

A Mixture of Coalesced Generalized Hyperbolic Distributions

Cristina Tortora*, Brian C. Franczak**, Ryan P. Browne†, Paul D. McNicholas††

*Department of Mathematics & Statistics, San José State University, CA, USA.

**Department of Mathematics & Statistics, MacEwan University, Edmonton, AB, Canada.

†Department of Statistics and Actuarial Sciences, University of Waterloo, ON, Canada.

††Department of Mathematics and Statistics, McMaster University, ON, Canada.

Abstract

A mixture of multiple scaled generalized hyperbolic distributions (MMSGHDs) is introduced. Then, a coalesced generalized hyperbolic distribution (CGHD) is developed by joining a generalized hyperbolic distribution with a multiple scaled generalized hyperbolic distribution. After detailing the development of the MMSGHDs, which arises via implementation of a multi-dimensional weight function, the density of the mixture of CGHDs is developed. A parameter estimation scheme is developed using the ever-expanding class of MM algorithms and the Bayesian information criterion is used for model selection. The issue of cluster convexity is examined and a special case of the MMSGHDs is developed that is guaranteed to have convex clusters. These approaches are illustrated and compared using simulated and real data. The identifiability of the MMSGHDs and the mixture of CGHDs is discussed in an appendix.

Keywords: clustering; coalesced distributions; convexity; finite mixture models; generalized hyperbolic distribution; mixture of mixtures; MM algorithm; multiple scaled distributions.

1 Introduction

Finite mixture models have been linked with clustering since the idea of defining a cluster in terms of a component in finite mixture model was put forth more than 60 years ago (see McNicholas, 2016a, Section 2.1). Nowadays, mixture model-based clustering is a popular approach to clustering. A random vector \mathbf{X} arises from a finite mixture model if, for all $\mathbf{x} \in \mathbf{X}$, its density can be written

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

where $\pi_g > 0$ such that $\sum_{g=1}^G \pi_g = 1$ are the mixing proportions, $f_g(\mathbf{x} \mid \boldsymbol{\theta}_g)$ is the g th component density, and $\boldsymbol{\vartheta} = (\boldsymbol{\pi}, \boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_G)$ denotes the vector of parameters with $\boldsymbol{\pi} = (\pi_1, \dots, \pi_G)$. The component densities $f_1(\mathbf{x} \mid \boldsymbol{\theta}_1), \dots, f_G(\mathbf{x} \mid \boldsymbol{\theta}_G)$ are typically taken to be of the same type, most commonly multivariate Gaussian. In fact, until a few years after the turn of the century, almost all work on clustering and classification using mixture models had been based on Gaussian mixture models (e.g., Banfield and Raftery, 1993; Celeux and Govaert, 1995; Ghahramani and Hinton, 1997; Tipping and Bishop, 1999; McLachlan and Peel, 2000; Fraley and Raftery, 2002).

Early work on non-Gaussian mixtures was on mixtures of multivariate t -distributions (e.g., Peel and McLachlan, 2000). A little beyond the turn of the century, work on t -mixtures burgeoned into a substantial subfield of mixture model-based classification (e.g., McLachlan et al., 2007; Andrews and McNicholas, 2011a,b, 2012; Baek and McLachlan, 2011; Steane et al., 2012; Lin et al., 2014; Pesevski et al., 2018). Around the same time, work on mixtures of skewed distributions took off, including work on skew-normal mixtures (e.g., Lin, 2009), skew- t mixtures (e.g., Lin, 2010; Vrbik and McNicholas, 2012, 2014; Lee and McLachlan, 2013a,b; Murray et al., 2014), Laplace mixtures (e.g., Franczak et al., 2014), variance-gamma mixtures (McNicholas et al., 2017), generalized hyperbolic mixtures (Browne and McNicholas, 2015), and other non-elliptically contoured distributions (e.g., Karlis and Santourian, 2009; Murray et al., 2017; Tang et al., 2018). A thorough review of work on model-based clustering is given by McNicholas (2016b).

More recently, mixtures of multiple scaled distributions have been considered (Section 2). In the present manuscript, a multiple scaled generalized hyperbolic distribution is introduced (Section 3). Because the (multivariate) generalized hyperbolic distribution is not a special case of the multiple scaled generalized hyperbolic distribution, a coalesced generalized hyperbolic distribution is also developed (Section 4). The issue of cluster convexity, which has essentially been ignored in work to date on multiple scaled distributions is discussed and a special case of the multiple scaled generalized hyperbolic distribution is developed to guarantee that the components, i.e., the clusters, are convex (Section 5).

2 Multiple Scaled Distributions

The distribution of a p -dimensional random variable \mathbf{X} is said to be a normal variance-mean mixture if its density can be written in the form

$$f(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \boldsymbol{\theta}) = \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma}) h(w \mid \boldsymbol{\theta}) dw, \quad (1)$$

where $\phi_p(\mathbf{x} \mid \boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma})$ is the density of a p -dimensional Gaussian distribution with mean $\boldsymbol{\mu} + w\boldsymbol{\alpha}$ and covariance matrix $w\boldsymbol{\Sigma}$, and $h(w \mid \boldsymbol{\theta})$ is the density of a univariate random variable $W > 0$ that has the role of a weight function (see Barndorff-Nielsen et al., 1982; Gneiting, 1997). This weight function can take on many forms, some of which lead to density representations for well-known non-Gaussian distributions, e.g., if $h(w \mid \boldsymbol{\theta})$ is the density of

an inverse-gamma random variable with parameters $(\nu/2, \nu/2)$, then (1) is a representation of the skew- t distribution with ν degrees of freedom (see Demarta and McNeil, 2005; Murray et al., 2014). Further details on, and examples of, normal variance-mean mixtures are given by Barndorff-Nielsen (1978), Kotz et al. (2001), and Kotz and Nadarajah (2004), amongst others.

Now, the density of $W > 0$ from an inverse-gamma distribution with parameters (α, β) is given by

$$h(w \mid \alpha, \beta) = w^{-\alpha-1} \frac{\beta^\alpha}{\Gamma(\alpha)} \exp \left\{ -\frac{\beta}{w} \right\}, \quad (2)$$

where $\Gamma(\cdot)$ is the Gamma function. Setting $\alpha = \beta = \nu/2$ in (2) gives

$$h(w \mid \nu/2, \nu/2) = w^{-\nu/2-1} \frac{(\nu/2)^{\nu/2}}{\Gamma(\nu/2)} \exp \left\{ -\frac{\nu}{2w} \right\}, \quad (3)$$

Setting $\boldsymbol{\alpha} = \mathbf{0}$ in (1), and using (3) for $h(w \mid \boldsymbol{\theta})$, it follows that the density of the multivariate t -distribution with ν degrees of freedom can be written

$$\begin{aligned} f_t(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu) &= \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu}, w\boldsymbol{\Sigma}) h(w \mid \nu/2, \nu/2) dw \\ &= \frac{\Gamma([\nu + p]/2) |\boldsymbol{\Sigma}|^{-1/2}}{(\pi\nu)^{p/2} \Gamma(\nu/2) [1 + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})/\nu]^{(\nu+p)/2}} \end{aligned} \quad (4)$$

where $\delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})$ is the squared Mahalanobis distance between \mathbf{x} and $\boldsymbol{\mu}$. Forbes and Wraith (2014) show that a multi-dimensional weight variable

$$\boldsymbol{\Delta}_{\mathbf{W}} = \text{diag}(w_1^{-1}, \dots, w_p^{-1})$$

can be incorporated into (1) via an eigen-decomposition of the symmetric positive-definite matrix $\boldsymbol{\Sigma}$. Specifically, they set $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'$, where $\boldsymbol{\Gamma}$ is a $p \times p$ matrix of eigenvectors and $\boldsymbol{\Phi}$ is a $p \times p$ diagonal matrix containing the eigenvalues of $\boldsymbol{\Sigma}$. It follows that the density of \mathbf{X} becomes

$$\begin{aligned} f(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\theta}) &= \\ \int_0^\infty \cdots \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu} + \boldsymbol{\Delta}_{\mathbf{W}}\boldsymbol{\alpha}, \boldsymbol{\Gamma}\boldsymbol{\Delta}_{\mathbf{W}}\boldsymbol{\Phi}\boldsymbol{\Gamma}') h_{\mathbf{W}}(w_1, \dots, w_p \mid \boldsymbol{\theta}) dw_1 \dots dw_p, \end{aligned} \quad (5)$$

where

$$h_{\mathbf{W}}(w_1, \dots, w_p \mid \boldsymbol{\theta}) = h(w_1 \mid \boldsymbol{\theta}_1) \times \cdots \times h(w_p \mid \boldsymbol{\theta}_p)$$

is a p -dimensional density such that the random variables W_1, \dots, W_p are independent, i.e., the weights are independent. The density given in (5) adds flexibility to normal variance-mean mixtures because the parameters $\boldsymbol{\theta}_1, \dots, \boldsymbol{\theta}_p$ are free to vary in each dimension. Using the density in (5), Forbes and Wraith (2014) derive the density of a multiple scaled

multivariate- t distribution, Wraith and Forbes (2015) derive a multiple scaled normal-inverse Gaussian distribution, and Franczak et al. (2015) develop a multiple scaled shifted asymmetric Laplace distribution.

Setting $\Sigma = \mathbf{\Gamma}\mathbf{\Phi}\mathbf{\Gamma}'$, it follows from (5) that the density of a multiple scaled analogue of (4) can be written

$$f_{tMS}(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\nu}) = \int_0^\infty \cdots \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Gamma}\mathbf{\Phi}\mathbf{\Delta}_w\mathbf{\Gamma}') h_w(w_1, \dots, w_p \mid \boldsymbol{\nu}) dw_1 \dots dw_p, \quad (6)$$

where $\mathbf{\Delta}_w = \text{diag}(w_1^{-1}, \dots, w_p^{-1})$ and the weight function

$$h_w(w_1, \dots, w_p \mid \boldsymbol{\nu}) = h(w_1 \mid \nu_1/2, \nu_1/2) \times \cdots \times h(w_p \mid \nu_p/2, \nu_p/2)$$

is a p -dimensional Gamma density, where $h(w_j \mid \nu_j/2, \nu_j/2)$ is given by (3). Note that the scaled Gaussian density in (6) can be written

$$\phi_p(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Gamma}\mathbf{\Phi}\mathbf{\Delta}_w\mathbf{\Gamma}') = \prod_{j=1}^p \phi_1([\mathbf{\Gamma}'\mathbf{x}]_j \mid [\mathbf{\Gamma}'\boldsymbol{\mu}]_j, \Phi_j w_j^{-1}) = \prod_{j=1}^p \phi_1([\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})]_j \mid 0, \Phi_j w_j^{-1}), \quad (7)$$

where $\phi_1([\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})]_j \mid 0, \Phi_j w_j^{-1})$ is the density of a univariate Gaussian distribution with mean 0 and variance $\Phi_j w_j^{-1}$, $[\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})]_j$ is the j th element of $\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})$, and Φ_j is the j th eigenvalue of $\mathbf{\Phi}$, i.e., the j th diagonal element of the matrix $\mathbf{\Phi}$. It follows that (6) can be written

$$f_{tMS}(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\nu}) = \prod_{j=1}^p \int_0^\infty \phi_1([\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})]_j \mid 0, \Phi_j w_j^{-1}) h(w_j \mid \nu_j/2, \nu_j/2) dw_j. \quad (8)$$

Solving the integral in (8) gives the density of a multiple scaled multivariate- t distribution,

$$f_{tMS}(\mathbf{x} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\nu}) = \prod_{j=1}^p \frac{\Gamma([\nu_j + 1]/2)}{\Gamma(\nu_j/2)(\Phi_j \nu_j \pi)^{1/2}} \left[1 + \frac{[\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})]_j^2}{\Phi_j \nu_j} \right]^{-(\nu_j + 1)/2}, \quad (9)$$

where Φ_j is the j th eigenvalue of $\mathbf{\Phi}$, $\mathbf{\Gamma}$ is a matrix of eigenvectors, $\boldsymbol{\mu}$ is a location parameter, and $[\mathbf{\Gamma}'(\mathbf{x} - \boldsymbol{\mu})]_j^2 / \Phi_j$ can be regarded as the squared Mahalanobis distance between \mathbf{x} and $\boldsymbol{\mu}$.

The main difference between the traditional multivariate- t density given in (4) and the multiple scaled multivariate- t density given in (9) is that the degrees of freedom can now be parameterized separately in each dimension j . Therefore, unlike the standard multivariate- t distribution, the multiple scaled density in (9) can account for different tail weight in each dimension (Forbes and Wraith, 2014).

3 Mixture of Multiple Scaled Generalized Hyperbolic Distributions

There are different ways to formulate the density of a generalized hyperbolic distribution (GHD; see McNeil et al., 2005, for example). Browne and McNicholas (2015) use a mixture of GHDs (MGHDs) for clustering and they use the following formulation for the density of a p -dimensional random vector \mathbf{X} from a GHD:

$$f_{\text{GH}}(\mathbf{x} \mid \boldsymbol{\theta}) = \left[\frac{\omega + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})}{\omega + \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha}} \right]^{(\lambda - p/2)/2} \frac{K_{\lambda - p/2} \left(\sqrt{[\omega + \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha}] [\omega + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma})]} \right)}{(2\pi)^{p/2} |\boldsymbol{\Sigma}|^{1/2} K_{\lambda}(\omega) \exp \left\{ -(\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha} \right\}}, \quad (10)$$

where $\boldsymbol{\mu} \in \mathbb{R}^p$ is the location parameter, $\boldsymbol{\alpha} \in \mathbb{R}^p$ is the skewness parameter, $\boldsymbol{\Sigma} \in \mathbb{R}^{p \times p}$ is the scale matrix, $\lambda \in \mathbb{R}$ is the index parameter, $\omega \in \mathbb{R}^+$ is the concentration parameter, and K_{λ} is the modified Bessel function of the third kind with index λ . Now, $\mathbf{X} \mid w \sim \text{N}(\boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma})$ and $W \mid \mathbf{x} \sim \text{GIG}(\omega + \boldsymbol{\alpha}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha}, \omega + \delta(\mathbf{x}, \boldsymbol{\mu} \mid \boldsymbol{\Sigma}), \lambda - p/2)$, where $W \sim \text{GIG}(a, b, \lambda)$ denotes that W follows a generalized inverse Gaussian (GIG) distribution with density formulated as

$$q(w \mid a, b, \lambda) = \frac{(a/b)^{\lambda/2} w^{\lambda-1}}{2K_{\lambda}(\sqrt{ab})} \exp \left\{ -\frac{aw + b/w}{2} \right\}, \quad (11)$$

for $w > 0$, where $a, b \in \mathbb{R}^+$, and $\lambda \in \mathbb{R}$. The GIG distribution has some attractive properties including the tractability of the following expected values:

$$\mathbb{E}[W] = \sqrt{\frac{b}{a}} \frac{K_{\lambda+1}(\sqrt{ab})}{K_{\lambda}(\sqrt{ab})}, \quad \mathbb{E}[1/W] = \sqrt{\frac{a}{b}} \frac{K_{\lambda+1}(\sqrt{ab})}{K_{\lambda}(\sqrt{ab})} - \frac{2\lambda}{b}, \quad (12)$$

$$\mathbb{E}[\log W] = \log \left(\sqrt{\frac{b}{a}} \right) + \frac{1}{K_{\lambda}(\sqrt{ab})} \frac{\partial}{\partial \lambda} K_{\lambda}(\sqrt{ab}). \quad (13)$$

Write $\mathbf{X} \sim \text{GHD}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \omega, \lambda)$ to denote that the p -dimensional random variable \mathbf{X} has the density in (10).

