

A Mixture of Matrix Variate Bilinear Factor Analyzers

Michael P.B. Gallagher

Paul D. McNicholas*

Abstract

Over the years data has become increasingly higher dimensional, which has prompted an increased need for dimension reduction techniques. This is perhaps especially true for clustering (unsupervised classification) as well as semi-supervised and supervised classification. Although dimension reduction in the area of clustering for multivariate data has been quite thoroughly discussed within the literature, there is relatively little work in the area of three-way, or matrix variate, data. Herein, we develop a mixture of matrix variate bilinear factor analyzers (MMVBFA) model for use in clustering high-dimensional matrix variate data. This work can be considered both the first matrix variate bilinear factor analysis model as well as the first MMVBFA model. Parameter estimation is discussed, and the MMVBFA model is illustrated using simulated and real data.

Key Words: Clustering, factor analysis, matrix variate, mixture model

1. Introduction

Dimensionality is an ever present concern with data becoming increasingly higher dimensional over the last few years. To combat this issue, dimension reduction techniques have become very important tools, especially in the area of clustering (unsupervised classification) as well as semi-supervised and supervised classification. For multivariate data, the mixture of factor analyzers model has proved to be very useful in this regard as the model performs clustering and dimension reduction simultaneously, details in Section 2. However, there is relative paucity in the area of dimension reduction for use in model-based clustering for matrix variate data. Matrix variate distributions have been shown to be useful for modelling three-way data such as images and multivariate longitudinal data; however, the methods presented in the literature suffer from dimensionality concerns. In this paper, we present a mixture of matrix variate bilinear factor analyzers (MMVBFA) model for use in clustering higher dimensional matrix variate data. The matrix variate bilinear factor analyzers model can be viewed as a generalization of bilinear principal component analysis (BPCA; Zhao et al., 2012), and contains BPCA as a special case. An alternating expectation conditional maximization (AECM) algorithm (Meng and van Dyk, 1997) is used for parameter estimation. The proposed method is illustrated using both simulated and real datasets.

2. Background

2.1 Model-Based Clustering

Model-based clustering makes use of a finite mixture model. A G -component finite mixture model assumes a random variate \mathbf{X} has density

$$f(\mathbf{x} | \boldsymbol{\vartheta}) = \sum_{g=1}^G \pi_g f_g(\mathbf{x} | \boldsymbol{\theta}_g),$$

*Dept. of Mathematics and Statistics, McMaster University, Hamilton, ON, Canada L8S 4K1.

where $\boldsymbol{\vartheta} = (\pi_1, \pi_2, \dots, \pi_G, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_G)$, $f_g(\cdot)$ is the g th component density, and $\pi_g > 0$ is the g th mixing proportion such that $\sum_{i=1}^G \pi_g = 1$. The association between clustering and mixture models, as discussed in McNicholas (2016a), can be traced all the way back to Tiedeman (1955). The earliest use of a mixture model, specifically a Gaussian mixture model, for model-based clustering can be found in Wolfe (1965). Other early work in this area can be found in Baum et al. (1970) and Scott and Symons (1971) with a recent review given by McNicholas (2016b).

The Gaussian mixture model is well-established for clustering both multivariate and matrix variate data because of its mathematical tractability; however, there are many examples of non-Gaussian distributions used for clustering. In the multivariate case, work has been done using symmetric component densities that parameterize concentration (tail weight), e.g., the t distribution (Peel and McLachlan, 2000; Andrews and McNicholas, 2011, 2012; Steane et al., 2012; Morris and McNicholas, 2013a; Lin et al., 2014) and the power exponential distribution (Dang et al., 2015). Furthermore, there has been some work in the area of multivariate skewed-distributions, for example the normal-inverse Gaussian distribution (Karlis and Santourian, 2009; Subedi and McNicholas, 2014; O'Hagan et al., 2016), the skew- t distribution (Lin, 2010; Vrbik and McNicholas, 2012, 2014; Murray, Browne and McNicholas, 2014; Murray, McNicholas and Browne, 2014; Lee and McLachlan, 2014, 2016; Murray et al., 2017), the shifted asymmetric Laplace distribution (Morris and McNicholas, 2013b; Franczak et al., 2014), the generalized hyperbolic distribution (Browne and McNicholas, 2015; Wei et al., 2019), the variance-gamma distribution (McNicholas et al., 2017), and the joint generalized hyperbolic distribution (Tang et al., 2018).

In the area of matrix variate data, Viroli (2011) considers a mixture of matrix variate normal distributions for clustering, and Dođru et al. (2016) consider a mixture of matrix variate t distributions. More recently, Gallagher and McNicholas (2018a) consider mixtures of four skewed matrix variate distributions, specifically the matrix variate skew- t distribution (Gallagher and McNicholas, 2017) as well as generalized hyperbolic, variance-gamma and normal-inverse Gaussian distributions (Gallagher and McNicholas, 2018b), with an application in handwritten digit recognition. As pointed out by Gallagher and McNicholas (2018a), these approaches are limited by the dimensionality of the data and the work described herein aims to help address that limitation.

2.2 Matrix Variate Normal Distribution

An $n \times p$ random matrix \mathcal{X} follows a matrix variate normal distribution with location parameter \mathbf{M} and scale matrices $\boldsymbol{\Delta}$ and $\boldsymbol{\Omega}$ of dimensions $n \times n$ and $p \times p$, respectively, denoted by $\mathcal{N}_{n \times p}(\mathbf{M}, \boldsymbol{\Delta}, \boldsymbol{\Omega})$, if the density of \mathcal{X} can be written

$$\varphi_{n \times p}(\mathbf{X} | \mathbf{M}, \boldsymbol{\Delta}, \boldsymbol{\Omega}) = \frac{1}{(2\pi)^{\frac{np}{2}} |\boldsymbol{\Delta}|^{\frac{p}{2}} |\boldsymbol{\Omega}|^{\frac{n}{2}}} \exp \left\{ -\frac{1}{2} \text{tr} (\boldsymbol{\Delta}^{-1} (\mathbf{X} - \mathbf{M}) \boldsymbol{\Omega}^{-1} (\mathbf{X} - \mathbf{M})') \right\}.$$

One notable property of the matrix variate normal distribution (Harrar and Gupta, 2008) is

$$\mathcal{X} \sim \mathcal{N}_{n \times p}(\mathbf{M}, \boldsymbol{\Delta}, \boldsymbol{\Omega}) \iff \text{vec}(\mathcal{X}) \sim \mathcal{N}_{np}(\text{vec}(\mathbf{M}), \boldsymbol{\Omega} \otimes \boldsymbol{\Delta}), \quad (1)$$

where $\mathcal{N}_{np}(\cdot)$ is the multivariate normal density with dimension np , $\text{vec}(\cdot)$ is the vectorization operator, and \otimes is the Kronecker product.

