

## On Fractionally-Supervised Classification: Weight Selection and Extension to the Multivariate $t$ -Distribution

Michael P.B. Gallagher

McMaster University, Canada

Paul D. McNicholas

McMaster University, Canada

**Abstract:** Recent work on fractionally-supervised classification (FSC), an approach that allows classification to be carried out with a fractional amount of weight given to the unlabelled points, is further developed in two respects. The primary development addresses a question of fundamental importance over how to choose the amount of weight given to the unlabelled points. The resolution of this matter is essential because it makes FSC more readily applicable to real problems. Interestingly, the resolution of the weight selection problem opens up the possibility of a different approach to model selection in model-based clustering and classification. A secondary development demonstrates that the FSC approach can be effective beyond Gaussian mixture models. To this end, an FSC approach is illustrated using mixtures of multivariate  $t$ -distributions.

**Keywords:** Fractionally-supervised classification; Weight selection; Multivariate  $t$ -distribution.

---

The authors are most grateful to two anonymous reviewers and the Editor for their very helpful comments. This work is supported by CGS-M and Vanier CGS-D scholarships from the Natural Sciences and Engineering Research Council of Canada (Gallagher) and the Canada Research Chairs program (McNicholas).

Corresponding Author's Address: P.D. McNicholas, Department of Mathematics and Statistics, McMaster University, Hamilton, Ontario, Canada, email: [mcnicholas@math.mcmaster.ca](mailto:mcnicholas@math.mcmaster.ca).

Published online: 16 November 2018

## 1. Introduction

In a typical classification application, some of the observations are unlabelled and the objective is to predict the labels of the unlabelled points, for details see McNicholas (2016a). In such situations, classification is generally semi-supervised or supervised (also called discriminant analysis). These two species of classification differ in whether any weight is given to the unlabelled points in the prediction of their labels. In semi-supervised classification, the labelled and unlabelled points are given equal weight; however, in supervised classification, the unlabelled points are given zero weight. Furthermore, it is possible to either give all the weight to the unlabelled points or treat all the points as unlabelled. This third, and well known, species of classification is called unsupervised classification or cluster analysis. These three species of classification are well established; yet, in any given scenario, it might be the case that labelled or unlabelled observations are more important when building a classifier.

Vrbik and McNicholas (2015) introduce a general approach, called fractionally-supervised classification (FSC), where classification can be carried out with a fractional amount of weight—anything between none and all—being given to the unlabelled points. This approach allows for an intermediate solution between the three different species of classification. Moreover, although it was conceived in the model-based paradigm with the use of Gaussian mixture models and weighted likelihood, discussed in detail in Section 2, it is more generally applicable and will be illustrated herein for  $t$ -mixtures.

Vrbik and McNicholas (2015) show that FSC oftentimes improves classification performance when compared to the three different species of classification; however, the problem over how to choose the appropriate amount of weight to give the unlabelled points remains unanswered. Vrbik and McNicholas (2015) discussed a few different options to choose the appropriate weight but all of these procedures were deemed undesirable. Vrbik and McNicholas (2015) ultimately decided to use the adjusted Rand index (ARI; Hubert and Arabie, 1985) to choose the weight; however, while this approach was sufficient to illustrate that FSC can be very effective, it is not viable in practice because it assumes knowledge of the labels that are treated as unknown in the analysis. The main contribution of the present work is to determine a weight selection criterion that can be used in real problems, where there are genuinely unlabelled points. The secondary contribution of this paper is the demonstration of FSC for non-Gaussian mixture models, in particular mixtures of multivariate  $t$ -distributions.

The remainder of this paper is laid out as follows. In Section 2, a detailed discussion of mixture models, FSC and weighted likelihood, as well

as a brief discussion of the multivariate  $t$ -distribution, is presented. Then, FSC with the multivariate  $t$ -distribution is laid out (Section 3) and a detailed discussion on weight selection criteria is presented (Section 4). In Section 5, simulations and demonstrations using real data are presented and we conclude with a discussion and suggestions for future work (Section 6).

## 2. Background

### 2.1 Finite Mixture Models and Model-Based Clustering

McNicholas (2016a) traces the relationship between mixture models and clustering back as far as Tiedeman (1955). The first use of finite mixture models for model-based clustering is generally regarded to be by Wolfe (1965) and, in the intervening years, model-based clustering has become a popular approach for clustering (a recent review is given by McNicholas, 2016b). A finite mixture model assumes that an observation  $\mathbf{x}$  comes from a population with  $G$  subgroups. The density function of  $\mathbf{x}$  is given by

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

where  $\pi_g > 0$ , with  $\sum_{g=1}^G \pi_g = 1$ , are called the mixing proportions,  $f_g(\cdot)$  are the component densities, and  $\boldsymbol{\vartheta} = (\pi_1, \pi_2, \dots, \pi_G, \boldsymbol{\theta}_1, \boldsymbol{\theta}_2, \dots, \boldsymbol{\theta}_G)$ .

Because of its mathematical tractability, the Gaussian mixture model has been looked at extensively in the literature. In addition to Wolfe (1965), other examples of earlier work in the area of model-based clustering using Gaussian mixtures include Baum et al. (1970), Scott and Symons (1971) and Orchard and Woodbury (1972). For more details on the history of model-based clustering, see McNicholas (2016a). More recently, there has also been a fair amount of work using non-Gaussian mixtures such as the  $t$ -distribution (e.g., Peel and McLachlan, 2000; Andrews and McNicholas, 2011a,b, 2012; Steane et al., 2012; Lin et al., 2014) and skewed distributions (Lin, 2010; Vrbik and McNicholas, 2012, 2014; Lee and McLachlan, 2013, 2014; Franczak et al., 2014, 2015; Dang et al., 2015; Lin et al., 2016; Murray et al., 2014a,b, 2017a,b). Related to this work, an interesting vein of work has been carried out on cluster-weighted models (CWMs; e.g., Ingrassia et al., 2012; Subedi et al., 2013; Ingrassia et al., 2015; Subedi et al., 2015; Punzo and McNicholas, 2017).

### 2.2 Three Species of Classification

Let the  $N \times D$  matrix  $\mathbb{X} = (\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_N)'$  be a data matrix, where the  $\mathbf{x}_i$  are  $D$ -dimensional vectors and  $N$  is the number of data points. We can then split  $\mathbb{X}$  into two sub-matrices  $\mathbb{X}_1$  and  $\mathbb{X}_2$ , where  $\mathbb{X}_1 = (\mathbf{x}'_{11}, \mathbf{x}'_{12}, \dots,$

$\mathbf{x}'_{1n_1})'$  are data points with known labels, and  $\mathbb{X}_2 = (\mathbf{x}'_{21}, \mathbf{x}'_{22}, \dots, \mathbf{x}'_{2n_2})'$  are observations with unknown labels. Then write  $\mathbb{X} = (\mathbb{X}_1, \mathbb{X}_2)'$ . Also, define  $\mathbb{Z} = (\mathbb{Z}_1, \mathbb{Z}_2)'$  to be a matrix of indicator vectors. Specifically, we define  $\mathbb{Z}_1 = (\mathbf{z}_1^{(1)'}, \mathbf{z}_2^{(1)'}, \dots, \mathbf{z}_{n_1}^{(1)'})'$ , where  $\mathbf{z}_i^{(1)}$  are  $G$ -dimensional vectors with elements 0 or 1. For convenience, we will denote element  $g$  of  $\mathbf{z}_j^{(1)}$  by  $z_{jg}^{(1)}$  where

$$z_{jg}^{(1)} = \begin{cases} 1 & \text{if } \mathbf{x}_{1j} \text{ is in group } g, \\ 0 & \text{otherwise.} \end{cases}$$

We can likewise define  $\mathbb{Z}_2$  in the same manner. Furthermore,  $z_{jg}^{(2)}$  for  $j = 1, 2, \dots, n_2$  are analogous to  $z_{jg}^{(1)}$  for the unlabelled observations. Define  $D_o = \{\mathbb{X}, \mathbb{Z}_1\}$  to be our set of observed data, and  $D_c = \{\mathbb{X}, \mathbb{Z}\}$  to be our complete-data. We can furthermore denote the observed data corresponding to labelled observations by  $D_L = \{\mathbb{X}_1, \mathbb{Z}_1\}$ , and the data corresponding to unlabelled observations by  $D_U = \{\mathbb{X}_2\}$ .

Using the above notation, we can now describe the three species of classification. The first species is discriminant analysis, which makes use of only labelled data to build a classifier. The likelihood function in the case of a discriminant analysis can be written as

$$\mathcal{L}_{\text{DA}}(\boldsymbol{\vartheta} \mid D_L) = \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} \mid \boldsymbol{\theta}_g)]^{z_{jg}^{(1)}}. \quad (2)$$

The second species is cluster analysis, and can take on one of two forms. The first form is the one that we will primarily consider, and makes use of only unlabelled data points and ignores the labelled points. In this case, the likelihood function is given by

$$\mathcal{L}_{\text{clust}}(\boldsymbol{\vartheta} \mid D_U) = \prod_{j=1}^{n_2} \sum_{g=1}^G \pi_g f_g(\mathbf{x}_{2j} \mid \boldsymbol{\theta}_g). \quad (3)$$

The second form of the cluster analysis utilizes both labelled and unlabelled points, but treats the labelled points as unlabelled.

The third species is semi-supervised classification. This makes use of all of the observed data  $D_o$  and treats labelled and unlabelled points equally when building a classifier. The likelihood function for semi-supervised classification is given by the product of  $\mathcal{L}_{\text{DA}}(\boldsymbol{\vartheta} \mid D_L)$  and  $\mathcal{L}_{\text{clust}}(\boldsymbol{\vartheta} \mid D_U)$  to give

$$\mathcal{L}_{\text{semi}}(\boldsymbol{\vartheta} \mid D_o) = \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} \mid \boldsymbol{\theta}_g)]^{z_{jg}^{(1)}} \prod_{j=1}^{n_2} \sum_{g=1}^G \pi_g f_g(\mathbf{x}_{2j} \mid \boldsymbol{\theta}_g). \quad (4)$$

### 2.3 Fractionally-Supervised Classification

Introduced by Vrbik and McNicholas (2015), FSC allows for a solution intermediate to the three species of classification. This is achieved by introducing the weight  $\alpha_1 = \alpha$  to labelled observations, and  $\alpha_2 = 1 - \alpha$  to unlabelled observations, where  $0 \leq \alpha \leq 1$ . Using these weights, the most natural form of the weighted observed likelihood can be written as

$$\begin{aligned} \mathcal{L}_{\text{FSC}}(\boldsymbol{\vartheta} \mid D_o, \alpha) &= [\mathcal{L}_{\text{DA}}(\boldsymbol{\vartheta} \mid D_L)]^\alpha [\mathcal{L}_{\text{clust}}(\boldsymbol{\vartheta} \mid D_U)]^{1-\alpha} \\ &= \left[ \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} \mid \boldsymbol{\theta}_g)]^{z_{jg}^{(1)}} \right]^\alpha \left[ \prod_{j'=1}^{n_2} \sum_{h=1}^H \pi_h f_g(\mathbf{x}_{2j'} \mid \boldsymbol{\theta}_h) \right]^{1-\alpha}, \end{aligned} \tag{5}$$

where  $z_{jg}^{(1)}$  is the  $g$ th element of  $\mathbf{z}_j^{(1)}$ . Although  $H$  does not necessarily have to equal  $G$ , we will make the assumption that  $H = G$ . We can then write the complete-data log-likelihood function as

$$\ell(\boldsymbol{\vartheta} \mid D_c) = \sum_{i=1}^2 \sum_{j=1}^{n_i} \sum_{g=1}^G \alpha_i z_{jg}^{(i)} [\log(\pi_g) + \log(f_g(\mathbf{x}_{ij} \mid \boldsymbol{\theta}_g))]. \tag{6}$$

The expectation-maximization (EM) algorithm (Dempster et al., 1977) can then be used to maximize (6). The EM algorithm is an iterative algorithm that, on each iteration, consists of a conditional expectation (E-) step and the subsequent maximization of this expectation (M-step). We first initialize the parameters, and we denote this by  $\boldsymbol{\vartheta}^{(0)}$ . Iteration  $t + 1$  of the EM algorithm proceeds as follows.