Now, note that the formulation of the GHD in (10) can be written as a normal variance-mean mixture where the univariate density is GIG, i.e.,

$$\mathbf{X} = \boldsymbol{\mu} + W\boldsymbol{\alpha} + \sqrt{W}\mathbf{V}, \quad (14)$$

where $\mathbf{V} \sim \text{N}(\mathbf{0}, \boldsymbol{\Sigma})$ and W has density

$$h(w \mid \omega, 1, \lambda) = \frac{w^{\lambda-1}}{2K_{\lambda}(\omega)} \exp \left\{ -\frac{\omega}{2} \left(w + \frac{1}{w} \right) \right\}, \quad (15)$$

for $w > 0$, where ω and λ are as previously defined. Note that (15) is just an alternative parameterization of the GIG distribution. From (14) and (15), it follows that the generalized

hyperbolic density can be written

$$f(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \omega, \lambda) = \int_0^\infty \phi_p(\mathbf{x} \mid \boldsymbol{\mu} + w\boldsymbol{\alpha}, w\boldsymbol{\Sigma}) h(w \mid \omega, 1, \lambda) dw. \quad (16)$$

We can use (5), (16) and an alternative parameterization to write the density of a multiple scaled generalized hyperbolic distribution (MSGHD) as

$$f_{\text{MSGHD}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}) = \int_0^\infty \cdots \int_0^\infty \phi_p(\boldsymbol{\Gamma}'\mathbf{x} - \boldsymbol{\mu} - \boldsymbol{\Delta}_w\boldsymbol{\alpha} \mid \mathbf{0}, \boldsymbol{\Delta}_w\boldsymbol{\Phi}) h_w(w_1, \dots, w_p \mid \boldsymbol{\omega}, \mathbf{1}, \boldsymbol{\lambda}) dw_1 \dots dw_p, \quad (17)$$

where $\boldsymbol{\omega} = (\omega_1, \dots, \omega_p)'$, $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_p)'$, $\mathbf{1}$ is a p -vector of 1s, and

$$h_w(w_1, \dots, w_p \mid \boldsymbol{\omega}, \mathbf{1}, \boldsymbol{\lambda}) = h(w_1 \mid \omega_1, 1, \lambda_1) \times \cdots \times h(w_p \mid \omega_p, 1, \lambda_p).$$

From (6) and (7), it follows that (17) can be written

$$f_{\text{MSGHD}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}) = \prod_{j=1}^p \int_0^\infty \phi_1([\boldsymbol{\Gamma}'\mathbf{x} - \boldsymbol{\mu} - \boldsymbol{\Delta}_w\boldsymbol{\alpha}]_j \mid 0, \Phi_j w_j) h_W(w_j \mid \omega_j, 1, \lambda_j) dw_j \\ = \prod_{j=1}^p \left\{ \left[\frac{\omega_j + \Phi_j^{-1}([\boldsymbol{\Gamma}'\mathbf{x}]_j - \mu_j)^2}{\omega_j + \alpha_j^2 \Phi_j^{-1}} \right]^{\frac{\lambda_j - 1/2}{2}} \frac{K_{\lambda_j - 1/2} \left(\sqrt{[\omega_j + \alpha_j^2 \Phi_j^{-1}] [\omega_j + \Phi_j^{-1}([\boldsymbol{\Gamma}'\mathbf{x}]_j - \mu_j)^2]} \right)}{(2\pi)^{1/2} \Phi_j^{1/2} K_{\lambda_j}(\omega_j) \exp \left\{ -([\boldsymbol{\Gamma}'\mathbf{x}]_j - \mu_j) \Phi_j^{-1} \alpha_j \right\}} \right\},$$

where $[\boldsymbol{\Gamma}'\mathbf{x}]_j$ is the j th element of the vector $\boldsymbol{\Gamma}'\mathbf{x}$, μ_j is the j th element of the location parameter $\boldsymbol{\mu}$, α_j is the j th element of the skewness parameter $\boldsymbol{\alpha}$, $\boldsymbol{\Gamma}$ is a $p \times p$ matrix of eigenvectors, Φ_j is the j th eigenvalue of the diagonal matrix $\boldsymbol{\Phi}$, $\boldsymbol{\omega} = (\omega_1, \dots, \omega_p)'$ controls the concentration in each dimension p , and $\boldsymbol{\lambda} = (\lambda_1, \dots, \lambda_p)'$ is a p -dimensional index parameter. Write $\mathbf{X} \sim \text{MSGHD}(\boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda})$ to indicate that the random vector \mathbf{X} follows an MSGHD with density $f_{\text{MSGHD}}(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda})$. Then, a mixture of MSGHDs (MMSGHDs) has density

$$f(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^G \pi_g f_{\text{MSGHD}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\alpha}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g). \quad (18)$$

The identifiability of the MMSGHD is discussed in Appendix C.

4 Mixture of Coalesced Generalized Hyperbolic Distributions

Note that the generalized hyperbolic distribution is not a special or limiting case of the MSGHD under any parameterization with $p > 1$. Motivated by this, consider a coalesced

generalized hyperbolic distribution (CGHD) that contains both the generalized hyperbolic distribution and MSGHD as limiting cases. The CGHD arises through the introduction of a random vector

$$\mathbf{R} = U\mathbf{X} + (1 - U)\mathbf{S}, \quad (19)$$

where $\mathbf{X} = \mathbf{\Gamma}\mathbf{Y}$, $\mathbf{Y} \sim \text{GHD}(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\alpha}, \omega_0, \lambda_0)$, $\boldsymbol{\Sigma} = \mathbf{\Gamma}\mathbf{\Phi}\mathbf{\Gamma}'$, $\mathbf{S} \sim \text{MSGHD}(\boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda})$, and U is an indicator variable such that

$$U = \begin{cases} 1 & \text{if } \mathbf{R} \text{ follows a generalized hyperbolic distribution, and} \\ 0 & \text{if } \mathbf{R} \text{ follows a MSGHD.} \end{cases}$$

It follows that $\mathbf{X} = \mathbf{\Gamma}\boldsymbol{\mu} + W\mathbf{\Gamma}\boldsymbol{\alpha} + \sqrt{W}\mathbf{\Gamma}\mathbf{V}$, where $\mathbf{\Gamma}\mathbf{V} \sim N_p(\mathbf{0}, \mathbf{\Gamma}\mathbf{\Phi}\mathbf{\Gamma}')$, $\mathbf{S} = \mathbf{\Gamma}\boldsymbol{\mu} + \mathbf{\Gamma}\boldsymbol{\alpha}\boldsymbol{\Delta}_{\mathbf{w}} + \mathbf{\Gamma}\mathbf{A}$, where $\mathbf{\Gamma}\mathbf{A} \sim N_p(\mathbf{0}, \mathbf{\Gamma}\boldsymbol{\Delta}_{\mathbf{w}}\mathbf{\Phi}\mathbf{\Gamma}')$, and the density of \mathbf{R} can be written

$$f_{\text{CGHD}}(\mathbf{r} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \omega_0, \lambda_0, \varpi) = \varpi f_{\text{GHD}}(\mathbf{r} \mid \boldsymbol{\mu}, \mathbf{\Gamma}\mathbf{\Phi}\mathbf{\Gamma}', \boldsymbol{\alpha}, \omega_0, \lambda_0) + (1 - \varpi) f_{\text{MSGHD}}(\mathbf{r} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}), \quad (20)$$

where $f_{\text{GHD}}(\cdot)$ is the density of a generalized hyperbolic random variable, $f_{\text{MSGHD}}(\cdot)$ is the density of a MSGHD random variable, and $\varpi \in (0, 1)$ is a mixing proportion. Note that the random vector \mathbf{R} would be distributed generalized hyperbolic if $\varpi = 1$ and would be distributed MSGHD if $\varpi = 0$. However, we must restrict $\varpi \in (0, 1)$ for identifiability reasons. Specifically, if $\varpi = 0$ then the value of $f_{\text{CGHD}}(\mathbf{r} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \omega_0, \lambda_0, \varpi)$ will be the same for any ω_0 and λ_0 whereas, if $\varpi = 1$, then the value of $f_{\text{CGHD}}(\mathbf{r} \mid \boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \omega_0, \lambda_0, \varpi)$ will be the same for any $\boldsymbol{\omega}$ and $\boldsymbol{\lambda}$. The parameters $\boldsymbol{\mu}$, $\boldsymbol{\alpha}$, $\mathbf{\Gamma}$, and $\mathbf{\Phi}$ are the same for both densities, the parameters ω_0 and λ_0 are univariate values peculiar to the generalized hyperbolic distribution, and the p -dimensional parameters $\boldsymbol{\omega}$ and $\boldsymbol{\lambda}$ are peculiar to the MSGHD. Write $\mathbf{R} \sim \text{CGHD}(\boldsymbol{\mu}, \mathbf{\Gamma}, \mathbf{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}, \omega_0, \lambda_0, \varpi)$ to indicate that the random vector \mathbf{R} follows a CGHD with density in (20).

A mixture of CGHDs (MCGHDs) has density

$$f(\mathbf{x} \mid \boldsymbol{\vartheta}) = \sum_{g=1}^G \pi_g f_{\text{CGHD}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \mathbf{\Gamma}_g, \mathbf{\Phi}_g, \boldsymbol{\alpha}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g, \omega_{0g}, \lambda_{0g}, \varpi_g).$$

Parameter estimation can be carried out via a generalized expectation-maximization (GEM) algorithm (Dempster et al., 1977). There are four sources of missing data: the latent w_{0ig} , the multi-dimensional weights $\boldsymbol{\Delta}_{\mathbf{w}ig} = \text{diag}(w_{1ig}, \dots, w_{pig})$, the component membership labels z_{ig} , and the inner component labels u_{ig} , for $i = 1, \dots, n$ and $g = 1, \dots, G$. As usual, $z_{ig} = 1$ if observation i belongs to component g and $z_{ig} = 0$ otherwise. Similarly, $u_{ig} = 1$ if observation i , in component g , is distributed generalized hyperbolic and $u_{ig} = 0$ if observation i , in component g , is distributed MSGHD. It follows that the complete-data

log-likelihood for the MCGHD is

$$\begin{aligned}
l_c = \sum_{i=1}^n \sum_{g=1}^G \Big\{ & z_{ig} \log \pi_g + z_{ig} u_{ig} \log \varpi_g + z_{ig} (1 - u_{ig}) \log (1 - \varpi_g) + z_{ig} u_{ig} \log h(w_{0ig} \mid \omega_{0g}, 1, \lambda_{0g}) \\
& + z_{ig} (1 - u_{ig}) \sum_{j=1}^p \log h(w_{jig} \mid \omega_{jg}, 1, \lambda_{jg}) + z_{ig} u_{ig} \log \phi_p(\mathbf{\Gamma}'_g \mathbf{x}_i \mid \boldsymbol{\mu}_g + w_{0ig} \boldsymbol{\alpha}_g, w_{0ig} \boldsymbol{\Phi}) \\
& + z_{ig} (1 - u_{ig}) \sum_{j=1}^p \log \phi_1([\mathbf{\Gamma}'_g \mathbf{x}_i]_j \mid \mu_{jg} + w_{jig} \alpha_{jg}, \omega_{jg} \phi_{jg}) \Big\}.
\end{aligned}$$

Further details on parameter estimation (Appendix A) and the identifiability of the MCGHD (Appendix C) are discussed in appendices.

5 Cluster Convexity

The definition of clusters has been discussed quite extensively in the literature. Recently, Hennig (2015) provides a very interesting discussion of some potential characteristics that clusters may have. Even though they cannot always be observed — and, in some situations, desired characteristics may even conflict — such characteristics provide important links and contrasts between different definitions that are used for a cluster. In fact, the idea of desirable characteristics is not a new one, e.g., Cormack (1971) gives internal cohesion and external isolation as two “basic ideas” in this direction. In this paper, we follow the definition given by McNicholas (2016a): “a cluster is a unimodal component within an appropriate finite mixture model”. The term appropriate means that the model has the flexibility to fit the data and, in many cases, this also means that each cluster is convex (see McNicholas, 2016a, Section 9.1). The development of flexible models outlined herein may make it a little easier to find an “appropriate” component density.

The MSGHD is more flexible than the GHD; however, similar to the multiple scaled multivariate t -distribution of Forbes and Wraith (2014), the MSGHD can have contours that are not convex. Accordingly, the MMSGHD can have components that are non-convex, leading to non-convex clusters. Consider the data in Figure 1. How many clusters are there? The most plausible answer to this question is two overlapping clusters: one with positive correlation between the variables and another with negative correlation between the variables. There may also be an argument for four or five. Note that the data are generated from a $G = 2$ component mixture of multivariate t -distributions.

The MMSGHD is fitted to these data for $G = 1, \dots, 5$ and, as is common in model-based clustering applications, the Bayesian information criterion (BIC; Schwarz, 1978) is used to select G . Note that the BIC is given by

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

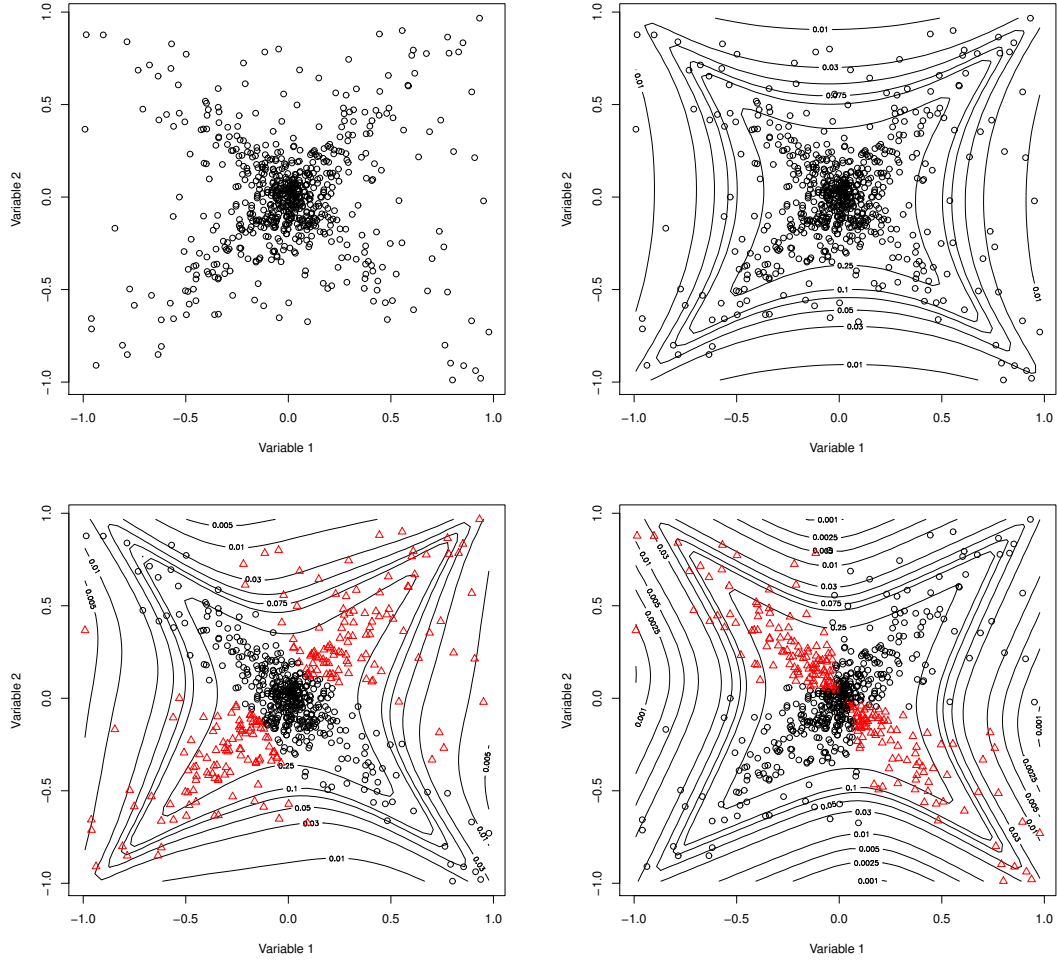


Figure 1: Scatter plot of data generated from a two-component t -mixture (top-left) along with contours from fitted $G = 1$ component MMSGHD (top-right), $G = 2$ component MMSGHD (bottom-left), and $G = 2$ component McMSGHD (bottom-right) models, where plotting symbol and colour represent predicted classifications.

where $l(\hat{\boldsymbol{\theta}})$ is the maximized log-likelihood, ρ is the number of free parameters, and n is the number of observations. For the MMSGHD, the BIC selects a $G = 1$ component model (Figure 1). Furthermore, looking at the $G = 2$ component MMSGHD solution (Figure 1) confirms that the problem is not just one of model selection; the $G = 2$ component MMSGHD solution selects one component that is roughly elliptical and another that is not convex. On the other hand, forcing the MMSGHD to be convex — which can be done by imposing the constraint $\lambda_j > 1$, for $j = 1, \dots, p$ — leads to what we call the convex MMSGHD (cMSGHD). Fitting the corresponding mixture of cMSGHDs (McMSGHDs) ensures that, if each component is associated with a cluster, then convex clusters are guaranteed. Results in a $G = 2$

