FIM) wrt $\beta$

$$
\mathcal{I}(\theta)=-\frac{\partial^{2} \mathcal{L} \mathcal{L}\left(\theta ; x_{\mathrm{obs}}, y\right)}{\partial \theta \partial \theta^{T}}
$$

Louis formula

$$
\begin{aligned}
\mathcal{I}(\theta)= & -\mathbb{E}\left(\left.\frac{\partial^{2} \mathcal{L} \mathcal{L}(\theta ; x, y)}{\partial \theta \partial \theta^{T}} \right\rvert\, x_{\mathrm{obs}}, y ; \theta\right) \\
& -\mathbb{E}\left(\left.\frac{\partial \mathcal{L} \mathcal{L}(\theta ; x, y)}{\partial \theta} \frac{\partial \mathcal{L} \mathcal{L}(\theta ; x, y)^{T}}{\partial \theta} \right\rvert\, x_{\mathrm{obs}}, y ; \theta\right) \\
& +\mathbb{E}\left(\left.\frac{\partial \mathcal{L} \mathcal{L}(\theta ; x, y)}{\partial \theta} \right\rvert\, x_{\mathrm{obs}}, y ; \theta\right) \mathbb{E}\left(\left.\frac{\partial \mathcal{L} \mathcal{L}(\theta ; x, y)}{\partial \theta} \right\rvert\, x_{\mathrm{obs}}, y ; \theta\right)^{T} .
\end{aligned}
$$

Given the MH samples of unobserved data $\left(x_{i, \text { mis }}^{(m)}, 1 \leq i \leq n, 1 \leq m \leq M\right)$, and the SAEM estimate $\hat{\theta}$
$\Rightarrow$ Estimate FIM by empirical means.

## Model selection: criterion BIC

With $\tilde{p}_{\theta}$ the number of estimated parameters in a given model $\mathcal{M}$, model selection criterion (penalized likelihood) :

$$
\operatorname{BIC}(\mathcal{M})=-2 \mathcal{L} \mathcal{L}\left(\hat{\theta}_{\mathcal{M}} ; x_{\mathrm{obs}}, y\right)+\log (n) d(\mathcal{M})
$$

How to estimate observed likelihood ?

## Model selection: criterion BIC

With $\tilde{p}_{\theta}$ the number of estimated parameters in a given model $\mathcal{M}$, model selection criterion (penalized likelihood) :

$$
\operatorname{BIC}(\mathcal{M})=-2 \mathcal{L} \mathcal{L}\left(\hat{\theta}_{\mathcal{M}} ; x_{\mathrm{obs}}, y\right)+\log (n) d(\mathcal{M})
$$

How to estimate observed likelihood ?

$$
\begin{aligned}
\mathrm{p}\left(y_{i}, x_{i, \text { obs }} ; \theta\right) & =\int \mathrm{p}\left(y_{i}, x_{i, \text { obs }} \mid x_{i, \text { mis }} ; \theta\right) \mathrm{p}\left(x_{i, \text { mis }} ; \theta\right) d x_{i, \text { mis }} \\
& =\int \mathrm{p}\left(y_{i}, x_{i, \text { obs }} \mid x_{i, \text { mis }} ; \theta\right) \frac{\mathrm{p}\left(x_{i, \text { mis }} ; \theta\right)}{g_{i}\left(x_{i, \text { mis }}\right)} g_{i}\left(x_{i, \text { mis }}\right) d x_{i, \text { mis }} \\
& =\mathbb{E}_{g_{i}}\left(\mathrm{p}\left(y_{i}, x_{i, \text { obs }} \mid x_{i, \text { mis }} ; \theta\right) \frac{\mathrm{p}\left(x_{i, \text { mis }} ; \theta\right)}{g_{i}\left(x_{i, \text { mis }}\right)}\right)
\end{aligned}
$$

Sample from $g_{i}$ (the proposal distribution in SAEM)
$\Rightarrow$ Empirical mean.

## Comparison with competitors: estimates

$x: d=5, n=1000 / n=10000 \Rightarrow y \in\{0,1\}$
percentage of missingness $=10 \%$.
Repeat 1000 times for each setting.


## Comparison with competitors: coverage

Table 1: Coverage (\%) for $n=10000$, calculated over 1000 simulations.

| parameter | no NA | CC | mice | SAEM |
| :--- | :---: | :---: | :---: | :---: |
| $\beta_{0}$ | 95.2 | 94.4 | 95.2 | 94.9 |
| $\beta_{1}$ | 96.0 | 94.7 | 93.9 | 95.1 |
| $\beta_{2}$ | 95.5 | 94.6 | 94.0 | 94.3 |
| $\beta_{3}$ | 94.9 | 94.3 | 86.5 | 94.7 |
| $\beta_{4}$ | 94.6 | 94.2 | 96.2 | 95.4 |
| $\beta_{5}$ | 95.9 | 94.4 | 89.6 | 94.7 |



## Comparison with competitors: execution time

Table 2: Comparison of execution time between no NA, MCEM, mice, and SAEM with $n=1000$ calculated over 1000 simulations.

| Execution time (seconds) | no NA | MCEM | mice | SAEM |
| :--- | :---: | :---: | :---: | ---: |
| min | $2.87 \times 10^{-3}$ | 492 | 0.64 | 9.96 |
| mean | $4.65 \times 10^{-3}$ | 773 | 0.70 | 13.50 |
| $\max$ | $43.50 \times 10^{-3}$ | 1077 | 0.76 | 16.79 |

## Application on TraumaBase

- 6384 patients, 14 variables, percentage of NA from 0 to $60 \%$
- Prediction of hemoragic shock
- Selection of 8 variables, interpretation of coefficients (age, low glasgow score positive effect)

```
> library(misaem)
> reg <- miss.glm(y~., data = don)
> regBIC <- miss.glm.model.select(don$y, subset(don,-c("y")))
> pr.saem <- predict(reg, newdata = dontest)
```


## Take home message inference/imputation

- Few implementation of EM strategies
"The idea of imputation is both seductive and dangerous". It is seductive because it can lull the user into the pleasurable state of believing that the data are complete after all, and it is dangerous because it lumps together situations where the problem is sufficiently minor that it can be legitimately handled in this way and situations where standard estimators applied to the imputed data have substantial biases." (Dempster \& Rubin, 1983)
- Single imputation aims at completing a dataset as best as possible
- Multiple imputation aims at estimating the parameters and their variability taking into account the uncertainty of the missing values
- Single imputation can be appropriate for point estimates
- Both \% of NA \& structure matter ( $5 \%$ of NA can be an issue)


## Take home message inference/imputation

$\Rightarrow$ Challenges with multiple imputation

- Multiple imputation in high dimension ?
- Aggregating lasso regressions
- Aggregating different models
- Theory with other asymptotic small $n$, large $p$ ?
$\Rightarrow$ Other contributions:
Bogdan, J. et al. 2020. Adaptive Bayesian SLOPE - High dimensional Model Selection with Missing Values. JCGS.

Muzelec, Cuturi, Boyer, J. 2020. Missing Data Imputation using Optimal Transport. ICML.

## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## PCA (complete)

Find the subspace that best represents the data


Figure 2: Camel or dromedary?
$\Rightarrow$ Best approximation with projection
$\Rightarrow$ Best representation of the variability
$\Rightarrow$ Do not distort the distances between observations

## PCA (complete)

Find the subspace that best represents the data


Figure 2: Camel or dromedary? source J.P. Fénelon
$\Rightarrow$ Best approximation with projection
$\Rightarrow$ Best representation of the variability
$\Rightarrow$ Do not distort the distances between observations

## PCA reconstruction


$\Rightarrow$ Minimizes distance between observations and their projection
$\Rightarrow$ Approx $X_{n \times p}$ with a low rank matrix $S<p\|A\|_{2}^{2}=\operatorname{tr}\left(A A^{\top}\right)$ :

$$
\operatorname{argmin}_{\mu}\left\{\|X-\mu\|_{2}^{2}: \operatorname{rank}(\mu) \leq S\right\}
$$

## PCA reconstruction


$\Rightarrow$ Minimizes distance between observations and their projection
$\Rightarrow$ Approx $X_{n \times p}$ with a low rank matrix $S<p\|A\|_{2}^{2}=\operatorname{tr}\left(A A^{\top}\right)$ :

$$
\operatorname{argmin}_{\mu}\left\{\|X-\mu\|_{2}^{2}: \operatorname{rank}(\mu) \leq S\right\}
$$

$$
\begin{aligned}
\text { SVD } X: \quad \hat{\mu}^{P C A} & =U_{n \times S} \Lambda_{S \times S}^{\frac{1}{2}} V_{p \times S}^{\prime} & & F=U \Lambda^{\frac{1}{2}} \\
& =F_{n \times S} V_{p \times S}^{\prime} & & V \text { principal axes - loadings }
\end{aligned}
$$

## Missing values in PCA

$\Rightarrow$ PCA: least squares

$$
\operatorname{argmin}_{\mu}\left\{\left\|X_{n \times p}-\mu_{n \times p}\right\|_{2}^{2}: \operatorname{rank}(\mu) \leq S\right\}
$$

$\Rightarrow$ PCA with missing values: weighted least squares

$$
\operatorname{argmin}_{\mu}\left\{\left\|W_{n \times p} \odot(X-\mu)\right\|_{2}^{2}: \operatorname{rank}(\mu) \leq S\right\}
$$

with $W_{i j}=0$ if $X_{i j}$ is missing, $W_{i j}=1$ otherwise; $\odot$ elementwise multiplication

Many algorithms: weighted alternating least squares (Gabriel \& Zamir, 1979); iterative PCA (Kiers, 1997)

## Iterative PCA

$$
\begin{array}{rr}
\mathrm{x} 1 & \mathrm{x} 2 \\
-2.0 & -2.01 \\
-1.5 & -1.48 \\
0.0 & -0.01 \\
1.5 & \mathrm{NA} \\
2.0 & 1.98
\end{array}
$$



## Iterative PCA



Initialization $\ell=0: X^{0}$ (mean imputation)

## Iterative PCA



PCA on the completed data set $\rightarrow\left(U^{\ell}, \wedge^{\ell}, V^{\ell}\right)$;

## Iterative PCA



Missing values imputed with the fitted matrix $\hat{\mu}^{\ell}=U^{\ell} \Lambda^{1 / 2} V^{\ell \prime}$

## Iterative PCA



The new imputed dataset is $\hat{X}^{\ell}=W \odot X+(\mathbf{1}-W) \odot \hat{\mu}^{\ell}$

## Iterative PCA



## Iterative PCA

$$
\begin{array}{rr}
\mathrm{x} 1 & \mathrm{x} 2 \\
-2.0 & -2.01 \\
-1.5 & -1.48 \\
0.0 & -0.01 \\
1.5 & \mathrm{NA} \\
2.0 & 1.98 \\
& \\
\mathrm{x} 1 & \mathrm{x} 2 \\
-2.0 & -2.01 \\
-1.5 & -1.48 \\
0.0 & -0.01 \\
1.5 & 0.57 \\
2.0 & 1.98 \\
& \\
\mathrm{x} 1 & \mathrm{x} 2 \\
-2.00 & -2.01 \\
-1.47 & -1.52 \\
0.09 & -0.11 \\
1.20 & 0.90 \\
2.18 & 1.78 \\
& \\
\mathrm{x} 1 & \mathrm{x} 2 \\
-2.0 & -2.01 \\
-1.5 & -1.48 \\
0.0 & -0.01 \\
1.5 & 0.90 \\
2.0 & 1.98
\end{array}
$$



## Iterative PCA



Steps are repeated until convergence

## Iterative PCA

$$
\begin{array}{rr}
\mathrm{x} 1 & \mathrm{x} 2 \\
-2.0 & -2.01 \\
-1.5 & -1.48 \\
0.0 & -0.01 \\
1.5 & \mathrm{NA} \\
2.0 & 1.98 \\
& \\
& \\
& \\
& \\
& \\
& \\
& \\
& \\
& \\
& \\
\hline 2.0 & -2.01 \\
-1.5 & -1.48 \\
0.0 & -0.01 \\
1.5 & 1.46 \\
2.0 & 1.98
\end{array}
$$



PCA on the completed data set $\rightarrow\left(U^{\ell}, \Lambda^{\ell}, V^{\ell}\right)$
Missing values imputed with the fitted matrix $\hat{\mu}^{\ell}=U^{\ell} \Lambda^{1 / 2} V^{\ell \prime}$

## Iterative PCA

## Iterative PCA/SVD algorithm

(1) initialization $\ell=0: X^{0}$ (mean imputation)
(2) step $\ell$ :
(a) PCA on the completed data $\rightarrow\left(U^{\ell}, \wedge^{\ell}, V^{\ell}\right)$;
$S$ dimensions kept
(b) missing values are imputed with $\left(\hat{\mu}^{S}\right)^{\ell}=U^{\ell} \Lambda^{1 / 2^{\ell}} V^{\ell \prime}$ the new imputed data is $\hat{X}^{\ell}=W \odot X+(1-W) \odot\left(\hat{\mu}^{S}\right)^{\ell}$
(3) steps of estimation and imputation are repeated

## Iterative PCA

## Iterative PCA/SVD algorithm

(1) initialization $\ell=0: X^{0}$ (mean imputation)
(2) step $\ell$ :
(a) PCA on the completed data $\rightarrow\left(U^{\ell}, \Lambda^{\ell}, V^{\ell}\right)$;
$S$ dimensions kept
(b) missing values are imputed with $\left(\hat{\mu}^{S}\right)^{\ell}=U^{\ell} \Lambda^{1 / 2^{\ell}} V^{\ell \prime}$ the new imputed data is $\hat{X}^{\ell}=W \odot X+(1-W) \odot\left(\hat{\mu}^{S}\right)^{\ell}$

3 steps of estimation and imputation are repeated
$\Rightarrow \hat{\mu}$ from incomplete data: EM algo $X=\mu+\varepsilon, \varepsilon_{i j}{ }^{\text {iid }} \sim \mathcal{N}\left(0, \sigma^{2}\right)$ with $\mu$ of low rank, $x_{i j}=\sum_{s=1}^{S} \sqrt{\tilde{\lambda}_{s}} \tilde{u}_{i s} \tilde{j}_{j s}+\varepsilon_{i j}$
$\Rightarrow$ Completed data: good imputation (matrix completion, Netflix)

19 J. \& Husson, 2012. Selecting the number of components in PCA using cross-validation approximations. CSDA.

## Iterative PCA

## Iterative PCA/SVD algorithm

(1) initialization $\ell=0: X^{0}$ (mean imputation)
(2) step $\ell$ :
(a) PCA on the completed data $\rightarrow\left(U^{\ell}, \Lambda^{\ell}, V^{\ell}\right)$;
$S$ dimensions kept
(b) missing values are imputed with $\left(\hat{\mu}^{S}\right)^{\ell}=U^{\ell} \Lambda^{1 / 2^{\ell}} V^{\ell \prime}$
the new imputed data is $\hat{X}^{\ell}=W \odot X+(1-W) \odot\left(\hat{\mu}^{S}\right)^{\ell}$
3 steps of estimation and imputation are repeated
$\Rightarrow \hat{\mu}$ from incomplete data: EM algo $X=\mu+\varepsilon, \varepsilon_{i j}{ }^{\text {iid }} \sim \mathcal{N}\left(0, \sigma^{2}\right)$ with $\mu$ of low rank, $x_{i j}=\sum_{s=1}^{S} \sqrt{\tilde{\lambda}_{s}} \tilde{u}_{i s} \tilde{j}_{j s}+\varepsilon_{i j}$
$\Rightarrow$ Completed data: good imputation (matrix completion, Netflix)
Reduction of variability (imputation by $U \Lambda^{1 / 2} V^{\prime}$ )
Selecting S? Generalized cross-validation ${ }^{19}$
${ }^{19}$ J. \& Husson, 2012. Selecting the number of components in PCA using cross-validation approximations. CSDA.

## Overfitting

Overfitting when:

- many parameters $\left(U_{n \times S}, V_{S \times p}\right)$ ) the number of observed values: $S$ large, many NA
- data are very noisy
$\Rightarrow$ "Trust too much the relationship between variables"

Remarks:

- missing values: special case of small data set
- iterative PCA: prediction method

Solution:
$\Rightarrow$ Regularization

## Soft thresholding iterative SVD

$\Rightarrow$ Init - estimation - imputation steps:
The imputation step

$$
\hat{\mu}_{i j}^{P C A}=\sum_{s=1}^{s} \sqrt{\lambda_{s}} u_{i s} v_{j s}
$$

is replaced by ${ }^{20}$

$$
\begin{gathered}
\hat{\mu}_{i j}^{\text {Soft }}=\sum_{s=1}^{p}\left(\sqrt{\lambda_{s}}-\lambda\right)_{+} u_{i s} v_{j s} \\
X=\mu+\varepsilon \quad \operatorname{argmin}_{\mu}\left\{\|W \odot(X-\mu)\|_{2}^{2}+\lambda\|\mu\|_{\star}\right\},
\end{gathered}
$$

with $\|\mu\|_{\star}$, the nuclear norm, i.e. the sum of its singular values.
Implemented in softImpute
${ }^{20}$ T. Hastie, R. Mazumber, 2015, Matrix Completion and Low-Rank SVD via Fast Alternating Least Squares. JMLR.

