



Matrix Normal Cluster-Weighted Models

Salvatore D. Tomarchio¹ · Paul D. McNicholas² · Antonio Punzo¹

Accepted: 27 April 2021 / Published online: 2 June 2021
© The Author(s) 2021

Abstract

Finite mixtures of regressions with fixed covariates are a commonly used model-based clustering methodology to deal with regression data. However, they assume assignment independence, i.e., the allocation of data points to the clusters is made independently of the distribution of the covariates. To take into account the latter aspect, finite mixtures of regressions with random covariates, also known as cluster-weighted models (CWMs), have been proposed in the univariate and multivariate literature. In this paper, the CWM is extended to matrix data, e.g., those data where a set of variables are simultaneously observed at different time points or locations. Specifically, the cluster-specific marginal distribution of the covariates and the cluster-specific conditional distribution of the responses given the covariates are assumed to be matrix normal. Maximum likelihood parameter estimates are derived using an expectation-conditional maximization algorithm. Parameter recovery, classification assessment, and the capability of the Bayesian information criterion to detect the underlying groups are investigated using simulated data. Finally, two real data applications concerning educational indicators and the Italian non-life insurance market are presented.

Keywords Mixture models · Matrix-variate · Classification · Clustering

1 Introduction

Finite mixture models have seen increasing use over the last decades (for a recent survey, see McNicholas 2016b). Because of their flexibility, they are a suitable statistical tool for modeling a wide range of phenomena characterized by unobserved heterogeneity, and constitute a powerful device for clustering and classification. Specifically, in a typical “direct application,” each mixture component represents a group (or cluster) within data, and the scope is to identify these groups and estimate the parameters of the conditional-group distributions (McLachlan and Peel 2000; McNicholas 2016a). If no exogenous variables explain the means and the variances of each component, they are also called unconditional mixture models. However, when there is a linear relationship between some variables, important

✉ Salvatore D. Tomarchio
daniele.tomarchio@unict.it

¹ Department of Economics and Business, University of Catania, Catania, Italy

² Department of Mathematics and Statistics, McMaster University, Hamilton, Canada

insight can be gained by accounting for functional dependencies between them. For this reason, finite mixtures of regression models with fixed covariates (FMR) have been proposed in the literature (see DeSarbo and Cron 1988; Frühwirth-Schnatter 2006, for examples). Finite mixtures of regression models with concomitant variables (FMRC; Dayton and Macready 1988) are an extension of FMR where the mixing weights depend on some concomitant variables (which are often the same covariates) and are usually modeled by a multinomial logistic model (see Ingrassia and Punzo 2016, 2020 and Mazza et al. 2019 for details). Unfortunately, these methodologies do not explicitly use the distribution of the covariates for clustering, i.e., the assignment of data points to clusters does not directly utilize any information from the distribution of the covariates.

As an alternative to these approaches, finite mixtures of regressions with random covariates (Gershensfeld 1997; Gershensfeld et al. 1999), also known as cluster-weighted models (CWMs), allow for such functional dependency. This occurs because, for each mixture component, CWMs decompose the joint distribution of responses and covariates into the product between the marginal distribution of the covariates and the conditional distribution of the responses given the covariates. Several CWMs have been introduced in the univariate and multivariate literature. Most of them consider a univariate response variable, along with a set of covariates, modeled by a univariate and a multivariate distribution, respectively (see Ingrassia et al. 2012, 2014; Punzo 2014, for examples). Fewer CWMs exist in the case of a multivariate response (Punzo and McNicholas 2017; Dang et al. 2017).

In recent years, there has been an increasing interest in applications involving matrix-variate (three-way) data, e.g., Viroli (2011), Anderlucci et al. (2015), Gallagher and McNicholas (2017, 2018), Melnykov and Zhu (2018, 2019), and Tomarchio et al. (2020). This data structure can occur in several and different application domains, such as longitudinal data on multiple response variables, spatial multivariate data, multivariate repeated measures, or spatio-temporal data. Nevertheless, there exists a limited number of contributions involving matrix-variate regression models. First introduced by Viroli (2012), these models are a natural generalization of the multivariate-multiple regression (see also Anderlucci et al. 2014). Within the mixture framework, finite mixtures of matrix-variate regressions (MN-FMR) have been recently considered by Melnykov and Zhu (2019). There are no matrix-variate CWMs in the literature and this paper aims to fill this gap by introducing and discussing a matrix-variate CWM in which the cluster-specific marginal distribution of the covariates and the cluster-specific conditional distribution of the responses given the covariates are assumed to be matrix normal.

The remainder of the paper is organized as follows. In Section 2, some preliminary aspects are described. The matrix normal cluster-weighted model (MN-CWM) and the expectation-conditional maximization (ECM) algorithm (Meng and Rubin 1993) for parameter estimation are discussed in Section 3. Computational and operational aspects are laid out in Section 4. In the simulation study outlined in Section 5, the parameter recovery and the classification performance of the MN-CWM are investigated as well as the capability of the Bayesian information criterion (BIC; Schwarz and et al. 1978) to detect the underlying group structure. The MN-CWM is also therein compared to the MN-FMR and to the multivariate-multiple normal CWM (MMN-CWM). The application of the MN-CWM to two real datasets concerning educational indicators and the Italian non-life insurance market is therefore analyzed in Section 6, whereas some conclusions and ideas for future developments are drawn in Section 7.

2 Background

2.1 Matrix Normal Distribution

A $p \times r$ continuous random matrix $Y \in \mathbb{R}^{p \times r}$ has a matrix normal (MN) distribution, denoted by $\mathcal{N}_{p \times r}(\mathbf{M}, \Phi, \Psi)$, if its density is

$$\phi_{p \times r}(Y; \mathbf{M}, \Phi, \Psi) = (2\pi)^{-\frac{pr}{2}} |\Phi|^{-\frac{r}{2}} |\Psi|^{-\frac{p}{2}} \exp \left\{ -\frac{1}{2} \text{tr} \left[\Phi^{-1} (Y - \mathbf{M}) \Psi^{-1} (Y - \mathbf{M})^\top \right] \right\}, \quad (1)$$

where \mathbf{M} is the $p \times r$ mean matrix, and Φ and Ψ are the $p \times p$ and $r \times r$ covariance matrices associated with the p rows and r columns, respectively. An equivalent definition specifies the $(p \times r)$ -matrix normal distribution as a special case of the pr -multivariate normal distribution. Specifically,

$$Y \sim \mathcal{N}_{p \times r}(\mathbf{M}, \Phi, \Psi) \Leftrightarrow \text{vec}(Y) \sim \mathcal{N}_{pr}(\text{vec}(\mathbf{M}), \Psi \otimes \Phi), \quad (2)$$

where $\mathcal{N}_{pr}(\cdot)$ denotes the pr -variate normal distribution, $\text{vec}(\cdot)$ is the vectorization operator, and \otimes denotes the Kronecker product. However, a MN distribution has the desirable feature of simultaneously modeling and identify the between- and within-variables variabilities as well as reducing the number of free parameters from pr ($pr + 1$)/2 to $[r(r + 1)/2] + [p(p + 1)/2] - 1$.

2.2 Matrix-Variate Regression Model

Let $Y \in \mathbb{R}^{p \times r}$ be a continuous random matrix of dimension $p \times r$, containing p responses measured in r occasions. Suppose we observe a set of q covariates for each occasion, inserted in a matrix X of dimension $q \times r$. A generic matrix-variate regression model for Y has the form

$$Y = \beta \mathbf{w}^\top + \mathbf{B}X + U, \quad (3)$$

where β is the $p \times 1$ vector consisting in the parameters related with the intercept, \mathbf{w} is a $r \times 1$ vector of ones, \mathbf{B} is the $p \times q$ matrix containing the parameters related to the q covariates, and U is the $p \times r$ error term matrix. Model (3) can be expressed in compact notation as

$$Y = \mathbf{B}^* X^* + U, \quad (4)$$

where \mathbf{B}^* is the $p \times (1 + q)$ matrix involving all the parameters to be estimated and X^* is the $(1 + q) \times r$ matrix containing the information about the intercept and q covariates (Viroli 2012). If we assume $U \sim \mathcal{N}_{p \times r}(\mathbf{0}, \Phi, \Psi)$, then $Y|X^* \sim \mathcal{N}_{p \times r}(\mathbf{B}^* X^*, \Phi, \Psi)$. Therefore, a matrix-variate regression can be viewed as an encompassing framework containing as special cases the multivariate-multiple regression, when $r = 1$, and the univariate-multiple regression when $r = 1$ and $p = 1$.