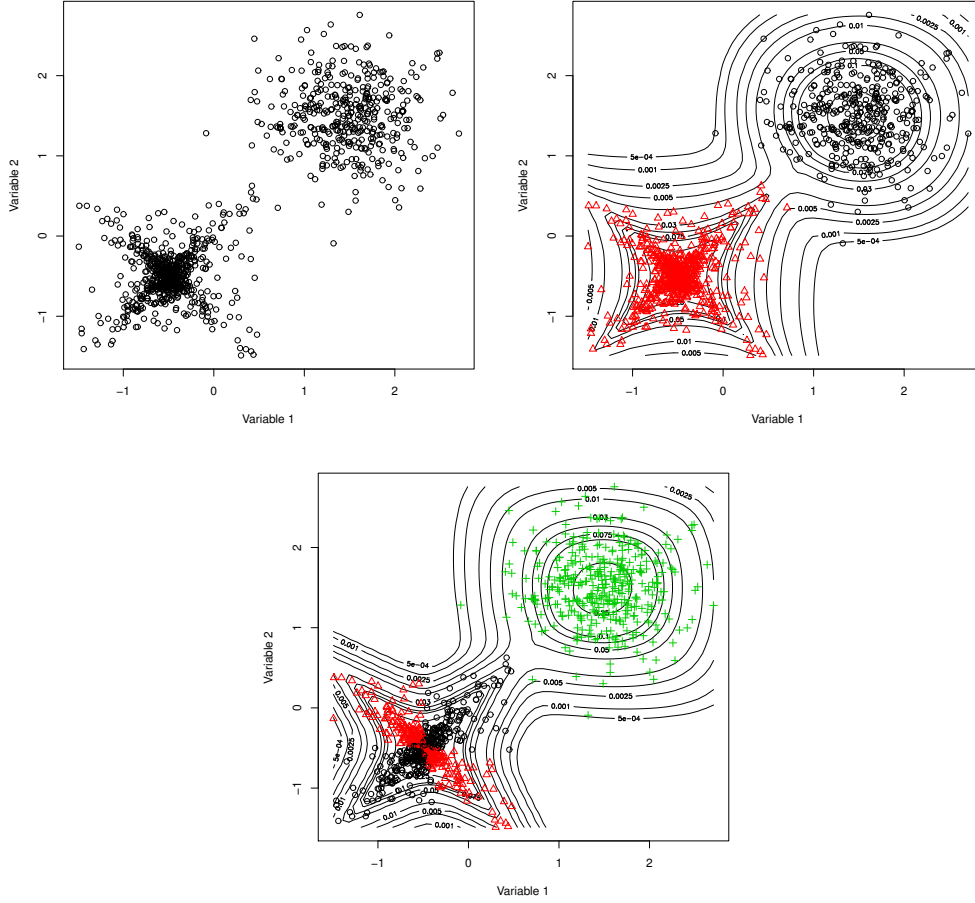


Figure 2: Scatter plots for model-based clustering results on data simulated from a three-component t -mixture (top-left), with contours from the selected MMSGHD (top-right) and McMSGHD (bottom) models, respectively, where plotting symbol and colour represent predicted classifications.

component model being selected (Figure 1). Formally, this amounts to insuring that the MSGHD is quasi-convex; see Appendix B. The general point here is that if convexity is not enforced, then the MSGHD can give components that contain multiple clusters. While it is easy to spot this in two dimensions, e.g., Figure 1, this phenomenon may go unrecognized in higher dimensions, possibly resulting in greatly misleading results. Of course, the issue of non-convex clusters does not arise with most model-based approaches; however, when multiple scaled mixtures are considered, the issue can crop up. Another example in a similar vein is given in Figure 2, where data are generated from a $G = 3$ component mixture of multivariate t -distributions. The selected MMSGHD has $G = 2$ components, including one clearly non-convex cluster, while the McMSGHD gives sensible clustering results.

The intention behind the introduction of the McMSGHD is not that it should supplant the

MMSGHD, but rather that it provides a convenient check on the number of components. In a higher dimensional application, where visualization is difficult or impossible, situations where the selected MMSGHD has fewer components than the selected McMSGHD will deserve special attention. Of course, this is not to say that the selected McMSGHD will always have more components in situations where the MSGHD has too few, but rather that it will help to avoid the sort of situations depicted in Figures 1 and 2. Note that parameter estimation for the McMSGHD is analogous to the algorithm for the MMSGHD algorithm but λ_{kj} ($k = 0, 1, \dots, G$, $j = 1, \dots, p$) is updated only if the value of its update exceeds 1.

6 Illustrations

6.1 Implementation and Evaluation

In the following illustrations, we fit the MGHs, the MMSGHDs, the McMSGHDs, and the MCGHDs using the corresponding functions available in the `MixGHD` package (Tortora et al., 2017) for R (R Core Team, 2017). We use k -means and k -medoids clustering to initialize the \hat{z}_{ig} . The adjusted Rand index (ARI; Hubert and Arabie, 1985) is used to compare predicted classifications with true classes. The ARI corrects the Rand index (Rand, 1971) for chance, its expected value under random classification is 0, and it takes a value of 1 when there is perfect class agreement. Steinley (2004) gives guidelines for interpreting ARI values.

Several mixtures of skewed distributions have been proposed for model-based clustering and classification. Among them, three different formulations of the multivariate skew- t distribution have been used for clustering in the mixture setting. The formulation used by Murray et al. (2014) is a special case of the generalized hyperbolic distributions. Azzalini et al. (2016) refers to the other two formulations as the classical and SDB formulations (for the initials of the authors' names), respectively. Lee and McLachlan (2013b, 2014) compare the mixture of classical skew- t distributions to the mixture of SDB skew- t distributions (MSDBST) and to skew-normal analogues of both formulations; their results lead one to believe that the MSDBST is generally preferable (see Azzalini et al., 2016, for an alternative viewpoint). The MSDBST approach is implemented using the `EMMIXuskew` package (Lee and McLachlan, 2013a) for R and it is used for comparison herein.

6.2 Simulation Study

We use a simulation study to measure the performance of the proposed methods. The interest is to observe the ARI under different circumstances, the data are generated using two-component mixtures of Gaussian (Scenario 1), generalized hyperbolic (Scenario 2), and multiple scaled generalized hyperbolic distributions (Scenario 3), using the R function `mvrnorm` and the stochastic relationships given in (14) and Section 4. For each scenario a three factors full factorial design was used, where the factors are: medium (M) or high (H) correlation, 25% or 50% overlapping, and same (S), $n_1 = n_2 = 100$, or different (D),

$n_1 = 150$ and $n_2 = 50$, number of elements per component. In all the scenarios $\mu_1 = \mathbf{1}$. Table 1 shows the average ARI, with standard deviation, obtained on 10 data sets generated from a Gaussian distribution with $p = 10$, $G = 2$, correlation equal to 0.33 for M and 0.66 for H, $\mu_2 = \mathbf{2}$ for 25% overlapping and $\mu_2 = \mathbf{3}$ for 50% overlapping. Table 2 shows the average ARI, with standard deviation, obtained on 10 data sets generated from a MGHD distribution with $p = 10$, $G = 2$, $\alpha = \mathbf{1}$, $\lambda = 0.5$, $\omega = 1$ and correlation equal to 0.25 for M and 0.5 for H, $\mu_2 = \mathbf{10}$ for 25% overlapping and $\mu_2 = \mathbf{15}$ for 50% overlapping. Table 3 shows the average ARI, with 1 standard deviation, obtained on 10 data sets generated from a MMSGHD distribution with $p = 10$, $G = 2$, $\alpha_1 = \mathbf{1}$, $\alpha_2 = -\mathbf{1}$, $\lambda = \mathbf{0.5}$, $\omega = \mathbf{1}$ and correlation equal to 0.25 for M and 0.35 for H, $\mu_2 = \mathbf{19}$ for 25% overlapping and $\mu_2 = \mathbf{23}$ for 50% overlapping. For the data sets generated using the MGHD and the MMSGHD the values of the correlation had to be reduced in order to maintain 25% and 50% overlap, higher correlation lead to higher overlap. Figure 3 shows the scatterplots of the data simulated using the two-component MGDs, MGHDs, and MMSGHDs with high correlation, 50% overlapping and cluster of different size, where plotting symbol and colour represent the true classifications.

Table 1: Average ARI values, with standard deviations, for each data set generated from MGDs, with different starting parameters.

| Correlation | Overlapping | n_g | MCGHD | MGHD | MMSGHD | McMSGHD |
|-------------|-------------|-------|---------------|---------------|---------------|---------------|
| M | 25% | S | 0.964 (0.020) | 0.932 (0.056) | 0.943 (0.025) | 0.941 (0.026) |
| H | 25% | S | 0.850 (0.044) | 0.730 (0.132) | 0.848 (0.050) | 0.817 (0.062) |
| M | 50% | S | 0.734 (0.064) | 0.642 (0.085) | 0.736 (0.062) | 0.721 (0.084) |
| H | 50% | S | 0.565 (0.116) | 0.408 (0.179) | 0.595 (0.059) | 0.553 (0.090) |
| M | 25% | D | 0.981 (0.016) | 0.979 (0.017) | 0.981 (0.021) | 0.975 (0.022) |
| H | 25% | D | 0.831 (0.047) | 0.771 (0.047) | 0.780 (0.082) | 0.787 (0.059) |
| M | 50% | D | 0.715 (0.053) | 0.663 (0.064) | 0.705 (0.069) | 0.696 (0.073) |
| H | 50% | D | 0.494 (0.133) | 0.354 (0.151) | 0.446 (0.122) | 0.447 (0.130) |

The MCGHDs perform at least comparably to the best of the MMSGHDs and MGHDs. As expected, the MGHDs over performs the MMSGHDs when the data are generated using a MGHDs model and *vice versa* when the data are generated from a MMSGHDs. The two methods perform similarly on normally distributed clusters. The generated clusters are all convex and that explains the similarity of the results obtained using MMSGHDs and McMSGHDs. The level of overlapping has a notable impact on the ARI, as expected. The performance of all the methods are slightly better when the clusters are of the same size and when there is a lower correlation.

Table 2: Average ARI values, with standard deviations, for each data set generated from MGHs, with different starting parameters and $\alpha = 1$.

| Correlation | Overlapping | n_g | MCGHD | MGHD | MMSGHD | McMSGHD |
|-------------|-------------|-------|---------------|---------------|---------------|---------------|
| M | 25% | S | 0.933 (0.045) | 0.922 (0.041) | 0.765 (0.228) | 0.773 (0.211) |
| H | 25% | S | 0.842 (0.040) | 0.770 (0.119) | 0.463 (0.348) | 0.399 (0.371) |
| M | 50% | S | 0.712 (0.093) | 0.680 (0.107) | 0.510 (0.229) | 0.477 (0.221) |
| H | 50% | S | 0.475 (0.176) | 0.416 (0.104) | 0.188 (0.180) | 0.159 (0.113) |
| M | 25% | D | 0.836 (0.058) | 0.858 (0.058) | 0.579 (0.233) | 0.557 (0.233) |
| H | 25% | D | 0.754 (0.106) | 0.713 (0.124) | 0.257 (0.250) | 0.318 (0.283) |
| M | 50% | D | 0.630 (0.092) | 0.622 (0.098) | 0.144 (0.134) | 0.194 (0.198) |
| H | 50% | D | 0.427 (0.090) | 0.360 (0.086) | 0.121 (0.109) | 0.127 (0.093) |

Table 3: Average ARI, with standard deviations, for each data set generated from MMSGHDs, with different starting parameters and $\alpha = 1$.

| Correlation | Overlapping | n_g | MCGHD | MGHD | MMSGHD | McMSGHD |
|-------------|-------------|-------|---------------|---------------|---------------|---------------|
| M | 25% | S | 0.903 (0.047) | 0.801 (0.077) | 0.912 (0.043) | 0.922 (0.045) |
| H | 25% | S | 0.806 (0.085) | 0.697 (0.069) | 0.793 (0.073) | 0.810 (0.060) |
| M | 50% | S | 0.751 (0.080) | 0.614 (0.081) | 0.811 (0.053) | 0.822 (0.044) |
| H | 50% | S | 0.592 (0.062) | 0.475 (0.121) | 0.603 (0.090) | 0.634 (0.078) |
| M | 25% | D | 0.925 (0.042) | 0.838 (0.108) | 0.950 (0.042) | 0.950 (0.037) |
| H | 25% | D | 0.813 (0.049) | 0.740 (0.088) | 0.824 (0.066) | 0.815 (0.070) |
| M | 50% | D | 0.786 (0.090) | 0.616 (0.104) | 0.742 (0.126) | 0.746 (0.120) |
| H | 50% | D | 0.605 (0.111) | 0.511 (0.118) | 0.391 (0.231) | 0.402 (0.224) |

6.3 Real Data Analysis

To assess the classification performance of the MGHs, the MMSGHDs, the McMSGHDs, the MCGHDs, and the MSDBST, we consider four real data sets that are commonly used within the model-based clustering literature (Table 4).

Table 4: Summary details for four real data sets commonly used within the model-based clustering literature.

| | Classes | n | p | Source | Original source |
|------------|---------|-----|-----|----------------------------------|-----------------------------|
| Bankruptcy | 2 | 66 | 2 | R package MixGHD | Altman (1968) |
| Banknote | 2 | 200 | 7 | R package MixGHD | Flury and Riedwyl (1988) |
| Seeds | 3 | 210 | 7 | UCI machine learning repository* | Charytanowicz et al. (2010) |
| AIS | 2 | 202 | 11 | R package EMMIXuskew | Cook and Weisberg (1994) |

*<http://archive.ics.uci.edu/ml/>

To compare the performance of the methods we set the number of components, G , equal to the true number of classes. In general, the BIC is used to select the best starting criterion

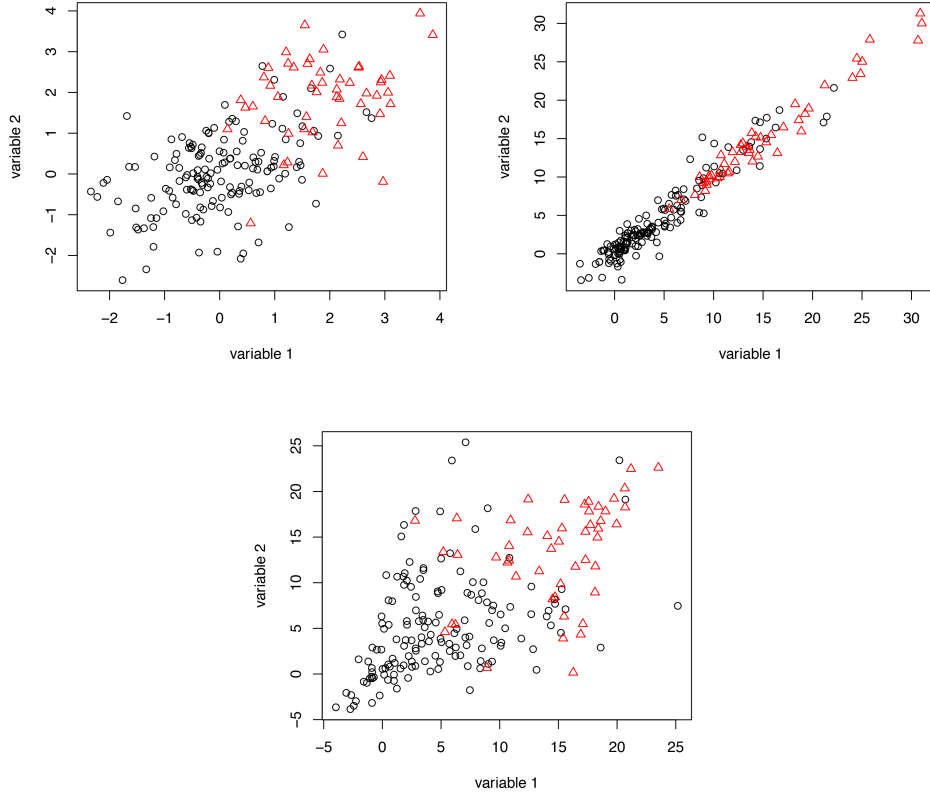


Figure 3: Scatter plots for data simulated from a two-component MGDs (top-left), MGHDs (top-right) and MMSGHDs (bottom) models, respectively, with high correlation, 50% overlapping and clusters of different size, where plotting symbol and colour represent the true classifications.

between k -means and k -medoids; however, because of the presence of outliers, only k -medoids starts are used for the bankruptcy data set. Table 5 displays the classification performance. Note that MSDBST is not used on the AIS data set because of the prohibitively high dimensionality ($p = 11$). The MCGHD generally performs comparably to the best of the other four approaches. Specifically, the MCGHD gives the best classification performance — either outright or jointly — for the four data sets. Interestingly, the MGHD gives very good classification performance on three of the four data sets; however, this must be taken in context with its very poor classification performance on the bankruptcy data set ($\text{ARI} \approx 0$). It is also interesting to compare the classification performance of the McMSGHD to the MMSGHD as well as that of the MGHD to the MMSGHD. The McMSGHD and the MGHD approaches both outperform MMSGHD for one data set, and give a similar performance on two of the other three data sets. This highlights the fact that a mixture of multiple scaled

distributions may well not outperform its single scaled analogue, and underlines the need for an approach with both the MSGHD and MGHD models as special cases. The results for the bankruptcy data illustrate that the MCGHD approach can give very good classification performance in situations where neither the MGHD nor the MMSGHD perform well. Finally, the MCGHD outperforms the MSDBST on two data sets and gives the same result on the third (recall that the MSDBST could not be fitted to the AIS data).