2.3 Mixture of Factor Analyzers Model

For the purpose of this subsection, we temporarily revert back to the notation where \mathbf{X}_i represents a p -dimensional random vector, with \mathbf{x}_i as its realization. The factor analysis

model for $\mathbf{X}_1, \dots, \mathbf{X}_n$ is given by

$$\mathbf{X}_i = \boldsymbol{\mu} + \boldsymbol{\Lambda}\mathbf{U}_i + \boldsymbol{\varepsilon}_i,$$

where $\boldsymbol{\mu}$ is a location vector, $\boldsymbol{\Lambda}$ is a $p \times q$ matrix of factor loadings with $q < p$, $\mathbf{U}_i \sim \mathcal{N}_q(\mathbf{0}, \mathbf{I})$ denotes the latent factors, $\boldsymbol{\varepsilon}_i \sim \mathcal{N}_p(\mathbf{0}, \boldsymbol{\Psi})$, where $\boldsymbol{\Psi} = \text{diag}(\psi_1, \psi_2, \dots, \psi_p)$, and the \mathbf{U}_i and the $\boldsymbol{\varepsilon}_i$ are each independently distributed and are independent of one another. Under this model, the marginal distribution of \mathbf{X}_i is $\mathcal{N}_p(\boldsymbol{\mu}, \boldsymbol{\Lambda}\boldsymbol{\Lambda}' + \boldsymbol{\Psi})$. Probabilistic principal component analysis (PPCA) arises as a special case with the isotropic constraint $\boldsymbol{\Psi} = \psi\mathbf{I}$, $\psi \in \mathbb{R}^+$ (Tipping and Bishop, 1999b).

Ghahramani and Hinton (1997) develop the mixture of factor analyzers model, which is a Gaussian mixture model with covariance structure $\boldsymbol{\Sigma}_g = \boldsymbol{\Lambda}_g\boldsymbol{\Lambda}_g' + \boldsymbol{\Psi}$. McLachlan and Peel (2000) utilize the more general structure $\boldsymbol{\Sigma}_g = \boldsymbol{\Lambda}_g\boldsymbol{\Lambda}_g' + \boldsymbol{\Psi}_g$. Tipping and Bishop (1999a) introduce the closely-related mixture of PPCAs with $\boldsymbol{\Sigma}_g = \boldsymbol{\Lambda}_g\boldsymbol{\Lambda}_g' + \psi_g\mathbf{I}$. McNicholas and Murphy (2008) construct a family of eight parsimonious Gaussian models by considering the constraint $\boldsymbol{\Lambda}_g = \boldsymbol{\Lambda}$ in addition to $\boldsymbol{\Psi}_g = \boldsymbol{\Psi}$ and $\boldsymbol{\Psi}_g = \psi_g\mathbf{I}$. There has also been work on extending the mixture of factor analyzers to other distributions, such as the skew- t distribution (Murray, Browne and McNicholas, 2014; Murray et al., 2017), the generalized hyperbolic distribution (Tortora et al., 2016), the skew-normal distribution (Lin et al., 2016), the variance-gamma distribution (McNicholas et al., 2017), and others (e.g., Murray et al., 2017).

2.4 Previous Work on Matrix Variate Factor Analysis

Xie et al. (2008) and Yu et al. (2008) consider a matrix variate extension of PPCA in a linear fashion. For N independent $n \times p$ random matrices $\mathcal{X}_1, \dots, \mathcal{X}_N$, the model assumes

$$\mathcal{X}_i = \mathbf{M} + \mathbf{A}\mathcal{U}_i\mathbf{B}' + \mathcal{E}_i, \quad (2)$$

where \mathbf{M} is an $n \times p$ location matrix, \mathbf{A} is an $n \times q$ matrix of column factor loadings, \mathbf{B} is a $p \times r$ matrix of row factor loadings, $\mathcal{U}_i \sim \mathcal{N}_{q \times r}(\mathbf{0}, \mathbf{I}_q, \mathbf{I}_r)$, and $\mathcal{E}_i \sim \mathcal{N}_{n \times p}(\mathbf{0}, \sigma\mathbf{I}_n, \sigma\mathbf{I}_p)$, with $\sigma \in \mathbb{R}^+$. Note that the \mathcal{U}_i and the \mathcal{E}_i are each independently distributed and are independent of one another. The main disadvantage of this model is that, in general, \mathcal{X}_i does not follow a matrix variate normal distribution.

Zhao et al. (2012) present bilinear probabilistic principal component analysis (BPPCA) which extends (2) by adding two projected error terms. The resulting model assumes

$$\mathcal{X}_i = \mathbf{M} + \mathbf{A}\mathcal{U}_i\mathbf{B}' + \mathbf{A}\mathcal{E}_i^B + \mathcal{E}_i^A\mathbf{B}' + \mathcal{E}_i, \quad (3)$$

where $\mathcal{E}_i^B \sim \mathcal{N}_{q \times p}(\mathbf{0}, \mathbf{I}_q, \sigma_B\mathbf{I}_p)$, $\mathcal{E}_i^A \sim \mathcal{N}_{n \times r}(\mathbf{0}, \sigma_A\mathbf{I}_n, \mathbf{I}_r)$, $\mathcal{E}_i \sim \mathcal{N}_{n \times p}(\mathbf{0}, \sigma_A\mathbf{I}_n, \sigma_B\mathbf{I}_p)$, with $\sigma_A \in \mathbb{R}^+$ and $\sigma_B \in \mathbb{R}^+$, and the other terms are as defined for (2). In this model, each of the \mathcal{U}_i , \mathcal{E}_i^B , \mathcal{E}_i^A and \mathcal{E}_i are independently distributed and all are independent of each other. It is important to note that the term ‘‘column factors’’ refers to reduction in the dimension of the columns, which is equivalent to the number of rows, and not a reduction in the number of columns. Likewise, the term ‘‘row factors’’ refers to the reduction in the dimension of the rows (i.e., in the number of columns). As discussed by Zhao et al. (2012), the interpretation of the \mathcal{E}_i^B and \mathcal{E}_i^A are the row and column noise, respectively, whereas \mathcal{E}_i is the common noise. It can be shown using property (1) that under this model $\mathcal{X}_i \sim \mathcal{N}_{n \times p}(\mathbf{M}, \mathbf{A}\mathbf{A}' + \sigma_A\mathbf{I}_n, \mathbf{B}\mathbf{B}' + \sigma_B\mathbf{I}_p)$. Note that the covariance structure for the two covariance matrices of this matrix variate normal distribution are analogous to the covariance structure for the (multivariate) factor analysis model.

