



Clustering via finite nonparametric ICA mixture models

Xiaotian Zhu¹ · David R. Hunter²

Received: 31 March 2017 / Revised: 14 August 2018 / Accepted: 16 August 2018 /
Published online: 28 August 2018
© Springer-Verlag GmbH Germany, part of Springer Nature 2018

Abstract

We propose a novel extension of nonparametric multivariate finite mixture models by dropping the standard conditional independence assumption and incorporating the independent component analysis (ICA) structure instead. This innovation extends nonparametric mixture model estimation methods to situations in which conditional independence, a necessary assumption for the unique identifiability of the parameters in such models, is clearly violated. We formulate an objective function in terms of penalized smoothed Kullback–Leibler distance and introduce the nonlinear smoothed majorization–minimization independent component analysis algorithm for optimizing this function and estimating the model parameters. Our algorithm does not require any labeled observations a priori; it may be used for fully unsupervised clustering problems in a multivariate setting. We have implemented a practical version of this algorithm, which utilizes the FastICA algorithm, in the R package *icamix*. We illustrate this new methodology using several applications in unsupervised learning and image processing.

Keywords Independent component analysis · Kernel density estimation · Nonparametric estimation · Penalized smoothed likelihood · Unsupervised learning

Mathematics Subject Classification 62H30 · 62G07

1 Introduction

Cluster analysis, or clustering, is one of several general approaches to the problem of unsupervised learning, that is, classification when no class labels are given. In

The authors gratefully acknowledge the support of Award No. DMS-1209007 from the National Science Foundation.

✉ David R. Hunter
dhunter@stat.psu.edu

¹ Natera Inc., San Carlos, CA, USA

² Department of Statistics, Pennsylvania State University, University Park, PA, USA

practice, clustering is often based on heuristic ideas and intuitive measures, such as hierarchical clustering or k-means clustering, that do not assume a probability model. By contrast, this paper focuses on model-based clustering, in which data are viewed as coming from a mixture of probability distributions, each representing a cluster. Typically, the distributions are assumed to come from a parametric family such as normal (Fraley and Raftery 1998, 2002; Banfield and Raftery 1993), and group membership is learned from data by estimation algorithms that are often variations of the expectation-maximization method described by Dempster et al. (1977). Since the early work of Wolfe (1963) and others, the literature on model-based clustering has expanded enormously. Indeed, there are several book-length treatments of mixture models, such as Frühwirth-Schnatter (2006) and McLachlan and Peel (2000).

The advent of easily accessible computing power has given rise to semi- and non-parametric methods that avoid the standard assumption that the cluster densities come from a known parametric family, and applications and extensions of these methods are growing more common in the literature. A semiparametric model-based clustering analysis for DNA microarray data can be found in Han and Davis (2006). Azzalini and Torelli (2007) propose nonparametric density estimation using Delaunay triangulation for clustering via identification of subpopulations with regions with high density of the underlying probability distribution. Li et al. (2007) develop a clustering approach based on mode identification by applying new optimization techniques to a nonparametric density estimator. Vichi (2008) fits semiparametric clustering models to dissimilarity data. In Zhang et al. (2009), a semiparametric model is introduced to account for varying impacts of factors over clusters by using cluster-level covariates. Mallapragada et al. (2010) propose a non-parametric mixture model (NMM) for data clustering. Guglielmi et al. (2014) fit Bayesian semiparametric logit models to grouped data of in-hospital survival outcomes of patients hospitalized with ST-segment Elevation Myocardial Infarction diagnosis. Certain mixtures of linear regressions also fall under the category of semiparametric model-based clustering. For instance, Hunter and Young (2012) present an algorithm for estimating parameters in a mixture-of-regressions model in which the errors are assumed to be independent and identically distributed but no other distributional assumption is made. Huang et al. (2013) propose nonparametric finite mixture-of-regression models for analysis of U.S. housing price index (HPI) data. Vandekerkhove (2013) studies estimation of a semiparametric mixture-of-regressions model of two components when one component is known. Bajari et al. (2011) views a game abstractly as a semiparametric mixture distribution and studies the semiparametric efficiency bound of this model. Finally, Butucea and Vandekerkhove (2014) consider a semiparametric mixture of two distributions that are equal up to a shift parameter.

The current article combines recent advances in methods for multivariate non-parametric finite mixture models under an assumption that we refer to as the *conditional independence assumption* with another method, called independent components analysis (ICA), that solves one of the main drawbacks of this assumption. To illustrate this drawback, let us first introduce the modeling framework: suppose that r -dimensional vectors $\mathbf{Y}_i = (Y_{i1}, Y_{i2}, \dots, Y_{ir})^\top$, $1 \leq i \leq n$, are a simple random sample from a finite mixture of $m > 1$ components with positive mixing propor-

tions $\lambda_1, \lambda_2, \dots, \lambda_m$ that sum up to 1, and (Lebesgue measurable) density functions q_1, q_2, \dots, q_m respectively. This gives the mixture density

$$d(\mathbf{y}) = \sum_{j=1}^m \lambda_j q_j(\mathbf{y}) \quad (1)$$

for $\mathbf{y} \in R^r$. The sole assumption imposed on the densities q_1, q_2, \dots, q_m is that the coordinates of \mathbf{y} are independent given the component from which \mathbf{y} is sampled, so that Eq. (1) becomes

$$d(\mathbf{y}) = \sum_{j=1}^m \lambda_j \prod_{k=1}^r q_{jk}(y_k). \quad (2)$$

The basic idea of conditional independence is outlined by Hall and Zhou (2003), and, notably, Allman et al. (2009) prove the generic identifiability of the parameters in Eq. (2) for $r \geq 3$ under some weak assumptions. Chauveau et al. (2015) present a survey of the growing literature on the theory and algorithmic treatment of model (1) under the conditional independence assumption. The algorithms in this article have their roots in the EM-like algorithm of Benaglia et al. (2009), which is later modified by Levine et al. (2011). An alternative estimation method to EM-like algorithms is the method of moments approach described in Anandkumar et al. (2012). More recently, Bonhomme et al. (2016a, b) provide new estimation algorithms for this model and prove consistency and asymptotic normality of the estimators, an important statistical innovation missing from earlier work. Related work on mixtures of nonparametric hidden Markov models (HMMs) is presented by Gassiat and Rousseau (2016), Gassiat et al. (2016), and Castro et al. (2016); these authors describe links between their HMM work and the model of Eq. (2).

Although the conditional independence assumption of Eq. (2) is important theoretically due to its guarantee of identifiable parameters despite essentially no assumptions other than $r \geq 3$, it is clearly not appropriate for some clustering problems. As a simple example, consider the well-known Fisher Iris data depicted in Fig. 1. In a model-based clustering scenario, the goal of estimation would be to learn the shapes of the three four-dimensional distributions, one for each species, without the benefit of the species labels. Under the conditional independence assumption of Eq. (2), no bivariate plot should exhibit correlation within any of the categories. Yet it is clear in Fig. 1a that nonzero within-species correlation exists, so any correct classifier of the three species would necessarily violate the conditional independence assumption.

To remedy this shortcoming, the algorithm we propose in this article combines existing work on nonparametric mixture models with the ideas of independent components analysis (ICA). The basic idea of ICA, as elucidated, for example, by Hyvarinen et al. (2002), is to find a linear transformation of a multivariate dataset under which its coordinates are as close to independent as possible.

