



HAL
open science

Model-based clustering of Gaussian copulas for mixed data

Matthieu Marbac, Christophe Biernacki, Vincent Vandewalle

► **To cite this version:**

Matthieu Marbac, Christophe Biernacki, Vincent Vandewalle. Model-based clustering of Gaussian copulas for mixed data. 2014. hal-00987760v3

HAL Id: hal-00987760

<https://hal.science/hal-00987760v3>

Preprint submitted on 29 Sep 2015 (v3), last revised 20 Dec 2016 (v4)

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.

Model-based clustering of Gaussian copulas for mixed data

Matthieu Marbac and Christophe Biernacki and Vincent Vandewalle

September 29, 2015

Abstract

Clustering task of mixed data is a challenging problem. In a probabilistic framework, the main difficulty is due to a shortage of conventional distributions for such data. In this paper, we propose to achieve the mixed data clustering with a Gaussian copula mixture model, since copulas, and in particular the Gaussian ones, are powerful tools for easily modelling the distribution of multivariate variables. Indeed, considering a mixing of continuous, integer and ordinal variables (thus all having a cumulative distribution function), this copula mixture model defines intra-component dependencies similar to a Gaussian mixture, so with classical correlation meaning. Simultaneously, it preserves standard margins associated to continuous, integer and ordered features, namely the Gaussian, the Poisson and the ordered multinomial distributions. As an interesting by-product, the proposed mixture model generalizes many well-known ones and also provides tools of visualization based on the parameters. At a practical level, the Bayesian inference is retained and it is achieved with a Metropolis-within-Gibbs sampler. Experiments on simulated and real data sets finally illustrate the expected advantages of the proposed model for mixed data: flexible and meaningful parametrization combined with visualization features.

keywords Clustering; Gaussian copula; Metropolis-within-Gibbs algorithm; Mixed data; Mixture models; Visualization.

1 Introduction

Clustering is an efficient tool for managing large data sets, as it groups individuals into a few specific classes. In a probabilistic framework, a class groups together individuals arising from the same distribution. In such a framework, the most popular approaches model the data distribution using finite mixture models of parametric distributions (?). The literature consecrated to homogeneous data (composed of variables of the same type) is extensive and presents Gaussian mixture models (?), multinomial mixture models (?) and Poisson mixture models (?) as the standard models used to cluster such data sets. The

success of these models is due to the use of conventional distributions for mixture components, thus allowing practitioners to easily interpret the components. However, even though many data sets contain mixed data (variables of different types), only a few multivariate distributions are available for such data sets. Moreover, these distributions can be difficult to interpret. We shall now present the three main models used to cluster mixed data. A more detailed overview is given by ?.

The *locally independent mixture model* (?) analyzes data sets by assuming that all variables are independent given the component. It can provide meaningful results (see the applications of ??), particularly when the one-dimensional marginal distributions of the components follow conventional distributions. However, this model can lead to severe bias when its main assumption is violated (see the application of ?). In such cases, two methods can be considered. The first method consists of selecting a subset of intra-class independent variables (?), but some information may be lost. The second method consists of using a model to relax the conditional independence assumption. We shall now present both methods.

The *location mixture model* (??) assumes that continuous variables follow a multivariate Gaussian distribution conditionally on both component and categorical variables. This model takes intra-component dependencies into account but requires too many parameters. It was therefore expanded by ? by grouping the variables into conditionally independent blocks, where a block comprises no more than one categorical variable and follows a location model. However, the interpretation of components can be complex. Indeed, for a given component, the one-dimensional marginal distribution of the continuous variables are itself Gaussian mixture models (not conventional distributions!). Moreover, the allocation of variables into blocks is a difficult problem that the authors achieve with an ascending method that is sub-optimal.

The *underlying variables mixture model* (?) analyzes data sets with continuous and ordinal variables. It assumes that each discrete variable arises from a latent continuous variable and that all continuous variables (observed and unobserved) follow a Gaussian mixture model. The distribution of observed variables is obtained by integrating each Gaussian component into the subset of latent variables. However, in practice, this computation is not feasible when there are more than two discrete variables. In an effort to study data sets with numerous binary variables, ? has expanded this model by estimating the scores of latent variables from those of binary variables. However, the interpretation of components is carried out for score-related parameters (not those related to observed variables).

Previous models illustrate the difficulty of managing mixed variables with a model whose interpretation and inference are easy. Moreover, they do not the widespread case of integer variables. The main difficulty is due to a shortage of conventional distributions for mixed data. However, copulas are standard

tools to define multivariate distributions in a systematic way, thus having good potentiality for providing a sensible answer.

Copulas (??) can be used to build a multivariate model by defining, on the one hand, the *one-dimensional marginal distributions*, and, on the other, the *dependency model*. Recently, ? and ? have modelled the distribution of mixed variables using one Gaussian copula. As pointed-out by ?, the maximum likelihood inference is very difficult for a Gaussian copula with discrete margins. So, it is often replaced by the *Inference Function for Margins* method performing the inference in two steps (see Chapter 10 of ?) but which is sub-optimal. When all the variables are continuous, the fixed-point-based algorithm proposed by ? achieves the maximum likelihood estimation, but this approach is not doable for mixed data. Therefore, as shown by ?, it is more convenient to work in a Bayesian framework since it simplifies the inference by using the latent structure of the model.

In this paper, we propose to achieve the mixed data clustering with a *Gaussian copula mixture model*. This new model assumes that each component follow a Gaussian copula (??). Thus, the *intra-component dependencies* are taken into account and are analyzed from the intra-component correlation matrices. Taking advantage of the properties of the copulas, the Gaussian copula mixture model provides *conventional distributions* for all one-dimensional marginal distributions of each component. It results a *three-level framework* which permits a user-friendly interpretation, *i.e.* the proportions indicate the component weights, the one-dimensional marginal parameters of each component roughly describe the classes, and the correlation matrices refine this description. Finally, using the continuous latent structure of the Gaussian copulas, a visualisation tool based on a Principal Component Analysis (PCA) per component is proposed. This visualization provides a summary of main intra-component dependencies and a scatterplot of individuals according to component parameters.

This paper is organized as follows. Section 2 introduces the Gaussian copula mixture model and its links to well-known models. Section 3 presents the Metropolis-within-Gibbs algorithm used to perform Bayesian inference. Section 4 illustrates algorithm behavior and model robustness through numerical experiments. Section 5 presents two applications of the new model by clustering two real data sets. Section 6 concludes the study.

2 Mixture model of Gaussian copulas

2.1 Finite mixture model

The vector $\mathbf{x} = (x^1, \dots, x^e) \in \mathbb{R}^c \times \mathcal{X}$ denotes the $e = c + d$ observed variables. Its first c elements are denoted by \mathbf{x}^c and correspond to the subset of the continuous variables defined on the space \mathbb{R}^c . Its last d elements are denoted by \mathbf{x}^d and correspond the subset of the discrete variables (integer, ordinal or

binary) defined on the space \mathcal{X} . Note that if x^j is an ordinal variable with m_j modalities, then it uses a numeric coding $\{1, \dots, m_j\}$.

Data \mathbf{x} are assumed to arise from the mixture model of g parametric distributions whose the probability distribution function (pdf) is written as follows

$$p(\mathbf{x}|\boldsymbol{\theta}) = \sum_{k=1}^g \pi_k p(\mathbf{x}|\boldsymbol{\alpha}_k), \quad (1)$$

where $\boldsymbol{\theta} = (\boldsymbol{\pi}, \boldsymbol{\alpha})$ denotes the whole parameters. Vector $\boldsymbol{\pi} = (\pi_1, \dots, \pi_g)$ is defined on the simplex of size g and groups together the component proportions, where π_k is the proportion of component k . Vector $\boldsymbol{\alpha} = (\boldsymbol{\alpha}_1, \dots, \boldsymbol{\alpha}_g)$ groups together the component parameters, where $\boldsymbol{\alpha}_k$ denotes the parameters of component k .

2.2 Component modelled by a Gaussian copula

The Gaussian copula mixture model assumes that each component follows a Gaussian copula. Thus, component k is parametrized by the correlation matrix $\boldsymbol{\Gamma}_k$ of size $e \times e$ and by the parameters of the one-dimensional margin distributions $\boldsymbol{\beta}_k = (\beta_{k1}, \dots, \beta_{ke})$ where β_{kj} denotes the parameters of the one-dimensional marginal distribution j . By grouping all the parameters of component k in $\boldsymbol{\alpha}_k = (\boldsymbol{\Gamma}_k, \boldsymbol{\beta}_k)$, the cumulative distribution function (cdf) of component k is written as

$$P(\mathbf{x}|\boldsymbol{\alpha}_k) = \Phi_e(\Phi_1^{-1}(u_k^1), \dots, \Phi_1^{-1}(u_k^e)|\mathbf{0}, \boldsymbol{\Gamma}_k), \quad (2)$$

where $u_k^j = P(x^j|\beta_{kj})$ is the value of the cdf of the one-dimensional marginal distribution of variable j for component k evaluated at x^j , where $\Phi_e(\cdot|\mathbf{0}, \boldsymbol{\Gamma}_k)$ is the cdf of the e -variate centred Gaussian distribution with correlation matrix $\boldsymbol{\Gamma}_k$ and where $\Phi_1^{-1}(\cdot)$ is the inverse cumulative distribution function of the standard univariate Gaussian $\mathcal{N}_1(0, 1)$.

In cluster analysis, the categorical variable $z \in \{1, \dots, g\}$ which indicates the individual's component membership is unobserved. Moreover, it follows the multinomial distribution $\mathcal{M}_g(\pi_1, \dots, \pi_g)$. Therefore, mixture models are often interpreted as the marginal distribution of \mathbf{x} based on the distribution of the variable pair (\mathbf{x}, z) . The Gaussian copula mixture model involves a second latent variable $\mathbf{y} = (y^1, \dots, y^e) \in \mathbb{R}^e$ such as $\mathbf{y}|z = k$ follows an e -variate centred Gaussian distribution $\mathcal{N}_e(\mathbf{0}, \boldsymbol{\Gamma}_k)$. Thus, this model can be interpreted as the marginal distribution of \mathbf{x} based on the distribution of the variable triplet $(\mathbf{x}, \mathbf{y}, z)$. Conditionally on $(\mathbf{y}, z = k)$, \mathbf{x} is defined by

$$x^j = P^{-1}(\Phi_1(y^j)|\beta_{kj}), \quad \forall j = 1, \dots, e. \quad (3)$$

Thus, the generative model of the Gaussian copula mixture model is written as

- Class membership *sampling*: $z \sim \mathcal{M}_g(\pi_1, \dots, \pi_g)$,
- Gaussian copula *sampling*: $\mathbf{y}|z = k \sim \mathcal{N}_e(\mathbf{0}, \boldsymbol{\Gamma}_k)$,
- Observed data *deterministic computation*: \mathbf{x} is obtained from (3).