3. Methodology

3.1 MMVBFA Model

A MMVBFA model is derived here by extending (3). Specifically, we remove the isotropic constraint and assume that

$$\mathcal{X}_i = \mathbf{M}_g + \mathbf{A}_g \mathcal{U}_{ig} \mathbf{B}'_g + \mathbf{A}_g \mathcal{E}_{ig}^B + \mathcal{E}_{ig}^A \mathbf{B}'_g + \mathcal{E}_{ig} \quad (4)$$

with probability π_g , for $g = 1, \dots, G$, where \mathbf{M}_g is an $n \times p$ location matrix, \mathbf{A}_g is an $n \times q$ column factor loading matrix, with $q < n$, \mathbf{B}_g is a $p \times r$ row factor loading matrix, with $r < p$, and

$$\begin{aligned} \mathcal{U}_{ig} &\sim \mathcal{N}_{q \times r}(\mathbf{0}, \mathbf{I}_q, \mathbf{I}_r), \\ \mathcal{E}_{ig}^B &\sim \mathcal{N}_{q \times p}(\mathbf{0}, \mathbf{I}_q, \mathbf{\Psi}_g), \\ \mathcal{E}_{ig}^A &\sim \mathcal{N}_{n \times r}(\mathbf{0}, \mathbf{\Sigma}_g, \mathbf{I}_r), \\ \mathcal{E}_{ig} &\sim \mathcal{N}_{n \times p}(\mathbf{0}, \mathbf{\Sigma}_g, \mathbf{\Psi}_g), \end{aligned}$$

each independently distributed and independent of each other, $\mathbf{\Sigma}_g = \text{diag}\{\sigma_{1g}, \dots, \sigma_{ng}\}$, with $\sigma_{ig} \in \mathbb{R}^+$, and $\mathbf{\Psi}_g = \text{diag}\{\psi_{1g}, \dots, \psi_{pg}\}$, with $\psi_{ig} \in \mathbb{R}^+$.

Let $\mathbf{z}_i = (z_{i1}, \dots, z_{iG})'$ denote the component membership for \mathbf{X}_i , where

$$z_{ig} = \begin{cases} 1 & \text{if } \mathbf{X}_i \text{ belongs to component } g, \\ 0 & \text{otherwise,} \end{cases}$$

for $i = 1, \dots, N$ and $g = 1, \dots, G$. Using the vectorization of \mathcal{X}_i , and property (1), it can be shown that

$$\mathcal{X}_i | z_{ig} = 1 \sim \mathcal{N}_{n \times p}(\mathbf{M}_g, \mathbf{\Sigma}_g + \mathbf{A}_g \mathbf{A}'_g, \mathbf{\Psi}_g + \mathbf{B}_g \mathbf{B}'_g).$$

Therefore, the density of \mathcal{X}_i can be written

$$f(\mathbf{X}_i | \boldsymbol{\vartheta}) = \sum_{g=1}^G \pi_g \varphi_{n \times p}(\mathbf{X}_i | \mathbf{M}_g, \mathbf{\Sigma}_g + \mathbf{A}_g \mathbf{A}'_g, \mathbf{\Psi}_g + \mathbf{B}_g \mathbf{B}'_g),$$

where $\varphi_{n \times p}(\cdot)$ denotes the $n \times p$ matrix variate normal density (see Section 2.2). Following a similar procedure to that described by Zhao et al. (2012), by introducing latent matrix variables \mathcal{Y}_{ig}^B and \mathcal{V}_{ig}^B , (4) can be written

$$\begin{aligned} \mathcal{X}_i &= \mathbf{M}_g + \mathbf{A}_g \mathcal{Y}_{ig}^B + \mathcal{V}_{ig}^B, \\ \mathcal{Y}_{ig}^B &= \mathcal{U}_{ig} \mathbf{B}'_g + \mathcal{E}_{ig}^B, \\ \mathcal{V}_{ig}^B &= \mathcal{E}_{ig}^A \mathbf{B}'_g + \mathcal{E}_{ig}. \end{aligned}$$

The two-stage interpretation of this formulation of the model is the same as that given by Zhao et al. (2012), where this can be viewed as first projecting \mathcal{X}_i in the column direction onto the latent matrix \mathcal{Y}_{ig}^B , and then \mathcal{Y}_{ig}^B and \mathcal{V}_{ig}^B are further projected in the row direction. Likewise, introducing \mathcal{Y}_{ig}^A and \mathcal{V}_{ig}^A , (4) can be written

$$\begin{aligned} \mathcal{X}_i &= \mathbf{M}_g + \mathcal{Y}_{ig}^A \mathbf{B}'_g + \mathcal{V}_{ig}^A, \\ \mathcal{Y}_{ig}^A &= \mathbf{A}_g \mathcal{U}_{ig} + \mathcal{E}_{ig}^A, \\ \mathcal{V}_{ig}^A &= \mathbf{A}_g \mathcal{E}_{ig}^B + \mathcal{E}_{ig}. \end{aligned}$$

The interpretation is the same as before only we project in the row direction first followed by the column direction. It can be shown that

$$\mathcal{Y}_{ig}^B | \mathbf{X}_i, z_{ig} = 1 \sim \mathcal{N}_{q \times p}(\mathbf{W}_g^A \mathbf{A}_g' \boldsymbol{\Sigma}_g^{-1} (\mathbf{X}_i - \mathbf{M}_g), \mathbf{W}_g^A, \boldsymbol{\Lambda}_{\mathbf{B}_g})$$

and

$$\mathcal{Y}_{ig}^A | \mathbf{X}_i, z_{ig} = 1 \sim \mathcal{N}_{n \times r}((\mathbf{X}_i - \mathbf{M}_g) \boldsymbol{\Psi}_g^{-1} \mathbf{B}_g \mathbf{W}_g^B, \boldsymbol{\Lambda}_{\mathbf{A}_g}, \mathbf{W}_g^B),$$

where $\mathbf{W}_g^A = \mathbf{I}_q + \mathbf{A}_g' \boldsymbol{\Sigma}_g^{-1} \mathbf{A}_g$, $\mathbf{W}_g^B = \mathbf{I}_r + \mathbf{B}_g' \boldsymbol{\Psi}_g^{-1} \mathbf{B}_g$, $\boldsymbol{\Lambda}_{\mathbf{A}_g} = \boldsymbol{\Sigma}_g + \mathbf{A}_g \mathbf{A}_g'$, and $\boldsymbol{\Lambda}_{\mathbf{B}_g} = \boldsymbol{\Psi}_g + \mathbf{B}_g \mathbf{B}_g'$

3.2 Parameter Estimation

Suppose we observe N observations $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N$ then the log-likelihood is given by

$$\mathcal{L}(\boldsymbol{\vartheta}) = \sum_{i=1}^N \log \sum_{g=1}^G \pi_g \varphi_{n \times p}(\mathbf{X}_i | \mathbf{M}_g, \boldsymbol{\Sigma}_g + \mathbf{A}_g \mathbf{A}_g', \boldsymbol{\Psi}_g + \mathbf{B}_g \mathbf{B}_g'). \quad (5)$$

To maximize (5), the observed data is viewed as incomplete and an AECM is then to maximize (5). There are three different sources of missingness: the component memberships $\mathbf{z}_1, \dots, \mathbf{z}_n$ as well as the latent matrix variables \mathcal{Y}_{ig}^B and \mathcal{Y}_{ig}^A . A three-stage AECM algorithm is now described for parameter estimation.