We can see the result of applying our algorithm to the iris dataset in Fig. 1b, in which each mixture component is associated with its own linear transformation. The categorizations displayed in this figure are based on the highest probability of each point among the three possible categories, and we observe that 7 of the 150

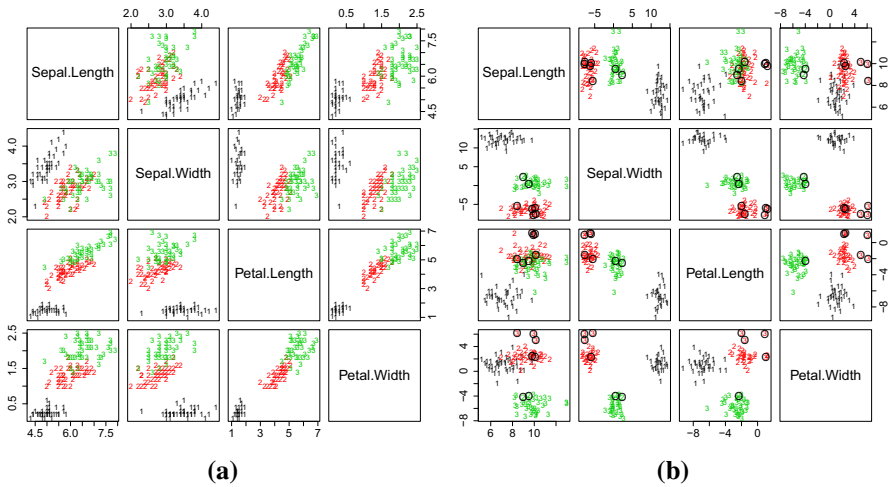


Fig. 1 **a** At left, the four-dimensional iris dataset with $n = 150$ and three distinct species. **b** At right, the same dataset classified according to our algorithm into three groups, where each group is transformed linearly in an attempt to achieve conditional independence. The numbers are the correct species labels, while the seven points misclassified by the algorithm are circled

points are incorrectly classified—recall that the mixture model parameter estimates are calculated without taking labels into account—but the correlation structure evident in the left-hand plots has been eliminated. In the remainder of this article, we describe the algorithm, discuss the potential issues of identifiability that arise, and illustrate the algorithm's performance on three datasets including one in which $n = 10,000$ and $r = 144$. The result of our work is an algorithm we call the NSMM-ICA algorithm, which we implement in the **icamix** package for R (R Core Team 2015), available at <http://cran.r-project.org/web/packages/icamix/index.html>.

The novelty of the current article is its combination of the non-parametric mixture structure with the linear transformations of ICA; previous work on model-based clustering using ICA has imposed parametric assumptions on the component density functions. Lee et al. (1999b) and Lee et al. (2000) propose parametric ICA mixture models with algorithms based on the infomax principle for various unsupervised classification problems. Shah et al. (2004) apply the ICA mixture model methodology to the problem of unsupervised classification of hyperspectral or multispectral imagery where image data are captured at multiple or a continuous range of frequencies across the electromagnetic spectrum. This is an important application of remote sensing and land cover classification. Palmer et al. (2008) derive an asymptotic Newton algorithm for Quasi-maximum likelihood estimation of the parametric ICA mixture model and presents its application to EEG segmentation. Finally, Salazar et al. (2010) extend the ICA mixture model methodology of Lee et al. (1999b) and others by positing that the component density functions actually have the form of kernel density estimates—i.e., the components are themselves mixtures of parametric functions of a fixed and known form (the kernel density itself), with one component per observation. Salazar et al.

(2015) apply a similar idea in an agglomerative clustering framework; however, in our work we assume that the number of mixture components is fixed and known.

As mentioned above, our work drops the parametric assumption of all the previous ICA mixture model literature that we are aware of. Yet it does bear some similarities to the work described in, for example, Salazar et al. (2010); Salazar et al. (2015), because we also employ kernel density estimation in order to approximate the unknown underlying univariate density functions that we describe below. There is the additional similarity that these papers minimize a Kullback–Leibler divergence, which has a similar form to our penalized and smoothed Kullback–Leibler divergence objective function. However, a crucial distinction is that we do not assume the knowledge of any of the category labels a priori; our algorithm is designed to handle completely unsupervised multivariate clustering problems. In addition, there are methods other than ICA in the literature based on the same idea of exploring multivariate data to determine coordinate systems having some desirable property such as inter-coordinate independence. Invariant co-ordinate selection (Tyler et al. 2009; Miettinen et al. 2015), or ICS, is one such method, and Peña et al. (2010) in particular uses ICS to search for cluster structure in the data based on the eigenvalues of a kurtosis matrix. This work does not assume an underlying non-parametric mixture structure as in the current article, yet in principle it would be possible to combine non-parametric mixtures with ICS instead of ICA.

2 The nonparametric ICA mixture model

Most previous work on model (2) assumes that we observe the random sample $\mathbf{Y}_1, \dots, \mathbf{Y}_n$. However, in the current article we generalize this previous work by adding the assumption that the observed data are $\mathbf{X}_1, \dots, \mathbf{X}_n$, where $\mathbf{X}_i = \mathbf{A}_j \mathbf{Y}_i$ for some invertible $r \times r$ matrix \mathbf{A}_j , conditional on \mathbf{Y}_i being generated from the j th component density q_j . In other words, we introduce additional parameters $\mathbf{A}_1, \dots, \mathbf{A}_m$, one for each mixture component, consisting of the matrices that linearly transform the latent \mathbf{Y}_i with independent coordinates into the observed \mathbf{X}_i . When there is no mixture structure, this assumption is exactly the independent component analysis (ICA) framework as described by Hyvarinen et al. (2002). NB: The word “component” in “ICA” is replaced by “coordinate” or “dimension” in the terminology of this article; here, “component” refers to one of the mixture densities.

To aid notation, let us define $q_{\mathbf{A}}$ for any nonnegative function q on R^r and invertible $r \times r$ matrix \mathbf{A} as

$$q_{\mathbf{A}}(\mathbf{x}) = q(\mathbf{A}^{-1}\mathbf{x})|\det \mathbf{A}|^{-1},$$

which is the density function of a linearly transformed random variable having density q after left-multiplication by \mathbf{A} .

Our ICA mixture model may thus be described formally as follows: We observe a random sample $\mathbf{X}_1, \dots, \mathbf{X}_n$ from the mixture density

$$g(\mathbf{x}) = \sum_{j=1}^m \lambda_j f_j(\mathbf{x}), \quad (3)$$

where

$$f_j(\mathbf{x}) = (q_j)_{A_j}(\mathbf{x}) \quad (4)$$

and

$$q_j(\mathbf{y}) = \prod_{k=1}^r q_{jk}(y_k). \quad (5)$$

For each observed \mathbf{X}_i , we shall define the usual latent variables

$$Z_{ij} = I\{\mathbf{X}_i \text{ is drawn from the } j\text{th mixture component}\}$$

and $\mathbf{Y}_i = \mathbf{A}_j^{-1}\mathbf{X}_i$ for the unique j such that $Z_{ij} = 1$. For estimation purposes, we write

$$(e_j)_{A_j} = \lambda_j f_j,$$

so $e_j(\mathbf{x}) = \lambda_j q_j(\mathbf{x})$. Since any constant multiple of \mathbf{Y}_i can be absorbed into the A_j matrices, we mitigate against non-identifiable parameters by further assuming

$$\text{Var } Y_{ik} = 1 \quad (6)$$

for all i and for $1 \leq k \leq r$. Finally, we assume for each j , $1 \leq j \leq m$, at most one of the density functions q_{j1}, \dots, q_{jr} is a normal density function. The reason for this last assumption is that ICA operates on standardized versions of the data, and thus if a subset of the linearly transformed \mathbf{Y} coordinates is multivariate normal, the standardized versions are always standard multivariate normal so there is no way to uniquely identify an ICA transformation A_j . Thus, the non-normality assumption, along with assumptions (4), (5), and (6), are commonly used in the literature on ICA (Hyvarinen et al. 2002). Since the q_{jk} are not assumed to follow any parametric form, we call the model

$$g(\mathbf{x}) = \sum_{j=1}^m \lambda_j |\det A_j|^{-1} \prod_{k=1}^r q_{jk} \left([A_j^{-1}\mathbf{x}]_k \right) \quad (7)$$

a nonparametric ICA mixture model.

It is not known whether the parameters in Eq. (7) are uniquely identifiable in the case of perfect information about the form of $g(\mathbf{x})$. Our empirical experience with our algorithm, of which we provide examples in Sect. 5, suggests that parameter estimation is well-behaved, yet this important theoretical question remains. For example, general identifiability does not follow directly from the facts that the q_{jk} and λ_j parameters are uniquely determined by Eq. (2) and the A_j parameters are uniquely determined by Eq. (5) together with $\mathbf{X} = A_j \mathbf{Y}$, even under the usual assumptions that are stated above. It may be possible to extend the methods of Allman et al. (2009) to prove identifiability, but for now we are in a situation analogous to the period just prior to the publication of that article, when estimation algorithms existed for cases in which only special cases of identifiability had been addressed in the literature and no general identifiability result had yet been established. It is also important to realize that even in a case where

parameters are theoretically within the set of identifiable parameters, estimation may still be difficult when the true parameters happen to be near the boundary of that set. However, here again we can point to our empirical experience in an example such as the iris dataset of Sect. 1. Quite often, the iris species have been modeled in the literature as multivariate normal distributions, suggesting that the data are generated from a set of parameters quite close to non-identifiability; yet in practice, our algorithm appears to find reasonable A_j estimates as shown in Fig. 1b.

