

2019

Classification with the matrix-variate-t distribution

Geoffrey Z. Thompson
Iowa State University, gzt@iastate.edu

Ranjan Maitra
Iowa State University, maitra@iastate.edu

William Q. Meeker
Iowa State University, wqmeeker@iastate.edu

Ashraf Bastawros
Iowa State University, bastaw@iastate.edu

Follow this and additional works at: https://lib.dr.iastate.edu/aere_pubs



Part of the [Multivariate Analysis Commons](#), and the [Statistical Methodology Commons](#)

The complete bibliographic information for this item can be found at https://lib.dr.iastate.edu/aere_pubs/148. For information on how to cite this item, please visit <http://lib.dr.iastate.edu/howtocite.html>.

This Article is brought to you for free and open access by the Aerospace Engineering at Iowa State University Digital Repository. It has been accepted for inclusion in Aerospace Engineering Publications by an authorized administrator of Iowa State University Digital Repository. For more information, please contact digirep@iastate.edu.

Classification with the matrix-variate-t distribution

Abstract

Matrix-variate distributions can intuitively model the dependence structure of matrix-valued observations that arise in applications with multivariate time series, spatio-temporal or repeated measures. This paper develops an Expectation-Maximization algorithm for discriminant analysis and classification with matrix-variate t-distributions. The methodology shows promise on simulated datasets or when applied to the forensic matching of fractured surfaces or the classification of functional Magnetic Resonance, satellite or hand gestures images.

Disciplines

Multivariate Analysis | Statistical Methodology

Comments

This is a pre-print of the article Thompson, Geoffrey Z., Ranjan Maitra, William Q. Meeker, and Ashraf Bastawros. "Classification with the matrix-variate- t distribution." *arXiv preprint arXiv:1907.09565* (2019).
Posted with permission.

Classification with the matrix-variate- t distribution

Geoffrey Z. Thompson, Ranjan Maitra, William Q. Meeker and Ashraf Bastawros

Abstract

Matrix-variate distributions can intuitively model the dependence structure of matrix-valued observations that arise in applications with multivariate time series, spatio-temporal or repeated measures. This paper develops an Expectation-Maximization algorithm for discriminant analysis and classification with matrix-variate t -distributions. The methodology shows promise on simulated datasets or when applied to the forensic matching of fractured surfaces or the classification of functional Magnetic Resonance, satellite or hand gestures images.

Index Terms

BIC, ECME, fMRI, fracture mechanics, LANDSAT, supervised learning

I. INTRODUCTION

Matrix-variate distributions [1] can conveniently model matrix-valued observations that arise, for instance, with multivariate time series or spatial datasets or when we observe p -variate responses at q different settings, yielding a $p \times q$ matrix of responses for each unit of observation. The matrix-variate normal distribution (abbreviated in this paper by MxVN) is also helpful for inference, but is sometimes inadequate for modeling populations where the matrix-variate- t distribution (henceforth MxVt) may be a better fit.

There exist discriminant analysis and classification methods for MxVN [2], [3] mixtures but for many applications, the MxVt distribution may model each group better. However, parameter estimation for the MxVt distribution requires special care because, unlike in the normal case, it can not be viewed [4] as simply a rearrangement of its vector-multivariate cousin for which several variants [5]–[7] of the Expectation-Maximization (EM) algorithm exist.

This paper develops, in Section II, methodology for parameter estimation in the MxVt distribution and extends it to discriminant analysis and classification using MxVt mixtures. Our methods are evaluated on simulated and real-life datasets in Section III. This paper concludes with some discussion. An online supplement explicitly detailing the derivations of our algorithm, with sections referenced using the prefix “S-”, and an R [8] package `MixMatrix` that implements the methodology are also included.

II. METHODOLOGY

A. Background and Preliminary Development

1) The Matrix-variate Normal Distribution:

Definition 1. A random matrix \mathbf{X} of p rows and q columns has the MxVN distribution with parameters \mathbf{M} , $\mathbf{\Sigma}$ and $\mathbf{\Omega}$ if it has the probability density function (PDF)

$$f(\mathbf{X}; \mathbf{M}, \mathbf{\Sigma}, \mathbf{\Omega}) = \frac{\exp\left(-\frac{1}{2} \text{tr}\left[\mathbf{\Omega}^{-1}(\mathbf{X} - \mathbf{M})^T \mathbf{\Sigma}^{-1}(\mathbf{X} - \mathbf{M})\right]\right)}{(2\pi)^{pq/2} |\mathbf{\Omega}|^{p/2} |\mathbf{\Sigma}|^{q/2}},$$

where \mathbf{M} is a $p \times q$ matrix that is the mean of \mathbf{X} , and $\mathbf{\Sigma}$ and $\mathbf{\Omega} > 0$ describing the covariances between, respectively, each of the p rows and the q columns of \mathbf{X} . We write $\mathbf{X} \sim \mathcal{N}_{p,q}(\mathbf{M}, \mathbf{\Sigma}, \mathbf{\Omega})$. For identifiability, we set the first element of $\mathbf{\Sigma}$ to be unity.

The MxVN distribution can be considered, after rearranging into a vector (denoted by $\text{vec}(\mathbf{X})$), to be from a multivariate normal (MVN) distribution with a Kronecker covariance structure [1]. So, if $\mathbf{X} \sim \mathcal{N}_{p,q}(\mathbf{M}, \mathbf{\Sigma}, \mathbf{\Omega})$, then $\text{vec}(\mathbf{X}) \sim \mathcal{N}_{pq}(\text{vec}(\mathbf{M}), \mathbf{\Sigma} \otimes \mathbf{\Omega})$. This reformulation allows us to readily obtain the maximum likelihood (ML) estimates. If \mathbf{X}_i , $i = 1, 2, \dots, n$ are independent identically distributed (IID) random matrices from the $\mathcal{N}_{p,q}(\mathbf{M}, \mathbf{\Sigma}, \mathbf{\Omega})$, then \mathbf{M} , $\mathbf{\Sigma}$, and $\mathbf{\Omega}$ have ML estimators $\widehat{\mathbf{M}} = n^{-1} \sum_{i=1}^n \mathbf{X}_i$, $\widehat{\mathbf{\Sigma}} = (np)^{-1} \sum_{i=1}^n (\mathbf{X}_i - \widehat{\mathbf{M}}) \widehat{\mathbf{\Omega}}^{-1} (\mathbf{X}_i - \widehat{\mathbf{M}})^T$, and $\widehat{\mathbf{\Omega}} = (nq)^{-1} \sum_{i=1}^n (\mathbf{X}_i - \widehat{\mathbf{M}})^T \widehat{\mathbf{\Sigma}}^{-1} (\mathbf{X}_i - \widehat{\mathbf{M}})$ under the constraint that the first element of $\widehat{\mathbf{\Sigma}}$ is set to unity. The matrices $\widehat{\mathbf{\Sigma}}$ and $\widehat{\mathbf{\Omega}}$ are obtained iteratively after initialization with the identity matrix \mathbf{I} . The ML estimates are unique if $n > \max(p, q)$ [9].

The authors are with Iowa State University. Email: {gzt,maitra,wqmeeker,bastaw}@iastate.edu.