AECM Stage 1: In the first stage, the complete-data is taken to be the observed matrices $\mathbf{X}_1, \dots, \mathbf{X}_N$ and the component memberships $\mathbf{z}_1, \dots, \mathbf{z}_N$, and the updates for π_g and \mathbf{M}_g are calculated. The complete-data log-likelihood in the first stage is

$$\ell_1 = C_1 + \sum_{g=1}^G \sum_{i=1}^N z_{ig} \left\{ \log \pi_g - \frac{1}{2} \text{tr}[\boldsymbol{\Lambda}_{\mathbf{A}_g}^{-1} (\mathbf{X}_i - \mathbf{M}_g) \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} (\mathbf{X}_i - \mathbf{M}_g)'] \right\},$$

where C_1 is a constant with respect to π_g and \mathbf{M}_g . In the E-Step, the updates for the component memberships z_{ig} are given by

$$\hat{z}_{ig} = \frac{\pi_g \varphi_{n \times p}(\mathbf{X}_i | \mathbf{M}_g, \boldsymbol{\Lambda}_{\mathbf{A}_g}, \boldsymbol{\Lambda}_{\mathbf{B}_g})}{\sum_{h=1}^G \pi_h \varphi_{n \times p}(\mathbf{X}_i | \mathbf{M}_h, \boldsymbol{\Lambda}_{\mathbf{A}_h}, \boldsymbol{\Lambda}_{\mathbf{B}_h})},$$

where $\varphi_{n \times p}(\cdot)$ denotes the $n \times p$ matrix variate normal density. As usual, these expectations \hat{z}_{ig} are calculated using the current estimates of the parameters. In the CM-step, the updates for π_g and \mathbf{M}_g are calculated using

$$\hat{\pi}_g = \frac{N_g}{N} \quad \text{and} \quad \hat{\mathbf{M}}_g = \frac{1}{N_g} \sum_{i=1}^N \hat{z}_{ig} \mathbf{X}_i,$$

respectively, where $N_g = \sum_{i=1}^N \hat{z}_{ig}$.

AECM Stage 2: In the second stage, the complete-data is taken to be the observed $\mathbf{X}_1, \dots, \mathbf{X}_N$, the component memberships $\mathbf{z}_1, \dots, \mathbf{z}_N$ and the $n \times q$ latent matrices \mathcal{Y}_{ig}^B . The complete-data log-likelihood is then

$$\begin{aligned} \ell_2 = & C_2 - \frac{N_g p}{2} \log |\boldsymbol{\Sigma}_g| - \frac{1}{2} \sum_{g=1}^G \sum_{i=1}^N z_{ig} \text{tr} [\boldsymbol{\Sigma}_g^{-1} (\mathbf{X}_i - \mathbf{M}_g) \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} (\mathbf{X}_i - \mathbf{M}_g)' \\ & - \boldsymbol{\Sigma}_g^{-1} \mathbf{A}_g \mathcal{Y}_{ig}^B \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} (\mathbf{X}_i - \mathbf{M}_g)' - \boldsymbol{\Sigma}_g^{-1} (\mathbf{X}_i - \mathbf{M}_g) \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} \mathcal{Y}_{ig}^B \mathbf{A}_g' \\ & + \boldsymbol{\Sigma}_g^{-1} \mathbf{A}_g \mathcal{Y}_{ig}^B \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} \mathcal{Y}_{ig}^B \mathbf{A}_g'], \end{aligned}$$

where C_2 is constant with respect to the parameters being updated. In the E-Step, the following expectations are calculated:

$$\begin{aligned} a_{ig}^B &:= \mathbb{E}[\mathcal{Y}_{ig}^B \mid \mathbf{X}_i, z_{ig} = 1] = \mathbf{W}_g^{A^{-1}} \mathbf{A}'_g \boldsymbol{\Sigma}_g^{-1} (\mathbf{X}_i - \mathbf{M}_g), \\ b_{ig}^B &:= \mathbb{E}[\mathcal{Y}_{ig}^B \hat{\boldsymbol{\Lambda}}_{\mathbf{B}_g}^{-1} \mathcal{Y}_{ig}^{B'} \mid \mathbf{X}_i, z_{ig} = 1] = p \mathbf{W}_g^{A^{-1}} + a_{ig}^B \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} a_{ig}^{B'}. \end{aligned}$$

As usual, these expectations are calculated using the current estimates of the parameters. In the CM-step, \mathbf{A}_g and $\boldsymbol{\Sigma}_g$ are updated via

$$\hat{\mathbf{A}}_g = \sum_{i=1}^N \hat{z}_{ig} (\mathbf{X}_i - \hat{\mathbf{M}}_g) \hat{\boldsymbol{\Lambda}}_{\mathbf{B}_g}^{-1} a_{ig}^{B'} \left(\sum_{i=1}^N \hat{z}_{ig} b_{ig}^B \right)^{-1} \quad \text{and} \quad \hat{\boldsymbol{\Sigma}}_g = \frac{1}{N_g p} \text{diag}\{\hat{\mathbf{S}}_g^B\},$$

respectively, where

$$\mathbf{S}_g^B = \sum_{i=1}^N \hat{z}_{ig} [(\mathbf{X}_i - \hat{\mathbf{M}}_g) \hat{\boldsymbol{\Lambda}}_{\mathbf{B}_g}^{-1} (\mathbf{X}_i - \hat{\mathbf{M}}_g)' - \hat{\mathbf{A}}_g a_{ig}^B \hat{\boldsymbol{\Lambda}}_{\mathbf{B}_g}^{-1} (\mathbf{X}_i - \hat{\mathbf{M}}_g)'].$$

AECM Stage 3: In the last stage of the AECM algorithm, the complete-data is taken to be the observed $\mathbf{X}_1, \dots, \mathbf{X}_N$, the component memberships $\mathbf{z}_1, \dots, \mathbf{z}_N$ and the $p \times r$ latent matrices \mathcal{Y}_{ig}^A . In this step, the complete-data log-likelihood is

$$\begin{aligned} \ell_3 = C_3 &- \frac{N_g n}{2} \log |\boldsymbol{\Psi}_g| - \frac{1}{2} \sum_{g=1}^G \sum_{i=1}^N z_{ig} \text{tr} [\boldsymbol{\Psi}_g^{-1} (\mathbf{X}_i - \mathbf{M}_g)' \boldsymbol{\Lambda}_{\mathbf{A}_g}^{-1} (\mathbf{X}_i - \mathbf{M}_g) \\ &- \boldsymbol{\Psi}_g^{-1} \mathbf{B}_g \mathcal{Y}_{ig}^{A'} \boldsymbol{\Lambda}_{\mathbf{A}_g}^{-1} (\mathbf{X}_i - \mathbf{M}_g) - \boldsymbol{\Psi}_g^{-1} (\mathbf{X}_i - \mathbf{M}_g)' \boldsymbol{\Lambda}_{\mathbf{A}_g}^{-1} \mathcal{Y}_{ig}^A \mathbf{B}_g' \\ &+ \boldsymbol{\Psi}_g^{-1} \mathbf{B}_g \mathcal{Y}_{ig}^{A'} \boldsymbol{\Lambda}_{\mathbf{A}_g}^{-1} \mathcal{Y}_{ig}^A \mathbf{B}_g'], \end{aligned}$$

where C_3 is constant with respect to the parameters being updated. In the E-Step, expectations similar to those in the second step are calculated, i.e.,

$$a_{ig}^A := \mathbb{E}[\mathcal{Y}_{ig}^A \mid \mathbf{X}_i, z_{ig} = 1] = (\mathbf{X}_i - \mathbf{M}_g) \boldsymbol{\Psi}_g^{-1} \mathbf{B}_g \mathbf{W}_g^{B^{-1}}$$

and

$$b_{ig}^A := \mathbb{E}[\mathcal{Y}_{ig}^{A'} \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} \mathcal{Y}_{ig}^A \mid \mathbf{X}_i, z_{ig} = 1] = n \mathbf{W}_g^{B^{-1}} + a_{ig}^A \boldsymbol{\Lambda}_{\mathbf{B}_g}^{-1} a_{ig}^A.$$