3 Parameter estimation

This section introduces an MM-like algorithm that seeks to estimate the parameters $\mathbf{e} = (e_1, e_2, \dots, e_m)$ and $\mathbf{A} = (A_1, A_2, \dots, A_m)$ by minimizing a function that gives in some sense the distance between the empirical data distribution and the theoretical mixture distribution determined by the parameters.

We begin by defining some operators that will aid notation. Much of the development of this section follows the recent work of Levine et al. (2011) and Zhu and Hunter (2016); the novelty here is in the incorporation of the A_j matrices into the usual conditional independence framework, which requires some delicacy.

First, we define the linear smoothing, or convolution, operators S_h and S_h^* on $L^1(R^r)$. Let $s_h(\cdot, \cdot) \in L^1(R \times R)$ be a nonnegative kernel function satisfying

$$\int s_h(v, z) dz = \int s_h(v, z) dv = 1 \tag{8}$$

for $v, z \in R$. Here, $h > 0$ is a user-specified tuning parameter often referred to as a bandwidth in smoothing contexts. For any $f \in L^1(R^r)$, define $S_h f$ and $S_h^* f$ by

$$(S_h f)(\mathbf{x}) = \int \tilde{s}_h(\mathbf{x}, \mathbf{u}) f(\mathbf{u}) d\mathbf{u} \quad \text{and} \quad (S_h^* f)(\mathbf{x}) = \int \tilde{s}_h(\mathbf{u}, \mathbf{x}) f(\mathbf{u}) d\mathbf{u},$$

where

$$\tilde{s}_h(\mathbf{x}, \mathbf{u}) = \prod_{k=1}^r s_h(x_k, u_k) \quad \text{for } \mathbf{x}, \mathbf{u} \in R^r. \tag{9}$$

Furthermore, let

$$(\mathcal{N}_h f)(\mathbf{x}) = \exp[(S_h^* \log f)(\mathbf{x})].$$

Notice that (\mathcal{N}_h) as an operator on $L^1(R^r)$ is nonlinear, as it is the exponentiation of the linear operator S_h applied to the logarithm of a function. This nonlinear smoothing operator plays an important role in the algorithm.

Finally, we reproduce the projection-multiplication operator of Zhu and Hunter (2016), defined as

$$(Pf)(\mathbf{x}) = \frac{\left[\prod_{k=1}^r \int f(\mathbf{x}) dx_1 dx_2 \cdots dx_{k-1} dx_{k+1} \cdots dx_r \right]}{[f f]^{(r-1)}}. \tag{10}$$

Zhu and Hunter (2016) point out that when f is a density on R^d , the right side of (10) simplifies because the denominator is 1, and also that the P and S_h operators commute, i.e., $(P \circ S_h)f = (S_h \circ P)f$.

Let us consider the hypothetical case of a known target density $g(\mathbf{x})$, which we sometimes call the infinite sample size case. To estimate the parameters $\mathbf{e} = (e_1, e_2, \dots, e_m)$ and $\mathbf{A} = (A_1, A_2, \dots, A_m)$, the idea is to minimize a measure of the distance between g and the mixture density determined by the parameters. Due to mathematical considerations explained in Levine et al. (2011), we wish to first apply the nonlinear smoother to the mixture density and then minimize the Kullback–Leibler distance between g and this nonlinearly smoothed density. We therefore propose in this hypothetical case to minimize

$$\int g(\mathbf{x}) \log \left[g(\mathbf{x}) / \sum_{j=1}^m [\mathcal{N}_h e_j]_{A_j}(\mathbf{x}) \right] d\mathbf{x} + \int \left[\sum_{j=1}^m (e_j)_{A_j}(\mathbf{x}) \right] d\mathbf{x}$$

with respect to \mathbf{e} and \mathbf{A} . To analyze an actual dataset, we would replace the g density by the empirical distribution of the data, which leads to the objective function

$$\ell(\mathbf{e}, \mathbf{A}) = - \sum_{i=1}^n \log \sum_{j=1}^m [\mathcal{N}_h e_j]_{A_j}(\mathbf{x}_i) + \int \left[\sum_{j=1}^m (e_j)_{A_j}(\mathbf{x}) \right] d\mathbf{x} \quad (11)$$

to be minimized with respect to \mathbf{e} and \mathbf{A} .

Two aspects of Eq. (11) are worth noticing. First, the second integral is part of a penalty term whose presence guarantees a convenient property of the functional parameters $\mathbf{e} = (e_1, e_2, \dots, e_m)$ that minimize $\ell(\mathbf{e}, \mathbf{A})$ for a fixed \mathbf{A} , and this property is explained below in Eq. (12). Second, the definition of $\ell(\mathbf{e}, \mathbf{A})$ uses $\sum_j [\mathcal{N}_h e_j]_{A_j}(\mathbf{x})$ instead of $\sum_j \mathcal{N}_h [(e_j)_{A_j}](\mathbf{x})$; that is, the nonlinear smoothing is applied *before* the linear transformation. The intuition is that after the transformation, the data are no longer conditionally independent and standardized, so the smoothing would affect each dimension very differently if it were applied after the transformation.

An advantage of using the \mathbf{e} parameters instead of $\boldsymbol{\lambda}$ and \mathbf{q} is that the latter parameterization requires the constraint that every q_j is a density function. With the \mathbf{e} parameters, such a constraint is unnecessary: As a straightforward corollary of Theorem 2.1 of Zhu and Hunter (2016), any minimizer $(\tilde{\mathbf{e}}, \tilde{\mathbf{A}})$ of (11) must satisfy

$$\int \left[\sum_{j=1}^m (e_j)_{A_j}(\mathbf{x}) \right] d\mathbf{x} = \int \sum_{j=1}^m \tilde{e}_j(\mathbf{x}) d\mathbf{x} = 1. \quad (12)$$

4 The NSMM-ICA algorithm

Here, we derive an iterative algorithm for solving the main problem of minimizing Eq. (11). The algorithm is based on the MM framework, which stands for

majorization-minimization (Hunter and Lange 2004) and which involves constructing and minimizing an alternative to the $\ell(\mathbf{e}, \mathbf{A})$ function with respect to \mathbf{e} and \mathbf{A} at each iteration.

4.1 Majorizing the objective function

Given the current estimate $\mathbf{e}^{(0)}$ and $\mathbf{A}^{(0)}$, let us define

$$w_j^{(0)}(\mathbf{x}) = \frac{[\mathcal{N}_h e_j^{(0)}]_{\mathbf{A}_j^{(0)}}(\mathbf{x})}{\sum_{j'=1}^m [\mathcal{N}_h e_{j'}^{(0)}]_{\mathbf{A}_{j'}^{(0)}}(\mathbf{x})}.$$

Since $\sum_j w_j^{(0)}(\mathbf{x}) = 1$, Jensen’s inequality gives

$$\begin{aligned} \ell(\mathbf{e}, \mathbf{A}) - \ell(\mathbf{e}^{(0)}, \mathbf{A}^{(0)}) &= - \int g(\mathbf{x}) \log \sum_{j=1}^m w_j^{(0)}(\mathbf{x}) \frac{(\mathcal{N}_h e_j)_{\mathbf{A}_j}(\mathbf{x})}{(\mathcal{N}_h e_j^{(0)})_{\mathbf{A}_j^{(0)}}(\mathbf{x})} \, d\mathbf{x} + \int \left(\sum_{j=1}^m (e_j)_{\mathbf{A}_j} - \sum_{j=1}^m (e_j^{(0)})_{\mathbf{A}_j^{(0)}} \right) \\ &\leq - \int g(\mathbf{x}) \sum_{j=1}^m w_j^{(0)}(\mathbf{x}) \log \frac{(\mathcal{N}_h e_j)_{\mathbf{A}_j}(\mathbf{x})}{(\mathcal{N}_h e_j^{(0)})_{\mathbf{A}_j^{(0)}}(\mathbf{x})} \, d\mathbf{x} + \int \left(\sum_{j=1}^m (e_j)_{\mathbf{A}_j} - \sum_{j=1}^m (e_j^{(0)})_{\mathbf{A}_j^{(0)}} \right). \end{aligned}$$