This research was supported in part by the National Institute of Justice (NIJ) under Grants No. 2015-DN-BX-K056 and 2018-R2-CX-0034. The research of the second author was also supported in part by the National Institute of Biomedical Imaging and Bioengineering (NIBIB) of the National Institutes of Health (NIH) under Grant R21EB016212, and the United States Department of Agriculture (USDA) National Institute of Food and Agriculture (NIFA) Hatch project IOW03617. The content of this paper is however solely the responsibility of the authors and does not represent the official views of the NIJ, the NIBIB, the NIH, the NIFA or the USDA.

2) The Matrix-variate t -distribution:

Definition 2. A random $p \times q$ matrix \mathbf{X} is MxVt-distributed with parameters $(\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ of similar order as in Definition 1 (with $\boldsymbol{\Sigma}$ and $\boldsymbol{\Omega} > 0$) and degrees of freedom (df) $\nu \geq 1$ if its PDF is

$$f(\mathbf{X}; \nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}) = \frac{\Gamma_p\left(\frac{\nu+p+q-1}{2}\right)}{(\pi)^{\frac{pq}{2}} \Gamma_p\left(\frac{\nu+p-1}{2}\right)} |\boldsymbol{\Omega}|^{-\frac{p}{2}} |\boldsymbol{\Sigma}|^{-\frac{q}{2}} \left| \mathbf{I}_p + \boldsymbol{\Sigma}^{-1}(\mathbf{X} - \mathbf{M})\boldsymbol{\Omega}^{-1}(\mathbf{X} - \mathbf{M})^T \right|^{-\frac{\nu+p+q-1}{2}}.$$

We use the notation $\mathbf{X} \sim t_{p,q}(\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ to indicate that \mathbf{X} has this density.

a) *Properties:* We mention some properties of the MxVt distribution relevant to this paper.

- 1) For $p = 1$ and $\boldsymbol{\Sigma} \equiv \nu$ (or $q = 1$ and $\boldsymbol{\Omega} \equiv \nu$), the MxVt distribution reduces to its vector-multivariate t (MVT) cousin. However, this reduction does not generally hold so additional development is needed for inference. We provide methods to do so in the next section.
- 2) Let the random matrix $\mathbf{S} \sim \mathcal{W}_p(\nu + p - 1, \boldsymbol{\Sigma}^{-1})$, where $\mathcal{W}_p(\kappa, \boldsymbol{\Psi})$ is the $p \times p$ -dimensional Wishart distribution with d.f. κ and scale matrix $\boldsymbol{\Psi}$. If $\mathbf{X} | \mathbf{S} \sim \mathcal{N}_{p,q}(\mathbf{M}, \mathbf{S}^{-1}, \boldsymbol{\Omega})$, then $\mathbf{X} \sim t_{p,q}(\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$. Further, $\mathbf{S} | \mathbf{X} \sim \mathcal{W}_p(\nu + p + q - 1, [(\mathbf{X} - \mathbf{M})\boldsymbol{\Omega}^{-1}(\mathbf{X} - \mathbf{M})^T + \boldsymbol{\Sigma}]^{-1})$.

B. ML Estimation of the MxVt parameters

The MxVt distribution does not have closed-form ML estimators so we provide an Expectation/Conditional Maximization Either (ECME) algorithm [7] in a manner that is similar to that used to find ML parameter estimates in the MVT distribution.

Let $\mathbf{X}_i, i = 1, 2, \dots, n$ be IID realizations from $t_{p,q}(\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$. Write $\boldsymbol{\Theta} \equiv \{\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}\}$. For each $i = 1, 2, \dots, n$, let \mathbf{S}_i be (unobserved) random matrices that are as per Property 2. From the detailed development and derivations provided in the supplement, we get the expectation step (E-step) updates at the current value $\boldsymbol{\Theta}^{(t)}$ of $\boldsymbol{\Theta}$:

$$\begin{aligned} \mathbf{S}_i^{(t+1)} &\doteq \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\mathbf{S}_i | \mathbf{X}_i) = (\nu^{(t)} + p + q - 1)[(\mathbf{X}_i - \mathbf{M}^{(t)})\boldsymbol{\Omega}^{(t)-1}(\mathbf{X}_i - \mathbf{M}^{(t)})^T + \boldsymbol{\Sigma}^{(t)}]^{-1} \\ \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\log |\mathbf{S}_i| | \mathbf{X}_i) &= \psi_p\left(\frac{\nu^{(t)} + p + q - 1}{2}\right) + p \log 2 + \log \left| \frac{\mathbf{S}_i^{(t+1)}}{\nu^{(t)} + p + q - 1} \right| \end{aligned}$$

where $\psi_p(\cdot)$ is the p -variate digamma function – that is, $\psi_p(x) = \frac{d}{dx} \log \Gamma_p(x)$. Further computational reductions are possible by defining and storing the updates in terms of the sufficient statistics

$$\begin{aligned} \mathbf{S}_S^{(t+1)} &\doteq \sum_{i=1}^n \mathbf{S}_i^{(t+1)}, \\ \mathbf{S}_{SX}^{(t+1)} &\doteq \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\mathbf{S}_i \mathbf{X}_i | \mathbf{X}_i) = \sum_{i=1}^n \mathbf{S}_i^{(t+1)} \mathbf{X}_i, \\ \mathbf{S}_{XSX}^{(t+1)} &\doteq \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\mathbf{X}_i^T \mathbf{S}_i \mathbf{X}_i | \mathbf{X}_i) = \sum_{i=1}^n \mathbf{X}_i^T \mathbf{S}_i^{(t+1)} \mathbf{X}_i, \\ \mathbf{S}_{|S|}^{(t+1)} &\doteq \mathbb{E}_{\boldsymbol{\Theta}^{(t)}} \left[\sum_{i=1}^n \log |\mathbf{S}_i| | \mathbf{X}_i \right]. \end{aligned}$$

These statistics can be expressed with $(\nu^{(t)} + p + q - 1)$ factored out, and for convenience may be computed and stored as such when ν needs to be estimated.

The M-step updates are conceptually immediate and maximize Equation (S-5) in the supplement with respect to $\boldsymbol{\Theta}$ to yield $\boldsymbol{\Theta}^{(t+1)}$. The EM algorithm can, however, be sped up by maximizing $(\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ at the current value of ν and vice versa. We are led to an Expectation/Conditional Maximization Either (ECME) algorithm, with updates $\hat{\boldsymbol{\Theta}}^{(t+1)}$ conditional on $\nu^{(t)}$ as

$$\begin{aligned} \widehat{\mathbf{M}}^{(t+1)} &= \left(\sum_{i=1}^n \mathbf{S}_i^{(t+1)} \right)^{-1} \sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i = \mathbf{S}_S^{(t+1)-1} \mathbf{S}_{SX}^{(t+1)} \\ \widehat{\boldsymbol{\Omega}}^{(t+1)} &= \frac{1}{np} \sum_{i=1}^n (\mathbf{X}_i - \widehat{\mathbf{M}}^{(t)})^T \mathbf{S}_i^{(t+1)} (\mathbf{X}_i - \widehat{\mathbf{M}}^{(t)}) = \frac{1}{np} \left(\mathbf{S}_{XSX}^{(t+1)} - \mathbf{S}_{SX}^{(t+1)T} \mathbf{S}_S^{(t+1)-1} \mathbf{S}_{SX}^{(t+1)} \right) \\ \widehat{\boldsymbol{\Sigma}}_{(t+1)}^{-1} &= \frac{1}{n(\nu^{(t)} + p - 1)} \sum_{i=1}^n \mathbf{S}_i^{(t+1)} = \frac{\mathbf{S}_S^{(t+1)}}{n(\nu^{(t)} + p - 1)}, \end{aligned}$$

but the conditional maximization of ν given $(\mathbf{M}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \boldsymbol{\Omega}^{(t+1)})$ can be sped up substantially by maximizing it instead over the observed log-likelihood function given $(\mathbf{M}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \boldsymbol{\Omega}^{(t+1)})$. We get the ML estimating equation:

$$n \frac{d}{d\nu} \log \Gamma_p((\nu + p - 1)/2) - \frac{1}{2} (\mathbf{S}_{|S|} - np \log 2 + n \log |\widehat{\boldsymbol{\Sigma}}|) = 0. \quad (1)$$