3 Methodology

3.1 Matrix Normal CWM

Let (X, Y) be a pair of random matrices, defined as in Section 2.2, with joint distribution $p(X, Y)$. Then, a general matrix CWM has the following joint distribution:

$$p(X, Y) = \sum_{g=1}^G p_g(Y|X^*) p_g(X) \pi_g, \tag{5}$$

where $p_g(Y|X^*)$ is the cluster-specific conditional distribution of the responses, $p_g(X)$ is the cluster-specific marginal distribution of the covariates, and $\pi_g > 0$ is the mixing weight (with $\sum_{g=1}^G \pi_g = 1$). Furthermore, we assume that in each group the conditional expectation $\mathbb{E}(Y|X^*)$ is a linear function of X^* depending on some parameters.

In this paper, we focus on model (5) by assuming that both $p_g(Y|X^*)$ and $p_g(X)$ are matrix normal densities, and $\mathbb{E}(Y|X^*) = B^* X^*$, as described in Section 2.2. Thus, model (5) can be rewritten as

$$p(X, Y; \Theta) = \sum_{g=1}^G \phi_{p \times r}(Y|X^*; B_g^* X^*, \Phi_{Y_g}, \Psi_{Y_g}) \phi_{q \times r}(X; M_g, \Phi_{X_g}, \Psi_{X_g}) \pi_g, \tag{6}$$

where Θ denotes the set of all parameters. For ease of understanding, a subscript with the variable name is added to the respective covariance matrices. Notice that there is an identifiability issue concerning the covariance matrices. Indeed, for a MN distribution,

$$\Psi \otimes \Phi = \Psi^* \otimes \Phi^*$$

if $\Phi^* = a\Phi$ and $\Psi^* = a^{-1}\Psi$. As a result, Φ and Ψ are identifiable up to a multiplicative constant a (Melnykov and Zhu 2019). According to Gallagher and McNicholas (2017, 2018) a way to obtain a unique solution is to fix the first diagonal element of the row covariance matrix to 1. Therefore, we adopt this approach in model (6) by setting the first diagonal element of Φ_{Y_g} and Φ_{X_g} to 1.

If the MN-CWM was not available, a possible approach would be to vectorize the matrices and consider the MMN-CWM, of which the MN-CWM is a special case; see Eq. 2. However, such a procedure leads to two principal concerns. The first one is the overparameterization of the vectorized model. Secondly, this increased number of free parameters may directly impact model selection, as will be better analyzed in Section 5.3.

3.2 Parameter Estimation

Parameter estimation is carried out via the ECM algorithm, which differs from the expectation-maximization (EM) algorithm (Dempster et al. 1977) in that the M-step is replaced by a sequence of simpler and computationally convenient CM-steps. The EM algorithm cannot be directly implemented because there is no closed form solution for the covariance matrices of the MN distribution, i.e., one of the two depends on the value of the other at the previous iteration (Dutilleul 1999).

Let $\mathcal{S} = \{(X_i, Y_i)\}_{i=1}^N$ be a sample of N independent observations from model (6). Then, the incomplete-data likelihood function is

$$L(\Theta|\mathcal{S}) = \prod_{i=1}^N p(X_i, Y_i; \Theta) = \prod_{i=1}^N \left[\sum_{g=1}^G \phi_{p \times r}(Y_i | X_i^*; B_g^* X_i^*, \Phi_{Y_g}, \Psi_{Y_g}) \phi_{q \times r}(X_i; M_g, \Phi_{X_g}, \Psi_{X_g}) \pi_g \right]. \tag{7}$$

Within the formulation of mixture models, \mathcal{S} is viewed as being incomplete because, for each observation, we do not know its component membership. Let $z_i = (z_{i1}, \dots, z_{iG})^\top$ be the component membership vector such that $z_{ig} = 1$ if (X_i, Y_i) comes from group (i.e., component) g and $z_{ig} = 0$ otherwise. Now, the complete-data are $\mathcal{S}_c = \{(X_i, Y_i, z_i)\}_{i=1}^N$, and the complete-data likelihood is

$$L_c(\Theta|\mathcal{S}_c) = \prod_{i=1}^N \prod_{g=1}^G \left[\phi_{p \times r}(Y_i | X_i^*; B_g^* X_i^*, \Phi_{Y_g}, \Psi_{Y_g}) \phi_{q \times r}(X_i; M_g, \Phi_{X_g}, \Psi_{X_g}) \pi_g \right]^{z_{ig}}. \tag{8}$$

Therefore, the corresponding complete-data log-likelihood can be written as

$$l_c(\Theta|\mathcal{S}_c) = \sum_{i=1}^N \sum_{g=1}^G z_{ig} \ln(\pi_g) + \sum_{i=1}^N \sum_{g=1}^G z_{ig} \ln \left[\phi_{p \times r}(Y_i | X_i^*; B_g^* X_i^*, \Phi_{Y_g}, \Psi_{Y_g}) \right] + \sum_{i=1}^N \sum_{g=1}^G z_{ig} \ln \left[\phi_{q \times r}(X_i; M_g, \Phi_{X_g}, \Psi_{X_g}) \right]. \tag{9}$$

In the following, by adopting the notation used in Tomarchio et al. (2020), the quantities marked with one dot correspond to the updates at the previous iteration and those marked with two dots represent the updates at the current iteration.

E-Step The E-step requires calculation of the conditional expectation of Eq. 9, given the observed data and the current estimate of the parameters $\hat{\Theta}$. To do this, we need to calculate

$$\begin{aligned} \ddot{z}_{ig} &= \mathbb{E}_{\hat{\Theta}} [Z_{ig} | X_i, Y_i] \\ &= \frac{\dot{\pi}_g \phi_{p \times r}(Y_i | X_i^*; \dot{B}_g^* X_i^*, \dot{\Phi}_{Y_g}, \dot{\Psi}_{Y_g}) \phi_{q \times r}(X_i; \dot{M}_g, \dot{\Phi}_{X_g}, \dot{\Psi}_{X_g})}{\sum_{j=1}^G \dot{\pi}_j \phi_{p \times r}(Y_i | X_i^*; \dot{B}_j^* X_i^*, \dot{\Phi}_{Y_j}, \dot{\Psi}_{Y_j}) \phi_{q \times r}(X_i; \dot{M}_j, \dot{\Phi}_{X_j}, \dot{\Psi}_{X_j})}, \end{aligned} \tag{10}$$

which corresponds to the posterior probability that the unlabeled observation (X_i, Y_i) belongs to the g th component of the mixture.

CM-Step 1 At the first CM-step, we maximize the expectation of the complete-data log-likelihood with respect to $\Theta_1 = \{\pi_g, M_g, \Phi_{X_g}, B_g^*, \Phi_{Y_g}\}_{g=1}^G$, fixing $\Theta_2 =$

$\{\Psi_{X_g}, \Psi_{Y_g}\}_{g=1}^G$ at $\hat{\Theta}_2$. In particular, we obtain

$$\ddot{\pi}_g = \frac{1}{N} \sum_{i=1}^N \ddot{z}_{ig}, \quad \ddot{M}_g = \frac{1}{\sum_{i=1}^N \ddot{z}_{ig}} \sum_{i=1}^N \ddot{z}_{ig} X_i, \tag{11}$$

$$\ddot{B}_g^* = \left[\sum_{i=1}^N \ddot{z}_{ig} Y_i (\dot{\Psi}_{Y_g})^{-1} X_i^{*\top} \right] \left[\sum_{i=1}^N \ddot{z}_{ig} X_i^* (\dot{\Psi}_{Y_g})^{-1} X_i^{*\top} \right]^{-1}, \tag{12}$$

$$\ddot{\Phi}_{X_g} = \frac{1}{r \sum_{i=1}^N \ddot{z}_{ig}} \sum_{i=1}^N \ddot{z}_{ig} (X_i - \ddot{M}_g) (\dot{\Psi}_{X_g})^{-1} (X_i - \ddot{M}_g)^\top, \tag{13}$$

$$\ddot{\Phi}_{Y_g} = \frac{1}{r \sum_{i=1}^N \ddot{z}_{ig}} \sum_{i=1}^N \ddot{z}_{ig} (Y_i - \ddot{B}_g^* X_i^*) (\dot{\Psi}_{Y_g})^{-1} (Y_i - \ddot{B}_g^* X_i^*)^\top. \tag{14}$$

CM-Step 2 At the second CM-step, we maximize the expectation of the complete-data log-likelihood with respect to Θ_2 , keeping fixed Θ_1 at $\hat{\Theta}_1$. Therefore, we have

$$\ddot{\Psi}_{X_g} = \frac{1}{q \sum_{i=1}^N \ddot{z}_{ig}} \sum_{i=1}^N \ddot{z}_{ig} (X_i - \ddot{M}_g)^\top (\ddot{\Phi}_{X_g})^{-1} (X_i - \ddot{M}_g), \tag{15}$$

$$\ddot{\Psi}_{Y_g} = \frac{1}{p \sum_{i=1}^N \ddot{z}_{ig}} \sum_{i=1}^N \ddot{z}_{ig} (Y_i - \ddot{B}_g^* X_i^*)^\top (\ddot{\Phi}_{Y_g})^{-1} (Y_i - \ddot{B}_g^* X_i^*). \tag{16}$$

4 Computational and Operational Aspects

The code for the ECM algorithm previously described, and for the analysis in the following sections, is written within the R computing environment (R Core Team 2018).