Thus, if we let

$$b^{(0)}(\mathbf{e}, \mathbf{A}) = - \int g(\mathbf{x}) \sum_{j=1}^m w_j^{(0)}(\mathbf{x}) \cdot \log (\mathcal{N}_h e_j)_{\mathbf{A}_j}(\mathbf{x}) \, d\mathbf{x} + \int \left(\sum_{j=1}^m (e_j)_{\mathbf{A}_j} \right),$$

then

$$\ell(\mathbf{e}, \mathbf{A}) - \ell(\mathbf{e}^{(0)}, \mathbf{A}^{(0)}) \leq b^{(0)}(\mathbf{e}, \mathbf{A}) - b^{(0)}(\mathbf{e}^{(0)}, \mathbf{A}^{(0)}).$$

Therefore $b^{(0)}$ majorizes ℓ at $(\mathbf{e}^{(0)}, \mathbf{A}^{(0)})$ up to an additive constant. We conclude that minimizing $b^{(0)}(\mathbf{e}, \mathbf{A})$ will create an MM algorithm, as explained by Hunter and Lange (2004), and taking the next estimate in the iterative algorithm to be the minimizer will guarantee that the algorithm possesses a descent property.

4.2 Minimizing the majorizer

For each j , $1 \leq j \leq m$, we wish to minimize

$$b_j^{(0)}(e_j, \mathbf{A}_j) = - \int g(\mathbf{x}) w_j^{(0)}(\mathbf{x}) \cdot \log (\mathcal{N}_h e_j)_{\mathbf{A}_j}(\mathbf{x}) \, d\mathbf{x} + \int (e_j)_{\mathbf{A}_j}(\mathbf{x}) \, d\mathbf{x} \quad (13)$$

with respect to e_j and A_j . Instead of finding a global minimizer for $b_j^{(0)}$, we first hold A_j fixed and minimize with respect to e_j , then plug in the resulting update to e_j and minimize with respect to A_j . The resulting algorithm, which mimics the multiple “conditional maximization” steps of the ECM algorithm (Meng and Rubin 1993), does not actually minimize $b_j^{(0)}$, but it does ensure that the next iteration achieves a smaller value of $b_j^{(0)}$. This property is enough to guarantee the descent property, which states that the value of the objective function decreases at each iteration of the algorithm.

We find that Eq. (13) has a closed-form minimizer as a function of e_j when A_j is held fixed.

Proposition 1 *The minimizer of Eq. (13) with respect to e_j , with A_j held fixed, is*

$$\hat{e}_j(\mathbf{u}) = \frac{|\det A_j|}{\left[\int g(A_j \mathbf{x}) w_j^{(0)}(A_j \mathbf{x}) d\mathbf{x} \right]^{r-1}} \cdot \prod_{k=1}^r \int g(A_j \mathbf{y}) w_j^{(0)}(A_j \mathbf{y}) s_h(u_k, y_k) d\mathbf{y}. \quad (14)$$

A proof of Proposition 1 is provided in “Appendix A”.

Equation (14) can be rewritten as

$$\hat{e}_j(\mathbf{u}) = \left[P \circ S_h(|\det A_j| \cdot (g \cdot w_j^{(0)}) \circ A_j) \right](\mathbf{u})$$

using the P operator of Eq. (10). In general, for any nonnegative function f on R^r ,

$$S_h(f \circ A_j) = (S_h)_{A_j}(f) \circ A_j,$$

where

$$(S_h)_{A_j} f(\mathbf{x}) = \int |\det A_j|^{-1} \tilde{s}_h(A_j^{-1} \mathbf{x}, A_j^{-1} \mathbf{u}) f(\mathbf{u}) d\mathbf{u}.$$

Thus, we may also write

$$(\hat{e}_j)_{A_j}(\mathbf{u}) = \left[P_{A_j} \circ (S_h)_{A_j}(g \cdot w_j^{(0)}) \right](\mathbf{u}),$$

where

$$P_{A_j} f(\mathbf{u}) = [P(f_{A_j^{-1}})]_{A_j}(\mathbf{u}) = [P(f \circ A_j)](A_j^{-1} \mathbf{u}).$$

Now let us turn to the minimization of Eq. (13) with respect to A_j . We first define

$$\hat{q}_{jk}(u_k) = \frac{|\det A_j|}{\int g(A_j \mathbf{x}) w_j^{(0)}(A_j \mathbf{x}) d\mathbf{x}} \int g(A_j \mathbf{y}) w_j^{(0)}(A_j \mathbf{y}) s_h(u_k, y_k) d\mathbf{y}. \quad (15)$$

If we apply the change of variables $\mathbf{x} = A_j \mathbf{y}$ to Eq. (13) and then plug in $\hat{e}_j(\mathbf{u})$ into the resulting expression for $b_j^{(0)}(e_j, A_j)$, we find that minimizing the result with respect to A_j is equivalent to minimizing

$$\log |\det A_j| + \sum_{k=1}^r \int \hat{q}_{jk}(u) \log \hat{q}_{jk}(u) du \tag{16}$$

with respect to A_j , where \hat{q}_{jk} depends on A_j through (15).

In Expression (16), \hat{q}_{jk} is the k th margin of the kernel smoothed version of $(g \cdot w_j^{(0)})_{A_j^{-1}} / \int g \cdot w_j^{(0)}$. In the discrete case where $dG(\mathbf{x})$ is the empirical distribution, \hat{q}_{jk} is the k th margin of the kernel density estimate based on the linearly transformed (by A_j^{-1}) weighted observed data set, where the weight for the data point \mathbf{x}_i is $w_j^{(0)}(\mathbf{x}_i)$. Let us denote this weighted data set by $D_j^{(0)}$ and hence its linear transformation by $A_j^{-1}D_j^{(0)}$. By (15), the optimization mechanism at the current step views $A_j^{-1}D_j^{(0)}$ as a weighted sample generated from the unknown density function q_j , where $D_j^{(0)}$ is a weighted sample from the j th mixing component and A_j^{-1} is the matrix that recovers the associated ICA transformations. Let us call A_j^{-1} a recovering matrix.

By Eq. (6), we may treat $|\det A_j|$ as fixed given the weighted data $A_j^{-1}D_j^{(0)}$. The second term in (16) is an estimate of the sum of marginal entropies of q_j , which is equal, up to a term that does not involve A_j , to the mutual information of marginals of q_j . According to Hyvarinen et al. (2002), minimizing mutual information in this setting—that is, minimizing the mutual information of $A_j^{-1}\mathbf{S}$ given a randomly chosen weighted sample from \mathbf{S} —can be achieved by existing ICA algorithms such as the fastICA algorithm described in Sect. 4.3.

To summarize, the NSMM-ICA iterative algorithm will iterate as follows, where the parameters at the t th iteration will be denoted by $(\mathbf{e}^{(t)}, \mathbf{A}^{(t)})$:

Majorization Step: For $1 \leq j \leq m$, compute

$$w_j^{(t)}(\mathbf{x}) = \frac{(\mathcal{N}_h e_j^{(t)})_{A_j^{(t)}}(\mathbf{x})}{\sum_{j=1}^m (\mathcal{N}_h e_j^{(t)})_{A_j^{(t)}}(\mathbf{x})}.$$

ICA Step: Use the fastICA technique of Sect. 4.3 to find $A_j^{(t+1)}$ subject to (6) that minimizes

$$\sum_{k=1}^r \int \hat{q}_{jk}^{(t+1)}(u) \log \hat{q}_{jk}^{(t+1)}(u) du,$$

for $j = 1, \dots, m$, where $\hat{q}_{jk}^{(t+1)}(u_k)$ is defined in Eq. (15).

