



Robust low-rank covariance matrix estimation with a general pattern of missing values

A. Hippert-Ferrer, M.N. El Korso, Arnaud Breloy, Guillaume Ginolhac

► To cite this version:

A. Hippert-Ferrer, M.N. El Korso, Arnaud Breloy, Guillaume Ginolhac. Robust low-rank covariance matrix estimation with a general pattern of missing values. *Signal Processing*, 2022, 195, pp.108460. 10.1016/j.sigpro.2022.108460 . hal-03761376

HAL Id: hal-03761376

<https://hal.science/hal-03761376>

Submitted on 26 Aug 2022

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

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

See discussions, stats, and author profiles for this publication at: <https://www.researchgate.net/publication/357832494>

Robust low-rank covariance matrix estimation with a general pattern of missing values

Article in *Signal Processing* · January 2022

DOI: 10.1016/j.sigpro.2022.108460

CITATIONS

4

READS

21

4 authors, including:



[Arnaud Breloy](#)

University Paris Nanterre

74 PUBLICATIONS 268 CITATIONS

[SEE PROFILE](#)



[Guillaume Ginolhac](#)

Université Savoie Mont Blanc

145 PUBLICATIONS 841 CITATIONS

[SEE PROFILE](#)

Some of the authors of this publication are also working on these related projects:



G4M - Geodetection Multi Materials Multi Profession [View project](#)



ANR MARGARITA [View project](#)

Robust low-rank covariance matrix estimation with a general pattern of missing values

A. Hippert-Ferrer^{a,d}, M. N. El Korso^b, A. Breloy^b, G. Ginolhac^c

^a*Université Paris-Saclay, CNRS, CentraleSupélec, Laboratoire des signaux et systèmes, 91190, Gif-sur-Yvette, France.*

^b*LEME, Paris-Nanterre University, 50 rue de Sèvres, 92410 Ville d'Avray, France*

^c*LISTIC, Savoie Mont Blanc University, 5 chemin de Bellevue, Annecy, France*

^d*Corresponding author: alexandre.hippert-ferrer@centralesupelec.fr*

Abstract

This paper tackles the problem of robust covariance matrix estimation when the data is incomplete. Classical statistical estimation methodologies are usually built upon the Gaussian assumption, whereas existing robust estimation ones assume unstructured signal models. The former can be inaccurate in real-world data sets in which heterogeneity causes heavy-tail distributions, while the latter does not profit from the usual low-rank structure of the signal. Taking advantage of both worlds, a covariance matrix estimation procedure is designed on a robust (mixture of scaled Gaussian) low-rank model by leveraging the observed-data likelihood function within an expectation-maximization algorithm. It is also designed to handle general pattern of missing values. The proposed procedure is first validated on simulated data sets. Then, its interest for classification and clustering applications is assessed on two real data sets with missing values, which include multispectral and hyperspectral time series.

Keywords: Missing data, covariance matrix, mixture of scaled Gaussian, low-rank, expectation-maximization algorithm, classification.

1. Introduction

Missing data appear when no value of the data is available for a given variable and a given observation. This classical problem [1, 2] is a pitfall in statistical signal processing and its related fields, as statistical inference [3, 4] and data analysis [5, 6]. To name a few applications where missing data has drawn significant attention, we can cite biomedical

studies, chemometrics [7] or remote sensing where missing values created by poor atmospheric conditions or sensor failure can dramatically hamper the understanding of the physical phenomenon under observation [8]. Covariance matrix (CM) estimation theory, which is a fundamental issue in signal processing and machine learning problems, has witnessed particular efforts focusing in the case of incomplete data.

In this scope, it is known that any efficient estimation algorithm should be able to exploit the source of information offered by missing values, which is formerly called *informative missingness* [9]. One approach to estimate the CM with missing values is to rely on maximum likelihood (ML) estimation with a prior assumption on the probability distribution of the data. Within this framework, the Expectation-Maximization (EM) algorithm [10] is a handy iterative procedure to obtain ML estimates as it is based on the expectation of the conditional probability $p(\mathbf{z}|\mathbf{x}, \boldsymbol{\theta})$ of the latent (missing) variables \mathbf{z} given the observed variables \mathbf{x} and the parameters $\boldsymbol{\theta}$ under estimation. Extensive work has been put into CM estimation with missing values using the EM algorithm by assuming independent and identically distributed (iid) samples drawn from the Gaussian distribution when the sample size n exceeds the dimension p ($n > p$) [11, 12] or in high-dimensional regime [13, 14, 15] ($p > n$). Other models have been considered in which the covariance is assumed to have a low-rank plus identity structure [14], which will be referred in the following to as LR structure¹. This structure is closely related to signal subspace inference or principal component analysis (PCA) with missing values [16, 17].

To overcome robustness issues associated with the classical Gaussian assumption, more general distributions have been considered, such as the multivariate t -distribution [18, 19] and its regularized version for small sample size using an improved EM algorithm [20]. These distributions are encompassed by elliptical symmetric (ES) distributions, which are directly linked via ML estimation to the so-called M-estimators in the complete data case [21, 22, 23]. The robustness of these tools to non-Gaussianity [24] has been illustrated in a wide range of applications including radar processing, hyperspectral imagery and classification [25, 26, 27, 28, 29, 30]. Interestingly, this family can be accounted for in a robust (distribution-free) manner by considering the so-called mixture of scaled Gaussian (MSG) [31], which models the data as Gaussian conditionally to an

¹It is sometimes named spiked or factor model.



Figure 1: Illustration of three rectangular data sets with different missing data patterns (black=observed, white=missing): monotone, general and random. Note that the random pattern is a special case of the general pattern, where missing values appear on individual observations rather than blocks. See [35] (p. 5) for a precise description of missing data patterns.

unknown deterministic scale for each sample.

This paper proposes to take advantage of both robust estimation and LR structure models by proposing a EM-based procedure to estimate the CM in the case of incomplete data drawn from a MSG distribution. As cited above, existing robust covariance estimation algorithms based on MSG distributions are not designed to deal with missing values, and, to the best of our knowledge, existing works which use the LR structure are only based on Gaussian assumptions. In the full rank case, the work of [32] has extended generalized elliptical distributions (GES), from which MSG distributions are a sub-class, to incomplete data by proposing an adapted form of the Tyler’s M-estimator.

Furthermore, existing estimation algorithms that handle missing data assume that the missingness pattern, *i.e.*, the pattern describing which values are missing with respect to the observed data, is a monotone pattern [18, 12, 20, 32] (see the illustration of missing data patterns in Fig. 1). This pattern can be of interest, *e.g.*, in longitudinal studies [33] or sensor failure [34]. However, in other applications such as remote sensing, missing values can take very diverse patterns because of unpredictable events (clouds, snowfall, etc.) [8], which leads to the so-called general pattern [3].

With the previous points in mind concerning *i)* robustness to non-Gaussianity, *ii)* low-rank structured models and *iii)* missingness patterns, the contributions of this paper are summarized below:

- 1) A generic algorithm for CM estimation is developed for structured signals with a non-Gaussian distribution. A procedure for non-structured / Gaussian is also obtained as a special case. In the structured configuration, the covariance matrix is supposed to have a LR structure as in [36], which will be detailed in the next section;

- 2) The analysis of missingness patterns is extended to the general pattern, which fills a gap in robust estimation;
- 3) The proposed estimators are tested on both simulated and real data sets with two applications in machine learning, namely supervised classification and unsupervised clustering with missing values, which has rarely been unfolded (outside of simulated data) in the literature of CM estimation with missing values.

The rest of the paper is organized as follows: Section 2 formulates the problem by framing the MSG distribution, the LR CM structure and the missing data model. Section 3 describes the proposed robust CM estimation procedures for unstructured and structured models. Section 4 illustrates the performance of the proposed method on simulated data sets in terms of CM estimation and data imputation. Finally, Section 5 shows the interest of the proposed procedure in real data applications using covariance-based machine learning methods.

Notations. a indicates a scalar quantity, \mathbf{a} represents a vector quantity and \mathbf{A} a matrix. $\{a_i\}_{i=1}^n$ denotes the set of elements a_i with $i \in [1, n]$. The transpose operator is \top , whereas $\text{tr}\{\cdot\}$ and $|\cdot|$ are respectively the trace and the determinant operators. The eigenvalue decomposition of \mathbf{A} into matrices \mathbf{U} and $\mathbf{\Lambda}$ is denoted by $\mathbf{A} \stackrel{\text{EVD}}{=} \mathbf{U}\mathbf{\Lambda}\mathbf{U}^\top$. The symbol \succeq indicates positive definiteness, \mathcal{S}_{++}^p is the set of $p \times p$ symmetric positive definite (SPD) matrices² and $\text{St}_{p,r}$ is the real Stiefel manifold of $p \times r$ orthogonal matrices³. \propto stands for “proportional to”. $Q(\cdot)$ is the score function. Finally, $\mathbb{E}[\cdot]$ denotes the expectation operator.

2. Problem formulation

2.1. Mixture of scaled Gaussian distributions

Most of covariance matrix estimation procedures with incomplete data use the Gaussian assumption. However, this assumption can be inaccurate in remote sensing applications, where images often include heterogeneous areas. The mixture of scaled Gaussian

² $\mathcal{S}_{++}^p = \{\mathbf{\Sigma} \in \mathcal{S}_p : \forall \mathbf{x} \in \mathbb{R}^p \setminus \{0\}, \mathbf{x}^\top \mathbf{\Sigma} \mathbf{x} > 0\}$.

³ $\text{St}_{p,r} = \{\mathbf{U} \in \mathbb{R}^{p \times r} : \mathbf{U}^\top \mathbf{U} = \mathbf{I}_r\}$.

(MSG) distributions can tackle this issue by managing heavier tails, which offers a better fit to empirical data [31]. A real p -vector $\mathbf{y} \in \mathbb{R}^p$ follows a zero-mean multivariate MSG distribution⁴ if it admits the stochastic representation

$$\mathbf{y} = \sqrt{\tau} \mathbf{n} \quad (1)$$

with $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \mathbf{\Sigma})$ and for some scalar $\tau \in \mathbb{R}_{>0}$, called the *texture*, which is strictly positive, deterministic and unknown for each sample.