2.3 Specific distributions for mixed-type variables

We remind that the first c variables of \mathbf{x} are continuous and denoted by \mathbf{x}^c while the last d variables are discrete variables \mathbf{x}^d (integer and ordinal). So, the pdf of component k can be decomposed as

$$p(\mathbf{x}|\boldsymbol{\alpha}_k) = p(\mathbf{x}^c|\boldsymbol{\alpha}_k) \times p(\mathbf{x}^d|\mathbf{x}^c, \boldsymbol{\alpha}_k). \quad (4)$$

The expression of the pdf of component k can be deduced from the cdf of component k defined by (2) only when the one-dimensional marginal distributions of component k (*i.e.* distribution of $x^j|z = k$) are set. For facilitating the model interpretation, we consider that $x^j|z = k$ follows a conventional distribution. More precisely, $x^j|z = k$ follows a *Gaussian* (respectively *Poisson* and *ordered multinomial*) distribution if variable x^j is continuous (respectively integer and ordinal). We have $\beta_{kj} = (\mu_{kj}, \sigma_{kj})$ when $x^j|z = k$ follows the Gaussian distribution of mean μ_{kj} and variance σ_{kj}^2 , while $\beta_{kj} \in \mathbb{R}^{+*}$ when x^j is integer and β_{kj} is defined on the simplex of size m_j when x^j is ordinal. Even if the marginal distribution of an ordinal variable does not take into account of the order between modalities, this order is crucial since it allows to define the marginal cdf of the variable (it corresponds to what we previously called the “ordered multinomial”). Obviously, the Gaussian copula mixture model cannot manage data set with categorical variables due to the lack of order between modalities.

Conditionally on $z = k$ and β_{kj} , the knowledge of variable x^j fully determines the variable y^j when x^j is continuous by $y^j = \frac{x^j - \mu_{kj}}{\sigma_{kj}}$. When x^j is discrete, its knowledge only determines an interval $\mathcal{S}_k^j(x^j) =]b_k^{\ominus}(x^j), b_k^{\oplus}(x^j)]$ where y^j provides the same observed variable x^j by (3) with $b_k^{\ominus}(x^j) = \Phi_1^{-1}(P(x^j - 1|\boldsymbol{\beta}_{kj}))$ and $b_k^{\oplus}(x^j) = \Phi_1^{-1}(P(x^j|\boldsymbol{\beta}_{kj}))$. Thus, the pdf of component k is written as follows (in the same order as Equation (4))

$$p(\mathbf{x}|\boldsymbol{\alpha}_k) = \frac{\phi_c(\mathbf{y}^c|\mathbf{0}, \boldsymbol{\Gamma}_{kCC})}{\prod_{j=1}^c \sigma_{kj}} \times \int_{\mathcal{S}_k(\mathbf{x}^d)} \phi_d(\mathbf{u}|\boldsymbol{\mu}_k^d, \boldsymbol{\Sigma}_k^d) d\mathbf{u}, \quad (5)$$

where:

- the correlation matrix $\boldsymbol{\Gamma}_k = \begin{bmatrix} \boldsymbol{\Gamma}_{kCC} & \boldsymbol{\Gamma}_{kCD} \\ \boldsymbol{\Gamma}_{kDC} & \boldsymbol{\Gamma}_{kDD} \end{bmatrix}$ is decomposed into sub-matrices, for instance $\boldsymbol{\Gamma}_{kCC}$ is the sub-matrix of $\boldsymbol{\Gamma}_k$ composed by the rows and the columns related to the observed continuous variables,
- $\mathbf{y}^c = (\frac{x^j - \mu_{kj}}{\sigma_{kj}}; j = 1, \dots, c)$, $\phi_c(\mathbf{y}^c|\mathbf{0}, \boldsymbol{\Gamma}_{kCC})$ denotes the pdf of c -variate centred Gaussian distribution with correlation matrix $\boldsymbol{\Gamma}_{kCC}$,
- $\mathcal{S}_k(\mathbf{x}^d) = \mathcal{S}_k^{c+1}(x^{c+1}) \times \dots \times \mathcal{S}_k^e(x^e)$, $\phi_d(\cdot|\boldsymbol{\mu}_k^d, \boldsymbol{\Sigma}_k^d)$ denotes the pdf of a d -variate Gaussian distribution with mean $\boldsymbol{\mu}_k^d$ and covariance matrix $\boldsymbol{\Sigma}_k^d$, $\boldsymbol{\mu}_k^d = \boldsymbol{\Gamma}_{kDC} \boldsymbol{\Gamma}_{kCC}^{-1} \boldsymbol{\Psi}(\mathbf{x}^c; \boldsymbol{\alpha}_k)$ is the conditional mean of \mathbf{y}^d and $\boldsymbol{\Sigma}_k^d = \boldsymbol{\Gamma}_{kDD} - \boldsymbol{\Gamma}_{kDC} \boldsymbol{\Gamma}_{kCC}^{-1} \boldsymbol{\Gamma}_{kCD}$ is its conditional covariance matrix.

Remark 2.1 (Model identifiability). The Gaussian copula mixture model is identifiable if, at least, one variable is continuous or integer (see Appendix A).

2.4 Strengths of the Gaussian copula mixture model

2.4.1 Related models

The Gaussian copula mixture model can be used to generalize many conventional mixture models, including the four cases mentioned below.

- If the correlation matrices are diagonal (*i.e.* $\mathbf{\Gamma}_k = \mathbf{I}, \forall k = 1, \dots, g$), then the model is equivalent to the locally independent mixture model.
- If all the variables are continuous (*i.e.* $c = e$ and $d = 0$), then the model is equivalent to the Gaussian mixture model without constraints among parameters (?). In the spirit of the Gaussian homoscedastic model, we also propose a parsimonious version of the Gaussian copula mixture model by assuming the equality between the correlation matrices. This model is named *homoscedastic* since the covariance matrices of the latent Gaussian variables are equal between components (*i.e.* $\mathbf{\Gamma}_1 = \dots = \mathbf{\Gamma}_g$). The free correlation model will be now called *heteroscedastic* model).
- The model is linked to the binned Gaussian mixture model. For example, when variables are ordinal, it is equivalent to the mixture model presented by ?. In such cases, the model is stable through fusion of modalities.
- If the variables are both continuous and ordinal, then the model is a new parametrization of the model proposed by ?. It should be noted that Everitt directly estimates the space $\mathcal{S}_k(\mathbf{x}^D)$ containing the antecedents of \mathbf{x}^D . Moreover, he uses a simplex algorithm to perform maximum likelihood inference, but this method dramatically limits the number of ordinal variables. The new parametrization of the proposed mixture allows to directly estimate the one-dimensional marginal parameters β_{kj} of each component (see details in Section 3), while the parametrization of Everitt implies the difficult estimation of the space $\mathcal{S}_k(\mathbf{x}^D)$. Thus, the parameter inference is easier.

2.4.2 Standardized coefficient of correlation per class

The Gaussian copula provides a user-friendly correlation coefficient for each pair of variables. Indeed, when both variables are continuous, it is equal to the upper boundary of the correlation coefficients obtained by monotonic transformation of the variables (?). Furthermore, when both variables are discrete, it is equal to the polychoric correlation coefficient (?).

2.4.3 Data visualization per component: a by-product of Gaussian copulas

By using the latent vectors of the Gaussian copulas $\mathbf{y}|z$, a PCA-type method allows *visualization* of the individuals *per component* which permits an identification of main intra-component dependencies. The visualization of component k is performed by computing the coordinates $\mathbb{E}[\mathbf{y}|\mathbf{x}, z = k; \boldsymbol{\alpha}_k]$ and then projecting

them onto the PCA region associated with the Gaussian copula of component k . This space is obtained directly through spectral decomposition of $\mathbf{\Gamma}_k$. The individuals arising from component k follow a centred Gaussian distribution on this factorial map. Those arising from another component have an expectation not equal to zero. Therefore, an individual located far away from the origin arises from a distribution significantly different from the distribution of component k . Finally, the correlation circle summarizes intra-component correlations and avoids direct interpretation of the correlation matrix, which can be fastidious if e is large. The following example illustrates these properties.

Example 2.2. Let one continuous, one integer and one binary variables arising, in this order, from the bi-component Gaussian copula mixture model parametrized by

$$\boldsymbol{\pi} = (0.5, 0.5), \boldsymbol{\beta}_{11} = (-2, 1), \boldsymbol{\beta}_{12} = 5, \boldsymbol{\beta}_{13} = \boldsymbol{\beta}_{23} = (0.5, 0.5), \boldsymbol{\beta}_{21} = (2, 1),$$

$$\boldsymbol{\beta}_{22} = 15, \mathbf{\Gamma}_1 = \begin{pmatrix} 1 & -0.4 & 0.4 \\ -0.4 & 1 & 0.4 \\ 0.4 & 0.4 & 1 \end{pmatrix} \text{ and } \mathbf{\Gamma}_2 = \begin{pmatrix} 1 & 0.8 & 0.1 \\ 0.8 & 1 & 0.1 \\ 0.1 & 0.1 & 1 \end{pmatrix}.$$

Figure 1 provides an example of the data visualization. Figure 1(a) shows the scatterplot of the individuals in their native space. Figure 1(b) presents the scatterplot of the individuals in the first PCA-map of the second component (red). It permits to easily distinguish two classes: a centred one (red) and a second one (black) located on the left side. More precisely, the first axis (explained by the continuous and the integer variables) is strongly discriminative while the second axis (exclusively explained by the binary variable) is not discriminative. Figure 1(c) shows the correlation circle of the first PCA-map of the red component. It allows to point-out a strong correlation, for the red component, between the continuous and the integer variables.