As usual, these expectations are calculated using the current estimates of the parameters. In the CM-step, we update \mathbf{B}_g and $\boldsymbol{\Psi}_g$ by

$$\hat{\mathbf{B}}_g = \sum_{i=1}^N \hat{z}_{ig} (\mathbf{X}_i - \hat{\mathbf{M}}_g)' \hat{\boldsymbol{\Lambda}}_{\mathbf{A}_g}^{-1} a_{ig}^A \left(\sum_{i=1}^N \hat{z}_{ig} b_{ig}^A \right)^{-1} \quad \text{and} \quad \hat{\boldsymbol{\Psi}}_g = \frac{1}{N_g n} \text{diag}\{\hat{\mathbf{S}}_g^A\},$$

respectively, where

$$\mathbf{S}_g^A = \sum_{i=1}^N \hat{z}_{ig} [(\mathbf{X}_i - \hat{\mathbf{M}}_g)' \hat{\boldsymbol{\Lambda}}_{\mathbf{A}_g}^{-1} (\mathbf{X}_i - \hat{\mathbf{M}}_g) - \hat{\mathbf{B}}_g a_{ig}^A \hat{\boldsymbol{\Lambda}}_{\mathbf{A}_g}^{-1} (\mathbf{X}_i - \hat{\mathbf{M}}_g)].$$

3.3 Semi-Supervised Classification

The MMVBFA model presented herein for clustering may also be used for semi-supervised classification. Suppose N matrices are observed, and K of these observations have known labels from one of G classes. Following McNicholas (2010), without loss of generality,

order the matrices so that the first K have known labels and the remaining observations have unknown labels. The observed likelihood is then

$$L(\boldsymbol{\vartheta}) = \prod_{i=1}^K \prod_{g=1}^G [\pi_g \varphi_{n \times p}(\mathbf{X}_i | \mathbf{M}_g, \boldsymbol{\Sigma}_g + \mathbf{A}_g \mathbf{A}'_g, \boldsymbol{\Psi}_g + \mathbf{B}_g \mathbf{B}'_g)]^{z_{ig}} \\ \times \prod_{j=K+1}^N \sum_{h=1}^H \pi_h \varphi_{n \times p}(\mathbf{X}_j | \mathbf{M}_h, \boldsymbol{\Sigma}_h + \mathbf{A}_h \mathbf{A}'_h, \boldsymbol{\Psi}_h + \mathbf{B}_h \mathbf{B}'_h).$$

It is possible for $H \neq G$; however, for our analyses we assume that $H = G$. Parameter estimation then proceeds in a similar manner for the clustering scenario. For more information on semi-supervised classification refer to McNicholas (2016a).

3.4 Model Selection, Initialization and Convergence

For a typical dataset the number of components and/or the number of factors will not be known a priori and, therefore, we will have to select them. One common selection criterion is the Bayesian information criterion (BIC; Schwarz, 1978) and is given by

$$\text{BIC} = 2\ell(\hat{\boldsymbol{\vartheta}}) - \rho \log N,$$

where $\ell(\hat{\boldsymbol{\vartheta}})$ is the maximized log-likelihood and ρ is the number of free parameters. The BIC is used as the selection criterion for all of our analyses.

To initialize the AECM algorithm, we employ an alternating emEM strategy (Biernacki et al., 2003). This consists of running the AECM algorithm for a small number of iterations for different random starting values of the parameters and then using the parameters that maximize the likelihood to continue with the AECM algorithm until convergence.

The simplest convergence criterion would be to use lack of progress in the log-likelihood, however; it is possible for the log-likelihood to “plateau” and then increase again thus terminating the algorithm prematurely (see McNicholas, 2016a). One alternative is to use a criterion based on the Aitken acceleration (Aitken, 1926). The acceleration at iteration t is

$$a^{(t)} = \frac{l^{(t+1)} - l^{(t)}}{l^{(t)} - l^{(t-1)}},$$

where $l^{(t)}$ is the observed likelihood at iteration t . Now,

$$l_{\infty}^{(t+1)} = l^{(t)} + \frac{(l^{(t+1)} - l^{(t)})}{1 - a^{(t)}},$$

is an estimate of the observed log-likelihood after many iterations, at iteration $t + 1$ (see Böhning et al., 1994; Lindsay, 1995). As in McNicholas et al. (2010), we terminate the algorithm when $l_{\infty}^{(k+1)} - l^{(k)} \in (0, \epsilon)$. It is important to note that, in each AECM algorithm run for the analyses herein, we make the choice of ϵ based on the magnitude of the log-likelihood. Specifically, after running the 10 iterations of the emEM algorithm, we choose ϵ to be four orders of magnitude lower than the log-likelihood.

3.5 Reduction in Number of Free Covariance Parameters

Because the covariance structure of both covariance matrices in the MVVBFA model is equivalent to the covariance structure in the (multivariate) mixture of factor analyzers model, many of the results on the number of free covariance parameters may be used here. Specifically there are $nq + n - q(q - 1)/2$ free covariance parameters in $\boldsymbol{\Lambda}_{\mathbf{A}_g}$ and

$pr + p - r(r - 1)/2$ free covariance parameters in Λ_{B_g} (Lawley and Maxwell, 1962). Therefore, reduction in the number of free covariance parameters for the row covariance matrix is

$$\frac{1}{2}n(n + 1) - nq - n + \frac{1}{2}q(q - 1) = \frac{1}{2}[(n - q)^2 - (n + q)],$$

which is positive for $(n - q)^2 > n + q$. Likewise for the column covariance matrix the reduction in the number of parameters is

$$\frac{1}{2}p(p + 1) - pr - p + \frac{1}{2}r(r - 1) = \frac{1}{2}[(p - r)^2 - (p + r)],$$

which is positive for $(p - r)^2 > p + r$. In applications herein, the model is fit for a range of row factors and column factors. If the number of row or column factors chosen by the BIC is the maximum in that range, the relevant number of factors will be increased so long as the aforementioned conditions are met.