**E-Step:** Calculate  $Q(\boldsymbol{\vartheta} \mid \boldsymbol{\vartheta}^{(t)}) = \mathbb{E}_{\mathbb{Z}_2^{\mathbb{X}}}[\ell(\boldsymbol{\vartheta} \mid D_c) \mid D_o, \boldsymbol{\vartheta}^{(t)}]$  (7a)

**M-Step:** Find  $\arg \max_{\boldsymbol{\vartheta}} Q(\boldsymbol{\vartheta} \mid \boldsymbol{\vartheta}^{(t)})$  (7b)

Check for convergence. If the convergence criterion was not met, set  $t = t + 1$  and return to (7a). (7c)

As was shown in Vrbik and McNicholas (2015), in the case of a Gaussian mixture model, steps (7a) and (7b) simplify to the following.

**E-Step:** Update

$$\hat{z}_{jg}^{(2)} = \frac{\pi_g^{(t)} \phi(\mathbf{x}_{2j} \mid \boldsymbol{\mu}_g^{(t)}, \boldsymbol{\Sigma}_g^{(t)})}{\sum_{g=1}^G \pi_g^{(t)} \phi(\mathbf{x}_{2j} \mid \boldsymbol{\mu}_g^{(t)}, \boldsymbol{\Sigma}_g^{(t)})}.$$

Because the  $z_{jg}^{(1)}$  are known, we set  $\hat{z}_{jg}^{(1)} = z_{jg}^{(1)}$ .

**M-Step:** Update the estimates of  $\pi_g$ ,  $\mu_g$  and  $\Sigma_g$  by calculating

$$\pi_g^{(t+1)} = \frac{S_g}{\sum_{g=1}^G S_g}, \quad \mu_g^{(t+1)} = \frac{\sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)} \mathbf{x}_{ij}}{S_g},$$

$$\Sigma_g^{(t+1)} = \frac{\sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)} (\mathbf{x}_{jg} - \mu_g^{(t+1)}) (\mathbf{x}_{jg} - \mu_g^{(t+1)})'}{S_g},$$

where  $S_g = \sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)}$ .

This simplified form of the EM algorithm will prove useful when we discuss the EM algorithm in the case of a FSC with mixture of multivariate  $t$ -distributions (Section 3).

We note that the three different species of classification fall out naturally as special cases of FSC. If  $\alpha = 1$ , then all of the weight is given to the labelled observations, and the unlabelled observations are ignored. In this case, we are performing discriminant analysis. If  $\alpha = 0.5$ , then the labelled and unlabelled observations are given equal weight, and we are then performing semi-supervised classification. Finally, if  $\alpha = 0$ , then no weight is given to the labelled observations, and thus we are performing a cluster analysis (on the unlabelled observations). As mentioned in Section 1, the main unresolved issue with FSC is the selection of the weight  $\alpha$ .

### 2.4 The Multivariate $t$ -Distribution

The  $p$ -dimensional  $t$ -distribution with  $\nu$  degrees of freedom, location parameter  $\mu$  and scale matrix  $\Sigma$ , arises as a special case of a normal scale mixture (Peel and McLachlan, 2000). Specifically, we can write the normal scale mixture as

$$\epsilon \phi(\mathbf{x} \mid \mu, \Sigma) + (1 - \epsilon) \phi(\mathbf{x} \mid \mu, \nu \Sigma), \tag{8}$$

where  $\phi(\cdot)$  denotes the multivariate Gaussian density with mean  $\mu$  and covariance matrix  $\Sigma$ , and  $\epsilon$  is small. We can then rewrite (8) as

$$\int \phi(\mathbf{x} \mid \mu, \nu \Sigma) dH(w),$$

where

$$H(w) = \frac{1}{\Gamma(\frac{\nu}{2}) (\frac{\nu}{2})^{\frac{\nu}{2}}} w^{\frac{\nu}{2}-1} \exp \left\{ -\frac{2w}{\nu} \right\}, \tag{9}$$

$w > 0$  and  $\Gamma(\cdot)$  is the gamma function. Note that (9) is the probability density function of a  $\text{gamma}(\nu/2, \nu/2)$  random variable. The resulting density for the multivariate  $t$ -distribution is

$$f_t(\mathbf{x} \mid \boldsymbol{\mu}, \boldsymbol{\Sigma}, \nu) = \frac{\Gamma\left(\frac{\nu+p}{2}\right) |\boldsymbol{\Sigma}|^{-\frac{1}{2}}}{(\pi\nu)^{\frac{1}{2}p} \Gamma\left(\frac{\nu}{2}\right) \left[1 + \frac{\delta(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma})}{\nu}\right]^{\frac{1}{2}(\nu+p)}}, \quad (10)$$

where  $\delta(\mathbf{x}, \boldsymbol{\mu}, \boldsymbol{\Sigma}) = (\mathbf{x} - \boldsymbol{\mu})' \boldsymbol{\Sigma}^{-1} (\mathbf{x} - \boldsymbol{\mu})$  is the squared Mahalanobis distance. Maximum likelihood estimation for  $t$ -mixtures, in the context of model-based clustering, utilizes the introduction of latent variables  $W_{ig}$  such that

$$W_{ig} \mid z_{ig} = 1 \sim \text{gamma}(\nu_g/2, \nu_g/2).$$

## 2.5 Parsimonious Models

The eigen-decomposition of a matrix is widely used in both mathematics and multivariate statistics. In the context of mixture models, we can write a covariance, or scale, matrix in the form  $\boldsymbol{\Sigma}_g = \lambda_g \boldsymbol{\Lambda}_g \mathbf{D}_g \boldsymbol{\Lambda}_g'$ , where  $\lambda_g$  is a constant,  $\mathbf{D}_g$  is a diagonal matrix with entries that are proportional to the eigenvalues, and  $\boldsymbol{\Lambda}_g$  is a matrix of eigenvectors. We can then impose the following constraints:  $\lambda_g = \lambda$ ,  $\boldsymbol{\Lambda}_g = \boldsymbol{\Lambda}$ ,  $\boldsymbol{\Lambda}_g = \mathbf{I}$ ,  $\mathbf{D}_g = \mathbf{D}$ ,  $\mathbf{D}_g = \mathbf{I}$ , where  $\mathbf{I}$  is the identity matrix (Banfield and Raftery, 1993; Celeux and Govaert, 1995). Celeux and Govaert (1995) employ combinations of the above constraints to the covariance matrices in a Gaussian mixture model to form a family of 14 Gaussian parsimonious clustering models (GPCMs). Of these 14 models, 12 are extended to the  $t$ -distribution by Andrews and McNicholas (2012), with the result known as the tEIGEN family. These 12 models, together with the option to constrain  $\nu_g = \nu$ , leads to 24 different models in the tEIGEN family. The current form of the tEIGEN package (Andrews et al., 2016) in R (R Core Team, 2016) supports all 14 GPCM scale structures and hence a family of 28 tEIGEN models, which are summarized in Table 3 (Appendix A).

## 2.6 Model Selection Criteria

We now discuss a couple of criteria that are commonly used to select an appropriate parsimonious model. The most common approach is the Bayesian information criterion (BIC; Schwarz, 1978), which is given by

$$\text{BIC} = 2\ell_{\text{obs}}(\boldsymbol{\vartheta} \mid D_o) - p \log N,$$

where  $\ell_{\text{obs}}$  is the maximized observed likelihood,  $p$  is the number of free parameters, and  $N$  is the total number of data points. The BIC has been

frequently used for parsimonious model selection, e.g., Fraley and Raftery (1998) and McNicholas and Murphy (2008). Another criterion that is widely used is the integrated completed likelihood (ICL; Biernacki et al., 2000), which penalizes the BIC for classification uncertainty. The ICL can be approximated using the BIC:

$$\text{ICL} \approx \text{BIC} + 2 \sum_{i=1}^{n_g} \sum_{g=1}^G \text{MAP}(\hat{z}_{ig}) \log \hat{z}_{ig},$$

where

$$\text{MAP}(\hat{z}_{ig}) = \begin{cases} 1 & \text{if } \arg \max_{h=1, \dots, G} \{\hat{z}_{ih}\} = g, \\ 0 & \text{otherwise.} \end{cases}$$

### 3. FSC for $t$ -Mixtures

Before we discuss FSC for  $t$ -mixtures, we note that there is an alternative form of the weighted likelihood; for completeness, this is discussed in Appendix B. The main complication when using  $t$ -mixtures, compared to using Gaussian mixtures, is the update for the degrees of freedom. This update, unfortunately, has no closed form and has to be calculated using numerical methods. The incomplete weighted observed likelihood when using multivariate  $t$  component densities is

$$\begin{aligned} \mathcal{L}_{\text{obs}}(\boldsymbol{\vartheta} \mid D_o, \alpha) &= \left[ \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_t(\mathbf{x}_{1j} \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, \nu_g)] \right]^\alpha \\ &\times \left[ \prod_{j'=1}^{n_2} \sum_{g=1}^G \pi_g f_t(\mathbf{x}_{2j'} \mid \boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g) \right]^{1-\alpha}, \end{aligned}$$

where  $f_t(\cdot)$  is the density for the multivariate  $t$ -distribution defined in (10). To find  $\arg \max_{\boldsymbol{\vartheta}} \mathcal{L}_{\text{obs}}$ , we use a multicyle ECM algorithm similar to Andrews and McNicholas (2012). After initializing  $z_{jg}^{(i)}$  and  $w_{jg}^{(i)}$ , iteration  $t + 1$  of the multicyle ECM algorithm would proceed as follows:



**E-Step:** Update

$$\hat{z}_{jg}^{(2)} = \frac{\hat{\pi}_g f_t(\mathbf{x}_{2j} \mid \hat{\boldsymbol{\mu}}_g^{(t)}, \hat{\boldsymbol{\Sigma}}_g^{(t)}, \hat{\nu}_g^{(t)})}{\sum_{g=1}^G \hat{\pi}_g f_t(\mathbf{x}_{2j} \mid \hat{\boldsymbol{\mu}}_g^{(t)}, \hat{\boldsymbol{\Sigma}}_g^{(t)}, \hat{\nu}_g^{(t)})}, \quad (11a)$$

$$\hat{w}_{jg}^{(i)} = \frac{\hat{\nu}_g^{(t)} + p}{\hat{\nu}_g^{(t)} + \delta(\mathbf{x}_{ij}, \hat{\boldsymbol{\mu}}_g^{(t)}, \hat{\boldsymbol{\Sigma}}_g^{(t)}, \hat{\nu}_g^{(t)})}. \quad (11b)$$

**First CM-Step:** Update  $\hat{\pi}_g$ ,  $\hat{\boldsymbol{\mu}}_g$  and  $\hat{\nu}_g$ . The updates for  $\hat{\pi}_g$  and  $\hat{\boldsymbol{\mu}}_g$  are given in closed form as

$$\hat{\pi}_g^{(t+1)} = \frac{\sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)}}{\sum_{i=1}^2 \sum_{j=1}^{n_i} \sum_{g=1}^G \alpha_i \hat{z}_{jg}^{(i)}} \quad \text{and} \quad \hat{\boldsymbol{\mu}}_g^{(t+1)} = \frac{\sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)} \hat{w}_{jg}^{(i)} \mathbf{x}_{ij}}{\sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)} \hat{w}_{jg}^{(i)}}.$$

The updates for the degrees of freedom  $\nu_g$ , as mentioned before, do not have a closed form and have to be calculated using numerical methods.