3 Bayesian inference

3.1 Sampling layout on data and parameters

We observe the sample $\mathbf{x} = (\mathbf{x}_1, \dots, \mathbf{x}_n)$ composed of n independent individuals $\mathbf{x}_i \in \mathbb{R}^c \times \mathcal{X}$ assumed to arise from a Gaussian copula mixture model. As pointed-out by ?, the Bayesian framework considerably simplifies the inference since it uses the latent structure of the model (\mathbf{y}, z) .

Without prior information about the data, we assume independence between the prior distributions. The proportions and the parameters of the one-dimensional marginal distributions of each component $\boldsymbol{\beta}_{kj}$ follow the classical conjugate prior distributions (?). Finally, the conjugate prior of the covariance matrices is derived from an Inverse Wishart distribution as proposed by (?). Details on the prior distributions are given in Appendix B.

3.2 Gibbs and Metropolis-within-Gibbs samplers

The Bayesian estimation is managed by a Gibbs sampler (described in Algorithm 3.1) which is the most popular approach to infer mixture models since it uses the latent structure of the data. Its stationary distribution is $p(\boldsymbol{\theta}, \mathbf{y}, \mathbf{z}|\mathbf{x})$ where $\mathbf{z} = (z_1, \dots, z_n)$ denotes the class memberships of \mathbf{x} and where $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$ denotes the Gaussian vector related to \mathbf{x} . Thus, the sequence of the generated parameters is sampled from the marginal posterior distribution $p(\boldsymbol{\theta}|\mathbf{x})$.

Note that the Gaussian variable \mathbf{y} is twice sampled during one iteration of the algorithm to manage the strong dependencies between \mathbf{y} and \mathbf{z} , and between $\mathbf{y}_{[rk]}^j = \{y_i^j : z_i^{(r)} = k\}$ and $\boldsymbol{\beta}_{kj}$. Obviously, the stationary distribution stays unchanged.

Algorithm 3.1 (The Gibbs sampler). *Starting from an initial value $\boldsymbol{\theta}^{(0)}$, its iteration (r) consists in the following four steps ($k \in \{1, \dots, g\}, j \in \{1, \dots, e\}$)*

$$\mathbf{z}^{(r)}, \mathbf{y}^{(r-1/2)} \sim \mathbf{z}, \mathbf{y}|\mathbf{x}, \boldsymbol{\theta}^{(r-1)} \quad (6)$$

$$\boldsymbol{\beta}_{kj}^{(r)}, \mathbf{y}_{[rk]}^{j(r)} \sim \boldsymbol{\beta}_{kj}, \mathbf{y}_{[rk]}^j|\mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \boldsymbol{\beta}_{k\bar{j}}^{(r)}, \boldsymbol{\Gamma}_k^{(r-1)} \quad (7)$$

$$\boldsymbol{\pi}^{(r)} \sim \boldsymbol{\pi}|\mathbf{z}^{(r)} \quad (8)$$

$$\boldsymbol{\Gamma}_k^{(r)} \sim \boldsymbol{\Gamma}_k|\mathbf{y}^{(r)}, \mathbf{z}^{(r)} \quad (9)$$

where $\mathbf{y}_{[rk]} = \mathbf{y}_{\{i:z_i^{(r)}=k\}}$, $\mathbf{y}_i^{\bar{j}(r)} = (y_i^{1(r)}, \dots, y_i^{j-1(r)}, y_i^{j+1(r-1/2)}, \dots, y_i^{e(r-1/2)})$ and $\boldsymbol{\beta}_{k\bar{j}}^{(r)} = (\boldsymbol{\beta}_{k1}^{(r)}, \dots, \boldsymbol{\beta}_{k(j-1)}^{(r)}, \boldsymbol{\beta}_{k(j+1)}^{(r-1/2)}, \dots, \boldsymbol{\beta}_{ke}^{(r-1/2)})$.

The sampling according to (8) and (9) are classical. However, both samplings from (6) and (7) are difficulty performed. Therefore, they are replaced by one iteration of a Metropolis-Hastings algorithm that does not change the stationary distribution. The resulting algorithm is a Metropolis-within-Gibbs sampler whose the properties are detailed by ?. Details about the four steps of the Metropolis-within-Gibbs sampler are given in Appendix C.

By taking the mean of the parameters generated by the Metropolis-within-Gibbs sampler, we obtain a consistent estimate of $\boldsymbol{\theta}$.

3.3 Label switching problem

The label switching problem is generally solved by specific procedures (?). However, based on the argument of ?, these techniques are principally impacting when g is known.

However, when the model is used to cluster, the number of classes is unknown, and the model selection is performed by the BIC criterion (?) which simultaneously avoids the label switching phenomenon. Indeed, on the one hand, this criterion selects quite separated classes when the sample size is small, so the label switching is not present (with high probability) in practice because of the class separability. On the other hand, even if it can select more classes when

the sample size increases, the label switching problem does not occur since this phenomenon vanishes asymptotically.

Obviously, when the number of classes is fixed and the size of sample is small, the label switching problem can occur. In such a case, our advice is naturally to use the procedures of ?.

4 Simulations

In this section, two simulations are used to illustrate the new model. The first simulation shows the relevance of the estimation procedure. The second simulation illustrates the robustness of the proposed model by analyzing data sampled from a mixture of Poisson distributions.

Experiment conditions For each situation, 100 samples are generated. Parameters are estimated by taking the mean of the parameters sampled by 10^3 iteration of Algorithm 3.1 after a burn-in period of 10^2 iterations. Algorithm 3.1 is initialized with the maximum likelihood estimator of the locally independent model (particularly relevant when intra-class dependencies are small). The Kullback-Leibler divergence (?) is used to compare the estimated distribution and the distribution used to sample the data. This divergence is approximated via 10^4 iterations of a Monte Carlo simulation.

4.1 Estimation efficiency

Data sets are composed of one continuous variable, one integer variable and one binary variable. They are sampled from the distribution described in the example of Section 2.4. Results are presented by Figure 2. According to Figure 2(a), the estimated distribution converges to the true distribution when the sample size increases. Indeed, the Kullback-Leibler divergence of the estimated model from the true model decreases as a function of sample size and converges to zero. Moreover, as shown by Figure 2(b), the misclassification error rate converges to the theoretical misclassification error rate (equal to 0.005) when n increases. This simulation illustrates the convergence of the estimator computed by averaging the parameters sampled using the algorithm. Finally, the estimation procedure is not too much time consuming since samples of size $n = 100$ and $n = 1600$ require respectively 15 and 64 seconds, on an Intel Core i5-3320M processor.

4.2 Robustness

Samples are generated from the bivariate Poisson mixture model (?) with $\boldsymbol{\pi} = (1/3, 2/3)$, whose one-dimensional marginal parameters $\boldsymbol{\alpha}_k = (\lambda_{k1}, \lambda_{k2}, \lambda_{k3})$ take on the following values: $\lambda_{1h} = h$ and $\lambda_{2h} = 3 + h$, for $h = 1, 2, 3$ (see notation detailed in ?). Figure 3 presents the results and shows the robustness

of the Gaussian copula mixture model since it efficiently manages such data sets, as detailed now.

As shown by Figure 3(a), the resulting misclassification error rate converges to the theoretical misclassification error rate (equal to 0.0967). Moreover, Figure 3(b) shows that the Kullback-Leibler divergence almost vanishes when the sample size increases, thus demonstrating the flexibility of the Gaussian copula mixture model. Moreover, the resulting parameters reflect the main properties of the true distribution. Indeed, Figure 3(c) shows that the correlation coefficient between both variables, for component 1, converges to its theoretical value (equal to $\lambda_{11} + \lambda_{13} = 4$). In the same way, Figure 3(d) shows that the one-dimensional margin parameter of variable 1, for component 1, converges to its theoretical value (equal to $\lambda_{13}/\sqrt{(\lambda_{11} + \lambda_{13})(\lambda_{12} + \lambda_{13})} = 3/\sqrt{20} \simeq 0.67$).

Finally, the estimation procedure is not too much time consuming since samples of size $n = 100$ and $n = 1600$ require respectively 12 and 54 seconds, on an Intel Core i5-3320M processor.

5 Applications

In this section, we analyze two real data sets with the proposed Gaussian copula mixture model. For each number of components, 10 runs of Gibbs sampler are performed under the conditions described in the previous section.

Model selection is performed by using two information criteria (BIC criterion (?), ICL criterion (?)). These asymptotic criteria are computed with the estimate provided by the Gibbs sampler and not with the maximum likelihood estimate. However, this procedure remains valid since the BIC criterion can be computed with the estimate of the maximum *a posteriori* (?). Moreover, the ICL criterion can be computed by penalizing the BIC criterion with a term of entropy.

These criteria require the computation of the number of parameters. The locally independent model involves $\nu_{Loc} = (g - 1) + g \sum_{j=1}^e \nu_j$ parameters where ν_j is the number of parameters involved by the one-dimensional marginal distribution of one component (*i.e.* $\nu_j = 2$ if x^j is continuous, $\nu_j = 1$ if x^j is integer and $\nu_j = m_j - 1$ if x^j is ordinal with m_j modalities). The heteroscedastic Gaussian copula mixture model involves $\nu_{He} = \nu_{Loc} + g \frac{e(e-1)}{2}$ parameters and the homoscedastic Gaussian copula mixture model requires $\nu_{Ho} = \nu_{Loc} + \frac{e(e-1)}{2}$ parameters.

5.1 South African Hearth data set

The data

Data are available at <http://sci2s.ugr.es/keel/dataset.php?cod=184>. This data set is a retrospective sample of males in a heart-disease high-risk region of the Western Cape, South Africa. Many of the coronary heart disease (CHD) positive men have undergone blood pressure reduction treatment and other

programs to reduce their risk factors after their CHD event. In some cases the measurements were made after these treatments. The class label indicates if the person has a coronary heart disease (negative or positive) and is hidden for our analysis. Individuals are described by the following nine variables. The continuous variables are systolic blood pressure (sbp), cumulative tobacco (tobacco), low density lipoprotein cholesterol (ldl), adiposity, obesity and current alcohol consumption (alcohol). The integer variables are type-A behavior (typea) and age at onset (age). Finally, the binary variable indicates the presence or not of heart disease in the family history (famhist).