4.1 ECM Initialization

The choice of the starting values is an important aspect for EM-based algorithms (see, e.g., Maitra and Melnykov 2010; Michael and Melnykov 2016). A common way to start an EM-based algorithm consists in providing initial values for the z_{ig} in Eq. 10, during the first E-step of the algorithm (McLachlan and Peel 2000). In our case, we also need to provide initial values for Ψ_{X_g} and Ψ_{Y_g} in Eqs. 13 and 14, respectively, during the first CM-step of the algorithm. Therefore, the following initialization strategy is implemented:

1. Generate G random positive definite matrices for both Ψ_{X_g} and Ψ_{Y_g} . This is done via the `genPositiveDefMat()` function of the **clusterGeneration** package, by using the “eigen” method. For further details, see Qiu and Joe (2015).
2. Generate N random vectors $z_i = (z_{i1}, \dots, z_{iG})^\top, i = 1, \dots, N$. This is done by using the following three approaches:
 - 2.1 In a “soft” way, by generating G positive random values from a uniform distribution on $[0,1]$ for each observation, that are subsequently normalized to have a unitary sum. Being purely random, this procedure is repeated 15 times, and

- the solution maximizing the observed-data log-likelihood among these runs is considered;
- 2.2 In a “hard” way, by using the classification produced by the k -means algorithm on the vectorized and merged data. Specifically, after computing $\{\text{vec}(\mathbf{X}_i)\}_{i=1}^N$ and $\{\text{vec}(\mathbf{Y}_i)\}_{i=1}^N$, the data are merged so that for each observation we have a vector of dimension $(pr + qr) \times 1$;
 - 2.3 In a “hard” way, by using the classification produced by mixtures of matrix-normal distributions, computed on the merged data. In detail, for each observation, we have a $(p + q) \times r$ matrix.

The approach producing the largest (observed) log-likelihood is finally selected.

4.2 Spurious Clusters

A well-known issue in mixture models is related to the possibility for EM-based algorithms to converge to spurious local maxima. This is not a failing of these algorithms, and such solutions reflect a random pattern in the data rather than an underlying group structure. They can usually be detected by the presence of data groups with very few observations or small eigenvalues compared to the others. Consequently, these solutions often have a high likelihood, but are of little practical use or real-world interpretation (McLachlan and Peel 2000). Various approaches to mitigate this problem have been proposed in the literature. For example, Leisch (2004) and Dang and McNicholas (2015) impose a minimum size for the clusters by forcing $\pi_g \geq 0.05$. Melnykov and Zhu (2019) tackle this issue by excluding all the solutions that involved clusters consisting of fewer than five points. Herein, we deal with this problem by removing the solutions that included clusters with estimated $\pi_g \approx 0.05$ or less, and with eigenvalues close to 0, in one or more of its estimated covariance matrices.

4.3 Model Selection and Performance Evaluation

It is often the case that the number of groups G is not known in advance, and model selection criteria are commonly used for estimating it. Among them, the BIC is one of the most popular, and will be used in the following. Furthermore, when the true classification of the data is known, the adjusted Rand index (ARI; Hubert and Arabie 1985) can be used to evaluate the clustering performance of a model. Specifically, the ARI calculates the agreement between the true classification and the one predicted by the model. An ARI of 1 indicates perfect agreement between the two partitions, whereas the expected value of the ARI under a random classification is 0. The ARI will be used in the manuscript along with the misclassification rate η , which is the percentage of units misclassified.

5 Simulation Studies

5.1 Simulation 1: a Focus on the Matrix-Normal CWM

In this study, several aspects related to our model are analyzed. First of all, since the ECM algorithm is used to fit the model, it is desirable to evaluate its parameter recovery, i.e., whether it can recover the generating parameters accurately. For this reason, data are generated from a four-component MN-CWM with $p = q = r = 3$. Two scenarios are then evaluated, according to the different levels of overlap of the mixture components. In the first

scenario (labeled as “Scenario A₁”), the mixture components are well-separated both in X , by assuming relatively distant mean matrices, and in $Y|X^*$, by using different intercepts and slopes. On the contrary, in the second scenario (labeled as “Scenario B₁”), there is a certain amount of overlap because the intercepts are all equal among the mixture components, while the slopes and the mean matrices assume approximately the same values among the mixture components. The parameters used for Scenario A₁ are displayed in Appendix 1. Under Scenario B₁, the set of parameters $\{\pi_g, \Phi_{X_g}, \Psi_{X_g}, \Phi_{Y_g}, \Psi_{Y_g}\}_{g=1}^4$, M_1 and the slopes in B_1^* and B_3^* , are the same of Scenario A₁. The other mean matrices are obtained by adding a constant c to each element of the corresponding mean matrices used for Scenario A₁. In detail, c is equal to -5 , 5 , and -10 for M_2 , M_3 , and M_4 , respectively. The intercept column of all the mixture components is equal to $(7, 2, 5)^\top$, whereas the slopes in B_2^* and B_4^* are all multiplied by -1 , with respect to those used in Scenario A₁. Lastly, two sample sizes are considered within each scenario, i.e., $N = 200$ and $N = 500$.

On all the generated datasets, the MN-CWM is fitted directly with $G = 4$, and the bias and mean squared error (MSE) of the parameter estimates are computed. For brevity’s sake, and as also supported by the existing CWM literature (see, e.g., Punzo 2014; Ingrassia et al. 2015; Punzo and Ingrassia 2016; Punzo and McNicholas 2017), attention will be focused only on the regression coefficients B_1^*, \dots, B_G^* . Before showing the obtained results, it is important to underline the well-known label switching issue, caused by the invariance of a mixture distribution to relabeling the components (Frühwirth-Schnatter 2006). There are no generally accepted labeling methods. Herein, to assign the correct labels, an analysis of the overall estimated parameters is conducted on each generated dataset to properly identify each mixture component.

Table 1 summarizes the estimated bias and MSE of the parameter estimates for Scenario A₁, over one hundred replications for each sample size N , after fitting the MN-CWM with $G = 4$. The same is reported for Scenario B₁ in Table 2. The first and most immediate result is that the biases and the MSEs take very small values in both scenarios. This is particularly relevant for Scenario B₁ because of the presence of overlap between the mixture components. Furthermore, within each scenario, an increase in the sample size leads to a rough improvement of the parameter estimates, whereas it systematically reduces the MSE.

Other aspects that are investigated consist in the evaluation of the classification produced by our model, as well as the capability of the BIC in identifying the correct number of groups in the data. For this reason, under each of the considered scenarios, the MN-CWM is fitted to the generated datasets for $G \in \{1, 2, 3, 4, 5\}$, and the results are reported in Table 3.

It is easy to see that, under Scenario A₁, a perfect classification is always obtained regardless of the sample size. Additionally, the BIC regularly detects the correct number of groups in the data. Under Scenario B₁, because of the larger overlap, the $\overline{\text{ARI}}$ assumes lower, but in any case good, values. Relatedly, the percentage of misclassified units stands at around the 3% for both sample sizes. Concerning the BIC, also in this case, it properly identifies the underlying group structure, with only one exception (when $N = 200$).

A final aspect that is evaluated in this study concerns the initialization strategy. Specifically, Table 4 displays the number of times each strategy for the z_i produces the highest log-likelihood at convergence, within each scenario and for both sample sizes. The initial G random matrices for Ψ_{X_g} and Ψ_{Y_g} are assumed to be the same.

