Tilburg University

A hierarchical mixture model for clustering three-way data sets

Vermunt, J.K.

Published in:
Computational Statistics and Data Analysis

Publication date:
2007

Document Version
Publisher's PDF, also known as Version of record

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):
Vermunt, J. K. (2007). A hierarchical mixture model for clustering three-way data sets. *Computational Statistics and Data Analysis*, 51(11), 5368-5376.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

A hierarchical mixture model for clustering three-way data sets

Jeroen K. Vermunt*

Department of Methodology and Statistics, Tilburg University, P.O. Box 90153, 5000 LE Tilburg, The Netherlands

Available online 24 August 2006

Abstract

Three-way data sets occur when various attributes are measured for a set of observational units in different situations. Examples are genotype by environment by attribute data obtained in a plant experiment, individual by time point by response data in a longitudinal study, and individual by brand by attribute data in a market research survey. Clustering observational units (genotypes/individuals) by means of a special type of the normal mixture model has been proposed. An implicit assumption of this approach is, however, that observational units are in the same cluster in all situations. An extension is presented that makes it possible to relax this assumption and that because of this may yield much simpler clustering solutions. The proposed extension—which includes the earlier model as a special case—is obtained by adapting the multilevel latent class model for categorical responses to the three-way situation, as well as to the situation in which responses include continuous variables. An efficient EM algorithm for parameter estimation by maximum likelihood is described and two empirical examples are provided.

© 2006 Elsevier B.V. All rights reserved.

Keywords: Clustering; Three-way data; Finite mixture model; Longitudinal data; EM algorithm; Multilevel latent class model

1. Introduction

An example of a three-way data set is data collected in plant experiments where various attributes are measured on genotypes grown in several environments. This would be a genotype by environment by attribute data set. [Basford and McLachlan \(1985\)](#) proposed a variant of the normal mixture model for the analysis of such three-way data, where the aim is to cluster genotypes by explicitly taking into account the information on attributes and environments simultaneously. This is achieved by a multivariate normal mixture model with cluster-, environment-, and attribute-specific means, and with non-zero cluster-specific covariances between attributes within environments. More recently, [Hunt and Basford \(1999, 2001\)](#) extended the approach to cases with categorical attributes and with not all attributes observed on all genotypes. [Meulders et al. \(2002\)](#) proposed a restricted latent class model for the analysis of three-way dichotomous attribute data.

Other examples of three-way data include longitudinal data on multiple response variables—person by time point by response data—or data from experiments in which individuals provide multiple ratings for multiple objects (products, brands) or report on possible behaviors shown in multiple situations, yielding person by object by attribute and person by situation by behavior data, respectively. Other examples consist of data sets in which objects are rated on multiple attributes by multiple experts, such as exams with multiple questions corrected by multiple raters or products evaluated on multiple attributes by multiple raters. In the remaining, I will refer to the three ways of the data sets as cases,

* Tel.: +31 13 4662748; fax: +31 13 4663002.

E-mail address: j.k.vermunt@uvt.nl.

situations, and attributes, respectively. The aim of the application of a mixture model is to cluster cases based on measured attributes in various situations. Clusters will also be referred to as (latent) classes and groups.

An important characteristic of the Basford and McLachlan (B&M) mixture model for three-way data, as well as of the other variants mentioned above, is that cases are assumed to belong to the same cluster in all investigated situations. I propose an alternative mixture model for three-way data that relaxes this assumption: cases may be in a different latent class depending on the situation or, more specifically, cases are clustered with respect to the probability of being in a particular latent class at a certain situation. The basic idea is to treat the three ways as hierarchically nested levels and assume that there is a mixture distribution at each of the two higher levels; i.e., one at the case and one at the case-in-situation level. The proposed model is an adaptation of the multilevel latent class model by Vermunt (2003) to continuous responses, as well as to the specific model structures needed for dealing with three-way data. A nice feature is that it has the B&M three-way mixture model as a special case.

An important advantage of the proposed modelling approach is that it may yield more parsimonious solutions—solutions with less clusters—with an even better description of the data than the B&M model. Moreover, interpretation of results may be easier and the model may be more in agreement with reality and thus more meaningful. For example, in a longitudinal data application it is unrealistic to assume that individuals are in the same latent class at each time point or in a multiple experts study it is unrealistic to assume that each expert classifies an object in the same latent class.

Böhning et al. (2000) proposed a state–space mixture model in which, in fact, two ways (case by time point) are collapsed into one way. A standard mixture model is subsequently adopted, which implies that observations of the same case at different time points are assumed to be independent of one another. An advantage of the hierarchical mixture model described below is that it can take into account dependencies between repeated observations within cases. It should be noted that the hierarchical mixture model has the state–space mixture model of Böhning et al. (2000) as a special case; that is, as the limiting case in which there is only one higher-level mixture component.