Minimization Step: Let

$$e_j^{(t+1)}(\mathbf{u}) = \hat{\lambda}_j^{(t+1)} \hat{q}_j^{(t+1)}(\mathbf{u}) = \hat{\lambda}_j^{(t+1)} \prod_{k=1}^r \hat{q}_{jk}^{(t+1)}(u_k),$$

where

$$\hat{\lambda}_j^{(t+1)} = \int (g \cdot w_j^{(t)}).$$

4.3 Practical implementation of NSMM-ICA

Section 4.2 suggests alternating NSMM and ICA methods to form an iterative algorithm for the estimation of the nonparametric ICA mixture model. This section describes the practical considerations that went into the development of a package for R (R Core Team 2015), called **icamix**, that implements these ideas.

Empirical evidence suggests that NSMM and the npEM algorithm of Benaglia et al. (2009) tend to give very similar estimates (Levine et al. 2011). The reason is that usually $\mathcal{N}_h f$ is close to f itself. This suggests that the smoothed version of the algorithm can reasonably be replaced by the non-smoothed version because the former is more computationally burdensome than the latter. The decision to implement this non-smoothed version affects only Step 1 of the algorithm below. The result is an algorithm that fails to achieve the provable descent property of the smoothed version but which is much faster and which appears to result in nearly identical results for most test problems.

Among the many ICA techniques available in the literature, here we use the efficient and well-tested FastICA of Hyvarinen et al. (2002). At each iteration, FastICA will be applied to a weighted dataset, where the weight on observation i for component j is determined as the estimate, given the information available at the present iteration, of the probability that observation i falls into component j .

Assume we are given raw data as a matrix $\mathbf{X}^\top = \{\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n\}^\top$, where $\mathbf{x}_i = (x_{i1}, x_{i2}, x_{i3}, \dots, x_{ir})^\top$ for $1 \leq i \leq n$. We first choose a set of starting parameter values. Since our algorithm, like any MM algorithm, finds at best a local minimum, it is possible that different starting values will lead to different solutions. In the **icamix** package, we begin with a k -means clustering solution, which assigns each data point to a distinct mixture component (and which itself has a stochastic element that allows for ease in choosing multiple starting points). Given this initial partition, straight-forward estimation (e.g., using FastICA) on the separate components leads to initial parameter values. Then, our algorithm iterates through Steps 1 through 4 below until a convergence criterion is met.

Step 1. For all i and j , estimate the j th component weight for the i th observation using the non-smoothed densities:

$$p_{ij}^{(t)} = \frac{\lambda_j^{(t)} f_j^{(t)}(\mathbf{x}_i)}{\sum_{j'=1}^m \lambda_{j'}^{(t)} f_{j'}^{(t)}(\mathbf{x}_i)} = \frac{\lambda_j^{(t)} |\det \mathbf{A}_j^{(t)}|^{-1} \prod_{k=1}^r q_{jk}^{(t)} \left(\left[(\mathbf{A}_j^{(t)})^{-1} \mathbf{x}_i \right]_k \right)}{\sum_{j'=1}^m \lambda_{j'}^{(t)} |\det \mathbf{A}_{j'}^{(t)}|^{-1} \prod_{k=1}^r q_{j',k}^{(t)} \left(\left[(\mathbf{A}_{j'}^{(t)})^{-1} \mathbf{x}_i \right]_k \right)}.$$

Steps 2, 3a, 3b, 3c, and 4 are now repeated separately for each value of j , $1 \leq j \leq m$:

Step 2. Update the λ parameters:

$$\lambda_j^{(t+1)} = \frac{1}{n} \sum_{i=1}^n p_{ij}^{(t)}.$$

Step 3a. Centering FastICA step for component j : For each i , define

$$\tilde{\mathbf{x}}_i \leftarrow \mathbf{x}_i - \frac{\sum_{i=1}^n \mathbf{x}_i p_{ij}^{(t)}}{\sum_{i=1}^n p_{ij}^{(t)}}.$$

Step 3b. Decorrelating FastICA step for component j : We first obtain the eigenvalue decomposition as

$$\frac{\sum_{i=1}^n \tilde{\mathbf{x}}_i \tilde{\mathbf{x}}_i^\top p_{ij}^{(t)}}{\sum_{i=1}^n p_{ij}^{(t)}} = E_j D_j E_j^\top,$$

then let $V_j = E_j D_j^{-1/2} E_j^\top$ and $\mathbf{z}_{ij} = V_j \tilde{\mathbf{x}}_i$ for $i = 1, \dots, n$. Here, the notation \mathbf{z}_{ij} refers to the j th component version of the i th observation of the \mathbf{z} vector, which is r -dimensional. Therefore,

$$\frac{\sum_{i=1}^n \mathbf{z}_{ij} \mathbf{z}_{ij}^\top p_{ij}^{(t)}}{\sum_{i=1}^n p_{ij}^{(t)}} = V_j E_j D_j E_j^\top V_j^\top = I. \tag{17}$$

The transformed data \mathbf{z}_{ij} with weights $p_{ij}^{(t)}$, $1 \leq i \leq n$, thus have their coordinates uncorrelated and standardized according to (17). Since $\mathbf{Z}_j = V_j \tilde{\mathbf{X}} = V_j \mathbf{A}_j \mathbf{S}$, we need to first estimate $(V_j \mathbf{A}_j)^{-1}$, of which the i th row is the same as \mathbf{w}_{ij} in (18) below, and multiply it by V_j on the right to get an update of \mathbf{A}_j^{-1} .

Step 3c. Symmetric orthogonalization FastICA step for component j : At this step, we enter an internal loop that ultimately results in the update of the \mathbf{A}_j matrix. Let $W_j = [\mathbf{w}_{1j}, \mathbf{w}_{2j}, \dots, \mathbf{w}_{rj}]^\top$ for r -dimensional unit length vectors $\mathbf{w}_{1j}, \mathbf{w}_{2j}, \dots, \mathbf{w}_{rj}$. The first time we enter Step 3c, we may simply take \mathbf{w}_{ij} to be the i th standard basis vector for each j , and at succeeding iterations we take these \mathbf{w}_{ij} to be rows of the most recent W_j matrix.

The inner loop then proceeds by updating the \mathbf{w}_{ij} from their previous values according to

$$\mathbf{w}_{ij} \leftarrow \frac{\sum_{i=1}^n \mathbf{z}_{ij} g(\mathbf{w}_{ij}^\top \mathbf{z}_{ij}) p_{ij}^{(t)}}{\sum_{i=1}^n p_{ij}^{(t)}} - \mathbf{w}_{ij} \frac{\sum_{i=1}^n g'(\mathbf{w}_{ij}^\top \mathbf{z}_{ij}) p_{ij}^{(t)}}{\sum_{i=1}^n p_{ij}^{(t)}}, \tag{18}$$

where g may be chosen to be either $g(y) = \tanh(\alpha_1 y)$ for some $1 \leq \alpha_1 \leq 2$ or $g(y) = y \exp(-y^2/2)$ (Hyvarinen et al. 2002). Let $W_j = [\mathbf{w}_{1j}, \mathbf{w}_{2j}, \dots, \mathbf{w}_{rj}]^\top$ and then symmetrize and orthogonalize by

$$W_j \leftarrow (W_j W_j^\top)^{-1/2} W_j. \tag{19}$$

Iteratively update the \mathbf{w}_{ij} , $i = 1, \dots, r$ using Eqs. (18) and (19) until convergence is achieved. More precisely, we choose a tolerance τ and stop updating when

$$\max_{1 \leq i \leq r} \left\{ \left| \left(\mathbf{w}_{ij}^{(\text{previous})} \right)^\top \cdot \mathbf{w}_{ij}^{(\text{current})} - 1 \right| \right\} \leq \tau.$$

Finally, set

$$\mathbf{A}_j^{(t+1)} = \mathbf{V}_j^{-1} \mathbf{W}_j^{-1}.$$

Step 4. Non-parametric density estimation step: For all j and k , let

$$q_{jk}^{(t+1)}(\mathbf{u}) = \left(\sum_{i=1}^n p_{ij}^{(t+1)} \right)^{-1} \sum_{i=1}^n p_{ij}^{(t+1)} \frac{1}{h} K \left(\frac{\mathbf{u} - [(\mathbf{A}_j^{(t+1)})^{-1} \mathbf{x}_i]_k}{h} \right).$$

The R package **icamix** makes use of the **Rcpp** (Eddelbuettel and François 2011; Eddelbuettel 2013) and **RcppArmadillo** (Eddelbuettel and Sanderson 2014) packages for compiling and calling the core algorithms implemented in C++ code to speed up the calculations.