The first result suggests the importance of considering multiple initialization strategies because none of them is preferred in all the generated datasets. However, the random strategy is quite close to this target because it only fails in three datasets under Scenario B₁. Very similar performances are obtained when the mixture initialization is used. On the contrary,

Table 1 Estimated bias and MSE of the regression coefficients $\{B_g^*\}_{g=1}^4$, over 100 replications, under Scenario A₁

		$N = 200$	$N = 500$
Group 1	Bias	$\begin{pmatrix} 0.032 & -0.001 & 0.005 & 0.002 \\ -0.025 & -0.010 & -0.008 & -0.004 \\ -0.028 & 0.006 & -0.014 & -0.004 \end{pmatrix}$	$\begin{pmatrix} 0.001 & -0.003 & -0.002 & 0.007 \\ -0.033 & 0.002 & -0.007 & 0.004 \\ 0.003 & -0.004 & -0.002 & 0.005 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.343 & 0.011 & 0.018 & 0.014 \\ 0.337 & 0.010 & 0.017 & 0.018 \\ 0.365 & 0.011 & 0.020 & 0.016 \end{pmatrix}$	$\begin{pmatrix} 0.111 & 0.004 & 0.006 & 0.007 \\ 0.114 & 0.004 & 0.006 & 0.006 \\ 0.104 & 0.004 & 0.006 & 0.005 \end{pmatrix}$
Group 2	Bias	$\begin{pmatrix} 0.039 & -0.004 & 0.001 & -0.001 \\ -0.005 & -0.004 & 0.002 & 0.006 \\ -0.004 & -0.008 & -0.004 & 0.009 \end{pmatrix}$	$\begin{pmatrix} 0.001 & -0.003 & -0.002 & -0.001 \\ 0.019 & -0.000 & -0.000 & -0.000 \\ 0.042 & -0.004 & -0.002 & -0.001 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.252 & 0.003 & 0.003 & 0.004 \\ 0.204 & 0.002 & 0.003 & 0.004 \\ 0.170 & 0.002 & 0.003 & 0.005 \end{pmatrix}$	$\begin{pmatrix} 0.084 & 0.001 & 0.001 & 0.001 \\ 0.089 & 0.001 & 0.001 & 0.001 \\ 0.099 & 0.001 & 0.001 & 0.001 \end{pmatrix}$
Group 3	Bias	$\begin{pmatrix} 0.005 & -0.008 & 0.005 & 0.001 \\ -0.020 & -0.010 & -0.002 & -0.000 \\ -0.051 & -0.008 & -0.003 & -0.001 \end{pmatrix}$	$\begin{pmatrix} 0.002 & -0.000 & -0.001 & -0.000 \\ 0.055 & 0.001 & 0.004 & 0.001 \\ 0.018 & 0.003 & 0.002 & -0.001 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.229 & 0.005 & 0.002 & 0.003 \\ 0.244 & 0.005 & 0.002 & 0.003 \\ 0.235 & 0.006 & 0.002 & 0.004 \end{pmatrix}$	$\begin{pmatrix} 0.104 & 0.002 & 0.001 & 0.001 \\ 0.122 & 0.002 & 0.001 & 0.001 \\ 0.111 & 0.002 & 0.001 & 0.001 \end{pmatrix}$
Group 4	Bias	$\begin{pmatrix} 0.097 & -0.008 & 0.011 & -0.005 \\ 0.027 & -0.006 & 0.006 & 0.002 \\ -0.017 & -0.005 & 0.006 & 0.005 \end{pmatrix}$	$\begin{pmatrix} -0.041 & 0.003 & 0.001 & -0.002 \\ -0.045 & 0.003 & 0.005 & -0.002 \\ -0.006 & 0.002 & 0.003 & -0.004 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.412 & 0.003 & 0.005 & 0.004 \\ 0.412 & 0.003 & 0.007 & 0.004 \\ 0.397 & 0.003 & 0.006 & 0.004 \end{pmatrix}$	$\begin{pmatrix} 0.242 & 0.001 & 0.001 & 0.001 \\ 0.200 & 0.001 & 0.001 & 0.001 \\ 0.209 & 0.001 & 0.002 & 0.001 \end{pmatrix}$

the k -means strategy provides the worst performances, even if it produces the best solution in approximately the 80% of the datasets.

5.2 Simulation 2: a Comparison Between the Matrix-Normal CWM and the Matrix-Normal FMR

In this study, the matrix-normal CWM is compared to the matrix-normal FMR. Specifically, three scenarios with $N = 200$, $p = 2$, $q = 3$, and $r = 4$ are considered, and in each of them 30 datasets from a matrix-normal CWM with $G = 2$ are generated. The first scenario (hereafter simply referred to as “Scenario A₂”) is characterized by the fact that the two groups differ only for the intercepts and the covariance matrices. This implies that they have totally overlapping mean matrices, which should make the distribution of the covariates $p_g(X)$ not very important for clustering. The parameters used to generate the datasets are displayed in Appendix 2. In the second scenario (“Scenario B₂”), the two groups have the same B_g^* and π_g . The parameters used to generate the datasets are the same as for Scenario A₂, but with only two differences: a value $c = 5$ is added to each element of M_2 and we

Table 2 Estimated bias and MSE of the regression coefficients $\{B_g^*\}_{g=1}^G$, over 100 replications, under Scenario B₁

		$N = 200$	$N = 500$
Group 1	Bias	$\begin{pmatrix} -0.058 & -0.011 & -0.015 & 0.015 \\ -0.037 & -0.008 & -0.011 & 0.018 \\ -0.082 & -0.002 & -0.027 & 0.010 \end{pmatrix}$	$\begin{pmatrix} 0.052 & -0.001 & 0.008 & -0.002 \\ 0.034 & -0.001 & 0.004 & 0.002 \\ -0.018 & 0.000 & -0.006 & 0.001 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.368 & 0.014 & 0.018 & 0.021 \\ 0.410 & 0.014 & 0.021 & 0.022 \\ 0.361 & 0.011 & 0.019 & 0.022 \end{pmatrix}$	$\begin{pmatrix} 0.118 & 0.004 & 0.007 & 0.006 \\ 0.117 & 0.004 & 0.007 & 0.007 \\ 0.124 & 0.004 & 0.006 & 0.007 \end{pmatrix}$
Group 2	Bias	$\begin{pmatrix} -0.046 & 0.002 & -0.013 & -0.001 \\ -0.037 & 0.005 & 0.001 & 0.001 \\ -0.013 & 0.008 & 0.004 & 0.005 \end{pmatrix}$	$\begin{pmatrix} -0.014 & 0.006 & -0.002 & 0.005 \\ -0.030 & 0.004 & -0.006 & 0.008 \\ -0.008 & 0.000 & -0.002 & 0.009 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.046 & 0.003 & 0.006 & 0.004 \\ 0.046 & 0.004 & 0.003 & 0.005 \\ 0.043 & 0.004 & 0.004 & 0.005 \end{pmatrix}$	$\begin{pmatrix} 0.015 & 0.001 & 0.001 & 0.002 \\ 0.013 & 0.001 & 0.001 & 0.001 \\ 0.013 & 0.001 & 0.001 & 0.002 \end{pmatrix}$
Group 3	Bias	$\begin{pmatrix} 0.035 & -0.017 & 0.011 & 0.016 \\ 0.011 & -0.006 & 0.012 & 0.008 \\ 0.023 & -0.010 & 0.005 & 0.010 \end{pmatrix}$	$\begin{pmatrix} 0.007 & -0.004 & 0.002 & -0.000 \\ 0.015 & -0.004 & 0.001 & 0.000 \\ 0.028 & -0.005 & 0.004 & -0.000 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.078 & 0.006 & 0.003 & 0.003 \\ 0.073 & 0.005 & 0.003 & 0.003 \\ 0.080 & 0.005 & 0.002 & 0.003 \end{pmatrix}$	$\begin{pmatrix} 0.027 & 0.002 & 0.001 & 0.001 \\ 0.025 & 0.002 & 0.001 & 0.001 \\ 0.030 & 0.002 & 0.001 & 0.001 \end{pmatrix}$
Group 4	Bias	$\begin{pmatrix} 0.039 & 0.002 & 0.002 & -0.003 \\ 0.005 & -0.001 & -0.004 & -0.007 \\ -0.051 & -0.003 & 0.004 & 0.004 \end{pmatrix}$	$\begin{pmatrix} -0.043 & 0.004 & 0.001 & -0.003 \\ -0.014 & -0.000 & 0.003 & 0.002 \\ -0.008 & -0.002 & 0.002 & 0.003 \end{pmatrix}$
	MSE	$\begin{pmatrix} 0.147 & 0.003 & 0.005 & 0.004 \\ 0.160 & 0.006 & 0.007 & 0.006 \\ 0.132 & 0.003 & 0.007 & 0.006 \end{pmatrix}$	$\begin{pmatrix} 0.061 & 0.001 & 0.002 & 0.002 \\ 0.060 & 0.001 & 0.002 & 0.002 \\ 0.069 & 0.001 & 0.002 & 0.002 \end{pmatrix}$