## Regularized iterative PCA

The imputation step

$$
\hat{\mu}_{i j}^{\mathrm{PCA}}=\sum_{s=1}^{S} \sqrt{\lambda_{s}} u_{i s} v_{j s}
$$

is replaced by ${ }^{21},{ }^{22},{ }^{23}$ :

$$
\hat{\mu}_{i j}^{\mathrm{rPCA}}=\sum_{s=1}^{S}\left(\frac{\lambda_{s}-\hat{\sigma}^{2}}{\lambda_{s}}\right) \sqrt{\lambda_{s}} u_{i s} v_{j s}=\sum_{s=1}^{S}\left(\sqrt{\lambda_{s}}-\frac{\hat{\sigma}^{2}}{\sqrt{\lambda_{s}}}\right) u_{i s} v_{j s}
$$

$\sigma^{2}$ small $\rightarrow$ regularized PCA $\approx$ PCA
$\sigma^{2}$ large $\rightarrow$ mean imputation

$$
\hat{\sigma}^{2}=\frac{R S S}{\mathrm{ddl}}=\frac{n \sum_{s=S+1}^{p} \lambda_{s}}{n p-p-n S-p S+S^{2}+S} \quad\left(X_{n \times p} ; U_{n \times S} ; V_{p \times S}\right)
$$

Implemented in missMDA (Youtube link)
21 J., Husson. 2012. Handling missing values in exploratory multivariate data analysis. JSFDS.
${ }^{22}$ Verbank, J., Husson. 2013. Regularised PCA to denoise and visualise data Stat \& Computing.
${ }^{23}$ Rationale: L2+L0 penalty, empirical bayes Efron Moris, 1979, PPCA

## Properties

$\Rightarrow$ Powerful methods for matrix completion used in recommandation systems (ex Netflix prize: $99 \%$ missing)
$\Rightarrow$ Very good quality of imputation. Using similarities between observations and relationship between variables + reduction of dim

Model makes sense ${ }^{24}$ : Data $=$ structure of rank $S+$ noise
$\Rightarrow$ Different noise regime ${ }^{25},{ }^{26}$

- low noise: iterative PCA (tuning $S$ : CV - GCV)
- moderate: iterative regularized PCA (tuning $S$ : CV - GCV, $\sigma$ )
- high noise (SNR low, $S$ large): soft thresholding (tuning $\lambda$ : CV, $\sigma$ ) Implemented in denoiseR ${ }^{27}$

Imputed data should be analysed with caution by other methods

[^10]
## Random Forests versus PCA

Feat1 Feat2 Feat3 Feat4 Feat5...

| C1 | 1 | 1 | 1 | 1 | 1 |
| :--- | ---: | ---: | ---: | ---: | ---: |
| C2 | 1 | 1 | 1 | 1 | 1 |
| C3 | 2 | 2 | 2 | 2 | 2 |
| C4 | 2 | 2 | 2 | 2 | 2 |
| C5 | 3 | 3 | 3 | 3 | 3 |
| C6 | 3 | 3 | 3 | 3 | 3 |
| C7 | 4 | 4 | 4 | 4 | 4 |
| C8 | 4 | 4 | 4 | 4 | 4 |
| C9 | 5 | 5 | 5 | 5 | 5 |
| C10 | 5 | 5 | 5 | 5 | 5 |
| C11 | 6 | 6 | 6 | 6 | 6 |
| C12 | 6 | 6 | 6 | 6 | 6 |
| C13 | 7 | 7 | 7 | 7 | 7 |
| C14 | 7 | 7 | 7 | 7 | 7 |
| Igor | 8 | $N A$ | $N A$ | 8 | 8 |
| Frank | 8 | $N A$ | $N A$ | 8 | 8 |
| Bertrand | 9 | $N A$ | $N A$ | 9 | 9 |
| Alex | 9 | $N A$ | $N A$ | 9 | 9 |
| Yohann | 10 | $N A$ | $N A$ | 10 | 10 |
| Jean | 10 | $N A$ | $N A$ | 10 | 10 |

## Random forests versus PCA

|  | Feat1 | Feat2 | Feat3 | Feat4 | Feat5... |  | Feat | 1 Fe | at2 Feat3 | Feat4 | Feat5 | Feat1 | Feat2 | Feat3 | Feat4 | Feat5 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| C1 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1.0 | 1.00 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| C2 | 1 | 1 | 1 | 1 | 1 | 1 | 1 | 1.0 | 1.00 | 1 | 1 | 1 | 1 | 1 | 1 | 1 |
| C3 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2.0 | 2.00 | 2 | 2 | 2 | 2 | 2 | 2 | 2 |
| C4 | 2 | 2 | 2 | 2 | 2 | 2 | 2 | 2.0 | 2.00 | 2 | 2 | 2 | 2 | 2 | 2 | 2 |
| C5 | 3 | 3 | 3 | 3 | 3 | 3 | 3 | 3.0 | 3.00 | 3 | 3 | 3 | 3 | 3 | 3 | 3 |
| C6 | 3 | 3 | 3 | 3 | 3 | 3 | 3 | 3.0 | 3.00 | 3 | 3 | 3 | 3 | 3 | 3 | 3 |
| C7 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4.0 | 4.00 | 4 | 4 | 4 | 4 | 4 | 4 | 4 |
| C8 | 4 | 4 | 4 | 4 | 4 | 4 | 4 | 4.0 | 4.00 | 4 | 4 | 4 | 4 | 4 | 4 | 4 |
| C9 | 5 | 5 | 5 | 5 | 5 | 5 | 5 | 5.0 | 5.00 | 5 | 5 | 5 | 5 | 5 | 5 | 5 |
| C10 | 5 | 5 | 5 | 5 | 5 | 5 | 5 | 5.0 | 5.00 | 5 | 5 | 5 | 5 | 5 | 5 | 5 |
| C11 | 6 | 6 | 6 | 6 | 6 | 6 | 6 | 6.0 | 6.00 | 6 | 6 | 6 | 6 | 6 | 6 | 6 |
| C12 | 6 | 6 | 6 | 6 | 6 | 6 | 6 | 6.0 | 6.00 | 6 | 6 | 6 | 6 | 6 | 6 | 6 |
| C13 | 7 | 7 | 7 | 7 | 7 | 7 | 7 | 7.0 | 7.00 | 7 | 7 | 7 | 7 | 7 | 7 | 7 |
| C14 | 7 | 7 | 7 | 7 | 7 | 7 | 7 | 7.0 | 7.00 | 7 | 7 | 7 | 7 | 7 | 7 | 7 |
| Igor | 8 | NA | NA | 8 | 8 | 8 | 8 | 6.87 | 6.87 | 8 | 8 | 8 | 8 | 8 | 8 | 8 |
| Frank | 8 | NA | NA | 8 | 8 | 8 | 8 | 6.87 | 6.87 | 8 | 8 | 8 | 8 | 8 | 8 | 8 |
| Bertrand | 9 | NA | NA | 9 | 9 | 9 | 9 | 6.87 | 6.87 | 9 | 9 | 9 | 9 | 9 | 9 | 9 |
| Alex | 9 | NA | NA | 9 | 9 | 9 | 9 | 6.87 | 6.87 | 9 | 9 | 9 | 9 | 9 | 9 | 9 |
| Yohann | 10 | NA | NA | 10 | 10 |  | 10 | 6.87 | 6.87 | 10 | 10 | 10 | 10 | 10 | 10 | 10 |
| Jean | 10 | NA | NA | 10 | 10 |  | 10 | 6.87 | 6.87 | 10 | 10 | 10 | 10 | 10 | 10 | 10 |
| Missing |  |  |  |  |  | missForest |  |  |  |  |  | imputePCA |  |  |  |  |

$\Rightarrow$ Imputation inherits from the method: RF (computationaly costly) good for non linear relationships / PCA good for linear relationships

## Multiple imputation with Bootstrap PCA

$$
x_{i j}=\mu_{i j}+\varepsilon_{i j}=\sum_{s=1}^{S} \sqrt{\tilde{\lambda}_{s}} \tilde{u}_{i s} \tilde{v}_{j s}+\varepsilon_{i j}, \varepsilon_{i j} \sim \mathcal{N}\left(0, \sigma^{2}\right)
$$

(1) Variability of the parameters, $M$ plausible: $\left(\hat{\mu}_{i j}\right)^{1}, \ldots,\left(\hat{\mu}_{i j}\right)^{M}$
(2) Noise: for $m=1, \ldots, M$, missing values $x_{i j}^{m}$ drawn $\mathcal{N}\left(\hat{\mu}_{i j}^{m}, \hat{\sigma}^{2}\right)$

Implemented in missMDA (website)


François Husson
${ }^{28}$ J. Pages. Husson. 2011. Multiple imputation in principal component analysis. ADAC.

## Visualization of the imputed values

| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 3 | 20 | 10 | s |
| -6 | 45 | 6 | s |
| 0 | 4 | 30 | no s |
| -4 | 32 | 35 | s |
| -2 | 15 | 12 | no s |
| 1 | 63 | 40 | s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| -7 | 20 | 10 | s |
| -6 | 45 | 9 | s |
| 0 | 12 | 30 | no s |
| 13 | 32 | 35 | s |
| -2 | 10 | 12 | no s |
| 1 | 63 | 40 | s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | Y |
| :---: | :---: | :---: | :---: |
| 7 | 20 | 10 | s |
| -6 | 45 | 12 | s |
| 0 | -5 | 30 | no s |
| 2 | 32 | 35 | s |
| -2 | 20 | 12 | no s |
| 1 | 63 | 40 | s |



# library(missMDA) <br> MIPCA (traumadata) 

Percentage of NA?

## Joint, conditional and PCA

$\Rightarrow$ Good estimates of the parameters and their variance from an incomplete data (coverage close to 0.95)
The variability due to missing values is well taken into account

Amelia \& mice have difficulties with large correlations or $n<p$ missMDA does not but requires a tuning parameter: number of dim.

Amelia \& missMDA are based on linear relationships mice is more flexible (one model per variable)

MI based on PCA works in a large range of configuration, $n<p, n>p$ strong or weak relationships, low or high percentage of missing values

## Simulations

The simulated data $\mathcal{N}(\mu, \Sigma)$

- vary number of obs. $n$, variables $p$, correlation $\rho$
- vary \%NA, missing values mechanism (MCAR, MAR)

$\Rightarrow$ Multiple imputation $M=100$ imputed tables with PCA, JM, CM

$\Rightarrow$ Analysis model: estimate $\theta_{1}=\mathbb{E}[Y], \theta_{2}=\beta_{1}$ (regression coefficient)
$\Rightarrow$ Combine with Rubin's rule: $\hat{\theta}=\frac{1}{M} \sum_{m=1}^{M} \hat{\theta}_{m}$
$T=\frac{1}{M} \sum_{m} \widehat{\operatorname{Var}}\left(\hat{\theta}_{m}\right)+\frac{1}{M-1} \sum_{m}\left(\hat{\theta}_{m}-\hat{\theta}\right)^{2}$
Assess Bias, CI width \& coverage - 1000 simulations


## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Incomplete ozone

|  | maxO3 | T9 | T12 | T15 | Ne 9 | Ne 12 | Ne15 | V $\times 9$ | V $\times 12$ | V $\times 15$ | maxO3v |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 0601 | 87 | 15.6 | 18.5 | 18.4 | 4 | 4 | 8 | NA | -1.7101 | -0.6946 | 84 |
| 0602 | 82 | NA | 18.4 | 17.7 | 5 | 5 | 7 | NA | NA | NA | 87 |
| 0603 | 92 | NA | 17.6 | 19.5 | 2 | 5 | 4 | 2.9544 | 1.8794 | 0.5209 | 82 |
| 0604 | 114 | 16.2 | NA | NA | 1 | 1 | 0 | NA | NA | NA | 92 |
| 0605 | 94 | 17.4 | 20.5 | NA | 8 | 8 | 7 | -0.5 | NA | -4.3301 | 114 |
| 0606 | 80 | 17.7 | NA | 18.3 | NA | NA | NA | -5.6382 | -5 | -6 | 94 |
| 0607 | NA | 16.8 | 15.6 | 14.9 | 7 | 8 | 8 | -4.3301 | -1.8794 | -3.7588 | 80 |
| 0610 | 79 | 14.9 | 17.5 | 18.9 | 5 | 5 | 4 | 0 | -1.0419 | -1.3892 | NA |
| 0611 | 101 | NA | 19.6 | 21.4 | 2 | 4 | 4 | -0.766 | NA | -2.2981 | 79 |
| 0612 | NA | 18.3 | 21.9 | 22.9 | 5 | 6 | 8 | 1.2856 | -2.2981 | -3.9392 | 101 |
| 0613 | 101 | 17.3 | 19.3 | 20.2 | NA | NA | NA | -1.5 | -1.5 | -0.8682 | NA |
| : | : | : | : | : | : | : | : | : |  |  |  |
| - | , | . | . | . | . | . | , | , | . | , |  |
| 0919 | NA | 14.8 | 16.3 | 15.9 | 7 | 7 | 7 | -4.3301 | -6.0622 | -5.1962 | 42 |
| 0920 | 71 | 15.5 | 18 | 17.4 | 7 | 7 | 6 | -3.9392 | -3.0642 | 0 | NA |
| 0921 | 96 | NA | NA | NA | 3 | 3 | 3 | NA | NA | NA | 71 |
| 0922 | 98 | NA | NA | NA | 2 | 2 | 2 | 4 | 5 | 4.3301 | 96 |
| 0923 | 92 | 14.7 | 17.6 | 18.2 | 1 | 4 | 6 | 5.1962 | 5.1423 | 3.5 | 98 |
| 0924 | NA | 13.3 | 17.7 | 17.7 | NA | NA | NA | -0.9397 | -0.766 | -0.5 | 92 |
| 0925 | 84 | 13.3 | 17.7 | 17.8 | 3 | 5 | 6 | 0 | -1 | -1.2856 | NA |
| 0927 | NA | 16.2 | 20.8 | 22.1 | 6 | 5 | 5 | -0.6946 | -2 | -1.3681 | 71 |
| 0928 | 99 | 16.9 | 23 | 22.6 | NA | 4 | 7 | 1.5 | 0.8682 | 0.8682 | NA |
| 0929 | NA | 16.9 | 19.8 | 22.1 | 6 | 5 | 3 | -4 | -3.7588 | -4 | 99 |
| 0930 | 70 | 15.7 | 18.6 | 20.7 | NA | NA | NA | 0 | -1.0419 | -4 | NA |

## Complete ozone

|  | $\max 03$ | T9 | T12 | T15 | Ne 9 | Ne12 | Ne15 | Vx9 | Vx12 | Vx15 |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 20010601 | 87.000 | 15.600 | 18.500 | 20.471 | 4.000 | 4.000 | 8.000 | 0.695 | -1.710 | -0.695 | 84.000 |
| 20010602 | 82.000 | 18.505 | 20.870 | 21.799 | 5.000 | 5.000 | 7.000 | -4.330 | 0 | 0 | 0 |
| 20010603 | 92.000 | 15.300 | 17 | 19 | 2.000 | 3.984 | 3.812 | 4 | 1 | 0.521 | 0 |
| 200 | 11 | 16 | 19 | 2 | 0 | 1.000 | 0 | 2.044 | 7 | -0.174 | 92.000 |
| 2001060 | 94.00 | 18.96 | 20 | 20 | 5.294 | 5.272 | 5.056 | -0.500 | 4 | 0 | 114.000 |
| 20010606 | 80.000 | 17.700 | 19.800 | 18.300 | 6.000 | 7.020 | 7.000 | $-5.638$ | 00 | 00 | 94.000 |
| 20010607 | 79.000 | 16.800 | 15.600 | 14.900 | 7.000 | 8.000 | 6.556 | -4.330 | 879 | $-3.759$ | 80.000 |
| 20010610 | 79.000 | 14.90 | 17.500 | 18.900 | 5.000 | 5.000 | 5.016 | 0.000 | 1.042 | $-1.389$ | 99.000 |
| 20010611 | 101.000 | 16.100 | 19.600 | 21.400 | 2.000 | 4.691 | 4.000 | -0.766 | -1.026 | -2.298 | 79.000 |
| 20010612 | 106.000 | 18.300 | 22.494 | 22.900 | 5.000 | 4.627 | 4.495 | 1.286 | -2.298 | -3.939 | 101.000 |
| 20010613 | 101.000 | 17.300 | 19.300 | 20.200 | 7.000 | 7.000 | 3.000 | -1.500 | -1.500 | -0.868 | 106.000 |


| 20010915 | 69.000 | 17.100 | 17.700 | 17.500 | 6.000 | 7.000 | 8.000 | -5.196 | -2.736 | -1.042 | 71.000 |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 20010916 | 71.000 | 15.400 | 18.091 | 16.600 | 4.000 | 5.000 | 5.000 | -3.830 | 0.000 | 1.389 | 69.000 |
| 20010917 | 60.000 | 15.283 | 18.565 | 19.556 | 4.000 | 5.000 | 4.000 | 0.000 | 3.214 | 0.000 | 71.000 |
| 20010918 | 42.000 | 14.091 | 14.300 | 14.900 | 8.000 | 7.000 | 7.000 | -2.500 | -3.214 | -2.500 | 60.000 |
| 20010919 | 65.000 | 14.800 | 16.425 | 15.900 | 7.000 | 7.982 | 7.000 | -4.341 | -6.062 | -5.196 | 42.000 |
| 20010920 | 71.000 | 15.500 | 18.000 | 17.400 | 7.000 | 7.000 | 6.000 | -3.939 | -3.064 | 0.000 | 65.000 |
| 20010924 | 76.000 | 13.300 | 17.700 | 17.700 | 5.631 | 5.883 | 5.453 | -0.940 | -0.766 | -0.500 | 65.139 |
| 20010925 | 75.573 | 13.300 | 18.434 | 17.800 | 3.000 | 5.000 | 5.001 | 0.000 | -1.000 | -1.286 | 76.000 |
| 20010927 | 77.000 | 16.200 | 20.800 | 20.499 | 5.368 | 5.495 | 5.177 | -0.695 | -2.000 | -1.473 | 71.000 |
| 20010928 | 99.000 | 18.074 | 22.169 | 23.651 | 3.531 | 3.610 | 3.561 | 1.500 | 0.868 | 0.868 | 93.135 |
| 20010929 | 83.000 | 19.855 | 22.663 | 23.847 | 5.374 | 5.000 | 3.000 | -4.000 | -3.759 | -4.000 | 99.000 |
| 20010930 | 70.000 | 15.700 | 18.600 | 20.700 | 7.000 | 6.405 | 7.000 | -2.584 | -1.042 | -4.000 | 83.000 |

## > library (missMDA)

> res.comp <- imputePCA(ozo[, 1:11])
> res.comp\$comp

## Pattern visualization





```
> library(VIM)
> aggr(don, sortVar = TRUE)
```


## Visualization


> library(VIM)
> library(VIM)
> matrixplot(don, sortby = 2)
> matrixplot(don, sortby = 2)
> marginplot(don[ ,c("T9", "max03")])
> marginplot(don[ ,c("T9", "max03")])


## Visualization with Multiple Correspondence Analysis

$\Rightarrow$ Create the missingness matrix

```
> mis.ind <- matrix("o", nrow = nrow(don), ncol = ncol(don))
> mis.ind[is.na(don)] = "m"
 dimnames(mis.ind) = dimnames(don)
> mis.ind
```



## Visualization with Multiple Correspondence Analysis

MCA graph of the categories


```
> library(FactoMineR)
> resMCA <- MCA(mis.ind)
> plot(resMCA, invis = "ind", title = "MCA graph of the categories")
```


## Imputation with PCA in practice

```
=> Step 1: Estimation of the number of dimensions
> library(missMDA)
> nb <- estim_ncpPCA(don, method.cv = "Kfold")
> nb$ncp #2
> plot(0:5, nb$criterion, xlab = "nb dim", ylab ="MSEP")
```



## Imputation with PCA in practice

$\Rightarrow$ Step 2: Imputation of the missing values
> res. comp <- imputePCA(don, ncp = 2)
> res.comp\$completeObs[1:3, ]
$\operatorname{max03} \mathrm{T} 9 \mathrm{~T} 12 \mathrm{~T} 15 \mathrm{Ne} 9 \mathrm{Ne} 12 \mathrm{Ne} 15 \mathrm{~V} 9$ Vx12 Vx15 max03v
$0601 \quad 87 \quad 15.6018 .5020 .47 \quad 4 \quad 4.00 \quad 8.00 \quad 0.69-1.71-0.69 \quad 84$
$0602 \quad 8218.5120 .88 \quad 21.81 \quad 5 \quad 5.00 \quad 7.00-4.33-4.00-3.00 \quad 87$
$\begin{array}{llllllllllll}0603 & 92 & 15.30 & 17.60 & 19.50 & 2 & 3.98 & 3.81 & 2.95 & 1.97 & 0.52 & 82\end{array}$

## Cherry on the cake: PCA on incomplete data!

## Individuals factor map (PCA)



Variables factor map (PCA)


```
> imp <- cbind.data.frame(res.comp$completeObs, ozo[, 12])
> res.pca <- PCA(imp, quanti.sup = 1, quali.sup = 12)
> plot(res.pca, hab = 12, lab = "quali"); plot(res.pca, choix = "var")
> res.pca$ind$coord #scores (principal components)
```


## Imputation for continuous data

```
> library(softImpute)
> fit1 <- softImpute(XNA, rank = , lambda = )
> X.soft <- complete(XNA, fit1)
> library(denoiseR)
> adaNA <- imputeada(XNA, gamma = 1) ## time consuming...
> X.ada <- adaNA$completeObs
```