In the discrete algorithm we have developed, a single fixed bandwidth calculated from the data will not be sensible, especially because the scale is now changing according to the ICA framework. Thus we propose an iterative scheme for choosing the bandwidth similar to that of Benaglia et al. (2011), whereby

$$h_{jk}^{t+1} = 0.5 \cdot \text{SD}_{jk}^{t+1} \cdot (n\lambda_j^{t+1})^{-0.2} = 0.5 \cdot (n\lambda_j^{t+1})^{-0.2}. \quad (20)$$

For simplicity, we replace $\min \{\text{SD}_{j,l}^{t+1}, \text{IQR}_{j,l}^{t+1}/1.349\}$ in the original Benaglia et al. (2011) formulation by $\text{SD}_{j,l}^{t+1} = 1$. We also propose the ad hoc coefficient of 0.5 rather than Silverman's 0.9 used by Benaglia et al. (2011) in order to capture fine features of the density for better performance in the classification task. Our experience is that using 0.9 tends to oversmooth the estimated density. Simulation studies and applications we have run suggest that Eq. (20) works well in practice.

When running the algorithm, we have determined that Step 4 dominates the computing time. By making use of Gaussian kernels and utilizing certain symmetric structure in evaluating some of Gaussians, we are able to lower the computing cost for the kernel density estimation step by about 50% with respect to the **mixtools** package. Further improvement might be possible via the Fast Gauss Transform (Raykar et al. 2005) and related techniques, though we have not implemented these improvements.

5 Applications

Here, we describe our experience applying the modified NSMM-ICA algorithm implemented in the **icamix** package to several datasets of varying size.

5.1 Italian wine classification

The Italian wine data set is another popular data set used for comparing various classifiers (Forina et al. 1988; Aeberhard et al. 1992). It contains results of a chemical analysis of wines grown in Italy but derived from three different cultivars. A total of 178 observations are recorded, each with 13 continuous attributes such as color

Table 1 Wine data

classifications by
PCA + NSMM-ICA algorithm

	Class 1	Class 2	Class 3
Barolo	59	0	0
Grignolino	6	61	4
Barbera	0	0	48




Fig. 2 Wine data: comparison of true species information (far left) and two results from our unsupervised learning algorithms

intensity, magnesium and malic acid. There are 59, 71 and 48 instances in the first (Barolo), second (Grignolino) and third (Barbera) wine classes, respectively (Table 1).

If we feed the unlabeled data directly to the NSMM-ICA algorithm, we obtain a classification error rate equal to 28.65%, prompting us to consider remedies. It seems that given the small number of observations, the relatively large number of attributes may be somewhat challenging as there are too many parameters to estimate and some of the attributes may consist of noise. So instead of using all 13 attributes, we first run principal component analysis (PCA) on the attributes and then select the 5 PCA scores that explain the largest proportion of variance in the attributes. Finally, we run the NSMM-ICA algorithm on the data set with the chosen PCA scores as attributes. In this way, the classification performance improves quite a lot, giving a classification error rate equal to 5.62%. Hence, in situations with relatively large numbers of coordinates, it might be worthwhile to utilize a dimension reduction technique followed by the NSMM-ICA algorithm. Figure 2 shows a comparison of true species information and results from our unsupervised learning algorithms.

5.2 Tone data

The tone perception experiment and data were first introduced by Cohen (1984). These data have been analyzed by Veaux (1989), Viele and Tong (2002), and Hunter and Young (2012) in the context of mixtures of regressions. In each trial of the experiment, a musician is presented with a fundamental tone plus a series of overtones determined by a stretching ratio. Then the musician is asked to tune an adjustable tone to one octave above the fundamental tone. Both the stretching ratio and the ratio of the adjusted tone to the fundamental are reported for each trial. There are five musicians involved in the experiment. However, the tone data set only contains 150 trials with the




Fig. 3 Comparison of three algorithms for fitting the tone dataset of Cohen (1984). Only the plot labeled SP EM explicitly assumes a mixture of regressions. The npEM algorithm does not utilize ICA and therefore misses the two lines entirely

Table 2 Comparison of mixtures of least squares fits for the tone dataset of Cohen (1984)

	SP EM	NSMM-ICA/weighted LS
Component 1		
$\hat{\beta}_0$	1.77533	1.82215
$\hat{\beta}_1$	0.11954	0.09076
Component 2		
$\hat{\beta}_0$	0.02121	-0.12111
$\hat{\beta}_1$	0.97929	1.05584
$\hat{\lambda}_1$	0.67653	0.46779

same musician. The problem of interest in conducting this experiment is to investigate the theory that the musician would either tune the tones to the nominal octave at a ratio of 2:1 to the fundamental tone (i.e., the interval memory hypothesis) or use the overtone to tune the tone to the stretching ratio (i.e., the partial matching hypothesis). The findings by Hunter and Young (2012) via modeling through a semi-parametric mixture of regressions conforms with the latter theory.

For this unsupervised learning task, we run both the npEM algorithm by Benaglia et al. (2009) and our NSMM-ICA algorithm on the tone data. Figure 3 shows that our NSMM-ICA algorithm does a good job of classification, very close to the mixture-of-regressions results obtained by Hunter and Young (2012), despite omitting any explicit assumption of regression structure. The reason why the results of the npEM algorithm shown in Fig. 3 do not look nearly as good is because with regression lines that have nonzero slopes the mixture is far from being conditionally independent. Thus, the additional ICA step in our algorithm is essential.

For comparing our result with that of Hunter and Young (2012), we can use the estimated mixing weights obtained from NSMM-ICA to fit a weighted ordinary least squares model to obtain the regression coefficients. The results, summarized in Table 2, reflect the difference in estimated membership between SP EM and our NSMM-ICA primarily at the intersection of the two components, which is responsible for the noticeable difference in the estimated mixing weights.

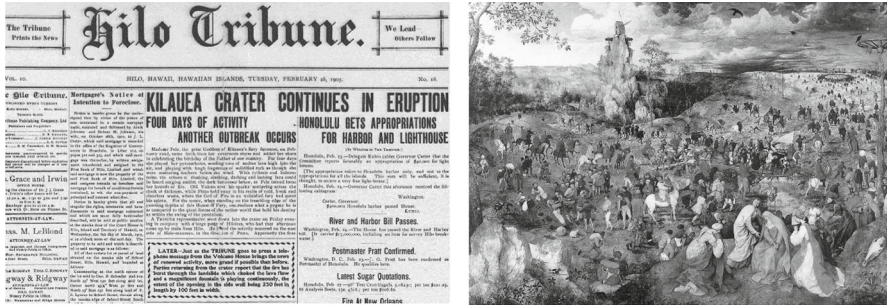


Fig. 4 Two images as sources for data: newspaper and painting

5.3 Clustering images

Learning efficient codes for images obtained from different sources or contexts is an important problem in the area of image processing. The task involves extracting intrinsic structure in images by clustering and finding a complete set of efficient linear basis functions for each image source, which results in coefficient values being as statistically independent as possible. Techniques that utilize a parametric form of ICA mixture models have been proposed in Lee et al. (1999a; 1999b; 2000). Here, we apply the NSMM-ICA algorithm, which eliminates the parametric assumptions, to the task of unsupervised learning of image codes.

The two images shown in Fig. 4 will be used as sources for the data set of the application. One is a painting image (2508×1808 pixels) and the other is a newspaper image (2057×1365 pixels). The images are transformed to grey scale: Each pixel consists of a pixel intensity value ranging from 0 (black) to 1 (white).

We select 5000 12×12 pixel patches randomly from each image. So the complete data set is of dimension $10,000 \times 144$. The NSMM-ICA algorithm converges after 19 iterations, which lasts a little less than 8 hours. Again each bandwidth is automatically learned by the iterative scheme we implemented. The result shows a very good recovery of the class-membership information, with a classification error rate of 1.2%.

The learned basis functions (i.e., a basis for the linear space of the pixel patches) show interesting patterns. Figures 5 and 6 show the basis functions for each image. The ones for the painting image appear smoother but more irregular, while the ones for the newspaper image look spottier and more regular.