In the unconstrained case, one has to solve (11c) for  $\hat{\nu}_g^{\text{new}}$  :

$$\begin{aligned} -\Psi\left(\frac{1}{2}\hat{\nu}_g^{\text{new}}\right) + \log\left(\frac{1}{2}\hat{\nu}_g^{\text{new}}\right) - \Psi\left(\frac{\hat{\nu}_g + p}{2}\right) - \log\left(\frac{\hat{\nu}_g + p}{2}\right) + 1 \\ + \frac{1}{m_g} \sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)} \left(\log \hat{\omega}_{jg}^{(i)} - \hat{\omega}_{jg}^{(i)}\right) = 0 \end{aligned} \quad (11c)$$

where

$$m_g = \sum_{i=1}^2 \sum_{j=1}^{n_i} \alpha_i \hat{z}_{jg}^{(i)}$$

and  $\Psi(\cdot)$  is the digamma function. Then, set  $\hat{\nu}_g^{(t+1)} = \hat{\nu}_g^{\text{new}}$ . Note that we used the `uniroot` function in `R` to solve (11c).

**E-Step:** Update  $\hat{z}_{jg}^{(2)}$  and  $\hat{w}_{jg}^{(i)}$  using (11a) and (11b) with current parameter estimates.

**Second CM Step:** Update  $\boldsymbol{\Sigma}_g$ . In the completely unconstrained case, the update is

$$\hat{\boldsymbol{\Sigma}}_g^{(t+1)} = \frac{1}{m_g} \sum_{i=1}^2 \sum_{j=1}^{n_i} \sum_{g=1}^G \alpha_i \hat{z}_{jg}^{(i)} \hat{\omega}_{jg}^{(i)} (\mathbf{x}_{ij} - \hat{\boldsymbol{\mu}}_g^{(t+1)})(\mathbf{x}_{ij} - \hat{\boldsymbol{\mu}}_g^{(t+1)})'.$$

We take this time to note that, except for the inclusion of the weights, the multicycle ECM algorithm described here is exactly the same as that described in Andrews and McNicholas (2012).

We perform k-means clustering (MacQueen, 1967) with 50 random starts to initialize the ECM algorithm, and use the Aitken acceleration (Aitken, 1926) procedure described in McNicholas et al. (2010) as our convergence criteria. Because of the updates for the degrees of freedom, fitting FSC with a  $t$ -mixture becomes more computationally expensive than fitting a Gaussian mixture. However, because of the heavier tails of the  $t$ -distribution, the  $t$ -mixture is more robust to outlying observations.

#### 4. Weight Selection Criteria

The ARI compares two different partitions of a dataset and, in the classification paradigm, a value of 1 corresponds to perfect classification, whereas a value of 0 indicates that the classification solution is as would be expected if the labels were randomly assigned. In Section 1, we point out that Vrbik and McNicholas (2015) use the ARI as a weight selection criteria for FSC. However, this is only useful when exploring the overall performance of FSC in simulations and datasets where all the labels are known (but some are treated as unknown). In a real classification scenario, not all the labels will be known and hence the ARI could not be used to select the weight  $\alpha$ . We, therefore, try other criteria for weight selection.

The first criteria we consider are the BIC and ICL. The results are not shown here but suffice it to say that various analyses revealed both of these criteria to be monotone in  $\alpha$  and a boundary point was always chosen. Three different classification-based criteria are considered: the entropy, an alternative form of the entropy (Celeux and Soromenho, 1996), and the  $U$  criterion (Bensmail et al., 1997).

In our case, the entropy  $E$  can be written

$$E = \sum_{i=1}^2 \sum_{j=1}^{n_i} \sum_{g=1}^G \text{MAP}(\hat{z}_{jg}^{(i)}) \log \hat{z}_{jg}^{(i)} = \sum_{j=1}^{n_2} \sum_{g=1}^G \text{MAP}(\hat{z}_{jg}^{(2)}) \log \hat{z}_{jg}^{(2)}, \quad (12)$$

where

$$\text{MAP}(\hat{z}_{jg}^{(i)}) = \begin{cases} 1 & \text{if } \hat{z}_{jg}^{(i)} = \max_{h=1,2,\dots,G}(\hat{z}_{jh}^{(i)}), \\ 0 & \text{otherwise,} \end{cases}$$

and taking  $0 \log 0 = 0$ . The entropy in (12) is always negative, unless there is no uncertainty in the clustering solution, in which case it is 0. When using this criterion, we choose the optimal weight to correspond to the maximum value of  $E$ .

An alternative form of the entropy is sometimes used that eliminates the MAP. The resulting criterion, in our case, is given by

$$A = \sum_{j=1}^{n_2} \sum_{g=1}^G \hat{z}_{jg}^{(2)} \log \hat{z}_{jg}^{(2)}.$$

Once again, we choose the optimal weight to correspond to the maximum value of  $A$ . The third, and final, classification-based criterion that we consider is the  $U$  criterion. In our case, this is given by

$$U = \sum_{i=1}^2 \sum_{j=1}^{n_i} \min_{g=1,2,\dots,G} (1 - \hat{z}_{jg}^{(i)}) = \sum_{j=1}^{n_2} \min_{g=1,2,\dots,G} (1 - \hat{z}_{jg}^{(2)}).$$

We observe that  $U$  is always positive and, if there is no uncertainty in the classification solution, then  $U = 0$ . Again, we choose the optimal weight to correspond to the maximum value of  $U$ .

In addition to these three classification-based criteria, we consider two non-parametric criteria. Before the BIC became popular, the sum of squares matrix was used as a basis for criteria to choose the number of groups in a model (see Gordon, 1981, Sec. 3.3, for discussion). Assuming that our data matrix  $\mathbf{X}$  has been partitioned into  $G$  groups, we can define the total sum of squares matrix to be

$$\mathbf{S} = \sum_{g=1}^G \sum_{i=1}^{n_g} (\mathbf{x}_{ig} - \bar{\mathbf{x}})(\mathbf{x}_{ig} - \bar{\mathbf{x}})',$$

where  $n_g$  is the number of points classified into group  $g$  and  $\bar{\mathbf{x}}$  is the grand mean. Using a decomposition of  $\mathbf{S}$  we can write

$$\mathbf{S} = \mathbf{W} + \mathbf{B},$$

where  $\mathbf{W}$  is the within cluster sum of squares matrix defined as

$$\mathbf{W} = \sum_{g=1}^G \sum_{i=1}^{n_g} (\mathbf{x}_{ig} - \bar{\mathbf{x}}_g)(\mathbf{x}_{ig} - \bar{\mathbf{x}}_g)',$$

where  $\bar{\mathbf{x}}_g$  is the sample mean of group  $g$ , and  $\mathbf{B}$  is the between cluster sum of squares matrix defined as

$$\mathbf{B} = \sum_{g=1}^G n_g (\bar{\mathbf{x}}_g - \bar{\mathbf{x}})(\bar{\mathbf{x}}_g - \bar{\mathbf{x}})'$$

Although the principle of using the sum of squares matrix was considered all the way back in the 1960s (e.g., Edwards and Cavalli-Sforza, 1965; Friedman and Rubin, 1967), it is still visible within the modern literature (e.g., Andrews and McNicholas, 2014). Herein (Section 5.5), two different criteria that use the within cluster sum of squares matrix  $\mathbf{W}$  are tried. The first criterion is based on minimizing the trace of  $\mathbf{W}$ , i.e.,  $\text{tr}(\mathbf{W})$ , and the second criterion is based on minimizing the determinant of  $\mathbf{W}$ , i.e.,  $\det(\mathbf{W})$ .

## 5. Analyses

### 5.1 Specifying the Number of Groups

For the purposes of our simulations and data analyses, we assume that the number of groups is equal to the number of components or classes present in the labelled points. However, this could be potentially problematic. For one, there could be a group present in the population that is not represented in the labelled data—this may be more likely if only a small proportion of the data points are labelled. Although perhaps less likely, it is also possible for the true number of groups to be less than that indicated by the labels. The former problem can be handled by fitting FSC with a different number of groups  $H \geq G$  in the cluster analysis component of the likelihood, and then using a criterion such as the BIC or ICL to choose the number of groups. The latter case, however, would need to be treated more carefully; likely in conjunction with a subject matter (data) expert.

### 5.2 Simulations

Simulations are performed, similar to those in Vrbik and McNicholas (2015), to demonstrate FSC with the multivariate  $t$ -distribution. In all, 100 datasets are simulated, each with 200 data points and two groups. The first group follows a  $t_2(\mathbf{0}, \Sigma_1, \nu_1)$  distribution, where  $\nu_1 = 3$ , and

$$\Sigma_1 = \begin{bmatrix} 1 & 0.7 \\ 0.7 & 1 \end{bmatrix}.$$

The second group is taken from a  $t_2(\Delta, \Sigma_2, \nu_2)$  distribution, where  $\Delta = [0, \Delta]'$ ,  $\nu_2 = 70$ , and

$$\Sigma_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

In this case, one group has a multivariate  $t$ -distribution, while the other group is approximately Gaussian, i.e.,  $\nu_2$  is quite large. This time, we take

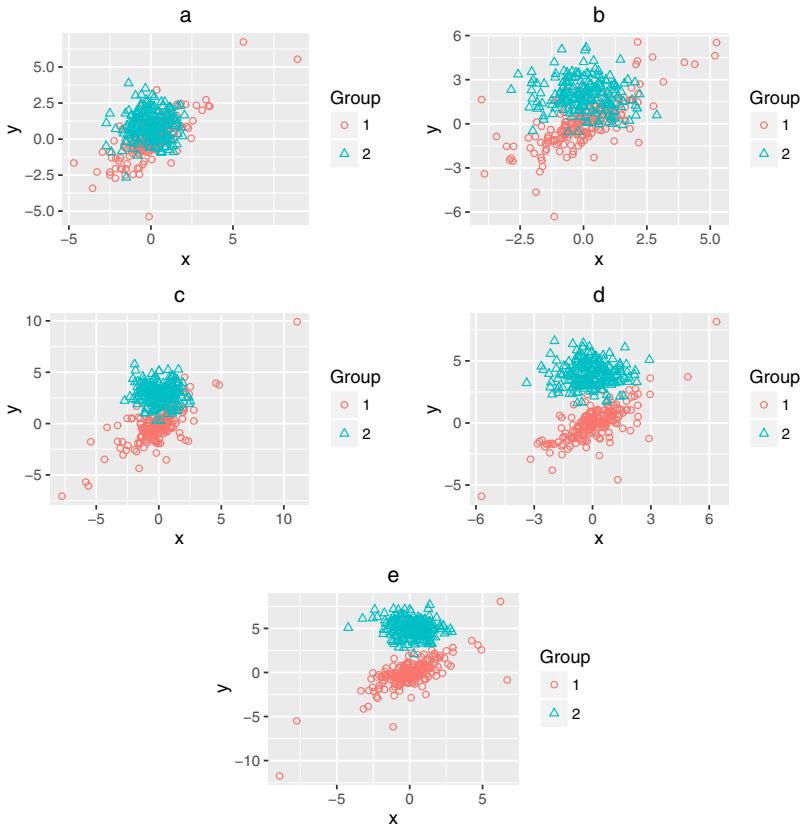


Figure 1. Typical datasets for (a)  $\Delta = 1$ , (b)  $\Delta = 2$ , (c)  $\Delta = 3$ , (d)  $\Delta = 4$ , (e)  $\Delta = 5$ .

$\Delta \in \{1, 2, 3, 4, 5\}$  and percentage of labelled data  $p \in \{10\%, 20\%, \dots, 80\%, 90\%\}$ . In Figure 1, we show example datasets for each  $\Delta$ .