4. Data Analyses

4.1 Simulations

Simulation 1

In the first simulation, $G = 2$ groups are considered with 10×7 matrices. The mixing proportions are taken to be $\pi_1 = \pi_2 = 0.5$, and we set $N \in \{200, 400, 800\}$. Observations are simulated from (4) with $q = 2$ column factors and $r = 3$ row factors. For each value of N , 50 datasets are simulated. For each dataset, for each N , the correct number of groups, column and row factors are selected. In addition, perfect classification is achieved (ARI = 1). Note that the adjusted Rand index (ARI; Rand, 1971; Hubert and Arabie, 1985) is often used to assess agreement between true and predicted classes; it takes a value of 1 for perfect class agreement and has expected value 0 under random class assignment. In Table 1, we show the average value of $\|\mathbf{M}_g - \hat{\mathbf{M}}_g\|_1$, for $g = 1, 2$ and for each value of N , over the 50 datasets. Note that if \mathbf{W} is an $n \times p$ matrix then

$$\|\mathbf{W}\|_1 = \max_{1 \leq j \leq p} \sum_{i=1}^n |w_{ij}|.$$

As expected, the estimates of \mathbf{M}_g get closer to the true values as the sample size N is increased. Moreover, the variability of $\|\mathbf{M}_g - \hat{\mathbf{M}}_g\|_1$ decreases as the sample size increases.

Table 1: Average $\|\mathbf{M}_g - \hat{\mathbf{M}}_g\|_1$ values over 50 datasets, for $g = 1, 2$ and $N = 200, 400, 800$, in Simulation 1, with standard deviations in parentheses.

g	N		
	200	400	800
1	13.97(3.61)	9.66(2.65)	6.48(1.69)
2	12.08(3.25)	7.45(1.79)	5.69(1.32)

Simulation 2

The second simulation considers $G = 3$ groups with 28×17 matrices. The mixing proportions are $\pi_1 = \pi_3 = 0.4$ and $\pi_2 = 0.2$, and $N \in \{250, 500, 1000\}$. Again, 50

datasets are simulated for each N with $q = 2$ column factors and $r = 3$ row factors. As in Simulation 1, the correct number of groups, column and row factors are chosen and perfect classification is achieved. In Table 2, we again show the average 1-norms for the differences between the true and estimated location parameters.

Table 2: Average $\|\mathbf{M}_g - \hat{\mathbf{M}}_g\|_1$ values over 50 datasets, for $g = 1, 2, 3$ and $N = 250, 500, 1000$, in Simulation 2, with standard deviations in parentheses.

g	N		
	250	500	1000
1	36.28(7.95)	26.36(5.12)	19.37(4.62)
2	55.23(11.75)	40.42(9.64)	29.30(6.26)
3	39.10(8.89)	27.09(6.37)	19.99(4.45)

4.2 MNIST Digit Recognition

We consider the 28×28 MNIST digit dataset (LeCun et al., 1998), which contains over 60,000 greyscale images of handwritten Arabic digits 0 to 9. The images are represented by 28×28 pixel matrices with greyscale intensities ranging from 0 to 255. Because of the lack of variability in the outer rows and columns, some random noise is added while adding 50 to each of the non-zero elements to avoid confusing the noise with a true signal. We are interested in comparing digit 1 to digit 7, as was considered in Gallagher and McNicholas (2018a). Similar to Gallagher and McNicholas (2018a), we consider semi-supervised classification with 25%, 50% and 75% supervision. In each case, 25 datasets are considered, each consisting of 200 observations from each digit, and we fit the model for 10 to 20 column and row factors.

In Table 3, we show an aggregated classification table between the true and predicted classifications at each level of supervision for the points considered unlabelled. As expected, slightly better classification performance is obtained when the level of supervision is increased. Moreover, there is a more substantial difference when going from 25% supervision to 50% supervision than from 50% to 75%.

Table 3: Cross-tabulations of true (1,7) versus predicted (P1, P7) classifications for the observations considered unlabelled in the MNIST data at each level of supervision, aggregated over all runs.

	25% Supervision		50% Supervision		75% Supervision	
	P1	P7	P1	P7	P1	P7
1	3550	173	2449	53	1232	26
7	200	3577	51	2447	18	1221

Table 4 shows the average ARI and misclassification rate (MCR) over the 25 datasets, with the respective standard deviations, for each level of supervision. We note that we obtain better results than Gallagher and McNicholas (2018a) even with a lower level of supervision; however, the results in Gallagher and McNicholas (2018a) were based on resized images due to dimensionality constraints whereas this analysis was performed on the original images.

In Table 5 the frequency of the number of factors chosen for each level of supervision over the 25 datasets is shown. For the majority of the datasets, the number of row and column factors lie between 13 and 15.

Table 4: Average ARI values and misclassification rates (MCR), with associated standard deviations in parentheses, for each level of supervision for the points considered unlabelled for the MNIST data, aggregated over all runs.

	ARI (std. dev.)	MCR(std. dev.)
25%	0.82(0.15)	0.050(0.046)
50%	0.92 (0.056)	0.021 (0.015)
75%	0.93 (0.056)	0.018 (0.015)

Table 5: Numbers of row and columns factors chosen for the MNIST dataset for 25%, 50% and 75% supervision.

	10	11	12	13	14	15	16	17	18	19	20
25% Supervision											
Row Factors	0	0	0	2	7	6	4	3	2	1	0
Column Factors	0	0	2	6	7	6	3	1	0	0	0
50% Supervision											
Row Factors	0	0	0	4	6	10	2	0	1	1	1
Column Factors	0	0	2	9	7	5	1	1	0	0	0
75% Supervision											
Row Factors	0	0	0	1	9	9	3	3	0	0	0
Column Factors	0	0	0	9	11	4	0	0	0	0	1

Finally, in Figure 1, heatmaps are displayed for the average estimates of the location matrices over the 25 runs for each level of supervision for both digits. We see a slight increase in quality when going from 25% to 50% supervision for digit 7 with the centre of the digit being a little smoother with 50% supervision. There is no noticeable difference when going from 50% to 75% supervision. This similarity across the three levels of supervision illustrates the power of semi-supervised classification.

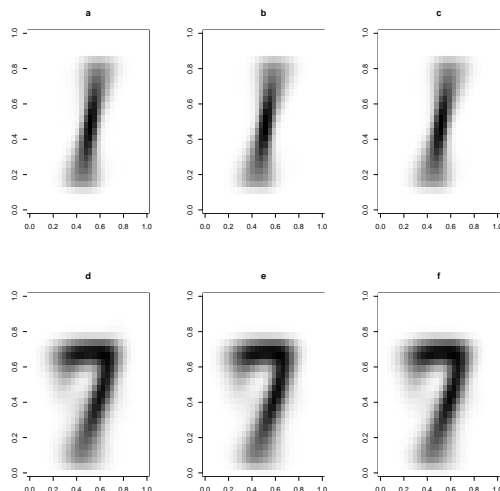


Figure 1: Heatmaps for the average estimated location matrices taken over the 25 runs for digit 1 at 25%, 50% and 75% supervision, respectively (a, b, c), and digit 7 at 25%, 50% and 75% supervision, respectively (d, e, f).

4.3 Olivetti Faces Dataset

Finally, consider the Olivetti faces dataset from the R package `RnavGraphImageData` (Waddell and Oldford, 2013). The dataset consists of greyscale images of faces that were taken between 1992 and 1994 at AT&T laboratories in Cambridge. There were 40 individuals with 10 images of each individual for a total of 400 64×64 images. The images were taken with varied lighting, expressions (eyes open/closed, smile/frown etc.), and glasses or no glasses. We fit the model for 15 to 30 column and row factors, and for $G = 1, \dots, 9$ components. The BIC chooses three components with 23 column factors and 26 row factors. The estimated mixing proportions are $\pi_1 = 0.22, \pi_2 = 0.49, \pi_3 = 0.29$. In Figure 2, we show a heatmap of the estimated location parameters for each component. The heatmap for component 3 arguably shows the clearest image and appears to display the glasses feature.