Model selection

Three mixture models (locally independent, heteroscedastic and homoscedastic mixture of Gaussian copulas) are fitted for various numbers of components. Table 1 presents the values of information criteria used to select the homoscedastic tri-component Gaussian copula mixture model. Indeed, this model obtains the best results for fitting the data distribution (BIC) and for fitting the best partition (ICL). Note that the locally independent model with five components obtains a value of the BIC criterion close to the tri-component homoscedastic model, but the ICL criterion favors this latter. Moreover, this model detects less components than the locally independent model, thus its interpretation is easier.

	g	1	2	3	4	5
BIC	loc. indpt.	-14127.26	-13131.88	-12813.92	-12829.68	-12738.66
	homo.	-14724.98	-13016.09	-12739.94	-12774.15	-12927.45
	hetero.	-14724.98	-13076.93	-12971.72	-13071.92	-13253.06
ICL	loc. indpt.	-14127.26	-13144.21	-12832.12	-12887.19	-12805.68
	homo.	-14724.98	-13028.07	-12762.79	-12816.44	-12979.06
	hetero.	-14724.98	-13085.52	-12989.06	-13103.61	-13299.16

Table 1: Values of the BIC and ICL criteria obtained on the South African Hearth data set (best values are in bold).

Partition study

The three models overestimate the number of components (true number is two). However, the Gaussian copula mixture model finds number of components more relevant than the locally independent model. If we assume that the number of component is known, the misclassification error rates is slightly better for both copulas models (homo. and hetero.) in comparison to the locally independence model (see Table 2). Moreover, the equality constraint of the covariance matrix increases the value of the BIC criterion obtained by the homoscedastic model, but it also impacts its resulting partition according to the partition resulting from the heteroscedastic model (see Table 3).

CHD	loc. indpt.		homo.		hetero.	
	Class 1	Class 2	Class 1	Class 2	Class 1	Class 2
negative	187	149	191	150	195	148
positive	115	11	111	10	107	12

Table 2: Confusion matrices between the CHD status and the partitions provided by the three bi-component mixture models.

		hetero.		
		Class 1	Class 2	Class 3
homo.	Class 1	46	11	0
	Class 2	0	92	14
	Class 3	0	3	296

Table 3: Confusion matrices between the tri-components homoscedastic model (row) and the tri-component homoscedastic model (column).

Figure 4 shows the PCA visualization for the component 3. The tri-component homoscedastic Gaussian copula mixture model provides three well-separated classes as shown by the factorial representation of Figure 4(a). We can see that the class 2 (red triangles) is an "intermediate" class, while class 1 (red dots) is strongly different to class 3 (black dots). We now detail the model interpretation.

Interpretation of best-fit model

The PCA visualization (Figure 4) shows that individuals of class 3 have a riskier behaviour than the others (large consumption of tobacco, alcohol, older population, large level of obesity).

A three-level interpretation (proportions, one-dimensional marginal distributions and intra-class dependencies) is feasible using the parameters. The main characteristic variables are summarized in Figure 5:

- Class 1 (*weak-risk behaviours*): this is the minority class ($\pi_1 = 0.07$). It groups the young individuals having a short consumption alcohol and tobacco. This class groups 57 individuals where only one has a coronary heart disease.
- Class 2 (*moderate-risk behaviours*): this class have a moderated proportion ($\pi_2 = 0.24$). This class is composed with individuals having a strong alcohol consumption. The other variables take intermediate values. Among the 106 individuals belonging to this class, 20 have a coronary heart disease.
- Class 3 (*high-risk behaviours*): this is the majority class ($\pi_3 = 0.69$). It

groups the individuals having the more dangerous behaviour. Among the 299 individuals belonging to this class, 139 have a coronary heart disease.

Finally, for all the components, age and the consumptions of alcohol and tobacco are strongly linked (see Figure 4(b)). Moreover, adiposity and obesity are also strongly linked.

Conclusion

For such data, the Gaussian copula mixture model reduces the drawbacks of the locally independent model. By decreasing the number of components, it yields a more interpretable model that better fits the data (BIC criterion) and provides a pertinent partition (ICL criterion). Finally, the estimation of main intra-class dependencies, based on PCA outputs per component, is an efficient tool for refining the interpretation.

5.2 Forest fire data set

The data

Data are composed of 517 forest fires that have occurred in the northeast region of Portugal (?). These forest fires are described by the following meteorological variables: seven continuous variables (four fire weather index (FWI) system variables, *i.e.* fine fuel moisture code (FFMC), duff moisture code (DMC), drought code (DC), initial season index (ISI), and three meteorological variables, *i.e.* temperature (Temp), relative humidity (RH) and wind) and three binary variables indicating the presence of rain, the season (summer or other) and the day of the week (weekend or other).

Model selection

Table 4 presents the values of information criteria used to distinctly select the heteroscedastic tri-component Gaussian copula mixture model. Note that the locally independent model degenerates with five components.

	g	1	2	3	4	5
BIC	loc. indpt.	-15152.95	-14164.51	-13990.27	-14068.92	NA
	homo.	-14401.80	-13751.82	-13927.05	-13986.90	-13755.69
	hetero	-14401.80	-13781.86	-13680.67	-13846.63	-13745.84
ICL	loc. indpt.	-15152.95	-14170.97	-14022.49	-14131.22	NA
	homo.	-14401.80	-13756.76	-13956.68	-14070.49	-13774.41
	hetero	-14401.80	-13785.33	-13682.68	-13885.11	-13776.81

Table 4: Values of the BIC and ICL criteria obtained on the forest fire data set (best values are in bold).

The heteroscedastic model with three components obtains better values of the information criteria than the locally independent model since it models the

intra-component dependencies. Moreover, as shown by Table 5, these intra-component dependencies influence the resulting partition since four individuals are affiliated in different classes by both models.

		hetero.		
		Class 1	Class 2	Class 3
	Class 1	33	1	0
loc. indpt.	Class 2	1	402	0
	Class 3	2	0	78

Table 5: Confusion matrix between the partition provides by the locally independent model with three components (rows) and the heteroscedastic Gaussian copula mixture with three components (columns).

Interpretation of best-fit model

The three-step interpretation of the heteroscedastic tri-component Gaussian copula mixture model is presented by using the parameters summarized in Figure 6:

- Class 1 (*unpredictable fires*): this is the minority class ($\pi_1 = 0.09$). It groups fires difficultly predictable since they occur with small values of the four FWI system variables (especially ISI). These fires occur during all the year only with dry weather.
- Class 2 (*predictable fires of summer*): this is the majority class ($\pi_2 = 0.78$). This class is composed with fires occurring mainly in summer. They appear with high values of the four FWI system variables and a high temperature.
- Class 3 (*fires of winter*): this is a moderate class ($\pi_3 = 0.13$). It groups fires occurring in winter, so with small temperature. They occur with a dry weather and very small values of the four FWI systems (except ISI).

The correlation matrices highlight dependencies between the summer period and high temperatures, and between FFMC and DMC values (see Figure 7(b) for component 3). Finally, it should be noted that the variable indicating the day of the week roughly follows the same distribution for all three classes.

The results of the PCA done according to component 3 is shown in Figure 7. Obviously, Figure 7(a) shows that the individuals belonging to component 3 are strongly different to the other ones. Indeed, few individuals belonging to component 2 are visible in this map. The other individuals are too far away from the origin. Thus, the distribution of component 3 is strongly different from the other distributions.

Conclusion

The cluster analysis obtained with the Gaussian copula mixture model is more accurate than that obtained with the locally independent model. Indeed, it provides a meaningful model which needs less components. This model puts the light on three kinds of fires: the fires predictable with the FWI system (class 2), the fires occurring in winter (class 3) and the unpredictable fires (class 1).

6 Conclusion and future extensions

A Gaussian copula mixture model has been introduced and used to cluster mixed data. Using Gaussian copulas, the one-dimensional marginal distributions of each component follow conventional distributions, and intra-class dependencies are effectively modelled. Thus the model results can be easily interpreted in three steps, as in the case of models developed for data sets composed of a single type of variable. Using the continuous latent variables of Gaussian copulas, a PCA-type method displays for component-based visualization of individuals. Moreover, this approach provides a summary of intra-component dependencies, which avoids fastidious interpretation of correlation matrices.

In the description of numerical experiments and applications, we pointed out that this model is sufficiently robust to fit data obtained from another model. Furthermore, it can reduce the bias produced by the locally independent model (*e.g.* reduction of the number of component).

The number of parameters increases with the number of components and number of variables, particularly due to the correlation matrices of the Gaussian copulas. In order to overcome this drawback, we have proposed an homoscedastic version of the model assuming equality between correlation matrices. However, the number of parameters required by this model can stay large when the number of variables increases. Therefore, more parsimonious correlation matrices could be proposed to overcome this drawback in future studies.

Since the distribution of all the variables is modeled, this model could be used to manage data sets with missing values. By assuming that values are missing at random, the Gibbs sampler could also be adapted, but the underlying principle would remain roughly the same.

Finally, the proposed model cannot cluster non-ordinal categorical variables having more than two modalities. In such cases, the cumulative distribution function is not defined. An artificial order between modalities could be added to define a cumulative distribution function, but this method presents three potential difficulties that require attention: it assumes regular dependencies between the modalities of two variables, its estimation would slow down the estimation algorithm, and its stability would have to be verified.

MixCluster (https://r-forge.r-project.org/R/?group_id=1939) is an R package which performs the cluster analysis method described in the article. It also contains the data sets used in this paper.