The remaining of this article is organized as follows. Using B&M's model as the starting point, I first describe the simplest form of the new model, and subsequently introduce variants such as restricted multivariate normals, models for categorical and mixed responses, and models with covariates and regression type constraints. Subsequently, I show how parameter estimation can be performed using a special variant of the EM algorithm which is implemented in the Latent GOLD mixture modelling software (Vermunt and Magidson, 2005). The new approach is illustrated with two empirical examples.

2. Mixture models for three-way data

2.1. Basford and McLachlan's mixture model

Following a similar notation as in McLachlan and Peel (2000, p. 114) and Hunt and Basford (2001), suppose that the responses on P attributes were recorded in N cases, each of which was observed in R situations. Let $\mathbf{y}_{ir}$ be a $P \times 1$ vector containing the values of the P attributes of case i in situation r , for $i = 1, \dots, N$; $r = 1, \dots, R$. The $RP \times 1$ observation vector $\mathbf{y}_i$ is given by

$$\mathbf{y}_i = (\mathbf{y}'_{i1}, \dots, \mathbf{y}'_{iR})',$$

where $\mathbf{y}_i$ contains the multi-attribute responses of the i th case in all R situations. Under the mixture model proposed by Basford and McLachlan (1985), it is assumed that cases belong to one of K possible groups or latent classes $G_1, \dots, G_K$ in proportions $\pi_1, \dots, \pi_K$, respectively, where $\sum \pi_k = 1$ and $\pi_k \geq 0$ for $k = 1, \dots, K$. The responses of case i in situation r have a multivariate normal distribution conditional on group G_k ; i.e., $\mathbf{y}_{ir} \sim N(\boldsymbol{\mu}_{kr}, \boldsymbol{\Sigma}_k)$. The mixture model for three-way data proposed by Basford and McLachlan (1985) has the following form:

$$f(\mathbf{y}_i) = \sum_{k=1}^K \pi_k \prod_{r=1}^R f_k(\mathbf{y}_{ir}; \boldsymbol{\mu}_{kr}, \boldsymbol{\Sigma}_k). \quad (1)$$

Note that the values of the within-class covariance matrices are constant across situations, whereas the class-specific attribute means differ across situations. An important assumption is that conditional on the class membership of case i the responses in the different situations are independent of one another. It is, however, impossible to relax that

assumption because too many covariances would have to be estimated in a more general model with free covariances across situations.

Another important assumption of the B&M model is that cases are in the same latent class in each of the investigated situations. The more extended model described in the next subsection relaxes this assumption.

2.2. The hierarchical mixture model

As under the model described in Eq. (1), under the hierarchical mixture model for three-way data, it is assumed that cases belong to one of K possible groups $G_1, \dots, G_K$ in proportions $\pi_1, \dots, \pi_K$, respectively, where $\sum \pi_k = 1$ and $\pi_k \geq 0$ for $k = 1, \dots, K$. A new element is that conditional on belonging to G_k , in situation r cases are assumed to belong to one of L groups $H_1, \dots, H_L$ in proportions $\theta_{1|k}, \dots, \theta_{L|k}$, respectively, where $\sum \theta_{\ell|k} = 1$ and $\theta_{\ell|k} \geq 0$ for $\ell = 1, \dots, L$ and $k = 1, \dots, K$, which yields a two-layer structure similar to the model proposed by Li (2005). The responses in situation r have a multivariate normal distribution conditional on group H_ℓ , i.e., $\mathbf{y}_{ir} \sim N(\boldsymbol{\mu}_{\ell r}, \boldsymbol{\Sigma}_\ell)$. As in the B&M model, the within-class covariance matrix $\boldsymbol{\Sigma}_\ell$ is constant across situations, whereas the class-specific means differ across situations. The hierarchical mixture model has the following form:

$$f(\mathbf{y}_i) = \sum_{k=1}^K \pi_k \prod_{r=1}^R \sum_{\ell=1}^L \theta_{\ell|k} f_\ell(\mathbf{y}_{ir}; \boldsymbol{\mu}_{\ell r}, \boldsymbol{\Sigma}_\ell). \quad (2)$$

It should be noted that this model is equivalent to the B&M model described in Eq. (1) if $L = K$ and if $\theta_{\ell|k}$ is restricted to be equal to 1 for $\ell = k$ and to 0 for $\ell \neq k$; that is, if cases belong to the same class in each situation. This shows that the hierarchical model extends the standard model by allowing cases to be in a different latent class per situation with a certain probability. Higher-level mixture components differ with respect to these prior class membership probabilities, which is captured by the $K(L - 1)$ extra model parameters $\theta_{\ell|k}$.

The model described in Eq. (2) is similar to the multilevel latent class model proposed by Vermunt (2003). An important difference is that this was a model for categorical rather than continuous responses, as well as that it could not deal with parameters that differ across situations. In one particular aspect the multilevel latent class model is more general than the model presented here; namely, in that the number of lower-level units may differ across higher-level units or, translated into the three-way terminology, that there is no need that all cases have been observed in the same (number of) situations.