6 Discussion

Although the conditional independence model for multivariate data is promising due to the fact that its parameters are provably identifiable under weak conditions, even some simple datasets such as the well-known iris dataset clearly violate this conditional independence assumption. The current article loosens this assumption, positing instead that each multivariate component can be linearly transformed to having independent coordinates. This gives a much more flexible model-based clustering framework than




Fig. 5 Learned basis functions for the newspaper image of Fig. 4. Each basis function, which is a 144-dimensional vector, is standardized to be within 0 and 1, then displayed as a 12×12 image patch

the conditional independence assumption alone. Among the particular favorable features of this extended model is that it allows for linear feature extraction, as illustrated by the tone data application of Sect. 5.2.

Yet much work remains to be done on this and similar methods. Despite the favorable-looking results that we have obtained on the datasets in this article, the theoretically important question of the identifiability of the parameters remains unresolved. A related issue is the fact that ICA cannot operate in the setting of multivariate normal data, since in that case the standardized data are already standard multivariate normal, so that there is no way to identify an ICA transformation. Thus, with the increased flexibility of our current framework come additional questions regarding conditions under which parameter identifiability holds.

In addition, large-sample behavior of the NSMM estimator such as convergence rates is still not fully known; however, recent work on alternative algorithms such as those of Bonhomme et al. (2016a), for which asymptotic properties have been



Fig. 6 Learned basis functions for the painting image of Fig. 4. Each basis function, which is a 144-dimensional vector, is standardized to be within 0 and 1, then displayed as a 12×12 image patch

proven, might provide suitable alternatives for estimation of the nonparametric mixture structure. Similarly, alternatives to the ICA algorithm we employ here, such as ICS, are also possible, and it is possible that exploiting the ability of ICS to search for interesting clustering features as in Peña et al. (2010) could be exploited.

In addition, there is the question of computational efficiency given the burden of estimating multiple univariate densities, particularly for large datasets. Our implementation of the algorithm currently replaces the smoothed NSMM portion of the algorithm by the non-smoothed npEM of Benaglia et al. (2009), since empirical comparisons have suggested that these two algorithms often result in nearly the same solutions. The **icamix** package for R (R Core Team 2015), available on CRAN, interweaves npEM with a weighted version of the FastICA algorithms (Hyvarinen et al. 2002). The package also implements an automated and adaptive scheme for bandwidth selection that is based on the work of Benaglia et al. (2011). Further computing efficiencies may be attainable through the use of the Fast Gauss Transform.

It is important to remember that although this article compares some clustering solutions obtained via our algorithm with those obtained using other techniques, such as k-means clustering, model-based clustering yields much more than mere cluster memberships: Not only does it assign each data point a probability vector of component membership, but the statistical modeling of the individual components is often of great interest beyond the assignment of individuals to groups. We did not explore the component density estimates obtained via our algorithm in this article, but for an example that does so in the nonparametric mixture modeling literature, consider the water-level dataset as analyzed in Section 5.2 of Chauveau et al. (2015).

All in all, given its flexibility and hence wide applicability, we believe that the novel approach to model-based clustering presented here has the potential to be a useful alternative to existing approaches based on parametric mixtures or mixtures that assume conditional independence.

A Proof of Proposition 1

We assume that for each $1 \leq j \leq m$, there exists $\theta_j > 0$ such that

$$e_j(\mathbf{x}) = \theta_j \prod_{k=1}^r e_{jk}(x_k), \quad (21)$$

where for each k , $1 \leq k \leq r$, $e_{jk} \in L^1(R)$ is positive. This overparameterization is employed for the sake of convenience and does not influence identifiability because we will never estimate θ_j separately.

The change of variables $\mathbf{x} = \mathbf{A}_j \mathbf{y}$ transforms Eq. (13) into

$$\begin{aligned} b_j^{(0)}(e_j, \mathbf{A}_j) &= - \int g(\mathbf{A}_j \mathbf{y}) w_j^{(0)}(\mathbf{A}_j \mathbf{y}) \cdot \log \left\{ (\mathcal{N}_h e_j)(\mathbf{y}) |\det \mathbf{A}_j|^{-1} \right\} |\det \mathbf{A}_j| \, d\mathbf{y} \\ &\quad + \int e_j(\mathbf{y}) \, d\mathbf{y}. \end{aligned} \quad (22)$$

Ignoring the term involving $\log |\det \mathbf{A}_j|^{-1}$ since it does not involve e_j , we find that minimizing $b_j^{(0)}(e_j, \mathbf{A}_j)$ as a function of e_j with \mathbf{A}_j fixed is equivalent to minimizing

$$- \int g(\mathbf{A}_j \mathbf{y}) w_j^{(0)}(\mathbf{A}_j \mathbf{y}) \int \bar{s}_h(\mathbf{u}, \mathbf{y}) \log e_j(\mathbf{u}) |\det \mathbf{A}_j| \, d\mathbf{u} \, d\mathbf{y} + \int e_j(\mathbf{u}) \, d\mathbf{u}, \quad (23)$$

which by Eqs. (8), (9), and (21) equals

$$\begin{aligned} &- \sum_{k=1}^r \iint g(\mathbf{A}_j \mathbf{y}) w_j^{(0)}(\mathbf{A}_j \mathbf{y}) s_h(u_k, y_k) \log e_{jk}(u_k) |\det \mathbf{A}_j| \, du_k \, d\mathbf{y} \\ &\quad + \theta_j \int \prod_{k=1}^r e_{jk}(u_k) \, du_k \end{aligned} \quad (24)$$

plus a term involving $\log \theta_j$ but none of the e_{jk} .

Picking a specific k and viewing Expression (24) as an integral with respect to du_k , we minimize it by minimizing the value of its integrand at each point. Differentiating with respect to $e_{jk}(u_k)$ and setting it equal to zero gives

$$-\frac{\int g(\mathbf{A}_j \mathbf{y}) w_j^{(0)}(\mathbf{A}_j \mathbf{y}) s_h(u_k, y_k) |\det \mathbf{A}_j| \, d\mathbf{y}}{e_{jk}(u_k)} + \theta_j \left[\prod_{l \neq k} \int e_{jk}(u_l) \, du_l \right] = 0,$$

yielding

$$\hat{e}_{jk}(u_k) \propto \int g(\mathbf{A}_j \mathbf{y}) w_j^{(0)}(\mathbf{A}_j \mathbf{y}) s_h(u_k, y_k) \, d\mathbf{y},$$

which implies

$$\hat{e}_j(\mathbf{u}) = \alpha_j \prod_{k=1}^r \int g(\mathbf{A}_j \mathbf{y}) w_j^{(0)}(\mathbf{A}_j \mathbf{y}) s_h(u_k, y_k) \, d\mathbf{y} \quad (25)$$

for some constant α_j . To find α_j , we plug (25) into (23), differentiate with respect to α_j , and set the result equal to zero to obtain

$$\alpha_j = \frac{|\det \mathbf{A}_j|}{\left[\int g(\mathbf{A}_j \mathbf{x}) w_j^{(0)}(\mathbf{A}_j \mathbf{x}) \, d\mathbf{x} \right]^{r-1}},$$

which implies Eq. (14).