References

- Banfield, J., Raftery, A., 1993. Model-based Gaussian and non-Gaussian clustering. *Biometrics*, 803–821.
- Barnard, J., McCulloch, R., Meng, X., 2000. Modeling covariance matrices in terms of standard deviations and correlations, with application to shrinkage. *Statistica Sinica* 10 (4), 1281–1312.
- Biernacki, C., Celeux, G., Govaert, G., 2000. Assessing a mixture model for clustering with the integrated completed likelihood. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 22 (7), 719–725.
- Cortez, P., Morais, A., 2007. A data mining approach to predict forest fires using meteorological data. *Associação Portuguesa para a Inteligência Artificial (APPIA)*.
- Everitt, B., 1988. A finite mixture model for the clustering of mixed-mode data. *Statistics & Probability Letters* 6 (5), 305–309.
- Frühwirth-Schnatter, S., 2006. *Finite mixture and Markov switching models*. Springer.
- Goodman, L., 1974. Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* 61 (2), 215–231.
- Gouget, C., 2006. Utilisation des modèles de mélange pour la classification automatique de données ordinales. PhD thesis, Université de Technologie de Compiègne.
- Hand, D., Yu, K., 2001. Idiot’s Bayes - Not So Stupid after All? *International Statistical Review* 69 (3), 385–398.
- Hoff, P., 2007. Extending the rank likelihood for semiparametric copula estimation. *The Annals of Applied Statistics*, 265–283.
- Hoff, P., Niu, X., Wellner, J., 2011. Information bounds for Gaussian copulas. arXiv preprint arXiv:1110.3572.
- Hunt, L., Jorgensen, M., 1999. Theory & Methods: Mixture model clustering using the MULTIMIX program. *Australian & New Zealand Journal of Statistics* 41 (2), 154–171.
- Hunt, L., Jorgensen, M., 2011. Clustering mixed data. *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 1 (4), 352–361.
- Jacques, J., Biernacki, C., 2014. Model-based clustering for multivariate partial ranking data. *Journal of Statistical Planning and Inference* 149, 201–217.
- Joe, H., 1997. *Multivariate models and multivariate dependence concepts*. Vol. 73. CRC Press.

- Karlis, D., Tsiamyrtzis, P., 2008. Exact Bayesian modeling for bivariate Poisson data and extensions. *Statistics and Computing* 18 (1), 27–40.
- Klaassen, C., Wellner, J., 1997. Efficient estimation in the bivariate normal copula model: normal margins are least favourable. *Bernoulli* 3 (1), 55–77.
- Krzanowski, W., 1993. The location model for mixtures of categorical and continuous variables. *Journal of Classification* 10 (1), 25–49.
- Kullback, S., Leibler, R. A., 1951. On information and sufficiency. *The Annals of Mathematical Statistics*, 79–86.
- Lebarbier, E., Mary-Huard, T., 2006. Une introduction au critère BIC : fondements théoriques et interprétation. *Journal de la SFdS* 147 (1), 39–57.
- Lewis, D., 1998. Naive (Bayes) at forty: The independence assumption in information retrieval. In: *Machine learning: ECML-98*. Springer, pp. 4–15.
- Maugis, C., Celeux, G., Martin-Magniette, M., 2009. Variable selection in model-based clustering: A general variable role modeling. *Computational Statistics & Data Analysis* 53 (11), 3872–3882.
- McLachlan, G., Peel, D., 2000. *Finite mixture models*. Wiley Series in Probability and Statistics: Applied Probability and Statistics, Wiley-Interscience, New York.
- Morlini, I., 2012. A latent variables approach for clustering mixed binary and continuous variables within a Gaussian mixture model. *Advances in Data Analysis and Classification* 6 (1), 5–28.
- Moustaki, I., Papageorgiou, I., 2005. Latent class models for mixed variables with applications in archaeometry. *Computational Statistics & Data Analysis* 48 (3), 659–675.
- Murray, J., Dunson, D., Carin, L., Lucas, J., 2013. Bayesian Gaussian copula factor models for mixed data. *Journal of the American Statistical Association* 108 (502), 656–665.
- Nelsen, R., 1999. *An introduction to copulas*. Springer.
- Olsson, U., 1979. Maximum likelihood estimation of the polychoric correlation coefficient. *Psychometrika* 44 (4), 443–460.
- Pitt, M., Chan, D., Kohn, R., 2006. Efficient Bayesian inference for Gaussian copula regression models. *Biometrika* 93 (3), 537–554.
- Raftery, A., 1996. Hypothesis testing and model selection. In: *Markov chain Monte Carlo in practice*. Springer, pp. 163–187.
- Robert, C., 2007. *The Bayesian choice: from decision-theoretic foundations to computational implementation*. Springer.

- Robert, C., Casella, G., 2004. Monte Carlo statistical methods. Springer Verlag.
- Schwarz, G., 1978. Estimating the dimension of a model. *Annals of Statistics* 6, 461–464.
- Smith, M., Khaled, M., 2012. Estimation of copula models with discrete margins via Bayesian data augmentation. *Journal of the American Statistical Association* 107 (497), 290–303.
- Song, P. X.-K., Fan, Y., Kalbfleisch, J. D., 2005. Maximization by parts in likelihood inference. *Journal of the American Statistical Association* 100 (472), 1145–1158.
- Stephens, M., 2000. Dealing with label switching in mixture models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 62 (4), 795–809.
- Teicher, H., 1963. Identifiability of Finite Mixtures. *The Annals of Mathematical Statistics*, 1265–1269.
- Van Hattum, P., Hoijsink, H., 2009. Market Segmentation Using Brand Strategy Research: Bayesian Inference with Respect to Mixtures of Log-Linear Models. *Journal of Classification* 26 (3), 297–328.
- Willse, A., Boik, R., 1999. Identifiable finite mixtures of location models for clustering mixed-mode data. *Statistics and Computing* 9 (2), 111–121.
- Yakowitz, S., Spragins, J., et al., 1968. On the identifiability of finite mixtures. *The Annals of Mathematical Statistics* 39 (1), 209–214.

A Proof of the model identifiability

The model identifiability is proved by two propositions. The first proposition proves the model identifiability when the variables are continuous and/or integer. This proposition presents the reasoning in a simple case since it does not consider the ordinal variables. The second proposition proves that the model requires at least one continuous or integer variable to be identifiable.

Proposition A.1 (Identifiability with continuous and integer variables). *The Gaussian copula mixture model is weakly identifiable (?) if the variables are continuous and integer ones (i.e. the one-dimensional marginal distributions of the components are Gaussian or Poisson distributions). Thus,*

$$\forall \mathbf{x} \in \mathbb{R}^c \times \mathbb{N}^d, \quad \sum_{k=1}^g \pi_k p(\mathbf{x} | \boldsymbol{\alpha}_k) = \sum_{k=1}^{g'} \pi'_k p(\mathbf{x} | \boldsymbol{\alpha}'_k) \quad (10)$$

$$\Rightarrow g = g', \quad \boldsymbol{\pi} = \boldsymbol{\pi}', \quad \boldsymbol{\alpha} = \boldsymbol{\alpha}'. \quad (11)$$

Proof. The identifiability of the multivariate Gaussian mixture models and of the univariate Poisson mixture model (??) involves that (10) implies

$$g = g', \quad \boldsymbol{\pi} = \boldsymbol{\pi}', \quad \boldsymbol{\beta}_{kj} = \boldsymbol{\beta}'_{kj} \text{ and } \boldsymbol{\Gamma}_{kCC} = \boldsymbol{\Gamma}'_{kCC}. \quad (12)$$

We now show that $\boldsymbol{\Gamma}_{kCD} = \boldsymbol{\Gamma}'_{kCD}$ and $\boldsymbol{\Gamma}_{kDD} = \boldsymbol{\Gamma}'_{kDD}$.

Let $j \in \{1, \dots, c\}$ and $h \in \{c+1, \dots, e\}$. We denote by $\rho_k = \boldsymbol{\Gamma}_k(j, h)$, $\rho'_k = \boldsymbol{\Gamma}'_k(j, h)$, $v_k = \Phi_1^{-1}(P(x^j; \boldsymbol{\beta}_{kj}))$, $\varepsilon_k(x^j) = \pi_k \frac{\phi(v_k)}{\sigma_{kj}}$, $a_k = \frac{b_k^{\oplus}(x^j) - \rho_k v_k}{\sqrt{1 - \rho_k^2}}$ and $a'_k = \frac{b_k^{\oplus}(x^j) - \rho'_k v_k}{\sqrt{1 - \rho_k'^2}}$. Without loss of generality, we order the components as such $\sigma_{kj} > \sigma_{k+1j}$ and if $\sigma_{kj} = \sigma_{k+1j}$ then $\mu_{kj} > \mu_{k+1j}$, then (10) implies that

$$1 + \sum_{k=2}^g (\varepsilon_k(x^j) \Phi(a_k)) / (\varepsilon_1(x^j) \Phi(a_1)) = \sum_{k=1}^g \varepsilon_k(x^j) \Phi(a'_k) / (\varepsilon_1(x^j) \Phi(a_1)).$$

Let $\gamma_t = \{(x^j, x^h) \in \mathbb{R} \times \mathbb{N} : a_1 = t\}$. Then, letting $x^h \rightarrow \infty$ as such $(x^j, x^h) \in \gamma_t$,

$$\forall t, \quad \frac{\int_t^{a'_1} \phi(u) du}{\Phi(t)} = 0. \quad (13)$$

Thus $a'_1 = a_1$, so $\rho'_1 = \rho_1$. Repeating this argument for $k = 2, \dots, g$ and for all the couples (j, h) , we conclude that $\boldsymbol{\Gamma}_{kCD} = \boldsymbol{\Gamma}'_{kCD}$.

When both variables are integer, we use the same argument with $\gamma_{(t, \xi)} = \{(x^j, x^h) \in \mathbb{N} \times \mathbb{N} : a_1 \in B(t, \xi)\}$. Note that if $\rho_1 \neq \rho'_1$ then $\exists n_0$ as such $\forall x^j > n_0$ $a'_1 > t + \xi$. Letting $x^h \rightarrow \infty$ as such $(x^j, x^h) \in \gamma_{(t, \xi)}$, we obtain the following contradiction $\frac{\int_t^{a'_1} \phi(u) du}{\Phi(t - \xi)} = 0$ and $\frac{\int_t^{a'_1} \phi(u) du}{\Phi(t - \xi)} > 0$. So, $a'_1 = a_1$ then $\rho_1 = \rho'_1$. Repeating this argument for $k = 2, \dots, g$ and for all the couples (j, h) , we conclude that $\boldsymbol{\Gamma}_{kDD} = \boldsymbol{\Gamma}'_{kDD}$. \square

Proposition A.2 (Identifiability of the Gaussian copula mixture model). *The Gaussian copula mixture model is weakly identifiable (?) if at least one variable is continuous or integer.*

Proof. In this proof, we consider only one continuous variable and two binary variables. Obviously, the same reasoning can be extend to the other cases. We now show that $\boldsymbol{\Gamma}_{kCD} = \boldsymbol{\Gamma}'_{kCD}$ and $\boldsymbol{\Gamma}_{kDD} = \boldsymbol{\Gamma}'_{kDD}$.