To choose the weights for FSC, we consider 11 candidate values of  $\alpha$ ; specifically,  $\alpha \in \alpha_{\text{ARI}}$ , where  $\alpha_{\text{ARI}} = \{0, 0.1, 0.2, \dots, 1\}$ . Then, the ARI is calculated for each of these weights for the 100 datasets and the average ARI is computed for each weight. The weight with the highest average ARI is chosen. The resulting FSC solution for each weight  $\alpha$  is denoted by  $\text{FSC}_{\alpha}$ . Furthermore, for the FSC solution with the chosen weight resulting from the highest average ARI, the notation  $\text{FSC}_{\text{ARI}}$  is used. Finally, in the special cases corresponding to the three species of classification  $\alpha = 0, 0.5, 1$ , the FSC solution is denoted by  $\text{FSC}_{\text{clust}}$ ,  $\text{FSC}_{\text{class}}$  and  $\text{FSC}_{\text{DA}}$ , respectively.

In Figure 2, we give line plots for the case when  $\Delta = 1$ , where the average ARI is plotted against the percentage of labelled data  $p$  for each candidate weight. Further, a black dotted line is used to show the result for  $\text{FSC}_{\text{ARI}}$  with the corresponding chosen weight shown above each point. The left plot




Figure 2. For  $\Delta = 1$ : a)  $FSC_\alpha$  and  $FSC_{ARI}$  for  $\alpha \in \alpha_{ARI}$ , b)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$ .

shows the results when using all the weights and the right plot singles out the three different species of classification and  $FSC_{ARI}$ . The standard errors are calculated by taking the ARI for all 100 datasets of the chosen weight of  $FSC_{ARI}$  and calculating one (darker grey) and two (lighter grey) standard deviations from the mean ARI.

For  $\Delta = 1$ , we notice that the line for  $FSC_{clust}$  does not appear because the average ARI for each percentage of labelled data is quite small in comparison to the other weights (see Figure 2). Furthermore, for all other values of  $\Delta$ ,  $FSC_{clust}$  has the worst performance at higher percentages of labelled data, which is somewhat expected. We also see that all of the chosen weights correspond to a non-species solution. Furthermore, it is interesting to point out that, for lower percentages of labelled data, more weight is given to the labelled points and, at higher percentages, with the exception of 80%, less weight is given to the labelled observations. Similar results are given in Figures 3–6, where similar plots are shown for the other values of  $\Delta$ . For the remaining values of  $\Delta$ , of the 36 different cases, the chosen weight corresponds to a species of classification only nine times. Of these nine occurrences, eight of them correspond to semi-supervised classification, one corresponds to a discriminant analysis, and none of them correspond to a cluster analysis.

*Estimation*

In addition to classification performance, we also consider the accuracy of the parameter estimates. Parameter estimates for  $FSC_{ARI}$  from our most re-




Figure 3. For  $\Delta = 2$ : a)  $FSC_{\alpha}$  and  $FSC_{ARI}$  for  $\alpha \in \alpha_{ARI}$ , b)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$ .




Figure 4. For  $\Delta = 3$ : a)  $FSC_{\alpha}$  and  $FSC_{ARI}$  for  $\alpha \in \alpha_{ARI}$ , b)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$ .

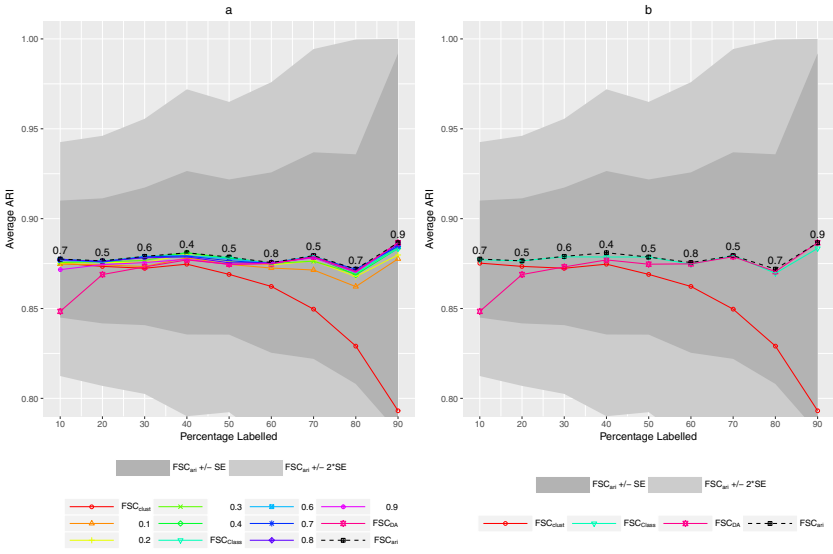


Figure 5. For  $\Delta = 4$ : a)  $FSC_{\alpha}$  and  $FSC_{ARI}$  for  $\alpha \in \alpha_{ARI}$ , b)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$ .

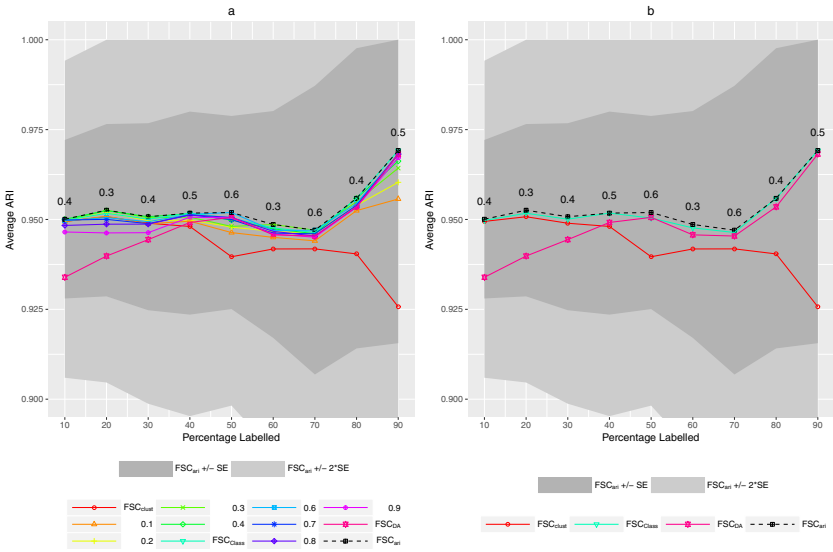


Figure 6. For  $\Delta = 5$ : a)  $FSC_{\alpha}$  and  $FSC_{ARI}$  for  $\alpha \in \alpha_{ARI}$ , b)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$ .



Table 1. Average parameter estimates for  $\Delta = 3$  for 20%, 50% and 80% of points labelled with component wise standard deviations in brackets

$\nu_1$ (sd)	$\mu_1$ (sd)	$\Sigma_1$ (sd)	$\nu_2$ (sd)	$\mu_2$ (sd)	$\Sigma_2$ (sd)
20%					
( $\alpha = 0.6$ )					
3.21 (0.766)	$\begin{bmatrix} -0.00698 \\ -0.00352 \end{bmatrix}$	$\begin{bmatrix} 1.01 & 0.703 \\ 0.7031.01 \end{bmatrix}$	63.2 (57.0)	$\begin{bmatrix} 0.00535 \\ 2.99 \end{bmatrix}$	$\begin{bmatrix} 0.988 & -0.00720 \\ -0.007200.978 \end{bmatrix}$
	$\left( \begin{bmatrix} 0.100 \\ 0.100 \end{bmatrix} \right)$	$\left( \begin{bmatrix} 0.2000.154 \\ 0.1540.184 \end{bmatrix} \right)$		$\left( \begin{bmatrix} 0.0772 \\ 0.0845 \end{bmatrix} \right)$	$\left( \begin{bmatrix} 0.133 & 0.0881 \\ 0.08810.138 \end{bmatrix} \right)$
50%					
( $\alpha = 0.6$ )					
3.19 (0.742)	$\begin{bmatrix} 0.00186 \\ 0.00476 \end{bmatrix}$	$\begin{bmatrix} 1.03 & 0.716 \\ 0.7161.03 \end{bmatrix}$	67.3 (57.4)	$\begin{bmatrix} -0.00270 \\ 3.00 \end{bmatrix}$	$\begin{bmatrix} 0.990 & 0.000940 \\ 0.0009400.980 \end{bmatrix}$
	$\left( \begin{bmatrix} 0.0956 \\ 0.0913 \end{bmatrix} \right)$	$\left( \begin{bmatrix} 0.1950.143 \\ 0.1430.170 \end{bmatrix} \right)$		$\left( \begin{bmatrix} 0.0760 \\ 0.0799 \end{bmatrix} \right)$	$\left( \begin{bmatrix} 0.127 & 0.0811 \\ 0.08110.140 \end{bmatrix} \right)$
80%					
( $\alpha = 0.4$ )					
3.20 (0.716)	$\begin{bmatrix} -0.00242 \\ 0.00122 \end{bmatrix}$	$\begin{bmatrix} 1.02 & 0.712 \\ 0.7121.02 \end{bmatrix}$	67.1 (52.7)	$\begin{bmatrix} -0.00254 \\ 3.00 \end{bmatrix}$	$\begin{bmatrix} 0.995 & -0.00215 \\ -0.002150.978 \end{bmatrix}$
	$\left( \begin{bmatrix} 0.0951 \\ 0.0906 \end{bmatrix} \right)$	$\left( \begin{bmatrix} 0.1940.149 \\ 0.1490.170 \end{bmatrix} \right)$		$\left( \begin{bmatrix} 0.0730 \\ 0.0737 \end{bmatrix} \right)$	$\left( \begin{bmatrix} 0.124 & 0.0804 \\ 0.08040.120 \end{bmatrix} \right)$

cent simulation are considered, for  $p = 20\%$ ,  $50\%$ , and  $80\%$  of points labelled and  $\Delta = 3$ . The results (Table 1) show that the estimates are very close to the actual values in all cases. We note that there is a lot of variability in the estimate for  $\nu_2$ —this is to be expected because the second component is approximately Gaussian.

### 5.3 Simulation with Three Groups

Finally, we perform a simulation with three groups. We follow the same procedure as the simulations previously discussed, this time with 100 observations in each group for a total of 300 observations for each of the 100 datasets, once again all from bivariate  $t$ -mixtures. The first two groups are simulated from exactly the same distributions as the previous simulations with  $\Delta = 2$ . For the third group, we took  $\mu_3 = (2, 2)$ ,  $\nu_3 = 10$  and

$$\Sigma_3 = \begin{bmatrix} 1 & -0.7 \\ -0.7 & 1 \end{bmatrix}.$$

A typical dataset is shown in Figure 7, where the three groups are moderately well separated but there this is still some overlap. In Figure 8, we show line plots, as before, and see that one of the three species is selected in only two of the nine cases.



Figure 7. Typical dataset for simulation with three groups.

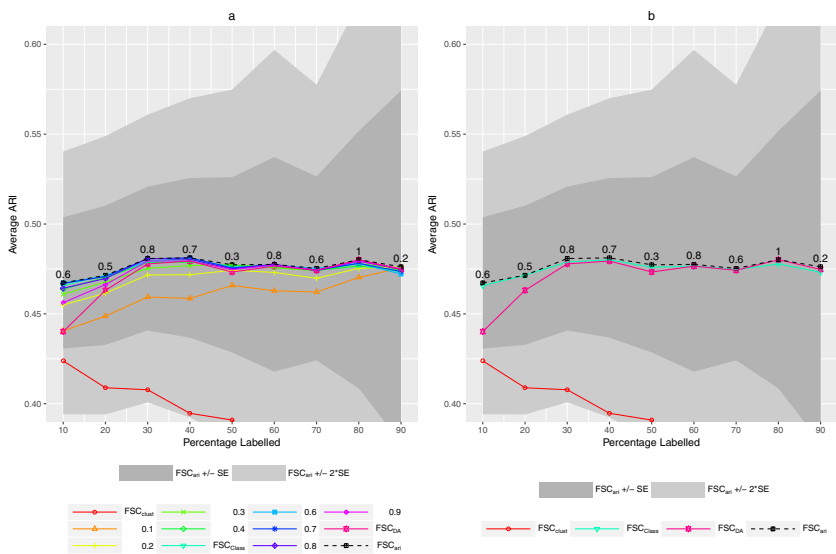


Figure 8. For the simulation with three groups: a)  $FSC_{\alpha}$  and  $FSC_{ARI}$  for  $\alpha \in \alpha_{ARI}$ , b)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$ .