Table 5: ARI values for the MCGHD, MGHD, MMSGHD, McMSGHD, and MSDBST approaches on four real data sets.

| Data | G | MCGHD | MGHD | MMSGHD | McMSGHD | MSDBST |
|------------|-----|-------|-------|--------|---------|--------|
| Bankruptcy | 2 | 0.824 | 0.019 | 0.255 | 0.170 | 0.085 |
| Bank note | 2 | 0.980 | 0.980 | 0.980 | 0.980 | 0.980 |
| Seeds | 3 | 0.775 | 0.617 | 0.519 | 0.533 | 0.723 |
| AIS | 2 | 0.903 | 0.884 | 0.884 | 0.865 | NA |

In this analysis, we took the number of components to be known. However, in a true clustering scenario, we would not have *a priori* information about the number of groups. Therefore, the analysis was repeated without this assumption and, for all but the bankruptcy data set, all approaches were run for $G = 1, \dots, 5$ components. Because of the small number of observations, the bankruptcy data were run for $G = 1, 2, 3$. Table 6 gives the number of components selected by the BIC.

Table 6: Selected number of components using the BIC for each real data set and corresponding ARI in brackets, where the correct number of components is highlighted in bold face font.

| Data | Classes | MCGHD | MGHD | MMSGHD | McMSGHD | MSDBST |
|------------|---------|------------------|------------------|------------------|------------------|------------------|
| Bankruptcy | 2 | 1(0.000) | 1(0.000) | 1(0.000) | 2 (0.170) | 1(0.000) |
| Bank note | 2 | 2 (0.980) | 2 (0.980) | 2 (0.980) | 2 (0.980) | 2 (0.980) |
| Seeds | 3 | 2(0.502) | 4(0.484) | 2(0.530) | 2(0.530) | 1(0.000) |
| AIS | 2 | 4(0.435) | 3(0.615) | 1(0.000) | 1(0.000) | NA |

For the bank note data, the BIC selects the correct number of components for every method; however, for the bankruptcy data, it picks the right number of components only for McMSGHDs. For the seed and AIS data sets, the BIC does not select the correct number of components for any approach. Recall that the MSDBST could not be fitted to the AIS data; also, it could not be fitted to the seed data for $G = 4$. These results illustrate that the BIC is not necessarily reliable for selecting the number of components in real data examples for any of the five approaches.

6.4 Computation time

For the data sets listed in Table 4, we measured the elapsed user times, in seconds, to perform 100 (G)EM iterations. Note that all code was run in R version 3.0.2 on a 32-core

Intel Xeon E5 server with 256GB RAM running 64-bit CentOS. Figure 4 displays the average elapsed time for two replications of each algorithm using a k -means and a k -medoids starting partition for $G = 1, \dots, 5$ components. The EM algorithm for the MSDBST is significantly slower than that of the hyperbolic-based approaches (Figure 4). In fact, on the banknote data set, it takes the MSDBST more than 11 hours to perform the required number of EM iterations when $G = 2$ and about 63 hours when $G = 5$. For the seeds data set, it takes the MSDBST more than 5 hours when $G = 2$ and more than 10 hours when $G = 5$, whereas the $G = 5$ component MMSGHD, McMSGHD, and MCGHD approaches need less than 25 seconds. For the bankruptcy data set, the MSDBST requires approximately 80 seconds when $G = 3$, whereas the hyperbolic distributions need less than 20 seconds each.

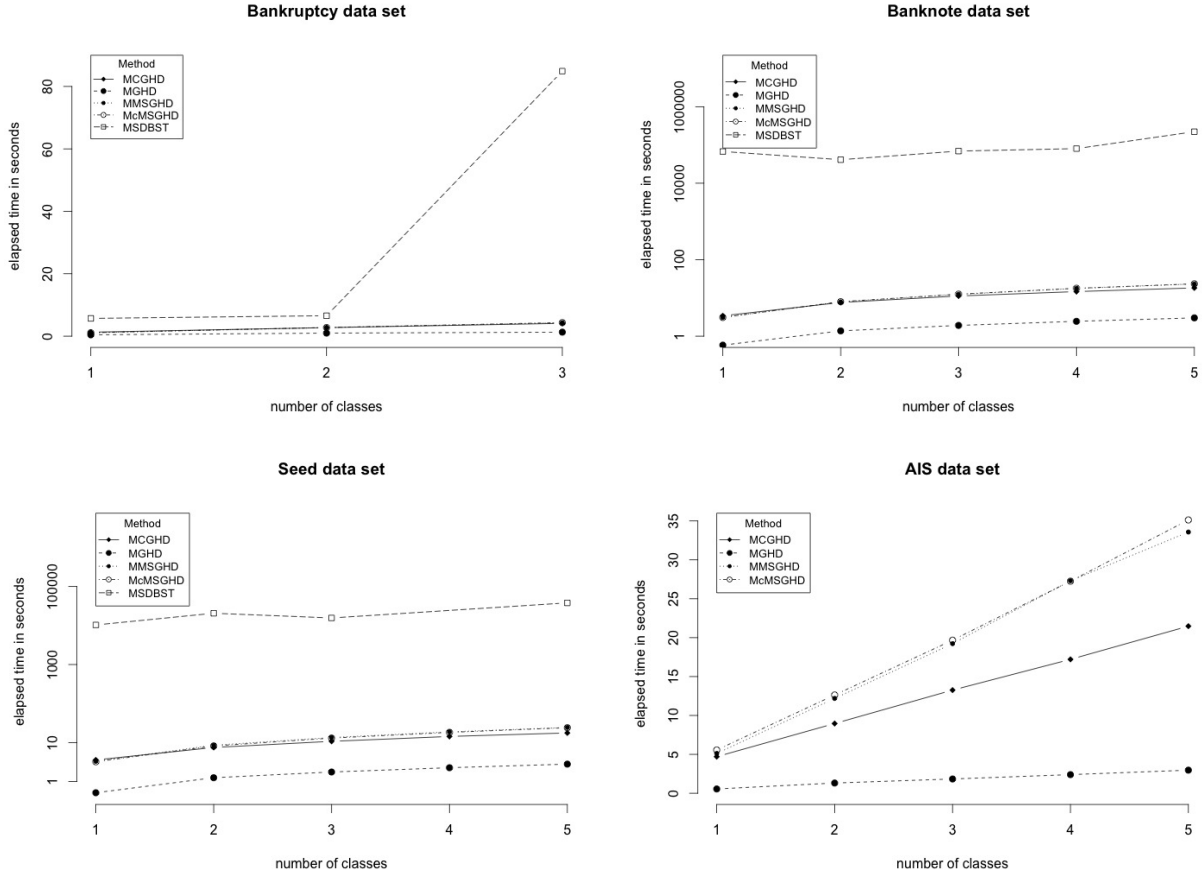


Figure 4: The average elapsed time to perform 100 iterations of the (G)EM algorithm when varying the number of components for MCGHD, MGHD, MMSGHD, McMSGHD, and MSDBST models on the bankruptcy, banknote, seeds, and AIS data sets, respectively.

7 Discussion

Novel MCGHDs, MMSGHDs, and McMSGHDs models have been introduced and applied for model-based clustering. The GHD is a flexible distribution, capable of handling skewness and heavy tails, and has many well known distributions as special or limiting cases. Furthermore, it is a normal variance-mean mixture, arising via a relationship between a multivariate Gaussian and an univariate GIG distribution. The MSGHD extends the GHD to include a multivariate GIG distribution, increasing the flexibility of the model. However, the GHD is not a special case of the MSGHD; hence, we created MCGHDs, which has both the GHD and MSGHD as special cases. The McMSGHD approach was introduced as a convex version of the MMSGHD, and this point deserves some further discussion. The extension of the multivariate- t distribution to multiple scale was carried out by Forbes and Wraith (2014); as discussed in the Appendix B, the multiple scaled multivariate t -distribution cannot be quasi-concave, i.e., the clusters associated with a mixture of multiple scaled multivariate t -distributions cannot be convex. We have seen examples where the MMSGHD can put multiple clusters into one component, and the McMSGHD has an important role in helping to prevent this; if both approaches are fitted and lead to different numbers of components, then further attention is warranted.

Comparing the MGHD, MMSGHD, McMSGHD, and MCGHD approaches yielded some interesting results. Amongst them, we see that the MMSGHD does not necessarily outperform the MGHD; far from it, in fact, only the MGHD approach gave better clustering performance than the MMSGHD approach on one of the four real data sets we considered, as well as identical performance on a other two. This underlines the fact that a mixture of multiple scaled distributions may well not outperform its single scaled analogue, and highlights the benefit of approaches with both a multiple scaled distribution and its single scaled analogue as special cases. The MCGHDs represent one such approach, with the MSGHD and MGHD models as special cases. The approaches introduced herein, as well as the MGHDs, have been made freely available via the `MixGHD` package for R.

Future work will focus in several directions. For one, it will be interesting to study the performance of the approaches introduced herein within the fractionally-supervised classification framework (Vrbik and McNicholas, 2015; Gallagher and McNicholas, 2019a). The extension of these approaches to the matrix-variate paradigm is also of interest and may proceed in an analogous fashion to the work of Gallagher and McNicholas (2018, 2019b). Finally, the models introduced herein can be extended to account for missing data (see Wei et al., 2019).

Acknowledgements

This work was supported by a grant-in-aid from Compusense Inc., a Collaborative Research and Development Grant from the Natural Sciences and Engineering Research Council of Canada, and an Early Researcher Award from the Ontario Ministry of Research and Innovation.

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.
- Altman, E. (1968). Financial ratios, discriminant analysis and the prediction of corporate bankruptcy. *Journal of Finance* 23(4), 589–609.
- Andrews, J. L. and P. McNicholas (2012). Model-based clustering, classification, and discriminant analysis via mixtures of multivariate t -distributions. *Statistics and Computing* 22(5), 1021–1029.
- Andrews, J. L. and P. D. McNicholas (2011a). Extending mixtures of multivariate t -factor analyzers. *Statistics and Computing* 21(3), 361–373.
- Andrews, J. L. and P. D. McNicholas (2011b). Mixtures of modified t -factor analyzers for model-based clustering, classification, and discriminant analysis. *Journal of Statistical Planning and Inference* 141(4), 1479–1486.
- Azzalini, A., R. P. Browne, M. G. Genton, and P. D. McNicholas (2016). On nomenclature for, and the relative merits of, two formulations of skew distributions. *Statistics and Probability Letters* 110, 201–206.
- Baek, J. and G. J. McLachlan (2011). Mixtures of common t -factor analyzers for clustering high-dimensional microarray data. *Bioinformatics* 27, 1269–1276.
- Banfield, J. D. and A. E. Raftery (1993). Model-based Gaussian and non-Gaussian clustering. *Biometrics* 49(3), 803–821.
- Barndorff-Nielsen, O. (1978). Hyperbolic distributions and distributions on hyperbolae. *Scandinavian Journal of Statistics* 5(3), 151–157.
- Barndorff-Nielsen, O., J. Kent, and M. Sørensen (1982). Normal variance-mean mixtures and z distributions. *International Statistical Review / Revue Internationale de Statistique* 50(2), 145–159.
- Böhning, D., E. Dietz, R. Schaub, P. Schlattmann, and B. Lindsay (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 P. D. McNicholas (2014). Estimating common principal components in high dimensions. *Advances in Data Analysis and Classification* 8(2), 217–226.
- Browne, R. P. and P. D. McNicholas (2015). A mixture of generalized hyperbolic distributions. *Canadian Journal of Statistics* 43(2), 176–198.
- Celeux, G. and G. Govaert (1995). Gaussian parsimonious clustering models. *Pattern Recognition* 28(5), 781–793.

- Charytanowicz, M., J. Niewczas, P. Kulczycki, P. A. Kowalski, S. Łukasik, and S. Żak (2010). A complete gradient clustering algorithm for features analysis of x-ray images. In E. Piętko and J. Kawa (Eds.), *Information Technologies in Biomedicine: Volume 2*, pp. 15–24. Berlin, Heidelberg: Springer.
- Cook, R. D. and S. Weisberg (1994). *An Introduction to Regression Graphics*. John Wiley & Sons, New York.
- Cormack, R. M. (1971). A review of classification (with discussion). *Journal of the Royal Statistical Society: Series A* 34, 321–367.
- Debreu, G. and T. C. Koopmans (1982). Additively decomposed quasiconvex functions. *Mathematical Programming* 24(1), 1–38.
- Demarta, S. and A. J. McNeil (2005). The t copula and related copulas. *International Statistical Review* 73(1), 111–129.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society: Series B* 39(1), 1–38.
- Flury, B. and H. Riedwyl (1988). *Multivariate Statistics: A Practical Approach*. London: Chapman & Hall.
- Forbes, F. and D. Wraith (2014). A new family of multivariate heavy-tailed distributions with variable marginal amounts of tailweights: Application to robust clustering. *Statistics and Computing* 24(6), 971–984.
- Fraley, C. and A. E. Raftery (2002). Model-based clustering, discriminant analysis, and density estimation. *Journal of the American Statistical Association* 97(458), 611–631.
- Franczak, B. C., R. P. Browne, and P. D. McNicholas (2014). Mixtures of shifted asymmetric Laplace distributions. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36(6), 1149–1157.
- Franczak, B. C., C. Tortora, R. P. Browne, and P. D. McNicholas (2015). Unsupervised learning via mixtures of skewed distributions with hypercube contours. *Pattern Recognition Letters* 58(1), 69–76.
- Gallaughier, M. P. B. and P. D. McNicholas (2018). Finite mixtures of skewed matrix variate distributions. *Pattern Recognition* 80, 83–93.
- Gallaughier, M. P. B. and P. D. McNicholas (2019a). On fractionally-supervised classification: Weight selection and extension to the multivariate t-distribution. *Journal of Classification* 36. In press.
- Gallaughier, M. P. B. and P. D. McNicholas (2019b). Three skewed matrix variate distributions. *Statistics and Probability Letters* 145, 103–109.

- Ghahramani, Z. and G. E. Hinton (1997). The EM algorithm for factor analyzers. Technical Report CRG-TR-96-1, University Of Toronto, Toronto.
- Gneiting, T. (1997). Normal scale mixtures and dual probability densities. *Journal of Statistical Computation and Simulation* 59(4), 375–384.
- Hennig, C. (2015). What are the true clusters? *Pattern Recognition Letters* 63, 53–62.
- Holzmann, H., A. Munk, and T. Gneiting (2006). Identifiability of finite mixtures of elliptical distributions. *Scandinavian Journal of Statistics* 33, 753–763.
- Hubert, L. and P. Arabie (1985). Comparing partitions. *Journal of Classification* 2(1), 193–218.
- Hunter, D. R. and K. Lange (2000). Quantile regression via an MM algorithm. *Journal of Computational and Graphical Statistics* 9(1), 60–77.
- Karlis, D. and A. Santourian (2009). Model-based clustering with non-elliptically contoured distributions. *Statistics and Computing* 19(1), 73–83.
- Kent, J. T. (1983). Identifiability of finite mixtures for directional data. *The Annals of Statistics* 11, 984–988.
- Kiers, H. A. (2002). Setting up alternating least squares and iterative majorization algorithms for solving various matrix optimization problems. *Computational Statistics and Data Analysis* 41(1), 157–170.
- Kotz, S., T. J. Kozubowski, and K. Podgorski (2001). *The Laplace Distribution and Generalizations: A Revisit with Applications to Communications, Economics, Engineering, and Finance* (1st ed.). Burkhauser Boston.
- Kotz, S. and S. Nadarajah (2004). *Multivariate t-distributions and their applications*. Cambridge University Press.
- Lee, S. X. and G. J. McLachlan (2013a). *EMMIXuskew: Fitting Unrestricted Multivariate Skew t Mixture Models*. R package version 0.11-5.
- Lee, S. X. and G. J. McLachlan (2013b). On mixtures of skew normal and skew t-distributions. *Advances in Data Analysis and Classification* 7(3), 241–266.
- Lee, S. X. and G. J. McLachlan (2014). Finite mixtures of multivariate skew t-distributions: some recent and new results. *Statistics and Computing* 24(2), 181–202.
- Lin, T. I. (2009). Maximum likelihood estimation for multivariate skew normal mixture models. *Journal of Multivariate Analysis* 100(2), 257–265.
- Lin, T. I. (2010). Robust mixture modeling using multivariate skew t distributions. *Statistics and Computing* 20(3), 343–356.