Table 3 $\overline{\text{ARI}}$ and $\overline{\eta}$ values, along with the number of times the true G is selected by the BIC, over 100 replications, for scenarios A₁ and B₁

	$N = 200$			$N = 500$		
	$\overline{\text{ARI}}$	$\overline{\eta}$	True G	$\overline{\text{ARI}}$	$\overline{\eta}$	True G
Scenario A ₁	1.00	0.00%	100	1.00	0.00%	100
Scenario B ₁	0.91	3.04%	99	0.92	2.71%	100

Table 4 Number of times, over 100 replications, the considered initialization strategies produced the highest log-likelihood at convergence

	$N = 200$, for scenarios A ₁ and B ₁			$N = 500$		
	Random	k -Means	Mixture	Random	k -Means	Mixture
Scenario A ₁	100	77	98	100	79	95
Scenario B ₁	97	74	87	100	83	100

set $B_2^* = B_1^*$. Lastly, in the third scenario (“Scenario C₂”), the two groups have only the same slopes and π_g . Here, with respect to the parameters used under Scenario B₂, the only difference is in the intercept vectors, which are $(-3, -4)^\top$ and $(-7, -8)^\top$ for the first and the second groups, respectively.

The MN-CWM and the MN-FMR are then fitted to the datasets of each scenario for $G \in \{1, 2, 3\}$, and the results in terms of model selection and clustering are reported in Table 5. It is possible to see that, in Scenario A₂, the BIC correctly selects two groups for both models and the classifications produced are perfect. Therefore, even if the two groups have the same means and are strongly overlapping, the MN-CWM seems able to properly identify the true underlying group structure. However, under such a scenario, the MN-FMR should be preferred because the distribution of the covariates $p_g(\mathbf{X})$ is not useful for clustering and it is more parsimonious than the MN-CWM. On the contrary, Scenarios B₂ and C₂ represent typical examples of the usefulness of $p_g(\mathbf{X})$. Specifically, the BIC always identifies just one group under both scenarios for the MN-FMR, with obvious consequences in terms of the classification produced. Notice that, even if the MN-FMR had been fitted directly with $G = 2$, the resulting classifications would lead to almost identical $\overline{\text{ARI}}$ and $\overline{\eta}$ for Scenario B₂, and slightly better performance for Scenario C₂, because $\overline{\text{ARI}} = 0.15$ and $\overline{\eta} = 32.48\%$. This underlines how, regardless of the BIC, the MN-FMR is not able to properly model such data structures.

5.3 Simulation 3: a Comparison Between the Matrix-Normal CWM and the Multivariate-Multiple Normal CWM

In this study, the MN-CWM is compared to the MMN-CWM. To show the effects of data vectorization, we consider two experimental factors: the matrix dimensionality and the number of groups G . About the dimensionality, we assume square matrices having the same dimensions both for the responses and the covariates, i.e., $p = q = r \in \{2, 3, 4\}$. Similarly, situations with three different number of groups are evaluated, i.e., $G \in \{2, 3, 4\}$. By combining both experimental factors, nine scenarios are obtained and, for each of them, 30 datasets are generated from a MN-CWM. The parameters used to generate the data come from Section 5.1 and are shown in Appendix 1. In detail, when $p = q = r = 2$ they are obtained by taking the submatrix in the upper-left corner of each parameter, when $p = r = q = 3$ they are exactly as displayed, whereas when $p = r = q = 4$ a row and a column on each parameter matrix are added, which for brevity’s sake are not reported here. About the number of groups, and by considering Appendix 1, when $G = 2$ and $G = 3$ the first two and three groups are selected, respectively, while when $G = 4$ all of them are considered.

Table 5 $\overline{\text{ARI}}$ and $\overline{\eta}$ values, along with the number of times the true G is selected by the BIC, over 30 replications, for scenarios A₂, B₂, and C₂

	MN-CWM			MN-FMR		
	$\overline{\text{ARI}}$	$\overline{\eta}$	True G	$\overline{\text{ARI}}$	$\overline{\eta}$	True G
Scenario A ₂	1.00	0.00%	100	1.00	0.00%	100
Scenario B ₂	0.99	0.03%	100	0.00	47.22%	0
Scenario C ₂	1.00	0.01%	100	0.00	47.18%	0

The MN-CWM is then fitted to each dataset for $G \in \{1, 2, 3, 4, 5\}$. The same is done for the MMN-CWM after data vectorization, and the results of both models in terms of model selection via the BIC are shown in Table 6.

As we can see from Table 6, when the MN-CWM is considered, regardless of the data dimensionality and the number of groups, the BIC always selects the correct number of groups. The same also holds for the MMN-CWM when $p = q = r = 2$ or, regardless of the data dimensionality, when $G = 2$. However, when $p = q = r = 3$, the BIC starts to face issues for $G = 3$ because the true number of groups is detected only 11 times (the other 19 times, $G = 2$ is selected) and it systematically fails when $G = 4$. This problem gets even worse when $p = r = q = 4$ (with the exclusion of $G = 2$). The reason for such failures is related to the increased number of parameters with respect to the MN-CWM. Therefore, on the one hand, we have a model that can become seriously overparameterized with negative effects also on model selection (the MMN-CWM) and, on the other hand, we have a model (the MN-CWM) which is able to fit the same data in a far more parsimonious way and without causing problems on model selection.

6 Real Data Applications

6.1 Education Data

The first dataset comes from the Italian national agency for the evaluation of universities and research institutes, which makes available to Italian universities quantitative indicators related to the academic careers of the students and the results of the training activities. For this application, the following $p = 2$ responses that measure the level of completion of studies by students are considered:

1. The percentage of students that graduate within $T + 1$ years (Complete), and
2. The percentage of students that drop after $T + 1$ years (Drop),

where T is the normal duration of the study program. Moreover, the following $q = 2$ covariates that may be helpful in explaining this progress are taken into account:

1. The percentage of course credits earned in the first year over the total to be achieved (Credits), and
2. The percentage of students that have earned at least 40 course credits during the solar year (Students).

For the sake of simplicity, these variables will hereafter be referred to by the names given in the associated parentheses. All the measurements refer to $N = 75$ study programs in the

Table 6 The number of times, over 30 replications, the true G is selected by the BIC in each of the nine scenarios, for the MN-CWM and MMN-CWM

	MN-CWM			MMN-CWM		
	$G = 2$	$G = 3$	$G = 4$	$G = 2$	$G = 3$	$G = 4$
$p = q = r = 2$	30	30	30	30	30	30
$p = r = q = 3$	30	30	30	30	11	0
$p = r = q = 4$	30	30	30	30	0	0

non-telematic Italian universities over $r = 3$ years. Each study program is measured at the national level, i.e., it is the average value of all the study programs of the same type across the country, for the reference period.

There are two groups in the data, namely $N_1 = 33$ bachelor's degrees and $N_2 = 42$ master's degrees. The MN-CWM and the MN-FMR are fitted to the data for $G \in \{1, 2, 3\}$ and their results are reported in Table 7. The BIC selects a two-component MN-CWM that yields a perfect classification of the data. On the contrary, a three-component MN-FMR is chosen by the BIC and has 6.67% misclassified units (ARI = 0.88). Therefore, our model is able to completely recognize the underlying group structure, differently from the MN-FMR. Notice that, even if we consider the MN-FMR directly with $G = 2$, this will produce a similar classification with an ARI of 0.89. Additionally, if we vectorize the data and fit the MMN-CWM for $G \in \{1, 2, 3\}$, then the BIC chooses a one-component model. This is probably due to the overparameterization issue discussed in Section 5.3. In any case, if we would ignore this issue and fit the MMN-CWM directly with $G = 2$, the ARI would be 0.84, which is lower than those of matrix-variate models.

The estimated regression coefficients by the MN-CWM are displayed in Table 8, for the two groups, respectively. As it is plausible to expect, both covariates have a positive effect on the successful completion of the study programs, even if their magnitude is different in the two groups. The credits obtained during the first year of study might be more important for bachelor's students, considering the difficulties that arise in the transition from high school to university (Krause 2006). At the same time, obtaining at least 40 course credits per year should be easier for masters' students, resulting in a greater importance for the completion of the studies. Conversely, both covariates have a negative impact on the drop rates, with the exception of the Credits variable that surprisingly turns out to have a positive sign for the master's courses.