## Multiple imputation in practice

## $\Rightarrow$ Step 1: Generate $M$ imputed data sets

```
> library(Amelia)
> res.amelia <- amelia(don, m = 100)
> library(mice)
> res.mice <- mice(don, m = 100, defaultMethod = "norm.boot")
> library(missMDA)
> res.MIPCA <- MIPCA(don, ncp = 2, nboot = 100)
> res.MIPCA$res.MI
```


## Multiple imputation in practice

## $\Rightarrow$ Step 2: visualization

Observed and Imputed values of T12


Observed versus Imputed Values of maxO3


```
> library(Amelia)
> res.amelia <- amelia(don, m = 100)
> compare.density(res.amelia, var = "T12")
> overimpute(res.amelia, var = "max03")
```


## Multiple imputation in practice

$\Rightarrow$ Step 2: visualization
$\Rightarrow$ Individuals position (and variables) with other predictions


Regularized iterative PCA
$\Rightarrow$ reference configuration

## Multiple imputation in practice

$\Rightarrow$ Step 2: visualization
$\Rightarrow$ Individuals position (and variables) with other predictions


Regularized iterative PCA
$\Rightarrow$ reference configuration

## Multiple imputation in practice

$\Rightarrow$ Step 2: visualization
$\Rightarrow$ Individuals position (and variables) with other predictions


Regularized iterative PCA
$\Rightarrow$ reference configuration

## Multiple imputation in practice

```
# Step 2: visualization
> res.MIPCA <- MIPCA(don, ncp = 2)
> plot(res.MIPCA, choice = "ind.supp"); plot(res.MIPCA, choice = "var")
```


## Supplementary projection



Variable representation


## Multiple imputation in practice

$\Rightarrow$ Step 3. Regression on each table and pool the results

$$
\begin{aligned}
& \hat{\beta}=\frac{1}{M} \sum_{m=1}^{M} \hat{\beta}_{m} \\
& T=\frac{1}{M} \sum_{m} \widehat{\operatorname{Var}}\left(\hat{\beta}_{m}\right)+\left(1+\frac{1}{M}\right) \frac{1}{M-1} \sum_{m}\left(\hat{\beta}_{m}-\hat{\beta}\right)^{2}
\end{aligned}
$$

```
> library(mice)
```

> res.mice <- mice(don, $m=100$ )
> imp.micerf <- mice(don, m = 100, defaultMethod = "rf")
$>\operatorname{lm}$. mice.out <- with (res.mice, $\operatorname{lm}(\operatorname{max03} \sim \mathrm{T} 9+\mathrm{T} 12+\mathrm{T} 15+\mathrm{Ne} 9+\ldots+\mathrm{Vx} 15+\max 03 \mathrm{v}))$
> pool.mice <- pool(lm.mice.out)
$>$ summary (pool.mice)

|  | est | se | t | df | $\operatorname{Pr}(>\|\mathrm{t}\|)$ | lo 95 | hi 95 | nmis | fmi | lambda |  |
| :--- | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| (Intercept) | 19.31 | 16.30 | 1.18 | 50.48 | 0.24 | -13.43 | 52.05 | NA | 0.46 | 0.44 |  |
| T9 | -0.88 | 2.25 | -0.39 | 26.43 | 0.70 | -5.50 | 3.75 | 37 | 0.71 | 0.69 |  |
| T12 | 3.29 | 2.38 | 1.38 | 27.54 | 0.18 | -1.59 | 8.18 | 33 | 0.70 | 0.68 |  |
| $\ldots \ldots$ |  |  |  |  |  |  |  |  |  |  |  |
| Vx15 | 0.23 | 1.33 | 0.17 | 39.00 | 0.87 | -2.47 | 2.93 | 21 | 0.57 | 0.55 |  |
| $\operatorname{max03v}$ | 0.36 | 0.10 | 3.65 | 46.03 | 0.00 | 0.16 | 0.56 | 12 | 0.50 | 0.48 |  |

## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Low rank estimation with MNAR data

$X \in \mathbb{R}^{n \times p}$ noisy realisation of a low-rank matrix $\mu \in \mathbb{R}^{n \times p}$ :

$$
X=\mu+\epsilon \text {, where }\left\{\begin{array}{l}
\mu \text { with rank } S<\min \{n, p\}, \\
\epsilon_{i} \stackrel{\Perp}{\sim} \mathcal{N}\left(0_{n}, \sigma^{2} I_{n \times n}\right), \forall i \in[1, n] .
\end{array}\right.
$$

$\rightarrow$ Access only to the missing-data matrix $Y \odot M$,

- How to estimate $\mu$ ?
- How to impute the unknown entries of $X$ ?

Data distribution

$$
p\left(x_{i j} ; \mu_{i j}\right)=\left(2 \pi \sigma^{2}\right)^{-1 / 2} \exp \left(-\frac{1}{2}\left(\frac{x_{i j}-\mu_{i j}}{\sigma}\right)^{2}\right) .
$$

## MNAR missing-data mechanism via a Logistic Model

$\forall i \in[1, n], \phi_{j}=\left(\phi_{1 j}, \phi_{2 j}\right)$ denoting a parameter vector:

$$
p\left(M_{i j} \mid x_{i j} ; \phi\right)=\left[\left(1+e^{-\phi_{1 j}\left(x_{i j}-\phi_{2 j}\right)}\right)^{-1}\right]^{\left(1-M_{i j}\right)}\left[1-\left(1+e^{-\phi_{1 j}\left(x_{i j}-\phi_{2 j}\right)}\right)^{-1}\right]^{M_{i j}}
$$

$\rightsquigarrow$ self-masked MNAR : the lack only depends on the value itself.

## Method 1: EM algo with MNAR (self-mask logistic) ${ }^{29}$

MAR (ignorable): maximize the observed penalized log-likelihood

$$
\hat{\mu} \in \operatorname{argmin}_{\mu}\|(X-\mu) \odot M\|_{2}^{2}+\lambda\|\mu\|_{\star},
$$

Algo: iterative soft-thresholding SVD (ISTA), accelerated version: FISTA
MNAR (non ignorable) $L\left(\mu, \phi ; x_{\text {obs }}, m\right)=\int p(x ; \mu) p(m \mid x ; \phi) d x_{\text {mis }}$.

- E-step:

$$
Q\left(\mu, \phi \mid \hat{\mu}^{(/ e l l)}, \hat{\phi}^{(\ell)}\right)=-\mathbb{E}_{X_{\text {mis }}}\left[\ell(\mu, \phi ; x, \mu) \mid X_{\text {obs }}, M ; \mu=\hat{\mu}^{(\ell)}, \phi=\hat{\phi}^{(t)}\right]
$$

- M-step:

$$
\hat{\mu}^{(\ell+1)}, \hat{\phi}^{(\ell+1)} \in \operatorname{argmin}_{\mu, \phi} Q\left(\mu, \phi \mid \hat{\mu}^{(t)}, \hat{\phi}^{(\ell)}\right)+\lambda\|\mu\|_{\star}
$$

- E-step: Monte-Carlo approximation and SIR algorithm.
- M-step: Separability of Q:
- $\mu$ : softImpute, FISTA.
- $\phi$ : Newton-Raphson algorithm.
$\Rightarrow$ Computationally costly, few variables with MNAR.

[^11]
## Method 2: implicitly modelling the mechanism

Add the mask!


Solve the classical MAR optimization problem

$$
\hat{\Theta} \in \operatorname{argmin}_{\Theta} \frac{1}{2}\|[(1-M) \odot X \mid M]-[M \mid 1] \odot \Theta\|_{2}^{2}+\lambda\|\Theta\|_{\star},
$$

- softImpute, FISTA.
- taking into account the mask binary type, with a Penalized Iteratively Reweighted Least Squares algorithm ${ }^{30}$.

Computationally efficient but no theoretical guaranties.

[^12]
## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Categorical data

Questionnaire data from health institute http://www.inpes.sante.fr


Principal components method to explore categorical data: Multiple Correpondence Analysis ${ }^{31}$
${ }^{31}$ M. Greenacre's books, MCA and related methods. 2006. Chapman and Hall/CRC.

## Multiple Correspondence Analysis (MCA)

$X_{n \times m} m$ categorical variables coded with dummies in $A_{n \times C_{j}}$, with $C_{j}$ the tot number of categories. For a category $c$, its frequency: $p_{c}=n_{c} / n$.


MCA: A SVD on weighted matrix: $Z=\frac{1}{\sqrt{m n}}\left(A-1 p^{T}\right) D_{p}^{-1 / 2}=U \wedge V^{\prime}$
The principal component $\left(F=U \Lambda^{1 / 2}\right)$ satisfies:

$$
\begin{gathered}
\underset{F \in \mathbb{R}^{n}}{\arg \max } \frac{1}{m} \sum_{j=1}^{m} \eta^{2}\left(F, X_{j}\right) \\
\eta^{2}\left(F, X_{j}\right)=\frac{\sum_{c=1}^{C_{j}} n_{c}\left(\bar{F}_{c}-\bar{F}\right)^{2}}{\sum_{i=1}^{n} \sum_{c=1}^{C_{j}}\left(F_{i c}-\bar{F}\right)^{2}}=\frac{\text { Between variance }}{\text { Total variance }}
\end{gathered}
$$

Benzecri, 1973 : "In data analysis the mathematical problems reduces to computing eigenvectors;
all the science (the art) is in finding the right matrix to diagonalize"

## Iterative MCA algorithm:

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | NA | g | $\ldots$ | u |
| ind 2 | NA | f | g |  | u |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | v |
| ind 5 | b | f | h |  | u |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | NA |  | v |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


|  | V1_a | V1_b | V1_c | V2_e | V2_f | V3_g | V3_h | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | 1 | 0 | 0 | NA | NA | 1 | 0 | $\ldots$ |
| ind 2 | NA | NA | NA | 0 | 1 | 1 | 0 | $\ldots$ |
| ind 3 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 4 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 5 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 6 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 7 | 0 | 0 | 1 | 0 | 1 | NA | NA | $\ldots$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| ind 1232 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |

library(missMDA); ?imputeMCA

32 J . et al. 2012. Handling Missing Values with Regularized Iterative Multiple Correspondence Analysis. Journal of classification.

## Regularized iterative MCA

## Iterative MCA algorithm:

(1) initialization: imputation of the indicator matrix (proportion)

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | NA | g | $\ldots$ | u |
| ind 2 | NA | f | g |  | u |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | $v$ |
| ind 5 | b | f | h |  | $u$ |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | NA |  | $v$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


|  | V1_a | V1_b | V1_c | V2_e | V2_f | V3_g | V3_h | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | 1 | 0 | 0 | 0.41 | 0.59 | 1 | 0 | $\ldots$ |
| ind 2 | 0.20 | 0.30 | 0.50 | 0 | 1 | 1 | 0 | $\ldots$ |
| ind 3 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 4 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 5 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 6 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 7 | 0 | 0 | 1 | 0 | 1 | 0.27 | 0.78 | $\ldots$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| ind 1232 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |

library(missMDA); ?imputeMCA

[^13]
## Regularized iterative MCA

## Iterative MCA algorithm:

(1) initialization: imputation of the indicator matrix (proportion)
(2) iterate until convergence (a) estimation: MCA on the completed data $\rightarrow U, \Lambda, V$

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | NA | g | $\ldots$ | u |
| ind 2 | NA | f | g |  | u |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | $v$ |
| ind 5 | b | f | h |  | u |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | NA |  | v |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


|  | V1_a | V1_b | V1_c | V2_e | V2_f | V3_g | V3_h | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | 1 | 0 | 0 | 0.41 | 0.59 | 1 | 0 | $\ldots$ |
| ind 2 | 0.20 | 0.30 | 0.50 | 0 | 1 | 1 | 0 | $\ldots$ |
| ind 3 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 4 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 5 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 6 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 7 | 0 | 0 | 1 | 0 | 1 | 0.27 | 0.78 | $\ldots$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| ind 1232 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |

library(missMDA); ?imputeMCA

[^14]
## Regularized iterative MCA

## Iterative MCA algorithm:

(1) initialization: imputation of the indicator matrix (proportion)
(2) iterate until convergence
(a) estimation: MCA on the completed data $\rightarrow U, \Lambda, V$
(b) imputation with the fitted matrix $\hat{\mu}=U_{S} \Lambda_{S}^{1 / 2} V_{S}^{\prime}$

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | NA | g | $\ldots$ | u |
| ind 2 | NA | $f$ | g |  | u |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | $v$ |
| ind 5 | b | f | h |  | $u$ |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | NA |  | $v$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


library(missMDA); ?imputeMCA

[^15]
## Regularized iterative MCA

## Iterative MCA algorithm:

(1) initialization: imputation of the indicator matrix (proportion)
(2) iterate until convergence
(a) estimation: MCA on the completed data $\rightarrow U, \Lambda, V$
(b) imputation with the fitted matrix $\hat{\mu}=U_{S} \Lambda_{S}^{1 / 2} V_{S}^{\prime}$
(c) column margins are updated

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | NA | g | $\ldots$ | u |
| ind 2 | NA | $f$ | g |  | $u$ |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | $v$ |
| ind 5 | b | f | h |  | $u$ |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | NA |  | $v$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


library(missMDA); ?imputeMCA

[^16]
## Regularized iterative MCA

## Iterative MCA algorithm:

(1) initialization: imputation of the indicator matrix (proportion)
(2) iterate until convergence
(a) estimation: MCA on the completed data $\rightarrow U, \Lambda, V$
(b) imputation with the fitted matrix $\hat{\mu}=U_{S} \Lambda_{S}^{1 / 2} V_{S}^{\prime}$
(c) column margins are updated

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | NA | g | $\ldots$ | u |
| ind 2 | NA | f | g |  | u |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | v |
| ind 5 | b | f | h |  | u |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | NA |  | v |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


|  | V1_a | V1_b | V1_c | V2_e | V2_f | V3_g | V3_h | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | 1 | 0 | 0 | 0.71 | 0.29 | 1 | 0 | $\ldots$ |
| ind 2 | 0.12 | 0.29 | 0.59 | 0 | 1 | 1 | 0 | $\ldots$ |
| ind 3 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 4 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 5 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 6 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 7 | 0 | 0 | 1 | 0 | 1 | 0.37 | 0.63 | $\ldots$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| ind 1232 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |

$\Rightarrow$ the imputed values can be seen as degree of membership
library(missMDA); ?imputeMCA

## Regularized iterative MCA

## Iterative MCA algorithm:

(1) initialization: imputation of the indicator matrix (proportion)
(2) iterate until convergence
(a) estimation: MCA on the completed data $\rightarrow U, \Lambda, V$
(b) imputation with the fitted matrix $\hat{\mu}=U_{S} \Lambda_{S}^{1 / 2} V_{S}^{\prime}$
(c) column margins are updated

|  | V1 | V2 | V3 | $\ldots$ | V14 |
| :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | a | e | g | $\ldots$ | u |
| ind 2 | c | f | g |  | u |
| ind 3 | a | e | h |  | $v$ |
| ind 4 | a | e | h |  | v |
| ind 5 | b | f | h |  | $u$ |
| ind 6 | c | f | h |  | u |
| ind 7 | c | f | g |  | v |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |  | $\ldots$ |
| ind 1232 | c | f | h |  | v |


|  | V1_a | V1_b | V1_c | V2_e | V2_f | V3_g | V3_h | $\ldots$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| ind 1 | 1 | 0 | 0 | 0.71 | 0.29 | 1 | 0 | $\ldots$ |
| ind 2 | $\mathbf{0 . 1 2}$ | 0.29 | 0.59 | 0 | 1 | 1 | 0 | $\ldots$ |
| ind 3 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\cdots$ |
| ind 4 | 1 | 0 | 0 | 1 | 0 | 0 | 1 | $\ldots$ |
| ind 5 | 0 | 1 | 0 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 6 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |
| ind 7 | 0 | 0 | 1 | 0 | 1 | 0.37 | 0.63 | $\ldots$ |
| $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ | $\ldots$ |
| ind 1232 | 0 | 0 | 1 | 0 | 1 | 0 | 1 | $\ldots$ |

Two ways to obtain categories: majority or draw
library(missMDA); ?imputeMCA

## Multiple imputation with MCA ${ }^{33}$

(1) Variability of the parameters: $M$ sets $\left(U_{n \times s}, \Lambda_{s \times s}, V_{m \times s}^{\top}\right)$ using a non-parametric bootstrap

(2) Categories drawn from multinomial disribution using the values in $\left(\hat{X}_{m}\right)_{1 \leq m \leq M}$

| y | $\cdots$ | Attack |
| :--- | :--- | :--- |
| y | $\cdots$ | Attack |
| y | $\cdots$ | Suicide |
|  |  |  |
| n | $\cdots$ | Accident |
| n | $\cdots$ | S |


| y | $\cdots$ | Attack |
| :--- | :--- | :--- |
| y | $\cdots$ | Attack |
| y | $\cdots$ | Attack |
|  |  |  |
| n | $\cdots$ | Accident |
| n | $\cdots$ | B |


| y | $\cdots$ | Attack |
| :---: | :---: | :---: |
| y | $\cdots$ | Attack |
| y | $\cdots$ | Suicide |
|  |  |  |
| n | $\cdots$ | Accident |
| n | $\cdots$ | Suicide |

library (missMDA) ; MIMCA()
${ }^{33}$ Audigier, Husson, J. MIMCA: Multiple imputation for categorical variables with multiple correspondence analysis. 2017. Statistic \& Computing.

## Multiple imputation for categorical data

## Joint modeling

- Log-linear model (Schafer, 1997) (cat): pb many levels
- Latent class models (Vermunt, 2014) - nonparametric Bayesian (Si \& Reiter, 2014, Murray \& Reiter, 2016) (MixedDataImpute, NPBayesImpute, NestedCategBayesImpute)


## Conditional modeling

- logistic, multinomial logit, forests (mice)
$\Rightarrow$ MIMCA provides valid inference (ex. logistic reg with missing) applied to data of various size (many levels, rare levels)

| Time (seconds) | Titanic | Galetas | Income |
| :--- | ---: | ---: | ---: |
| rows-variables-levels | $(2000-4-4)$ | $(1000-4-11)$ | $(6000-14-9)$ |
| MIMCA | 2.750 | 8.972 | 58.729 |
| Loglinear | 0.740 | 4.597 | NA |
| Nonparametric bayes | 10.854 | 17.414 | 143.652 |
| Cond logistic | 4.781 | 38.016 | 881.188 |
| Cond forests | 265.771 | 112.987 | 6329.514 |

## Categorical imputation in practice

- 1232 respondents, 14 questions, 35 categories, $9 \%$ of missing values concerning $42 \%$ of respondents

In missMDA (Youtube)

```
data(vnf)
summary(vnf)
MCA(vnf)
#1) select the number of components
nb <- estim_ncpMCA(vnf, ncp.max = 5) #Time-consuming, nb = 4
#2) Impute the indicator matrix
res.impute <- imputeMCA(vnf, ncp = 4)
res.impute$tab.disj
res.impute$comp
summary(res.impute$comp)
# MCA on the incomplete data vnf
res.mca <- MCA(vnf, tab.disj = res.impute$tab.disj)
plot(res.mca, invisible=c("var"))
plot(res.mca,invisible=c("ind"),autoLab="yes", selectMod="cos2 5", cex = 0.6)
```


## Categorical imputation in practice

- 1232 respondents, 14 questions, 35 categories, $9 \%$ of missing values concerning $42 \%$ of respondents



## Categorical imputation in practice

- 1232 respondents, 14 questions, 35 categories, $9 \%$ of missing values concerning $42 \%$ of respondents



## Comparison with respect to the imputation

- Mixed data: imputation with Factorial Analysis for Mixed Data ${ }^{34}$ FAMD. ${ }^{35}$
- Comparison with Random Forest imputation with RMSE for continuous data \& proportion of falsely classified entries for categorical data.