Writing $\kappa = \nu + p + q - 1$ for notational compactness, we have,

$$\begin{aligned} 0 &= n\psi_p((\nu + p - 1)/2) - \left\{ n\psi_p\left(\frac{\kappa}{2}\right) + \sum_{i=1}^n \log \left| \frac{\mathbf{S}_i^{(t+1)}}{\kappa} \right| - n \log \left| \frac{\mathbf{S}_S^{(t+1)}}{n(\nu + p - 1)} \right| \right\} \\ &= \psi_p((\nu + p - 1)/2) - \left\{ \psi_p\left(\frac{\kappa}{2}\right) + \frac{1}{n} \sum_{i=1}^n \log \left| \mathbf{Z}_i^{(t+1)} \right| + p \log \frac{n(\nu + p - 1)}{\kappa} - \log \left| \mathbf{Z}_S^{(t+1)} \right| \right\}, \end{aligned} \quad (2)$$

where $\mathbf{Z}_*^{(t+1)}$ is the appropriate $\mathbf{S}_*^{(t+1)}$ statistic with $(\nu^{(t)} + p + q - 1)$ factored out. The ML estimating equation can be solved using a one-dimensional search, yielding a ECME algorithm with the steps:

1. **E-step:** Update \mathbf{S}_i weights and statistics based on $\Theta^{(t)}$ and \mathbf{X}_i .
2. **CME-step:** Update $\Theta_1^{(t+1)} = (\mathbf{M}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \boldsymbol{\Omega}^{(t+1)})$.
3. **CME-step:** Update $\Theta_2^{(t+1)} = \nu^{(t+1)}$ using the observed log-likelihood given $(\mathbf{M}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \boldsymbol{\Omega}^{(t+1)})$.

Repeat these steps until convergence.

We conclude here by noting that restrictions on the parameter set [10], such as imposing an m th-order auto-regressive structure (AR(m)) on either $\boldsymbol{\Sigma}$ or $\boldsymbol{\Omega}$ or both as in our applications in Section III can be easily incorporated within our algorithm (see Section S-5.22).

C. Discrimination and Classification

Linear (LDA) and Quadratic Discriminant Analysis (QDA) for matrix-variate populations follow a similar approach as for the multivariate case, with the MxVN (but not MxVt) cases affording substantial reductions in the computations. We provide here the general framework for matrix-variate distributions and then discuss reductions for the cases of the MxVN models.

Suppose that there are two populations π_1 and π_2 , with prior probabilities η_1 and η_2 for an observation belonging to either. Let $\mathbb{P}(1|2)$ be the probability of classifying a member of π_2 to π_1 (and vice versa). As usual, the *total probability of misclassification* (TPM) is defined to be $\mathbb{P}(2|1)\eta_1 + \mathbb{P}(1|2)\eta_2$. A Bayes optimal classification rule that minimizes the TPM assigns a matrix-valued observation \mathbf{X} to π_1 if $\frac{f_1(\mathbf{X})}{f_2(\mathbf{X})} \geq \frac{\eta_2}{\eta_1}$, where $f_i(\mathbf{X})$ is the PDF for group π_i evaluated at \mathbf{X} [11]. The classification rule can be easily extended to the case when there are G groups $\pi_1, \pi_2, \dots, \pi_G$, each with prior probabilities of membership $\eta_1, \eta_2, \dots, \eta_G$ and densities f_1, f_2, \dots, f_G . Then the Bayes optimal classification for a matrix-valued observation \mathbf{X} is $\operatorname{argmax}_{i \in \{1, 2, \dots, G\}} R_i(\mathbf{X})$, where the cost function $R_i(\mathbf{X})$ is defined as $\log \eta_i f_i(\mathbf{X})$.

Unlike for the MxVt distributions, the MxVN case has closed-form solutions analogous to that of LDA or QDA in multivariate statistics. For the MxVN populations, the closed-form classification rule assigns \mathbf{X} to the g th group where $g = \operatorname{argmax}_{i=1, 2, \dots, G} R_i(\mathbf{X})$, with

$$\begin{aligned} R_i(\mathbf{X}) &= \operatorname{trace} \left\{ -\frac{1}{2} (\boldsymbol{\Omega}_i^{-1} \mathbf{X}^T \boldsymbol{\Sigma}_i^{-1} \mathbf{X}) + \boldsymbol{\Omega}_i^{-1} \mathbf{M}_i^T \boldsymbol{\Sigma}_i^{-1} \mathbf{X} - \frac{1}{2} \boldsymbol{\Omega}_i^{-1} \mathbf{M}_i^T \boldsymbol{\Sigma}_i^{-1} \mathbf{M}_i \right\} \\ &\quad - \frac{1}{2} (p \log |\boldsymbol{\Sigma}_i| + q \log |\boldsymbol{\Omega}_i|). \end{aligned} \quad (3)$$

The first and last term disappear when the G MxVN populations have common covariances, yielding a linear decision rule. Many adaptations [12]–[21] of LDA exist for homogeneous MxVN populations, but our development provides a natural and direct approach that is also flexible enough to include a range of assumptions. Assuming homogeneity does not yield a linear rule for MxVT populations where we still get a quadratic rule. Finally, in all cases, the parameters in $R(\mathbf{X})$ can be estimated using ML on the training set (with the ECME methodology of Section II-B for MxVt populations) and incorporated into the decision rule.

III. PERFORMANCE EVALUATIONS

This section evaluates performance of the ECME algorithm in recovering the MxVt parameters and also classification performance of our methodology on some real-life datasets.

A. Simulation Study

Our simulation study generated 200 datasets from the $t_{5,3}(\nu, \mathbf{M}, \Sigma, \Omega)$ distribution with $\nu = 5, 10, 20$ and $n \in \{35, 50, 100\}$, with the smallest n chosen to be large enough for simulations to successfully converge and the larger sample sizes chosen to give an idea of consistency of parameter estimation. The ECME algorithm in Section II-B, with unconstrained $(\mathbf{M}, \Sigma, \Omega)$ was used to estimate the parameters. Figure 1 summarizes the estimated $\hat{\nu}$ over the 200 samples for each ν . (We constrain $\hat{\nu}$ to be in $(2, 1000)$ in the Either step of the ECME algorithm.)

As expected, higher n improves both accuracy and precision of the estimates. For all nine cases, the peak of the distribution of $\hat{\nu}$ was close to the true ν value. Lower values for ν were more easily estimated in the sense that for any n , the $\hat{\nu}$ values are closer to the true ν (Figure 1). This may be because for larger true ν , the distributions are similar in a wider range (after all, as $\nu \rightarrow \infty$, the distribution reduces to the MxVN). For one aberrant sample with $\nu = 20$ and $n = 35$, the optimizer attained the upper bound and did not converge to an interior point. On the whole, however, our simulation results indicate good performance of the ECME algorithm in recovering the MxVt parameters.