6.2 Insurance Data

For this second real data application, the “Insurance” dataset included in the **splm** package (Millo and Piras 2012) is used. This dataset was introduced by Millo and Carmeci (2011) to study the consumption of non-life insurance across the $N = 103$ Italian provinces in the years 1998–2002 ($r = 5$).

In this application we select, as responses, the following $p = 2$ variables that are related to the consumption and the presence of insurance products in the market:

1. Real per-capita non-life premiums in 2000 Euros (PPCD), and
2. Density of insurance agencies per 1000 inhabitants (AGEN),

as well as the following $q = 3$ financial covariates:

1. Real per-capita GDP (RGDP),
2. Real per-capita bank deposits (BANK), and
3. Real interest rate on lending to families and small enterprises (RIRS).

Table 7 ARI and η values for the MN-CWM and MN-FMR selected by the BIC for the education data

Model	G	ARI	η
MN-CWM	2	1.00	0.00%
MN-FMR	3	0.88	6.67%

Table 8 Estimated regression coefficients for the MN-CWM for the education data

	Intercept	Credits	Students
(a) Bachelor's degrees			
Complete	-0.087	0.973	0.102
Drop	0.727	-0.557	-0.178
(b) Master's degrees			
Complete	0.309	0.101	0.587
Drop	0.049	0.044	-0.011

There are two reasons why we focused on this subset of covariates. Firstly, they are regularly used in the literature and their relevant effects on the consumption or development of insurance products have been widely discussed (see the references in Millo and Carmeci 2011, for further details). Indeed, they are commonly used as proxies for income and general level of economic activity (RGDP), stock of wealth (BANK), and opportunity cost of allocate funds in insurance policies (RIRS). Secondly, we wish to avoid excessive parametrization of the models. Notice that, for a better interpretation of the regression coefficients, the variables RGDP and BANK are divided by 1000; thus, thousands of euros are considered.

Differently from the previous application, we do not have a “ground truth” classification of the data. However, because we are using $p(X)$ in Eq. 6, the sampling distribution of each covariate could provide useful insights. They are reported in Fig. 1, for the full 5-year data, as done by Melnykov and Zhu (2019). The bimodality in all these histograms seems to suggest the presence of two groups in the data, validating the findings of Millo and Carmeci (2011). They underline the existence of two macro areas, namely Central-Northern Italy, characterized by an insurance penetration level relatively close to the European averages, and Southern Italy, where a general economic underdevelopment has long been standing as a fundamental social and political problem.

The MN-CWM and the MN-FMR models are hence fitted to the data for $G \in \{1, 2, 3\}$, and the BIC selects a two-component MN-CWM and a three-component MN-FMR, respectively. These two partitions are illustrated in Fig. 2 using the Italian political map. Specifically, the Italian regions are bordered in yellow (islands excluded), while the internal provinces are delimited with the black lines and colored according to the estimated group membership both for the MN-CWM and the MN-FMR.

Because we do not have a “ground truth” classification for these data, we cannot compute either ARI or η . Nevertheless, the partition produced by the MN-CWM seems in line with the findings of Millo and Carmeci (2011), with a clear separation between Central-Northern Italy and Southern-Insular Italy. Furthermore, with the exclusion of three cases, all the provinces belonging to the same region are clustered together. The only exceptions concern the following: the province of Rome (in the Lazio region), which due to its social-economic development is reasonably assigned to the Central-Northern Italy group; the province of Ascoli-Piceno (in the Marche region); and the province of Massa-Carrara (in the Toscana region). On the contrary, the three groups detected by the MN-FRM are not supported by the literature and are difficult to interpret, even because they put together provinces spanning all over the country without a straightforward and reasonable justification.

Notice that, also in this case, if we vectorize the data and fit the MMN-CWM for $G \in \{1, 2, 3\}$, then the BIC selects a one-component model. This is in contrast to all the evidence discussed above, and it is plausibly driven by the overparameterization that for this

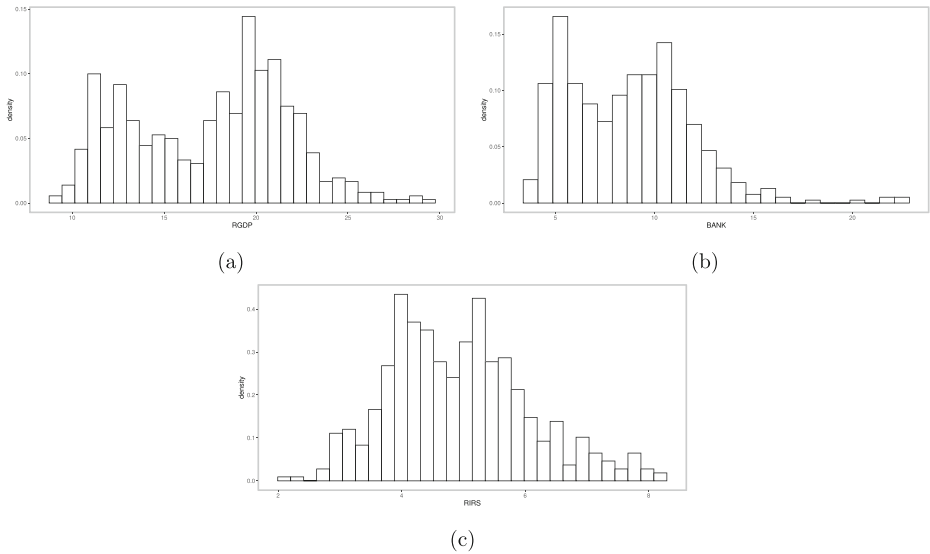


Fig. 1 Sampling distributions of the covariates for the insurance data

dataset is even higher than the previous application, given that the matrices have a greater dimensionality.

As in Section 6.1, the estimated regression coefficients by the MN-CWM are briefly presented in Table 9 for the two groups. Overall, the coefficients are quite different between the two groups, highlighting the divergences that characterize these two macro areas. The RIRS has a negative impact on the PPCD because the higher the interest rate on lending, the higher the opportunity cost of investing in insurance products. Conversely, an increase in the RIRS has a positive effect on the AGEN because, from the insurance companies’ point of

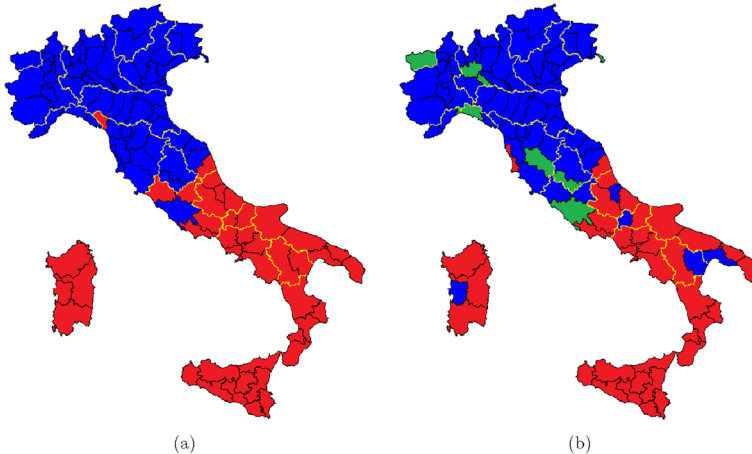


Fig. 2 Partitions produced by the MN-CWM (a) and MN-FMR (b) for the insurance data

Table 9 Estimated regression coefficients by the MN-CWM for the insurance data

	Intercept	RGDP	BANK	RIRS
Central-Northern Italy				
PPCD	85.7936	9.1932	1.8513	-7.3079
AGEN	0.5343	-0.0085	0.0071	0.0073
Southern Italy				
PPCD	-3.6968	6.0029	4.2062	-1.2885
AGEN	0.0307	0.0041	0.0279	0.0039

view, it raises the gains of investing the premiums on financial markets in the time between premium collection and claims settling. Therefore, these increased revenues might lead to an expansion of the insurance companies over the territory. Similarly, and in accordance with existing literature, income (RGDP) and wealth (BANK) have a positive influence on the PPCD (Millo and Carmeci 2011). It would be reasonable to expect that they also have a positive effect on AGEN. However, this is not the case at least for Central-Northern Italy that shows lower coefficients than Southern Italy. In the southern part of the country, where the insurance density and penetration are much lower, an increased income and wealth can give more space to growth and investment opportunities.