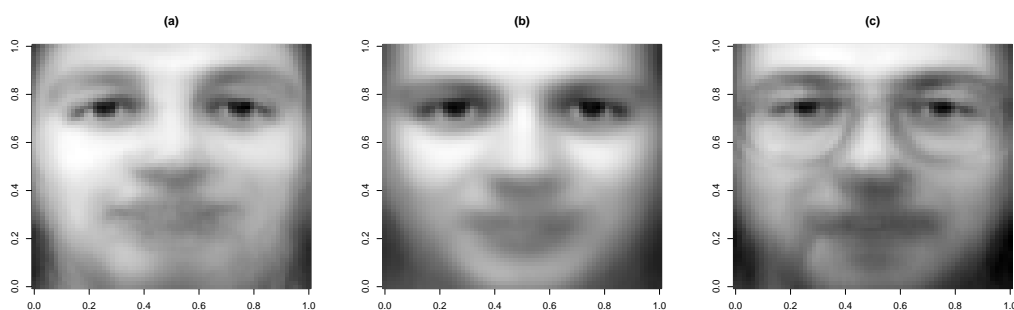


Figure 2: Estimated location matrices for (a) component 1, (b) component 2, and (c) component 3 for the faces dataset.

Upon looking at individual faces classified to component 3 (Figure 2), all the faces have glasses. Moreover, all faces with glasses are classified to component 3 with the exception of two which are classified to component 2. The faces with closed eyes are scattered throughout the three different components and are not classified to any one component. Although it is difficult to determine the main feature that differentiates component 1 from component 2, it is apparent that the eyebrows for the faces classified to component 1 tend to be more prominent and higher above the eyelid. Of course, a semi-supervised approach to these data could be used to detect specific classes, similar to the MNIST analysis (Section 4.2). However, the unsupervised analysis here has shown that the MMVBFA approach can be effective at detecting subgroups without training.

5. Summary

In this paper, we developed a MMVBFA model for use in clustering and classification of matrix variate data. Two simulations as well as two real data examples were used for illustration. For each of the simulations, the correct number of components and column/row factors were chosen by the BIC for all of the datasets. Perfect classification performance was also obtained in the simulations. In the MNIST digit application, even with a lower level of supervision, we obtained better results than Gallagher and McNicholas (2018a). However, this is probably due to the fact that the MMVBFA model could use the full 28×28 image. In the faces application, the BIC chooses three groups with the third group being defined by the presence of the glasses facial feature. The matrix normality of \mathcal{X} in the MMVBFA model will allow for direct extensions to mixtures of matrix variate t factor analyzers, as well as skewed matrix variate factor analyzers analogous to their multivariate counterparts.

References

- Aitken, A. C. (1926), "A series formula for the roots of algebraic and transcendental equations", *Proceedings of the Royal Society of Edinburgh* 45, 14–22.
- Andrews, J. L. and McNicholas, P. D. (2011), "Extending mixtures of multivariate t-factor analyzers", *Statistics and Computing* 21(3), 361–373.
- Andrews, J. L. and McNicholas, P. D. (2012), "Model-based clustering, classification, and discriminant analysis via mixtures of multivariate t -distributions: The $tEIGEN$ family", *Statistics and Computing* 22(5), 1021–1029.
- Baum, L. E., Petrie, T., Soules, G. and Weiss, N. (1970), "A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains", *Annals of Mathematical Statistics* 41, 164–171.
- Biernacki, C., Celeux, G. and Govaert, G. (2003), "Choosing starting values for the EM algorithm for getting the highest likelihood in multivariate Gaussian mixture models", *Computational Statistics and Data Analysis* 41, 561–575.
- Böhning, D., Dietz, E., Schaub, R., Schlattmann, P. and Lindsay, B. (1994), "The distribution of the likelihood ratio for mixtures of densities from the one-parameter exponential family", *Annals of the Institute of Statistical Mathematics* 46, 373–388.
- Browne, R. P. and McNicholas, P. D. (2015), "A mixture of generalized hyperbolic distributions", *Canadian Journal of Statistics* 43(2), 176–198.
- Dang, U. J., Browne, R. P. and McNicholas, P. D. (2015), "Mixtures of multivariate power exponential distributions", *Biometrics* 71(4), 1081–1089.
- Doğru, F. Z., Bulut, Y. M. and Arslan, O. (2016), "Finite mixtures of matrix variate t distributions", *Gazi University Journal of Science* 29(2), 335–341.
- Franczak, B. C., Browne, R. P. and McNicholas, P. D. (2014), "Mixtures of shifted asymmetric Laplace distributions", *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36(6), 1149–1157.
- Gallaughar, M. P. B. and McNicholas, P. D. (2018a), "Finite mixtures of skewed matrix variate distributions", *Pattern Recognition* 80, 83–93.
- Gallaughar, M. P. B. and McNicholas, P. D. (2018b), "Three skewed matrix variate distributions", *Statistics and Probability Letters*. To appear, DOI: 10.1016/j.spl.2018.08.012.
- Gallaughar, M. P. B. and McNicholas, P. D. (2017), "A matrix variate skew- t distribution", *Stat* 6(1), 160–170.
- Ghahramani, Z. and Hinton, G. E. (1997), The EM algorithm for factor analyzers, Technical Report CRG-TR-96-1, University of Toronto, Toronto, Canada.
- Harrar, S. W. and Gupta, A. K. (2008), "On matrix variate skew-normal distributions", *Statistics* 42(2), 179–194.
- Hubert, L. and Arabie, P. Gupta (1985), "Comparing partitions", *Journal of Classification* 2(1), 193–218.
- Karlis, D. and Santourian, A. (2009), "Model-based clustering with non-elliptically contoured distributions", *Statistics and Computing* 19(1), 73–83.
- Lawley, D. N. and Maxwell, A. E. (1962), "Factor analysis as a statistical method", *Journal of the Royal Statistical Society: Series D* 12(3), 209–229.
- LeCun, Y., Bottou, L., Bengio, Y. and Haffner, P. (1998), "Gradient-based learning applied to document recognition", *Proceedings of the IEEE* 86(11), 2278–2324.
- Lee, S. and McLachlan, G. J. (2014), "Finite mixtures of multivariate skew t -distributions: some recent and new results", *Statistics and Computing* 24, 181–202.
- Lee, S. and McLachlan, G. J. (2016), "Finite mixtures of canonical fundamental skew t -distributions: the unification of the restricted and unrestricted skew t -mixture models", *Statistics and Computing* 26, 573–589.
- Lin, T.-I. (2010), "Robust mixture modeling using multivariate skew t distributions", *Statistics and Computing* 20(3), 343–356.