## 5.4 Application to Datasets

We now look at a few datasets and compare the performance of FSC using a  $t$ -mixture and FSC with a Gaussian mixture. We took 100 random splits for each dataset for each percentage of labelled data,  $p \in \{10\%, 20\%, \dots, 80\%, 90\%\}$ . We used the same criterion as in the simulations, i.e., the ARI, to choose the optimal weight. As with the simulations we use a completely unconstrained model for both the covariance structure and, in the case of the  $t$ -mixtures, the degrees of freedom. For completeness, we note that we are not necessarily able to perform a discriminant analysis when the percent labelled is low or a cluster analysis when the percent labelled is high.

### *Iris Data*

The Anderson Iris data contains four different attributes of three different species of iris and is available in the R package `datasets`. The measurements (in centimetres) are the sepal length and width, and the petal length and width. The results are depicted in Figure 9. On the left hand side, we show the results for the  $t$ -mixture, and on the right hand side we show the results for the Gaussian mixture. Comparing these two plots, we see that the overall classification performance is similar between the  $t$ -mixture and the Gaussian mixture. Moreover, except at  $p = 60\%$ , the weights chosen for both the  $t$  and Gaussian mixtures are very similar if not exactly the same.

### *Crabs Data*

The crabs dataset consists of 5 measurements on four different types of rock crabs (two species, male and female in each species) and are available in the R package `MASS` (Venables and Ripley, 2002). These measurements are the frontal lobe size, carapace length and width, and the rear length and width. The results (Figure 10) show that, as for the iris data, the classification performance for the  $t$  and Gaussian mixtures are similar. Moreover, the weights chosen are very similar. It is interesting to note that almost all the weights are around 0.5

### *Wine Data*

The wine dataset from the R package `gclus` (Hurley, 2004) considers 13 characteristics of three different classes of wine. One interesting aspect of the results (Figure 11) is that, until one gets to the higher proportions of labelled data, the  $t$ -mixture performs slightly better than the Gaussian mixture. Another thing to note is that, similar to the crabs data, the cluster analysis

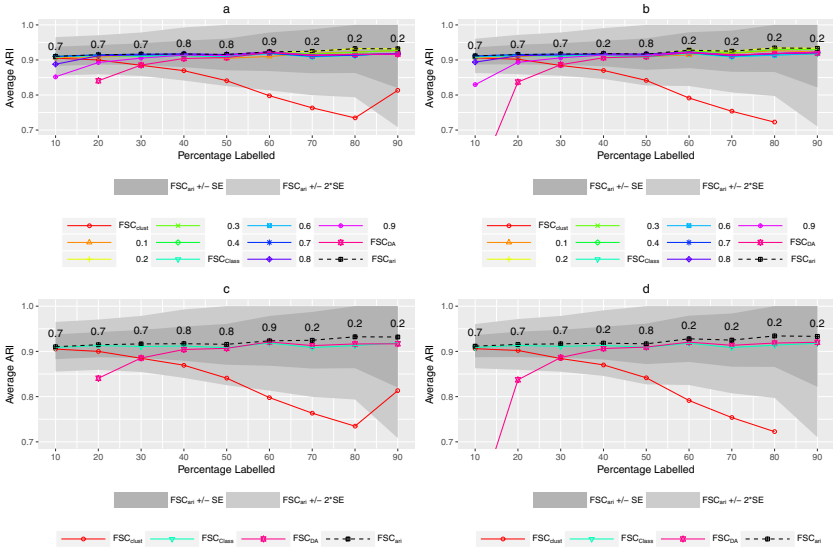


Figure 9.  $FSC_\alpha$  for  $\alpha \in \alpha_{ARI}$  and  $FSC_{ARI}$  for the iris data for: a) the  $t$ -mixture and b) for the Gaussian mixture.  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$  for c) the  $t$ -mixture, and d) the Gaussian mixture.

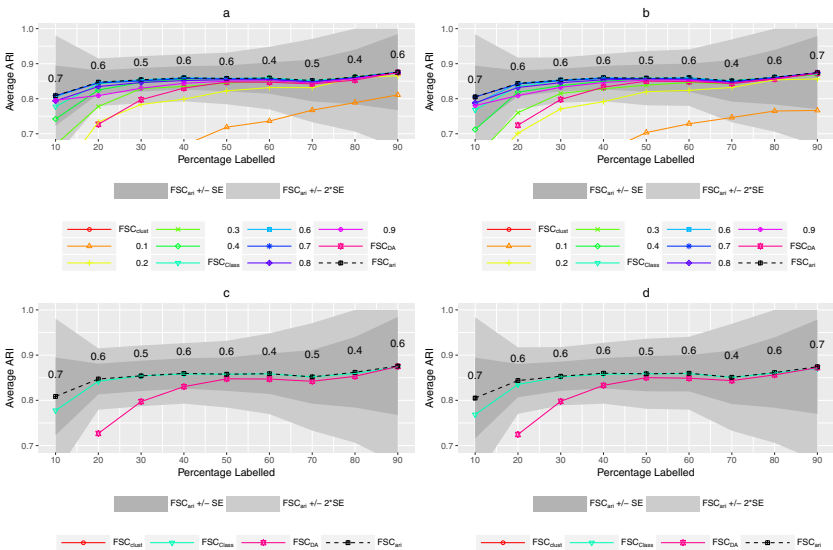


Figure 10.  $FSC_\alpha$  for  $\alpha \in \alpha_{ARI}$  and  $FSC_{ARI}$  for the crabs data for: a) the  $t$ -mixture and b) for the Gaussian mixture.  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$  for c) the  $t$ -mixture, and d) the Gaussian mixture.

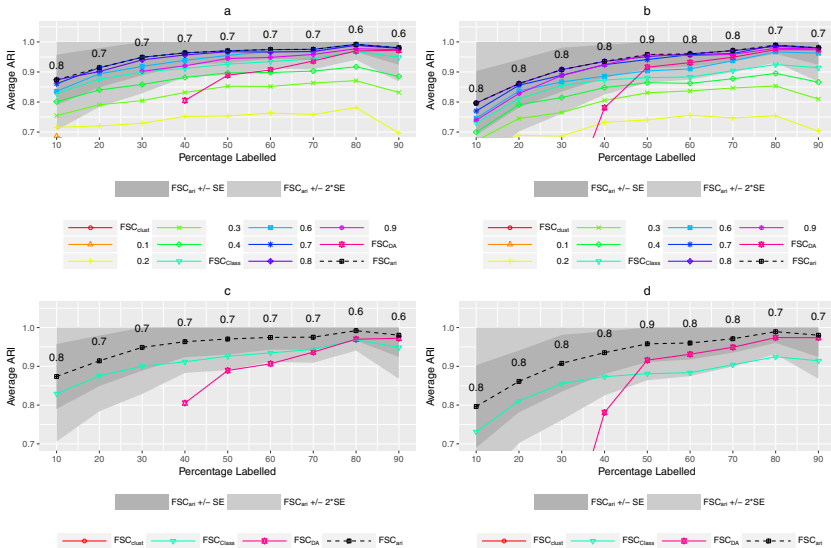


Figure 11.  $FSC_{\alpha}$  for  $\alpha \in \alpha_{ARI}$  and  $FSC_{ARI}$  for the wine data for: a) the  $t$ -mixture and b) for the Gaussian mixture.  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$  for c) the  $t$ -mixture, and d) the Gaussian mixture.

does not perform well in comparison to the other values of  $\alpha$ . Finally, the chosen weights for the  $t$ - and Gaussian mixtures are fairly similar and tend to choose larger weights for the labelled observations at all proportions.

### Bankruptcy Data

The bankruptcy data, found in the R package `mixGHD` (Tortora et al., 2015), consider the financial situation of 66 American firms: each firm was labelled as either bankrupt or financially sound. The results (Figure 12) show a greater difference between the  $t$ - and Gaussian mixtures when compared to the other datasets we have looked at. First, note the chosen weights. The weights chosen using a  $t$ -mixture are very different than those chosen when using the Gaussian mixture. The second item to note is that, similar to the wine data, the  $t$ -mixture gives better classification performance at lower percentages of labelled points. Finally, we note the difference in variability. For the Gaussian mixture, at lower percentages, we see a lot more variability in the error bars than for the  $t$ -mixture. Also, in general, there is more variability between the different weights for the Gaussian mixture. This could suggest that the selection of the weight should be treated a bit more carefully for the Gaussian mixture in this case, as the selection of a non-optimal weight can result in decreased classification performance. This is especially true, once again, at lower percentages of labelled points.

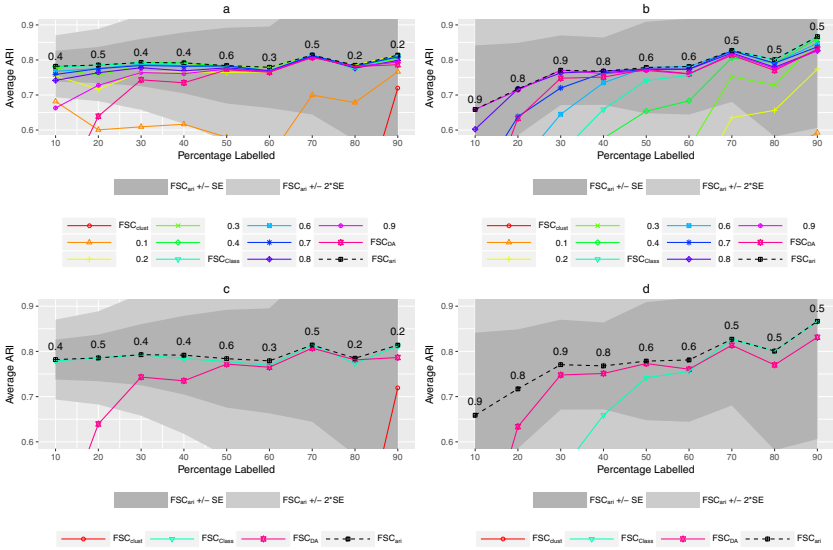


Figure 12.  $FSC_{\alpha}$  for  $\alpha \in \alpha_{\text{ARI}}$  and  $FSC_{\text{ARI}}$  for the bankruptcy data for: a) the  $t$ -mixture and b) for the Gaussian mixture.  $FSC_{\text{clust}}$ ,  $FSC_{\text{class}}$ ,  $FSC_{\text{DA}}$  and  $FSC_{\text{ARI}}$  for c) the  $t$ -mixture, and d) the Gaussian mixture.

### 5.5 Weight Selection Criteria for Parsimonious Models

In Section 4, five different weight selection criteria are discussed. In this section, we compare the performance of these criteria by considering FSC on  $t$ -mixtures for the wine, bankruptcy, crabs and iris datasets. We take 50 different splits for each dataset, with 80% of data labelled and use a mixture of multivariate  $t$ -distributions. We take the same candidate weights as before (see Section 5.2). For each candidate weight, we choose the model—i.e., the value of  $G$  and the covariance structure (Table 3, Appendix A)—using the BIC, and then calculate each of weight selection criteria mentioned earlier. We then choose the optimal weight, based on each of the selection criteria, and calculate the ARI. Also, we consider the highest ARI of all the weights after choosing the model to evaluate the overall performance of each of the criteria. In Figure 13, we show box plots of the resulting ARI values using each of the criteria, as well as the box plot for the distribution of the highest ARI.