B. Classification Examples

We evaluate MxVt classification and discrimination on datasets from four different applications.

1) *Matching Fractured Surfaces*: Our first example is on the potential ability of our classification algorithm to distinguish between pairs of fractured surfaces into matches or non-matches, with implications in forensics to decide on, say, whether a knife blade fragment found at a crime scene is a match to something that visually appears to be the remainder of the blade. Because of the novelty of this application, we discuss it at some length here. Our investigation is a formal proof-of-concept conducted in the lab where a set of 38 stainless steel knives had their blades broken under similar conditions, resulting in each of them having a base and a tip. The cross-sectional fractured surfaces were then scanned using a standard non-contact 3D optical interferometer at 9 regularly-spaced locations to get 9 successive 1024×1024 images (with 75% overlap, in order to get a reasonable number of replications while also imaging the entire length of the exposed surfaces).

Cross-correlations between matching knife base-tip image pairs in the $5\text{-}10 \mu\text{m}^{-1}$ and $10\text{-}20 \mu\text{m}^{-1}$ two-dimensional (2D) Fourier frequencies were computed, yielding, for each knife, a 2×9 matrix of measurements describing the similarity of the base of the knife to its tip (2 measured cross-correlations per image and 9 images). Similar cross-correlations between all possible knife base-tip pairs (regardless of origin) yielded a sample from the population of similarity matrices coming from known matching (KM) and known non-matching (KNM) base-tip pairs.

Figure 2 shows the scatterplot of the Fisher's Z -transformed [22] cross-correlation data to be fairly elliptical. The two classes are almost but not completely separated when looking at individual image pairs. Classification using



ν	n	Range	Median	Mean	SD
5	35	(2.62, 15.12)	5.24	5.45	1.43
	50	(3.46, 12.78)	5.32	5.45	1.28
	100	(3.76, 7.01)	5.14	5.18	0.56
10	35	(5.41, 395.12)	11.57	16.28	28.26
	50	(5.28, 106.72)	10.44	11.94	7.91
	100	(6.99, 18.10)	10.19	10.56	1.84
20	35	(9.93, 999.83)	29.94	89.01	149.12
	50	(11.38, 495.67)	24.45	46.97	73.09
	100	(12.68, 147.67)	21.98	25.52	14.91

(b)

Fig. 1: (a) Density plots and (b) numerical summaries of $\hat{\nu}$ for datasets of size $n = 35, 50, 100$ with true $\nu = 5, 10, 20$ (vertical line).




Fig. 2: Cross-correlations for individual images, along with 99% confidence ellipses under bivariate normal assumptions, in known matching (KM) and known non-matching (KNM) surfaces on Fisher-transformed axes. Known matches and known non-matches can be distinguished, but not perfectly, in this example by features in these two frequency ranges.

only one pair of images per surface rather than a set of multiple images is potentially ambiguous. We remove this potential ambiguity by considering multiple images on each surface. These multiple sets of images on each knife are not independent and have a natural multivariate repeated measures (*i.e.* matrix-variate) structure because of the 75% overlap between successive images so a model incorporating such structure may improve classification accuracy.

We model each match/non-match dataset in terms of the $MxVt$ distribution with group-specific mean matrix and matrix dispersion structures, with an AR(1) correlation structure for the Fourier domain correlations at the same frequency band between successive (overlapping) image pairs. The AR(1) structure is appropriate because of the overlap between successive images: this correlation structure also has the best Bayesian Information Criterion (BIC) among the correlation structures tested on the data [23]. The mean across images for each frequency band was constrained to be constant. Because there are only 9-10 observations for the cases where the knife tip-base have the same origin, we forgo estimating ν and instead investigate classification with the $MxVt$ distribution with $\nu = 5$ and $\nu = 10$ (in addition to the $MxVN$).

Figure 3 displays the distribution of the log-odds of being a match for the models based on each of the four training sets. The models trained on each set were then tested on the data from all four sets of surfaces. In this figure, positive log-odds indicates a higher probability of being a KM and negative log-odds indicates a higher probability of being a KNM.

With equal priors, there is a 0% false exclusion (false negative) rate and a 0.003% false identification (false positive) rate (1 FP). The only FP is from the $MxVN$ model, which is also overly confident about the matches it produces. It predicts some surfaces being a match with log-odds greater than 200, which is extremely implausible. The $MxVt$ distribution accounts for uncertainty better and results in more plausible log-odds ratios. Perfect discrimination is attained with $MxVt$ for all four training sets, suggesting that the results generalize well to out-of-sample data despite the relatively small sample size.

For comparison, we also obtained predictions using the penalized likelihood approach of [21] which works only when at least two sets of knives are used as training sets, two sets as a validation set for the tuning parameters, and the rest as the test set. We were able to obtain perfect classification for all permutations of the six sets when an appropriate grid of tuning parameters is used (two additional sets of images taken from one set of knives were used to make a total of six sets). However, the method forced the covariance matrices to be diagonal, which is unlikely to be reasonable given the 75% overlap between successive images.

2) *Finger-tapping Experiment*: [24] provided 12 functional Magnetic Resonance Imaging (fMRI) scans of the brain of a right-hand dominant male subject during a right-hand finger-thumb opposition activity and 12 similar scans using the left-hand, with each pair of scan collected at regular intervals over a 2-month period. We restrict attention to the 20th slice of the image volume, with 128×128 pixels, that previous work [25], [26] indicated as adequate to distinguish activation between the left- and right-hand finger-tapping. With only 12 observations per class, we are limited in the types of correlation matrices that we may consider, so we selected a 20×20 section of the 20th slice having the left-topmost pixel at (33, 67), which was the 20×20 section of the slice displaying the highest average activation in the left-hand activation images as determined by [27]’s FAST-fMRI algorithm. We then trained and tested the classifiers using the leave-one-out method with an AR(1) covariance structure and a compound symmetry covariance structure in the $MxVN$ and $MxVt$ distributions with $\nu = 5$ or $\nu = 10$ (for the $MxVt$). The BIC on the fitted models indicated that a compound symmetry covariance structure was the best model. In all cases, except that of the $MxVt$ distribution with $\nu = 10$ and an AR(1) covariance structure, 23 out of 24 images were correctly classified. The one mislabeled case was the same one that was previously identified by [26] as an outlier. Using the $MxVt$ distribution with $\nu = 10$ had one more misclassification. The number of cases for this reduced dataset is not enough for `MatrixLDA` to estimate the correlation structure so we forgo that comparison here.