Let us now define a rectangular data set $\mathbf{Y} \in \mathbb{R}^{p \times n}$ represented by $\mathbf{Y} = \{\mathbf{y}_i = (y_{1,i}, y_{2,i}, \dots, y_{p,i})^\top\}$ where $\{\mathbf{y}_i\}_{i=1}^n$ are modeled as n iid vectors of dimension p drawn from a MSG distribution. This leads to the following model:

$$\mathbf{y}_i | \tau_i \sim \mathcal{N}(\mathbf{0}, \tau_i \mathbf{\Sigma}), \quad \mathbf{\Sigma} \subseteq \mathcal{S}_{++}^p, \quad \tau_i > 0 \quad (2)$$

The loglikelihood function of model (2) is given by

$$\log \ell(\{\mathbf{y}_i\} | \mathbf{\Sigma}, \{\tau_i\}) \propto -n \log |\mathbf{\Sigma}| - p \sum_{i=1}^n \log \tau_i - \sum_{i=1}^n \mathbf{y}_i^\top (\tau_i \mathbf{\Sigma})^{-1} \mathbf{y}_i \quad (3)$$

In this model, the texture can be seen as a *scale* setting of the Gaussian model [37]. Such model is also intertwined with the class of compound Gaussian (CG) distributions (which is a subclass of ES distributions), which assumes the texture independent from \mathbf{n} in (1) and with a given PDF $f_\tau(\cdot)$. Note that the type of distribution within the class of CG distributions is mainly guided by the assumption on the PDF function [31].

Considering deterministic $\{\tau_i\}$ instead of assuming a PDF f_τ has proven its convenience in terms of robustness⁵ [21, 23, 31]. This distribution is also more robust than the purely Gaussian one because the scales allow flexibility in the presence of heterogeneous data, *e.g.*, noisy data, possible outliers or inconsistencies in the data (see, *e.g.*, [40] for Synthetic Aperture Radar data).

⁴Note that this distribution can be easily transposed to complex-valued vectors [31].

⁵This model follows a parametrization of the covariance matrix of the real elliptical model $\mathcal{E}(\mathbf{0}, \tau \boldsymbol{\xi})$, where $\boldsymbol{\xi}$ is the shape matrix and where τ has only one value. In our case, the scale parameter varies at each observation $i = 1, \dots, n$. More details can be found in [38, 24, 39].

Remark 1. As it is clear that $\tau_i = 1, \forall i$ in the deterministic MSG distribution (2) gives the Gaussian distribution, results regarding CM estimation will be given by considering the Gaussian distribution as a special case of the MSG distribution.

Here we assume that Σ is characterized by a LR structure which can be modeled by the well-known factor model [41] (also known as spiked model [42]):

$$\begin{cases} \Sigma = \sigma^2 \mathbf{I}_p + \mathbf{H} \\ \mathbf{H} \succeq \mathbf{0} \\ \text{rank}(\mathbf{H}) = r \end{cases} \quad (4)$$

where \mathbf{I}_p denotes the p -dimensional identity matrix and \mathbf{H} is a $p \times p$ low-rank signal covariance matrix of rank r . As in many works, the latter is considered to be known from prior physical assumptions [36] or pre-estimated as in model order selection techniques [43]. Model (4) is directly related to principal component analysis (PCA) and subspace recovery [44].

2.2. Data model

As each vector \mathbf{y}_i might have missing elements, it is necessary to design a model that takes into account incompleteness. Thereby, each of its observed and missing elements can be grouped in vectors denoted by \mathbf{y}_i^o and \mathbf{y}_i^m respectively, which are stacked in vectors $\tilde{\mathbf{y}}_i$ such that

$$\tilde{\mathbf{y}}_i = \mathbf{P}_i \mathbf{y}_i = \begin{pmatrix} \mathbf{y}_i^o \\ \mathbf{y}_i^m \end{pmatrix}, \quad i = 1, \dots, n \quad (5)$$

where $\mathbf{P}_i \in \mathbb{R}^{p \times p}$ denotes a permutation matrix⁶. As illustrated by Fig. 2, this set of operations consists in permuting the missing elements at the bottom of each vector \mathbf{y}_i . Then, the covariance matrix of $\tilde{\mathbf{y}}_i$ becomes

$$\tilde{\Sigma}_i = \begin{pmatrix} \tilde{\Sigma}_{i,oo} & \tilde{\Sigma}_{i,mo} \\ \tilde{\Sigma}_{i,om} & \tilde{\Sigma}_{i,mm} \end{pmatrix} = \mathbf{P}_i \Sigma \mathbf{P}_i^\top \quad (6)$$

where $\tilde{\Sigma}_{i,mm}$, $\tilde{\Sigma}_{i,mo}$, $\tilde{\Sigma}_{i,oo}$ are the block CM of \mathbf{y}_i^m , of \mathbf{y}_i^m and \mathbf{y}_i^o , and of \mathbf{y}_i^o .

⁶Note that \mathbf{P}_i is invertible and $\mathbf{P}_i^{-1} = \mathbf{P}_i^\top$.




Figure 2: Missing elements (colored) of each \mathbf{y}_i are permuted to fall at the bottom of each $\tilde{\mathbf{y}}_i$. Notice that a completely observed \mathbf{y}_i implies $\mathbf{P}_i = \mathbf{I}_p$.

In a similar manner, the recent work of [45] models (5) by $\tilde{\mathbf{y}}_i = \mathbf{A}_i \mathbf{y}_i$, where \mathbf{A}_i is a $n_i \times p$ selection matrix constructed from the extraction of the $n_i \leq n$ rows of \mathbf{I}_p corresponding to the available observations at the i -th snapshot. This is only an alternative view of the problem. However, the derived results of [45] hold for the Gaussian distribution. In our case, it is worth mentioning that since each \mathbf{y}_i follows a MSG distribution, their permuted version $\tilde{\mathbf{y}}_i$ also follow a MSG distribution because only elements within each p -vector are permuted.

2.3. The EM algorithm: a brief reminder

The EM algorithm is a widely employed iterative scheme for ML estimation in incomplete data problems [10]. This algorithm offers a rigorous and formal approach to the intuitive *ad hoc* idea of filling in missing values⁷: at the E-step, the following conditional expectation of the complete data loglikelihood is found given the observed data and the current deterministic estimated parameters $\boldsymbol{\theta}$:

$$\mathcal{L}_c(\tilde{\mathbf{Y}}|\boldsymbol{\theta}) = \log \ell(\{\mathbf{y}_i\}|\boldsymbol{\Sigma}, \{\tau_i\}) \quad (7)$$

where $\boldsymbol{\theta}$ is the vector of unknown parameters (which are τ_i and $\boldsymbol{\Sigma}$ in our case) and $\tilde{\mathbf{Y}} = [\tilde{\mathbf{y}}_1, \dots, \tilde{\mathbf{y}}_n]$. Note that $\mathcal{L}_c(\tilde{\mathbf{Y}}|\boldsymbol{\theta})$ is equivalent to $\mathcal{L}_c(\mathbf{Y}|\boldsymbol{\theta})$ since both $\tilde{\mathbf{Y}}$ and \mathbf{Y} contain the same missing values, ordered and unordered, respectively. $\tilde{\mathbf{Y}}$ (or \mathbf{Y}) are

⁷Indeed, it is not the missing values themselves that are filled, but the function of the missing values (the sufficient statistics) that are computed [35].

the so-called *complete data* which are the combination of observed and missing data. At the M-step, the parameters are updated by maximizing the expected complete data loglikelihood. To summarize, if $\boldsymbol{\theta}^{(t)}$ is the current estimate of the parameter $\boldsymbol{\theta}$ and $f(\cdot)$ the density function, the E-step computes

$$Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) = \int \mathcal{L}_c(\tilde{\mathbf{Y}}|\boldsymbol{\theta})f(\mathbf{y}^m|\mathbf{y}^o, \boldsymbol{\theta} = \boldsymbol{\theta}^{(t)})d\mathbf{y}^m \quad (8)$$

and the M-step find $\boldsymbol{\theta}^{(t+1)}$ by maximizing (8):

$$Q(\boldsymbol{\theta}^{(t+1)}|\boldsymbol{\theta}^{(t)}) \geq Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \quad (9)$$

The goal is then to repeat E and M-steps until a stopping criteria, such as the distance $\|\boldsymbol{\theta}^{(m+1)} - \boldsymbol{\theta}^{(m)}\|$, converges to a pre-defined threshold.

3. Covariance estimation under non-Gaussian environment

In this section, covariance matrix estimation procedures are developed in the presence of missing data under the MSG distribution. The rank constraint given by (4) is also resolved. For both configurations, the EM algorithm is adopted. In the following, we consider the determinant-based normalization [46] which leads to the estimation of the normalized CM $\boldsymbol{\Sigma}/|\boldsymbol{\Sigma}|^{\frac{1}{p}}$, called the shape matrix [47].

3.1. Robust full-rank estimation

In this case, the unknown parameters to estimate are $\boldsymbol{\theta} = [\boldsymbol{\zeta}^\top, \{\tau_i\}_{i=1}^n]^\top$, where $\boldsymbol{\zeta}$ contains the elements of the lower triangular matrix of $\boldsymbol{\Sigma}$ including its diagonal and where $\{\tau_i\}$ are the texture parameters. As we shall see hereafter, the differences with the purely Gaussian case ($\tau_i = 1$) are that 1) the unknown scales τ_i must be taken into account in the formulation and computation of the expectation at the E-step and 2) a closed-form expression $\hat{\tau}_i$ must be derived to update $\hat{\boldsymbol{\Sigma}}$ at the M-step. These estimators can be found by maximization the loglikelihood of the incomplete data, which is formulated onward. Replacing \mathbf{y}_i by its permuted version $\tilde{\mathbf{y}}_i$ in (3), the complete data loglikelihood of the MSG distribution is alternatively given by:

$$\mathcal{L}_c(\tilde{\mathbf{Y}}|\boldsymbol{\theta}) \propto -n \log |\boldsymbol{\Sigma}| - p \sum_{i=1}^n \log \tau_i - \sum_{i=1}^n \tilde{\mathbf{y}}_i^\top (\tau_i \tilde{\boldsymbol{\Sigma}}_i)^{-1} \tilde{\mathbf{y}}_i \quad (10)$$

At the E-step of the algorithm, the expectation of the complete loglikelihood (10) is computed by using the so-called Q -function. This function is the expectation of the missing data conditioned by the estimation of the parameters at the t -th iteration of the algorithm, that is:

$$Q_i(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) = \mathbb{E}_{\mathbf{y}_i^m|\mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathcal{L}_c(\mathbf{y}_i^o, \mathbf{y}_i^m|\boldsymbol{\theta})] \quad (11)$$

Due to the iid of the observations, one obtains:

$$Q(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) = \sum_{i=1}^n Q_i(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \quad (12)$$

where

$$\begin{aligned} Q_i(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) &= \mathbb{E}_{\mathbf{y}_i^m|\mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} \left[-n \log |\boldsymbol{\Sigma}| - p \log \tau_i - \tilde{\mathbf{y}}_i^\top (\tau_i \tilde{\boldsymbol{\Sigma}}_i)^{-1} \tilde{\mathbf{y}}_i \right] \\ &= \mathbb{E}_{\mathbf{y}_i^m|\mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} \left[-n \log |\boldsymbol{\Sigma}| - p \log \tau_i - \begin{pmatrix} \mathbf{y}_i^o \\ \mathbf{y}_i^m \end{pmatrix}^\top (\tau_i \tilde{\boldsymbol{\Sigma}}_i)^{-1} \begin{pmatrix} \mathbf{y}_i^o \\ \mathbf{y}_i^m \end{pmatrix} \right] \end{aligned} \quad (13)$$