${ }^{34}$ F. Husson, et. al. 2017. Exploratory Multivariate Analysis by Example Using R. Chapman \& Hall.


## Comparison with respect to the imputation

## Imputations with PC methods are appropriate

- for strong linear relationships
- for categorical variables
- especially for rare categories (weights of MCA)
$\Rightarrow$ Tuning: number of components $S$ (Cross-Validation)


## Imputations with RF are appropriate

- for strong non-linear relationships between continuous variables
- when there are interactions
$\Rightarrow$ No tuning parameters?
Rq: categorical data improve the imputation on continuous data and continuous data improve the imputation on categorical data


## Comparison with respect to the imputation

## Imputations with PC methods are appropriate

- for strong linear relationships
- for categorical variables
- especially for rare categories (weights of MCA)
$\Rightarrow$ Tuning: number of components S (Cross-Validation)


## Imputations with RF are appropriate

- for strong non-linear relationships between continuous variables (cutting continuous variables into categories)
- when there are interactions (creating interactions)
$\Rightarrow$ No tuning parameters?
Rq: categorical data improve the imputation on continuous data and continuous data improve the imputation on categorical data


## Mixed imputation in practice

```
> library(missMDA)
> res.ncp <- estim_ncpFAMD(ozo)
> res.famd <-imputeFAMD(ozo, ncp = 2)
> res.famd$completeObs
> library(missForest)
> res.rf <- missForest(ozo)
> res.rf$ximp
```


## Missing values in multi-source, multi-scale data



## Ex of missing values per group of variables: Journal impact

## factors

Data from journalmetrics.com
443 journals (Computer Science, Statistics, Probability and Mathematics),
15 years,
3 types of measures:

- IPP - Impact Per Publication: like the ISI impact factor but for 3 (rather than 2) years.
- SNIP - Source Normalized Impact Per Paper: Tries to weight by the number of citations per subject field to adjust for different citation cultures.
- SJR - SCImago Journal Rank: Tries to capture average prestige per publication.

Many missing values per block of years.

## Multiple Factor Analysis (MFA) with missing values


${ }^{36}$ Husson, J. 2013. Handling missing values in Multiple Factor Analysis. FQP.

## Multiple Factor Analysis (MFA) with missing values


${ }^{36}$ Husson, J. 2013. Handling missing values in Multiple Factor Analysis. FQP.

## Multiple Factor Analysis (MFA) with missing values


${ }^{36}$ Husson, J. 2013. Handling missing values in Multiple Factor Analysis. FQP.

## Multiple Factor Analysis (MFA) with missing values

ACM Transactions on Networking trajectory

Individual factor map

${ }^{36}$ Husson, J. 2013. Handling missing values in Multiple Factor Analysis. FQP.

## MFA imputation in practice

```
> library(denoiseR)
> library(missMDA)
> data(impactfactor)
> year=NULL; for (i in 1: 15) year= c(year, seq(i,45,15))
> res.imp <- imputeMFA(impactfactor, group = rep(3, 15), type = rep("s", 15))
##
> res.mfa <-MFA(res.imp$completeObs, group=rep(3,15), type=rep("s",15),
name.group=paste("year", 1999:2013,sep="_"),graph=F)
plot(res.mfa, choix = "ind", select = "contrib 15", habillage = "group", cex = 0.7)
points(res.mfa$ind$coord[c("Journal of Statistical Software",
"Journal of the American Statistical Association", "Annals of Statistics"),
1:2], col=2, cex=0.6)
text(res.mfa$ind$coord[c("Journal of Statistical Software"), 1],
res.mfa$ind$coord[c("Journal of Statistical Software"), 2],cex=1,
labels=c("Journal of Statistical Software"), pos=3, col=2)
plot.MFA(res.mfa,choix="var", cex=0.5,shadow=TRUE, autoLab = "yes")
plot(res.mfa, select="IEEE/ACM Transactions on Networking",
partial="all",
habillage="group",unselect=0.9,chrono=TRUE)
```


## Multilevel component analysis for group of observations

Ex: inhabitants nested within countries $X \in \mathbb{R}^{K \times J}$

- similarities between countries? level 1
- similarities between inhabitants within each country? level 2
- relationship between variables at each level


$$
\begin{aligned}
x_{i j k_{i}}=x_{. j .}+\left(x_{i j .}-x_{. j .}\right) & +\left(x_{i j k_{i}}-x_{i j .}\right) \\
\text { Between } & + \text { Within }
\end{aligned}
$$

Analysis of variance: split the sum of squares for each variable $j$

$$
\sum_{i=1}^{\prime} \sum_{k=1}^{k_{i}}\left(x_{i j k_{i}}\right)^{2}=\sum_{i=1}^{l} k_{i}\left(x_{. j .}\right)^{2}+\sum_{i=1}^{l} k_{i}\left(x_{i j .}-x_{. j}\right)^{2}+\sum_{i=1}^{l} \sum_{k=1}^{k_{i}}\left(x_{i j k_{i}}-x_{i j .}\right)^{2}
$$

## Multilevel PCA (MLPCA)

$\Rightarrow$ Model for the between and within part $i=1, \ldots, I$ groups, $J$ var

$$
X_{i_{\left(k_{i} \times J\right)}}=1_{k_{i}} m^{\prime}+1_{k_{i}} F_{i}^{b^{\prime}} V^{b^{\prime}}+F_{i}^{w} V^{w^{\prime}}+E_{i}
$$

- $F_{i}^{b}\left(Q_{b} \times 1\right)$ between component scores of group $i$
- $V^{b}\left(J \times Q_{b}\right)$ between loading matrix
- $F_{i}^{w}\left(k_{i} \times Q_{w}\right)$ within component scores of group $i$
- $V_{w}\left(J \times Q_{w}\right)$ within loading matrix. Constant across groups

Fitted by minimizing the least squares ${ }^{37}$

$$
\operatorname{argmin}_{\left(m, F_{i}^{b}, V^{b}, F_{i}^{w}, V^{w}\right)} \sum_{i=1}^{I}\left\|X_{i}-1_{k_{i}} m^{\prime}-1_{k_{i}} F_{i}^{b^{\prime}} V^{b^{\prime}}-F_{i}^{w} V^{w^{\prime}}\right\|^{2}
$$

$\sum_{i=1}^{l} k_{i} F_{i}^{b}=0_{Q_{b}}$ and $1_{k_{i}}^{\prime} F_{i}^{w}=0_{Q_{w}}, \forall i$ for identifiability.

[^17]
## MLPCA - quantitative data

$i=1, \ldots, /$ groups, $J$ var, $k_{i}$ nb obs in group $i$

$$
\operatorname{argmin}_{\left(m, F_{i}^{b}, V^{b}, F_{i}^{w}, V^{w}\right)} \sum_{i=1}^{\prime}\left\|X_{i}-1_{k_{i}} m^{\prime}-1_{k_{i}} F_{i}^{b^{\prime}} V^{b^{\prime}}-F_{i}^{w} V^{w^{\prime}}\right\|^{2},
$$

$\sum_{i=1}^{\prime} k_{i} F_{i}^{b}=0_{Q_{b}}$ and $1_{k_{i}}^{\prime} F_{i}^{w}=0_{Q_{w}}, \forall i$ for identifiability.
$\left(\hat{F}^{b}, \hat{V}^{b}\right)$ : Weigthed PCA on the between part: SVD on $D_{w} X_{m} ; X_{m}(I \times J)$ the means of the variables per group, $D_{w}(I \times I) D_{w i i}=\sqrt{k_{i}}$
$\left(\hat{F}^{w}, \hat{V}^{w}\right)$ PCA on the within part: SVD on the centered data per group $X^{w}(K \times J), K=\sum_{i} k_{i}$
$\Rightarrow$ With missing values: Weighted Least Squares
$\Rightarrow$ Iterative imputation algorithm (imputation - estimation)

## Iterative MLPCA

2. iteration $\ell$ : estimation of the between structure

- SVD $D_{w} X_{m}^{\ell}=P D Q^{\prime} ; Q_{b}$ eigenvectors are kept:

$$
\begin{aligned}
& \hat{F}_{i}^{b}=\left[D_{w}^{-1} P_{Q_{b}}\right]_{i}, \hat{F}^{b} \text { concatenation by row of }\left[1_{k_{i}} \hat{F}_{i}^{b}\right] \\
& \hat{V}^{b}=Q_{Q_{b}} D_{Q_{b}},\left(J \times Q_{b}\right)
\end{aligned}
$$

- the between hat matrix is computed: $\left(\hat{X}^{b}\right)^{\ell}=\hat{F}^{b} \hat{V}^{b^{\prime}}$

3. iteration $\ell$ : imputation of the missing values with the fitted values

- $\hat{X}^{\ell}=\mathbf{1}_{K} \hat{m}^{(\ell-1)^{\prime}}+\left(\hat{X}^{b}\right)^{\ell}+\left(\hat{X}^{W}\right)^{(\ell-1)}$. The newly imputed dataset is $X^{\ell}=W \odot X+\left(\mathbf{1}_{K} \times \mathbf{1}_{J}^{\prime}-W\right) \odot \hat{X}^{\ell}$
- $\hat{m}^{\ell}$ is computed on $X^{\ell}$

4. iteration $\ell$ : estimation of the within structure

- $\operatorname{SVD}\left(X^{w}\right)^{\ell}=P D Q^{\prime} ; Q_{w}$ eigenvectors are kept:

$$
\begin{aligned}
& F^{w}=P_{Q_{w}}\left(K \times Q_{w}\right) \\
& V^{w}=Q_{Q_{w}} D_{Q_{w}}\left(J \times Q_{w}\right)
\end{aligned}
$$

- the within hat matrix is computed $\left(\hat{X}^{w}\right)^{\ell}=\hat{F}^{w} \hat{V}^{w^{\prime}}$

5. iteration $\ell$ : imputation of the missing values with the fitted values

- $X^{\ell+1}=W \odot X+\left(\mathbf{1}_{K} \times \mathbf{1}_{J}^{\prime}-W\right) \odot\left(\mathbf{1}_{K} \hat{m}^{(\ell)^{\prime}}+\left(\hat{X}^{b}\right)^{\ell}+\left(\hat{X}^{w}\right)^{\ell}\right)$
- $\hat{m}^{\ell+1}$ is computed on $X^{\ell+1}$


## Simulations design

The simulated data:

- $X_{i_{\left(k_{i} \times J\right)}}=1_{k_{i}} m^{\prime}+1_{k_{i}} F_{i}^{b^{\prime}} V^{b^{\prime}}+F_{i}^{w} V^{w^{\prime}}+E_{i}$, with $E_{i j k_{i}} \sim \mathcal{N}(0, \sigma)$
- 5 groups, 10 variables, $Q_{b}=2, Q_{w}=2$

Many scenarios are considered:

- number of individuals per group: 10-20, 70-100
- level of noise: low, strong
- percentage of missing values: $10 \%, 25 \%, 40 \%$
- missing values mechanism: MCAR, MAR
$\Rightarrow$ Prediction error: $\frac{1}{k J} \sum\left(x_{i j k_{i}}-x_{i j k_{i}}\right)^{2}$


## Comparison with competitors in terms of imputation

- Conditional model with random effect regression ${ }^{38}$, implemented in micemd
- Random forests imputation
- Global PCA - Separate PCA on each table
- Global mixed PCA (FAMD) with hospital as a variable

${ }^{38}$ Audigier, White, Jolani, Debray, Quartagno, Carpenter, van Buuren, S. \& Resche-Rigon. 2018.


## Comparison with competitors in terms of imputation

- PCA mixed give similar results than Random Forest
- mice (random effect model): difficulties with large dimensions
- Separate PCA: pb with many missing values
- Multilevel PCA is equivalent to global PCA when no group effect
- Other methods do not handle categorical variables
$\Rightarrow$ Multilevel PCA Computationaly fast in comparison to mice or RF. Implemented R package missMDA
- Numbers of components $Q_{b}$ and $Q_{w}$ ?
- Inference after imputation. Underestimation of the variance with single imputation


## Comparison with competitors in terms of imputation

- PCA mixed give similar results than Random Forest
- mice (random effect model): difficulties with large dimensions
- Separate PCA: pb with many missing values
- Multilevel PCA is equivalent to global PCA when no group effect
- Other methods do not handle categorical variables
$\Rightarrow$ Multilevel PCA Computationaly fast in comparison to mice or RF. Implemented R package missMDA
- Numbers of components $Q_{b}$ and $Q_{w}$ ? cross-validation?
- Inference after imputation. Underestimation of the variance with single imputation


## Comparison with competitors in terms of imputation

- PCA mixed give similar results than Random Forest
- mice (random effect model): difficulties with large dimensions
- Separate PCA: pb with many missing values
- Multilevel PCA is equivalent to global PCA when no group effect
- Other methods do not handle categorical variables
$\Rightarrow$ Multilevel PCA Computationaly fast in comparison to mice or RF. Implemented R package missMDA
- Numbers of components $Q_{b}$ and $Q_{w}$ ? cross-validation?
- Inference after imputation. Underestimation of the variance with single imputation
Multiple imputation: bootstrap + drawn from the predictive distribution $\mathcal{N}\left(\mathbf{1}_{K} \hat{m}^{\prime}+\hat{F}^{b} \hat{B}^{b^{\prime}}+\hat{F}^{w} \hat{B}^{w^{\prime}}, \hat{\sigma}^{2}\right)$


## Aggregation of medical data

Combining data from different institutional databases promises many advantages in personalizing medical care (large $n$, more chance for finding patients like me)

## Aggregation of medical data

Combining data from different institutional databases promises many advantages in personalizing medical care (large $n$, more chance for finding patients like me)
$\Rightarrow$ The problem: high barriers to aggregation of medical data

- lack of standardization of ontologies
- privacy concerns
- proprietary attitude towards data, reluctance to cede control
- complexity/size of aggregated data, updates problems


## Solution: distributed computation

$\Rightarrow$ Data aggregation is not always necessary
$\Rightarrow$ NIH splits the storage of aggregated data across several centers

$\Rightarrow$ Data can stay at site
$\Rightarrow$ Computations can be distributed (share burden)
$\Rightarrow$ Hospitals only share intermediate results instead of the raw data

## Topology: master-workers (Wolfson, et. al (2010))


$\Rightarrow$ Ex: Each site share the sum of age $\tilde{X}_{i}$ and the number of patients $n_{i}$.
The master computes $\bar{X}=\sum n_{i} \tilde{X}_{i} / \sum n_{i}$

## Solution: distributed computation

$\Rightarrow$ Many models fitting can be implemented:

- Maximizing a likelihood. Intermediate computations break up into sums of quantities computed on local data at sites. Log-likelihood, score function and Fisher information can partition into sums. (OK for logistic regression)
- Singular Value Decomposition (ex power method involve inner product and sum). Iterative algorithms available for SVD using quantities computed on local data at sites.
- And more.

Implemented in the R package discomp ${ }^{39}$
${ }^{39}$ Narasimhan et. al. 2017. Software for Distributed Computation on Medical
Databases: A Demonstration Project.

## Privacy preserving rank k SVD

Data: each worker has private data $X_{i} \in \mathcal{R}^{n_{i} \times p}$
Result: $V \in \mathcal{R}^{p \times k}$, and $d_{1} \geq \ldots d_{k} \geq 0$
$V \leftarrow 0, d \leftarrow 0$ foreach worker site $j$ do
$U^{[j]}=0$;
transmit $n_{j}$ to master;
end
for $i \leftarrow 1$ to $k$ do
foreach worker site $j$ do $u^{[j]} \leftarrow(1,1, \ldots, 1)$ of length $n_{j}$;
$\|u\| \leftarrow \sqrt{\sum_{j} n_{j}} ;$
transmit $\|u\|, V$, and $D$ to workers;
repeat
foreach worker site $j$ do
$u^{[j]} \leftarrow u^{[j]} /\|u\| ;$
calculate $v^{[j]} \leftarrow\left(\boldsymbol{X}^{[j]}-U^{[j]} D V^{\top}\right)^{\top} u^{[j]}$;
transmit $v^{[j]}$ to master;
end
$v \leftarrow \sum_{j} v^{[j]} ;$
$v \leftarrow v /\|v\|$;
transmit $v$ to workers;
foreach worker site $j$ do
calculate $u^{[j]} \leftarrow \boldsymbol{X}^{[j]} v$;
transmit $\left\|u^{[j]}\right\|$ to master;
end
$\|u\| \leftarrow \sum_{j}\left\|u^{[j]}\right\| ;$
transmit $\|u\|$ to workers;
$d_{i} \leftarrow\|u\| ;$
until convergence;
$V \leftarrow \operatorname{cbind}(V, v)$;
foreach worker site $j$ do $U^{[j]} \leftarrow \operatorname{cbind}\left(U^{[j]}, u^{[j]}\right)$;

## Iterative multilevel distributed imputation (distributed iterative MLPCA)


$\Rightarrow$ Impute the data of one hospital using the data of the others
$\Rightarrow$ Incentive to encourage the hospitals to participate in the project
${ }^{40}$ Robin, Husson, Narasimhan, J. (2018). Imputation of mixed data with multilevel singular value
decomposition JCGS

## Low rank matrix completion for heterogeneous (count data)

- Robin, J., Moulines \& Sardy. Low-rank model with covariates for count data with missing values. 2019. Journal of Multivariate Analysis (slides)
- Robin, Klopp, J., Moulines \& Tibshirani. Main effects and interactions in mixed and incomplete data frames. 2019. JASA.
- Sportisse, Boyer, J. Estimation and imputation in Probabilistic Principal Component Analysis with Missing Not At Random data. 2020. NeurIPS.

Works of Madeleine Udell:

- Missing Value Imputation for Mixed Data Through Gaussian Copula. 2020. ACM SIGKDD conference.
- Matrix Completion with Quantified Uncertainty through Low Rank Gaussian Copula. 2020. NeurIPS.


## Take home message: estimation/imputation with low rank methods

- Principal component methods powerful for single \& multiple imputation of quanti \& categorical data (rare categories): dimensionality reduction \& capture similarities between obs and variables.
$\Rightarrow$ Correct inferences for analysis model based on relationships between pairs of variables
$\Rightarrow$ Requires to choose the number of dimensions $S$
- SVD can be distributed
- Handling missing values in PCA, MCA, FAMD, MFA, Correspondence analysis for contingency tables
- Preprocessing before clustering - clustering with missing values


## Ressources implementation

Package missMDA:
http://factominer.free.fr/missMDA/index.html

Youtube: https://www.youtube.com/watch?v=00M8_FH6_8o\&list= PLnZgp6epRBbQzxFnQrcxg09kRt-PA66T_playlist

Article JSS: https://www.jstatsoft.org/article/view/v070i01

MOOC Exploratory Multivariate Data Analysis
Package FactoShiny (Shiny interface), Factolnvestigate (for automatic reporting)

## Outline

1. Introduction
2. Inference and Imputation with missing values

## Multiple imputation

Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low ran'k estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Outline

1. Introduction
2. Inference and Imputation with missing values

## Multiple imputation

Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low ran'k estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Collaborators on supervised learning with missing values

- M. Le Morvan, Researcher, INRIA, Paris.
- E. Scornet, Asso. Pr., Ecole Polytechnique, Paris. Topic: random forests.
- G. Varoquaux, Researcher, INRIA, Paris. Topic: machine learning/ Scikitlearn

$\Rightarrow$ Random Forests with missing values

1. Consistency of supervised learning with missing values. (2019). Revis JMLR.
$\Rightarrow$ Linear regression with missing values - MultiLayer perceptron
2. Linear predictor on linearly-generated data with missing values: non consistency and solutions. AISTAT2020.
3. Neumiss networks: differential programming for supervised learning with missing values. Neurips2020. Oral.
$\Rightarrow$ Impute then regress: What's a good imputation to predict with missing values? Neurips2021. Spotlight.