- Lin, T.-I., P. D. McNicholas, and J. H. Hsiu (2014). Capturing patterns via parsimonious t mixture models. *Statistics and Probability Letters* 88, 80–87.
- Lindsay, B. (1995). Mixture models: Theory, geometry and applications. In *NSF-CBMS Regional Conference Series in Probability and Statistics*, Volume 5, California: Institute of Mathematical Statistics: Hayward.
- McLachlan, G. J., R. W. Bean, and L. B.-T. Jones (2007). Extension of the mixture of factor analyzers model to incorporate the multivariate t-distribution. *Computational Statistics and Data Analysis* 51(11), 5327–5338.
- McLachlan, G. J. and T. Krishnan (2008). *The EM Algorithm and Extensions*. New York: Wiley.
- McLachlan, G. J. and D. Peel (2000). Mixtures of factor analyzers. In *Proceedings of the Seventh International Conference on Machine Learning*, San Francisco, pp. 599–606. Morgan Kaufmann.
- McNeil, A. J., R. Frey, and P. Embrechts (2005). *Quantitative risk management: concepts, techniques and tools*. Princeton university press.
- McNicholas, P. D. (2016a). *Mixture Model-Based Classification*. Boca-Raton: Chapman & Hall/CRC press.
- McNicholas, P. D. (2016b). Model-based clustering. *Journal of Classification* 33(3), 331–373.
- McNicholas, P. D., T. B. Murphy, A. F. McDaid, and D. Frost (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., P. D. McNicholas, and R. P. Browne (2017). A mixture of variance-gamma factor analyzers. In S. E. Ahmed (Ed.), *Big and Complex Data Analysis: Methodologies and Applications*, pp. 369–385. Cham: Springer International Publishing.
- Murray, P. M., R. B. Browne, and P. D. McNicholas (2014). Mixtures of skew-t factor analyzers. *Computational Statistics and Data Analysis* 77, 326–335.
- Murray, P. M., R. B. Browne, and P. D. McNicholas (2017). Hidden truncation hyperbolic distributions, finite mixtures thereof, and their application for clustering. *Journal of Multivariate Analysis* 161, 141–156.
- Niculescu, C. and L. Persson (2006). *Convex Functions and Their Applications*. New York: Springer.
- Ortega, J. M. and W. C. Rheinboldt (1970). *Iterative Solutions of Nonlinear Equations in Several Variables*. New York: Academic Press.
- Peel, D. and G. J. McLachlan (2000). Robust mixture modelling using the t distribution. *Statistics and Computing* 10(4), 339–348.

- Pesevski, A., B. C. Franczak, and P. D. McNicholas (2018). Subspace clustering with the multivariate-t distribution. *Pattern Recognition Letters* 112(1), 297–302.
- R Core Team (2017). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Rand, W. M. (1971). Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association* 66(336), 846–850.
- Rockafellar, R. T. and R. J. B. Wets (2009). *Variational Analysis*. New York: Springer-Verlag.
- Schwarz, G. (1978). Estimating the dimension of a model. *Annals of Statistics* 6(2), 461–464.
- Steane, M. A., P. D. McNicholas, and R. Yada (2012). Model-based classification via mixtures of multivariate t-factor analyzers. *Communications in Statistics – Simulation and Computation* 41(4), 510–523.
- Steinley, D. (2004). Properties of the Hubert-Arable adjusted Rand index. *Psychological methods* 9(3), 386.
- Tang, Y., R. P. Browne, and P. D. McNicholas (2018). Flexible clustering of high-dimensional data via mixtures of joint generalized hyperbolic distributions. *Stat* 7(1), e177.
- Tipping, M. E. and C. M. Bishop (1999). Mixtures of probabilistic principal component analysers. *Neural Computation* 11(2), 443–482.
- Tortora, C., R. P. Browne, B. C. Franczak, and P. D. McNicholas (2017). MixGHD: Model based clustering, classification and discriminant analysis using the mixture of generalized hyperbolic distributions. R package version 2.1.
- Vrbik, I. and P. D. McNicholas (2012). Analytic calculations for the EM algorithm for multivariate skew-mixture models. *Statistics and Probability Letters* 82(6), 1169–1174.
- Vrbik, I. and P. D. McNicholas (2014). Parsimonious skew mixture models for model-based clustering and classification. *Computational Statistics and Data Analysis* 71, 196–210.
- Vrbik, I. and P. D. McNicholas (2015). Fractionally-supervised classification. *Journal of Classification* 32(3), 359–381.
- Wei, Y., Y. Tang, and P. D. McNicholas (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.
- Wraith, D. and F. Forbes (2015). Clustering using skewed multivariate heavy tailed distributions with flexible tail behaviour. arXiv preprint arXiv:1408.0711.
- Yakowitz, S. J. and J. Spragins (1968). On the identifiability of finite mixtures. *Ann. Math. Statist.* 39, 209–214.

A Parameter estimation

We use the EM algorithm to estimate the parameters of the MCGHDs. The EM algorithm belongs to a larger class of algorithms known as MM algorithms (Ortega and Rheinboldt, 1970; Hunter and Lange, 2000) and is well-suited for problems involving missing data. ‘MM’ stands for ‘minorize-maximize’ or ‘majorize-minimize,’ depending on the purpose of the algorithm; in the EM context, the minorizing function is the expected value of the complete-data log-likelihood. The EM algorithm iterates between two steps, an E-step and a M-step, and has been used to estimate the parameters of mixture models in many experiments (McLachlan and Krishnan, 2008). On each E-step, the expected value of the complete-data log-likelihood, \mathcal{Q} , is calculated and on each M-step is maximized with respect to $\pi_g, \boldsymbol{\mu}_g, \boldsymbol{\Phi}_g, \boldsymbol{\alpha}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g, \omega_{0g}, \lambda_{0g}, \varpi_g$. However, in each M-step \mathcal{Q} increases with respect to Γ_g rather than maximize; accordingly, the algorithm is formally a generalized EM (GEM) algorithm. For our MCGHDs, there are four sources of missing data: the latent variable W_{0ig} , the multi-dimensional weight variable $\boldsymbol{\Delta}_{wig}$, the group component indicator labels z_{ig} , and inner component labels u_{ig} , for $i = 1, \dots, n$ and $g = 1, \dots, G$. For each observation i , $z_{ig} = 1$ if observation i is in component g and $z_{ig} = 0$ otherwise. Similarly, for each observation i , $u_{ig} = 1$ if observation i , in component g , is distributed generalized hyperbolic and $u_{ig} = 0$ if observation i , in component g , is distributed multiple scaled generalized hyperbolic. It follows that the complete-data log-likelihood for the MCGHDs is given by

$$l_c = \sum_{i=1}^n \sum_{g=1}^G \left\{ z_{ig} \log \pi_g + z_{ig} u_{ig} \log \varpi_g + z_{ig} (1 - u_{ig}) \log (1 - \varpi_g) + z_{ig} u_{ig} \log h(w_{0ig} \mid \omega_{0g}, 1, \lambda_{0g}) \right. \\ \left. + z_{ig} (1 - u_{ig}) \sum_{j=1}^p \log h(w_{jig} \mid \omega_{jg}, 1, \lambda_{jg}) + z_{ig} u_{ig} \log \phi_p(\boldsymbol{\Gamma}'_g \mathbf{x}_i \mid \boldsymbol{\mu}_g + w_{0ig} \boldsymbol{\alpha}_g, w_{0ig} \boldsymbol{\Phi}) \right. \\ \left. + z_{ig} (1 - u_{ig}) \sum_{j=1}^p \log \phi_1([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j \mid \mu_{jg} + w_{jig} \alpha_{jg}, \omega_{jg} \phi_{jg}) \right\},$$

where $\phi_p(\cdot)$ represents a p -dimensional Gaussian density function, $\phi_1(\cdot)$ is a unidimensional Gaussian density function, and $h(\cdot)$ is the density of a GIG distribution given in (15).

We are now prepared to outline the calculations for our GEM algorithm for the MCGHDs. On the E-step, the expected value of the complete-data log-likelihood, \mathcal{Q} , is computed by replacing the sufficient statistics of the missing data by their expected values. For each component indicator label z_{ig} and inner component label u_{ig} , for $i = 1, \dots, n$ and $g = 1, \dots, G$, we require the expectations

$$\mathbb{E}[Z_{ig} \mid \mathbf{x}_i] = \frac{\pi_g f_{\text{CGHD}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\alpha}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g, \omega_{0g}, \lambda_{0g}, \varpi_g)}{\sum_{h=1}^G \pi_h f_{\text{CGHD}}(\mathbf{x} \mid \boldsymbol{\mu}_h, \boldsymbol{\Gamma}_h, \boldsymbol{\Phi}_h, \boldsymbol{\alpha}_h, \boldsymbol{\omega}_h, \boldsymbol{\lambda}_h, \omega_{0h}, \lambda_{0h}, \varpi_h)} =: \hat{z}_{ig} \quad (21)$$

and

$$\mathbb{E}[U_{ig} \mid \mathbf{x}_i, z_{ig} = 1] = \frac{\varpi_g f_{\text{GHD}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\alpha}_g, \omega_{0g}, \lambda_{0g})}{\varpi_g f_{\text{GHD}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\alpha}_g, \omega_{0g}, \lambda_{0g}) + (1 - \varpi_g) f_{\text{MSGHD}}(\mathbf{x} \mid \boldsymbol{\mu}_g, \boldsymbol{\Gamma}_g, \boldsymbol{\Phi}_g, \boldsymbol{\alpha}_g, \boldsymbol{\omega}_g, \boldsymbol{\lambda}_g)} =: \hat{u}_{ig}, \quad (22)$$

where f_{CGHD} is given in (20), f_{GHD} is given in (10) and f_{MSGHD} is given in (17). For the latent variable W_{0ig} , we use the expected value given in Browne and McNicholas (2015). The authors show that, given the density in (15), the following is true

$$W_{0ig} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 1 \sim \mathcal{GIG}(\omega_{0g} + \boldsymbol{\alpha}'_g(\boldsymbol{\Gamma}_g \boldsymbol{\Phi}_g \boldsymbol{\Gamma}'_g)^{-1} \boldsymbol{\alpha}_g, \omega_{0g} + \delta(\mathbf{x}_i, \boldsymbol{\mu}_g \mid \boldsymbol{\Gamma}_g \boldsymbol{\Phi}_g \boldsymbol{\Gamma}'_g), \lambda_{0g} - p/2).$$

For the MCGHDs, the maximization of \mathcal{Q} requires the expected values of W_{0ig} , W_{0ig}^{-1} and $\log W_{0ig}$, i.e.,

$$\begin{aligned} \mathbb{E}[W_{0ig} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 1] &= \sqrt{\frac{e_{ig}}{d_g}} \frac{K_{\lambda_{0g}-p/2+1}(\sqrt{d_g e_{ig}})}{K_{\lambda_{0g}-p/2}(\sqrt{d_g e_{ig}})} =: a_{ig}, \\ \mathbb{E}[W_{0ig}^{-1} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 1] &= \sqrt{\frac{d_g}{e_{ig}}} \frac{K_{\lambda_{0g}-p/2+1}(\sqrt{d_g e_{ig}})}{K_{\lambda_{0g}-p/2}(\sqrt{d_g e_{ig}})} - \frac{2\lambda_{0g} - p}{e_{ig}} =: b_{ig}, \\ \mathbb{E}[\log W_{0ig} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 1] &= \log \sqrt{\frac{e_{ig}}{d_g}} + \frac{\partial}{\partial v} \log \left\{ K_v(\sqrt{d_g e_{ig}}) \right\} \Big|_{v=\lambda_{0g}-p/2} =: c_{ig}, \end{aligned}$$

where $d_g = \omega_{0g} + \boldsymbol{\alpha}'_g(\boldsymbol{\Gamma}_g \boldsymbol{\Phi}_g \boldsymbol{\Gamma}'_g)^{-1} \boldsymbol{\alpha}_g$ and $e_{ig} = \omega_{0g} + \delta(\mathbf{x}_i, \boldsymbol{\mu}_g \mid \boldsymbol{\Gamma}_g \boldsymbol{\Phi}_g \boldsymbol{\Gamma}'_g)$.

The maximization of \mathcal{Q} also requires the expected values of the multidimensional weight variables $\boldsymbol{\Delta}_{\mathbf{wig}}$, $\boldsymbol{\Delta}_{\mathbf{wig}}^{-1}$, and $\log \boldsymbol{\Delta}_{\mathbf{wig}}$. Given the density in (17), it follows that

$$W_{ijg} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 0 \sim \mathcal{GIG}(\omega_{jg} + \alpha_{jg}^2 \boldsymbol{\Phi}_{jg}^{-1}, \omega_{jg} + ([\boldsymbol{\Gamma}' \mathbf{x}]_j - \boldsymbol{\mu}_{gj})^2 / \phi_{jg}, \lambda_{jg} - 1/2).$$

Each multidimensional weight variable is replaced by its expected value and so we need to compute $\mathbf{E}_{1ig} = \text{diag}\{E_{1i1g}, \dots, E_{1ipg}\}$, $\mathbf{E}_{2ig} = \text{diag}\{E_{2i1g}, \dots, E_{2ipg}\}$, and $\mathbf{E}_{3ig} = \text{diag}\{E_{3i1g}, \dots, E_{3ipg}\}$, where

$$\begin{aligned} \mathbb{E}[W_{ijg} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 0] &= \sqrt{\frac{\bar{e}_{ijg}}{\bar{d}_{jg}}} \frac{K_{\lambda_{jg}+1/2}(\sqrt{\bar{d}_{jg} \bar{e}_{ijg}})}{K_{\lambda_{jg}-1/2}(\sqrt{\bar{d}_{jg} \bar{e}_{ijg}})} =: E_{1ijg}, \\ \mathbb{E}[W_{ijg}^{-1} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 0] &= \sqrt{\frac{\bar{d}_{jg}}{\bar{e}_{ijg}}} \frac{K_{\lambda_{jg}+1/2}(\sqrt{\bar{d}_{jg} \bar{e}_{ijg}})}{K_{\lambda_{jg}-1/2}(\sqrt{\bar{d}_{jg} \bar{e}_{ijg}})} - \frac{2\lambda_{jg} - 1}{\bar{e}_{ijg}} =: E_{2ijg}, \\ \mathbb{E}[\log W_{ijg} \mid \mathbf{x}_i, z_{ig} = 1, u_{ig} = 0] &= \log \sqrt{\frac{\bar{e}_{ijg}}{\bar{d}_{jg}}} + \frac{\partial}{\partial v} \log \left\{ K_v(\sqrt{\bar{d}_{jg} \bar{e}_{ijg}}) \right\} \Big|_{v=\lambda_{jg}-1/2} =: E_{3ijg}, \end{aligned} \tag{23}$$

$\bar{d}_{jg} = \omega_{jg} + \alpha_{jg}^2 \boldsymbol{\Phi}_{jg}^{-1}$ and $\bar{e}_{ijg} = \omega_{jg} + ([\mathbf{x}_i - \boldsymbol{\mu}_g]_j)^2 / \phi_{jg}$. Let $n_g = \sum_{i=1}^n \hat{z}_{ig}$, $A_g = (1/n_g) \sum_{i=1}^n \hat{z}_{ig} a_{ig}$, $B_g = (1/n_g) \sum_{i=1}^n \hat{z}_{ig} b_{ig}$, $C_g = (1/n_g) \sum_{i=1}^n \hat{z}_{ig} c_{ig}$, $\bar{E}_{1jg} = (1/n_g) \sum_{i=1}^n \hat{z}_{ig} E_{1ijg}$, $\bar{E}_{2jg} = (1/n_g) \sum_{i=1}^n \hat{z}_{ig} E_{2ijg}$, and $\bar{E}_{3jg} = (1/n_g) \sum_{i=1}^n \hat{z}_{ig} E_{3ijg}$.