3) *Landsat Satellite Data*: Multi-spectral satellite imagery allows for multiple observations over a spatial grid, yielding matrix-valued observations. We examine a set of satellite images [28] that are in two visible and two infrared bands. The subset of images under consideration [29] consists of a training and a test set of 3×3 pixel segments labeled according to the terrain type (gray soil, damp gray soil or soil with vegetation stubble) of their middle pixel. Regarding the data as a 4×9 matrix and with a $MxVN$ classifier and unconstrained covariance matrices yielded an error rate of 0.116 [29], while `MatrixLDA` with tuning parameters selected by 5-fold CV [21] yielded a 0.118 error rate. Our $MxVN$ and $MxVt$ models (the latter with $\nu = 10$ and 20) with unconstrained covariance matrices and prior probabilities equal to the class representation in the training set

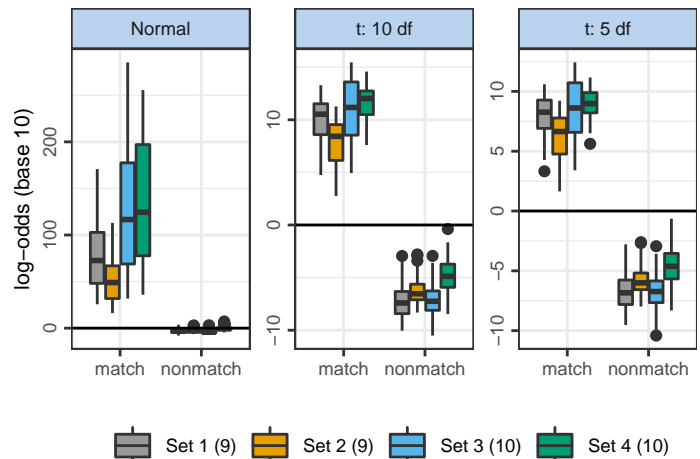


Fig. 3: Positive results indicate match is more probable than non-match. There are 9 known matches and 72 known non-matches in Set 1 and Set 2 and 10 known matches and 90 known non-matches in Sets 3 and 4. The results are from training on the indicated set and testing on all four sets.

yielded error rates of 0.126, 0.116, and 0.109, in line with previous results. BIC indicated that using unconstrained covariance matrices and means constrained to be equal within rows as a better model, with error rates of 0.123, 0.121, and 0.107.

4) *Cambridge Hand Gestures Data*: We tested our method using leave-one-out cross-validation (LOOCV) on the set of 80 images extracted from the Cambridge hand gestures database [30] as processed by [21] into 80×60 pixel gray-scale images. There are four classes in this problem: the images show a hand gesture in one of two shapes and one of two orientations: in each image, the hand is either in a flat or “V” shape and is located either in the center of the image or to the left side of the image. Using an AR(1) covariance structure on both dimensions, which was suggested using BIC to determine the best structure, and a $MxVt$ distributions with 5 and 10 degrees of freedom, we were able to obtain a 100% classification rate on the dataset. [21] report a 90% correct classification rate using LOOCV on this dataset.

IV. CONCLUSIONS

We have provided an ECME method for fitting the parameters of the $MxVt$ distribution that can be used on three-way data sets such as multivariate repeated measures, image or spatial data, and have demonstrated the method on simulation datasets and on classification and discrimination in four real-world applications where the new method using the $MxVt$ -distribution outperforms that using the $MxVN$. The ECME algorithm and the discriminant analysis are implemented in the R package `MixMatrix`. The package also includes functions for sampling from and computing the density of the $MxVN$ and $MxVt$ distributions and includes the datasets used in this paper.

Our model can be extended beyond supervised learning to mixture model-based clustering and can be made to accommodate more specialized covariance structures such as those described in [31] and [32]. It may also be readily extended to cases with incomplete records. Determining the convergence and uniqueness properties would also be desirable. For instance, we know how many observations are required to have unique ML estimates of the parameters in the $MxVN$ distribution with unconstrained mean and covariance matrices but such results may be useful to develop for the $MxVt$ or the constrained $MxVN$. Nevertheless, the EM algorithm is guaranteed to converge to a local maximum, provided it is initialized where the log-likelihood function is finite. Another area which could benefit is the extension of `MatrixLDA` to include the $MxVt$ distribution, where we believe our development in this paper will be helpful.

REFERENCES

- [1] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*. CRC Press, 1999, vol. 104.
- [2] C. Viroli, “Model based clustering for three-way data structures,” *Bayesian Analysis*, vol. 6, no. 4, pp. 573–602, 12 2011. [Online]. Available: <https://doi.org/10.1214/11-BA622>
- [3] L. Anderlucci and C. Viroli, “Covariance pattern mixture models for the analysis of multivariate heterogeneous longitudinal data,” *The Annals of Applied Statistics*, vol. 9, no. 2, pp. 777–800, 06 2015. [Online]. Available: <https://doi.org/10.1214/15-AOAS816>
- [4] J. M. Dickey, “Matrix-variate generalizations of the multivariate t distribution and the inverted multivariate t distribution,” *The Annals of Mathematical Statistics*, vol. 38, no. 2, pp. 511–518, 04 1967. [Online]. Available: <https://doi.org/10.1214/aoms/1177698967>
- [5] A. P. Dempster, N. M. Laird, and D. B. Rubin, “Maximum likelihood from incomplete data via the EM algorithm,” *Journal of the Royal Statistical Society. Series B (methodological)*, pp. 1–38, 1977.
- [6] X.-L. Meng and D. B. Rubin, “Maximum likelihood estimation via the ECM algorithm: A general framework,” *Biometrika*, vol. 80, no. 2, pp. 267–278, 1993. [Online]. Available: <http://dx.doi.org/10.1093/biomet/80.2.267>
- [7] C. Liu and D. B. Rubin, “The ECME algorithm: A simple extension of EM and ECM with faster monotone convergence,” *Biometrika*, vol. 81, no. 4, pp. 633–648, 1994. [Online]. Available: <http://www.jstor.org/stable/2337067>
- [8] R Core Team, *R: A Language and Environment for Statistical Computing*, R Foundation for Statistical Computing, Vienna, Austria, 2019. [Online]. Available: <https://www.R-project.org/>
- [9] P. Dutilleul, “The MLE algorithm for the matrix normal distribution,” *Journal of Statistical Computation and Simulation*, vol. 64, no. 2, pp. 105–123, 1999.
- [10] A. Roy and R. Khattree, “On discrimination and classification with multivariate repeated measures data,” *Journal of Statistical Planning and Inference*, vol. 134, no. 2, pp. 462 – 485, 2005. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0378375804002186>
- [11] T. W. Anderson and R. R. Bahadur, “Classification into two multivariate normal distributions with different covariance matrices,” *The Annals of Mathematical Statistics*, vol. 33, no. 2, pp. 420–431, 1962. [Online]. Available: <http://www.jstor.org/stable/2237521>
- [12] K. Inoue and K. Urahama, “Non-iterative two-dimensional linear discriminant analysis,” in *18th International Conference on Pattern Recognition (ICPR’06)*, Aug 2006, pp. 540–543.
- [13] K. Inoue, K. Hara, and K. Urahama, “Non-iterative symmetric two-dimensional linear discriminant analysis,” *IEICE Transactions on Information and Systems*, vol. E94.D, no. 4, pp. 926–929, 2011.
- [14] M. Li and B. Yuan, “2d-LDA: A statistical linear discriminant analysis for image matrix,” *Pattern Recognition Letters*, vol. 26, no. 5, pp. 527–532, Apr. 2005. [Online]. Available: <http://dx.doi.org/10.1016/j.patrec.2004.09.007>
- [15] H. Lu, K. N. Plataniotis, and A. N. Venetsanopoulos, “Uncorrelated multilinear discriminant analysis with regularization and aggregation for tensor object recognition,” *IEEE Transactions on Neural Networks*, vol. 20, no. 1, pp. 103–123, Jan 2009.
- [16] M. S. Mahanta and K. N. Plataniotis, “Ranking 2DLDA features based on Fisher discriminance,” in *2014 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, May 2014, pp. 8307–8311.
- [17] S. Yan, D. Xu, Q. Yang, L. Zhang, X. Tang, and H.-J. Zhang, “Multilinear discriminant analysis for face recognition,” *IEEE Transactions on Image Processing*, vol. 16, no. 1, 2007.
- [18] J. Ye, R. Janardan, and Q. Li, “Two-dimensional linear discriminant analysis,” in *Advances in Neural Information Processing Systems 17 - Proceedings of the 2004 Conference, NIPS 2004*. Neural information processing systems foundation, 1 2005.
- [19] J. Zhao, P. L. Yu, L. Shi, and S. Li, “Separable linear discriminant analysis,” *Computational Statistics and Data Analysis*, vol. 56, no. 12, pp. 4290 – 4300, 2012. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167947312001636>
- [20] W.-S. Zheng, J.-H. Lai, and S. Z. Li, “1D-LDA vs. 2D-LDA: When is vector-based linear discriminant analysis better than matrix-based?” *Pattern Recognition*, vol. 41, no. 7, pp. 2156–2172, 2008.