In terms of structure the proposed model is also similar to hierarchical mixtures-of-experts models (Jordan and Jacobs, 1994; Vermunt and Magidson, 2003). An important difference is that the hierarchical mixtures-of-experts architecture is not used with three-way but with standard two-way data sets. Other differences are that in these models the parameters of the component distributions may also depend on G_k and that explanatory variables may enter in the various model parts. But as is shown below, similar types of extensions can be defined for the model proposed in this article.

The model described in Eq. (2) also shares some similarities with the hierarchical latent class model proposed by Zhang (2004) for the exploratory analysis of data sets with large numbers of response variables. This is a model for two-way data sets that allows for a hierarchy of latent variables with as many levels as needed to get a good description of the data set at hand. The EM algorithm used by Zhang is similar to the one presented in the next section.

2.3. Variants and extensions of the hierarchical model

Various variants and extensions of the above model can be defined. For instance, a more parsimonious variant is obtained by assuming that the class-specific means do not vary across situations, which involves replacing $\boldsymbol{\mu}_{\ell r}$ by $\boldsymbol{\mu}_\ell$. The fact that means were allowed to differ across situations was in fact specific to the type of application for which Basford and McLachlan (1985) developed their model. In other applications, it may be more natural to assume homogeneity across situations; for example, in longitudinal data applications, we will most likely not wish to allow class-specific means to differ across time points.

An intermediate variant in terms of parsimony is obtained by defining an analysis-of-variance type of linear model for $\boldsymbol{\mu}_{\ell r}$, with main effects for class and situation but without an interaction effect: $\boldsymbol{\mu}_{\ell r} = \boldsymbol{\beta}_0 + \boldsymbol{\beta}_\ell^H + \boldsymbol{\beta}_r^S$, where H stands for latent class and S for situation. What we are saying here is that means are situation specific, but the way

responses vary across situations is the same for all classes, a simplifying assumption that seems to make sense in many application types.

In a regression model for the class-specific means, we could also include other case- and situation-specific predictors. We could even include attribute-specific predictors, yielding a mixture regression model structure (Wedel and DeSarbo, 1994). Finally, the means could also be allowed to depend on G_k —the case-level classes.

Not only can the class-specific means be further restricted, but also the covariance matrices. Interesting constraints are homogeneity across classes, diagonal covariance matrices, and lower-dimensional representations using factor-analytic structures. In fact, all the restricted covariance structures that have been proposed for standard multivariate normal mixture models (see, e.g., McLachlan and Peel, 2000; Vermunt and Magidson, 2002) can be applied within the context of the proposed hierarchical mixture model for three-way data.

Rather than assuming that the attribute means depend on situation, we could also allow the probability of belonging to group H_ℓ given membership of group G_k to depend on situation, which involves replacing $\theta_{\ell|k}$ by $\theta_{\ell|kr}$. As for the means, to eliminate interaction terms, these probabilities could be restricted by means of a regression model, in this case by a multinomial logistic regression model containing only the main effects for case-level classes G_k and situations. In this regression model we could also include other case- and situation-specific predictors. Also the probability of belonging to group G_k can be allowed to depend on (case-specific) covariates. Note that the use of covariates yields models which are similar to the concomitant variable latent class model by Dayton and Macready (1988).

The last variant I would like to mention is relevant when there are categorical or mixed responses. As in standard mixture models, for categorical responses, we will typically use multinomial (Goodman, 1974) or Poisson (Böhning et al., 2000; Knorr-Held and Raßer, 2000) within-class distributions. Taking the more general case in which the class-specific densities can take on other forms than multivariate normal, the three-way mixture model is formulated as follows:

$$f(\mathbf{y}_i) = \sum_{k=1}^K \pi_k \prod_{r=1}^R \sum_{\ell=1}^L \theta_{\ell|k} f_{\ell}(\mathbf{y}_{ir}; \phi_{\ell r}),$$

where $f_{\ell}(\mathbf{y}_{ir}; \phi_{\ell r})$ is the density for situation r conditional on class H_ℓ , and $\phi_{\ell r}$ is the vector of unknown parameters defining this density. In the case of local independence, we will in addition assume that

$$f_{\ell}(\mathbf{y}_{ir}; \phi_{\ell r}) = \prod_{p=1}^P f_{\ell}(y_{irp}; \phi_{\ell rp});$$

that is, that the multi-attribute density can be obtained as a product of the univariate densities corresponding to the P attributes. A special case of the multilevel latent class model proposed by Vermunt (2003) is obtained when the P responses can be assumed to come from locally independent multinomial distributions, an example of which is presented below.