Let $j = 1$ and let $h \in \{2, 3\}$. We note $\rho_k = \boldsymbol{\Gamma}_k(j, h)$, $\rho'_k = \boldsymbol{\Gamma}'_k(j, h)$, $v_k = \Phi_1^{-1}(P(x^j; \boldsymbol{\beta}_{kj}))$, $\varepsilon_k(x^j) = \pi_k \frac{\phi(v_k; 0, 1)}{\sigma_{kj}}$, $a_k = \frac{b_k^{\oplus}(x^j) - \rho_k v_k}{\sqrt{1 - \rho_k^2}}$ and $a'_k = \frac{b_k^{\oplus}(x^j) - \rho'_k v_k}{\sqrt{1 - \rho_k'^2}}$. Without loss of generality, we order the components as such $\sigma_{kj} > \sigma_{[k+1]j}$ and if $\sigma_{kj} = \sigma_{[k+1]j}$ then $\mu_{kj} > \mu_{[k+1]j}$. Note that (10) implies that

$$1 + \sum_{k=2}^g (\varepsilon_k(x^j) \Phi(a_k)) / (\varepsilon_1(x^j) \Phi(a_1)) = \sum_{k=1}^g \varepsilon_k(x^j) \Phi(a'_k) / (\varepsilon_1(x^j) \Phi(a_1)).$$

Letting $x^1 \rightarrow \infty$ and assuming that $\rho_k > 0$ then $\frac{\Phi(a'_k)}{\Phi(a_k)} = 1$. So, $\text{sign}(\rho_k) = \text{sign}(\rho'_k)$. By denoting $\kappa = \lim_{a \rightarrow \infty} \frac{\phi(a)}{\Phi(a)}$ and letting $x^1 \rightarrow \infty$ $\kappa \frac{1}{\kappa} \frac{\phi(a'_k)}{\phi(a_k)} = 1$. Thus $a'_1 = a_1$, so $\rho'_1 = \rho_1$ and $b_k^\oplus(x^j) = b'_k{}^\oplus(x^j)$ so $\beta_{kh} = \beta'_{kh}$.

Note that the same result can be obtained by tending x^1 to $-\infty$ is $\rho_k < 0$. Repeating this argument for $k = 2, \dots, g$ and for all the couples (j, h) , we conclude that $\mathbf{\Gamma}_{k\text{CD}} = \mathbf{\Gamma}'_{k\text{CD}}$ then $\mathbf{\Gamma}_{k\text{DD}} = \mathbf{\Gamma}'_{k\text{DD}}$. \square

B Prior distributions

We assume independence between the parameters as follows

$$p(\boldsymbol{\theta}) = p(\boldsymbol{\pi}) \prod_{k=1}^g \left(p(\mathbf{\Gamma}_k) \prod_{j=1}^e p(\boldsymbol{\beta}_{kj}) \right). \quad (14)$$

The classical conjugate prior distribution of the proportion vector is the Jeffreys non informative one which is the following Dirichlet distribution

$$\boldsymbol{\pi} \sim \mathcal{D}_g \left(\frac{1}{2}, \dots, \frac{1}{2} \right). \quad (15)$$

If x^j is *continuous*, then $\boldsymbol{\beta}_{kj}$ denotes the parameters of a univariate Gaussian distribution so $p(\boldsymbol{\beta}_{kj}) = p(\mu_{kj} | \sigma_{kj}^2) p(\sigma_{kj}^2)$ with

$$\sigma_{kj}^2 \sim \mathcal{G}^{-1}(c_0, C_0) \text{ and } \mu_{kj} | \sigma_{kj}^2 \sim \mathcal{N}_1(b_0, \sigma_{kj}^2 / N_0), \quad (16)$$

where $\mathcal{G}^{-1}(\cdot, \cdot)$ denotes the inverse gamma distribution. With an empirical Bayesian approach, the hyper-parameters (c_0, C_0, b_0, N_0) are fixed as proposed by ?, so $c_0 = 1.28$, $C_0 = 0.36 \text{Var}(\mathbf{x}^j)$, $b_0 = \frac{1}{n} \sum_{i=1}^n x_i^j$ and $N_0 = \frac{2.6}{\text{argmax } \mathbf{x}^j - \text{argmin } \mathbf{x}^j}$.

If x^j is *integer*, $\boldsymbol{\beta}_{kj}$ denotes the parameter of a Poisson distribution and

$$\boldsymbol{\beta}_{kj} \sim \mathcal{G}(a_0, A_0). \quad (17)$$

According to ?, the values of hyper-parameters a_0 and A_0 are empirically fixed to $a_0 = 1$ and $A_0 = a_0 n / \sum_{i=1}^n x_i^j$.

If x^j is *ordinal*, $\boldsymbol{\beta}_{kj}$ denotes the parameter of a multinomial distribution and its Jeffreys non informative conjugate prior involves that

$$\boldsymbol{\beta}_{kj} \sim \mathcal{D}_{m_j} \left(\frac{1}{2}, \dots, \frac{1}{2} \right). \quad (18)$$

The conjugate prior of a covariance matrix is the Inverse Wishart distribution denoted by $\mathcal{W}^{-1}(\cdot, \cdot)$. Therefore, it is natural to define the prior of the correlation matrix $\mathbf{\Gamma}_k$ from the prior of the correlation matrix $\mathbf{\Lambda}_k$. Indeed, $\mathbf{\Gamma}_k | \mathbf{\Lambda}_k$ is deterministic (?). So,

$$\mathbf{\Lambda}_k \sim \mathcal{W}^{-1}(s_0, S_0) \text{ and } \forall 1 \leq h, \ell \leq e, \mathbf{\Gamma}_k[h, \ell] = \frac{\mathbf{\Lambda}_k[h, \ell]}{\sqrt{\mathbf{\Lambda}_k[h, h] \mathbf{\Lambda}_k[\ell, \ell]}}, \quad (19)$$

where (s_0, S_0) are two hyper-parameters. However, these parameters can not be fitted by an empirical Bayesian approach since \mathbf{y} is not observed. Uniform distribution on $] - 1, 1[$ is also obtained for the margin distributions of each correlation coefficient by setting $s_0 = e + 1$ and S_0 equal to the identity matrix (?).

C Metropolis-within-Gibbs sampler

This section details the four samplings used in Algorithm 3.1. The first two samplings are difficult to perform directly, so they are done by one iteration of two Metropolis-Hastings algorithms. For both Metropolis-Hastings algorithms, the instrumental distributions assume conditional independence between parameters. So, the smaller are the intra-class dependencies, the closer of the stationary distributions are the instrumental distributions of both algorithms. Finally, the last two samplings used in Algorithm 3.1 are classical.

C.1 Class membership and Gaussian vector sampling

The aim is to sample from (6) but the sampling from $\mathbf{z}, \mathbf{y} | \mathbf{x}, \boldsymbol{\theta}^{(r-1)}$ cannot be performed directly. So, it is achieved by one iteration of following Metropolis-Hastings algorithms. The sampling from (6) is performed in two steps by using independence between the individuals which involves that

$$p(\mathbf{z}, \mathbf{y} | \mathbf{x}, \boldsymbol{\theta}^{(r-1)}) = \prod_{i=1}^n p(z_i | \mathbf{x}_i, \boldsymbol{\theta}^{(r-1)}) p(\mathbf{y}_i | \mathbf{x}_i, z_i, \boldsymbol{\theta}^{(r-1)}). \quad (20)$$

Firstly, each $z_i^{(r)}$ is independently sampled from the multinomial distribution

$$z_i | \mathbf{x}_i, \boldsymbol{\theta}^{(r-1)} \sim \mathcal{M}_g(t_{i1}(\boldsymbol{\theta}^{(r-1)}), \dots, t_{ig}(\boldsymbol{\theta}^{(r-1)})), \quad (21)$$

where $t_{ik}(\boldsymbol{\theta}^{(r-1)}) = \frac{\pi_k^{(r-1)} p(\mathbf{x}_i | \boldsymbol{\alpha}_k^{(r-1)})}{p(\mathbf{x}_i | \boldsymbol{\theta}^{(r-1)})}$. Note that $t_{ik}(\boldsymbol{\theta}^{(r-1)})$ is the posterior probability that \mathbf{x}_i arises from component k with the parameters $\boldsymbol{\theta}^{(r-1)}$.

Secondly, each $\mathbf{y}_i^{(r-1/2)}$ is independently sampled given $(\mathbf{x}_i, z_i^{(r)}, \boldsymbol{\theta}^{(r-1)})$. Its first c elements, denoted by $\mathbf{y}_i^{c(r-1/2)}$, are deterministically defined by $\mathbf{y}_i^{c(r-1/2)} = \Psi(\mathbf{x}_i^c; \boldsymbol{\alpha}_{z_i^{(r)}}^{(r-1)})$. Its last d elements, denoted by $\mathbf{y}_i^{d(r-1/2)}$, are sampled from the d -variate Gaussian distribution $\mathcal{N}_d(\mathbf{0}, \boldsymbol{\Gamma}_{z_i^{(r)}}^{(r-1)})$ truncated on the space $\mathcal{S}_{z_i^{(r)}}(\mathbf{x}_i^d)$

$$p(\mathbf{y}_i^d | \mathbf{x}_i, \mathbf{y}_i^{c(r-1/2)}, z_i^{(r)}, \boldsymbol{\theta}^{(r-1)}) \propto \phi_d(\mathbf{y}_i^d; \boldsymbol{\mu}_{z_i^{(r)}}^{d(r-1)}, \boldsymbol{\Sigma}_{z_i^{(r)}}^{d(r-1)}) \mathbb{1}_{\{\mathbf{y}_i^d \in \mathcal{S}_{z_i^{(r)}}(\mathbf{x}_i^d)\}}, \quad (22)$$

where $\boldsymbol{\mu}_{z_i^{(r)}}^{d(r-1)} = \boldsymbol{\Gamma}_{z_i^{(r)} \text{DC}}^{(r-1)} \boldsymbol{\Gamma}_{z_i^{(r)} \text{CC}}^{-1(r-1)} \mathbf{y}_i^{c(r-1/2)}$.