The computation of (13) is hastened as both first and second terms in the expectation are deterministic, which means that only the expectation of the last term has to be computed. To calculate this expectation, the trace $\text{tr}\{\cdot\}$ is used:

$$\begin{aligned} &\mathbb{E}_{\mathbf{y}_i^m|\mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} \left[\begin{pmatrix} \mathbf{y}_i^o \\ \mathbf{y}_i^m \end{pmatrix}^\top (\tau_i \tilde{\boldsymbol{\Sigma}}_i)^{-1} \begin{pmatrix} \mathbf{y}_i^o \\ \mathbf{y}_i^m \end{pmatrix} \right] \\ &= \mathbb{E}_{\mathbf{y}_i^m|\mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} \left[\text{tr} \left\{ \begin{pmatrix} \mathbf{y}_i^o \\ \mathbf{y}_i^m \end{pmatrix} \begin{pmatrix} \mathbf{y}_i^{o\top} & \mathbf{y}_i^{m\top} \end{pmatrix} (\tau_i \tilde{\boldsymbol{\Sigma}}_i)^{-1} \right\} \right] \\ &= \tau_i^{-1} \text{tr} \left\{ \mathbb{E}_{\mathbf{y}_i^m|\mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} \left[\begin{pmatrix} \mathbf{y}_i^o \mathbf{y}_i^{o\top} & \mathbf{y}_i^o \mathbf{y}_i^{m\top} \\ \mathbf{y}_i^m \mathbf{y}_i^{o\top} & \mathbf{y}_i^m \mathbf{y}_i^{m\top} \end{pmatrix} \right] \tilde{\boldsymbol{\Sigma}}_i^{-1} \right\} \\ &= \tau_i^{-1} \text{tr} \{ \mathbf{B}_i^{(t)} \tilde{\boldsymbol{\Sigma}}_i^{-1} \} \end{aligned} \quad (14)$$

where $\mathbf{B}_i^{(t)} = \begin{pmatrix} \mathbf{D}_i^{(t)} & \mathbf{E}_i^{(t)} \\ \mathbf{F}_i^{(t)} & \mathbf{G}_i^{(t)} \end{pmatrix}$ is a $p \times p$ matrix at iteration t of the EM algorithm, with blocks given by

$$\mathbf{D}_i^{(t)} = \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^o \mathbf{y}_i^{o\top}] = \mathbf{y}_i^o \mathbf{y}_i^{o\top} \quad (15)$$

$$\mathbf{E}_i^{(t)} = \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^o \mathbf{y}_i^{m\top}] = \mathbf{y}_i^o \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^{m\top}] \quad (16)$$

$$\mathbf{F}_i^{(t)} = \mathbf{E}_i^{(t)\top} \quad (17)$$

$$\mathbf{G}_i^{(t)} = \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^m \mathbf{y}_i^{m\top}] \quad (18)$$

Thus, the expectation to compute are $\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^m]$ and $\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^m \mathbf{y}_i^{m\top}]$, which are the expectations of the missing data conditioned by the observed data. Note that \mathbf{y}_i^m and $\mathbf{y}_i^m \mathbf{y}_i^{m\top}$ are the sufficient statistics for the complete data $\tilde{\mathbf{Y}}$, which have a distribution from the regular exponential family [35]. Here, using a classical result on conditional distributions (see Theorem 2.5.1 in [48], p. 35), the conditional distribution of \mathbf{y}_i^m given \mathbf{y}_i^o is

$$\mathbf{y}_i^m | \mathbf{y}_i^o \sim \mathcal{N}(\tilde{\boldsymbol{\mu}}_{i,m|o}, \tilde{\boldsymbol{\Sigma}}_{i,mm|o}) \quad (19)$$

where $\tilde{\boldsymbol{\mu}}_{i,m|o} = \tilde{\boldsymbol{\Sigma}}_{i,mo} \tilde{\boldsymbol{\Sigma}}_{i,oo}^{-1} \tilde{\mathbf{y}}_i^o$ and $\tilde{\boldsymbol{\Sigma}}_{i,mm|o} = \tau_i (\tilde{\boldsymbol{\Sigma}}_{i,mm} - \tilde{\boldsymbol{\Sigma}}_{i,mo} \tilde{\boldsymbol{\Sigma}}_{i,oo}^{-1} \tilde{\boldsymbol{\Sigma}}_{i,om})$ are the conditional mean and covariance matrix, respectively. Note that τ_i is absent from $\tilde{\boldsymbol{\mu}}_{i,m|o}^{(t)}$ as it is annihilated by its inverse. Following [35], at the t -th iteration, the E-step of the EM algorithm consists in calculating

$$\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^m] = \tilde{\boldsymbol{\mu}}_{i,m|o}^{(t)} \quad (20)$$

$$\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathbf{y}_i^m \mathbf{y}_i^{m\top}] = \tilde{\boldsymbol{\Sigma}}_{i,mm|o}^{(t)} + \tilde{\boldsymbol{\mu}}_{i,m|o}^{(t)} \tilde{\boldsymbol{\mu}}_{i,m|o}^{(t)\top} \quad (21)$$

Finally, $\boldsymbol{\theta}^{(t+1)}$ is obtained at the M-step of the algorithm as the solution of the following maximization problem:

$$\begin{aligned} & \max_{\boldsymbol{\theta}} \quad Q_i(\boldsymbol{\theta} | \boldsymbol{\theta}^{(t)}) \\ & \text{subject to} \quad \boldsymbol{\Sigma} \succeq \mathbf{0} \\ & \quad \quad \tau_i > 0, \quad \forall i \end{aligned} \quad (22)$$

Proposition 3.1. *The ML estimates $\widehat{\Sigma}$ and $\widehat{\tau}_i$ of problem (22) are given by the following closed-form expressions:*

$$\widehat{\tau}_i = \frac{\text{tr}\{B_i^{(t)} \bar{\Sigma}_i^{-1}\}}{p} \quad \text{for } i \in [1, n] \quad (23)$$

$$\widehat{\Sigma} = \frac{p}{n} \sum_{i=1}^n \frac{C_i^{(t)\top}}{\text{tr}\{C_i^{(t)} \widehat{\Sigma}^{-1}\}} \triangleq \mathcal{H}(\widehat{\Sigma}) \quad (24)$$

with $\bar{\Sigma}_i = P_i \widehat{\Sigma} P_i^\top$, $C_i^{(t)} = P_i^\top B_i^{(t)} P_i$ and where $\mathcal{H}(\cdot)$ is the fixed point equation.

Proof. Let us start by writing the expectation of the loglikelihood $\mathcal{L}_c(\tilde{Y}|\theta)$ given by (10):

$$\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \theta^{(t)}} [\mathcal{L}_c(\tilde{Y}|\theta)] \propto \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \theta^{(t)}} \left[-n \log |\Sigma| - p \sum_{i=1}^n \log \tau_i - \sum_{i=1}^n \tilde{\mathbf{y}}_i^\top (\tau_i \tilde{\Sigma}_i)^{-1} \tilde{\mathbf{y}}_i \right] \quad (25)$$

Notwithstanding equation (25) is similar to (13), it differs as it is the expectation conditioned on all observations contained in \tilde{Y} , whereas the latter stands for only one observation. By using the result (14) obtained from (13), it follows that:

$$\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \theta^{(t)}} [\mathcal{L}_c(\tilde{Y}|\theta)] \propto -n \log |\Sigma| - p \sum_{i=1}^n \log \tau_i - \sum_{i=1}^n \tau_i^{-1} \text{tr}\{B_i^{(t)} \tilde{\Sigma}_i^{-1}\} \quad (26)$$

Again, the block matrix $B_i^{(t)}$ contains the expectation of the sufficient statistics $\mathbf{y}_i^m \mathbf{y}_i^{m\top}$ to be computed at the E-step using (15). Now, let differentiate (26) with respect to τ_i and resolve the equality:

$$\frac{\partial \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \theta^{(t)}} [\mathcal{L}_c(\tilde{Y}|\theta)]}{\partial \tau_i} = \mathbf{0}$$

This calculus is trivial and finally gives

$$\widehat{\tau}_i = \frac{\text{tr}\{B_i^{(t)} \bar{\Sigma}_i^{-1}\}}{p}$$

with $\bar{\Sigma}_i = P_i \widehat{\Sigma} P_i^\top$, which is the desired expression for $\widehat{\tau}_i$. To find $\widehat{\Sigma}$, we first replace τ_i by $\widehat{\tau}_i$ in (26) and obtain the following loglikelihood:

$$\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \theta^{(t)}} [\mathcal{L}_c(\tilde{Y}|\theta)] \propto -n \log |\Sigma| - p \sum_{i=1}^n \log \frac{\text{tr}\{B_i^{(t)} \tilde{\Sigma}_i^{-1}\}}{p} - np \quad (27)$$

Then, we derive (27) w.r.t. Σ , which leads to the resolution of the following equation:

$$\frac{\partial \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \boldsymbol{\theta}^{(t)}} [\mathcal{L}_c(\tilde{\mathbf{Y}} | \boldsymbol{\theta})]}{\partial \Sigma} = \mathbf{0}$$

By using the definition $\tilde{\Sigma}_i = \mathbf{P}_i \Sigma \mathbf{P}_i^\top$, one obtains first:

$$-n \Sigma^{-1} - p \sum_{i=1}^n \frac{\Sigma^{-1} \mathbf{C}_i^{(t)\top} \Sigma^{-1}}{\text{tr}\{\mathbf{C}_i^{(t)} \Sigma^{-1}\}} = \mathbf{0}$$

where $\mathbf{C}_i^{(t)} = \mathbf{P}_i^\top \mathbf{B}_i^{(t)} \mathbf{P}_i$. Arranging terms, then multiplying right and left terms by Σ gives the desired closed-form expression of $\hat{\Sigma}$ and concludes the proof. \square

Remark 2. For $n > p$, which is our case, this estimator can be computed using the fixed point algorithm $\Sigma_{m+1} = \mathcal{H}(\Sigma_m)$ where m refers to the iteration index of the fixed point.

Remark 3. In the Gaussian case, finding Σ that maximizes this expression requires to derive Q with respect to Σ and then to solve $\frac{\partial Q}{\partial \Sigma} = \frac{\partial \mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \Sigma^{(t)}} [\mathcal{L}_c(\tilde{\mathbf{Y}} | \Sigma)]}{\partial \Sigma} = 0$ where $\mathbb{E}_{\mathbf{y}_i^m | \mathbf{y}_i^o, \Sigma^{(t)}} [\mathcal{L}_c(\tilde{\mathbf{Y}} | \Sigma)] \propto -n \log |\Sigma| - \sum_{i=1}^n \text{tr}\{\mathbf{B}_i^{(t)} \tilde{\Sigma}_i^{-1}\}$. Based on the above and after some basic derivation calculus, the solution to be computed at each iteration of the EM algorithm is given by:

$$\hat{\Sigma} = \frac{1}{n} \mathbf{C}^{(t)\top} \quad (28)$$

with $\mathbf{C}^{(t)} = \sum_{i=1}^n \mathbf{P}_i^\top \mathbf{B}_i^{(t)} \mathbf{P}_i$. One can notice that in the case of no missing values, this result leads to the classical Sample Covariance Matrix (SCM) with $\mathbf{P}_i = \mathbf{I}_p$ and $\mathbf{B}_i^{(t)} = \mathbb{E}[\mathbf{y}_i \mathbf{y}_i^\top]$.