3. Parameter estimation by the EM algorithm

Let $\mathbf{z}_i = (z_{i1}, \dots, z_{iK})'$, for $i = 1, \dots, N$, be a vector of indicator variables, where z_{ik} equals 1 if case i belongs to group G_k and 0 otherwise, and let $\mathbf{w}_{ir} = (w_{ir1}, \dots, w_{irL})'$, for $i = 1, \dots, N$ and $r = 1, \dots, R$, be another vector of indicator variables, where $w_{ir\ell}$ equals 1 if case i belongs to group H_ℓ in situation r and 0 otherwise. The $\mathbf{z}_i$ are assumed to come from a multinomial distribution with parameters π_k , and, conditionally on the $\mathbf{z}_i$, the $\mathbf{w}_{ir}$ are assumed to come from a multinomial distribution with parameters $\theta_{\ell|k}$.

By treating these indicator variables as missing or unobserved, parameter estimation by maximum likelihood can be solved by means of the EM algorithm (Dempster et al., 1977). Because of the extremely high dimensionality of the missing data problem, in the implementation of the E step I use a similar trick as in the Baum–Welch algorithm for hidden Markov modelling (Baum et al., 1970). It should be noted that the model defined in Eq. (2) contains $1 + R$ unobserved variables with a total of $K \times L^R$ categories. This implies, for example, that with $R = 8$ and $K = L = 4$, we are dealing with a model with 262,144 entries in the joint distribution of the variables with missing values. It will be clear that this cannot be solved with a standard EM algorithm. In the graphical or Bayesian belief network modelling field, the hierarchical mixture model would be recognized as a single-connected network or polytree, for

which relevant marginal conditional probabilities can be obtained by propagation algorithms (Pearl, 1988). Both the forward–backward algorithm for hidden Markov models and the upward–downward algorithm discussed below are propagation algorithms.

Rather than repeating all the well-known details on the EM algorithm for the estimation of normal mixture models which can be found in, for example, McLachlan and Peel (2000), I will concentrate on the specific aspects associated with the estimation of the hierarchical mixture model described in Eq. (2). The complete data log-likelihood function for this model has the following form:

$$\begin{aligned} \log L_C(\boldsymbol{\phi}) = & \sum_{i=1}^N \sum_{k=1}^K z_{ik} \log \pi_k + \sum_{i=1}^N \sum_{k=1}^K \sum_{r=1}^R \sum_{\ell=1}^L z_{ik} w_{ir\ell} \log \theta_{\ell|k} \\ & + \sum_{i=1}^N \sum_{k=1}^K \sum_{r=1}^R \sum_{\ell=1}^L z_{ik} w_{ir\ell} \log f_{\ell}(\mathbf{y}_{ir}; \boldsymbol{\mu}_{\ell r}, \boldsymbol{\Sigma}_{\ell}), \end{aligned} \quad (3)$$

where $\boldsymbol{\phi}$ refers to the full set of unknown model parameters. Calculation of the expected value of the complete data log-likelihood—which is the E step of the EM algorithm—involves replacing the indicator variables z_{ik} and $w_{ir\ell}$ by their expected values $\hat{z}_{ik} = P(z_{ik} = 1 | \mathbf{y}_i; \boldsymbol{\phi})$ and $\hat{w}_{ir\ell|k} = P(w_{ir\ell} = 1 | \mathbf{y}_i, z_{ik} = 1; \boldsymbol{\phi})$, which are the estimated posterior probabilities that case i belongs to class G_k and that it belongs to class H_{ℓ} when it is in situation r given G_k , conditional on the observed data and the current parameter estimates. Note that $\hat{z}_{ik} \hat{w}_{ir\ell|k} = P(z_{ik} = 1, w_{ir\ell} = 1 | \mathbf{y}_i; \boldsymbol{\phi})$, which is the expected value of the product term $z_{ik} w_{ir\ell}$ appearing in Eq. (3).

Crucial in the implementation of the E step of the algorithm is that one can make use of the fact that lower-level (case-in-situation) observations are independent of one another given the higher-level (case) class memberships. More specifically, we make use of the fact that

$$\hat{w}_{ir\ell|k} = P(w_{ir\ell} = 1 | \mathbf{y}_i, z_{ik} = 1; \boldsymbol{\phi}) = P(w_{ir\ell} = 1 | \mathbf{y}_{ir}, z_{ik} = 1; \boldsymbol{\phi});$$

that is, that given class membership of the case (z_{ik}), class membership in a certain situation ($w_{ir\ell}$) is independent of the observed data at the other situations.