The computation of $t_{ik}(\boldsymbol{\theta}^{(r-1)})$ involves the calculation of the integral defined in (5) which can be time consuming if d is large ($d > 6$). In such a case,

the sampling from (6) is replaced by one iteration of the Metropolis-Hastings algorithm which independently samples each couple (z_i, \mathbf{y}_i) . Its stationary distribution is

$$p(z_i, \mathbf{y}_i | \mathbf{x}_i, \boldsymbol{\theta}^{(r-1)}) \propto \pi_{z_i} p(\mathbf{x}_i, \mathbf{y}_i | z_i, \boldsymbol{\theta}^{(r-1)}). \quad (23)$$

Note that $p(\mathbf{x}_i, \mathbf{y}_i | z_i, \boldsymbol{\theta}^{(r-1)}) = \phi_e(\mathbf{y}_i; \mathbf{0}, \boldsymbol{\Gamma}_{z_i}^{(r-1)}) \mathbb{1}_{\{\mathbf{y}_i^c = \Psi(\mathbf{x}_i^c; \boldsymbol{\alpha}_{z_i}^{(r-1)})\}} \mathbb{1}_{\{\mathbf{y}_i^D \in \mathcal{S}_{z_i}(\mathbf{x}_i^D)\}}$.

The Metropolis-Hastings algorithm samples a candidate (z_i^*, \mathbf{y}_i^*) by the instrumental distribution $q_1(\cdot | \mathbf{x}_i, \boldsymbol{\theta}^{(r-1)})$ which uniformly samples z_i^* then which samples $\mathbf{y}_i^* | z_i^*$ as follows. Its first c elements, denoted by \mathbf{y}_i^{*c} , are equal to $\mathbf{y}_i^{*c} = \Psi(\mathbf{x}_i^c; \boldsymbol{\alpha}_{z_i^*}^{(r-1)})$. Its last d elements, denoted by \mathbf{y}_i^{*D} , follow a *multivariate independent Gaussian* distribution truncated on $\mathcal{S}_{z_i^*}(\mathbf{x}_i^D)$. Thus,

$$q_1(z_i, \mathbf{y}_i | \mathbf{x}_i, \boldsymbol{\theta}^{(r-1)}) = \frac{1}{g} \frac{\phi_d(\mathbf{y}_i^D | \mathbf{0}, \mathbf{I}) \mathbb{1}_{\{\mathbf{y}_i^D \in \mathcal{S}_{z_i}(\mathbf{x}_i^D)\}}}{\prod_{j=c+1}^e p(x_i^j; \boldsymbol{\beta}_{z_{ij}}^{(r-1)})} \mathbb{1}_{\{\mathbf{y}_i^c = \Psi(\mathbf{x}_i^c | \boldsymbol{\alpha}_{z_i}^{(r-1)})\}}. \quad (24)$$

The candidate is accepted with the probability

$$\rho_{1i}^{(r)} = \min \left\{ \frac{q_1(z_i^{(r-1)}, \mathbf{y}_i^{(r-1)} | \mathbf{x}_i)}{q_1(z_i^*, \mathbf{y}_i^* | \mathbf{x}_i)} \frac{\pi_{z_i^*} \phi_e(\mathbf{y}_i^*; \mathbf{0}, \boldsymbol{\Gamma}_{z_i^*}^{(r-1)})}{\pi_{z_i^{(r-1)}} \phi_e(\mathbf{y}_i^{(r-1)}; \mathbf{0}, \boldsymbol{\Gamma}_{z_i^{(r-1)}}^{(r-1)})}; 1 \right\}. \quad (25)$$

Thus, at iteration (r) of Algorithm 3.1, the sampling according to (6) is performed via one iteration of the following Metropolis-Hastings algorithm having $p(z_i, \mathbf{y}_i | \mathbf{x}_i, \boldsymbol{\theta}^{(r-1)})$ as stationary distribution.

Algorithm C.1.

$$(z_i^*, \mathbf{y}_i^*) \sim q_1(z, \mathbf{y} | \mathbf{x}_i) \quad (26)$$

$$(z_i^{(r)}, \mathbf{y}_i^{(r-1/2)}) = \begin{cases} (z_i^*, \mathbf{y}_i^*) & \text{with probability } \rho_{1i}^{(r)} \\ (z_i^{(r-1)}, \mathbf{y}_i^{(r-1)}) & \text{with probability } 1 - \rho_{1i}^{(r)}. \end{cases} \quad (27)$$

C.2 Margin parameter and Gaussian vector sampling

The aim is to sample from (7) but the sampling from $\boldsymbol{\beta}_{kj} | \mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \boldsymbol{\beta}_{k\bar{j}}^{(r)}, \boldsymbol{\Gamma}_k$ cannot be performed directly. So, it is achieved by one iteration of the following Metropolis-Hastings algorithms. The sampling from (7) is performed by using the following decomposition

$$p(\boldsymbol{\beta}_{kj}, \mathbf{y}_{[rk]}^j | \mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \boldsymbol{\beta}_{k\bar{j}}^{(r)}, \boldsymbol{\Gamma}_k^{(r-1)}) = p(\boldsymbol{\beta}_{kj} | \mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \boldsymbol{\beta}_{k\bar{j}}^{(r)}, \boldsymbol{\Gamma}_k^{(r-1)}) \\ \times p(\mathbf{y}_{[rk]}^j | \mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \boldsymbol{\beta}_{k\bar{j}}^{(r)}, \boldsymbol{\beta}_{kj}, \boldsymbol{\Gamma}_k^{(r-1)}). \quad (28)$$

Parameter $\boldsymbol{\beta}_{kj}^{(r)}$ is firstly sampled. The full conditional distribution of $\boldsymbol{\beta}_{kj}$ is defined up to a normalizing constant such as

$$p(\boldsymbol{\beta}_{kj} | \mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \boldsymbol{\beta}_{k\bar{j}}^{(r)}, \boldsymbol{\Gamma}_k^{(r-1)}) \propto p(\boldsymbol{\beta}_{kj}) \prod_{\{i: z_i^{(r)}=k\}} p(x_i^j | \mathbf{y}_i^{\bar{j}(r)}, z_i^{(r)}, \boldsymbol{\Gamma}_k^{(r-1)}, \boldsymbol{\beta}_{kj}). \quad (29)$$

The distribution of $x_i^j | \mathbf{y}_i^{\bar{j}(r)}, z_i^{(r)}, \mathbf{\Gamma}_k^{(r-1)}$ with $z_i^{(r)} = k$ is defined by

$$p(x_i^j | \mathbf{y}_i^{\bar{j}(r)}, z_i^{(r)}, \mathbf{\Gamma}_k^{(r-1)}, \beta_{kj}) = \begin{cases} \phi_1\left(\frac{x_i^j - \mu_{kj}}{\sigma_{kj}}; \tilde{\mu}_i, \tilde{\sigma}_i^2\right) / \sigma_{kj} & \text{if } 1 \leq j \leq c \\ \Phi_1\left(\frac{b^\oplus(x_i^j) - \tilde{\mu}_i}{\tilde{\sigma}_i} - \Phi_1\left(\frac{b^\ominus(x_i^j) - \tilde{\mu}_i}{\tilde{\sigma}_i}\right)\right) & \text{otherwise,} \end{cases} \quad (30)$$

where the real $\tilde{\mu}_i = \mathbf{\Gamma}_k^{(r-1)}[j, \bar{j}] \mathbf{\Gamma}_k^{(r-1)}[\bar{j}, \bar{j}]^{-1} \mathbf{y}_i^{\bar{j}(r)}$ is the full conditional mean of y_i^j , $\mathbf{\Gamma}_k[j, \bar{j}]$ being the row j of $\mathbf{\Gamma}_k$ deprived of element j and $\mathbf{\Gamma}_k[\bar{j}, \bar{j}]$ being the matrix $\mathbf{\Gamma}_k$ deprived of the row and the column j , and where $\tilde{\sigma}_i^2$ is the full conditional variance of y_i^j defined by $\tilde{\sigma}_i^2 = 1 - \mathbf{\Gamma}_k^{(r-1)}[j, \bar{j}] \mathbf{\Gamma}_k^{(r-1)}[\bar{j}, \bar{j}]^{-1} \mathbf{\Gamma}_k^{(r-1)}[\bar{j}, j]$. As the normalizing constant of (29) is unknown, $\beta_{kj}^{(r)}$ cannot be directly sampled. This problem is avoided by one iteration of the Metropolis-Hastings algorithm. The instrumental distribution of this Metropolis-Hastings algorithm $q_2(\cdot | \mathbf{x}, \mathbf{z})$ samples a candidate β_{kj}^* according to the posterior distribution of β_{kj} under conditional independence assumption (this distribution is explicit since the conjugate prior distributions are used). So, $q_2(\cdot | \mathbf{x}, \mathbf{z}) = p(\beta_{kj} | \mathbf{x}, \mathbf{z}, \mathbf{\Gamma}_k = \mathbf{I})$. Thus, according to (29), the candidate β_{kj}^* is accepted with the probability

$$\rho_2^{(r)} = \min \left\{ \frac{q_2(\beta_{kj}^{(r-1)} | \mathbf{x}, \mathbf{z}) p(\beta_{kj}^*)}{q_2(\beta_{kj}^* | \mathbf{x}, \mathbf{z}) p(\beta_{kj}^{(r-1)})} \prod_{\{i: z_i^{(r)} = k\}} \frac{p(x_i^j | \mathbf{y}_i^{\bar{j}(r)}, z_i, \mathbf{\Gamma}_k^{(r-1)}, \beta_{kj}^*)}{p(x_i^j | \mathbf{y}_i^{\bar{j}(r)}, z_i, \mathbf{\Gamma}_k^{(r-1)}, \beta_{kj}^{(r-1)})}; 1 \right\}.$$

Thus, at iteration (r) of Algorithm 3.1, step (7) is performed via one iteration of the following Metropolis-Hastings algorithm whose the stationary distribution is $p(\beta_{kj} | \mathbf{x}_{[rk]}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}, \beta_{kj}^{(r)}, \mathbf{\Gamma}_k)$.