## Supervised learning framework

- A feature matrix $X$ and a response vector $Y$
- Find a prediction function that minimizes the expected risk

Bayes rule: $f^{\star} \underset{f: \mathcal{X} \rightarrow \mathcal{y}}{\operatorname{argmin}} \mathbb{E}[\ell(f(X), Y)] ; \quad f^{\star}(X)=\mathbb{E}[Y \mid X]$

A new data $\mathcal{D}_{n, \text { test }}$ to estimate the generalization error rate

- Bayes consistent: $\mathbb{E}\left[\ell\left(\hat{f}_{n}(X), Y\right)\right] \underset{n \rightarrow \infty}{\longrightarrow} \mathbb{E}\left[\ell\left(f^{\star}(\mathbf{X}), Y\right)\right]$


## Supervised learning framework

- A feature matrix $X$ and a response vector $Y$
- Find a prediction function that minimizes the expected risk

Bayes rule: $f^{\star} \in \underset{f: \mathcal{X} \rightarrow \mathcal{Y}}{\operatorname{argmin}} \mathbb{E}[\ell(f(X), Y)] ; \quad f^{\star}(X)=\mathbb{E}[Y \mid X]$

- Empirical risk: $\hat{\mathcal{F}}_{\mathcal{D}_{n, \text { train }}} \in \underset{f: \mathcal{X} \rightarrow \mathcal{Y}}{\operatorname{argmin}}\left(\frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(X_{i}\right), Y_{i}\right)\right)$

A new data $\mathcal{D}_{n, \text { test }}$ to estimate the generalization error rate

- Bayes consistent: $\mathbb{E}\left[\ell\left(\hat{f}_{n}(X), Y\right)\right] \underset{n \rightarrow \infty}{\longrightarrow} \mathbb{E}\left[\ell\left(f^{\star}(\mathbf{X}), Y\right)\right]$


## Differences with classical litterature

Aim: target an outcome $Y$ (not estimate parameters and their variance) Specificities: train \& test sets with missing values
$\Rightarrow$ Is it possible to use previous approaches (EM - impute), consistent?
$\Rightarrow$ Do we need to design new ones?

## Imputation prior to learning

## Imputation with the same model

Easy to implement for univariate imputation: The means $\left(\hat{\mu}_{1}, \ldots, \hat{\mu}_{d}\right)$ of each colum of the train. Also OK for Gaussian imputation.

Issue: Many methods are "black-boxes" and take as an input the incomplete data and output the completed data (mice, missForest)

## Separate imputation

Impute train and test separately (with a different model)
Issue: Depends on the size of the test set? one observation?

## Group imputation/ semi-supervised

Impute train and test simultaneously but the predictive model is learned only on the training imputed data set

Issue: Sometimes no training set at test time

## Mean imputation is bad for estimation



Ecological data: ${ }^{41} n=69000$ species -6 traits. Estimated correlation between Pmass \& Rmass $\approx 0$ (mean imputation) or $\approx 1$ (EM PCA)
${ }^{41}$ Wright, I. et al. (2004). The worldwide leaf economics spectrum. Nature.

## Constant (mean) imputation is consistent for prediction

$\tilde{X}=X \odot(1-M)+N A \odot M$. New feature space is $\widetilde{\mathbb{R}}^{d}=(\mathbb{R} \cup\{N A\})^{d}$.
$Y=\left(\begin{array}{l}4.6 \\ 7.9 \\ 8.3 \\ 4.6\end{array}\right) \quad \tilde{X}=\left(\begin{array}{lll}9.1 & \text { NA } & 1 \\ 2.1 & \text { NA } & 3 \\ \text { NA } & 9.6 & 2 \\ \text { NA } & 5.5 & 6\end{array}\right) \quad X=\left(\begin{array}{lll}9.1 & 8.5 & 1 \\ 2.1 & 3.5 & 3 \\ 6.7 & 9.6 & 2 \\ 4.2 & 5.5 & 6\end{array}\right) \quad M=\left(\begin{array}{lll}0 & 1 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0\end{array}\right)$

## Find a prediction function that minimizes the risk.

$$
\begin{aligned}
& \text { Bayes rule: } f^{*} \in \underset{f: \widetilde{\mathbb{R}}^{d} \rightarrow \mathbb{R}}{\operatorname{argmin}} \mathbb{E}\left[(Y-f(\tilde{X}))^{2}\right] . \\
& \begin{aligned}
f^{*}(\tilde{X}) & =\mathbb{E}[Y \mid \tilde{X}]=\mathbb{E}\left[Y \mid X_{o b s(M), M}\right] \\
& =\sum_{m \in\{0,1\}^{d}} \mathbb{E}\left[Y \mid X_{o b s(m)}, M=m\right] \mathbb{1}_{M=m}
\end{aligned}
\end{aligned}
$$

$\Rightarrow$ One model per pattern $\left(2^{d}\right)$ (Rubin, 1984, generalized propensity score)

## Constant (mean) imputation is consistent for prediction

$\tilde{X}=X \odot(1-M)+N A \odot M$. New feature space is $\widetilde{\mathbb{R}}^{d}=(\mathbb{R} \cup\{N A\})^{d}$.
$Y=\left(\begin{array}{l}4.6 \\ 7.9 \\ 8.3 \\ 4.6\end{array}\right) \quad \tilde{X}=\left(\begin{array}{lll}9.1 & \text { NA } & 1 \\ 2.1 & \text { NA } & 3 \\ \text { NA } & 9.6 & 2 \\ \text { NA } & 5.5 & 6\end{array}\right) \quad X=\left(\begin{array}{lll}9.1 & 8.5 & 1 \\ 2.1 & 3.5 & 3 \\ 6.7 & 9.6 & 2 \\ 4.2 & 5.5 & 6\end{array}\right) \quad M=\left(\begin{array}{lll}0 & 1 & 0 \\ 0 & 1 & 0 \\ 1 & 0 & 0 \\ 1 & 0 & 0\end{array}\right)$

## Find a prediction function that minimizes the risk.

$$
\begin{aligned}
& \text { Bayes rule: } f^{*} \in \underset{f: \widetilde{\mathbb{R}}^{d} \rightarrow \mathbb{R}}{\operatorname{argmin}} \mathbb{E}\left[(Y-f(\tilde{X}))^{2}\right] . \\
& \begin{aligned}
f^{*}(\tilde{X}) & =\mathbb{E}[Y \mid \tilde{X}]=\mathbb{E}\left[Y \mid X_{o b s(M), M}\right] \\
& =\sum_{m \in\{0,1\}^{d}} \mathbb{E}\left[Y \mid X_{o b s(m)}, M=m\right] \mathbb{1}_{M=m}
\end{aligned}
\end{aligned}
$$

$\Rightarrow$ One model per pattern $\left(2^{d}\right)$ (Rubin, 1984, generalized propensity score)

## Constant (mean) imputation is consistent

Framework - assumptions

- $Y=f(X)+\varepsilon$
- $X=\left(X_{1}, \ldots, X_{d}\right)$ has a continuous density $g>0$ on $[0,1]^{d}$
- $\|f\|_{\infty}<\infty$
- Missing data MAR on $X_{1}$ with $M_{1} \Perp X_{1} \mid X_{2}, \ldots, X_{d}$.
- $\left(x_{2}, \ldots, x_{d}\right) \mapsto \mathbb{P}\left[M_{1}=1 \mid X_{2}=x_{2}, \ldots, X_{d}=x_{d}\right]$ is continuous
- $\varepsilon$ is a centered noise independent of $\left(X, M_{1}\right)$
(remains valid when missing values occur for several variables $X_{1}, \ldots, X_{j}$ )


## Constant (mean) imputation is consistent

Constant imputed entry $x^{\prime}=\left(x_{1}^{\prime}, x_{2}, \ldots, x_{d}\right): x_{1}^{\prime}=x_{1} \mathbb{1}_{M_{1}=0}+\alpha \mathbb{1}_{M_{1}=1}$

## Theorem. (J. et al. 2019)

$$
\begin{aligned}
f_{\text {impute }}^{\star}\left(x^{\prime}\right)= & \mathbb{E}\left[Y \mid X_{2}=x_{2}, \ldots, X_{d}=x_{d}, M_{1}=1\right] \\
& \mathbb{1}_{\left.x_{1}^{\prime}=\alpha\right]} \mathbb{1}_{\mathbb{P}}\left[M_{1}=1 \mid X_{2}=x_{2}, \ldots, X_{d}=x_{d}\right]>0 \\
& +\mathbb{E}\left[Y \mid X=x^{\prime}\right] \mathbb{1}_{x_{1}^{\prime}=\alpha} \mathbb{1}_{\mathbb{P}}\left[M_{1}=1 \mid X_{2}=x_{2}, \ldots, X_{d}=x_{d}\right]=0 \\
& +\mathbb{E}\left[Y \mid X_{1}=x_{1}, X_{2}=x_{2}, \ldots, X_{d}=x_{d}, M_{1}=0\right] \mathbb{1}_{x_{1}^{\prime} \neq \alpha} .
\end{aligned}
$$

Prediction with mean is equal to the Bayes function almost everywhere

$$
f_{\text {impute }}^{\star}\left(X^{\prime}\right)=f^{\star}(\tilde{X})=\mathbb{E}[Y \mid \tilde{X}=\tilde{x}]
$$

Rq: pointwise equality if using a constant out of range.
$\Rightarrow$ Learn on the mean-imputed training data, impute the test set with the same means and predict is optimal if the missing data are MAR and the learning algorithm is universally consistent

## Consistency of supervised learning with NA: Rationale

- Specific value, systematic like a code for missing
- Need a lot of data (asymptotic result) and a super powerful learner
- The learner detects the code and recognizes it at the test time
- With categorical data, just code "Missing"
- With continuous data, any constant:


Train


Test

Mean imputation not bad for prediction; it is consistent; despite its drawbacks for estimation - Useful in practice!

Empirically good results for MNAR

## Consistency of supervised learning with NA: Rationale

- Specific value, systematic like a code for missing
- Need a lot of data (asymptotic result) and a super powerful learner
- The learner detects the code and recognizes it at the test time
- With categorical data, just code "Missing"
- With continuous data, any constant: out of range


Mean imputation not bad for prediction; it is consistent; despite its drawbacks for estimation - Useful in practice!

Empirically good results for MNAR

## CART (Breiman, 1984)

Built recursively by splitting the current cell into two children: Find the feature $j^{\star}$, the threshold $z^{\star}$ which minimises the (quadratic) loss

$$
\begin{aligned}
\left(j^{\star}, z^{\star}\right) \in \underset{(j, z) \in \mathcal{S}}{\operatorname{argmin}} \mathbb{E} & {\left[\left(Y-\mathbb{E}\left[Y \mid X_{j} \leq z\right]\right)^{2} \cdot \mathbb{1}_{X_{j} \leq z}\right.} \\
& \left.+\left(Y-\mathbb{E}\left[Y \mid X_{j}>z\right]\right)^{2} \cdot \mathbb{1}_{X_{j}>z}\right] .
\end{aligned}
$$


root

## CART (Breiman, 1984)

Built recursively by splitting the current cell into two children: Find the feature $j^{\star}$, the threshold $z^{\star}$ which minimises the (quadratic) loss

$$
\begin{aligned}
\left(j^{\star}, z^{\star}\right) \in \underset{(j, z) \in \mathcal{S}}{\operatorname{argmin}} \mathbb{E} & {\left[\left(Y-\mathbb{E}\left[Y \mid X_{j} \leq z\right]\right)^{2} \cdot \mathbb{1}_{X_{j} \leq z}\right.} \\
& \left.+\left(Y-\mathbb{E}\left[Y \mid X_{j}>z\right]\right)^{2} \cdot \mathbb{1}_{X_{j}>z}\right] .
\end{aligned}
$$



## CART (Breiman, 1984)

Built recursively by splitting the current cell into two children: Find the feature $j^{\star}$, the threshold $z^{\star}$ which minimises the (quadratic) loss

$$
\begin{aligned}
\left(j^{\star}, z^{\star}\right) \in \underset{(j, z) \in \mathcal{S}}{\operatorname{argmin}} \mathbb{E} & {\left[\left(Y-\mathbb{E}\left[Y \mid X_{j} \leq z\right]\right)^{2} \cdot \mathbb{1}_{X_{j} \leq z}\right.} \\
& \left.+\left(Y-\mathbb{E}\left[Y \mid X_{j}>z\right]\right)^{2} \cdot \mathbb{1}_{X_{j}>z}\right] .
\end{aligned}
$$



## CART with missing values

## root

|  | $X_{1}$ | $X_{2}$ | Y |
| :--- | :--- | :--- | :--- |
| 1 |  |  |  |
| 2 | NA |  |  |
| 3 | NA |  |  |
| 4 |  |  |  |

## CART with missing values

|  | $X_{1}$ | $X_{2}$ | $Y$ |
| ---: | :---: | :---: | :---: |
| 1 |  |  |  |
| 2 | NA |  |  |
| 3 | NA |  |  |
| 4 |  |  |  |

$$
X_{1} \leq s_{1}^{\text {root }} X_{1}>s_{1}
$$

1) Select variable and threshold on observed values ( $1 \& 4$ for $X_{1}$ )
$\mathbb{E}\left[\left(Y-\mathbb{E}\left[Y \mid X_{j} \leq z, M_{j}=0\right]\right)^{2} \cdot \mathbb{1}_{X_{j} \leq z, M_{j}=0}+\left(Y-\mathbb{E}\left[Y \mid X_{j}>z, M_{j}=0\right]\right)^{2} \cdot \mathbb{1}_{X_{j}>z, M_{j}=0}\right]$.

## CART with missing values

|  | $X_{1}$ | $X_{2}$ | Y |
| :---: | :---: | :---: | :---: |
| 1 |  |  |  |
| 2 | NA |  |  |
| 3 | NA |  |  |
| 4 |  |  |  |

$$
x_{1} \leq s_{1}^{\text {root }}
$$

1) Select variable and threshold on observed values ( $1 \& 4$ for $X_{1}$ )
$\mathbb{E}\left[\left(Y-\mathbb{E}\left[Y \mid X_{j} \leq z, M_{j}=0\right]\right)^{2} \cdot \mathbb{1}_{X_{j} \leq z, M_{j}=0}+\left(Y-\mathbb{E}\left[Y \mid X_{j}>z, M_{j}=0\right]\right)^{2} \cdot \mathbb{1}_{X_{j}>z, M_{j}=0}\right]$.
2) Propagate observations (2 \& 3) with missing values?


- Block: Send all to a side by minimizing the error (xgboost, lightgbm)
- Surrogate split: Search another variable that gives a close partition (rpart)


## Missing incorporated in attribute (Twala et al. 2008)

One step: select the variable, the threshold and propagate missing values
(1) $\left\{\widetilde{X}_{j} \leq z\right.$ or $\widetilde{X}_{j}=$ NA $\}$ vs $\left\{\widetilde{X}_{j}>z\right\}$
(2) $\left\{\widetilde{X}_{j} \leq z\right\}$ vs $\left\{\widetilde{X}_{j}>z\right.$ or $\left.\widetilde{X}_{j}=\mathrm{NA}\right\}$
(3) $\left\{\widetilde{X}_{j} \neq \mathrm{NA}\right\}$ vs $\left\{\widetilde{X}_{j}=\mathrm{NA}\right\}$.

- The splitting location $z$ depends on the missing values
- Missing values treated like a category (well to handle $\mathbb{R} \cup N A$ )
- Good for informative pattern ( $M$ explains $Y$ )

Targets one model per pattern:

$$
\mathbb{E}[Y \mid \tilde{X}]=\sum_{m \in\{0,1\}^{d}} \mathbb{E}\left[Y \mid X_{o b s(m)}, M=m\right] \mathbb{1}_{M=m}
$$

- Implementation ${ }^{42}$ : grf package, scikit-learn, partykit
$\Rightarrow$ Extremely good performances in practice for any mechanism.
${ }^{42}$ implementation trick, J. Tibshirani, duplicate the incomplete columns, and replace


## Consistency: 40\% missing values MCAR

Linear problem (high noise)


Sample size


- Surrogates (rpart)
- Mean imputation

Friedman problem (high noise)


Sample size


- Gaussian imputation - MIA

Non-linear problem (low noise)



## Outline

1. Introduction
2. Inference and Imputation with missing values

## Multiple imputation

Expectation Maximization
3. Low rank anproximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low ran'k estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Linear model with missing values

## Linear model:

$$
Y=\beta_{0}+\langle X, \beta\rangle+\varepsilon, \quad X \in \mathbb{R}^{d}, \varepsilon \text { gaussian. }
$$

## Existing solutions

- ML with EM algo. (available implementation struggles for large $d$ )
- Multiple imputation (few aggregation strategies for predictive models)
$\Rightarrow$ Mainly to estimate parameters in Missing At Random setting
Aim: Predict $Y$ (out of sample) with any missing value mechanism
$\tilde{X}=X \odot(1-M)+N A \odot M$. New feature space is $\widetilde{\mathbb{R}}^{d}=(\mathbb{R} \cup\{N A\})^{d}$.

$$
\begin{gathered}
\text { Bayes rule: } f^{*} \in \underset{f: \widetilde{\mathbb{R}^{d} \rightarrow \mathbb{R}}}{\arg \min } \mathbb{E}\left[(Y-f(\tilde{X}))^{2}\right] . \\
f^{*}(\tilde{X})=\mathbb{E}[Y \mid \tilde{X}]=\sum_{m \in\{0,1\}^{d}} \mathbb{E}\left[Y \mid X_{o b s(m)}, M=m\right] \mathbb{1}_{M=m}
\end{gathered}
$$

$\Rightarrow$ One model per pattern $\left(2^{d}\right)$ (Rubin, 1984, generalized propensity score)

## Linear model with missing values not necessarely linear

## Example

Let $Y=X_{1}+X_{2}+\varepsilon$, where $X_{2}=\exp \left(X_{1}\right)+\varepsilon_{1}$. Now, assume that only $X_{1}$ is observed. Then, the model can be rewritten as

$$
Y=X_{1}+\exp \left(X_{1}\right)+\varepsilon+\varepsilon_{1},
$$

where $f\left(X_{1}\right)=X_{1}+\exp \left(X_{1}\right)$ is the Bayes predictor. In this example, the submodel for which only $X_{1}$ is observed is not linear.
$\Rightarrow$ There exists a large variety of submodels for a same linear model.
Depend on the structure of $X$ and on the missing-value mechanism.