- [21] A. J. Molstad and A. J. Rothman, "A penalized likelihood method for classification with matrix-valued predictors," *Journal of Computational and Graphical Statistics*, vol. 28, no. 1, pp. 11–22, 2019. [Online]. Available: <https://doi.org/10.1080/10618600.2018.1476249>
- [22] R. A. Fisher, "Frequency distribution of the values of the correlation coefficient in samples from an indefinitely large population," *Biometrika*, vol. 10, no. 4, pp. 507–521, 1915. [Online]. Available: <http://www.jstor.org/stable/2331838>
- [23] G. Schwarz, "Estimating the dimension of a model," *The Annals of Statistics*, vol. 6, no. 2, pp. 461–464, 03 1978. [Online]. Available: <https://doi.org/10.1214/aos/1176344136>
- [24] R. Maitra, S. R. Roys, and R. P. Gullapalli, "Test-retest reliability estimation of functional MRI data," *Magnetic Resonance in Medicine*, vol. 48, no. 1, pp. 62–70, 2002. [Online]. Available: <https://onlinelibrary.wiley.com/doi/abs/10.1002/mrm.10191>
- [25] R. Maitra, "Assessing certainty of activation or inactivation in test-retest fMRI studies," *Neuroimage*, vol. 47, no. 1, pp. 88–97, 2009.
- [26] —, "A re-defined and generalized percent-overlap-of-activation measure for studies of fMRI reproducibility and its use in identifying outlier activation maps," *NeuroImage*, vol. 50, no. 1, pp. 124 – 135, 2010. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1053811909012567>
- [27] I. Almodóvar-Rivera and R. Maitra, "Fast adaptive smoothing and thresholding for improved activation detection in low-signal fMRI," *IEEE Transactions on Medical Imaging*, pp. 1–1, 2019. [Online]. Available: <https://doi.org/10.1109/TMI.2019.2915052>
- [28] D. Dua and C. Graff, "UCI machine learning repository," <http://archive.ics.uci.edu/ml>, 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [29] C. Viroli, "Finite mixtures of matrix normal distributions for classifying three-way data," *Statistics and Computing*, vol. 21, no. 4, pp. 511–522, Oct 2011. [Online]. Available: <https://doi.org/10.1007/s11222-010-9188-x>
- [30] T. Kim, S. Wong, and R. Cipolla, "Tensor canonical correlation analysis for action classification," in *2007 IEEE Conference on Computer Vision and Pattern Recognition*, June 2007, pp. 1–8.
- [31] C. Fraley and A. E. Raftery, "Model-based clustering, discriminant analysis, and density estimation," *Journal of the American Statistical Association*, vol. 97, no. 458, pp. 611–631, 2002. [Online]. Available: <https://www.tandfonline.com/doi/abs/10.1198/016214502760047131>
- [32] J. Andrews, J. Wickins, N. Boers, and P. McNicholas, "teigen: An r package for model-based clustering and classification via the multivariate t distribution," *Journal of Statistical Software, Articles*, vol. 83, no. 7, pp. 1–32, 2018. [Online]. Available: <https://www.jstatsoft.org/v083/i07>
- [33] K. L. Lange, R. J. Little, and J. M. Taylor, "Robust statistical modeling using the t distribution," *Journal of the American Statistical Association*, vol. 84, no. 408, pp. 881–896, 1989.

SUPPLEMENTARY MATERIALS

S-5. THE EM ALGORITHM FOR PARAMETER ESTIMATION IN THE $\text{MxV}t$ DISTRIBUTION

As mentioned in the paper, the $\text{MxV}t$ distribution does not have closed-form ML estimators so we develop an EM algorithm by augmenting the data, in similar spirit as done for the vector-multivariate t -distribution [33], and then present an ECME (Expectation/Conditional Maximization Either) algorithm [7] to improve the speed of convergence of the EM algorithm. Let $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n$ be independent realizations from the $t_{p,q}(\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ density. Then each \mathbf{X}_i can be augmented with latent weight matrices \mathbf{S}_i as follows:

$$\begin{aligned} \mathbf{X}_i | \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}, \nu, \mathbf{S}_i &\sim \mathcal{N}_{p,q}(\mathbf{M}, \mathbf{S}_i^{-1}, \boldsymbol{\Omega}) \\ \mathbf{S}_i | \mathbf{M}, \boldsymbol{\Omega}, \boldsymbol{\Sigma}, \nu &\sim \mathcal{W}_p(\nu + p - 1, \boldsymbol{\Sigma}^{-1}), \quad \text{for } i = 1, 2, \dots, n. \end{aligned} \quad (\text{S-4})$$

To show the benefits of using the latent \mathbf{S}_i s, we first derive ML estimators with the complete data and then use that to derive an EM algorithm using only the observed data. We then modify the EM algorithm to its more efficient ECME derivative.

A. ML Estimation of parameters with complete data

Suppose that we have $(\mathbf{X}_i, \mathbf{S}_i), i = 1, 2, \dots, n$ where each $\mathbf{S}_i \sim \mathcal{W}_p(\nu + p - 1, \boldsymbol{\Sigma}^{-1})$ and $\mathbf{X}_i | \mathbf{S}_i \sim \mathcal{N}_{p,q}(\mathbf{M}, \mathbf{S}_i^{-1}, \boldsymbol{\Omega})$ for each $i = 1, 2, \dots, n$. Then the complete log-likelihood function ℓ_c of the parameters $(\mathbf{M}, \boldsymbol{\Omega})$ given the data $(\mathbf{X}_i, \mathbf{S}_i), i = 1, 2, \dots, n$ can be written as a sum of (conditional) MxVN log-likelihood functions ℓ_N and a sum of Wishart log-likelihood functions ℓ_W :

$$\ell_c(\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}, \nu; \mathbf{X}, \mathbf{S}) = \ell_N(\mathbf{M}, \mathbf{S}^{-1}, \boldsymbol{\Omega}; \mathbf{X} | \mathbf{S}) + \ell_W(\nu, \boldsymbol{\Sigma}; \mathbf{S})$$

From the definitions of the MxVN and Wishart distributions, we have, after ignoring additive constants,

$$\begin{aligned} \ell_N(\mathbf{M}, \mathbf{S}^{-1}, \boldsymbol{\Omega}; \mathbf{X} | \mathbf{S}) &= -\frac{np}{2} \log |\boldsymbol{\Omega}| + \frac{q}{2} \sum_{i=1}^n \log |\mathbf{S}_i| \\ &\quad - \frac{1}{2} \text{tr} \left[\sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i \boldsymbol{\Omega}^{-1} \mathbf{X}_i^T + \left(\sum_{i=1}^n \mathbf{S}_i \right) \mathbf{M} \boldsymbol{\Omega}^{-1} \mathbf{M}^T - 2 \left(\sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i \right) \boldsymbol{\Omega}^{-1} \mathbf{M}^T \right] \end{aligned}$$

and

$$\begin{aligned} \ell_W(\nu, \boldsymbol{\Sigma}; \mathbf{S}) &= (\nu - 2)/2 \sum_{i=1}^n \log |\mathbf{S}_i| - \sum_{i=1}^n \text{tr}(\boldsymbol{\Sigma} \mathbf{S}_i)/2 - n\nu p/2 \log 2 \\ &\quad + n(\nu + p - 1)/2 \log |\boldsymbol{\Sigma}| - n \log \Gamma_p((\nu + p - 1)/2). \end{aligned}$$