The distributions of the ARI values for the three classification-based criteria show that the resulting ARI from the chosen weight is generally much lower than if we were to use the highest ARI. Moreover, the variability is generally much higher and especially so for the bankruptcy and crabs data. On the other hand,  $\text{tr}(\mathbf{W})$  performs well in comparison to the three

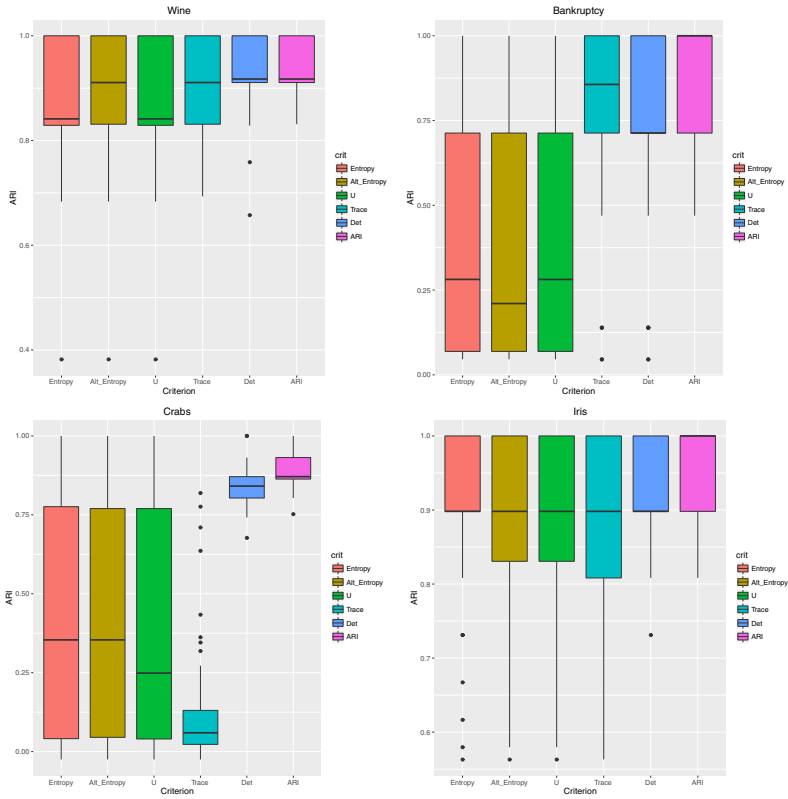


Figure 13. Distribution of ARI values for each of the criteria as well as the distribution of the highest ARI for the four datasets. The BIC was used to choose the model.

classification-based criteria for the wine and bankruptcy data. Furthermore, in the case of the bankruptcy data, it performs the best of all five criteria, when comparing the medians, and has a distribution closest to that of the highest ARI. However, in the case of the crabs data, it performs very poorly, and has the worst performance of the five criteria. For the iris data, the performance is similar to the alternative entropy and  $U$  criteria. Finally, we see that  $\det(\mathbf{W})$  performs well for all of the datasets. In the case of the wine data, except for a couple of outliers, the distribution is very similar to that for the highest ARI—this is quite remarkable when one considers that the ARI assumes knowledge of the true labels. Furthermore,  $\det(\mathbf{W})$  performs the best of all of the proposed criteria in all of the datasets except for the bankruptcy data. In this case,  $\text{tr}(\mathbf{W})$  performs better, but the inter-quartile ranges are very similar. Therefore, we propose  $\det(\mathbf{W})$  as a criterion to select the weight  $\alpha$  in FSC.

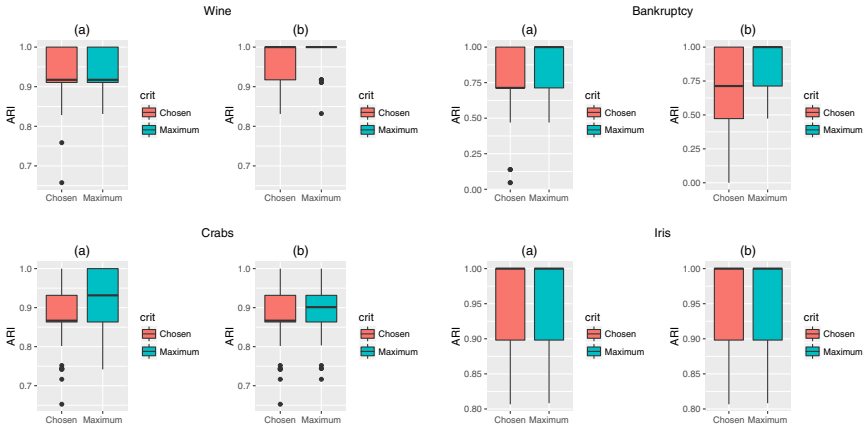


Figure 14. Distribution of ARI values for (a) the first procedure and (b) the second procedure for each of the four datasets.

*The Determinant as a Model Selection Criterion*

We have already seen that  $\det(\mathbf{W})$  appears to be an effective selection criterion for the weights in FSC. Now, we consider the possibility of using this criterion for model selection in general. To further explore this idea, we once again consider the four datasets and perform 50 random splits with 80% of the data points having known labels. This time, we consider two different procedures. In the first procedure, we proceed as before and choose the model based on the BIC, and then the weight using  $\det(\mathbf{W})$ . In the second, we choose the model based on  $\det(\mathbf{W})$  and then the weight also based on  $\det(\mathbf{W})$ . We once again take the ARI values after choosing the model and the weight using one of these two procedures, and we take the maximum ARI value amongst all of the weights. In Figure 14, we show box plots of the distributions of the ARI values. In (a) we show the results for the first procedure and, in (b), we show the results for the second procedure.

There are a few interesting items to note. First, for the wine dataset, we see that when using  $\det(\mathbf{W})$  to choose the model, the distribution of the maximum ARI has a lot less variability. Also, these maximum ARI values are generally larger after using  $\det(\mathbf{W})$  to choose the model. One final note on the wine dataset is that the median ARI values using procedure 2 is higher than those from procedure 1. For the bankruptcy data, we see that the distribution of the maximum ARI is the same regardless of using the BIC or  $\det(\mathbf{W})$  to choose the model. However, after choosing the weight, we see that the distribution of the ARI values for procedure 2 shows more variability than procedure 1. In the case of the crabs data, we see that the






Figure 15. Two different possible datasets with different organizations of labelled points with the true classification.

distribution of the ARI for the selected model and weight are approximately the same for both procedures; however, the maximum ARI is generally better when using the BIC to choose the model. Finally, for the Iris data, all of the distributions are very similar. The results are inconclusive in that neither procedure outperformed the other; however, the fact that the BIC did not outperform  $\det(\mathbf{W})$  for model selection is remarkable. In fact, the possibility of using  $\det(\mathbf{W})$  for model selection in model-based clustering, as alternative to the BIC, is worthy of further consideration.

### 5.6 Justification for a Cluster Analysis

If some of the points are labelled, it may not be immediately clear as to why a cluster analysis should even be considered. However, there are situations in which performing a cluster analysis is just as good, if not better, than putting more weight on the labelled observations. In Figure 15, we show two different situations where this would be the case. In Table 2, we look at the ARI and  $\det(\mathbf{W})$  for each of the weights for the two different cases. In the first case, only 10% of the points are labelled, and all labelled points are around the intersection of the two clusters. In this case, we see from the ARI and determinant values that we would only want give very little weight, or no weight, to the labelled observations.

In the first case, we see that a cluster analysis is actually better than using higher weights, and just as good as using smaller weights. In the second case, 90% of the points are labelled, and the unlabelled points lie on the outside of the two clusters. From the ARI and  $\det(\mathbf{W})$  values (Table 2), it is clear that all weights give perfect classification, including a cluster analysis, and thus a cluster analysis would perform just as well as the other weights in this case.

Table 2. ARI and determinant values for each candidate weight for both of the cases in Figure 15.

Weight	First Case		Second Case	
	ARI	Det.	ARI	Det.
0	0.9341	81006	1	82849
0.1	0.9341	81006	1	82849
0.2	0.9341	81006	1	82849
0.3	0.9126	81984	1	82849
0.4	0.9126	81984	1	82849
0.5	0.9126	81984	1	82849
0.6	0.8914	84250	1	82849
0.7	0.8914	84250	1	82849
0.8	0.8914	84250	1	82849
0.9	0.0075	178858	1	82849
1	-0.0016	187192	1	82849

## 6. Conclusions and Future Work

The major contribution of this paper is to encourage the use of  $\det(\mathbf{W})$  as a weight selection criterion in FSC. Although based on old ideas, and ideas that have not been fashionable for some time, this criteria is shown to outperform alternatives such as the near-ubiquitous BIC for weight selection. Furthermore, it performs comparably to the BIC in the model selection stage. As a secondary contribution, the FSC approach is shown to be mathematically tractable and effective for mixtures of multivariate  $t$ -distributions. For example, in our simulations, the selected weight very rarely corresponded to one of the three traditional species of classification. Furthermore, in our real data analyses, the use of a mixture of multivariate  $t$ -distributions was shown to either perform as well as or, in the case of the wine and bankruptcy datasets, better than the mixture of multivariate Gaussian distributions. This is likely due, at least in part, to the  $t$ -distribution being more robust to outliers than the Gaussian distribution. It is not unreasonable to expect that the FSC will also perform well with other non-Gaussian mixtures—the reader is referred to the recent review paper of McNicholas (2016b) for some discussion of non-Gaussian mixtures.

Future work will investigate using  $\det(\mathbf{W})$  as an alternative to the BIC for model selection in model-based clustering and classification in general. Using the FSC approach in a wider range of situations will also be explored. For example, FSC could be applied in the area of item response theory.

**Appendix A: tEIGEN Models**

Table 3. Model nomenclature and number of free covariance parameters of tEIGEN models with constrained (C), unconstrained (U) and identity (I) elements.

Model	$\lambda_g = \lambda$	$\Lambda_g = \Lambda$	$\mathbf{D}_g = \mathbf{D}$	$\nu_g = \nu$	No. of Free Covariance Parameters
CHC	C	I	I	C	2
CIU	C	I	I	U	$1 + G$
UIIC	U	I	I	C	$(G - 1) + 1$
UIIU	U	I	I	U	$(G - 1) + G$
CICC	C	I	C	C	$p + 1$
CICU	C	I	C	U	$p + G$
UICC	U	U	C	C	$p + (G - 1) + 1$
UICU	U	I	C	U	$p + (G - 1) + G$
CIUC	C	I	U	C	$Gp - (G - 1) + 1$
CIUU	C	I	U	U	$Gp - (G - 1) + G$
UIUC	U	I	U	C	$Gp + 1$
UIUU	U	I	U	U	$Gp + G$
CCCC	C	C	C	C	$[p(p + 1)/2] + 1$
CCCU	C	C	C	U	$[p(p + 1)/2] + G$
UCCC	U	C	C	C	$[p(p + 1)/2] + (G - 1) + 1$
UCCU	U	C	C	U	$[p(p + 1)/2] + (G - 1) + G$
CUCC	C	U	C	C	$G[p(p + 1)/2] - (G - 1)(p) + 1$
CUCU	C	U	C	U	$G[p(p + 1)/2] - (G - 1)(p) + G$
UUCU	U	U	C	C	$G[p(p + 1)/2] - (G - 1)(p - 1) + 1$
UUCU	U	U	C	U	$G[p(p + 1)/2] - (G - 1)(p - 1) + G$
CCUC	C	C	U	C	$[p(p + 1)/2] + (G - 1)(p - 1) + 1$
CCUU	C	C	U	U	$[p(p + 1)/2] + (G - 1)(p - 1) + G$
CUUC	C	U	U	C	$G[p(p + 1)/2] - (G - 1) + 1$
CUUU	C	U	U	U	$G[p(p + 1)/2] - (G - 1) + G$
UCUC	U	C	U	C	$G[p(p + 1)/2] + (G - 1)p + 1$
UCUU	U	C	U	U	$G[p(p + 1)/2] + (G - 1)p + G$
UUUC	U	U	U	C	$G[p(p + 1)/2] + 1$
UUUU	U	U	U	U	$G[p(p + 1)/2] + G$