## Explicit Bayes predictor with missing values

## Linear model:

$$
Y=\beta_{0}+\langle X, \beta\rangle+\varepsilon, \quad X \in \mathbb{R}^{d}, \varepsilon \text { gaussian. }
$$

Bayes predictor for the linear model:

$$
\begin{aligned}
f^{\star}(\tilde{X}) & =\mathbb{E}[Y \mid \tilde{X}]=\mathbb{E}\left[\beta_{0}+\beta^{\top} X \mid M, X_{o b s(M)}\right] \\
& =\beta_{0}+\beta_{o b s(M)}^{\top} X_{o b s(M)}+\beta_{\operatorname{mis}(M)}^{\top} \mathbb{E}\left[X_{\operatorname{mis}(M)} \mid M, X_{o b s(M)}\right]
\end{aligned}
$$

## Assumptions on covariates and missing values

1. Gaussian pattern mixture model, PMM: $X \mid(M=m) \sim \mathcal{N}\left(\mu_{m}, \Sigma_{m}\right)$
2. Gaussian assumption $X \sim \mathcal{N}(\mu, \Sigma)+$ MCAR and MAR
3. (Also for Gaussian assumption + MNAR self mask gaussian)

## Under Assump. the Bayes predictor is linear per pattern

$f^{\star}\left(X_{o b s}, M\right)=\beta_{0}^{\star}+\left\langle\beta_{o b s}^{\star}, X_{o b s}\right\rangle+\left\langle\beta_{m i s}^{\star}, \mu_{m i s}+\sum_{m i s, o b s}\left(\sum_{o b s}\right)^{-1}\left(X_{o b s}-\mu_{o b s}\right)\right\rangle$
use of obs instead of obs $(M)$ for lighter notations - Expression for 2.

## Expanded Bayes predictor

Under GPMM, bayes predictor is linear per pattern $\Leftrightarrow$ linear model in $W$

$$
f^{*}(\tilde{X})=\langle W, \delta\rangle
$$

$W$ an expansion ( $2^{d}$ blocks) \& parameters $\delta \in \mathbb{R}^{d}$ function of $\beta, \mu_{m}, \Sigma_{m}$

$$
\begin{gathered}
\tilde{X}=\left(\begin{array}{c|cc}
1 & x_{1,1} & x_{1,2} \\
1 & x_{2,1} & x_{2,2} \\
\hline 1 & x_{3,1} & \text { NA } \\
1 & x_{4,1} & \text { NA } \\
\hline 1 & \text { NA } & x_{5,2} \\
1 & \text { NA } & x_{6,2} \\
\hline 1 & \text { NA } & \text { NA } \\
1 & \text { NA } & \text { NA }
\end{array}\right) W=\left(\begin{array}{rrr|rr|rr|r}
1 & x_{1,1} & x_{1,2} & 0 & 0 & 0 & 0 & 0 \\
1 & x_{2,1} & x_{2,2} & 0 & 0 & 0 & 0 & 0 \\
\hline 0 & 0 & 0 & 1 & x_{3,1} & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & x_{4,1} & 0 & 0 & 0 \\
\hline 0 & 0 & 0 & 0 & 0 & 1 & x_{5,2} & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & x_{6,2} & 0 \\
\hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{array}\right) \\
W=\left(\mathbb{1}_{M=(0,0)}, x_{1} \mathbb{1}_{M=(0,0)}, x_{2} \mathbb{1}_{M=(0,0)}, \mathbb{1}_{M=(0,1)}, x_{1} \mathbb{1}_{M=(0,1)}, \mathbb{1}_{M=(1,0),}, x_{2} \mathbb{1}_{M=(1,0),}, \mathbb{1}_{M=(1,1)}\right) .
\end{gathered}
$$

## Expanded Bayes predictor

Under GPMM, bayes predictor is linear per pattern $\Leftrightarrow$ linear model in $W$

$$
f^{*}(\tilde{X})=\langle W, \delta\rangle
$$

$W$ an expansion ( $2^{d}$ blocks) \& parameters $\delta \in \mathbb{R}^{d}$ function of $\beta, \mu_{m}, \Sigma_{m}$

$$
\tilde{X}=\left(\begin{array}{l|ll}
1 & x_{1,1} & x_{1,2} \\
1 & x_{2,1} & x_{2,2} \\
\hline 1 & x_{3,1} & \text { NA } \\
1 & x_{4,1} & \text { NA } \\
\hline 1 & \text { NA } & x_{5,2} \\
1 & \text { NA } & x_{6,2} \\
\hline 1 & \text { NA } & \text { NA } \\
1 & \text { NA } & \text { NA }
\end{array}\right) \quad W=\left(\begin{array}{rrr|rr|rr|r}
1 & x_{1,1} & x_{1,2} & 0 & 0 & 0 & 0 & 0 \\
1 & x_{2,1} & x_{2,2} & 0 & 0 & 0 & 0 & 0 \\
\hline 0 & 0 & 0 & 1 & x_{3,1} & 0 & 0 & 0 \\
0 & 0 & 0 & 1 & x_{4,1} & 0 & 0 & 0 \\
\hline 0 & 0 & 0 & 0 & 0 & 1 & x_{5,2} & 0 \\
0 & 0 & 0 & 0 & 0 & 1 & x_{6,2} & 0 \\
\hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\
0 & 0 & 0 & 0 & 0 & 0 & 0 & 1
\end{array}\right)
$$

$W=\left(1_{M=(0,0)}, X_{1} 1_{M=(0,0)}, X_{2} 1_{M=(0,0)}, 1_{M=(0,1)}, X_{1} 1_{M=(0,1)}, 1_{M=(1,0)}, X_{2} 1_{M=(1,0)}, 1_{M=(1,1)}\right)$.
Problem: Dim of $W$ is $p=\sum_{k=0}^{d}\binom{d}{k} \times(k+1)=2^{d-1} \times(d+2)$.
Need to approximate it: Linear + MLP approximation + Neumiss

## Linear Approximation

The Bayes predictor can be expressed as a polynome of $X$ and $M$, which can be truncated to a lower-dimensional approximation.

$$
f_{\mathrm{approx}}^{\star}(\tilde{X})=\beta_{0,0}^{\star}+\sum_{j=1}^{d} \beta_{j, 0}^{\star} M_{j}+\sum_{j=1}^{d} \beta_{j}^{\star} X_{j}\left(1-M_{j}\right) .
$$

$\left(\begin{array}{c|rr|rr}1 & X_{1} \odot\left(1-M_{1}\right) & X_{2} \odot\left(1-M_{2}\right) & M_{1} & M_{2} \\ \hline 1 & x_{1,1} & x_{1,2} & 0 & 0 \\ 1 & x_{2,1} & x_{2,2} & 0 & 0 \\ \hline 1 & x_{3,1} & 0 & 0 & 1 \\ 1 & x_{4,1} & 0 & 0 & 1 \\ \hline 1 & 0 & x_{5,2} & 1 & 0 \\ 1 & 0 & x_{6,2} & 1 & 0 \\ \hline 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1\end{array}\right)$

Imputing $X$ by 0 and concatenate $M$

## Linear Approximation

Impute $X$ by 0 and concatenate $M \Leftrightarrow$ optimize an imputation constant.

$$
\text { Given } \left.\left(\begin{array}{cc}
X_{1} & X_{2} \\
1.1 & 3.2 \\
\text { NA } & 0.1 \\
4.6 & \text { NA } \\
4.0 & 0.9 \\
\text { NA } & 2.2
\end{array}\right), \quad \begin{array}{cccc}
X_{1} & X_{2} & M_{1} & M_{2} \\
1.1 & 3.2 & 0 & 0 \\
0 & 0.1 & 1 & 0 \\
4.6 & 0 & 0 & 1 \\
4.0 & 0.9 & 0 & 0 \\
0 & 2.2 & 1 & 0
\end{array}\right) \Leftrightarrow\left(\begin{array}{cc}
X_{1} & X_{2} \\
1.1 & 3.2 \\
C_{1} & 0.1 \\
4.6 & C_{2} \\
4.0 & 0.9 \\
C_{1} & 2.2
\end{array}\right)
$$

Indeed,

$$
\beta_{j}\left\{X_{j}\left(1-M_{j}\right)+c_{j} M_{j}\right\}=\beta_{j} X_{j}\left(1-M_{j}\right)+\left\{\beta_{j} c_{j}\right\} M_{j} .
$$

## Expanded model VS Linear approximation

$\left(\begin{array}{ccc|ccc|cc|c}1 & x_{1,1} & x_{1,2} & 0 & 0 & 0 & 0 & 0 \\ 1 & x_{2,1} & x_{2,2} & 0 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & x_{3,1} & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & x_{4,1} & 0 & 0 & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 & x_{5,2} & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & x_{6,2} & 0 \\ \hline 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1\end{array}\right)$ VS $\left(\begin{array}{c|ccc|cc}1 & x_{1,1} & x_{1,2} & 0 & 0 \\ 1 & x_{2,1} & x_{2,2} & 0 & 0 \\ \hline 1 & x_{3,1} & 0 & 0 & 1 \\ 1 & x_{4,1} & 0 & 0 & 1 \\ \hline 1 & 0 & x_{5,2} & 1 & 0 \\ 1 & 0 & x_{6,2} & 1 & 0 \\ \hline 1 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 1 & 1\end{array}\right)$

Two estimations strategies:

- Linear reg. to estimate the expanded bayes predictor: rich model, powerful in low dimension. Costly, large variance in high dimension
- Linear approximation: lower approximation capacity smaller variance since it contains fewer parameters

Finite sample bounds - Excess of risk

- Expanded: $\mathcal{O}\left(\frac{2^{d}}{n}\right)$
- Linear approximation: $\mathcal{O}\left(d^{2}+\frac{d}{n}\right)$

Comparing the upper bounds: Risk of expanded is lower than risk of approximation when $n \gg \frac{2^{d}}{d}$

## Bayes consistency of the MLP

## Theorem. Bayes consistency of a MLP. Le Morvan et al. (2020)

Under linear model + Gaussian pattern mixture model, a MLP:

- with one hidden layer containing $2^{d}$ hidden units
- ReLU activation functions
- fed with $[X \odot(1-M), M]$ ( $\tilde{X}$ imputed by 0 concatenated with mask) can achieve the Bayes rate.

Rationale: The MLP produces a prediction function piecewise affine. Since the Bayes predictor is linear per pattern, MLP good candidate.

We show that there exists a configuration of the parameters of the MLP so that the resulting predictor is the Bayes predictor.

Number of parameters: $(d+1) 2^{d+1}+1$.
$\Rightarrow$ Provides a natural way to reduce the model capacity by reducing the number of hidden units. (Trading off estimation and approximation error)

## Neumiss Networks to approximate the covariance matrix

## The Bayes predictor is linear per pattern (Gaussian+M(C)AR)

$f^{\star}\left(X_{o b s}, M\right)=\beta_{0}^{\star}+\left\langle\beta_{o b s}^{\star}, X_{o b s}\right\rangle+\left\langle\beta_{m i s}^{\star}, \mu_{m i s}+\Sigma_{m i s, o b s}\left(\Sigma_{o b s}\right)^{-1}\left(X_{o b s}-\mu_{o b s}\right)\right\rangle$
Order- $\ell$ approx of $\left(\Sigma_{o b s(m)}^{-1}\right)$ for any $m$ defined recursively:

$$
S_{o b s(m)}^{(\ell)}=\left(I d-\Sigma_{o b s(m)}\right) S_{o b s(m)}^{(\ell-1)}+I d .
$$

Neuman Series, $S^{(0)}=I d, \ell=\infty:\left(\Sigma_{o b s(m)}\right)^{-1}=\sum_{k=0}^{\infty}\left(I d-\Sigma_{o b s(m)}\right)^{k}$

## Neumiss Networks to approximate the covariance matrix

## Order- $\ell$ approx of the Bayes predictor in MAR

$f_{\ell}^{\star}\left(X_{o b s}, M\right)=\left\langle\beta_{o b s}, X_{o b s}\right\rangle+\left\langle\beta_{m i s}, \mu_{m i s}+\sum_{m i s, o b s} S_{o b s(m)}^{(\ell)}\left(X_{o b s}-\mu_{o b s}\right)\right\rangle$.
Order- $\ell$ approx of $\left(\Sigma_{o b s(m)}^{-1}\right)$ for any $m$ defined recursively:

$$
S_{o b s(m)}^{(\ell)}=\left(I d-\Sigma_{o b s(m)}\right) S_{o b s(m)}^{(\ell-1)}+I d .
$$

Neuman Series, $S^{(0)}=I d, \ell=\infty:\left(\Sigma_{o b s(m)}\right)^{-1}=\sum_{k=0}^{\infty}\left(I d-\Sigma_{o b s(m)}\right)^{k}$

## Proposition (Risk of the Order- $\ell$ approx)

Let $\nu$ be the smallest eigenvalue of $\Sigma$. Assume linear model with Gaussian covariates, $M(C) A R$, and that the spectral radius of $\Sigma$ is $<1$. Then, for all $\ell \geq 1$,
$\mathbb{E}\left[\left(f_{l}^{\star}\left(X_{o b s}, M\right)-f^{\star}\left(X_{o b s}, M\right)\right)^{2}\right] \leq \frac{(1-\nu)^{2 \ell}\left\|\beta^{\star}\right\|_{2}^{2}}{\nu} \mathbb{E}\left[\left\|I d-S_{o b s(M)}^{(0)} \Sigma_{o b s}(M)\right\|_{2}^{2}\right]$
The error of the order- $\ell$ approximation decays exponentially fast with $\ell$.

## Neumiss Networks to approximate the covariance matrix

## Order- $\ell$ approx of the Bayes predictor in MAR

$$
f_{\ell}^{\star}\left(X_{o b s}, M\right)=\left\langle\beta_{o b s}, X_{o b s}\right\rangle+\left\langle\beta_{m i s}, \mu_{m i s}+\sum_{m i s, o b s} S_{o b s(m)}^{(\ell)}\left(X_{o b s}-\mu_{o b s}\right)\right\rangle .
$$

Order- $\ell$ approx of $\left(\Sigma_{o b s(m)}^{-1}\right)$ for any $m$ defined recursively:

$$
S_{o b s(m)}^{(\ell)}=\left(I d-\Sigma_{o b s(m)}\right) S_{o b s(m)}^{(\ell-1)}+I d .
$$

Neuman Series, $S^{(0)}=I d, \ell=\infty:\left(\Sigma_{o b s(m)}\right)^{-1}=\sum_{k=0}^{\infty}\left(I d-\Sigma_{o b s(m)}\right)^{k}$
$\Rightarrow$ Neural network architecture to approximate the Bayes predictor


Figure 1: Depth of $3, \bar{m}=1-m$. Each weight matrix $W^{(k)}$ corresponds to a simple transformation of the covariance matrix indicated in blue.

## Neumiss Networks to approximate the covariance matrix

## Order- $\ell$ approx of the Bayes predictor in MAR

$f_{\ell}^{\star}\left(X_{o b s}, M\right)=\left\langle\beta_{o b s}, X_{o b s}\right\rangle+\left\langle\beta_{m i s}, \mu_{m i s}+\sum_{m i s, o b s} S_{o b s(m)}^{(\ell)}\left(X_{o b s}-\mu_{o b s}\right)\right\rangle$.
Order- $\ell$ approx of $\left(\Sigma_{o b s(m)}^{-1}\right)$ for any $m$ defined recursively:

$$
S_{o b s(m)}^{(\ell)}=\left(I d-\Sigma_{o b s(m)}\right) S_{o b s(m)}^{(\ell-1)}+I d .
$$

Neuman Series, $S^{(0)}=I d, \ell=\infty:\left(\Sigma_{o b s(m)}\right)^{-1}=\sum_{k=0}^{\infty}\left(I d-\Sigma_{o b s(m)}\right)^{k}$
$\Rightarrow$ Neural network architecture to approximate the Bayes predictor


Figure 1: Depth of $3, \bar{m}=1-m$. Each weight matrix $W^{(k)}$ corresponds to a simple transformation of the covariance matrix indicated in blue.

## Networks with missing values: $\odot M$ nonlinearity



- Implementing a network with the matrix weights $W^{(k)}=\left(I-\Sigma_{o b s(m)}\right)$ masked differently for each sample can be challenging
- Masked weights is equivalent to masking input \& output vector. Let $v$ a vector, $\bar{m}=1-m .\left(W \odot \bar{m} \bar{m}^{\top}\right) v=(W(v \odot \bar{m})) \odot \bar{m}$

Classic network with multiplications by the mask nonlinearities $\odot M$

## Networks with missing values: $\odot M$ nonlinearity



- Implementing a network with the matrix weights $W^{(k)}=\left(I-\Sigma_{o b s(m)}\right)$ masked differently for each sample can be challenging
- Masked weights is equivalent to masking input \& output vector. Let $v$ a vector, $\bar{m}=1-m .\left(W \odot \bar{m} \bar{m}^{\top}\right) v=(W(v \odot \bar{m})) \odot \bar{m}$

Classic network with multiplications by the mask nonlinearities $\odot M$

## Proposition (equivalence MLP - depth-0 Neumiss network)

A MLP with ReLU activations, one hidden layer of $d$ hidden units, and which operates on the $[X \odot(1-M), M]$, the input $X$ imputed by 0 concatenated with the mask $M$, is equivalent to the 0 -depth $N N$

## Experiments for linear regression with missing values

- $Y=X \beta^{*}+\varepsilon, \varepsilon$ chosen such as $\operatorname{SNR}=10$.
- $X \sim \mathcal{N}(\mu, \Sigma)$
- $\Sigma=U U^{\top}+\operatorname{diag}\left(\epsilon^{\prime}\right), U \in \mathbb{R}^{d \times \frac{d}{2}}, U_{i j} \sim \mathcal{N}(0,1) \epsilon^{\prime} \sim \mathcal{U}\left(10^{-2}, 10^{-1}\right)$
- $50 \%$ of MCAR, MAR, Probit self-masking.
- Max Likelihood: to estimate the parameters of the joint Gaussian distribution $\left(X_{1}, \ldots, X_{d}, Y\right)$ with EM. Predict by conditional expectation of $Y$ given $X_{\text {obs }}$.
- ICE + LR: conditional imputation with an iterative imputer followed by linear regression.
- MLP: take as input the data imputed by 0 concatenated with the mask $[X \odot(1-M), M]$ with ReLU nonlinearity,
- MLP-Wide: one hidden layer with width increased (between $d \& 2^{d}$ )
- MLP-Deep: 1 to 10 hidden layers of $d$ hidden units
- Neumiss: The Neumiss architecture with the $\odot M$, choosing the depth on a validation set.


## Results



Figure 2: Predictive performances in various scenarios - varying missing-value mechanisms, number of samples $n$, and number of features $d$.
$\Rightarrow$ Best performances for MNAR scenario ( $50 \%$ of NA on all variables)

- More effective to increase the capacity of the Neumiss network (depth) than to increase the capacity (width) of MLP Wide.

Discussion - challenges

## Take-home message. Supervised learning with missing values.

Supervised learning different from usual inferential probabilistic models. Solutions useful in practice robust to the missing-value mechanisms but needs powerful model.