To simplify computation of the ML estimators, we define the following complete data sufficient statistics for the parameters:

$$\mathbf{S}_{SX} = \sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i; \quad \mathbf{S}_S = \sum_{i=1}^n \mathbf{S}_i; \quad \mathbf{S}_{XSX} = \sum_{i=1}^n \mathbf{X}_i^T \mathbf{S}_i \mathbf{X}_i; \quad \mathbf{S}_{|S|} = \sum_{i=1}^n \log |\mathbf{S}_i|.$$

Taking derivatives of log-likelihoods yields the ML estimates:

$$\begin{aligned} \widehat{\mathbf{M}} &= \left(\sum_{i=1}^n \mathbf{S}_i \right)^{-1} \sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i = \mathbf{S}_S^{-1} \mathbf{S}_{SX}, \\ \widehat{\boldsymbol{\Omega}} &= \frac{1}{np} \sum_{i=1}^n (\mathbf{X}_i - \widehat{\mathbf{M}})^T \mathbf{S}_i (\mathbf{X}_i - \widehat{\mathbf{M}}) = \frac{1}{np} \left(\mathbf{S}_{XSX} - \mathbf{S}_{SX}^T \mathbf{S}_S^{-1} \mathbf{S}_{SX} \right), \\ \widehat{\boldsymbol{\Sigma}}^{-1} &= \frac{1}{n(\nu + p - 1)} \sum_{i=1}^n \mathbf{S}_i = \frac{\mathbf{S}_S}{n(\nu + p - 1)}. \end{aligned}$$

The ML estimate of ν can be obtained by finding the root of the equation:

$$n\psi_p((\nu + p - 1)/2) - (\mathbf{S}_{|S|} - np \log 2 + n \log |\boldsymbol{\Sigma}|) = 0$$

with $\psi_p(\cdot)$ the p -variate digamma function, defined as $\psi_p(x) = d \log \Gamma_p(x) / dx$. The ML estimate of ν may be obtained numerically by a one-dimensional search algorithm. We now use the development in this section in our EM algorithm for a sample from the $\text{MxV}t$ distribution.

B. Estimating parameters from a MxVt sample

1) *The EM algorithm:* Let $\mathbf{X}_i, i = 1, 2, \dots, n$ be independent identically distributed realizations from $t_{p,q}(\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$. As in the main article, we write $\boldsymbol{\Theta} \equiv \{\nu, \mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega}\}$. From the development in the introduction of this section, for each $i = 1, 2, \dots, n$, let \mathbf{S}_i be (unobserved) random matrices as per Equation (S-4) and Property 2. Then the expected complete log-likelihood function is

$$\begin{aligned} Q(\boldsymbol{\Theta}; \boldsymbol{\Theta}^{(t)}) &= -\frac{np}{2} \log |\boldsymbol{\Omega}| - \frac{np \log 2}{2} - n \log \Gamma_p \left(\frac{\nu + p - 1}{2} \right) + n \frac{\nu + p - 1}{2} \log |\boldsymbol{\Sigma}| \\ &+ \mathbb{E}_{\boldsymbol{\Theta}^{(t)}} \left\{ \left[-\frac{1}{2} \text{tr} \left(\sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i \boldsymbol{\Omega}^{-1} \mathbf{X}_i^T + \sum_{i=1}^n \mathbf{S}_i \mathbf{M} \boldsymbol{\Omega}^{-1} \mathbf{M}^T - 2 \sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i \boldsymbol{\Omega}^{-1} \mathbf{M}^T \right) \right. \right. \\ &\left. \left. + \frac{\nu - 2}{2} \sum_{i=1}^n \log |\mathbf{S}_i| - \frac{1}{2} \sum_{i=1}^n \text{tr}(\boldsymbol{\Sigma} \mathbf{S}_i) + \frac{q}{2} \sum_{i=1}^n \log |\mathbf{S}_i| \middle| \mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_n \right] \right\}. \end{aligned} \quad (\text{S-5})$$

a) *E-step:* Using Property 2, the expectation step (E-step) updates, at the current value $\boldsymbol{\Theta}^{(t)}$ of $\boldsymbol{\Theta}$, are:

$$\begin{aligned} \mathbf{S}_i^{(t+1)} &\doteq \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\mathbf{S}_i | \mathbf{X}_i) = (\nu^{(t)} + p + q - 1) [(\mathbf{X}_i - \mathbf{M}^{(t)}) \boldsymbol{\Omega}^{(t)-1} (\mathbf{X}_i - \mathbf{M}^{(t)})^T + \boldsymbol{\Sigma}^{(t)}]^{-1}, \\ \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\log |\mathbf{S}_i | \mathbf{X}_i) &= \psi_p \left(\frac{\nu^{(t)} + p + q - 1}{2} \right) + p \log 2 + \log \left| \frac{\mathbf{S}_i^{(t+1)}}{\nu^{(t)} + p + q - 1} \right|, \end{aligned}$$

with $\psi_p(\cdot)$ as the p -variate digamma function. We define and store the sufficient statistics to reduce computational calculations:

$$\begin{aligned} \mathbf{S}_S^{(t+1)} &\doteq \sum_{i=1}^n \mathbf{S}_i^{(t+1)}, \\ \mathbf{S}_{SX}^{(t+1)} &\doteq \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\mathbf{S}_i \mathbf{X}_i | \mathbf{X}_i) = \sum_{i=1}^n \mathbf{S}_i^{(t+1)} \mathbf{X}_i, \\ \mathbf{S}_{X SX}^{(t+1)} &\doteq \sum_{i=1}^n \mathbb{E}_{\boldsymbol{\Theta}^{(t)}}(\mathbf{X}_i^T \mathbf{S}_i \mathbf{X}_i | \mathbf{X}_i) = \sum_{i=1}^n \mathbf{X}_i^T \mathbf{S}_i^{(t+1)} \mathbf{X}_i, \\ \mathbf{S}_{|S|}^{(t+1)} &\doteq \mathbb{E}_{\boldsymbol{\Theta}^{(t)}} \left[\sum_{i=1}^n \log |\mathbf{S}_i| \middle| \mathbf{X}_i \right], \end{aligned}$$

with the last expression needed only when we are also estimating ν . In that case, these statistics can be expressed with $(\nu^{(t)} + p + q - 1)$ factored out, and for convenience may be computed and stored that way when ν needs to be estimated.