In order to simplify the formulas for $\hat{z}_{ik}$ and $\hat{w}_{ir\ell|k}$, let $h_{ir\ell|k} = \theta_{\ell|k} f_{\ell}(\mathbf{y}_{ir}; \boldsymbol{\mu}_{\ell r}, \boldsymbol{\Sigma}_{\ell})$ and $g_{ir|k} = \sum_{\ell=1}^L h_{ir\ell|k}$. The relevant terms are obtained as follows:

$$\begin{aligned} \hat{w}_{ir\ell|k} &= \frac{\theta_{\ell|k} f_{\ell}(\mathbf{y}_{ir}; \boldsymbol{\mu}_{\ell r}, \boldsymbol{\Sigma}_{\ell})}{\sum_{\ell=1}^L \theta_{\ell|k} f_{\ell}(\mathbf{y}_{ir}; \boldsymbol{\mu}_{\ell r}, \boldsymbol{\Sigma}_{\ell})} = \frac{h_{ir\ell|k}}{g_{ir|k}}, \\ \hat{z}_{ik} &= \frac{\pi_k \prod_{r=1}^R g_{ir|k}}{\sum_{k=1}^K \pi_k \prod_{r=1}^R g_{ir|k}}. \end{aligned}$$

As can be seen, for each case i , we first compute $h_{ir\ell|k}$ for each k, r , and ℓ combination and collapse these over ℓ to obtain $g_{ir|k}$, which amounts to marginalizing over the lower-level cluster variables. Combining the $g_{ir|k}$ for all r gives the posterior for the higher-level cluster variable. Analogous to the forward–backward recursion algorithm, Vermunt (2003) refers to this step as the upward step because information from the lower nodes of the tree is passed to the upper node. The downward step involves the computation of the bivariate joint posterior of z_{ik} and $w_{ir\ell}$, the term that enters in the expected complete data log-likelihood; that is,

$$P(z_{ik} = 1, w_{ir\ell} = 1 | \mathbf{y}_i; \boldsymbol{\phi}) = \hat{z}_{ik} \hat{w}_{ir\ell|k}.$$

The M step of the EM algorithm proceeds similarly as described by Basford and McLachlan (1985); i.e.,

$$\begin{aligned} \hat{\pi}_k &= \frac{\sum_{i=1}^N \hat{z}_{ik}}{\sum_{i=1}^N \sum_{k=1}^K \hat{z}_{ik}}, \\ \hat{\theta}_{\ell|k} &= \frac{\sum_{i=1}^N \sum_{r=1}^R \hat{z}_{ik} \hat{w}_{ir\ell|k}}{\sum_{i=1}^N \sum_{r=1}^R \sum_{\ell=1}^L \hat{z}_{ik} \hat{w}_{ir\ell|k}}, \end{aligned}$$

$$\hat{\mu}_{\ell r} = \frac{\sum_{i=1}^N \sum_{k=1}^K \hat{z}_{ik} \hat{w}_{ir\ell|k} \mathbf{y}_{ir}}{\sum_{i=1}^N \sum_{k=1}^K \hat{z}_{ik} \hat{w}_{ir\ell|k}},$$

$$\hat{\Sigma}_{\ell} = \frac{\sum_{i=1}^N \sum_{k=1}^K \sum_{r=1}^R \hat{z}_{ik} \hat{w}_{ir\ell|k} \left[(\mathbf{y}_{ir} - \hat{\mu}_{\ell r}) (\mathbf{y}_{ir} - \hat{\mu}_{\ell r})' \right]}{\sum_{i=1}^N \sum_{k=1}^K \sum_{r=1}^R \hat{z}_{ik} \hat{w}_{ir\ell|k}}.$$

These M step equations can easily be adapted to other distributions such as Poisson or multinomial distributions for discrete response variables.

The special variant of the EM algorithm described above has been implemented in version 4.0 of the Latent GOLD software package for latent class and mixture modelling (Vermunt and Magidson, 2005). Although not all specific structures for the class-specific means and covariance matrices that one may need for three-way data are in the current program, the new version will contain all relevant options.

An important issue in mixture modelling is identifiability (McLachlan and Peel, 2000, pp. 26–28). Apart from the label switching problem, as in standard mixture models, it is not straightforward to provide general conditions for identifiability. It can, however, easily be observed that the model described in Eq. (2) is, in fact, built up by two submodels: a latent class like model for the higher-level latent classes in which the R lower-level class memberships serve as categorical “response” variables, and a standard mixture model for the lower-level classes. A necessary condition for identification is that the upper part of the model has the structure of an identifiable latent class model. This requires, for example, that the number of situations should be at least three ($R \geq 3$) (Goodman, 1974). If the upper part is identifiable, (separate) identifiability of the lower part is a sufficient condition but not always necessary when $K > 1$. An example is the case in which the lower part is a standard latent class model for 2 response variables ($P = 2$). Whereas such model is not identified for $K = 1$, it is for $K > 1$ (of course, assuming that $R \geq 3$). This discussion shows that in the more typical applications, such as in the ones discussed below, identifiability is not more problematic for the hierarchical mixture model than for the standard mixture model. In practice, as I did in the examples presented below, one can check local identifiability by determining the rank of Jacobian (Goodman, 1974; Formann, 1992).