The complete estimation procedure of $\boldsymbol{\theta}$ is given in Algorithm 1. The stopping condition of the EM algorithm is ensured by the evaluation of the quantity $\|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\|_F^2$ at each iteration, while the convergence of the fixed point algorithm relies on $\|\Sigma_{m+1}^{(t)} - \Sigma_m^{(t)}\|_F^2$.

At the step $t = 0$, the estimate $\Sigma^{(0)}$ is initialized with Tyler's estimator from available observations in their full dimension p , denoted $\hat{\Sigma}_{\text{Tyl-obs}}$. Unlike Σ , incomplete observations make the direct estimation of $\boldsymbol{\tau}$ with the fixed point impracticable: thus, as an initialization, all $\tau_i^{(0)}$ are set to one.

Algorithm 1 EM-Tyl: Estimation of $\boldsymbol{\theta}$ under MSG distribution with missing values.

Require: $\{\tilde{\mathbf{y}}_i\}_{i=1}^n \sim \mathcal{N}(\mathbf{0}, \tau_i \boldsymbol{\Sigma}), \{\mathbf{P}_i\}_{i=1}^n$

Ensure: $\hat{\boldsymbol{\Sigma}}, \{\hat{\tau}_i\}_{i=1}^n$

1: Initialization:

$$\boldsymbol{\Sigma}^{(0)} = \hat{\boldsymbol{\Sigma}}_{\text{Tyl-obs}}$$

$$\boldsymbol{\tau}^{(0)} = \mathbf{1}_N^\top$$

2: **repeat**

▷ EM loop, t varies

3: Compute

$$\mathbf{E}_i^{(t)} = \tilde{\mathbf{y}}_i^o \tilde{\boldsymbol{\Sigma}}_{i,mo}^{(t)} \tilde{\boldsymbol{\Sigma}}_{i,oo}^{-1(t)} \tilde{\mathbf{y}}_i^o$$

$$\mathbf{G}_i^{(t)} = \tau_i^{(t)} (\tilde{\boldsymbol{\Sigma}}_{i,mm}^{(t)} - \tilde{\boldsymbol{\Sigma}}_{i,mo}^{(t)} \tilde{\boldsymbol{\Sigma}}_{i,oo}^{-1(t)} \tilde{\boldsymbol{\Sigma}}_{i,om}^{(t)}) + \tilde{\boldsymbol{\Sigma}}_{i,mo}^{(t)} \tilde{\boldsymbol{\Sigma}}_{i,oo}^{-1(t)} \tilde{\mathbf{y}}_i^o (\tilde{\boldsymbol{\Sigma}}_{i,mo}^{(t)} \tilde{\boldsymbol{\Sigma}}_{i,oo}^{-1(t)} \tilde{\mathbf{y}}_i^o)^\top$$

4: Compute $\mathbf{B}_i^{(t)} = \begin{pmatrix} \mathbf{y}_i^o \mathbf{y}_i^{o\top} & \mathbf{E}_i^{(t)} \\ \mathbf{E}_i^{(t)\top} & \mathbf{G}_i^{(t)} \end{pmatrix}$

5: Compute $\mathbf{C}_i^{(t)} = \mathbf{P}_i^\top \mathbf{B}_i^{(t)} \mathbf{P}_i$

6: **repeat**

▷ fixed point, m varies (optional loop)

7: $\hat{\boldsymbol{\Sigma}}_{m+1}^{(t)} = \mathcal{H}(\hat{\boldsymbol{\Sigma}}_m^{(t)})$

8: **until** $\|\boldsymbol{\Sigma}_{m+1}^{(t)} - \boldsymbol{\Sigma}_m^{(t)}\|_F^2$ converges

9: Compute $\hat{\tau}_i^{(t)}, \quad i = 1, \dots, n$

10: $t \leftarrow t + 1$

11: **until** $\|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\|_F^2$ converges

3.2. Robust low-rank estimation

Let us now consider the case of data whose covariance matrix lives in a subspace of dimension $r < p$. As we recall, this configuration is useful in many real signal applications and is closely related to principal component analysis (PCA) and signal subspace estimation [44].

Following model (4), the parameters to estimate are now given by $\boldsymbol{\theta} = [\boldsymbol{\zeta}^\top, \sigma^2, \{\tau_i\}_{i=1}^n]^\top$, where $\boldsymbol{\zeta}$ contains the elements of the lower triangular matrix of \mathbf{H} . The maximization problem (22) to find $\boldsymbol{\theta}^{(t+1)}$ at the M-step of the EM algorithm becomes the following low-rank estimation problem:

$$\begin{aligned} \max_{\boldsymbol{\theta}} \quad & Q_i(\boldsymbol{\theta}|\boldsymbol{\theta}^{(t)}) \\ \text{subject to} \quad & \boldsymbol{\Sigma} = \sigma^2 \mathbf{I}_p + \mathbf{H} \\ & \text{rank}(\boldsymbol{\Sigma}) = r \\ & \sigma > 0, \quad \tau_i > 0, \quad \forall i \end{aligned} \tag{29}$$

A general solution to this problem was found in the seminal work of [44], which is recalled subsequently using the formulations of [49] and [50]. Firstly, let $\boldsymbol{\Sigma}^{(t)} \stackrel{\text{EVD}}{=} \sum_{i=1}^p \lambda_i^{(t)} \mathbf{u}_i^{(t)} \mathbf{u}_i^{(t)\top}$ be the eigenvalue decomposition of $\boldsymbol{\Sigma}^{(t)}$ at the t -th iteration of the EM algorithm, where $\lambda_1^{(t)} < \dots < \lambda_p^{(t)}$ are the eigenvalues of $\boldsymbol{\Sigma}^{(t)}$ and $\mathbf{u}_1^{(t)}, \dots, \mathbf{u}_p^{(t)}$ the corresponding eigenvectors. Then, the solution of (29) is given by:

$$\hat{\boldsymbol{\Sigma}} = \hat{\sigma}^2 \mathbf{I}_p + \sum_{i=1}^r \hat{\lambda}_i \mathbf{u}_i^{(t)} \mathbf{u}_i^{(t)\top} \tag{30}$$

where

$$\hat{\sigma}^2 = \frac{1}{p-r} \sum_{i=r+1}^p \lambda_i^{(t)} \tag{31}$$

$$\hat{\lambda}_i = \lambda_i^{(t)} - \hat{\sigma}^2, \quad i = 1, \dots, r \tag{32}$$

The procedure to estimate $\boldsymbol{\theta}$ under the LR assumption is given in Algorithm 2. It uses the same form than Algorithm 1 where equations (30)–(32) are applied just after the fixed point algorithm to each newly estimated $\boldsymbol{\Sigma}$. As stated in [50], the proof of monotonicity of the low-rank estimation algorithm is ensured by standards convergence results of the

majorization-minimization (MM) algorithm [51]. Finally, the estimation of the scales $\boldsymbol{\tau}$ does not change from the former full-rank algorithm.

Algorithm 2 EM-Tyl-r: low-rank estimation of $\boldsymbol{\theta}$ under MSG distribution with missing values.

Require: $\{\tilde{\mathbf{y}}_i\}_{i=1}^n \sim \mathcal{N}(\mathbf{0}, \tau_i \boldsymbol{\Sigma}), \{\mathbf{P}_i\}_{i=1}^n, \text{rank } r < p$

Ensure: $\hat{\boldsymbol{\Sigma}}, \{\hat{\tau}_i\}_{i=1}^n$

- 1: Initialize $\boldsymbol{\Sigma}^{(0)}, \boldsymbol{\tau}^{(0)}$ as in Algorithm 1.
 - 2: **repeat** ▷ EM loop, t varies
 - 3: Compute $\mathbf{E}_i^{(t)}, \mathbf{G}_i^{(t)}, \mathbf{B}_i^{(t)}$ and $\mathbf{C}_i^{(t)}$ as in Algorithm 1.
 - 4: **repeat** ▷ fixed point, m varies (optional loop)
 - 5: $\hat{\boldsymbol{\Sigma}}_{m+1}^{(t)} = \mathcal{H}(\hat{\boldsymbol{\Sigma}}_m^{(t)})$
 - 6: $\hat{\boldsymbol{\Sigma}}_{m+1}^{(t)} \stackrel{\text{EVD}}{=} \sum_{i=1}^p \lambda_i^{(t)} \mathbf{u}_i^{(t)} \mathbf{u}_i^{\top(t)}$
 - 7: Update $\hat{\boldsymbol{\Sigma}}_{m+1}^{(t)}$ by computing (30), (31), (32)
 - 8: $\hat{\boldsymbol{\Sigma}}_{m+1}^{(t)} = \hat{\boldsymbol{\Sigma}}_{m+1}^{(t)} / \text{tr}(\hat{\boldsymbol{\Sigma}}_{m+1}^{(t)})$
 - 9: **until** $\|\boldsymbol{\Sigma}_{m+1}^{(t)} - \boldsymbol{\Sigma}_m^{(t)}\|_F^2$ converges
 - 10: Compute $\hat{\tau}_i^{(t)}, \quad i = 1, \dots, n$
 - 11: $t \leftarrow t + 1$
 - 12: **until** $\|\boldsymbol{\theta}^{(t+1)} - \boldsymbol{\theta}^{(t)}\|_F^2$ converges
-

3.3. Implementation details

In both Algorithm 1 and Algorithm 2, the set of permutations matrices $\{\mathbf{P}_i\}_{i=1}^n$ is needed as an input to compute the transformed covariance matrices $\tilde{\boldsymbol{\Sigma}}_i = \mathbf{P}_i \boldsymbol{\Sigma} \mathbf{P}_i^\top$. Indeed, if a vector \mathbf{y}_i is fully observed, one finds that $\mathbf{P}_i = \mathbf{I}_p$, which makes the computation of $\tilde{\boldsymbol{\Sigma}}$ pointless. To avoid these extra computations, one can decompose the loglikelihood (10) in the following way:

$$\mathcal{L}_c(\tilde{\mathbf{Y}}|\boldsymbol{\theta}) \propto -n \log |\boldsymbol{\Sigma}| - p \sum_{i=1}^n \log \tau_i - \sum_{i \in N_o} \mathbf{y}_i^\top (\tau_i \boldsymbol{\Sigma}_i)^{-1} \mathbf{y}_i - \sum_{i \in N_m} \tilde{\mathbf{y}}_i^\top (\tau_i \tilde{\boldsymbol{\Sigma}}_i)^{-1} \tilde{\mathbf{y}}_i \quad (33)$$

where $N_o = \{i_l\}_{l=1}^{n_o}$ and $N_m = \{i_l\}_{l=1}^{n_m}$ are the sets of indices corresponding to fully observed $\{\mathbf{y}_i\}$ and partially observed $\{\tilde{\mathbf{y}}_i\}$, respectively, with $n = n_o + n_m$. Thenceforth,

only the set of permutations matrices $\{\mathbf{P}_i\}_{i=1}^{n_m}$ need to be computed. An important consequence is that if \mathbf{y}_i is fully observed, $\mathbf{B}_i^{(t)} = \mathbf{y}_i^o \mathbf{y}_i^{o\top} = \mathbf{y}_i \mathbf{y}_i^\top$. $\mathbf{C}_i^{(t)}$ becomes:

$$\mathbf{C}_i^{(t)} = \begin{cases} \mathbf{I}_p^\top \mathbf{B}_i^{(t)} \mathbf{I}_p = \mathbf{y}_i \mathbf{y}_i^\top, & i \in N_o \\ \mathbf{P}_i^\top \mathbf{B}_i^{(t)} \mathbf{P}_i, & i \in N_m \end{cases} \quad (34)$$

Empirical experiments have shown that such care in the computation of $\mathbf{C}_i^{(t)}$ depending on $i \in N_o$ or $i \in N_m$ accelerates the running time of the EM algorithm.