Algorithm C.2.

$$\beta_{kj}^* \sim q_2(\beta_{kj} | \mathbf{x}, \mathbf{z}) \quad (31)$$

$$\beta_{kj}^{(r)} = \begin{cases} \beta_{kj}^* & \text{with probability } \rho_2^{(r)} \\ \beta_{kj}^{(r-1)} & \text{with probability } 1 - \rho_2^{(r)}. \end{cases} \quad (32)$$

Vector $\mathbf{y}_{[rk]}^j$ is easily sampled after $\beta_{kj}^{(r)}$. Indeed, independence between the individuals defines the full conditional distribution of $\mathbf{y}_{[rk]}^j$ by

$$p(\mathbf{y}_{[rk]}^j | \mathbf{x}, \mathbf{y}_{[rk]}^{\bar{j}(r)}, \mathbf{z}^{(r)}, \beta_{kj}^{(r)}, \beta_{kj}, \mathbf{\Gamma}_k^{(r-1)}) = \prod_{\{i: z_i^{(r)} = k\}} p(y_i^j | x_i^j, \mathbf{y}_i^{\bar{j}(r)}, z_i^{(r)}, \beta_{kj}, \mathbf{\Gamma}_k^{(r-1)}). \quad (33)$$

If x^j is a continuous variable (*i.e.* $1 \leq j \leq c$), when $z_i^{(r)} = k$, the full conditional distribution of y_i^j is a Dirac distribution in $\frac{x_i^j - \mu_{kj}^{(r)}}{\sigma_{kj}^{(r)}}$. If x^j is a discrete variable (*i.e.* $c+1 \leq j \leq e$), when $z_i^{(r)} = k$, the full conditional distribution of y_i^j is a truncated Gaussian distribution such as,

$$p(y_i^j | x_i^j, \mathbf{y}_i^{\bar{j}(r)}, z_i^{(r)}, \beta_{kj}^{(r)}, \mathbf{\Gamma}_k^{(r-1)}) = \frac{\phi_1(y_i^j; \tilde{\mu}_i, \tilde{\sigma}_i^2)}{p(x_i^j; \beta_{kj}^{(r)})} \mathbb{1}_{\{y_i^j \in [b_k^\ominus(x_i^j), b_k^\oplus(x_i^j)]\}}, \quad (34)$$

So, step (7) is performed in two steps. First, $\beta_{kj}^{(r)}$ is sampled via one iteration of the Metropolis-Hastings algorithm whose the stationary distribution is $p(\beta_{kj} | \mathbf{x}, \mathbf{y}_{[rk]}^{j(r)}, \mathbf{z}^{(r)}, \beta_{k\bar{j}}^{(r)}, \mathbf{\Gamma}_k)$. Second, $\mathbf{y}_{[rk]}^{j(r)}$ is sampled from (34).

C.3 Proportion vector sampling

The aim is to sample from (8). The sampling from (8) is classical. Indeed, the conjugate Jeffreys non informative prior involves that

$$\boldsymbol{\pi} | \mathbf{z}^{(r)} \sim \mathcal{D}_g \left(n_1^{(r)} + \frac{1}{2}, \dots, n_g^{(r)} + \frac{1}{2} \right), \quad (35)$$

where $n_k^{(r)} = \sum_{i=1}^n \mathbf{1}_{\{z_i^{(r)}=k\}}$.

C.4 Correlation matrix sampling


The aim is to sample from (9). To sample from (9), we use the approach proposed by ? in the case of semiparametric Gaussian copula. First, a covariance matrix is generated by its explicit posterior distribution, and second, the correlation matrix is deduced by normalizing the covariance matrix. As (\mathbf{y}, \mathbf{z}) are known in this step, we are in the well-known case of a multivariate Gaussian mixture model with known means. Thus, the sampling according to $\mathbf{\Gamma}_k | \mathbf{y}^{(r)}, \mathbf{z}^{(r)}$ is performed by the following two steps

$$\mathbf{\Lambda}_k | \mathbf{y}^{(r)}, \mathbf{z}^{(r)} \sim \mathcal{W}^{-1} \left(s_0 + n_k^{(r-1)}, S_0 + \sum_{\{i: z_i^{(r)}=k\}} \mathbf{y}_i^{(r)T} \mathbf{y}_i^{(r)} \right), \quad (36)$$


where $\forall 1 \leq h, \ell \leq e$, $\mathbf{\Gamma}_k[h, \ell] = \frac{\mathbf{\Lambda}_k[h, \ell]}{\sqrt{\mathbf{\Lambda}_k[h, h] \mathbf{\Lambda}_k[\ell, \ell]}}$. As the homoscedastic model assumes the equality between the correlation matrices, in such a case we only sample one $\mathbf{\Lambda}$ so (36) is replaced by

$$\mathbf{\Lambda} | \mathbf{y}^{(r)}, \mathbf{z}^{(r)} \sim \mathcal{W}^{-1} \left(s_0 + n, S_0 + \sum_{i=1}^n \mathbf{y}_i^{(r)T} \mathbf{y}_i^{(r)} \right), \quad (37)$$


and we put $\mathbf{\Lambda}_k = \mathbf{\Lambda}$ for $k = 1, \dots, g$.



(a) Scatterplot of the individuals described by three variables: one continuous (abscissa), one integer (ordinate) and one binary (symbol). Colors indicate the component memberships




(b) Scatterplot of the individuals in the first factorial map of component 2. Colors and symbols indicate the component memberships




(c) representation of the variables in the first factorial map of component 2

Figure 1: Example of data visualization.




(a) Kullback-Leibler divergence of the estimated model from the true model.




(b) Misclassification error rate.


Figure 2: Boxplots of indicators of the good behavior of the estimation procedure for different sample sizes.




(a) Misclassification error rate




(b) Kullback-Leibler divergence from the true model




(c) Correlation coefficient²⁷ between both variables for component 1





(a) Scatterplot of the individuals



(b) Correlation circle

Figure 4: Visualization throughout parameters of component 3 for the tri-component homoscedastic Gaussian copula mixture model.

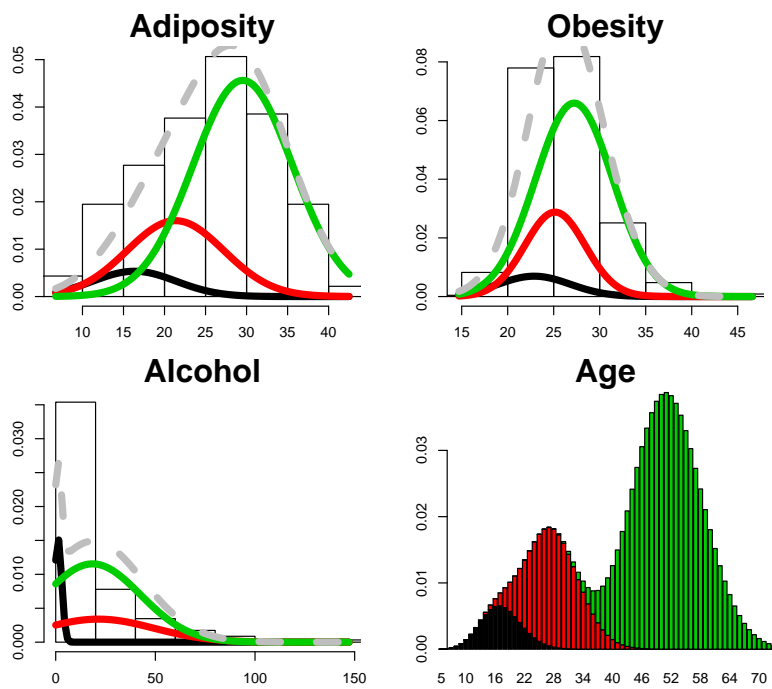


Figure 5: One dimensional marginal distributions for the South Africa Heart data: whole distribution (gray), component 1 (black), component 2 (red) and component 3 (green) of the tri-component homoscedastic Gaussian copula mixture model.

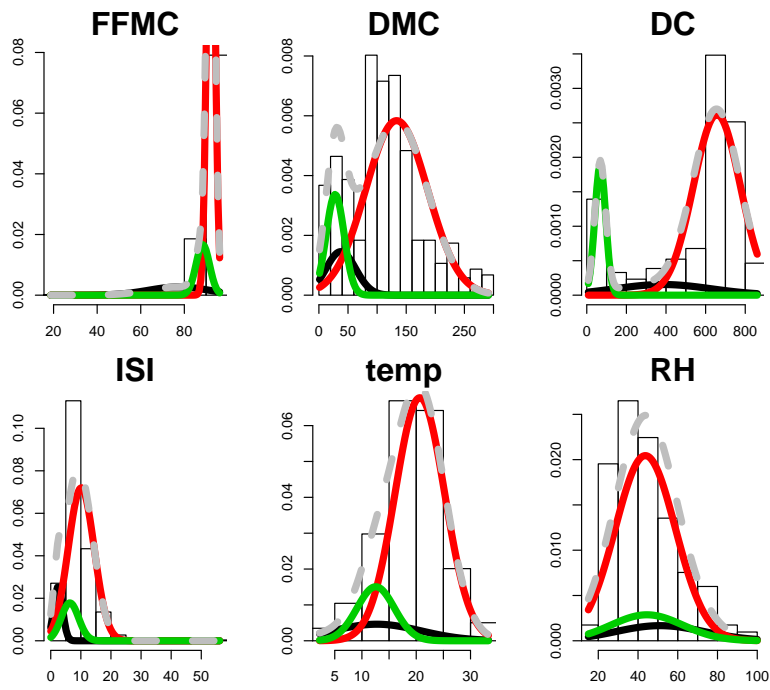




Figure 6: One dimensional marginal distributions for the forest fire data: whole distribution (gray), component 1 (black), component 2 (red) and component 3 (green) of the tri-component homoscedastic Gaussian copula mixture model.



(a) Scatterplot of the individuals



(b) Correlation circle

Figure 7: Visualization throughout parameters of component 3 for the tri-component heteroscedastic Gaussian copula mixture model.