In the M-step, we maximize the expected value of the complete-data log-likelihood with respect to the model parameters. The mixing proportions and inner mixing proportions are updated via

$\hat{\pi}_g = n_g/n$ and $\hat{\omega}_g = \sum_{i=1}^n \hat{u}_{ig} \hat{z}_{ig} / n_g$, respectively. The elements of the location parameter $\boldsymbol{\mu}_g$ and skewness parameter $\boldsymbol{\alpha}_g$ are replaced with

$$\hat{\mu}_{jg} = \frac{\sum_{i=1}^n \hat{z}_{ig} [\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j (\bar{s}_{1jg} s_{2ijg} - 1)}{\sum_{i=1}^n \hat{z}_{ig} (\bar{s}_{1jg} s_{2ijg} - 1)} \quad \text{and} \quad \hat{\alpha}_{jg} = \frac{\sum_{i=1}^n \hat{z}_{ig} [\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j (\bar{s}_{2jg} - s_{2ijg})}{\sum_{i=1}^n \hat{z}_{ig} (\bar{s}_{1jg} s_{2ijg} - 1)},$$

respectively, where $[\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j$ is the j th element of the matrix $\boldsymbol{\Gamma}'_g \mathbf{x}_i$, $s_{1ijg} = \hat{u}_{ig} a_{ig} + (1 - \hat{u}_{ig}) E_{1ijg}$, $s_{2ijg} = \hat{u}_{ig} b_{ig} + (1 - \hat{u}_{ig}) E_{2ijg}$, $\bar{s}_{1jg} = 1/n_g \sum_{i=1}^n \hat{z}_{ig} s_{1ijg}$, $\bar{s}_{2jg} = 1/n_g \sum_{i=1}^n \hat{z}_{ig} s_{2ijg}$. The diagonal elements of the matrix $\boldsymbol{\Phi}_g$ are updated using

$$\begin{aligned} \hat{\phi}_{jg} = & \frac{1}{n_g} \sum_{i=1}^n \left\{ \hat{z}_{ig} \hat{u}_{ig} \left[b_{ig} ([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg})^2 - 2 ([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg}) \hat{\alpha}_{jg} + a_{ig} \hat{\alpha}_{jg}^2 \right] \right. \\ & \left. + \hat{z}_{ig} (1 - \hat{u}_{ig}) \left[E_{2ijg} ([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg})^2 - 2 ([\boldsymbol{\Gamma}'_g \mathbf{x}_i]_j - \hat{\mu}_{jg}) \hat{\alpha}_{jg} + E_{1ijg} \hat{\alpha}_{jg}^2 \right] \right\}. \end{aligned}$$

To update the component eigenvector matrices $\boldsymbol{\Gamma}_g$, we wish to minimize the objective function

$$f(\boldsymbol{\Gamma}_g) = -\frac{1}{2} \text{tr} \left\{ \hat{z}_{ig} \hat{\boldsymbol{\Phi}}_g^{-1} \mathbf{V}_{ig} \boldsymbol{\Gamma}_g \mathbf{x}_i \mathbf{x}_i' \boldsymbol{\Gamma}_g' \right\} + \text{tr} \left\{ \hat{z}_{ig} \mathbf{x}_i (\mathbf{V}_{ig} \hat{\boldsymbol{\mu}}_g + \hat{\boldsymbol{\alpha}}_g)' \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{\Gamma}_g \right\} + C \quad (24)$$

with respect to $\boldsymbol{\Gamma}_g$, where $\mathbf{V}_{ig} = \hat{u}_{ig} b_{ig} \mathbf{I}_p + (1 - \hat{u}_{ig}) \mathbf{E}_{2ig}$. We employ an optimization routine that uses two simpler majorization-minimization algorithms. Our optimization routine exploits the convexity of the objective function in (24), providing a computationally stable algorithm for estimating $\boldsymbol{\Gamma}_g$. Specifically, we follow Kiers (2002) and Browne and McNicholas (2014) and use the surrogate function

$$f(\boldsymbol{\Gamma}_g) \leq C + \sum_{i=1}^n \text{tr} \{ \mathbf{F}_{rg} \boldsymbol{\Gamma}_g \}, \quad (25)$$

where C is a constant that does not depend on $\boldsymbol{\Gamma}_g$, $r \in \{1, 2\}$ is an index, and the matrices \mathbf{F}_{rg} are defined in (26) and (27).

Therefore, on each M -step, we calculate either

$$\mathbf{F}_{1g} = \sum_{i=1}^n \hat{z}_{ig} \left[-\mathbf{x}_i (\mathbf{V}_{ig} \hat{\boldsymbol{\mu}}_g + \hat{\boldsymbol{\alpha}}_g)' \hat{\boldsymbol{\Phi}}_g^{-1} + \mathbf{x}_i \mathbf{x}_i' \boldsymbol{\Gamma}_g' \hat{\boldsymbol{\Phi}}_g^{-1} \mathbf{V}_{ig} - \alpha_{1ig} \mathbf{x}_i \mathbf{x}_i' \boldsymbol{\Gamma}_g' \right] \quad (26)$$

or

$$\mathbf{F}_{2g} = \sum_{i=1}^n \hat{z}_{ig} \left[-\mathbf{x}_i (\mathbf{V}_{ig} \hat{\boldsymbol{\mu}}_g + \hat{\boldsymbol{\alpha}}_g)' \hat{\boldsymbol{\Phi}}_g^{-1} + \mathbf{x}_i \mathbf{x}_i' \boldsymbol{\Gamma}_g' \hat{\boldsymbol{\Phi}}_g^{-1} \mathbf{V}_{ig} - \alpha_{2ig} \mathbf{V}_{ig} \hat{\boldsymbol{\Phi}}_g^{-1} \boldsymbol{\Gamma}_g' \right], \quad (27)$$

where α_{1ig} is the largest eigenvalue of the diagonal matrix $\boldsymbol{\Phi}_g^{-1} \mathbf{V}_{ig}$, and α_{2ig} is equal to $\hat{z}_{ig} \mathbf{x}_i' \mathbf{x}_i$, which is the largest eigenvalue of the rank-1 matrix $\hat{z}_{ig} \mathbf{x}_i \mathbf{x}_i'$. Following this, we compute the singular value decomposition of \mathbf{F}_{rg} given by

$$\mathbf{F}_{rg} = \mathbf{P} \mathbf{B} \mathbf{R}'.$$

It follows that our update for $\boldsymbol{\Gamma}_g$ is given by

$$\hat{\boldsymbol{\Gamma}}_g = \mathbf{R} \mathbf{P}'.$$

The p -dimensional concentration and index parameters, i.e., ω_g and λ_g , are estimated by maximizing the function

$$q_{jg}(\omega_{jg}, \lambda_{jg}) = -\log K_{\lambda_{jg}}(\omega_{jg}) + (\lambda_{jg} - 1)\bar{E}_{3jg} - \frac{\omega_{jg}}{2}(\bar{E}_{1jg} + \bar{E}_{2jg}). \quad (28)$$

This leads to

$$\hat{\lambda}_{jg} = \bar{E}_{3jg} \lambda_{jg}^{\text{prev}} \left[\frac{\partial}{\partial v} \log K_v(\omega_{jg}^{\text{prev}}) \Big|_{v=\lambda_{jg}^{\text{prev}}} \right]^{-1}$$

and

$$\hat{\omega}_{jg} = \omega_{jg}^{\text{prev}} - \left[\frac{\partial}{\partial v} q_{jg}(v, \hat{\lambda}_{jg}) \Big|_{v=\omega_{jg}^{\text{prev}}} \right] \left[\frac{\partial^2}{\partial v^2} q_{jg}(v, \hat{\lambda}_{jg}) \Big|_{v=\omega_{jg}^{\text{prev}}} \right]^{-1},$$

where the superscript “prev” denotes that the estimate from the previous iteration is used. The univariate parameters ω_{0g} and λ_{0g} are estimated by maximizing the function

$$q_{0g}(\omega_{0g}, \lambda_{0g}) = -\log(K_{\lambda_{0g}}(\omega_{0g})) + (\lambda_{0g} - 1)C_g - \frac{\omega_{0g}}{2}(A_g + B_g), \quad (29)$$

giving

$$\hat{\lambda}_{0g} = C_g \lambda_{0g}^{\text{prev}} \left[\frac{\partial}{\partial v} \log K_v(\omega_{0g}^{\text{prev}}) \Big|_{v=\lambda_{0g}^{\text{prev}}} \right]^{-1}$$

and

$$\hat{\omega}_{0g} = \omega_{0g}^{\text{prev}} - \left[\frac{\partial}{\partial v} q_{0g}(v, \hat{\lambda}_{0g}) \Big|_{v=\omega_{0g}^{\text{prev}}} \right] \left[\frac{\partial^2}{\partial v^2} q_{0g}(v, \hat{\lambda}_{0g}) \Big|_{v=\omega_{0g}^{\text{prev}}} \right]^{-1}.$$

Our GEM algorithm is iterated until convergence, which is determined using the Aitken acceleration (Aitken, 1926). Formally, the Aitken acceleration is given by

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

where $l^{(k)}$ is the value of the log-likelihood at the iteration k and

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

is an asymptotic estimate of the log-likelihood on iteration $k + 1$. The algorithm can be considered to have converged when $l_{\infty}^{(k)} - l^{(k)} < \epsilon$, provided this difference is positive (Böhning et al., 1994; Lindsay, 1995; McNicholas et al., 2010). Herein, we set $\epsilon = 0.01$. When the algorithm converges we compute the maximum *a posteriori* (MAP) classification values using the posterior \hat{z}_{ig} , where $\text{MAP} \{\hat{z}_{ig}\} = 1$ if $g = \arg \max_h \{\hat{z}_{ih}\}$, and $\text{MAP} \{\hat{z}_{ig}\} = 0$ otherwise.

B Quasi-Concavity of the cMSGHD

In essence, we might want to consider only densities whose contours contain a set of points that are convex. Formally, such densities are quasi-concave. Extensive details on quasi-concavity, quasi-convexity, and related notions are given by Niculescu and Persson (2006) and Rockafellar and Wets (2009).

Definition 1. A function $f(\mathbf{x})$ is quasi-concave if each upper-level set $U_\alpha(f) = \{\mathbf{x} \mid f(\mathbf{x}) \geq \alpha\}$ is convex, for $\alpha \in \mathbb{R}$.

Definition 2. A function $f(\mathbf{x})$ is quasi-convex if each sub-level set $S_\alpha(f) = \{\mathbf{x} \mid f(\mathbf{x}) \leq \alpha\}$ is convex, for $\alpha \in \mathbb{R}$.

Lemma 1. The class of elliptical distributions, whose density functions have the form

$$f(\mathbf{x}) = \frac{1}{\sqrt{|\Sigma|}} g(\delta(\mathbf{x}, \boldsymbol{\mu} \mid \Sigma))$$

are quasi-concave if the generator function, g , is monotonic non-increasing.

Proof. Result follows from the fact that if the function $\delta(\mathbf{x}, \boldsymbol{\mu} \mid \Sigma)$ is convex since Σ is positive definite and the function g is monotonic non-increasing, then the function $f(\mathbf{x})$ is quasi-concave. \square

Theorem 1. The generalized hyperbolic distribution (GHD) is quasi-concave.

Proof. It is straightforward to show that the function

$$h(\mathbf{x}) = \sqrt{a + b \delta(\mathbf{x}, \boldsymbol{\mu} \mid \Sigma)}$$

is convex, where a and b are positive constants, and $\delta(\mathbf{x}, \boldsymbol{\mu} \mid \Sigma)$ is the Mahalanobis distance between \mathbf{x} and $\boldsymbol{\mu}$. Let $\tau = \lambda - p/2$. Then, the function

$$k(z) = \tau \log z + \log K_\tau(z),$$

where $z \in \mathbb{R}^+$, and K_τ is the modified Bessel function of the third kind with index τ , is monotonic decreasing (or non-increasing) because the first derivative

$$k'(z) = \frac{\tau}{z} + \frac{(\tau/z)K_\tau(z) - K_{\tau+1}(z)}{K_\tau(z)} = \frac{2\tau}{z} - \frac{K_{\tau+1}(z)}{K_\tau(z)} = -\frac{K_{\tau-1}(z)}{K_\tau(z)}$$

is negative for all $\tau \in \mathbb{R}$ and $z > 0$. In addition to being monotonic decreasing, $k(z)$ is convex for $\tau < 1/2$, concave and convex (linear) for $\tau = 1/2$, and concave for $\tau > 1/2$. Because $k(z)$ is a monotonic function, it satisfies the criteria for quasi-convexity and quasi-concavity, so it is simultaneously quasi-convex and quasi-concave. In this context, monotone functions are also known as quasi-linear or quasi-montone.

Recall that if the function U is quasi-convex and the function g is decreasing, then the function $f(x) = g(U(x))$ is quasi-concave. It follows that the composition $k(h(\mathbf{x}))$ is quasi-concave. Consider

the skewness part of the GHD density function, i.e. $a(\mathbf{x}) = -(\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} \boldsymbol{\alpha}$, which is a linear function. It follows that the function

$$\exp \{k(h(\mathbf{x})) + a(\mathbf{x})\} \quad (30)$$

is also quasi-concave, and the result follows from the fact that (30) is proportional to the density of the GHD. \square \square

Theorem 2. *The convex multiple scaled generalized hyperbolic distribution (cMSGHD) is quasi-concave. In other words, the multiple scaled generalized hyperbolic distribution (MSGHD) is quasi-concave provided that $\lambda_j > 1$ for all $j = 1, \dots, p$.*

Proof. A p -dimensional multiple scaled distribution is a product of p independent univariate densities. The density of the MSGHD has form

$$g_p(x_1, x_2, \dots, x_p) = g_1(x_1 \mid \boldsymbol{\theta}_1) g_1(x_2 \mid \boldsymbol{\theta}_2) \times \dots \times g_1(x_p \mid \boldsymbol{\theta}_p),$$

where $g_1(x_j \mid \boldsymbol{\theta}_j)$ is the density of the univariate hyperbolic distribution with parameters $\boldsymbol{\theta}_j$, $j = 1, \dots, p$. From Theorem 1, $\log g_1(x_j \mid \boldsymbol{\theta}_j)$ is a concave function for $\tau_j > 1/2$, i.e., for $\lambda_j > 1$ (because $p = 1$). Therefore, the function

$$\log g_p(x_1, x_2, \dots, x_p) = \log g_1(x_1 \mid \boldsymbol{\theta}_1) + \log g_1(x_2 \mid \boldsymbol{\theta}_2) + \dots + \log g_1(x_p \mid \boldsymbol{\theta}_p)$$

is concave provided that $\lambda_j > 1$ for all $j = 1, \dots, p$. Therefore, the function

$$g_p(x_1, x_2, \dots, x_p) = g_1(x_1 \mid \boldsymbol{\theta}_1) g_1(x_2 \mid \boldsymbol{\theta}_2) \times \dots \times g_1(x_p \mid \boldsymbol{\theta}_p)$$

is quasi-concave provided that $\lambda_j > 1$ for all $j = 1, \dots, p$. \square \square

Note that addition does not preserve quasi-convexity or quasi-concavity. The sum of two quasi-convex functions defined on different domains will be quasi-concave if they are additively decomposed (see Debreu and Koopmans, 1982). Debreu and Koopmans (1982) give necessary and sufficient conditions for the sum f of a set of functions f_1, \dots, f_m to be additively decomposed. These conditions depend on the convexity index $c(f)$ in which f is quasi convex if and only if either of the following hold: (i) $c(f_i) \geq 0$ for every i , or (ii) $c(f_j) < 0$ for some j , $c(f_i) > 0$ for every $i \neq j$, and $\sum_{i=1}^m \frac{1}{c(f_i)} \leq 0$. For differentiable functions, the convexity index satisfies the inequality $f''(x)/[f'(x)]^2 \geq c(f)$.

We have that a sufficient condition for the MSGHD to be quasi-concave is that all $\lambda_j > 1$. Furthermore, a sufficient condition for the MSGHD not to be quasi-concave is that all $\lambda_j < 1$ and finite. Interestingly, this means the multiple scaled t-distribution cannot provide convex level sets for any finite degrees of freedom. For large degrees of freedom, the multiple scaled t-distribution will behave similarly to a normal distribution near the mode; however, as one moves away from the mode, non-convex contours will be encountered. Finally, note that Debreu and Koopmans (1982) give necessary and sufficient conditions that suggest a quasi-concave MSGHD with some λ_j positive and others negative is possible, but going this route would greatly complicate the estimation procedure.

C Finite Mixture Identifiability

In this section we consider the notion of identifiability for finite mixtures of MSGHDs and coalesced generalized hyperbolic distributions (CGHDs). Herein, we take the term identifiability to mean finite mixture identifiability.

C.1 Background

Holzmann et al. (2006) prove identifiability of finite mixtures of elliptical distributions. They state that “finite mixtures are said to be identifiable if distinct mixing distributions with finite support correspond to distinct mixtures”. A finite mixture of the densities $f_p(\mathbf{x}|\Psi_1), \dots, f_p(\mathbf{x}|\Psi_G)$ is identifiable if the family $\{f_p(\mathbf{x}|\Psi) : \Psi \in \mathcal{A}^p\}$ is linearly independent. The founding work on finite mixture identifiability is by Yakowitz and Spragins (1968), who state that this linear independence is a necessary and sufficient condition for identifiability.

The GHD can be expressed as a normal variance-mean mixture. The stochastic relationship of the normal variance-mean mixture is given by

$$\mathbf{X} = \boldsymbol{\mu} + W\boldsymbol{\alpha} + \sqrt{W}\mathbf{U}, \quad (31)$$

where $\mathbf{U} \sim \mathcal{N}_p(\mathbf{0}, \boldsymbol{\Sigma})$ and W , independent of \mathbf{U} , is a positive univariate random variable with density $h(w|\boldsymbol{\theta})$. Browne and McNicholas (2015) proved identifiability for finite mixtures of GHDs through additivity of disjoint sets of identifiable distributions.