## Powerful learner with missing values

- Incomplete train and test $\rightarrow$ same imputation model
- Single constant imputation is consistent with a powerful learner
- Tree-based models: Missing Incorporated in Attribute
- To be done: nonasymptotic results, uncertainty, distributional shift: No NA in the test? Proofs in MNAR


## Linear regression with missing values

- The Bayes predictor is explicit under Gaussian assumptions/ MAR and gaussian self mask but high-dimensional.
- Approx include MLP which can be consistent and Neumiss Network
- New architecture for network with missing data: $\odot M$ nonlinearity.


## Outline

1. Introduction
2. Inference and Imputation with missing values

Multiple imputation
Expectation Maximization
3. Low rank approximation

PCA with missing values - (Multiple) Imputation with missing values
Practice
Low rank estimation with MNAR data
Categorical data/Mixed/Multi-Blocks/MultiLevel
4. Supervised learning with missing values

Random Forests with missing values
Linear regression with missing values
5. Causal Inference with missing values

## Collaborators

- Imke Mayer (Postdoc Charité Universittsmedizin Berlin)
- Stefan Wager, Erik Sverdrup (Stanford)
- Tobias Gauss, Jean-Denis Moyer (Assistance Publique Hopitaux de Paris, Traumabase)


Mayer, et al. Doubly robust treatment effect estimation with missing attributes. Annals of Applied Statistics, 14(3), 2020

## Traumabase

- 30000 patients
- 250 continuous and categorical variables: heterogeneous
- 24 hospitals
- 4000 new patients/ year

| Center | Accident | Age | Sex | Weight | Lactactes | BP | Acid Tran. | $Y$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Beaujon | fall | 54 | m | 85 | NM | 180 | treated | 0 |
| Pitie | gun | 26 | m | NR | NA | 131 | control | 1 |
| Beaujon | moto | 63 | m | 80 | 3.9 | 145 | treated | 1 |
| Pitie | moto | 30 | w | NR | Imp | 107 | control | 0 |
| HEGP | knife | 16 | m | 98 | 2.5 | 118 | treated | 1 |
| $\vdots$ |  |  |  |  |  |  |  | $\ddots$ |

## Traumabase

- 30000 patients
- 250 continuous and categorical variables: heterogeneous
- 24 hospitals
- 4000 new patients/ year

| Center | Accident | Age | Sex | Weight | Lactactes | BP | Acid Tran. | $Y$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Beaujon | fall | 54 | m | 85 | NM | 180 | treated | 0 |
| Pitie | gun | 26 | m | NR | NA | 131 | control | 1 |
| Beaujon | moto | 63 | m | 80 | 3.9 | 145 | treated | 1 |
| Pitie | moto | 30 | w | NR | Imp | 107 | control | 0 |
| HEGP | knife | 16 | m | 98 | 2.5 | 118 | treated | 1 |
| $\vdots$ |  |  |  |  |  |  |  | $\ddots$ |

$\Rightarrow$ Estimate causal effect: Administration of the treatment
"tranexamic acid" (within 3 hours after the accident) on the outcome mortality for traumatic brain injury patients.

## Missing values

Percentage of missing values


Different types of missing values
Multilevel data/ data integration: Systematic missing variable in one hospital

## Potential Outcome framework (Neyman, 1923, Rubin, 1974)

## Causal effect for a binary treatment

- $n$ i.i.d. obs $(\underbrace{X_{i}}_{\text {covariates }}, \overbrace{W_{i}}^{\text {treatment }}, \underbrace{Y_{i}(1), Y_{i}(0)}_{\text {potential outcomes }}) \in \mathbb{R}^{d} \times\{0,1\} \times \mathbb{R} \times \mathbb{R}$
- Individual causal effect of the treatment: $\Delta_{i} \triangleq Y_{i}(1)-Y_{i}(0)$

Missing problem: $\Delta_{i}$ never observed (only observe one outcome/indiv)

| Covariates |  |  | Treatment | Outcome(s) |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $X_{1}$ | $X_{2}$ | $X_{3}$ | W | $\mathrm{Y}(0)$ | $\mathrm{Y}(1)$ |
| 1.1 | 20 | F | 1 | $?$ | 200 |
| -6 | 45 | F | 0 | 10 | $?$ |
| 0 | 15 | M | 1 | $?$ | 150 |
|  | $\ldots$ |  | $\ldots$ | $\ldots$ | $\ldots$ |
| -2 | 52 | M | 0 | 100 | $?$ |


| Cov. |  |  | Treat. | Out. |
| :---: | :---: | :---: | :---: | :---: |
| $X_{1}$ | $X_{2}$ | $X_{3}$ | W | Y |
| 1.1 | 20 | F | 1 | 200 |
| -6 | 45 | F | 0 | 10 |
| 0 | 15 | M | 1 | 150 |
|  | $\ldots$ |  | $\ldots$ | $\ldots$ |
| -2 | 52 | M | 0 | 100 |

## Potential Outcome framework (Neyman, 1923, Rubin, 1974)

## Causal effect for a binary treatment

## treatment

- $n$ i.i.d. obs $(\underbrace{X_{i}}_{\text {covariates }}, \overbrace{W_{i}}, \underbrace{Y_{i}(1), Y_{i}(0)}_{\text {potential outcomes }}) \in \mathbb{R}^{d} \times\{0,1\} \times \mathbb{R} \times \mathbb{R}$
- Individual causal effect of the treatment: $\Delta_{i} \triangleq Y_{i}(1)-Y_{i}(0)$

Missing problem: $\Delta_{i}$ never observed (only observe one outcome/indiv)

| Covariates |  |  | Treatment | Outcome(s) |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $X_{1}$ | $X_{2}$ | $X_{3}$ | W | $\mathrm{Y}(0)$ | $\mathrm{Y}(1)$ |
| 1.1 | 20 | F | 1 | $?$ | 200 |
| -6 | 45 | F | 0 | 10 | $?$ |
| 0 | 15 | M | 1 | $?$ | 150 |
|  | $\ldots$ |  | $\ldots$ | $\ldots$ | $\ldots$ |
| -2 | 52 | M | 0 | 100 | $?$ |


| Cov. |  |  | Treat. | Out. |
| :---: | :---: | :---: | :---: | :---: |
| $X_{1}$ | $X_{2}$ | $X_{3}$ | W | Y |
| 1.1 | 20 | F | 1 | 200 |
| -6 | 45 | F | 0 | 10 |
| 0 | 15 | M | 1 | 150 |
|  | $\ldots$ |  | $\ldots$ | $\ldots$ |
| -2 | 52 | M | 0 | 100 |

Average Treatment Effect (ATE): $\tau=\mathbb{E}\left[\Delta_{i}\right]=\mathbb{E}\left[Y_{i}(1)-Y_{i}(0)\right]$
The ATE is the difference of the average outcome had everyone gotten treated and the average outcome had nobody gotten treatment

## Observational data: non random assignment

|  | survived | deceased | $\operatorname{Pr}($ survived $\mid$ treatment) | $\operatorname{Pr}$ (deceased $\mid$ treatment) |
| ---: | :---: | :---: | :---: | :---: |
| TA not administered | $6,238(76 \%)$ | $1,327(16 \%)$ | 0.82 | 0.18 |
| TA administered | $367(4 \%)$ | $316(4 \%)$ | 0.54 | 0.46 |

Mortality rate $20 \%$ - for treated $46 \%$ - not treated $18 \%$ : treatment kills?
Standardized mean differences between treated and control.

Severe patients (with higher risk of death) are more likely to be treated. If control group does not look like treatment group, difference in response may be confounded by differences between the groups.

## Assumption for ATE identifiability in observational data

## Unconfoundedness - selection on observables

$$
\left\{Y_{i}(0), Y_{i}(1)\right\} \Perp W_{i} \mid X_{i}
$$

Treatment assignment $W_{i}$ is random conditionally on covariates $X_{i}$
Measure enough covariates to capture dependence between $W_{i}$ and outcomes

## Overlap

Propensity score: probability of treatment given observed covariates.

$$
e(x) \triangleq \mathbb{P}\left(W_{i}=1 \mid X_{i}=x\right) \quad \forall x \in \mathcal{X}
$$

We assume overlap, i.e. $\eta<e(x)<1-\eta, \quad \forall x \in \mathcal{X}$ and some $\eta>0$
ATE not identifiable without assumptions: it is not a sample size problem.

## Assumption for ATE identifiability in observational data

## Unconfoundedness - selection on observables

$$
\left\{Y_{i}(0), Y_{i}(1)\right\} \Perp W_{i} \mid X_{i}
$$

Treatment assignment $W_{i}$ is random conditionally on covariates $X_{i}$
Measure enough covariates to capture dependence between $W_{i}$ and outcomes

## Overlap



Left: Non smoker and never treated
Right: Smokers and all treated If proba to be treated when smoker $e(x)=1$, how to estimate the outcome for smokers when not treated $Y(0)$ ? How to extrapolate if total confusion?

## Regression adjustment: g-estimator

$\mu_{(w)}(x) \triangleq \mathbb{E}[Y(w) \mid X=x]$
OLS model $w \in\{0,1\}$
$Y_{i}(w)=c_{(w)}+X_{i} \beta_{(w)}+\varepsilon_{i}(w)$


Identifiability (using $\left.\left\{Y_{i}(0), Y_{i}(1)\right\} \Perp W_{i} \mid X_{i}\right)$

$$
\begin{aligned}
\tau & =\mathbb{E}\left[\Delta_{i}\right]=\mathbb{E}\left[Y_{i}(1)-Y_{i}(0)\right] \\
& =\mathbb{E}\left[\mathbb{E}\left[Y_{i}(1)-Y_{i}(0) \mid X_{i}\right]=\mathbb{E}\left[\mu_{(1)}\left(X_{i}\right)-\mu_{(0)}\left(X_{i}\right)\right]\right. \\
& =\mathbb{E}\left[\mathbb{E}\left[Y_{i}(1) \mid W_{i}=1, X_{i}=x\right]-\mathbb{E}\left[Y_{i}(0)\left|W_{i}=0,\right| X_{i}=x\right]\right] \text { (uncounfoud) } \\
& =\mathbb{E}\left[\mathbb{E}\left[Y_{i} \mid W_{i}=1, X_{i}\right]-\mathbb{E}\left[Y_{i} \mid W_{i}=0, X_{i}\right]\right] \text { (consistency) }
\end{aligned}
$$

$\mathbb{E}\left[Y_{i} \mid W_{i}=1, X_{i}\right]$ can be estimated from data but $\mathbb{E}\left[Y_{i}(1) \mid X_{i}\right]$ not.
$\hat{\tau}_{O L S}=\frac{1}{n} \sum_{i=1}^{n}\left(\hat{\mu}_{1}\left(X_{i}\right)-\hat{\mu}_{0}\left(X_{i}\right)\right)=\frac{1}{n} \sum_{i=1}^{n}\left(\hat{c}_{(1)}+X_{i} \hat{\beta}_{(1)}\right)-\left(\hat{c}_{(0)}+X_{i} \hat{\beta}_{(0)}\right)$
$\Rightarrow$ Consistent if $\hat{\mu}_{(w)}$ consistent

## Inverse-propensity weighting estimator

Average treatment effect (ATE): $\tau \triangleq \mathbb{E}\left[\Delta_{i}\right]=\mathbb{E}\left[Y_{i}(1)-Y_{i}(0)\right]$
Propensity score (proba treated|covariates): $e(x) \triangleq \mathbb{P}\left(W_{i}=1 \mid X_{i}=x\right)$
IPW estimator (Horvitz-Thomson, survey)

$$
\hat{\tau}_{I P W} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\frac{W_{i} Y_{i}}{\hat{e}\left(X_{i}\right)}-\frac{\left(1-W_{i}\right) Y_{i}}{1-\hat{e}\left(X_{i}\right)}\right)
$$

$\Rightarrow$ Balance the differences between the two groups
$\Rightarrow$ Consistent estimator of $\tau$ when $\hat{e}(\cdot)$ consistent (logistic regression).
$\Rightarrow$ High variance (divide by probability)


## Doubly robust estimator

Define $\mu_{(w)}(x) \triangleq \mathbb{E}\left[Y_{i}(w) \mid X_{i}=x\right]$ and $e(x) \triangleq \mathbb{P}\left(W_{i}=1 \mid X_{i}=x\right)$.

## Augmented IPW - Double Robust (DR)

$$
\hat{\tau}_{A I P W} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{\mu}_{(1)}\left(X_{i}\right)-\hat{\mu}_{(0)}\left(X_{i}\right)+W_{i} \frac{Y_{i}-\hat{\mu}_{(1)}\left(X_{i}\right)}{\hat{( }\left(X_{i}\right)}-\left(1-W_{i}\right) \frac{Y_{i}-\hat{\mu}_{(0)}\left(X_{i}\right)}{1-\hat{e}\left(X_{i}\right)}\right)
$$

is consistent if either the $\hat{\mu}_{(w)}(x)$ are consistent or $\hat{e}(x)$ is consistent.

- $\hat{\tau}_{I P W} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\frac{W_{i} Y_{i}}{\hat{e}\left(X_{i}\right)}-\frac{\left(1-W_{i}\right) Y_{i}}{1-\hat{e}\left(X_{i}\right)}\right)$ : Treatment assignment $\sim$ covariates
- $\hat{\tau}_{O L S} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{\mu}_{1}\left(X_{i}\right)-\hat{\mu}_{0}\left(X_{i}\right)\right):$ Outcome $\sim$ covariates
$\Rightarrow$ Both sensitive to misspecification. DR: combine ols +ipw of residuals


## Doubly robust estimator

Define $\mu_{(w)}(x) \triangleq \mathbb{E}\left[Y_{i}(w) \mid X_{i}=x\right]$ and $e(x) \triangleq \mathbb{P}\left(W_{i}=1 \mid X_{i}=x\right)$.

## Augmented IPW - Double Robust (DR)

$$
\hat{\tau}_{A I P W} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{\mu}_{(1)}\left(X_{i}\right)-\hat{\mu}_{(0)}\left(X_{i}\right)+W_{i} \frac{Y_{i}-\hat{\mu}_{(1)}\left(X_{i}\right)}{\hat{e}\left(X_{i}\right)}-\left(1-W_{i}\right) \frac{Y_{i}-\hat{\mu}_{(0)}\left(X_{i}\right)}{1-\hat{e}\left(X_{i}\right)}\right)
$$

is consistent if either the $\hat{\mu}_{(w)}(x)$ are consistent or $\hat{e}(x)$ is consistent.

- $\hat{\tau}_{I P W} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\frac{W_{i} Y_{i}}{\hat{e}\left(X_{i}\right)}-\frac{\left(1-W_{i}\right) Y_{i}}{1-\hat{e}\left(X_{i}\right)}\right)$ : Treatment assignment $\sim$ covariates
- $\hat{\tau}_{O L S} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{\mu}_{1}\left(X_{i}\right)-\hat{\mu}_{0}\left(X_{i}\right)\right)$ : Outcome $\sim$ covariates
$\Rightarrow$ Both sensitive to misspecification. DR: combine ols +ipw of residuals

Rationale: makes group similar before extrapolation

$$
\sum_{i: W_{i}=1}\left(\widetilde{\hat{\mu}}_{(0)}\left(X_{i}\right)-\mu_{(0)}\left(X_{i}\right)\right)=\underbrace{\left(\bar{X}_{1}-\hat{\gamma}^{T} \bar{X}_{0}\right)}_{\text {covariate balancing }} \underbrace{\left(\hat{\beta}^{(0)}-\beta^{(0)}\right)}_{\text {extrapolation }}+\text { noise term }
$$

where $\hat{\gamma}=\left(1-\hat{e}\left(X_{j}\right)\right)^{-1}$

## Doubly robust ATE estimation

Model Treatment on Covariates $e(x) \triangleq \mathbb{P}\left(W_{i}=1 \mid X_{i}=x\right)$ Model Outcome on Covariates $\mu_{(w)}(x) \triangleq \mathbb{E}\left[Y_{i}(w) \mid X_{i}=x\right]$

## Augmented IPW - Double Robust (DR)

$$
\hat{\tau}_{A I P W} \triangleq \frac{1}{n} \sum_{i=1}^{n}\left(\hat{\mu}_{(1)}\left(X_{i}\right)-\hat{\mu}_{(0)}\left(X_{i}\right)+W_{i} \frac{Y_{i}-\hat{\mu}_{(1)}\left(X_{i}\right)}{\hat{e}\left(X_{i}\right)}-\left(1-W_{i}\right) \frac{Y_{i}-\hat{\mu}_{(0)}\left(X_{i}\right)}{1-\hat{e}\left(X_{i}\right)}\right)
$$

is consistent if either the $\hat{\mu}_{(w)}(x)$ are consistent or $\hat{e}(x)$ is consistent.
Possibility to use any (machine learning) procedure such as random forests, deep nets, etc. to estimate $\hat{e}(x)$ and $\hat{\mu}_{(w)}(x)$ without harming the interpretability of the causal effect estimation.

## Properties - Double Machine Learning (chernozhukov, et al. 2018)

If $\hat{e}(x)$ and $\hat{\mu}_{(w)}(x)$ converge at the rate $n^{1 / 4}$ then
$\sqrt{n}\left(\hat{\tau}_{D R}-\tau\right) \xrightarrow[n \rightarrow \infty]{d} \mathcal{N}\left(0, V^{*}\right), V^{*}$ semiparametric efficient variance.

## Causal inference with missing

 attributes?
## Missing (informative) values in the covariates

Straightforward - but often biased - solution is complete-case

| Covariates |  |  | Treatment | Outcome(s) |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $X_{1}^{*}$ | $X_{2}^{*}$ | $X_{3}^{*}$ | W | $\mathrm{Y}(0)$ | $\mathrm{Y}(1)$ |
| NA | 20 | F | 1 | $?$ | Survived |
| -6 | 45 | NA | 0 | Dead | $?$ |
| 0 | NA | M | 1 | $?$ | Survived |
| NA | 32 | F | 1 | $?$ | Dead |
| 1 | 63 | M | 1 | Dead | $?$ |
| -2 | NA | M | 0 | Survived | $?$ |

$\rightarrow$ Often not a good idea! What are the alternatives?

## Three families of methods - different assumptions

- Classical unconfoundedness + classical missing values mechanisms
- Unconfoundedness with missingness + (no) missing values mechanisms Mayer, J., Wager, Sverdrup, Moyer, Gauss. AOAS 2020.
- Latent unconfoundedness + classical missing values mechanisms Mayer, J., Raimundo, Vert. 2020.