**Appendix B: Alternative Form of the Likelihood**

**B.1 Alternative Likelihood**

We have already seen that the observed weighted likelihood can be written as in (5) and, analogous to (6), the associated complete-data weighted likelihood can be written as

$$\mathcal{L}_{\text{comp}}(\boldsymbol{\theta} | D_C, \alpha) = \prod_{i=1}^2 \left[ \prod_{j=1}^{n_i} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{ij} | \boldsymbol{\theta})]^{z_{jg}^{(i)}} \right]^{\alpha_i} \quad (13)$$

Dempster et al. (1977) state that when integrating the complete-data likelihood over the space of unknown quantities, in our case  $\mathbb{Z}_2$ , it is desired that the result should be the observed likelihood. The observed likelihood as given in (5), however, does not satisfy this property. Indeed,

$$\begin{aligned}
 & \int_{\mathbb{Z}_2} \mathcal{L}_{\text{comp}}(\vartheta | D_C, \alpha) d\mathbf{z}_2 = \\
 & \int_{\mathbb{Z}_2} \left\{ \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} | \boldsymbol{\theta}_g)]^{z_{jg}^{(1)} \alpha} \times \prod_{j=1}^{n_2} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{2j} | \boldsymbol{\theta}_g)]^{z_{jg}^{(2)} (1-\alpha)} \right\} d\mathbf{z}_2 \\
 & = \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} | \boldsymbol{\theta}_g)]^{z_{jg}^{(1)} \alpha} \prod_{j=1}^{n_2} \left\{ \int_{\mathbb{Z}_2} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{2j} | \boldsymbol{\theta}_g)]^{z_{jg}^{(2)} (1-\alpha)} d\mathbf{z}_2 \right\} \\
 & = \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} | \boldsymbol{\theta})]^{z_{jg}^{(1)} \alpha} \prod_{j=1}^{n_2} \left\{ \sum_{\mathbf{z}_j \in \mathfrak{B}} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{2j} | \boldsymbol{\theta})]^{z_{jg}^{(2)} (1-\alpha)} \right\} \\
 & = \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} | \boldsymbol{\theta})]^{z_{jg}^{(1)} \alpha} \prod_{j=1}^{n_2} \left\{ \sum_{g=1}^G [\pi_g f_g(\mathbf{x}_{2j} | \boldsymbol{\theta})]^{(1-\alpha)} \right\}, \tag{14}
 \end{aligned}$$

where

$$\mathfrak{B} = \left\{ \mathbf{z}_j = \left( z_{j1}^{(2)}, z_{j2}^{(2)}, \dots, z_{jG}^{(2)} \right) \mid z_{jg}^{(2)} \in \{0, 1\}, \forall g \in \{1, 2, \dots, G\}, \sum_{g=1}^G z_{jg}^{(2)} = 1 \right\}.$$

Clearly, this is not the same as the form given in (5). Therefore, to maintain the relationship between the complete and incomplete weighted likelihood as presented in Dempster et al. (1977), we consider using the form of the incomplete weighted likelihood given in (14) and denote this by  $\mathcal{L}_{\text{alt}}$ .

Note that there are two extreme cases that should be considered separately. The first extreme case is when  $\alpha = 0$ :

$$\begin{aligned}
 & \int_{\mathbb{Z}_2} \mathcal{L}_{\text{comp}}(\vartheta | D_C, \alpha = 0) d\mathbf{z}_2 \\
 & = \int_{\mathbb{Z}_2} \prod_{j=1}^{n_2} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{2j} | \boldsymbol{\theta}_g)]^{z_{jg}^{(2)} (1-\alpha)} d\mathbf{z}_2 = \prod_{j=1}^{n_2} \sum_{g=1}^G \pi_g f_g(\mathbf{x}_{2j} | \boldsymbol{\theta}),
 \end{aligned}$$

which is equivalent to (14) when  $\alpha = 0$ . The second extreme case, which turns out to be more interesting, is when  $\alpha = 1$ :

$$\mathcal{L}_{\text{alt}}(\boldsymbol{\vartheta} | D_o) = G^{m_2} \prod_{j=1}^{n_1} \prod_{g=1}^G [\pi_g f_g(\mathbf{x}_{1j} | \boldsymbol{\theta}_g)]^{z_{jg}^{(1)}} = G^{m_2} \mathcal{L}_{\text{DA}}(\boldsymbol{\vartheta} | D_L). \quad (15)$$

When  $\alpha = 1$  we are performing discriminant analysis; however, the resulting likelihood in (15) is not equal to  $\mathcal{L}_{\text{DA}}$  but is instead proportional to  $\mathcal{L}_{\text{DA}}$ .

For both the original and altered observed likelihoods, the complete-data likelihood is identical. Therefore, if we were to take a Gaussian mixture model, the updates in the M-step would be the same as those given in Vrbik and McNicholas (2015), regardless of whether the original or alternative likelihood were used. However, the updates for  $\hat{z}_{jg}^{(2)}$  in the E-step would become

$$\hat{z}_{jg}^{(2)} = \frac{\left[ \pi_g^{(t)} \phi(\mathbf{x}_{2j} | \boldsymbol{\mu}_g^{(t)}, \boldsymbol{\Sigma}_g^{(t)}) \right]^{(1-\alpha)}}{\sum_{g=1}^G \left[ \pi_g^{(t)} \phi(\mathbf{x}_{2j} | \boldsymbol{\mu}_g^{(t)}, \boldsymbol{\Sigma}_g^{(t)}) \right]^{(1-\alpha)}}.$$

## B.2 Simulation Comparing the Original and Altered Likelihoods

We perform simulations to compare the performance of the original and altered likelihoods. We simulate 100 datasets with 300 samples: 150 of these sample belong to one group which follows a  $\mathcal{N}_2(\mathbf{0}, \boldsymbol{\Sigma}_1)$ , and the remaining 150 belong to another group which follows a  $\mathcal{N}_2(\boldsymbol{\Delta}, \boldsymbol{\Sigma}_2)$ , where  $\boldsymbol{\Delta} = [0, \Delta]^t$ , and

$$\boldsymbol{\Sigma}_1 = \begin{bmatrix} 1 & 0.7 \\ 0.7 & 1 \end{bmatrix}, \quad \boldsymbol{\Sigma}_2 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}.$$

We take  $\Delta \in \{1, 5\}$  corresponding to different levels of clustering difficulty. For each dataset, we consider  $p \in \{10\%, 20\%, \dots, 80\%, 90\%\}$ , where  $p$  is the percentage of labelled data.

To choose the weights for FSC, we looked at 11 different values of  $\alpha$ . These values were taken to be  $\alpha \in \alpha_{\text{ARI}}$ , where  $\alpha_{\text{ARI}} = \{0, 0.1, 0.2, \dots, 1\}$ . We then calculate the ARI for each of these weights for the 100 datasets and take the average ARI for each weight. We then choose the weight that had the highest average ARI. We denote the resulting FSC solution for each weight  $\alpha$  by  $\text{FSC}_\alpha$ . Furthermore, denote by  $\text{FSC}_{\text{ARI}}$  the FSC solution with the chosen weight resulting from the highest average ARI. Finally, in the special cases corresponding to the three species of classification

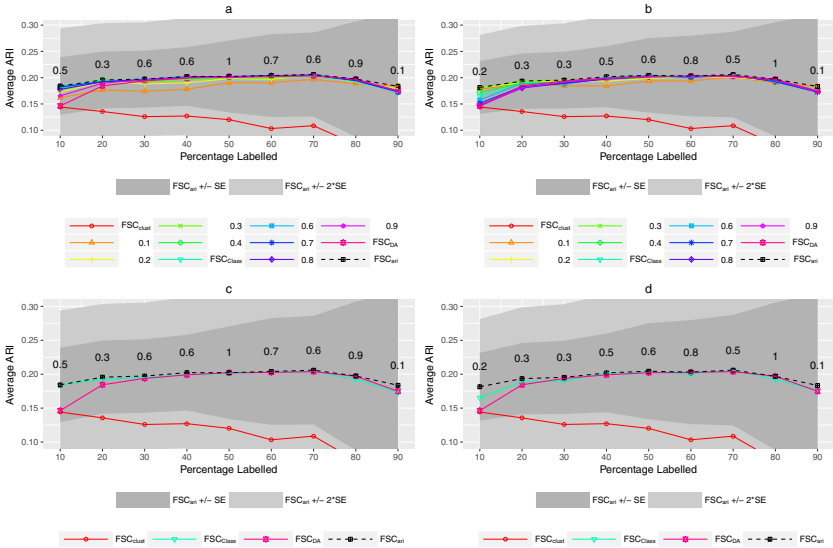


Figure 16. For  $\Delta = 1$ : (a) and (b)  $FSC_\alpha$  and  $FSC_{ARI}$  ( $\alpha \in \alpha_{ARI}$ ) for the original and altered likelihood respectively. (c) and (d)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$  for the original and altered likelihood respectively.

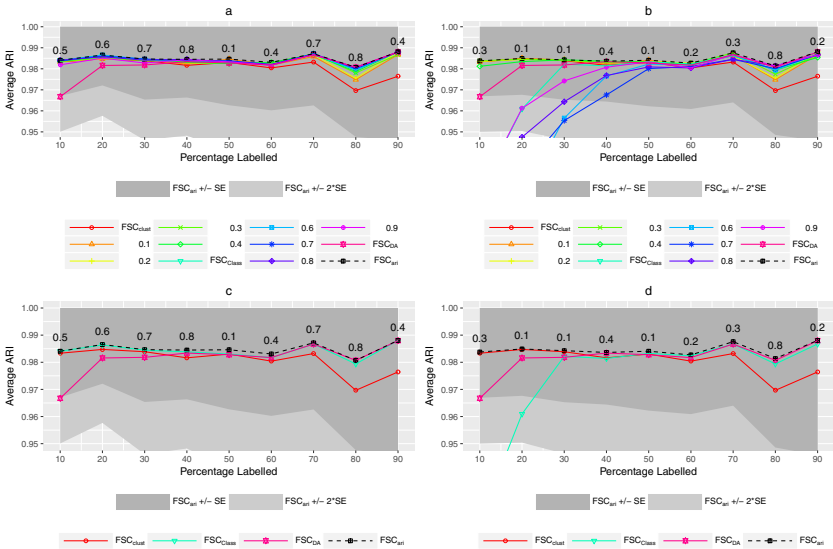


Figure 17. For  $\Delta = 5$ : (a) and (b)  $FSC_\alpha$  and  $FSC_{ARI}$  ( $\alpha \in \alpha_{ARI}$ ) for the original and altered likelihood respectively. (c) and (d)  $FSC_{clust}$ ,  $FSC_{class}$ ,  $FSC_{DA}$  and  $FSC_{ARI}$  for the original and altered likelihood respectively.

$\alpha = 0, 0.5, 1$ , we denote the FSC solution by  $FSC_{\text{clust}}$ ,  $FSC_{\text{class}}$  and  $FSC_{\text{DA}}$ , respectively.

In Figures 16 and 17, we show different line plots for  $\Delta = 1$  and  $\Delta = 5$ , respectively. In each plot, the average ARI is plotted against the percentage of labelled data  $p$ . A dotted black line is used to show the result for  $FSC_{\text{ARI}}$  with the corresponding chosen weight shown above each point. The first row in each plot shows the results when using all the weights, and the second row singles out the three different species of classification and  $FSC_{\text{ARI}}$ . The standard errors were calculated by taking the ARI for all 100 datasets of the chosen weight of  $FSC_{\text{ARI}}$  and calculating one (darker grey) and two (lighter grey) standard deviations from the mean ARI.