7 Conclusion

The matrix-variate CWM has been introduced. In the MN-CWM framework, the sets of responses and covariates may be simultaneously observed at different time points or locations. The matrix normal distribution was used for both the cluster-specific distributions of covariates and responses given covariates, and an ECM algorithm was used for parameter estimation. A simulation study analyzed parameter recovery and the classification performance of the proposed model as well as the capability of the BIC to detect the underlying group structure of the data. A comparison among the different initialization strategies has also been conducted. Additionally, the MN-CWM has been compared to the MN-FMR, discussing under which scenarios the cluster-specific distributions of covariates is useful, and to the MMN-CWM on the vectorized data, where the overparameterization issue and its consequences in model selection have been analyzed. Lastly, in the first real data application, the MN-CWM produced a perfect classification, differently from the MN-FMR model. Similarly, in the second real data analysis, our model seemed to provide a more reliable partition of the Italian provinces than the MN-FMR model.

Further model developments can be readily proposed. It is reasonable to assume that similar overparameterization concerns can affect also the MN-CWM when the dimensions of the matrices are quite high. To further increase the parsimony of our model, constrained parameterizations of the covariance matrices can be employed, by following the approaches of Sarkar et al. (2020), Subedi et al. (2013), Punzo and Ingrassia (2015), and Mazza et al. (2018). Furthermore, to accommodate skewness or mild outliers in the data, skewed or heavy tailed matrix-variate distributions could also be considered for the mixing components of the model (e.g., Melnykov and Zhu 2018; Gallagher and McNicholas 2018; Tomarchio et al. 2020).

Appendix 1

Table 10 Parameters used to generate the simulated datasets under Scenario A1

Parameter	Group 1	Group 2	Group 3	Group 4
π_g	0.30	0.30	0.20	0.20
M_g	$\begin{pmatrix} 1.00 & 2.00 & 0.00 \\ -4.00 & -3.00 & -3.00 \\ 1.00 & 2.00 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 6.00 & 8.00 & 6.00 \\ 2.00 & 1.00 & 3.00 \\ 5.00 & 6.00 & 6.00 \end{pmatrix}$	$\begin{pmatrix} -4.00 & -3.00 & -4.00 \\ -9.00 & -9.00 & -7.00 \\ -4.00 & -3.00 & -5.00 \end{pmatrix}$	$\begin{pmatrix} 12.00 & 12.00 & 11.00 \\ 6.00 & 7.00 & 7.00 \\ 10.00 & 11.00 & 11.00 \end{pmatrix}$
Φ_{X_g}	$\begin{pmatrix} 1.00 & 0.50 & 0.25 \\ 0.50 & 1.00 & 0.50 \\ 0.25 & 0.50 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 2.00 & 0.40 & 0.08 \\ 0.40 & 0.20 & 0.40 \\ 0.08 & 0.40 & 2.00 \end{pmatrix}$	$\begin{pmatrix} 1.50 & 0.75 & 0.38 \\ 0.75 & 1.50 & 0.75 \\ 0.38 & 0.75 & 1.50 \end{pmatrix}$	$\begin{pmatrix} 1.20 & 0.60 & 0.30 \\ 0.60 & 1.20 & 0.60 \\ 0.30 & 0.60 & 1.20 \end{pmatrix}$
Ψ_{X_g}	$\begin{pmatrix} 1.20 & 0.60 & 0.30 \\ 0.60 & 1.20 & 0.60 \\ 0.30 & 0.60 & 1.20 \end{pmatrix}$	$\begin{pmatrix} 1.40 & 0.70 & 0.35 \\ 0.70 & 1.40 & 0.70 \\ 0.35 & 0.70 & 1.40 \end{pmatrix}$	$\begin{pmatrix} 0.80 & 0.40 & 0.20 \\ 0.40 & 0.80 & 0.40 \\ 0.20 & 0.40 & 0.80 \end{pmatrix}$	$\begin{pmatrix} 1.60 & 0.80 & 0.40 \\ 0.80 & 1.60 & 0.80 \\ 0.40 & 0.80 & 1.60 \end{pmatrix}$
B_g^*	$\begin{pmatrix} 0.00 & 1.00 & 1.00 & 1.00 \\ -2.00 & 1.00 & 1.50 & 1.00 \\ 1.00 & 1.50 & 1.50 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 6.00 & -1.00 & -1.50 & -1.00 \\ 4.00 & -1.00 & -1.50 & -1.00 \\ 8.00 & -1.50 & -1.50 & -1.00 \end{pmatrix}$	$\begin{pmatrix} -5.00 & 1.00 & 1.00 & 1.00 \\ -3.00 & 1.50 & 1.00 & 1.00 \\ -6.00 & 1.50 & 1.50 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 1.00 & -1.00 & -1.00 & -1.00 \\ -5.00 & -1.00 & -1.50 & -1.50 \\ 0.00 & -1.50 & -1.00 & -1.50 \end{pmatrix}$
Φ_{Y_g}	$\begin{pmatrix} 1.40 & 0.84 & 0.50 \\ 0.84 & 1.40 & 0.84 \\ 0.50 & 0.84 & 1.40 \end{pmatrix}$	$\begin{pmatrix} 1.80 & 1.26 & 0.88 \\ 1.26 & 1.80 & 1.26 \\ 0.88 & 1.26 & 1.80 \end{pmatrix}$	$\begin{pmatrix} 1.20 & 0.84 & 0.59 \\ 0.84 & 1.20 & 0.84 \\ 0.59 & 0.84 & 1.20 \end{pmatrix}$	$\begin{pmatrix} 1.60 & 0.96 & 0.58 \\ 0.96 & 1.60 & 0.96 \\ 0.58 & 0.96 & 1.60 \end{pmatrix}$
Ψ_{Y_g}	$\begin{pmatrix} 2.00 & 0.60 & 0.18 \\ 0.60 & 0.20 & 0.60 \\ 0.18 & 0.60 & 2.00 \end{pmatrix}$	$\begin{pmatrix} 1.10 & 0.55 & 0.28 \\ 0.55 & 1.10 & 0.55 \\ 0.28 & 0.55 & 1.10 \end{pmatrix}$	$\begin{pmatrix} 1.90 & 1.71 & 1.54 \\ 1.71 & 1.90 & 1.71 \\ 1.54 & 1.71 & 1.90 \end{pmatrix}$	$\begin{pmatrix} 1.40 & 1.26 & 1.13 \\ 1.26 & 1.40 & 1.26 \\ 1.13 & 1.26 & 1.40 \end{pmatrix}$

Appendix 2

Table 11 Parameters used to generate the simulated datasets under Scenario A₂

Parameter	Group 1	Group 2
π_g	0.50	0.50
M_g	$\begin{pmatrix} 1.00 & 2.00 & 2.00 & 0.00 \\ -1.00 & 1.00 & 1.00 & 2.00 \\ 0.00 & 2.00 & 2.00 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 1.00 & 2.00 & 2.00 & 0.00 \\ -1.00 & 1.00 & 1.00 & 2.00 \\ 0.00 & 2.00 & 2.00 & 1.00 \end{pmatrix}$
Φ_{X_g}	$\begin{pmatrix} 1.00 & 0.50 & 0.25 \\ 0.50 & 1.00 & 0.50 \\ 0.25 & 0.50 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 2.00 & 0.40 & 0.08 \\ 0.40 & 0.20 & 0.40 \\ 0.08 & 0.40 & 2.00 \end{pmatrix}$
Ψ_{X_g}	$\begin{pmatrix} 1.70 & 0.85 & 0.42 & 0.21 \\ 0.85 & 1.70 & 0.85 & 0.42 \\ 0.42 & 0.85 & 1.70 & 0.85 \\ 0.21 & 0.42 & 0.85 & 1.70 \end{pmatrix}$	$\begin{pmatrix} 1.00 & 0.50 & 0.25 & 0.12 \\ 0.50 & 1.00 & 0.50 & 0.25 \\ 0.25 & 0.50 & 1.00 & 0.50 \\ 0.12 & 0.25 & 0.50 & 1.00 \end{pmatrix}$
B_g^*	$\begin{pmatrix} 2.00 & 1.00 & 1.00 & -1.00 \\ 3.00 & 1.00 & -1.00 & 1.00 \end{pmatrix}$	$\begin{pmatrix} -7.00 & 1.00 & 1.00 & -1.00 \\ -8.00 & 1.00 & -1.00 & 1.00 \end{pmatrix}$
Φ_{Y_g}	$\begin{pmatrix} 1.00 & 0.50 \\ 0.50 & 1.00 \end{pmatrix}$	$\begin{pmatrix} 2.00 & 1.20 \\ 1.20 & 2.00 \end{pmatrix}$
Ψ_{Y_g}	$\begin{pmatrix} 2.00 & 1.00 & 0.50 & 0.25 \\ 1.00 & 2.00 & 1.00 & 0.50 \\ 0.50 & 1.00 & 2.00 & 1.00 \\ 0.25 & 0.50 & 1.00 & 2.00 \end{pmatrix}$	$\begin{pmatrix} 1.70 & 0.75 & 0.38 & 0.19 \\ 0.75 & 1.50 & 0.75 & 0.38 \\ 0.38 & 0.75 & 1.50 & 0.75 \\ 0.19 & 0.39 & 0.75 & 1.50 \end{pmatrix}$

Acknowledgements We thank the three anonymous reviewers and the associate editor for their useful comments and suggestions. We also thank Professor Salvatore Ingrassia for providing the education dataset as well as for helpful discussions.