4. Two empirical examples

4.1. Soybean data

I will illustrate the new mixture model for three-way data using two empirical applications. The first one is a reanalysis of the classical soybean data used by Basford and McLachlan (1985) and McLachlan and Basford (1988) to illustrate their three-way normal mixture model. I obtained the data set from Pieter Kroonenberg’s website on three-way data analysis: <http://three-mode.leidenuniv.nl/>. The data originate from an experiment in which 58 soybean genotypes were evaluated at four locations in Queensland, Australia, at two time points, the eight combinations of which will be denoted as environments. Various attributes were measured on the genotypes. The two continuous attributes I used are “yield” and “protein percentage”.

I estimated multivariate mixture models of the forms (1) and (2) for K and L values ranging from 1 to 4. Note that for $K = 1$ the model of Eq. (2) reduces to a standard mixture that treats the observations of the same genotype at different environments as independent observations. Because combinations of $L = 1$ with $K > 1$ are not meaningful, these are omitted from the table. In the estimated models, the class- and environment-specific means were restricted by an ANOVA-like structure: $\mu_{\ell r} = \beta_0 + \beta_{\ell}^H + \beta_r^E$, where E refers to environment.

The encountered BIC values are reported in Table 1, where for the computation of BIC, I used 58, the number of genotypes, as the sample size. Conclusions would have been the same if model selection would have been based on AIC instead of BIC. As can be seen, for this data set, the B&M three-way mixture model performs much better than the hierarchical mixture model, which indicates that the assumption that genotypes are in the same class at each environment holds. Not surprisingly, in the hierarchical mixture models with $K = L$ the estimated values for the $\theta_{\ell|k}$ parameters were always close to 1 for $k = \ell$ and close to 0 otherwise, the values at which these parameters are fixed in B&M model.

To investigate whether this somewhat unexpected result is connected to this particular data set or whether it is an artifact of applying the proposed model with continuous response variables, I simulated 10 data sets with the same structure as the Soybean data set ($N = 58$, $R = 8$, and $P = 2$) using the estimated μ and Σ values of the model with

$K = 3$ and $L = 3$ as population values. I assumed $\pi_k = \frac{1}{3}$ for each k , and $\theta_{\ell|k} = 0.8$ for $\ell = k$ and $\theta_{\ell|k} = 0.1$ otherwise. The latter parameter values were very well recovered: across the 10 replications, I found an average estimate of 0.774 for the diagonal $\theta_{\ell|k}$. The fact that no single diagonal element was larger than 0.93 shows that boundary estimates are unlikely if such a model holds.

4.2. Anger data

The second application uses data from a psychological experiment described by Meulders et al. (2002) to illustrate another type of latent class model for three-way data. This data set is available at: <http://www.statisticalinnovations.com/>. It consists of the answers of 101 first-year psychology students who indicated whether or not they would show each of eight behaviors when angry at someone for six different situations. The eight behaviors consist of four pairs of reactions reflecting a particular way of dealing with anger: fighting [(1) fly off the handle, (2) quarrel], fleeing [(3) leave, (4) avoid], emotional sharing [(5) pour out one’s heart, (6) tell one’s story], and making up [(7) make up, (8) clear up the matter]. The six situations are whether one (1) likes, (2) dislikes or (3) is unfamiliar with the instigator of anger and whether the instigator has a (4) higher, (5) lower, or (6) equal status.

Because the reported behaviors come in four strongly overlapping pairs, it is not realistic to assume that the eight responses are locally independent (Goodman, 1974). Therefore, we allowed for local dependencies—or direct effects in the terminology of Hagenaaers (1988)—between pairs of behaviors connected with the same way of dealing with anger. More specifically, the joint distribution for a pair of 0/1 items—say the pair formed by y_{ir1} and y_{ir2} —within situations conditional on membership of group H_ℓ is multinomial: $(y_{ir1}, y_{ir2}) \sim \text{Mult}(\eta_{00\ell r}, \eta_{10\ell r}, \eta_{01\ell r}, \eta_{11\ell r})$. As in the B&M model, we allow responses to depend on situation, with constant effects across latent classes. This can be achieved by defining a logistic regression model for the item responses similar to the linear logistic latent class model by Formann (1992). The model for the log odds of the $(y_{ir1} = s, y_{ir2} = t)$ (for $s = 0, 1$ and $t = 0, 1$) versus the $(y_{ir1} = 0, y_{ir2} = 0)$ joint response is

$$\text{logit } \eta_{st\ell r} = \left(\beta_{01} + \beta_{\ell 1}^H + \beta_{r1}^S\right)^s + \left(\beta_{02} + \beta_{\ell 2}^H + \beta_{r2}^S\right)^t + (\beta_{012})^{st}.$$

The substantive interpretation of this specification is that certain reactions are more likely to occur in certain situations than in others, but that—on the logit scale—the amount by which the likelihood changes is equal across classes. The parameter β_{012} captures the within-cluster association between these two items.