Definition 3. *In the present context, a finite mixture of the multiple scale distributions $f(\mathbf{x}|\boldsymbol{\theta}_1), \dots, f(\mathbf{x}|\boldsymbol{\theta}_G)$ is identifiable if*

$$\sum_{g=1}^G \pi_g f(\mathbf{x}|\boldsymbol{\theta}_g) = \sum_{g=1}^G \pi_g^* f(\mathbf{x}|\boldsymbol{\theta}_g^*) \quad (32)$$

for $\mathbf{x} \in \mathbb{R}^p$, where G is a positive integer, $\sum_{g=1}^G \pi_g = \sum_{g=1}^G \pi_g^* = 1$ and $\pi_g, \pi_g^* > 0$ for $g = 1, \dots, G$, implies that there exists a permutation σ such that $(\pi_g, \boldsymbol{\theta}_g) = (\pi_{\sigma(g)}, \boldsymbol{\theta}_{\sigma(g)})$ for all g .

Browne and McNicholas (2015) prove identifiability for normal variance-mean mixtures, which includes the generalized hyperbolic. Here we view the results from a different vantage point to illustrate the concepts required for the identifiability of the multiple scaled distributions. We begin by noting the characteristic function for the generalized hyperbolic arises from the characteristic function of the normal variance-mean mixture,

$$\varphi_{\mathbf{X}}(\mathbf{v}) = \exp\{i\mathbf{v}'\boldsymbol{\mu}_g\} M_W \left(\boldsymbol{\beta}'_g \mathbf{v} i - \frac{1}{2} \mathbf{v}' \boldsymbol{\Sigma}_g \mathbf{v} \mid \boldsymbol{\Gamma}_g \right), \quad (33)$$

where

$$M_W(u) = \left[\frac{\omega}{\omega - 2u} \right]^{\frac{\lambda}{2}} \frac{K_{\lambda} \left(\sqrt{\omega(\omega - 2u)} \right)}{K_{\lambda}(\omega)} = \left[1 - 2\frac{u}{\omega} \right]^{-\frac{\lambda}{2}} \frac{K_{\lambda} \left(\sqrt{\omega(\omega - 2u)} \right)}{K_{\lambda}(\omega)}.$$

The characteristic function for the generalized hyperbolic is

$$\varphi_{\mathbf{X}}(\mathbf{v}) = \exp\{i\mathbf{v}'\boldsymbol{\mu}\} \left[1 + \frac{\mathbf{v}'\boldsymbol{\Sigma}\mathbf{v} - 2i\boldsymbol{\beta}'\mathbf{v}}{\omega} \right]^{-\frac{\lambda}{2}} \frac{K_{\lambda} \left(\sqrt{\omega [\omega + (\mathbf{v}'\boldsymbol{\Sigma}\mathbf{v} - 2i\boldsymbol{\beta}'\mathbf{v})]} \right)}{K_{\lambda}(\omega)}.$$

In context of a coalesced distribution, with a eigen-decomposed scale matrix, the characteristic function is

$$\varphi_{\mathbf{X}}(\mathbf{v}) = \exp\{i\mathbf{v}'\boldsymbol{\mu}\} \left[1 + \frac{\mathbf{v}'\boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'\mathbf{v} - 2i\boldsymbol{\beta}'\mathbf{v}}{\omega} \right]^{-\frac{\lambda}{2}} \frac{K_{\lambda} \left(\sqrt{\omega [\omega + (\mathbf{v}'\boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'\mathbf{v} - 2i\boldsymbol{\beta}'\mathbf{v})]} \right)}{K_{\lambda}(\omega)}.$$

Now, we let $\mathbf{v} = t\mathbf{z}$ and obtain

$$\varphi_{\mathbf{X}}(\mathbf{v} = t\mathbf{z}) = \exp\{it\mathbf{z}'\boldsymbol{\mu}\} \left[1 + \frac{t^2(\mathbf{z}'\boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'\mathbf{z}) - 2it(\boldsymbol{\beta}'\mathbf{z})}{\omega} \right]^{-\frac{\lambda}{2}} \frac{K_{\lambda} \left(\sqrt{\omega [\omega + t^2(\mathbf{z}'\boldsymbol{\Gamma}\boldsymbol{\Phi}\boldsymbol{\Gamma}'\mathbf{z}) - 2it(\boldsymbol{\beta}'\mathbf{z})]} \right)}{K_{\lambda}(\omega)}.$$

To prove identifiability of the generalized hyperbolic we could now use the results from Browne and McNicholas (2015) and Yakowitz and Spragins (1968, p. 211) that implies there exists \mathbf{z} such that the tuple $(\mathbf{z}'\boldsymbol{\Sigma}_g\mathbf{z}, \boldsymbol{\beta}'_g\mathbf{z}, \mathbf{z}'\boldsymbol{\mu}_g)$, where $\boldsymbol{\Sigma}_g = \boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}'_g$ is unique for all $g = 1, \dots, G$, allows a reduction to the univariate case. Now, we rewrite the term $\mathbf{z}'\boldsymbol{\Sigma}_g\mathbf{z}$ as

$$\mathbf{z}'\boldsymbol{\Sigma}_g\mathbf{z} = \mathbf{z}'\boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}'_g\mathbf{z} = \text{tr} [\mathbf{z}'\boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}'_g\mathbf{z}] = \text{tr} [\boldsymbol{\Gamma}'_g\mathbf{z}\mathbf{z}'\boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g] = \sum_{j=1}^p \Phi_{jg}[\boldsymbol{\Gamma}'_g\mathbf{z}]_j^2,$$

which implies the tuple

$$(\mathbf{z}'\boldsymbol{\Gamma}_g\boldsymbol{\Phi}_g\boldsymbol{\Gamma}'_g\mathbf{z}, \boldsymbol{\beta}'_g\mathbf{z}, \mathbf{z}'\boldsymbol{\mu}_g) \equiv \left(\sum_{j=1}^p \Phi_{jg}[\boldsymbol{\Gamma}'_g\mathbf{z}]_j^2, \boldsymbol{\beta}'_g\mathbf{z}, \mathbf{z}'\boldsymbol{\mu}_g \right)$$

is unique for all $g = 1, \dots, G$. A similar argument indicates there exists a \mathbf{z} such that the tuple

$$\left(\sum_{j=1}^p \Phi_{jg} |[\boldsymbol{\Gamma}'_g\mathbf{z}]_j|, \boldsymbol{\beta}'_g\mathbf{z}, \mathbf{z}'\boldsymbol{\mu}_g \right) \quad (34)$$

is unique. In fact, a more general statement indicates that there exists a \mathbf{z} such that the tuple

$$\left(\sum_{j=1}^p \Phi_{jg} \varphi([\boldsymbol{\Gamma}'_g\mathbf{z}]_j^2), \boldsymbol{\beta}'_g\mathbf{z}, \mathbf{z}'\boldsymbol{\mu}_g \right)$$

is unique for monotonic $\varphi : \mathbb{R}^+ \mapsto \mathbb{R}^+$. Deriving this unique set of tuples facilitates the reduction to the univariate case. This is useful because the univariate generalized hyperbolic density is identifiable (see Browne and McNicholas, 2015).

C.2 Identifiability of a Finite Mixture of Multiple Scaled Distributions

For a multiple scaled distribution, we only need to find a single direction where the distribution is finite mixture identifiable because, as noted in Remark 2 of Kent (1983), a distribution might be non-identifiable on a subset of \mathbb{R}^p but identifiability can endure over \mathbb{R}^p . In other words, for a distribution to be non-identifiable, a linear combination has to be equal to zero for all $x \in \mathbb{R}^p$. This is illustrated by the example given in Kent (1983):

“the polynomials $P(x_1, x_2) = 1$ and $P(x_1, x_2) = (x_1^2 + x_2)^3$, $x \in \mathbb{R}^2$, are equal on the unit circle, but are not the same on all of \mathbb{R}^2 .”

As a consequence, if a multivariate distribution is identifiable in some direction then it is identifiable over \mathbb{R}^p .

To begin, consider that if there is at one least direction or column of $\mathbf{\Gamma}_g$ that is equal across $g = 1, \dots, G$, then the identifiability of a multiple scaled distribution follows from the identifiability of the univariate distribution. Whereas if one column of $\mathbf{\Gamma}_g$ is unequal, that implies, by the nature of orthonormal matrices, that two columns of $\mathbf{\Gamma}_g$ are unequal. We will now illustrate how the bivariate multiple scaled distribution is identifiable, which implies identifiability for finite p .

When $\mathbf{\Gamma}_g$ differ, the identifiability of the multiple scaled distribution depends on the behaviour of the multiple scaled distribution's density and moment generating functions when we consider moving along directions other than the columns of $\mathbf{\Gamma}_g$. For example, a bivariate multiple scaled t -distribution behaves (by definition) like a t -distribution with ν_1 and ν_2 degrees of freedom along each of its principal axes, but along any other direction, a bivariate multiple scaled t -distribution behaves asymptotically like a t -distribution with $\nu_1 + \nu_2$ degrees of freedom.

Consider the following three orthonormal matrices in the context of an eigen-decomposition of a matrix;

$$\mathbf{\Gamma}_1 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{\Gamma}_2 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix} \quad \text{and} \quad \mathbf{\Gamma}_3 = \begin{bmatrix} -1 & 0 \\ 0 & 1 \end{bmatrix}.$$

If we have equal eigenvalues then we cannot distinguish between $\mathbf{\Gamma}_1$ and $\mathbf{\Gamma}_2$. In the same way, if we have the same distribution along the first and second axis, we cannot distinguish between them. However, if we have eigenvalue ordering we can distinguish between $\mathbf{\Gamma}_1$ and $\mathbf{\Gamma}_2$, but eigenvalue ordering will not allow us to distinguish between $\mathbf{\Gamma}_1$ and $\mathbf{\Gamma}_3$, since they yield the same basis or set of directions. Therefore, in general, $\mathbf{\Gamma}$ is unique up to multiplication by

$$\begin{bmatrix} \pm 1 & 0 \\ 0 & \pm 1 \end{bmatrix}.$$

One way to establish uniqueness is to require the largest value of each column of $\mathbf{\Gamma}$ to be positive. An equivalent requirement is for $\mathbf{\Gamma}_1 \neq \mathbf{\Gamma}_2$ which requires that

$$\mathbf{\Gamma}'_1 \mathbf{\Gamma}_2 \neq \mathbf{R} \quad \text{or} \quad [\mathbf{\Gamma}'_1 \mathbf{z}]_j \neq -[\mathbf{\Gamma}'_2 \mathbf{z}]_j \quad (35)$$

for $j = 1, \dots, p$, $\mathbf{z} \in \mathbb{R}^P$, $\mathbf{z} \neq \mathbf{0}_p$ and \mathbf{R} is a set of diagonal matrices such that $\text{diag}(\mathbf{R}) = (\pm 1, \dots, \pm 1)$ excluding the identity matrix. Note that $[\mathbf{a}]_j$ denotes the j th element of the vector \mathbf{a} . However, if we had two orthonormal matrices such that $\mathbf{\Gamma}'_1 \mathbf{\Gamma}_2 = \mathbf{I}$, then $\mathbf{\Gamma}_1 = \mathbf{\Gamma}_2$. If $\mathbf{\Gamma}'_1 \mathbf{\Gamma}_2 = \mathbf{R}$, then our orthonormal condition amounts to $\mathbf{\Gamma}'_1 = \mathbf{\Gamma}_2$ or equivalently, for all directions $\mathbf{z} \in \mathbb{R}^P$ and $\mathbf{z} \neq \mathbf{0}_p$

$$|[\mathbf{\Gamma}'_1 \mathbf{z}]_j| = |[\mathbf{\Gamma}'_2 \mathbf{z}]_j| \quad \text{for all } j = 1, \dots, p \quad \text{then} \quad \mathbf{\Gamma}_1 = \mathbf{\Gamma}_2. \quad (36)$$

This prevents the j th column of $\mathbf{\Gamma}_2$ from being in the opposite direction of the j th column of $\mathbf{\Gamma}_1$. This form of the condition is easier to incorporate into the identifiability illustration.

In the MSGHD, if we consider moving the amount t in a direction \mathbf{z} , which entails setting $\mathbf{x} = t\mathbf{z}$, we can write the density as

$$f_{\text{MSGHD}}(\mathbf{x} = t\mathbf{z} \mid \boldsymbol{\mu}, \boldsymbol{\Gamma}, \boldsymbol{\Phi}, \boldsymbol{\alpha}, \boldsymbol{\omega}, \boldsymbol{\lambda}) \\ = \prod_{j=1}^p \left[\frac{\omega_j + \Phi_j^{-1} (t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j)^2}{\omega_j + \alpha_j^2 \Phi_j^{-1}} \right]^{\frac{\lambda_j - \frac{1}{2}}{2}} \frac{K_{\lambda_j - \frac{1}{2}} \left(\sqrt{[\omega_j + \alpha_j^2 \Phi_j^{-1}] [\omega_j + \Phi_j^{-1} (t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j)^2]} \right)}{(2\pi)^{\frac{1}{2}} \Phi_j^{\frac{1}{2}} K_{\lambda_j}(\omega_j) \exp \{-(t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j) \alpha_j \Phi_j^{-1}\}}, \quad (37)$$

Note, if \mathbf{z} is equal to the k th eigenvector, which is the k th column of $\boldsymbol{\Gamma}$, then the density reduces to

$$c_k \left[\frac{\omega_k + \Phi_k^{-1} (t - \mu_k)^2}{\omega_k + \alpha_k^2 \Phi_k^{-1}} \right]^{\frac{\lambda_k - \frac{1}{2}}{2}} \frac{K_{\lambda_k - \frac{1}{2}} \left(\sqrt{[\omega_k + \alpha_k^2 \Phi_k^{-1}] [\omega_k + \Phi_k^{-1} (t - \mu_k)^2]} \right)}{(2\pi)^{\frac{1}{2}} \Phi_k^{\frac{1}{2}} K_{\lambda_k}(\omega_k) \exp \{-(t - \mu_k) \alpha_k \Phi_k^{-1}\}},$$

where

$$c_k = \prod_{j=1, j \neq k}^p \left[\frac{\omega_j + \Phi_j^{-1} \mu_j^2}{\omega_j + \alpha_j^2 \Phi_j^{-1}} \right]^{\frac{\lambda_j - \frac{1}{2}}{2}} \frac{K_{\lambda_j - \frac{1}{2}} \left(\sqrt{[\omega_j + \alpha_j^2 \Phi_j^{-1}] [\omega_j + \Phi_j^{-1} \mu_j^2]} \right)}{(2\pi)^{\frac{1}{2}} \Phi_j^{\frac{1}{2}} K_{\lambda_j}(\omega_j) \exp \{\mu_j \alpha_j \Phi_j^{-1}\}}.$$

Therefore, the density is simply proportional to

$$\propto \left[\frac{\omega_k + \Phi_k^{-1} (t - \mu_k)^2}{\omega_k + \alpha_k^2 \Phi_k^{-1}} \right]^{\frac{\lambda_k - \frac{1}{2}}{2}} \frac{K_{\lambda_k - \frac{1}{2}} \left(\sqrt{[\omega_k + \alpha_k^2 \Phi_k^{-1}] [\omega_k + \Phi_k^{-1} (t - \mu_k)^2]} \right)}{(2\pi)^{\frac{1}{2}} \Phi_k^{\frac{1}{2}} K_{\lambda_k}(\omega_k) \exp \{-(t - \mu_k) \alpha_k \Phi_k^{-1}\}}.$$

First, note that if the parameterizations are one-to-one, then if one parameterization is shown to be identifiable, the others are identifiable as well. Similar to Browne and McNicholas (2015), we let $\delta_j = \beta_j / \Phi_j$, $\alpha_j = \sqrt{\omega_j / \Phi_j + \beta_j^2 / \Phi_j^2}$ and $\kappa_j = \sqrt{\Phi_j \omega_j}$, where $\alpha_j \geq |\delta_j|$. Under this reparameterization, we now have

$$\Phi_j = \frac{\kappa_j}{\sqrt{\alpha_j^2 - \delta_j^2}}, \quad \omega_j = \kappa_j \sqrt{\alpha_j^2 - \delta_j^2} \quad \text{and} \quad \beta_j = \frac{\delta_j \kappa_j}{\sqrt{\alpha_j^2 - \delta_j^2}}. \quad (38)$$

For large z , the Bessel function can be approximated by

$$K_\lambda(z) = \sqrt{\frac{\pi}{2z}} e^{-z} \left[1 + O\left(\frac{1}{z}\right) \right],$$

which yields, using the alternative parameterization,

$$f(t \mid \boldsymbol{\theta}) \propto \left[1 + \frac{(t - \mu_j)^2}{\kappa_j^2} \right]^{\lambda_j/2} \exp \{ -\alpha_j |t - \mu_j| + \delta_j (t - \mu_j) \}. \quad (39)$$