b) *Maximization step:* Based on the updated weight matrices $\mathbf{S}_i^{(t+1)}$ and statistics based on $\boldsymbol{\Theta}^{(t)}$ and \mathbf{X} , we get the updates:

$$\begin{aligned} \widehat{\mathbf{M}} &= \left(\sum_{i=1}^n \mathbf{S}_i^{(t+1)} \right)^{-1} \sum_{i=1}^n \mathbf{S}_i \mathbf{X}_i = \mathbf{S}_S^{(t+1)-1} \mathbf{S}_{SX}^{(t+1)}, \\ \widehat{\boldsymbol{\Omega}} &= \frac{1}{np} \sum_{i=1}^n (\mathbf{X}_i - \mathbf{M}^{(t)})^T \mathbf{S}_i^{(t+1)} (\mathbf{X}_i - \mathbf{M}^{(t)}) \\ &= \frac{1}{np} \left(\mathbf{S}_{X SX}^{(t+1)} - \mathbf{S}_{SX}^{(t+1)T} \mathbf{S}_S^{(t+1)-1} \mathbf{S}_{SX}^{(t+1)} \right), \\ \widehat{\boldsymbol{\Sigma}}^{-1} &= \frac{1}{n(\nu^{(t)} + p - 1)} \sum_{i=1}^n \mathbf{S}_i^{(t+1)} = \frac{\mathbf{S}_S^{(t+1)}}{n(\nu^{(t)} + p - 1)}. \end{aligned}$$

And, again, treating the set of $\mathbf{S}_i^{(t+1)}$ as observed, the MLE of ν can be obtained:

$$n \frac{d}{d\nu} \log \Gamma_p((\nu + p - 1)/2) - \frac{1}{2} (\mathbf{S}_{|S|} - np \log 2 + n \log |\widehat{\boldsymbol{\Sigma}}|) = 0.$$

Defining $\kappa^{(t)} = \nu^{(t)} + p + q - 1$ for compactness:

$$\begin{aligned} 0 &= n\psi_p((\nu + p - 1)/2) - \left(n\psi_p\left(\frac{\kappa^{(t)}}{2}\right) + \sum_{i=1}^n \log \left| \frac{\mathbf{S}_i^{(t+1)}}{\kappa^{(t)}} \right| - n \log \left| \frac{\mathbf{S}_S^{(t+1)}}{n(\nu^{(t)} + p - 1)} \right| \right) \\ &= \psi_p((\nu + p - 1)/2) - \left(\psi_p\left(\frac{\kappa^{(t)}}{2}\right) + \frac{1}{n} \sum_{i=1}^n \log \left| \mathbf{Z}_i^{(t+1)} \right| + p \log \frac{n(\nu^{(t)} + p - 1)}{\kappa^{(t)}} - \log \left| \mathbf{Z}_S^{(t+1)} \right| \right) \end{aligned} \quad (\text{S-6})$$

where \mathbf{Z}_* is the appropriate \mathbf{S}_* statistic with $(\nu^{(t)} + p + q - 1)$ factored out and ψ_p is the p -dimensional digamma function. This can be solved using a 1-dimensional search.

c) *ML Estimation with the Expectation/Conditional Maximization Either (ECME) algorithm:* First we note that, if ν is known, there is no need to partition the M-step into multiple constrained maximization steps. If ν is required to be estimated, there is no difference between a standard EM and a standard ECM (Expectation/Conditional Maximization) algorithm in this setting, since, as in the case of the multivariate t distribution, the complete data likelihood function factorizes into $\Theta_1 = (\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ and $\Theta_2 = (\nu)$. However, by partitioning it in this way, it is possible, similarly to the case of the multivariate t , to find a more efficient method of maximization. This is desirable because the M-step for ν can be slow. Here we present an ECME (Expectation/Conditional Maximization Either) algorithm that first maximizes the expected log-likelihood for $(\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$ and then maximizes the actual log-likelihood over ν given the current values $(\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$, similar to [7].

Given $\Theta_1 = (\mathbf{M}, \boldsymbol{\Sigma}, \boldsymbol{\Omega})$, we can maximize for ν in Equation (1), yielding the set of equations provided in (2)

$$\begin{aligned} 0 &= n\psi_p((\nu + p - 1)/2) - \left\{ n\psi_p\left(\frac{\kappa}{2}\right) + \sum_{i=1}^n \log \left| \frac{\mathbf{S}_i^{(t+1)}}{\kappa} \right| - n \log \left| \frac{\mathbf{S}_S^{(t+1)}}{n(\nu + p - 1)} \right| \right\} \\ &= \psi_p((\nu + p - 1)/2) - \left\{ \psi_p\left(\frac{\kappa}{2}\right) + \frac{1}{n} \sum_{i=1}^n \log \left| \mathbf{Z}_i^{(t+1)} \right| + p \log \frac{n(\nu + p - 1)}{\kappa} - \log \left| \mathbf{Z}_S^{(t+1)} \right| \right\}. \end{aligned}$$

The difference is that the solution for $\nu^{(t+1)}$ no longer depends on $\nu^{(t)}$, Solving this equation is slightly more computationally complex than solving Equation (S-6) but this converges in fewer iterations. The ML estimating equation can be solved by a one-dimensional search, providing a ECME algorithm with the steps (as also provided in the main article):

1. **E-step:** Update \mathbf{S}_i weights and statistics based on $\Theta^{(t)}$ and \mathbf{X} .
2. **CME-step:** Update $\Theta_1^{(t+1)} = (\mathbf{M}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \boldsymbol{\Omega}^{(t+1)})$.
3. **CME-step:** Update $\Theta_2^{(t+1)} = \nu^{(t+1)}$ using the observed log-likelihood given the current values $(\mathbf{M}^{(t+1)}, \boldsymbol{\Sigma}^{(t+1)}, \boldsymbol{\Omega}^{(t+1)})$ by solving Equation (2).

Repeat these steps until convergence.

2) *Fitting with restrictions on the parameters:* In some settings, restrictions on the parametrization of the mean or covariance matrices are appropriate. In this section, we derive solutions in the cases of mean matrices that are constant across rows, columns, or the entire matrix. In [10] some results for restrictions on covariance matrices were derived and in this paper AR(1) covariance structures and compound symmetry (CS) variance structures were used; however, they were fit numerically as closed forms for the derivatives and determinants exist. Let $\mathbf{1}_{p,q}$ denote a $(p \times q)$ matrix consisting only of 1s. Then it can be shown that these are the appropriate M-step estimates for certain mean matrix constraints:

$$\begin{aligned} \mathbf{M} = \mathbf{1}_{p,q}\boldsymbol{\mu} : & \quad \widehat{\mathbf{M}} = \text{tr}(\mathbf{S}_{SX}\widehat{\boldsymbol{\Omega}}^{-1}\mathbf{1}_{q,p})/\text{tr}(\mathbf{S}_S\mathbf{1}_{p,q}\widehat{\boldsymbol{\Omega}}^{-1}\mathbf{1}_{q,p})\mathbf{1}_{p,q} \\ \mathbf{M} = \mathbf{1}_{p,1}\boldsymbol{\mu}_{1,q} : & \quad \widehat{\mathbf{M}} = \mathbf{1}_{p,p}\mathbf{S}_{SX}/(\mathbf{1}_{1,p}\mathbf{S}_S\mathbf{1}_{p,1}) \\ \mathbf{M} = \boldsymbol{\mu}_{p,1}\mathbf{1}_{1,q} : & \quad \widehat{\mathbf{M}} = \mathbf{S}_S^{-1}\mathbf{S}_{SX}\widehat{\boldsymbol{\Omega}}^{-1}\mathbf{1}_{q,q}/(\mathbf{1}_{1,q}\widehat{\boldsymbol{\Omega}}^{-1}\mathbf{1}_{q,1}) \end{aligned}$$

which can be used to simplify the ECME algorithms further.