Furthermore, the fixed point loop is optional as it can be seen as an inner EM where $\{\tau_i\}$ is the set of latent variables (hence a single update still increases the likelihood). Our empirical experiments evidenced that both versions of the algorithm achieve similar performance, while performing only a single fixed point iteration tends to achieve a faster convergence.

4. Numerical simulations

This section illustrates the validation of the proposed algorithms with numerical experiments on simulated data drawn from the multivariate Gaussian and mixture of scaled Gaussian distributions. The performances of the proposed covariance estimation procedure are evaluated in regard to three aspects corresponding to different experiments:

- 1) The missing data ratio and patterns in subsection 4.1;
- 2) The quantity of outliers corrupting the data in subsection 4.2;
- 3) The possibility to perform CM-based data imputation in the presence of outliers and missing data in subsection 4.3.

For the sake of the experiment, the CM and scales parameters are known. The CM \mathbf{R} is chosen to be Toeplitz⁸, which has the form:

$$(\mathbf{R})_{ij} = \rho^{|i-j|} \quad (35)$$

⁸Note that this information on the structure is not used in the estimation procedure.

Parameter ρ , which controls the correlation structure of the CM, is set to 0.7. Scales $\{\tau_i\}_{i=1}^n$ are drawn from a Gamma distribution with shape parameter α and scale parameter $\beta = \frac{1}{\alpha}$. In all experiments, we fix $\alpha = 1$. To generate a covariance matrix admitting the structure (4), we compute:

$$\mathbf{\Sigma} = \mathbf{I}_p + \sigma^2 \mathbf{U} \mathbf{U}^\top \quad (36)$$

where $\mathbf{U} \in \text{St}_{p,r}$ is the underlying signal subspace basis obtained from the eigenvalue decomposition of \mathbf{R} and σ is a free parameter corresponding to the signal to noise ratio.

In our experiments, the data dimension is $p = 15$ and $n = \{63, 109, 190, 331, 575, 1000\}$. Sets $\{\mathbf{y}_i\}_{i=1}^n$ are drawn from the MSG distribution with covariance $\mathbf{\Sigma}$. As the aim is to estimate the structured covariance matrix $\mathbf{\Sigma}$, the estimation is performed on 500 sets $\{\mathbf{y}_i\}_{i=1}^n$ simulated for each value of n . Indeed, for the sake of the experiment, missing data are also simulated, which allows a full control on their ratio and pattern. Importantly, as n increases from 63 to 1000, the missing data ratio decreases as $\{44, 22, 11, 5, 2, 1\}\%$.

The following covariance estimators, which are reported in Table ??, are considered for comparison:

- i) The covariance matrix for MSG distributions obtained from Algorithm 1 (full-rank) and from Algorithm 2 (low-rank), named $\hat{\mathbf{\Sigma}}_{\text{EM-Ty1}}$ and $\hat{\mathbf{\Sigma}}_{\text{EM-Ty1-r}}$, respectively.
- ii) The covariance matrix for the Gaussian distribution obtained from Algorithm 1 (full-rank) and from Algorithm 2 (low-rank) with the solution given by (28), named $\hat{\mathbf{\Sigma}}_{\text{EM-SCM}}$ and $\hat{\mathbf{\Sigma}}_{\text{EM-SCM-r}}$, respectively.
- iii) The *sample covariance matrix* (SCM) from the clairvoyant data (without missing data):

$$\hat{\mathbf{\Sigma}}_{\text{SCM-clair}} = \frac{1}{n} \sum_{i=1}^n \mathbf{y}_i \mathbf{y}_i^\top \quad (37)$$

and its low-rank version, named $\hat{\mathbf{\Sigma}}_{\text{SCM-clair-r}}$.

- iv) The SCM $\hat{\mathbf{\Sigma}}_{\text{SCM-obs}}$ estimated from the set of fully observed vectors $\{\mathbf{y}_{i_l}\}_{l=1}^{n_o} \in N_o$ where $N_o = \{i_l\}_{l=1}^{n_o}$.

v) The *fixed point* estimator or Tyler's estimator from the clairvoyant data:

$$\hat{\Sigma}_{\text{Ty1-clair}} = \frac{p}{n} \sum_{i=1}^n \frac{\mathbf{y}_i \mathbf{y}_i^\top}{\mathbf{y}_i \Sigma^{-1} \mathbf{y}_i^\top} \quad (38)$$

and its low-rank version, named $\hat{\Sigma}_{\text{Ty1-clair-r}}$.

vi) Tyler's estimator $\hat{\Sigma}_{\text{Ty1-obs}}$ from the set of fully observed vectors $\{\mathbf{y}_{i_l}\}_{l=1}^{n_o} \in N_o$ with $N_o = \{i_l\}_{l=1}^{n_o}$.

vii) Multiple imputation (MI) [52, 6] from which a robust version is proposed (RMI): for each \mathbf{y}_i with missing values \mathbf{y}_i^m , q vectors with imputed missing entries (referred to as *imputed vectors* hereafter) $(\hat{\mathbf{y}}_{i1}, \hat{\mathbf{y}}_{i2}, \dots, \hat{\mathbf{y}}_{iq})$ are generated with missing values drawn from a MSG distribution:

$$\hat{\mathbf{y}}_{ij}^m \sim \mathcal{N}(\boldsymbol{\mu}_{ij}^o, \sqrt{\tau_{ij}} \boldsymbol{\sigma}_{ij}^o), \quad j = 1, \dots, q \quad (39)$$

where $\boldsymbol{\mu}_{ij}^o$ and $\boldsymbol{\sigma}_{ij}^o$ are the mean and variance of the observed components of \mathbf{y}_{ij} , and τ_{ij} are the scales parameters drawn from the Gamma distribution with shape parameter $\alpha = 1$. The estimated covariance is the mean of the q Tyler's estimators from the q imputed vectors:

$$\hat{\Sigma}_{\text{RMI}} = \frac{1}{q} \sum_{j=1}^q \hat{\Sigma}_{\text{Ty1},j} = \frac{p}{nq} \sum_{j=1}^q \sum_{i=1}^n \frac{\hat{\mathbf{y}}_{ij} \hat{\mathbf{y}}_{ij}^\top}{\hat{\mathbf{y}}_{ij} \hat{\Sigma}_{\text{Ty1},j}^{-1} \hat{\mathbf{y}}_{ij}^\top} \quad (40)$$

The low-rank version is also estimated, named $\hat{\Sigma}_{\text{RMI-r}}$.

viii) Robust stochastic imputation (RSI): this procedure is tantamount to MI, but with $q = 1$.

ix) Mean imputation [6]: missing components of each vector \mathbf{y}_i are imputed by the mean of the observed components:

$$\hat{\mathbf{y}}_i^m = \boldsymbol{\mu}_i^o \quad (41)$$

Then, the covariance $\hat{\Sigma}_{\text{Mean-Ty1}}$ is estimated using Tyler's estimator. The low-rank version is also estimated, named $\hat{\Sigma}_{\text{Mean-Ty1-r}}$.

Type	Estimator	Description	Low-rank
EM	EM-Tyl	Robust covariance estimation (this study)	yes
	EM-SCM	Covariance estimation (this study)	yes
NO MISSING DATA	Tyl-clair	Tyler’s estimator on full data [21]	yes
	Tyl-obs	Tyler’s estimator on fully observed data [21]	no
	SCM-clair	Sample covariance matrix on full data	yes
	SCM-obs	Sample covariance matrix on fully observed data	no
IMPUTATION	RMI	Tyler’s estimator on multiple imputed data [52, 6]	yes
	RSI	Tyler’s estimator on stochastic imputed data	yes
	Mean-Tyl	Tyler’s estimator on data with mean imputation [6]	yes

Table 1: List of estimators used in numerical experiments. In the experiments, low-rank estimators keep their names with an additional “-r”, *e.g.*, EM-SCM-r.

4.1. Covariance estimation with missing data

This experiment shows the performances of covariance estimation as functions of the missing data ratio and missing data patterns. To compare the estimated covariance matrix to the true data covariance, the following geodesic distance is used [53]:

$$\delta_{\mathcal{S}_{++}^p}^2(\Sigma, \hat{\Sigma}) = \|\log(\Sigma^{-\frac{1}{2}} \hat{\Sigma} \Sigma^{-\frac{1}{2}})\|_2^2 \quad (42)$$

This distance, which is sufficient to measure estimation errors, emanates from the Fisher metric for the Gaussian distribution on \mathcal{S}_{++}^p [54].

Three missing data patterns are examined, which correspond to the following configurations (see Fig. 1):

- One block of missing data of size (7×20) . This is a special case of the monotone missing data pattern as studied in [12, 20]. This configuration corresponds to the case where one group of variables is missing at the same observations, *e.g.*, a group of sensors that are equally interrupted in time. This case is referred to as *monotone pattern* in the experiments.