- Lin, T.-I., McNicholas, P. D. and Hsiu, J. H. (2014), “Capturing patterns via parsimonious t mixture models”, *Statistics and Probability Letters* 88, 80–87.
- Lin, T., McLachlan, G. J. and Lee, S. X. (2016), “Extending mixtures of factor models using the restricted multivariate skew-normal distribution”, *Journal of Multivariate Analysis* 143, 398–413.
- Lindsay, B. G. (1995), Mixture models: Theory, geometry and applications, in “NSF-CBMS Regional Conference Series in Probability and Statistics”, Vol. 5, Hayward, California: Institute of Mathematical Statistics.
- McLachlan, G. J. and Peel, D. (2000), Mixtures of factor analyzers, in “Proceedings of the Seventh International Conference on Machine Learning”, Morgan Kaufmann, San Francisco, pp. 599–606.
- McNicholas, P. D. (2010), “Model-based classification using latent Gaussian mixture models”, *Journal of Statistical Planning and Inference* 140(5), 1175–1181.
- McNicholas, P. D. (2016a), *Mixture Model-Based Classification*, Chapman & Hall/CRC Press, Boca Raton.
- McNicholas, P. D. (2016b), “Model-based clustering”, *Journal of Classification* 33(3), 331–373.
- McNicholas, P. D. and Murphy, T. B. (2008), “Parsimonious Gaussian mixture models”, *Statistics and Computing* 18(3), 285–296.
- McNicholas, P. D., Murphy, T. B., McDaid, A. F. and Frost, D. (2010), “Serial and parallel implementations of model-based clustering via parsimonious Gaussian mixture models”, *Computational Statistics and Data Analysis* 54(3), 711–723.
- McNicholas, S. M., McNicholas, P. D. and Browne, R. P. (2017), A mixture of variance-gamma factor analyzers, in S. E. Ahmed, ed., “Big and Complex Data Analysis: Methodologies and Applications”, Springer International Publishing, Cham, pp. 369–385.
- Meng, X.-L. and van Dyk, D. (1997), “The EM algorithm — an old folk song sung to a fast new tune (with discussion)”, *Journal of the Royal Statistical Society: Series B* 59(3), 511–567.
- Morris, K. and McNicholas, P. D. (2013a), “Dimension reduction for model-based clustering via mixtures of multivariate t-distributions”, *Advances in Data Analysis and Classification* 7(3), 321–338.
- Morris, K. and McNicholas, P. D. (2013b), “Dimension reduction for model-based clustering via mixtures of shifted asymmetric Laplace distributions”, *Statistics and Probability Letters* 83(9), 2088–2093.
- Murray, P. M., Browne, R. B. and McNicholas, P. D. (2014), “Mixtures of skew-t factor analyzers”, *Computational Statistics and Data Analysis* 77, 326–335.
- Murray, P. M., Browne, R. B. and McNicholas, P. D. (2017), “A mixture of SDB skew-t factor analyzers”, *Econometrics and Statistics* 3, 160–168.
- Murray, P. M., McNicholas, P. D. and Browne, R. B. (2014), “A mixture of common skew-t factor analyzers”, *Stat* 3(1), 68–82.
- O’Hagan, A., Murphy, T. B., Gormley, I. C., McNicholas, P. D. and Karlis, D. (2016), “Clustering with the multivariate normal inverse Gaussian distribution”, *Computational Statistics and Data Analysis* 93, 18–30.
- Peel, D. and McLachlan, G. J. (2000), “Robust mixture modelling using the t distribution”, *Statistics and Computing* 10(4), 339–348.
- Rand, W. M. (1971), “Objective criteria for the evaluation of clustering methods”, *Journal of the American Statistical Association* 66(336), 846–850.
- Schwarz, G. (1978), “Estimating the dimension of a model”, *The Annals of Statistics* 6(2), 461–464.
- Scott, A. J. and Symons, M. J. (1971), “Clustering methods based on likelihood ratio criteria”, *Biometrics* 27, 387–397.
- Steane, M. A., McNicholas, P. D. and Yada, R. (2012), “Model-based classification via mixtures of multivariate t-factor analyzers”, *Communications in Statistics – Simulation and Computation* 41(4), 510–523.
- Subedi, S. and McNicholas, P. D. (2014), “Variational Bayes approximations for clustering via mixtures of normal inverse Gaussian distributions”, *Advances in Data Analysis and Classification* 8(2), 167–193.

- Tang, Y, Browne, R. P. and McNicholas, P. D. (2018), “Flexible clustering of high-dimensional data via mixtures of joint generalized hyperbolic distributions”, *Stat* 7(1), e177.
- Tiedeman, D. V. (1955), On the study of types, in S. B. Sells, ed., “Symposium on Pattern Analysis”, Air University, U.S.A.F. School of Aviation Medicine, Randolph Field, Texas.
- Tipping, M. E. and Bishop, C. M. (1999a), “Mixtures of probabilistic principal component analysers”, *Neural Computation* 11(2), 443–482.
- Tipping, M. E. and Bishop, C. M. (1999b), “Probabilistic principal component analysers”, *Journal of the Royal Statistical Society. Series B* 61, 611–622.
- Tortora, C., McNicholas, P. D. and Browne, R. P. (2016), “A mixture of generalized hyperbolic factor analyzers”, *Advances in Data Analysis and Classification* 10(4), 423–440.
- Viroli, C. (2011), “Finite mixtures of matrix normal distributions for classifying three-way data”, *Statistics and Computing* 21(4), 511–522.
- Vrbik, I. and McNicholas, P. D. (2012), “Analytic calculations for the EM algorithm for multivariate skew-t mixture models”, *Statistics and Probability Letters* 82(6), 1169–1174.
- Vrbik, I. and McNicholas, P. D. (2014), “Parsimonious skew mixture models for model-based clustering and classification”, *Computational Statistics and Data Analysis* 71, 196–210.
- Waddell, A. R. and Oldford, R. W. (2013), *RnavGraphImageData: Some image data used in the RnavGraph package demos*. R package version 0.0.3.
- Wei, Y., Tang, Y. and McNicholas, P. D. (2019), “Mixtures of generalized hyperbolic distributions and mixtures of skew-t distributions for model-based clustering with incomplete data”, *Computational Statistics and Data Analysis* 130, 18–41.
- Wolfe, J. H. (1965), A computer program for the maximum likelihood analysis of types, Technical Bulletin 65-15, U.S. Naval Personnel Research Activity.
- Xie, X., Yan, S., Kwok, J. T. and Huang, T. S. (2008), “Matrix-variate factor analysis and its applications”, *IEEE Transactions on Neural Networks* 19(10), 1821–1826.
- Yu, S., Bi, J. and Ye, J. (2008), Probabilistic interpretations and extensions for a family of 2D PCA-style algorithms, in “Workshop Data Mining Using Matrices and Tensors (DMMT ‘08): Proceedings of a Workshop held in Conjunction with the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (SIGKDD 2008)”.
- Zhao, J., Philip, L. and Kwok, J. T. (2012), “Bilinear probabilistic principal component analysis”, *IEEE Transactions on Neural Networks and Learning Systems* 23(3), 492–503.