Table 2 reports the BIC values obtained for the estimated models with the Anger data set. In the computation of BIC, I used 101, the number of students, as the sample size. Conclusions would have been the same if model selection would have been based on AIC instead of BIC. As can be seen, models that allow class membership to vary across situations

Table 1
BIC values for model estimated with the Soybean data set

	$K = 1$	$K = 2$	$K = 3$	$K = 4$	B&L
$L = 1$	3014	–	–	–	3014
$L = 2$	2999	2761	2768	2776	2752
$L = 3$	2996	2751	2690	2700	2665
$L = 4$	3004	2755	2698	2667	2618

Table 2
BIC values for models estimated with the Anger data set

	$K = 1$	$K = 2$	$K = 3$	$K = 4$	B&M
$L = 1$	5257	–	–	–	5257
$L = 2$	5119	5114	5118	5127	5217
$L = 3$	5127	5115	5117	5129	5209
$L = 4$	5142	5121	5111	5125	5208

Table 3

Estimated values for the π_k and $\theta_{\ell|k}$ parameters obtained with the model with $K = 3$ and $L = 4$

	G_1	G_2	G_3
π_k	0.65	0.23	0.12
H_1	0.71	0.02	0.06
H_2	0.20	0.11	0.68
H_3	0.10	0.45	0.00
H_4	0.00	0.42	0.26

Table 4

Estimated values for the (marginal) class-specific response probabilities obtained with the model with $K = 3$ and $L = 4$

	H_1	H_2	H_3	H_4
Y_1	0.39	0.02	0.80	0.29
Y_2	0.20	0.10	0.99	0.35
Y_3	0.16	0.92	0.00	0.39
Y_4	0.28	0.94	0.04	0.44
Y_5	0.58	0.56	0.58	0.42
Y_6	0.57	0.51	0.64	0.99
Y_7	0.46	0.30	0.55	0.36
Y_8	0.42	0.19	0.54	0.46

perform much better than the B&M specification with fixed class memberships. The model with the lowest BIC is the model with $K = 3$ and $L = 4$. Tables 3 and 4 report the parameter estimates obtained with this model.

The numbers in Table 3 show that G_1 —the largest class containing 65% of the cases—shows reaction type H_1 in most situations, but sometimes also types H_2 or H_3 . The second group selects types H_3 or H_4 in most situations, and class G_3 has preference for H_2 , but may also select H_4 . What these numbers indicate is that selecting a type of reaction given G_k is clearly a stochastic process and not deterministic as is assumed in the B&M three-way mixture model.

Table 4 provides the required information for labelling the types of reactions that one selects when angry at someone. Note that these are average response probabilities across situations and levels of the other variable in the locally dependent pair. Reaction types H_2 and H_3 are easiest to label; namely, fleeing and fighting. Classes H_1 and H_4 are similar, with the exception that the latter has a much higher probability for the second emotional sharing item and is also somewhat more likely to report the fleeing behaviors (Y_3 and Y_4). As far as the making up (Y_7 and Y_8) items is concerned, we do not see large differences across classes, except that class H_2 has somewhat lower probabilities for these reactions.

5. Discussion

A novel mixture clustering model was presented for the analysis of three-way data sets. The method—which is based on treating the three-way data as hierarchical data—is a variant of the multilevel latent class model proposed by Vermunt (2003). The proposed model is an extension of the model by Basford and McLachlan (1985) in the sense that it allows to relax the assumption that class membership does not change across situations.

The hierarchical model was illustrated by two empirical examples. In the first application for two continuous response variables, it did not perform better than the simpler B&M three-way mixture model, which indicates that the assumption of fixed class membership across situations holds for this data set. The contribution of the new approach was that it provided a test for the assumption of the B&M model. In the second application, the hierarchical mixture model performed much better than the B&M model. Even after taking into account that the situation may itself affect the responses, it was clearly not correct to assume that students use the same type of reaction for each situation.

In the section describing the hierarchical mixture model for three-way data sets, I already mentioned various possible variants and extensions of the proposed model, some of which were used in the two applications. I also discussed the connection to other types of mixture models, such as the hierarchical-mixtures-of-experts model and the

multilevel latent class model. However, I did not mention the connection to the grade-of-membership (GoM) model (Erosheva, 2004; Manton et al., 1994), which is sometimes referred to as the partial- or mixed-membership model. This is a not so well-known variant of the latent class model in which, as in the model proposed here, cases are allowed to belong to each of the latent classes with a certain probability or, in GoM terminology, cases have a certain GoM for each class. The difference between the GoM and the hierarchical mixture model are that in the former each case is assumed to have a unique set of membership probabilities coming from a particular distribution—for example, a Dirichlet distribution—whereas in the latter it is assumed that cases can be clustered based on their memberships probabilities. Actually, our approach can be seen as a nonparametric variant of the GoM model, provided that one either increases K up to a saturation point (Böhning, 2000; Lindsay, 1995) or assumes that the nonparametric maximum likelihood estimate of the mixing distribution has exactly K mass points according to some penalized likelihood criterion such as BIC (Keribin, 2000).