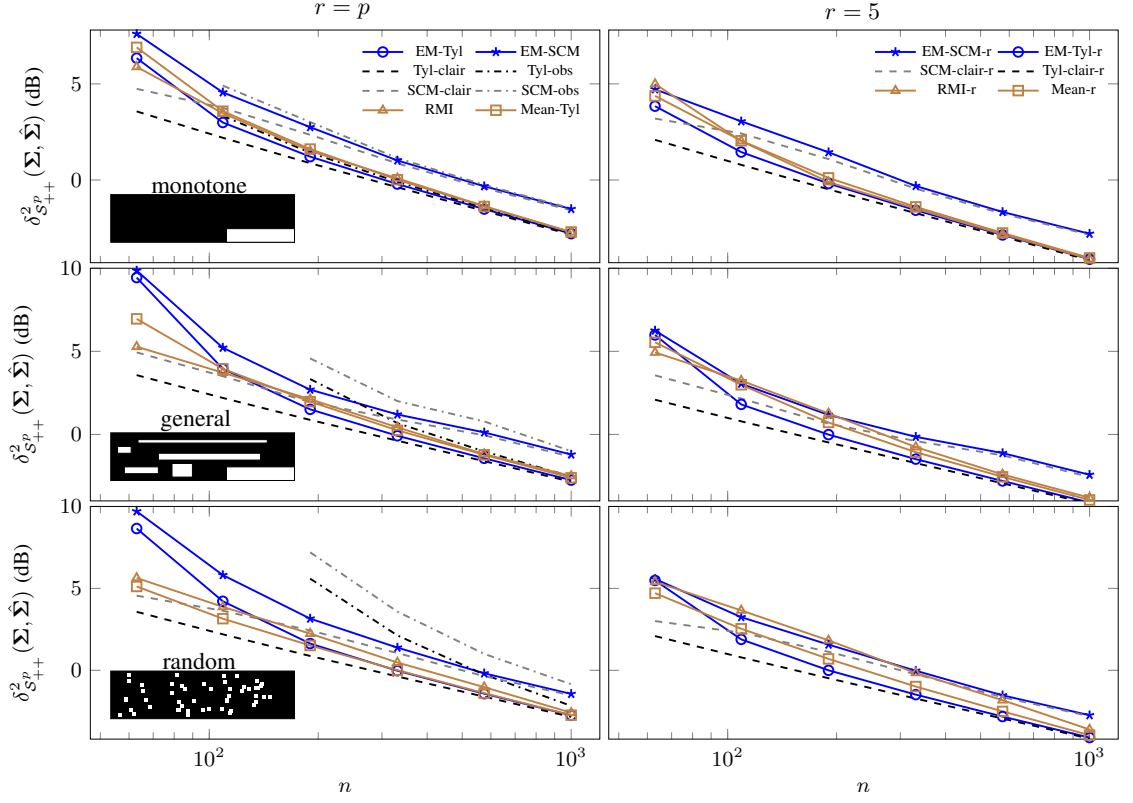


Figure 3: Mean of error measures 42 for $r = p$ (left) and $r = 5$ (right) of methods EM-Tyl, EM-Tyl-r, EM-SCM, EM-SCM-r, Tyl-clair, Tyl-obs, SCM-clair, SCM-obs, RMI and mean imputation as functions of the number of samples n . The mean are computed over 500 simulated sets $\{\mathbf{y}_i\}_{i=1}^n$ for monotone (top), general (middle) and random (bottom) missing data pattern with $p = 15$.

- Multiple blocks of missing data with various size and randomly scattered into the data set. This configuration corresponds to the general missing data pattern. It is most likely to happen in real-life applications (see, *e.g.*, [55, 8]). This case is referred to as *general pattern*.
- Randomly distributed missing data. This configuration is also a general missing data pattern except that multiple values for one variable at one observation are missing across the dataset. This case is referred to as *random pattern*.

Fig. 3 and 4 show the results in terms of mean distance for two rank settings (the full rank case $r = p$ and the low-rank case with $r = 5$) and the three aforementioned pattern




Figure 4: Boxplots showing geodesic distances of Tyler, Tyler observed, EM (ours), mean imputation, multiple imputation and stochastic imputation methods for $r = p$ and $n = 200$. Small dots show the computed distances over 500 simulated sets $\{\mathbf{y}_i\}_{i=1}^n$ corrupted by 20% of missing data with general pattern, which corresponds to the middle-left plot of Fig. 3.

configurations. In the full rank case, EM-Tyl shows better performances than all other estimators, but fails when the number of available observations is small, especially when it gets close to p .

For the random pattern configuration, this result is even more notable and when n increases, EM-Tyl exhibits similar errors than mean imputation. Good performances of the latter can be explained by the increasing availability of observations as n increases, which gives a better estimate of $\boldsymbol{\mu}_i^o$. Unsurprisingly, all estimators (except the ones based on the SCM) reach Tyler’s estimator on clairvoyant data as the missing data ratio decreases in large sample size. In the case of Gaussian estimators, the EM-SCM estimator shows very close performances than the clairvoyant SCM. As expected, Tyl-obs and SCM-obs display poor performances when the missing data ratio increases.

Results for the low-rank model are illustrated for $r = 5$. In all cases, we observe that EM-Tyl-r is a better estimator than EM-SCM-r in terms of distance. For small n , the gap between the EM-Tyl-r estimator and the imputation strategies (RMI-r and Mean-Tyl-r) diminishes compared to the full-rank case, and EM-Tyl-r significantly gives better estimates when the number of samples increases.

In conclusion, depending on the available sample size and the missing data pattern at hand, different estimators are preferable: imputation strategies are more advantageous for small amount of observed samples whereas our estimator performs better as the number of observed samples grows and when the covariance matrix has a low-rank structure.

4.2. Should outliers be discarded?

Data corrupted by outliers is one of the main motivation of robust estimation [56]. In this case, one can be interested to know whether Tyler’s estimator or EM-Tyl estimator gives higher performances. In the complete data case, the former can be directly estimated with outliers. To estimate the latter, outliers can be discarded or masked (*i.e.*, set as missing data) if their position are known, which is a common approach in various applications. In the following, a data set $\{\mathbf{y}_i\}_{i=1}^n$ is drawn from the MSG distribution with a fixed number of samples $n = 200$ and with $p = 15$. Outliers are generated by adding white Gaussian noise (WGN) $\mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \sigma_{wgn} \mathbf{I}_p)$ with an increasing variance σ_{wgn} on randomly selected observations \mathbf{y}_i .

Three data sets are considered to compare the CM estimation errors:

- 1) Corrupted data $\mathbf{y}_i + \mathbf{z}_i$;
- 2) Masked data where the outliers are set as missing data.
- 3) Data without outliers \mathbf{y}_i ;

Tyler's estimator is computed on data set 1), whereas the EM-Tyl algorithm is applied to data set 2). Finally, the clairvoyant SCM and Tyler's estimators are computed from data set 3).

Fig. 5 gives an overlook on what choice of the CM estimation procedure would be preferable. It shows that it mainly depends on the outlier variance: for a small σ_{wgn} , Tyler's estimator should be preferred, whereas EM-Tyl is more suitable when the outlier variance reaches the variance of the signal σ . Note also that EM-Tyl provides better performances than the clairvoyant SCM until the missing data ratio reaches $\sim 45\%$ of the data.

4.3. Robust imputation of missing values with outliers

In this experiment, we propose to apply our estimators to missing data imputation with possible outliers corrupting the data. Missing data imputation [6] concerns a wide range of applications, including remote sensing [8]. A recent procedure to deal with missing values was developed, namely the EM-EOF method [57], which uses the EM algorithm and empirical orthogonal functions (EOFs) to iteratively decompose the CM and reconstruct the incomplete data with a few number of selected EOFs $k \ll p$. The final number of components is the one giving the minimal error between the initial data and the imputed data. As shown in Fig. 6, instead of the SCM (which is used in the aforementioned study), the EM-Tyl-r estimator is used with $r = k$ at the last iteration of the EM-EOF algorithm. This is justified by the fact that this algorithm iteratively updates the SCM with the predicted missing values at each step, whereas the EM-Tyl-r estimator only needs the missing data pattern, which remains the same at each iteration. The RMI estimator (40) is also considered for comparison.




Figure 5: Mean geodesic distances over 200 simulated sets $\{\mathbf{y}_i\}_{i=1}^n$ as function of the missing data or outlier ratio (%) for $n = 200$. Tyl-clair and SCM-clair errors are shown with no outliers. Tyl-clair is displayed on corrupted data with different outlier variance σ_{wgn} w.r.t. the data variance σ . EM-Tyl and Tyl-obs errors are shown when outliers are masked as missing values.




Figure 6: Diagram of the EM-EOF method [57] for missing data imputation. After initializing the missing values, the CM matrix is estimated from the data matrix \mathbf{Y} , which is then reconstructed with a number of components $k \ll p$. Instead of using the SCM, the EM-Tyl-r estimator with $r = k$ is plugged in at the last iteration of the algorithm (in red).

The data is generated using a Haystack-type model [58] which draws samples $\{\mathbf{y}_i\}_{i=1}^n$ as inliers \mathbf{y}_i^{in} and outliers $\mathbf{y}_i^{\text{out}}$:

$$\{\mathbf{y}_i\}_{i=1}^n = \{\{\mathbf{y}_i^{\text{in}}\}_{i \in \{1, \dots, n_{\text{in}}\}}, \{\mathbf{y}_i^{\text{out}}\}_{i \in \{1, \dots, n_{\text{out}}\}}\} \quad (43)$$

$$\mathbf{y}_i^{\text{in}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p + \sigma_s^2 \mathbf{U} \mathbf{U}^\top) \quad (44)$$

$$\mathbf{y}_i^{\text{out}} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_p + \sigma_o^2 \mathbf{U}_\perp \mathbf{U}_\perp^\top) \quad (45)$$

where $\mathbf{U} \subset \text{St}_{p,k}$ is the signal subspace basis, σ_s^2 and σ_o^2 are respectively the signal to noise ratio (SNR) and the outlier to noise ratio (ONR), and $100 \times \frac{(n - n_{\text{in}})}{n}$ is the outlier ratio in the data set in percentage.

In the experiments, we choose $\sigma_s^2 = 10$, $\sigma_o^2 \in \{0, 1.5, 3, \dots, 30\}$ and the outlier ratio varies from 0 to 50% of the data. Note that inliers and outliers are whole vectors (not just entries) which are randomly chosen in the data set. 30% of the data is discarded as missing values under a general pattern.

To measure the performance of imputation, a set $\mathcal{Y} = \{y_l\}_{l=1}^s$ consisting of s random points over $\{y_{ij}\}_{1 \leq i \leq n, 1 \leq j \leq p}$ is selected. \mathcal{Y} is called the cross-validation (CV) data set. Then, we compute the root-mean-square error (RMSE) between \mathcal{Y} and its reconstruction $\hat{\mathcal{Y}} = \{\hat{y}_l\}_{l=1}^s$ after imputation:

$$\delta_{\text{imputation}} = \left[\frac{1}{s} \sum_{l=1}^s (\hat{y}_l - y_l)^2 \right]^{1/2} \quad (46)$$

In total, 1% of the data is chosen for the cross-validation data set. These values are artificially removed and copied before the process, and compared to the new values




Figure 7: Mean of cross-validation error (46) over 200 simulated sets $\{\mathbf{y}_i\}_{i=1}^n$ as function of the outlier ratio (%) and signal to outlier ratio σ_o^2 for the EM-EOF using EM-Tyl-r estimator, EM-EOF using the SCM and RMI procedures. 30 % of missing data under the general pattern are generated. Chosen parameters are $n = 200$ and $p = 15$.

after the imputation procedure. For model (43), CV errors illustrated by Fig. 7 show a substantial gain by replacing the SCM by the EM-Tyl-r estimator at the last step of the EM-EOF method, whereas RMI performs poorly. The gain is larger for important signal to outlier ratio σ_o^2 , which confirms the robustness of the proposed estimator for robust low-rank imputation under model (43).