## Under 1: Multiple Imputation

## Consistency of IPW with missing values (Seaman, White 2014)

Assume Missing At Random (MAR) mechanism. Multiple imputation (MICE using $\left(X^{*}, W, Y\right)$ ) with IPW on each imputed data is consistent when Gaussian covariates and logistic/linear treatment/oucome model

| $X_{1}^{*}$ | $X_{2}^{*}$ | $X_{3}^{*}$ | $\ldots$ | W | Y |
| :---: | :---: | :---: | :---: | :---: | :---: |
| NA | 20 | 10 | $\ldots$ | 1 | survived |
| -6 | 45 | NA | $\ldots$ | 1 | survived |
| 0 | NA | 30 | $\ldots$ | 0 | died |
| NA | 32 | 35 | $\ldots$ | 0 | survived |
| -2 | NA | 12 | $\ldots$ | 0 | died |
| 1 | 63 | 40 | $\ldots$ | 1 | survived |

1) Generate $M$ plausible values for each missing value

| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | W | Y |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 3 | 20 | 10 | $\ldots$ | 1 | s |
| -6 | 45 | 6 | $\ldots$ | 1 | s |
| 0 | 4 | 30 | $\ldots$ | 0 | d |
| -4 | 32 | 35 | $\ldots$ | 0 | s |
| -2 | 15 | 12 | $\ldots$ | 0 | d |
| 1 | 63 | 40 | $\ldots$ | 1 | s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | W | Y |
| :---: | :---: | :---: | :---: | :---: | :---: |
| -7 | 20 | 10 | $\ldots$ | 1 | s |
| -6 | 45 | 9 | $\ldots$ | 1 | s |
| 0 | 12 | 30 | $\ldots$ | 0 | d |
| 13 | 32 | 35 | $\ldots$ | 0 | s |
| -2 | 10 | 12 | $\ldots$ | 0 | d |
| 1 | 63 | 40 | $\ldots$ | 1 | s |


| $X_{1}$ | $X_{2}$ | $X_{3}$ | $\ldots$ | W | Y |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 7 | 20 | 10 | $\ldots$ | 1 | s |
| -6 | 45 | 12 | $\ldots$ | 1 | s |
| 0 | -5 | 30 | $\ldots$ | 0 | d |
| 2 | 32 | 35 | $\ldots$ | 0 | s |
| -2 | 20 | 12 | $\ldots$ | 0 | d |
| 1 | 63 | 40 | $\ldots$ | 1 | s |

2) Estimate ATE on each imputed data set: $\hat{\tau}_{m}, \widehat{\operatorname{Var}}\left(\hat{\tau}_{m}\right)$
3) Combine the results (Rubin's rules): $\hat{\tau}=\frac{1}{M} \sum_{m=1}^{M} \hat{\tau}_{m}$

$$
\widehat{\operatorname{Var}}(\hat{\tau})=\frac{1}{M} \sum_{m=1}^{M} \widehat{\operatorname{Var}}\left(\hat{\tau}_{m}\right)+\left(1+\frac{1}{M}\right) \frac{1}{M-1} \sum_{m=1}^{M}\left(\hat{\tau}_{m}-\hat{\tau}\right)^{2}
$$

## 2. Unconfoundedness with missing + (no) missing hypothesis



| Covariates |  |  | Treatment | Outcome(s) |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| $X_{1}^{*}$ | $X_{2}^{*}$ | $X_{3}^{*}$ | W | $\mathrm{Y}(0)$ | $\mathrm{Y}(1)$ |
| NA | 20 | F | 1 | $?$ | 200 |
| -6 | 45 | NA | 0 | 10 | $?$ |
| 0 | NA | M | 1 | $?$ | 150 |
| NA | 32 | F | 1 | $?$ | 100 |
| 1 | 63 | M | 1 | 15 | $?$ |
| -2 | NA | M | 0 | 20 | $?$ |

Unconfoundedness: $\left\{Y_{i}(1), Y_{i}(0)\right\} \Perp W_{i} \mid X$ not testable from the data. $\Rightarrow$ Doctors give us the DAG (covariates relevant for either treatment decision and for predicting the outcome)

Unconfoundedness with missing values: $\left\{Y_{i}(1), Y_{i}(0)\right\} \Perp W_{i} \mid X^{*}$ $X^{*} \triangleq(1-M) \odot X+M \odot N A$; with $M_{i j}=1$ if $X_{i j}$ is missing, 0 otherwise. $\Rightarrow$ Doctors decide to treat a patient based on what they observe/record. We have access to the same information as the doctors.

## Under 2: Double Robust with missing values

AIPW with missing values
$\hat{\tau}^{*} \triangleq \frac{1}{n} \sum_{i}\left(\widehat{\mu_{(1)}^{*}}\left(X_{i}\right)-\widehat{\mu_{(0)}^{*}}\left(X_{i}\right)+W_{i} \frac{Y_{i}-\widehat{\mu_{(1)}}\left(X_{i}\right)}{\left.\widehat{e^{*}( } X_{i}\right)}-\left(1-W_{i}\right) \frac{Y_{i}-\widehat{\mu_{(0)}^{*}}\left(X_{i}\right)}{1-\widehat{e^{*}}\left(X_{i}\right)}\right)$

## Generalized propensity score (Rosenbaum, Rubin JASA 1984)

$$
e^{*}\left(x^{*}\right) \triangleq \mathbb{P}\left(W=1 \mid X^{*}=x^{*}\right)
$$

One model per pattern: $\sum_{m \in\{0,1\}^{d}} \mathbb{E}\left[W \mid X_{o b s(m)}, M=m\right] \mathbb{1}_{M=m}$
$\Rightarrow$ Supervised learning with missing values. 123

- Mean imputation is consistent with a universally consistent learner.
- Missing Incorporate in Attributes (MIA) for trees methods.

[^18]
## Under 2: Double Robust with missing values

AIPW with missing values

$$
\hat{\tau}^{*} \triangleq \frac{1}{n} \sum_{i}\left(\widehat{\mu_{(1)}^{*}}\left(X_{i}\right)-\widehat{\mu_{(0)}^{*}}\left(X_{i}\right)+W_{i} \frac{Y_{i}-\widehat{\mu_{1(1)}^{*}}\left(X_{i}\right)}{\widehat{e^{*}\left(X_{i}\right)}}-\left(1-W_{i}\right) \frac{Y_{i}-\widehat{\mu_{0}^{*}}\left(X_{i}\right)}{1-\widehat{e^{*}}\left(X_{i}\right)}\right)
$$

Generalized propensity score (Rosenbaum, Rubin JASA 1984)

$$
e^{*}\left(x^{*}\right) \triangleq \mathbb{P}\left(W=1 \mid X^{*}=x^{*}\right)
$$

One model per pattern: $\sum_{m \in\{0,1\}^{d}} \mathbb{E}\left[W \mid X_{o b s(m)}, M=m\right] \mathbb{1}_{M=m}$
$\Rightarrow$ Supervised learning with missing values.

- Mean imputation is consistent with a universally consistent learner.
- Missing Incorporate in Attributes (MIA) for trees methods.

Implemented in grf package: combine two non-parametrics models, forests (conditional outcome and treatment assignment) adapted to any missing values with MIA.
$\hat{\tau}_{\text {AIPW* }}$ is $\sqrt{n}$-consistent, asymptotically normal given the product of RMSE of the nuisance estimates decay as $O\left(n^{-1 / 2}\right)$ Mayer, J. et al. AOAS 2020

## Methods to do causal inference with missing values

|  | Covariates |  | Missingness |  | Unconfoundedness |  |  | Models for$(W, Y)$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | multiva- <br> riate <br> normal | general | $\mathrm{M}(\mathrm{C}) \mathrm{AR}$ | general | Missing | Latent | Classical | logisticlinear | nonparam. |
| 1. (SA)EM ${ }^{4}$ | $\checkmark$ | $x$ | $\checkmark$ | $x$ | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $x$ |
| 1. Mean.GRF | $\checkmark$ | $\checkmark$ | $\checkmark$ | $(\checkmark)$ | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $\checkmark$ |
| 1. MIA.GRF | $\checkmark$ | $\checkmark$ | $\checkmark$ | $(\checkmark)$ | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $\checkmark$ |
| 2. Mult. Imp. | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | ( $X$ ) | $x$ | $\checkmark$ | $\checkmark$ | ( $X$ ) |
| 3. MatrixFact. | $\checkmark$ | $x$ | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $x$ | $\checkmark$ | ( $x$ ) |
| 3. MissDeepCausal | $\checkmark$ | $\checkmark$ | $\checkmark$ | $x$ | $x$ | $\checkmark$ | $x$ | $\checkmark$ | $\checkmark$ |

Methods \& assumptions on data generating process: models for covariates, missing values mechanism, identifiability conditions, models for treatment/outcome.
$\checkmark$ : can be handled $\boldsymbol{X}$ : not applicable in theory
$(\checkmark)$ : empirical results and ongoing work on theoretical guarantees
$(X)$ : no theoretical guarantees but heuristics.
${ }^{4}$ Use of EM algorithms for logistic regression with missing values. Jiang, et al. 2019

## Simulations: no overall best performing method.

- 10 covariates generated with Gaussian mixture model $X_{i} \sim \mathcal{N}_{d}\left(\mu_{\left(c_{i}\right)}, \Sigma_{\left(c_{i}\right)}\right) \mid C_{i}=c_{i}$,
$C$ from a multinomial distribution with three categories.
- Unconfoundedness on complete/observed covariates, 30\% NA
- Logistic-linear for $(W, Y)$, logit $\left(e\left(X_{i .}\right)\right)=\alpha^{T} X_{i .}, Y_{i} \sim \mathcal{N}\left(\beta^{T} X_{i}+\tau W_{i}, \sigma^{2}\right)$

Figure 1: Estimated with AIPW and true ATE $\tau=1$

$\rightarrow$ grf-MIA is asymptotically unbiased under unconfoundedness despite missingness.
$\rightarrow$ Multiple imputation requires many imputations to remove bias.

## Simulations: importance of unconfoundedness assumption and choice of estimator

## Setup

- Different data generating models (linear, nonlinear, latent, etc.)
- Different missingness mechanisms


## Results

- AIPW estimators outperform their IPW counterparts.
- For $\hat{\tau}_{\text {mia }}$, the unconfoundedness despite missingness is indeed necessary.
- $\hat{\tau}_{\text {mia }}$ unbiased for all missingness mechanisms, especially for MNAR.
- Multiple imputation (mice) only requires standard unconfoundedness, but needs MAR


## Results for Trauma Brain Injuries (TBI)

40 covariates, 18 confounders. 8,248 patients.
Overlap: cannot be tested but high level of uncertainty at diagnosing severe (internal bleeding) makes it likely
Many MNAR missing values
ATE estimations ( $\times 100$ ): effect of tranexamic acid on in-ICU mortality

( $y$-axis: estimation approach, solid: Double Robust AIPW, dotted: IPW), ( $x$-axis: ATE estimation with bootstrap Cl )
The obtained value corresponds to the difference in percentage points between mortality rates in treatment and control.

## Results for Trauma Brain Injuries (TBI)

40 covariates, 18 confounders. 8,248 patients.
Overlap: cannot be tested but high level of uncertainty at diagnosing severe (internal bleeding) makes it likely
Many MNAR missing values
ATE estimations $(\times 100)$ : effect of tranexamic acid on in-ICU mortality

( $y$-axis: estimation approach, solid: Double Robust AIPW, dotted: IPW), ( $x$-axis: ATE estimation with bootstrap CI )
Comparison with CRASH-3 study same conclusion of "no average treatment effect".

## Conclusion and perspectives

Take-away messages

- Missing attributes alter causal analyses.
- Additional assumptions on appropriate unconfoundedness.
- New proposals to handle missing values in causal inference.
- Prefer AIPW to IPW estimators, in theory and in practice.
- Heterogeneous treatment effects with missing values (causal forest) implemented in the grf R package


## Ongoing work

- Causal survival analysis, Policy learning (with missing values)
- Combine RCT and observational data to generalize the ATE to a (broader) target population ${ }^{5} 6$

|  | Set | S | $X_{1}$ | $X_{2}$ | $X_{3}$ | $W$ | $Y$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | $\mathcal{R}$ | 1 | 1.1 | 20 | NA | 1 | 24.1 |
| $\cdots$ | $\mathcal{R}$ | 1 |  | $\ldots$ |  | $\ldots$ | $\ldots$ |
| $n-1$ | $\mathcal{R}$ | 1 | -6 | 45 | 8.3 | 0 | 26.3 |
| $n$ | $\mathcal{R}$ | 1 | 0 | 15 | 6.2 | 1 | 23.5 |
| $n+1$ | $\mathcal{O}$ | $?$ | -2 | NA | 7.1 | NA | NA |
| $n+2$ | $\mathcal{O}$ | $?$ | -1 | NA | 2.4 | NA | NA |
| $\cdots \cdots$ | $\mathcal{O}$ | $?$ |  | $\cdots$ |  | NA | NA |
| $n+m$ | $\mathcal{O}$ | $?$ | -2 | NA | 3.4 | NA | NA |

Data with observed treatment $W$ and outcome $Y$ only in the RCT.

## CRASH3

- Multi-centric RCT over 29 counties
- No effect of TXA with difference in means ( -0.3 with [ $95 \% \mathrm{CI}-0.80 .2$ ])
ATE $=-0.035,95 \% \mathrm{Cl}[-0.380 .28]$ when generalizing with $g$-estimator.
Treatment effect modifiers "time to treatment" is missing in Traumabase

[^19]Traumabase

- Representative sample
- 8200 patients with TBI



## Missing value website

More information and details on missing values: R-miss-tastic platform. Mayer, J. et al., 2019
$\rightarrow$ Theoretical and practical tutorials, popular datasets, bibliography, workflows (in R and in python), active contributors/researchers in the community, etc.

```
rmisstastic.netlify.com
```

Interested in contribute to our platform? Feel free to contact us!

MERCI

## An active area of research! Join this exciting field!

Challenges:

- SGD with missing values for linear regression and MCAR ${ }^{43}$. Difficult to extend to logistic or MAR.
- Naively impute the missing values, get $\tilde{X}$,
- Adapt algorithm to account for the error \& apply this debiased version to the complete dataset $\tilde{X}$.
Naive imputation + debiasing also used for Lasso ${ }^{44}$


## Current works

- Times series with missing values for classification
- Model-based Clustering with Missing Not At Random Data
- MNAR missing values - CV with MNAR data? Contribution of causality for missing data

Mohan, Pearl. 2021. Graphical Models for Processing Missing Data. JASA.

Sportisse, Boyer, J. Estimation and imputation in Probabilistic Principal Component
Analysis with Missing Not At Random data. Neurips2020.

[^20]
## Ressources

R-miss-tastic https://rmisstastic.netlify.com/R-miss-tastic
J., I. Mayer, N. Tierney \& N. Vialaneix

Project funded by the R consortium (Infrastructure Steering
Committee) ${ }^{45}$
Aim: a reference platform on the theme of missing data management

- list existing packages
- available literature
- tutorials
- analysis workflows on data
- main actors
$\Rightarrow$ Federate the community
$\Rightarrow$ Contribute!
${ }^{45}$ https://www.r-consortium.org/projects/call-for-proposals


## Ressources

Examples:

- Lecture ${ }^{46}$ - General tutorial : Statistical Methods for Analysis with Missing Data (Mauricio Sadinle)
- Lecture - Multiple Imputation: mice by Nicole Erler ${ }^{47}$
- Longitudinal data, Time Series Imputation (Steffen Moritz - very active contributor of $r$-miss-tastic), Principal Component Methods ${ }^{48}$

```
46}\mathrm{ https://rmisstastic.netlify.com/lectures/
47}\mathrm{ https://rmisstastic.netlify.com/tutorials/erler_course_
multipleimputation_2018/erler_practical_mice_2018
48https://rmisstastic.netlify.com/tutorials/Josse_slides_imputation_PCA_2018.pdf
```


## Thank you



1) Those who can extrapolate from incomplete data

[^0]:    ${ }^{1}$ Doubly robust treatment effect estimation with incomplete confounders. Mayer, Wager, J. Annals Of Applied Statistics 2020.

[^1]:    ${ }^{1}$ Doubly robust treatment effect estimation with incomplete confounders. Mayer, Wager, J. Annals Of Applied Statistics 2020.

[^2]:    ${ }^{2}$ Robin, J, Moulines Sardy. 2019. Low-rank model with covariates for count data with missing values. Journal of Multivariate Analysis.
    ${ }^{3}$ Robin, Slop, J, Moulines Tibshirani. Main effects and interactions in mixed and incomplete data frames. 2019. JASA.

[^3]:    ${ }^{8} \mathrm{~J}$. et al. Main effects and interactions in mixed and incomplete data frames. (2018) JASA.
    ${ }^{9} \mathrm{~J}$. et al. 2018. Imputation of mixed data with multilevel SVD. . JCGS
    ${ }^{10} \mathrm{~J} .$, et al. https://cran.r-project.org/web/views/MissingData.html

[^4]:    ${ }^{11}$ The analysis model may be "in agreement" with the imputation model: congeniality.
    ${ }^{12}$ Little \& Rubin. 2019. Statistical Analysis with Missing Data, 3rd Edition. Wiley

[^5]:    ${ }^{11}$ The analysis model may be "in agreement" with the imputation model: congeniality.
    ${ }^{12}$ Little \& Rubin. 2019. Statistical Analysis with Missing Data, 3rd Edition. Wiley

[^6]:    Ofir Shalev (@ofirdi) May 2018

[^7]:    * Iterativelmputer by default does single imputation with iterative ridge regression

    14 Monte Carlo statistical methods (Robert, Casella, 2004) (p344),
    15 The EM algorithm and extensions (McLachlan, et al. 1998) (p243)

[^8]:    ${ }^{16}$ Stekhoven, Buhlmann. 2012. MissForest - non-parametric missing value imputation for mixed-type data. Bioinformatics

[^9]:    ${ }^{17}$ We assume separability of $\theta$ and $\phi$
    ${ }^{18} z_{\text {obs }}(m)$ is denoted $z_{\text {obs }}$

[^10]:    ${ }^{24}$ Udell \& Townsend. 2019. Why Are Big Data Matrices Approximately Low Rank? SIAM.
    ${ }^{25}$ J. \& Sardy. 2015. Adaptive Shrinkage of singular values. Stat \& Computing.
    ${ }^{26}$ J. \& Wager. 2016. Stable Autoencoding: A Flexible Framework for Regularized Low-Rank Matrix Estimation. JMLR.
    ${ }^{27}$ J. Wager, Sardy. 2016: denoiseR: A Package for Low Rank Matrix Estimation.

[^11]:    ${ }^{29}$ Sportisse, Boyer, J. 2018. Low-rank estimation with missing non at random data. Statistics \& Computing.

[^12]:    ${ }^{30}$ Robin, Klopp, J, Moulines Tibshirani. Main effects and interactions in mixed and incomplete data frames. 2019. JASA.

[^13]:    ${ }^{32} \mathrm{~J}$. et al. 2012. Handling Missing Values with Regularized Iterative Multiple Correspondence Analysis. Journal of classification.

[^14]:    ${ }^{32} \mathrm{~J}$. et al. 2012. Handling Missing Values with Regularized Iterative Multiple Correspondence Analysis. Journal of classification.

[^15]:    ${ }^{32} \mathrm{~J}$. et al. 2012. Handling Missing Values with Regularized Iterative Multiple Correspondence Analysis. Journal of classification.

[^16]:    ${ }^{32} \mathrm{~J}$. et al. 2012. Handling Missing Values with Regularized Iterative Multiple Correspondence Analysis. Journal of classification.

[^17]:    ${ }^{37}$ Timmerman. 2006. Multilevel component analysis. Br J Math Stat Psychol.

[^18]:    ${ }^{1}$ Consistency of supervised learning with missing values J., Prost, Scornet, Varoquaux. 2020
    ${ }^{2}$ Neumiss networks: differential programming for supervised learning with missing values. Le Morvan, J. et al. Neurips2020
    ${ }^{3}$ What's a good imputation to predict with missing values? Le Morvan, J. et al. Neurips 2021

[^19]:    ${ }^{5}$ Colnet, J. et al. (2021). Causal inference methods for combining RCT and observational studies: a review. In revision in Statistical Science - Causal effect on a target population: a sensitivity analysis to handle missing covariates. Submitted

[^20]:    ${ }^{43}$ Sportisse, Boyer, Dieuleveut, J. Debiasing Stochastic Gradient Descent to handle missing values. Neurips2020