References

- Basford, K.E., McLachlan, G.J., 1985. The mixture method for clustering applied to three-way data. *J. Classification* 2, 109–125.
- Baum, L.E., Petrie, T., Soules, G., Weiss, N., 1970. A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Ann. Math. Statist.* 41, 164–171.
- Böhning, D., 2000. *Computer-Assisted Analysis of Mixtures: Meta-Analysis, Disease Mapping and Others*. Chapman & Hall, London.
- Böhning, D., Dietz, E., Schlattmann, P., 2000. Space–time mixture modelling of public health data. *Statist. Med.* 19, 2333–2344.
- Dayton, C.M., Macready, G.B., 1988. Concomitant-variable latent-class models. *J. Amer. Statist. Assoc.* 83, 173–178.
- Dempster, A.P., Laird, N.M., Rubin, D.R., 1977. Maximum likelihood estimation from incomplete data via the EM algorithm (with discussion). *J. Roy. Statist. Soc. Ser. B* 39, 1–38.
- Erosheva, E.A., 2004. Partial membership models with application to disability survey data. In: Bozdogan, H. (Ed.), *Statistical Data Mining and Knowledge Discovery*. Chapman and Hall, CRC Press, Boca Raton, pp. 117–134.
- Formann, A.K., 1992. Linear logistic latent class analysis for polytomous data. *J. Amer. Statist. Assoc.* 87, 476–486.
- Goodman, L.A., 1974. Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* 61, 215–231.
- Hagenaars, J.A., 1988. Latent structure models with direct effects between indicators: local dependence models. *Sociol. Methods Res.* 16, 379–405.
- Hunt, L.A., Basford, K.E., 1999. Fitting a mixture model to three-mode three-way data with categorical and continuous missing information. *J. Classification* 18, 283–296.
- Hunt, L.A., Basford, K.E., 2001. Fitting a mixture model to three-mode three-way data with missing information. *J. Classification* 18, 209–226.
- Jordan, M.I., Jacobs, R.A., 1994. Hierarchical mixtures of experts and the EM algorithm. *Neural Comput.* 181–214.
- Keribin, C., 2000. Consistent estimation of the order of mixture models. *Sankhyā: Indian J. Statist.* 62, 49–66.
- Knorr-Held, L., Raßer, S., 2000. Bayesian detection of clusters and discontinuities in disease maps. *Biometrics* 56, 13–21.
- Li, J., 2005. Clustering based on a multi-layer mixture model. *J. Comput. Graph. Statist.* 14, 547–568.
- Lindsay, B.G., 1995. *Mixture model: theory, geometry, and applications*, NSF-CBMS Regional Conference Series in Probability and Statistics, Institute of Mathematical Statistics, Alexandria, Virginia, Amer. Statist. Assoc. 4.
- Manton, K.G., Woodbury, M.A., Tolley, H.D., 1994. *Statistical Applications Using Fuzzy Sets*. Wiley, New York.
- McLachlan, G.J., Basford, K.E., 1988. *Mixture Models: Inference and Applications to Clustering*. Marcel Dekker, New York.
- McLachlan, G.J., Peel, D., 2000. *Finite Mixture Models*. Wiley, New York.
- Meulders, M., De Boeck, P., Kuppens, P., Van Mechelen, I., 2002. Constrained latent class analysis of three-way three-mode data. *J. Classification* 19, 277–302.
- Pearl, J., 1988. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, Los Altos, CA.
- Vermunt, J.K., 2003. Multilevel latent class models. *Sociol. Methodol.* 33, 213–239.
- Vermunt, J.K., Magidson, J., 2002. Latent class cluster analysis. In: Hagenaars, J., McCutcheon, A. (Eds.), *Applied Latent Class Models*. Cambridge University Press, Cambridge, pp. 89–106.
- Vermunt, J.K., Magidson, J., 2003. Latent class models for classification. *Comput. Statist. Data Anal.* 41, 531–537.
- Vermunt, J.K., Magidson, J., 2005. *Latent GOLD 4.0 User's Guide*. Statistical Innovations Inc, Belmont, MA.
- Wedel, M., DeSarbo, W.S., 1994. A review of recent developments in latent class regression models. In: Bagozzi, R.P. (Ed.), *Advanced Methods of Marketing Research*. Blackwell Publishers, Cambridge, pp. 352–388.
- Zhang, N.L., 2004. Hierarchical latent class models for cluster analysis. *J. Mach. Learn. Res.* 5, 696–723.