5. Application to classification and clustering problems

5.1. Classification with missing values

In remote sensing, and more particularly in multispectral and hyperspectral imaging, missing data can arise for various reasons including:

- 1) Clouds when the sensor operates in the visible part of the spectrum [8];
- 2) Sensor failure, as the known problem on the scan-line corrector images of the Landsat ETM+ sensor [59] or Aqua MODIS band 6 [60];

In these cases, the data (or the entire band itself) is masked, downsampled to avoid temporal gaps or restored using specific gap-filling methods [59, 60]. When it comes to classification tasks, incomplete data is a challenge. In this scope, existing approaches

classify, mask, and interpolate values from cloudy observations in a pre-processing step [61].

To tackle this challenge, the proposed estimators are used as a set of descriptors⁹ $\{\theta_i\}$ to a classification problem on the *BreizhCrops* data set [61]. The data consists of Sentinel-2 multivariate time series of field crop reflectances on the Brittany region over 23 spectral bands. It is divided into four sub-regions called FRH01, FRH02, FRH03 and FRH04 corresponding to the four departments of Brittany. These sub-regions contain field labels which are gathered in 9 selected classes representing crop categories, *e.g.*, *barley*, *nuts*, *wheat*, etc. Following [61], each band is mean-aggregated over one field parcel to a feature vector $\mathbf{y}^t \in \mathbb{R}^p$, where p is the number of features (here the spectral bands) and t a timestamp. Thus, the whole data set is denoted $\{\{\mathbf{y}_k^t\}_{k=1}^K\}_{t=1}^T$, which corresponds to the aggregation of all reflectances at field parcels $k \in [1, K]$ and timestamps $t \in [1, T]$. The experiment is done on the L1C top-of-atmosphere product and on 13 bands selected by the *BreizhCrops* processing chain [61]. For the L1C product, the total number of parcels is $K = 608263$ and the number of observations per parcel is $T = 51$.

A supervised covariance-based classification is performed using a Minimum distance to Riemannian mean (MDRM) classifier [62], which works in a test-training form. The training step provides a set of $p \times p$ SPD matrices encoding field parcels for the available classes. Assuming that the data is complete, the set of SCMs would be computed as:

$$\Sigma_k = \frac{1}{T} \sum_{t=1}^T \mathbf{y}_k^t (\mathbf{y}_k^t)^\top, \quad k = 1, \dots, K \quad (47)$$

For each class, a center of mass of the available parcels is estimated. In the test step, a field parcel is also encoded as an SPD matrix and then assigned to the class whose center of mass is the closest according to the distance (42) acting on the manifold \mathcal{S}_{++}^p [63].

Before the experiment, missing data are artificially created in the data set. For this, successive bands are set as missing and the performances are evaluated as function of the missing data ratio. The classifier is trained first on the FRH01 region and tested on the FRH03 region. A second experiment is conducted, where the train data set is changed to the FRH04 region while the test set remains unchanged.

⁹In machine learning problems, statistical descriptors are classically used as they are often more discriminative than raw data.

Results in terms of overall accuracy (OA) for the FRH03 region with two different training sets versus the number of missing bands are displayed in Fig. 8. Three covariance estimation strategies are shown: covariance estimation for Gaussian data (EM-SCM), robust stochastic imputation (RSI) and robust covariance estimator (EM-Tyl) multiplied by the geometric mean of the estimated scales $\hat{\tau}_i$:

$$\text{EM-Tyl-}\Pi\tau_i = \text{EM-Tyl} \times \left(\prod_{i=1}^T \hat{\tau}_i \right)^{\frac{1}{T}} \quad (48)$$

Here, the geometric mean acts as an additional power information to the EM-Tyl estimator and better fits the average of textures that might vary on different scales. Unfortunately, experiments have shown that the low-rank structure does not improve the OA, which might be due to the mean-aggregation over each parcel, which essentially acts as a filter.

Results show that, for this data set, the MDRM classifier based on robust covariance matrix estimation (EM-Tyl- $\Pi\tau_i$ -based) is generally more robust to missing bands than the EM-SCM one, whereas the RSI-based classifier gives a rapidly decreasing OA when the gaps ratio increases. Furthermore, we observe that the EM-SCM-based classifier is more suited for large missing data ratio when the training phase is performed on the FRH04 region. Interestingly, the classifier based on the EM-Tyl- $\Pi\tau_i$ estimator provides an almost stable OA until the gap ratio reaches 28% of the data (4 missing bands) when FRH01 is the training set and 14% of the data (2 missing bands) when the classifier is trained on FRH04.

5.2. Image clustering with missing values

We apply the proposed CM estimators to a hyperspectral image clustering problem using the *K-means++* algorithm [64] on the *Indian Pines* data set [65], consisting of a 145×145 pixels image with $p = 200$ bands. As in the classification problem, the proposed estimators EM-SCM and EM-Tyl are used as descriptors $\{\theta_i\}$, as well as EM-SCM-r and EM-Tyl-r. The aim is to partition the descriptors in a number of clusters which correspond to the 16 classes dividing the *Indian Pines* image.

The image is first centered by subtracting the global mean. Then, a sliding window of size $w \times w$ is applied to the image. One descriptor $\{\theta_i\}$ is estimated using the $n = w^2$




Figure 8: Classification mean overall accuracy (%) versus the number of randomly missing bands (1 band $\sim 7\%$ of the data) over 50 run of the MDRM [62] classifier based on EM-SCM, EM-Tyl- $\prod \tau_i$ and RSI estimators.

observations in each window denoted $\mathbf{X}_i \in \mathbb{R}^{p \times n}$. Thus, we obtain a set of descriptors $\{\theta_i\}$ to cluster using a *K-means++* [64]. For the descriptors using the low-rank model (4), the first $r = 5$ components are kept which concur with the five principal eigenvectors of the SCM of *Indian Pines* containing more than 95% of the total cumulative variance¹⁰.

Due to the high dimensionality and a possible runtime overflow due to the creation of missing data, the image is subsampled regularly to reduce the dimension to 20 bands and cropped to get a final 85×70 image composed by 5 or the 16 original classes (see Fig. 9a). As shown by Fig. 9b, sensor failure is simulated by adding missing values on random columns of selected bands.

The averaged Overall Accuracy (OA) are reported in Fig. 10. We observe that in general, the descriptors based on RSI and RSI-r estimators give lower accuracies compared to the descriptors using the SCM, SCM-r, EM-Tyl and EM-Tyl-r estimators. As the missing data ratio increases, the performances of the EM-SCM estimator are undermined whereas EM-Tyl estimator accuracies remain stable around 50%. For large missing data ratio, EM-SCM-r performs surprisingly well compared to EM-Tyl and EM-Tyl-r. For lower missing data ratios, clustering achieves its best performances using the EM-Tyl-r estimator.

¹⁰This measure can be easily computed by the formula $\sum_{i=1}^r \lambda_i / \sum_{i=1}^p \lambda_i$.




Figure 9: Selected sub-image of *Indian Pines*: ground truth (left) and simulated sensor failure on band 10 (right) where white stripes are the missing values (20 columns of pixels are set as missing).




Figure 10: Clustering mean overall accuracy (%) versus the number of incomplete bands (as in Fig. 9b) over 10 run of the *Kmeans++* algorithm for the descriptors based on EM-SCM, EM-Tyl, RSI, EM-SCM-r, EM-Tyl-r and RSI-r estimators.

6. Conclusion

This article proposes a novel procedure based on the EM algorithm to perform robust low-rank estimation of the covariance matrix with missing data following a general pattern. The developed tools take advantage of the MSG distribution and of the LR structure of the covariance matrix. Closed-form expressions of the unknown CM and associated scales are derived and integrated to the M-step of the algorithm, which generalizes the Gaussian case. The performance of the proposed estimators are validated on simulated data sets with missing values and corrupted by outliers, as well as real-world incomplete data sets. Compared to the classical Gaussian assumption and to the unstructured model, experiments show the possibility to improve the results in terms of CM estimation and robust imputation, as well as supervised (classification) and unsupervised (image clustering) problems. Some potential extensions of this work include the extension to other classes of compound Gaussian distributions and experimenting classification tasks with temporal gaps rather than spectral gaps, as well as more spectral information by reducing the downsampling rate.

Acknowledgment

This work was supported by ANR ASTRID project MARGARITA (ANR-17-ASTR-0015). The authors would like to thank Antoine Collas and Ammar Mian for their precious help on the implementation pipeline of the MDRM and *Kmeans++* algorithms for real data, and Florent Bouchard for his deep insights in robust estimation and covariance-based classification.

References

- [1] T. W. Anderson, Maximum likelihood estimates for a multivariate normal distribution when some observations are missing, *Journal of the American Statistical Association* 52 (278) (1957) 200–203.
- [2] A. A. Afifi, R. M. Elashoff, Missing observations in multivariate statistics I. Review of the literature, *Journal of the American Statistical Association* 61 (315) (1966) 595–604.
- [3] R. J. A. Little, D. B. Rubin, *Statistical analysis with Missing Data*, Wiley, New York, 1987.
- [4] J. L. Schafer, *Analysis of incomplete multivariate data*, CRC press, 1997.
- [5] J.-P. Benzécri, et al., *L’analyse des données*, Vol. 2, Dunod Paris, 1973.
- [6] S. Van Buuren, *Flexible imputation of missing data*, CRC press, 2018.