In general, the overall classification performance between the altered and original likelihoods are similar. The chosen weights for  $FSC_{\text{ARI}}$ , however, differ between the two forms of the likelihood. For  $\Delta = 1$ , this difference is less pronounced than for  $\Delta = 5$ . More specifically, for  $\Delta = 1$ , the difference between the weights for all but 10%, 30% and 50% differ by at most 0.1 if they are not exactly the same. For  $\Delta = 5$ , however, the differences between the chosen weights are greater, and there are fewer proportions for which the difference is small. We also see that at lower percentages of labelled data, there is more variability in the average ARI between the different weights.

In conclusion, although the choice of the weights is different between the two likelihoods, the overall classification performance when using the chosen weight in each case is very similar. Moreover, the altered form is not strictly a likelihood. Accordingly, we use the original, and more natural form, form of the likelihood for FSC.

## References

- AITKEN, A.C. (1926), “A Series Formula for the Roots of Algebraic and Transcendental Equations”, *Proceedings of the Royal Society of Edinburgh*, 45, 14–22.
- ANDREWS, J.L., and MCNICHOLAS, P.D. (2011a), “Extending Mixtures of Multivariate t-Factor Analyzers”, *Statistics and Computing*, 21(3), 361–373.
- ANDREWS, J.L., and MCNICHOLAS, P.D. (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.
- ANDREWS, J.L., and MCNICHOLAS, P.D. (2012), “Model-Based Clustering, Classification, and Discriminant Analysis via Mixtures of Multivariate  $t$ -Distributions: The  $t$ EIGEN Family”, *Statistics and Computing*, 22(5), 1021–1029.
- ANDREWS, J.L., and MCNICHOLAS, P.D. (2014), “Variable Selection for Clustering and Classification”, *Journal of Classification*, 31(2), 136–153.
- ANDREWS, J.L., WICKINS, J.R., BOERS, N.M., and MCNICHOLAS, P.D. (2016), *teigen: Model-Based Clustering and Classification with the Multivariate  $t$  Distribution*, R Package Version 2.2.0.

- BANFIELD, J.D., and RAFTERY, A.E. (1993), "Model-Based Gaussian and Non-Gaussian Clustering", *Biometrics*, 49(3), 803–821.
- BAUM, L.E., PETRIE, T., SOULES, G., and WEISS, N. (1970), "A Maximization Technique Occurring in the Statistical Analysis of Probabilistic Functions of Markov Chains", *Annals of Mathematical Statistics*, 41, 164–171.
- BENSMAIL, H., CELEUX, G., RAFTERY, A., and ROBERT, C. (1997), "Inference in Model-Based Cluster Analysis", *Statistics and Computing*, 7, 1–10.
- BIERNACKI, C., CELEUX, G., and GOVAERT, G. (2000), "Assessing a Mixture Model for Clustering with the Integrated Completed Likelihood", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 22(7), 719–725.
- CELEUX, G., and GOVAERT, G. (1995), "Gaussian Parsimonious Clustering Models", *Pattern Recognition*, 28(5), 781–793.
- CELEUX, G., and SOROMENHO, G. (1996), "An Entropy Criterion for Assessing the Number of Clusters in a Mixture Model", *Journal of Classification*, 13, 195–212.
- DANG, U.J., BROWNE, R.P., and MCNICHOLAS, P.D. (2015), "Mixtures of Multivariate Power Exponential Distributions", *Biometrics*, 71(4), 1081–1089.
- DEMPSTER, A.P., LAIRD, N.M., and RUBIN, D.B. (1977), "Maximum Likelihood from Incomplete Data via the EM Algorithm", *Journal of the Royal Statistical Society: Series B*, 39(1), 1–38.
- EDWARDS, A.W.F., and CAVALLI-SFORZA, L.L. (1965), "A Method for Cluster Analysis", *Biometrics*, 21, 362–375.
- FRALEY, C., and RAFTERY, A.E. (1998), "How Many Clusters? Which Clustering Methods? Answers via Model-Based Cluster Analysis", *The Computer Journal*, 41(8), 578–588.
- FRANCZAK, B.C., BROWNE, R.P., and MCNICHOLAS, P.D. (2014), "Mixtures of Shifted Asymmetric Laplace Distributions", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(6), 1149–1157.
- FRANCZAK, B.C., TORTORA, C., BROWNE, R.P., and MCNICHOLAS, P.D. (2015), "Unsupervised Learning via Mixtures of Skewed Distributions with Hypercube Contours", *Pattern Recognition Letters*, 58(1), 69–76.
- FRIEDMAN, H.P., and RUBIN, J. (1967), "On Some Invariant Criteria for Grouping Data", *Journal of the American Statistical Association*, 62, 1159–1178.
- GORDON, A.D. (1981), *Classification*, London: Chapman and Hall.
- HUBERT, L., and ARABIE, P. (1985), "Comparing Partitions", *Journal of Classification*, 2(1), 193–218.
- HURLEY, C. (2004), "Clustering Visualizations of Multivariate Data", *Journal of Computational and Graphical Statistics*, 13(4), 788–806.
- INGRASSIA, S., MINOTTI, S.C., PUNZO, A., and VITTADINI, G. (2015), "The Generalized Linear Mixed Cluster-Weighted Model", *Journal of Classification*, 32(1), 85–113.
- INGRASSIA, S., MINOTTI, S.C., PUNZO, A., and VITTADINI, G. (2012), "Local Statistical Modeling via the Cluster-Weighted Approach with Elliptical Distributions", *Journal of Classification*, 29(3), 363–401.
- LEE, S., and MCLACHLAN, G.J. (2014), "Finite Mixtures of Multivariate Skew T-Distributions: Some Recent and New Results", *Statistics and Computing*, 24, 181–202.
- LEE, S.X., and MCLACHLAN, G.J. (2013), "On Mixtures of Skew Normal and Skew t-Distributions", *Advances in Data Analysis and Classification*, 7(3), 241–266.
- LIN, T.-I. (2010), "Robust Mixture Modeling Using Multivariate Skew t Distributions", *Statistics and Computing*, 20(3), 343–356.
- LIN, T.-I., MCLACHLAN, G.J., and LEE, S.X. (2016), "Extending Mixtures of Factor Models Using the Restricted Multivariate Skew-Normal Distribution", *Journal of Multivariate*



- Analysis*, 143, 398–413.
- LIN, T.-I., MCNICHOLAS, P.D., and HSIU, J.H. (2014), “Capturing Patterns via Parsimonious t Mixture Models”, *Statistics and Probability Letters*, 88, 80–87.
- MACQUEEN, J. (1967), “Some Methods for Classification and Analysis of Multivariate Observations”, in *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*, Berkeley, University of California Press, pp. 281–297.
- MCNICHOLAS, P.D. (2016a), *Mixture Model-Based Classification*, Boca Raton: Chapman and Hall/CRC Press.
- MCNICHOLAS, P.D. (2016b), “Model-Based Clustering”, *Journal of Classification*, 33(3), 331–373.
- MCNICHOLAS, P.D., and MURPHY, T.B. (2008), “Parsimonious Gaussian Mixture Models”, *Statistics and Computing*, 18(3), 285–296.
- MCNICHOLAS, P.D., MURPHY, T.B., MCDAID, A.F., and FROST, D. (2010), “Serial and Parallel Implementations of Model-Based Clustering via Parsimonious Gaussian Mixture Models”, *Computational Statistics and Data Analysis*, 54(3), 711–723.
- MURRAY, P.M., BROWNE, R.P., and MCNICHOLAS, P.D. (2014a), “Mixtures of Skew-t Factor Analyzers”, *Computational Statistics and Data Analysis*, 77, 326–335.
- MURRAY, P.M., BROWNE, R.P., and MCNICHOLAS, P.D. (2017a), “Hidden Truncation Hyperbolic Distributions, Finite Mixtures Thereof, and Their Application for Clustering”, *Journal of Multivariate Analysis*, 161, 141–156.
- MURRAY, P.M., BROWNE, R.P., and MCNICHOLAS, P.D. (2017b), “A Mixture of SDB Skew-t Factor Analyzers”, *Econometrics and Statistics*, 3, 160–168.
- MURRAY, P.M., MCNICHOLAS, P.D., and BROWNE, R.P. (2014b), “A Mixture of Common Skew-t Factor Analyzers”, *Stat*, 3(1), 68–82.
- ORCHARD, T., and WOODBURY, M.A. (1972), “A Missing Information Principle: Theory and Applications”, in *Proceedings of the Sixth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Theory of Statistics*, eds. L.M. Le Cam, J. Neyman, and E.L. Scott, Berkeley: University of California Press, pp. 697–715.
- PEEL, D., and MCLACHLAN, G.J. (2000), “Robust Mixture Modelling Using the t Distribution”, *Statistics and Computing*, 10(4), 339–348.
- PUNZO, A., and MCNICHOLAS, P.D. (2017), Robust Clustering in Regression Analysis via the Contaminated Gaussian Cluster-Weighted Model”, *Journal of Classification*, 34(2), 249–293.
- R CORE TEAM (2016), *R: A Language and Environment for Statistical Computing*, Vienna, Austria: R Foundation for Statistical Computing.
- SCHWARZ, G. (1978), “Estimating the Dimension of a Model”, *The Annals of Statistics*, 6(2), 461–464.
- SCOTT, A.J., and SYMONS, M.J. (1971), “Clustering Methods Based on Likelihood Ratio Criteria”, *Biometrics*, 27, 387–397.
- STEANE, M.A., MCNICHOLAS, P.D., and YADA, R. (2012), “Model-Based Classification via Mixtures of Multivariate t-Factor Analyzers”, *Communications in Statistics – Simulation and Computation*, 41(4), 510–523.
- SUBEDI, S., PUNZO, A., INGRASSIA, S., and MCNICHOLAS, P.D. (2013), “Clustering and Classification via Cluster-Weighted Factor Analyzers”, *Advances in Data Analysis and Classification*, 7(1), 5–40.
- SUBEDI, S., PUNZO, A., INGRASSIA, S., and MCNICHOLAS, P.D. (2015), “Cluster-Weighted t-Factor Analyzers for Robust Model-Based Clustering and Dimension Reduction”, *Statistical Methods and Applications*, 24(4), 623–649.
- TIEDEMAN, D.V. (1955), “On the Study of Types”, in *Symposium on Pattern Analysis*, ed.

- S.B. Sells, Randolph Field, Texas: Air University, U.S.A.F. School of Aviation Medicine.
- TORTORA, C., BROWNE, R.P., FRAN CZAK, B.C., and MCNICHOLAS, P.D. (2015), *MixGHD: Model Based Clustering, Classification and Discriminant Analysis Using the Mixture of Generalized Hyperbolic Distributions*, R Package Version 1.8.
- VENABLES, W.N., and RIPLEY, B.D. (2002), *Modern Applied Statistics with S* (4th ed.), New York: Springer.
- VRBIK, I., and MCNICHOLAS, P.D. (2012), “Analytic Calculations for the EM Algorithm for Multivariate Skew-t Mixture Models”, *Statistics and Probability Letters*, 82(6), 1169–1174.
- VRBIK, I., and MCNICHOLAS, P.D. (2014), “Parsimonious Skew Mixture Models for Model-Based Clustering and Classification”, *Computational Statistics and Data Analysis*, 71, 196–210.
- VRBIK, I., and MCNICHOLAS, P.D. (2015), “Fractionally-Supervised Classification”, *Journal of Classification*, 32(3), 359–381.
- WOLFE, J.H. (1965), “A Computer Program for the Maximum Likelihood Analysis of Types”, Technical Bulletin 65-15, U.S. Naval Personnel Research Activity.