References

- Aeberhard S, Coomans D, De Vel O (1992) Comparison of classifiers in high dimensional settings. Department of Mathematics and Statistics, James Cook University, North Queensland, Australia. Technical Report 92(02)
- Allman ES, Matias C, Rhodes JA (2009) Identifiability of parameters in latent structure models with many observed variables. *Ann Stat* 37(6A):3099–3132
- Anandkumar A, Hsu D, Kakade SM (2012) A method of moments for mixture models and hidden Markov models. In: Mannor S, Srebro N, Williamson RC (eds) Proceedings of the 25th annual conference on learning theory, vol 23, pp 33.1–33.34. PMLR, Edinburgh, Scotland
- Azzalini A, Torelli N (2007) Clustering via nonparametric density estimation. *Stat Comput* 17(1):71–80
- Bajari P, Hahn J, Hong H, Ridder G (2011) A note on semiparametric estimation of finite mixtures of discrete choice models with application to game theoretic models. *Int Econ Rev* 52(3):807–824
- Banfield JD, Raftery AE (1993) Model-based Gaussian and non-Gaussian clustering. *Biometrics* 49(3):803–821
- Benaglia T, Chauveau D, Hunter DR (2009) An EM-like algorithm for semi-and nonparametric estimation in multivariate mixtures. *J Comput Graph Stat* 18(2):505–526
- Benaglia T, Chauveau D, Hunter DR (2011) Bandwidth selection in an EM-like algorithm for nonparametric multivariate mixtures. In: Hunter DR, Richards DSP, Rosenberger JL (eds) Nonparametric statistics and mixture models: a festschrift in honor of Thomas P. Hettmansperger. World Scientific, Singapore, pp 15–27
- Bonhomme S, Jochmans K, Robin J-M (2016a) Estimating multivariate latent-structure models. *Ann Stat* 44(2):540–563

- Bonhomme S, Jochmans K, Robin J-M (2016b) Non-parametric estimation of finite mixtures from repeated measurements. *J R Stat Soc Ser B (Stat Methodol)* 78(1):211–229
- Butucea C, Vandekerckhove P (2014) Semiparametric mixtures of symmetric distributions. *Scand J Stat* 41(1):227–239
- Chauveau D, Hunter DR, Levine M (2015) Semi-parametric estimation for conditional independence multivariate finite mixture models. *Stat Surv* 9:1–31
- Cohen EA (1984) Some effects of inharmonic partials on interval perception. *Music Percept Interdiscip J* 1(3):323–349
- De Castro Y, Gassiat E, Lacour C (2016) Minimax adaptive estimation of nonparametric hidden Markov models. *J Mach Learn Res* 17(1):3842–3884
- De Veaux RD (1989) Mixtures of linear regressions. *Comput Stat Data Anal* 8(3):227–245
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Ser B (Methodol)* 39(1):1–38
- Eddelbuettel D (2013) *Seamless R and C++ integration with Rcpp*. Springer, New York
- Eddelbuettel D, François R (2011) Rcpp: seamless R and C++ integration. *J Stat Softw* 40(8):1–18
- Eddelbuettel D, Sanderson C (2014) RcppArmadillo: accelerating R with high-performance C++ linear algebra. *Comput Stat Data Anal* 71:1054–1063
- Forina M, Leardi R, Armanino C, Lanteri S, Conti P, Princi P (1988) Parvus: an extendable package of programs for data exploration, classification and correlation. *J Chemometr* 4(2):191–193
- Fraley C, Raftery AE (1998) How many clusters? Which clustering method? Answers via model-based cluster analysis. *Comput J* 41(8):578–588
- Fraley C, Raftery AE (2002) Model-based clustering, discriminant analysis, and density estimation. *J Am Stat Assoc* 97(458):611–631
- Frühwirth-Schnatter S (2006) *Finite mixture and markov switching models*. Springer Science & Business Media, LLC, New York
- Gassiat E, Cleynen A, Robin S (2016) Inference in finite state space non parametric hidden Markov models and applications. *Stat Comput* 26(1):61–71
- Gassiat E, Rousseau J (2016) Nonparametric finite translation hidden Markov models and extensions. *Bernoulli* 22(1):193–212
- Guglielmi A, Ieva F, Paganoni AM, Ruggeri F, Soriano J (2014) Semiparametric Bayesian models for clustering and classification in the presence of unbalanced in-hospital survival. *J R Stat Soc Ser C (Appl Stat)* 63(1):25–46
- Hall P, Zhou X-H (2003) Nonparametric estimation of component distributions in a multivariate mixture. *Ann Stat* 31(1):201–224
- Han B, Davis LS (2006) Semi-parametric model-based clustering for DNA microarray data. In: 18th International conference on pattern recognition (ICPR'06), vol 3, pp 324–327
- Huang M, Li R, Wang S (2013) Nonparametric mixture of regression models. *J Am Stat Assoc* 108(503):929–941
- Hunter DR, Lange K (2004) A tutorial on MM algorithms. *Am Stat* 58(1):30–37
- Hunter DR, Young DS (2012) Semiparametric mixtures of regressions. *J Nonparametr Stat* 24(1):19–38
- Hyvarinen A, Karhunen J, Oja E (2002) Independent component analysis. *Stud Inform Control* 11(2):205–207
- Lee T-W, Lewicki MS, Sejnowski TJ (1999a) ICA mixture models for image processing. In: Sixth joint symposium on neural computation proceedings, pp 79–86
- Lee T-W, Lewicki MS, Sejnowski TJ (1999b) Unsupervised classification with non-Gaussian mixture models using ICA. In: Kearns MJ, Solla SA, Cohn DA (eds) *Advances in neural information processing systems*, vol 11. MIT Press, Cambridge, pp 508–514
- Lee T-W, Lewicki MS, Sejnowski TJ (2000) ICA mixture models for unsupervised classification of non-Gaussian classes and automatic context switching in blind signal separation. *IEEE Trans Pattern Anal Mach Intell* 22(10):1078–1089
- Levine M, Hunter DR, Chauveau D (2011) Maximum smoothed likelihood for multivariate mixtures. *Biometrika* 98(2):403–416
- Li J, Ray S, Lindsay BG (2007) A nonparametric statistical approach to clustering via mode identification. *J Mach Learn Res* 8(8):1687–1723
- Mallapragada PK, Jin R, Jain A (2010) Non-parametric mixture models for clustering. In: Hancock ER, Wilson RC, Windeatt T, Ulusoy I, Escolano F (eds) *Structural, syntactic, and statistical pattern recognition*. Springer, Berlin, pp 334–343

- McLachlan G, Peel D (2000) *Finite mixture models*. Wiley, New York
- Meng X-L, Rubin DB (1993) Maximum likelihood estimation via the ECM algorithm: a general framework. *Biometrika* 80(2):267–278
- Miettinen J, Taskinen S, Nordhausen K, Oja H (2015) Fourth moments and independent component analysis. *Stat Sci* 30(3):372–390
- Palmer JA, Makeig S, Kreutz-Delgado K, Rao BD (2008) Newton method for the ICA mixture model. In: *Proceedings of the 2008 IEEE international conference on acoustics, speech, and signal processing*, pp 1805–1808
- Peña D, Prieto FJ, Viladomat J (2010) Eigenvectors of a kurtosis matrix as interesting directions to reveal cluster structure. *J Multivar Anal* 101(9):1995–2007
- R Core Team (2015) *R: a language and environment for statistical computing*. R Foundation for Statistical Computing, Vienna, Austria
- Raykar VC, Yang C, Duraiswami R, Gumerov N (2005) Fast computation of sums of Gaussians in high dimensions. Technical report, University of Maryland
- Salazar A, Igual J, Safont G, Vergara L, Vidal A (2015) Image applications of agglomerative clustering using mixtures of non-Gaussian distributions. In: *Proceedings of the 2015 international conference on computational science and computational intelligence (CSCI)*, pp 459–463
- Salazar A, Vergara L, Serrano A, Igual J (2010) A general procedure for learning mixtures of independent component analyzers. *Pattern Recognit* 43(1):69–85
- Shah CA, Arora MK, Varshney PK (2004) Unsupervised classification of hyperspectral data: an ICA mixture model based approach. *Int J Remote Sens* 25(2):481–487
- Tyler DE, Critchley F, Dümbgen L, Oja H (2009) Invariant co-ordinate selection (with discussion). *J R Stat Soc Ser B (Stat Methodol)* 71(3):549–592
- Vandekerkhove P (2013) Estimation of a semiparametric mixture of regressions model. *J Nonparametr Stat* 25(1):181–208
- Vichi M (2008) Fitting semiparametric clustering models to dissimilarity data. *Adv Data Anal Classif* 2(2):121–161
- Viele K, Tong B (2002) Modeling with mixtures of linear regressions. *Stat Comput* 12(4):315–330
- Wolfe JH (1963) *Object cluster analysis of social areas*. Ph.D. thesis, University of California
- Zhang W, Fan J, Sun Y (2009) A semiparametric model for cluster data. *Ann Stat* 37(5A):2377–2408
- Zhu X, Hunter DR (2016) Theoretical grounding for estimation in conditional independence multivariate finite mixture models. *J Nonparametr Stat* 28(1):683–701