- [7] B. Walczak, D. L. Massart, Dealing with missing data : Part I, *Chemom. Intell. Lab. Syst.* 58 (2001) 15–27.
- [8] H. Shen, X. Li, Q. Chen, C. Zeng, G. Yang, H. Li, L. Zhang, Missing information reconstruction of remote sensing data: A technical review, *IEEE Geosci. Remote Sens. Mag.* 3 (2015) 61–85. doi:10.1109/MGRS.2015.2441912.
- [9] D. B. Rubin, Inference and missing data, *Biometrika* 63 (3) (1976) 581–592.
- [10] A. P. Dempster, N. M. Laird, D. B. Rubin, Maximum likelihood from incomplete data via the EM algorithm, *J. Royal Statistical Society. Series B (Methodological)* 39 (1) (1977) 1–38.
- [11] M. Jamshidian, An EM algorithm for ML factor analysis with missing data, in: *Latent variable modeling and applications to causality*, Springer, 1997, pp. 247–258.
- [12] C. Liu, Efficient ML estimation of the multivariate normal distribution from incomplete data, *Journal of Multivariate Analysis* 69 (1999) 206–217. doi:10.1006/jmva.1998.1793.
- [13] T. Schneider, Analysis of incomplete climate data: Estimation of mean values and covariance matrices and imputation of missing values, *J. Climate* 14 (2001) 853–871.
- [14] K. Lounici, et al., High-dimensional covariance matrix estimation with missing observations, *Bernoulli* 20 (3) (2014) 1029–1058.
- [15] N. Städler, D. J. Stekhoven, P. Bühlmann, Pattern alternating maximization algorithm for missing data in high-dimensional problems., *J. Mach. Learn. Res.* 15 (1) (2014) 1903–1928.
- [16] T. Chen, E. Martin, G. Montague, Robust probabilistic PCA with missing data and contribution analysis for outlier detection, *Computational Statistics & Data Analysis* 53 (10) (2009) 3706–3716.
- [17] J. Josse, F. Husson, Handling missing values in exploratory multivariate data analysis methods, *Journal de la Société Française de Statistique* 153 (2) (2012) 79–99.
- [18] R. J. Little, Robust estimation of the mean and covariance matrix from data with missing values, *Journal of the Royal Statistical Society: Series C (Applied Statistics)* 37 (1) (1988) 23–38.
- [19] C. Liu, D. B. Rubin, ML estimation of the t distribution using EM and its extensions, ECM and ECME, *Statistica Sinica* (1995) 19–39.
- [20] J. Liu, D. P. Palomar, Regularized robust estimation of mean and covariance matrix for incomplete data, *Signal Processing* 165 (2019) 278 – 291. doi:https://doi.org/10.1016/j.sigpro.2019.07.009.
- [21] D. E. Tyler, A distribution-free M-estimator of multivariate scatter, *The Annals of Statistics* 15 (1) (1987) 234–251.
- [22] E. Conte, A. De Maio, G. Ricci, Recursive estimation of the covariance matrix of a compound-Gaussian process and its application to adaptive CFAR detection, *IEEE Transactions on signal processing* 50 (8) (2002) 1908–1915.
- [23] F. Pascal, Y. Chitour, J.-P. Ovarlez, P. Forster, P. Larzabal, Covariance structure maximum-likelihood estimates in compound Gaussian noise: Existence and algorithm analysis, *IEEE Transactions on Signal Processing* 56 (1) (2008) 34–48. doi:10.1109/TSP.2007.901652.
- [24] A. M. Zoubir, V. Koivunen, E. Ollila, M. Muma, *Robust statistics for signal processing*, Cambridge University Press, 2018.

- [25] F. Gini, M. V. Greco, M. Diani, L. Verrazzani, Performance analysis of two adaptive radar detectors against non-gaussian real sea clutter data, *IEEE Transactions on Aerospace and Electronic Systems* 36 (4) (2000) 1429–1439.
- [26] M. S. Greco, F. Gini, Statistical analysis of high-resolution SAR ground clutter data, *IEEE Transactions on Geoscience and Remote sensing* 45 (3) (2007) 566–575.
- [27] J. Theiler, B. Foy, A. Fraser, Characterizing non-gaussian clutter and detecting weak gaseous plumes in hyperspectral imagery, *Proceedings SPIE* 5806 (2005) 182–193.
- [28] P. Formont, F. Pascal, G. Vasile, J.-P. Ovarlez, L. Ferro-Famil, Statistical classification for heterogeneous polarimetric SAR images, *IEEE Journal of selected topics in Signal Processing* 5 (3) (2010) 567–576.
- [29] V. Ollier, M. N. El Korso, A. Ferrari, R. Boyer, P. Larzabal, Robust distributed calibration of radio interferometers with direction dependent distortions, *Signal Processing* 153 (2018) 348–354.
- [30] R. B. Abdallah, A. Breloy, M. N. El Korso, D. Lautru, Bayesian signal subspace estimation with compound Gaussian sources, *Signal Processing* 167 (2020) 107310.
- [31] E. Ollila, D. E. Tyler, V. Koivunen, H. V. Poor, Complex elliptically symmetric distributions: Survey, new results and applications, *IEEE Transactions on signal processing* 60 (11) (2012) 5597–5625.
- [32] G. Frahm, U. Jaekel, A generalization of Tyler’s M-estimators to the case of incomplete data, *Computational Statistics & Data Analysis* 54 (2) (2010) 374–393.
- [33] R. J. Little, D. B. Rubin, *Statistical analysis with missing data*, Vol. 793, John Wiley & Sons, 2019.
- [34] E. G. Larsson, P. Stoica, High-resolution direction finding: the missing data case, *IEEE transactions on signal processing* 49 (5) (2001) 950–958.
- [35] R. J. A. Little, D. B. Rubin, *Statistical analysis with Missing Data*, 2nd Edition, Wiley, New York, 2002.
- [36] N. Goodman, J. Stiles, On clutter rank observed by arbitrary arrays., *IEEE Transactions on Signal Processing* 55 (2007) 178–186.
- [37] A. Wiesel, Unified framework to regularized covariance estimation in scaled Gaussian models, *IEEE Transactions on Signal Processing* 60 (1) (2012) 29–38. doi:10.1109/TSP.2011.2170685.
- [38] K. Yao, A representation theorem and its applications to spherically-invariant random processes, *IEEE Transactions on Information Theory* 19 (5) (1973) 600–608.
- [39] R. A. Maronna, R. D. Martin, V. J. Yohai, M. Salibián-Barrera, *Robust statistics: theory and methods (with R)*, John Wiley & Sons, 2019.
- [40] A. Mian, G. Ginolhac, J.-P. Ovarlez, A. M. Atto, New robust statistics for change detection in time series of multivariate SAR images, *IEEE Transactions on Signal Processing* 67 (2) (2018) 520–534.
- [41] D. Ruppert, D. S. Matteson, *Statistics and data analysis for financial engineering*, Vol. 13, Springer, 2011.
- [42] I. M. Johnstone, On the distribution of the largest eigenvalue in principal components analysis, *The Annals of Statistics* 29 (2) (2001) 295 – 327.
- [43] P. Stoica, Y. Selen, Model-order selection: a review of information criterion rules, *IEEE Signal*

- Processing Magazine 21 (4) (2004) 36–47. doi:10.1109/MSP.2004.1311138.
- [44] M. E. Tipping, C. M. Bishop, Probabilistic principal component analysis, *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* 61 (3) (1999) 611–622.
 - [45] A. Aubry, A. De Maio, S. Marano, M. Rosamilia, Structured covariance matrix estimation with missing-(complex) data for radar applications via expectationmaximization, *IEEE Transactions on Signal Processing* (2021) 1–1doi:10.1109/TSP.2021.3111587.
 - [46] D. Paindaveine, A canonical definition of shape, *Statistics & probability letters* 78 (14) (2008) 2240–2247.
 - [47] X. Zhang, M. N. El Korso, M. Pesavento, MIMO radar target localization and performance evaluation under SIRP clutter, *Signal Processing* 130 (2017) 217–232.
 - [48] T. Anderson, *An Introduction to Multivariate Statistical Analysis*, Wiley, New York, 1965.
 - [49] B. Kang, V. Monga, M. Rangaswamy, Rank-constrained maximum likelihood estimation of structured covariance matrices, *IEEE Transactions on Aerospace and Electronic Systems* 50 (1) (2014) 501–515. doi:10.1109/TAES.2013.120389.
 - [50] Y. Sun, P. Babu, D. P. Palomar, Robust estimation of structured covariance matrix for heavy-tailed elliptical distributions, *IEEE Transactions on Signal Processing* 64 (14) (2016) 3576–3590. doi:10.1109/TSP.2016.2546222.
 - [51] M. Razaviyayn, M. Hong, Z.-Q. Luo, A unified convergence analysis of block successive minimization methods for nonsmooth optimization, *SIAM Journal on Optimization* 23 (2) (2013) 1126–1153.
 - [52] P. Royston, Multiple imputation of missing values, *The Stata Journal* 4 (3) (2004) 227–241.
 - [53] R. Bhatia, *Positive definite matrices*, Princeton university press, 2009.
 - [54] F. Bouchard, A. Mian, J. Zhou, S. Said, G. Ginolhac, Y. Berthoumieu, Riemannian geometry for compound Gaussian distributions: Application to recursive change detection, *Signal Processing* 176 (2020) 107716. doi:https://doi.org/10.1016/j.sigpro.2020.107716.
 - [55] J. P. Musial, M. M. Verstraete, N. Gobron, Technical note: Comparing the effectiveness of recent algorithms to fill and smooth incomplete and noisy time series, *Atmospheric Chemistry and Physics* 11 (15) (2011) 7905–7923. doi:10.5194/acp-11-7905-2011.
 - [56] A. M. Zoubir, V. Koivunen, Y. Chakhchoukh, M. Muma, Robust estimation in signal processing: A tutorial-style treatment of fundamental concepts, *IEEE Signal Processing Magazine* 29 (4) (2012) 61–80.
 - [57] A. Hippert-Ferrer, Y. Yan, P. Bolon, EM-EOF: Gap-filling in incomplete SAR displacement time series, *IEEE Trans. Geosci. Remote Sens.* 59 (7) (2021) 5794–5811.
 - [58] G. Lerman, M. McCoy, J. A. Tropp, T. Zhang, Robust computation of linear models, or how to find a needle in a haystack, *Tech. rep.*, California Inst of Tech Pasadena Dept of Computing and Mathematical Sciences (2012).
 - [59] J. Chen, X. Zhu, J. E. Vogelmann, F. Gao, S. Jin, A simple and effective method for filling gaps in Landsat ETM+ SLC-off images, *Remote sensing of environment* 115 (4) (2011) 1053–1064.
 - [60] P. Rakwatin, W. Takeuchi, Y. Yasuoka, Restoration of Aqua MODIS band 6 using histogram matching and local least squares fitting, *IEEE Transactions on Geoscience and Remote Sensing*

47 (2) (2008) 613–627.

- [61] M. Rußwurm, C. Pelletier, M. Zollner, S. Lefèvre, M. Körner, Breizhcrops: A time series dataset for crop type mapping, *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences ISPRS* (2020) (2020).
- [62] A. Barachant, S. Bonnet, M. Congedo, C. Jutten, Multiclass brain–computer interface classification by Riemannian geometry, *IEEE Transactions on Biomedical Engineering* 59 (4) (2012) 920–928. doi:10.1109/TBME.2011.2172210.
- [63] M. Congedo, P. L. C. Rodrigues, C. Jutten, The Riemannian minimum distance to means field classifier, in: *8th Graz Brain-Computer Interface Conference 2019*, 2019.
- [64] S. Vassilvitskii, D. Arthur, k-means++: The advantages of careful seeding, in: *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, 2006, pp. 1027–1035.
- [65] M. F. Baumgardner, L. L. Biehl, D. A. Landgrebe, 220 band aviris hyperspectral image data set: June 12, 1992 indian pine test site 3, *Purdue University Research Repository* 10 (2015).