If \mathbf{z} is not equal to the k th eigenvector, than, using the reparameterization given in (38), we have

$$\begin{aligned}
f(t \mid \boldsymbol{\theta}) &\propto \exp \left\{ -\sum_{j=1}^p \alpha_j \left| t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right| + \sum_{j=1}^p \delta_j \left(t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right) \right\} \prod_{j=1}^p \left[1 + \frac{\left(t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right)^2}{\kappa_j^2} \right]^{\frac{\lambda_j - \frac{1}{2}}{2}} \\
&\propto \exp \left\{ -\sum_{j=1}^p \alpha_j \left| t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right| + \sum_{j=1}^p \delta_j \left(t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right) \right\} t^{2 \sum_{j=1}^p I([\boldsymbol{\Gamma}'\mathbf{z}]_j \neq 0) \frac{\lambda_j - \frac{1}{2}}{2}} \\
&\propto \exp \left\{ \sum_{j=1}^p \left[-\alpha_j \left| t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right| + \delta_j \left(t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right) + 2I([\boldsymbol{\Gamma}'\mathbf{z}]_j \neq 0) \frac{\lambda_j - \frac{1}{2}}{2} \log(t) \right] \right\} \\
&\propto \prod_{j=1}^p \exp \left\{ -\alpha_j \left| t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right| + \delta_j \left(t [\boldsymbol{\Gamma}'\mathbf{z}]_j - \mu_j \right) + 2I([\boldsymbol{\Gamma}'\mathbf{z}]_j \neq 0) \frac{\lambda_j - \frac{1}{2}}{2} \log(t) \right\}.
\end{aligned} \tag{40}$$

The characteristic function for a multiple scaled distribution can be written as

$$\begin{aligned}
\varphi_{\mathbf{X}}(\mathbf{v}) &= \prod_{j=1}^P \exp\{i|[\boldsymbol{\Gamma}'\mathbf{v}]_j|\mu_j\} \left[1 + \frac{\Phi_j|[\boldsymbol{\Gamma}'\mathbf{v}]_j|^2 - 2\beta_j|[\boldsymbol{\Gamma}'\mathbf{v}]_j|i}{\omega_j} \right]^{-\frac{\lambda_j}{2}} \\
&\quad \times \frac{K_{\lambda_j} \left(\sqrt{\omega_j [\omega_j + (\Phi_j|[\boldsymbol{\Gamma}'\mathbf{v}]_j|^2 - 2\beta_j|[\boldsymbol{\Gamma}'\mathbf{v}]_j|i)]} \right)}{K_{\lambda_j}(\omega_j)},
\end{aligned}$$

which, under the alternative parameterization from equation (38), becomes

$$\begin{aligned}
\varphi_{\mathbf{X}}(\mathbf{v}) &= \prod_{j=1}^P \exp\{i|[\boldsymbol{\Gamma}'\mathbf{v}]_j|\mu_j\} \left[1 + \frac{|[\boldsymbol{\Gamma}'\mathbf{v}]_j|^2 - 2\delta_j|[\boldsymbol{\Gamma}'\mathbf{v}]_j|i}{\alpha_j^2 - \delta_j^2} \right]^{-\frac{\lambda_j}{2}} \\
&\quad \times \frac{K_{\lambda_j} \left(\sqrt{\kappa_j^2 \left[|[\boldsymbol{\Gamma}'\mathbf{v}]_j|^2 - 2\delta_j|[\boldsymbol{\Gamma}'\mathbf{v}]_j|i + \alpha_j^2 - \delta_j^2 \right]} \right)}{K_{\lambda_j} \left(\kappa_j \sqrt{\alpha_j^2 - \delta_j^2} \right)}.
\end{aligned} \tag{41}$$

Now if we consider moving t in the direction \mathbf{z}

$$\begin{aligned}
\varphi_{\mathbf{X}}(\mathbf{v} = t\mathbf{z}) &= \prod_{j=1}^P \exp\{it|[\boldsymbol{\Gamma}'\mathbf{z}]_j|\mu_j\} \left[1 + \frac{t^2|[\boldsymbol{\Gamma}'\mathbf{z}]_j|^2 - 2\delta_j t|[\boldsymbol{\Gamma}'\mathbf{z}]_j|i}{\alpha_j^2 - \delta_j^2} \right]^{-\frac{\lambda_j}{2}} \\
&\quad \times \frac{K_{\lambda_j} \left(\sqrt{\kappa_j^2 \left[t^2|[\boldsymbol{\Gamma}'\mathbf{z}]_j|^2 - 2\delta_j t|[\boldsymbol{\Gamma}'\mathbf{z}]_j|i + \alpha_j^2 - \delta_j^2 \right]} \right)}{K_{\lambda_j} \left(\kappa_j \sqrt{\alpha_j^2 - \delta_j^2} \right)},
\end{aligned}$$

and, for large t , the characteristic function is

$$\begin{aligned} \varphi_{\mathbf{X}}(\mathbf{v} = t\mathbf{z}) &\propto \exp \left\{ it \sum_{j=1}^P |[\mathbf{\Gamma}'\mathbf{z}]_j| \mu_j - t \sum_{j=1}^P \kappa_j |[\mathbf{\Gamma}'\mathbf{z}]_j| - \log(t) \sum_{j=1}^P \lambda_j I(|[\mathbf{\Gamma}'\mathbf{z}]_j| \neq 0) + O(1) \right\} \\ &\propto \exp \left\{ it \mathbf{z}'\mathbf{\Gamma}\boldsymbol{\mu} - t \sum_{j=1}^P \kappa_j |[\mathbf{\Gamma}'\mathbf{z}]_j| - \log(t) \sum_{j=1}^P \lambda_j I(|[\mathbf{\Gamma}'\mathbf{z}]_j| \neq 0) + O(1) \right\}. \end{aligned}$$

Therefore, from the condition given in (34), there exists \mathbf{z} such that the tuple $(\sum_{j=1}^P \kappa_j |[\mathbf{\Gamma}'\mathbf{z}]_j|, \mathbf{z}'\mathbf{\Gamma}\boldsymbol{\mu})$ is unique for all $g = 1, \dots, G$ and reduces to the univariate hyperbolic distribution, which is identifiable.

C.3 Identifiability of the Coalesced Generalized Hyperbolic Distribution

To prove the identifiability of the CGHD we only need to show that two sets of distributions, the multiple scaled and the generalized hyperbolic distribution are disjoint. Consider moving along the k th eigenvalue such that (λ_k, κ_k) is distinct from (λ_0, κ_0) and the proof easily follows from the identifiability of the univariate generalized hyperbolic distribution.

D Figures

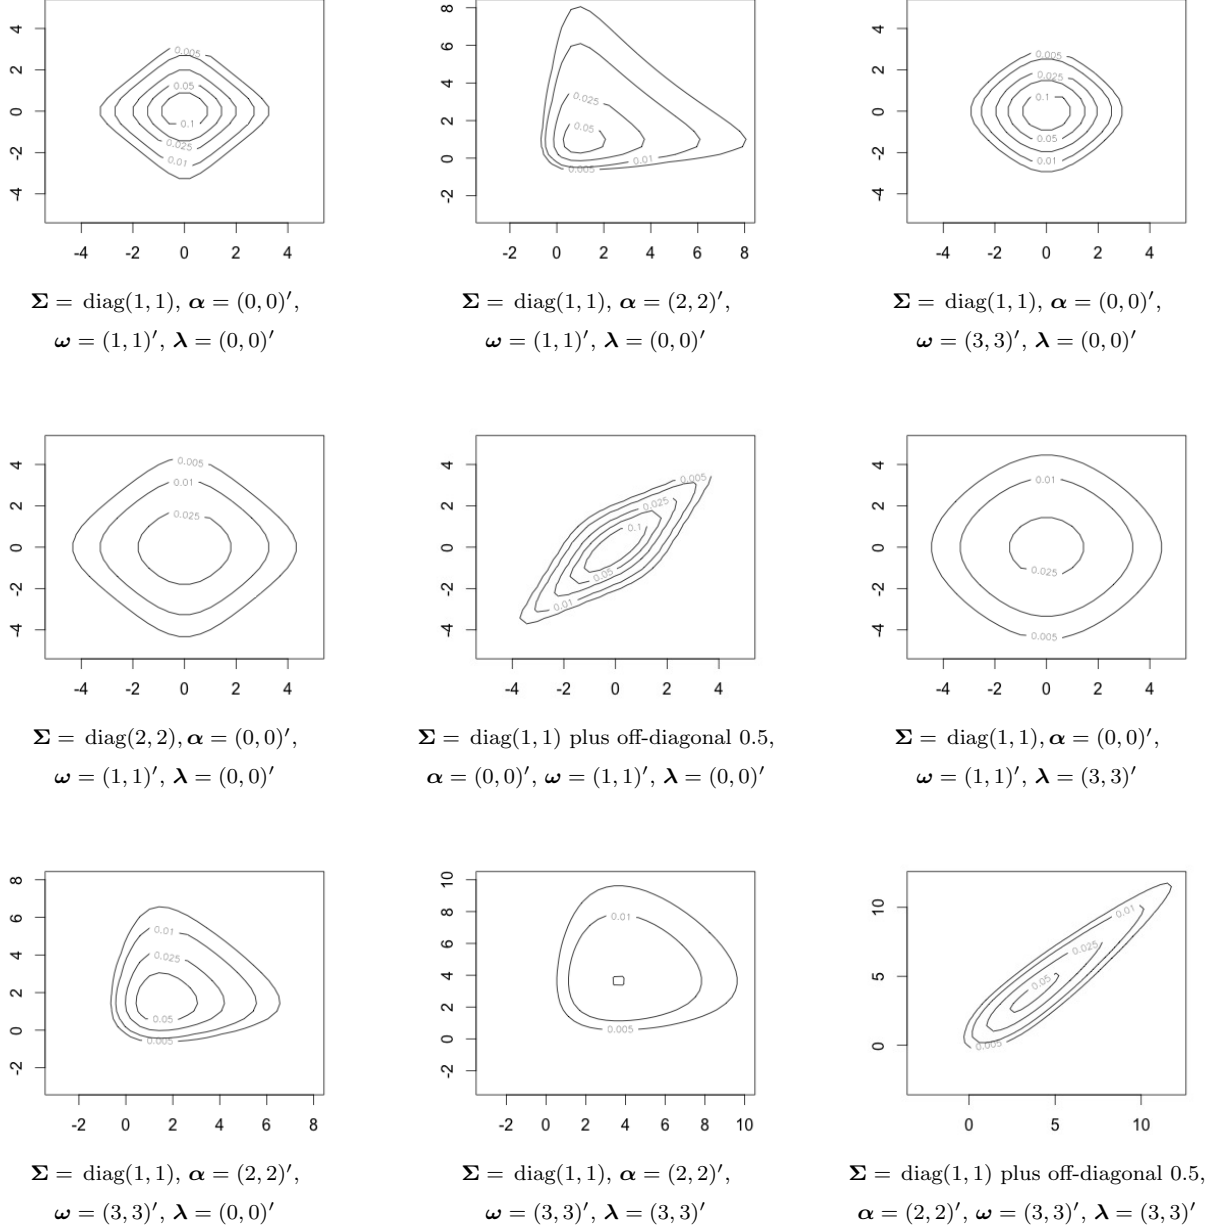
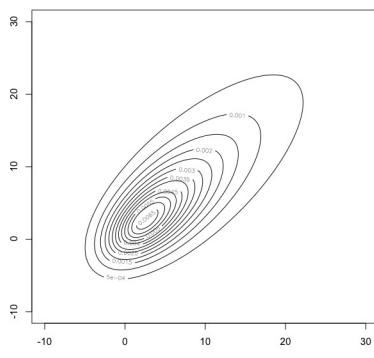
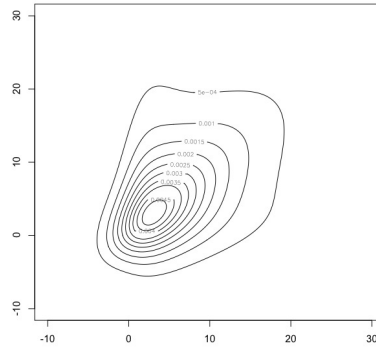


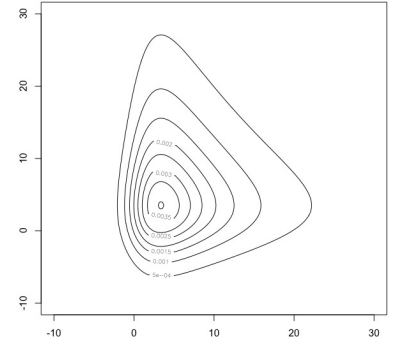
Figure 5: Bivariate contour plots of the MSGHD density with $\mu = (0, 0)'$ and varying Σ , α , ω , and λ .



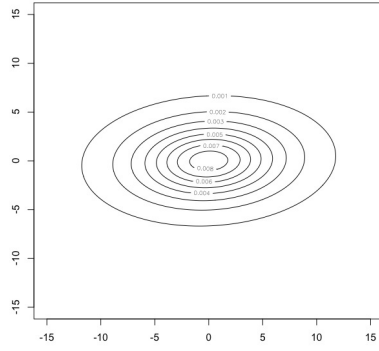
$$\Sigma = \text{diag}(1, 1), \alpha = (2, 2)', \varpi = 1$$



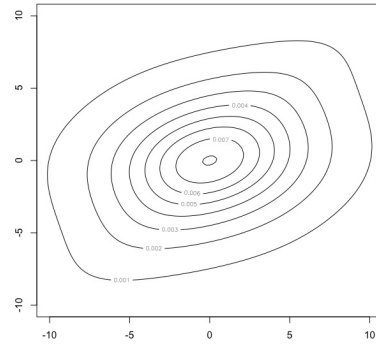
$$\Sigma = \text{diag}(1, 1), \alpha = (2, 2)', \varpi = 0.5$$



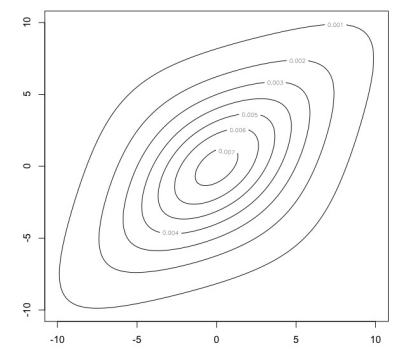
$$\Sigma = \text{diag}(1, 1), \alpha = (2, 2)', \varpi = 0$$



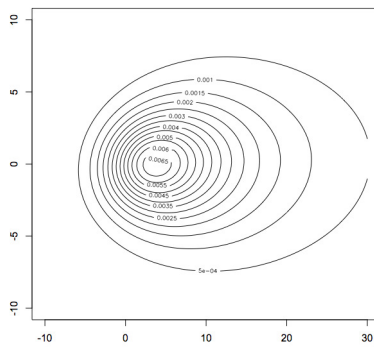
$$\Sigma = \text{diag}(1, 1) \text{ plus off-diagonal } 0.8, \\ \alpha = (0, 0)', \varpi = 1$$



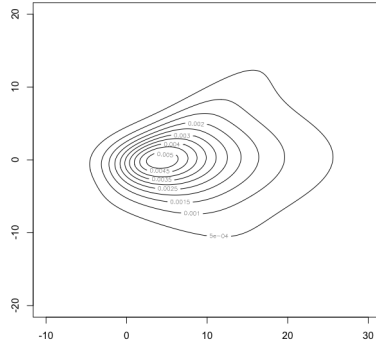
$$\Sigma = \text{diag}(1, 1) \text{ plus off-diagonal } 0.8, \\ \alpha = (0, 0)', \varpi = 0.5$$



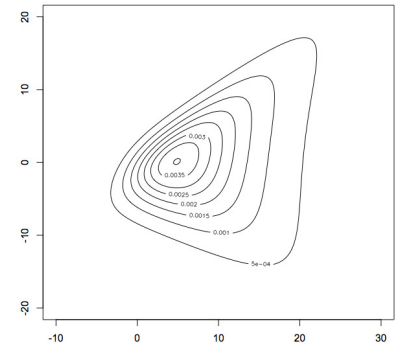
$$\Sigma = \text{diag}(1, 1) \text{ plus off-diagonal } 0.8, \\ \alpha = (0, 0)', \varpi = 0$$



$$\Sigma = \text{diag}(1, 1) \text{ plus off-diagonal } 0.8, \\ \alpha = (2, 2)', \varpi = 1$$



$$\Sigma = \text{diag}(1, 1) \text{ plus off-diagonal } 0.8, \\ \alpha = (2, 2)', \varpi = 0.5$$



$$\Sigma = \text{diag}(1, 1) \text{ plus off-diagonal } 0.8, \\ \alpha = (2, 2)', \varpi = 0$$

Figure 6: Bivariate contour plots of the MCGHD density varying Σ , α , and ϖ .