Funding Open access funding provided by Università degli Studi di Catania within the CRUI-CARE Agreement.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

Anderlucci, L., Montanari, A., Viroli, C. (2014). A matrix-variate regression model with canonical states: An application to elderly Danish twins. *Statistica*, 74(4), 367–381.

- Anderlucci, L., Viroli, C., et al. (2015). Covariance pattern mixture models for the analysis of multivariate heterogeneous longitudinal data. *The Annals of Applied Statistics*, 9(2), 777–800.
- Dang, U.J., & McNicholas, P.D. (2015). Families of parsimonious finite mixtures of regression models. In *Advances in statistical models for data analysis* (pp. 73–84): Springer.
- Dang, U.J., Punzo, A., McNicholas, P.D., Ingrassia, S., Browne, R.P. (2017). Multivariate response and parsimony for Gaussian cluster-weighted models. *Journal of Classification*, 34(1), 4–34.
- Dayton, C.M., & Macready, G.B. (1988). Concomitant-variable latent-class models. *Journal of the American Statistical Association*, 83(401), 173–178.
- Dempster, A.P., Laird, N.M., Rubin, D.B. (1977). Maximum likelihood from incomplete data via the em algorithm. *Journal of the Royal Statistical Society: Series B (Methodological)*, 39(1), 1–22.
- DeSarbo, W.S., & Cron, W.L. (1988). A maximum likelihood methodology for clusterwise linear regression. *Journal of Classification*, 5(2), 249–282.
- Dutilleul, P. (1999). The MLE algorithm for the matrix normal distribution. *Journal of Statistical Computation and Simulation*, 64(2), 105–123.
- Frühwirth-Schnatter, S. (2006). *Finite mixture and markov switching models*. New York: Springer Science & Business Media.
- Gallaugher, M.P.B., & McNicholas, P.D. (2017). A matrix variate skew-t distribution. *Stat*, 6(1), 160–170.
- Gallaugher, M.P.B., & McNicholas, P.D. (2018). Finite mixtures of skewed matrix variate distributions. *Pattern Recognition*, 80, 83–93.
- Gershenfeld, N. (1997). Nonlinear inference and cluster-weighted modeling. *Annals of the New York Academy of Sciences*, 808(1), 18–24.
- Gershenfeld, N., Schoner, B., Metois, E. (1999). Cluster-weighted modelling for time-series analysis. *Nature*, 397(6717), 329.
- Hubert, L., & Arabie, P. (1985). Comparing partitions. *Journal of Classification*, 2(1), 193–218.
- Ingrassia, S., Minotti, S.C., Punzo, A. (2014). Model-based clustering via linear cluster-weighted models. *Computational Statistics & Data Analysis*, 71, 159–182.
- Ingrassia, S., Minotti, S.C., Vittadini, G. (2012). Local statistical modeling via a cluster-weighted approach with elliptical distributions. *Journal of Classification*, 29(3), 363–401.
- Ingrassia, S., & Punzo, A. (2016). Decision boundaries for mixtures of regressions. *Journal of the Korean Statistical Society*, 45(2), 295–306.
- Ingrassia, S., & Punzo, A. (2020). Cluster validation for mixtures of regressions via the total sum of squares decomposition. *Journal of Classification*, 37(2), 526–547.
- Ingrassia, S., Punzo, A., Vittadini, G., Minotti, S.C. (2015). The generalized linear mixed cluster-weighted model. *Journal of Classification*, 32(1), 85–113.
- Krause, K. (2006). On being strategic about the first year. In *Keynote presentation, Queensland University of Technology First Year Forum*, Vol. 5.
- Leisch, F. (2004). Flexmix: A general framework for finite mixture models and latent class regression in R. *Journal of Statistical Software*, 11(8), 1–18.
- Maitra, R., & Melnykov, V. (2010). Simulating data to study performance of finite mixture modeling and clustering algorithms. *Journal of Computational and Graphical Statistics*, 19(2), 354–376.
- Mazza, A., Ingrassia, S., Punzo, A. (2019). Modeling return to education in heterogeneous populations: An application to Italy. In Greselin, F., Deldossi, L., Bagnato, L., Vichi, M. (Eds.) *Statistical learning of complex data, volume 88 of studies in classification, data analysis, and knowledge organization, Cham, Switzerland* (pp. 121–131): Springer.
- Mazza, A., Punzo, A., Ingrassia, S. (2018). flexCWM: A flexible framework for cluster-weighted models, (Vol. 86 pp. 1–30).
- McLachlan, G., & Peel, D. (2000). *Finite mixture models*. New York: Wiley.
- McNicholas, P.D. (2016a). *Mixture Model-Based classification*. Boca Raton: Chapman and Hall/CRC Press.
- McNicholas, P.D. (2016b). Model-based clustering. *Journal of Classification*, 33(3), 331–373.
- Melnykov, V., & Zhu, X. (2018). On model-based clustering of skewed matrix data. *Journal of Multivariate Analysis*, 167, 181–194.
- Melnykov, V., & Zhu, X. (2019). Studying crime trends in the USA over the years 2000–2012. *Advances in Data Analysis and Classification*, 13(1), 325–341.
- Meng, X.-L., & Rubin, D.B. (1993). Maximum likelihood estimation via the ECM algorithm: a general framework. *Biometrika*, 80, 267–278.
- Michael, S., & Melnykov, V. (2016). An effective strategy for initializing the EM algorithm in finite mixture models. *Advances in Data Analysis and Classification*, 10(4), 563–583.
- Millo, G., & Carmeci, G. (2011). Non-life insurance consumption in Italy: a sub-regional panel data analysis. *Journal of Geographical Systems*, 13(3), 273–298.

- Millo, G., & Piras, G. (2012). **splm**: Spatial panel data models in R. *Journal of Statistical Software*, 47(1), 1–38.
- Punzo, A. (2014). Flexible mixture modelling with the polynomial Gaussian cluster-weighted model. *Statistical Modelling*, 14(3), 257–291.
- Punzo, A., & Ingrassia, S. (2015). Parsimonious generalized linear Gaussian cluster-weighted models. In Morlini, I., Minerva, T., Vichi, M. (Eds.) *Advances in statistical models for data analysis, studies in classification, data analysis and knowledge organization, Switzerland* (pp. 201–209): Springer International Publishing.
- Punzo, A., & Ingrassia, S. (2016). Clustering bivariate mixed-type data via the cluster-weighted model. *Computational Statistics*, 31(3), 989–1013.
- Punzo, A., & McNicholas, P.D. (2017). Robust clustering in regression analysis via the contaminated Gaussian cluster-weighted model. *Journal of Classification*, 34(2), 249–293.
- Qiu, W., & Joe, H. (2015). clusterGeneration: Random Cluster Generation (with Specified Degree of Separation). R package version 1.3.4.
- R Core Team. (2018). *R: a language and environment for statistical computing vienna*. Austria: R Foundation for Statistical Computing.
- Sarkar, S., Zhu, X., Melnykov, V., Ingrassia, S. (2020). On parsimonious models for modeling matrix data. *Computational Statistics & Data Analysis*, 142, 106822.
- Schwarz, G., et al. (1978). Estimating the dimension of a model. *The Annals of Statistics*, 6(2), 461–464.
- Subedi, S., Punzo, A., Ingrassia, S., McNicholas, P.D. (2013). Clustering and classification via cluster-weighted factor analyzers. *Advances in Data Analysis and Classification*, 7(1), 5–40.
- Tomarchio, S.D., Punzo, A., Bagnato, L. (2020). Two new matrix-variate distributions with application in model-based clustering. *Computational Statistics & Data Analysis*, 152, 107050.
- Viroli, C. (2011). Finite mixtures of matrix normal distributions for classifying three-way data. *Statistics and Computing*, 21(4), 511–522.
- Viroli, C. (2012). On matrix-variate regression analysis. *Journal of Multivariate Analysis*, 111, 296–309.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.