

Latent Markov models: a review of a general framework for the analysis of longitudinal data with covariates

F. Bartolucci · A. Farcomeni · F. Pennoni

the date of receipt and acceptance should be inserted later

Abstract We provide a comprehensive overview of latent Markov (LM) models for the analysis of longitudinal data. Different types of LM models for longitudinal categorical data are reviewed. We illustrate the general version of the LM model which includes individual covariates, and several constrained versions. Constraints make the model more parsimonious and allow to consider and test hypotheses of interest. These constraints may be put on the conditional distribution of the response variables given the latent process (measurement model) or on the distribution of the latent process (latent model). We also illustrate in detail maximum likelihood estimation through the Expectation-Maximization algorithm, which may be efficiently implemented by recursions taken from the hidden Markov literature. We outline methods for obtaining standard errors for the parameter estimates. We also illustrate methods for selecting the number of states and for path prediction. Finally, we mention issues related to Bayesian inference of LM models. Possibilities for further developments are given in the concluding remarks.

Keywords EM algorithm · Bayesian framework · Forward-Backward recursions · Hidden Markov models · Measurement errors · Panel data · Unobserved heterogeneity

Francesco Bartolucci
Department of Economics, Finance and Statistics, University of Perugia, Perugia, Italy
E-mail: bart@stat.unipg.it

Alessio Farcomeni
Department of Public Health and Infectious Diseases, Sapienza - University of Rome, Rome, Italy
E-mail: alessio.farcomeni@uniroma1.it

Fulvia Pennoni
Department of Statistics and Quantitative Methods, University of Milano-Bicocca, Milan, Italy
E-mail: fulvia.pennoni@unimib.it

1 Introduction

In many applications involving longitudinal data the interest is focused on the evolution of a latent characteristic of a group of individuals over time, which is measured by one or more occasion-specific response variables. This characteristic may correspond, for instance, to the quality-of-life of subjects suffering from a certain disease, which is indirectly assessed on the basis of responses to a set of suitably designed items. These items may be repeatedly administered within a certain period of time.

Many models are proposed in the statistical literature for the analysis of longitudinal data; the choice mainly depends on the context of application. For a review see, among others, Fitzmaurice et al (2009). In this literature, the latent Markov (LM) models, on which the present article is focused, have a special role. A comprehensive overview on LM models may be found in Bartolucci et al (2013). These models assume the existence of a latent process which affects the distribution of the response variable. The main assumption behind this approach is that the response variables are conditionally independent given this latent process, which follows a Markov chain with a finite number of states. The basic idea related to this assumption, which is referred to as *local independence*, is that the latent process fully explains the observable behavior of a subject together with possibly available covariates. Therefore, in studying LM models, it is important to distinguish between two components: the *measurement model*, which concerns the conditional distribution of the response variables given the latent process, and the *latent model*, which concerns the distribution of the latent process.

From many points of view, longitudinal data are similar to time-series data. The main difference is that time-series usually consist of many repeated measures referred to a single unit, whereas only few repeated measures are typically available in a longitudinal dataset, but independently replicated for many statistical units. However, inferential approaches arising within the time-series literature cannot be directly extended to models for longitudinal data. In the context of time-series data analysis, for example, the asymptotic properties of an estimator are studied assuming that the number of repeated measures grows to infinity. On the contrary, in the context of longitudinal data analysis, asymptotic properties are studied assuming that the sample size tends to infinity, while the number of occasions of observation is held fixed.

In the statistical and econometric literatures, one of the most interesting model for the analysis of time-series data is the Hidden Markov (HM) model (Baum and Petrie, 1966). This model is based on the same assumptions and estimation methods of the LM model, but the structure of the data it aims to analyze is different. However, the terminology is not univocal and the name HM model is sometimes adopted even for models applied to the analysis of longitudinal data.

Many relevant results in the HM literature have been developed in the 60's and 70's; one of the most relevant paper is due to Baum et al (1970). Then, many advances have been developed in connection with engineering,

informatics, and bioinformatics applications; consider, for instance, the papers by Levinson et al (1983) and Ghahramani and Jordan (1997). For general reviews see the monographs of MacDonald and Zucchini (1997), Koski (2001), and Zucchini and MacDonald (2009), whereas for reviews and recent advances about the estimation methods see Bickel et al (1998), Cappé et al (2005), and Andersson and Rydén (2009).

In the following, we refer to the LM model based on a first-order Markov chain, non-homogeneous transition probabilities, and covariates as the *general* LM model. This model is presented in the case of multivariate data when we observe more response variables at each occasion. For the general LM model we discuss in detail maximum likelihood estimation through the Expectation-Maximization (EM) algorithm (Baum et al, 1970; Dempster et al, 1977). We also briefly outline Bayesian estimation methods even though we acknowledge that other estimation methods are available (Archer and Titterton, 2009; Cappé et al, 2005; Künsch, 2005; Turner, 2008). For the implementation of the EM algorithm, we illustrate suitable recursions which allow us to strongly reduce the computational effort.

The paper also focuses on some constrained versions of the general LM model. Constraints have the aim of making the model more parsimonious and easier to interpret. They are typically formulated to incorporate certain hypotheses that may be interesting to test. Constraints may be posed on the measurement model, that is, on the conditional distribution of the response variables given the latent process, or on the latent model, that is, on the distribution of the latent process. Regarding the former, we discuss in detail parametrizations which makes the latent states interpretable in terms of ability or propensity levels. Regarding the latter, we outline several simplifications of the transition matrix, mostly based on constraints of equality between certain elements of this matrix and/or on the constraint that some elements are equal to 0.

One of the main problems is how to test for model constraints. For this aim, we make use of the likelihood ratio (LR) statistic. When these restrictions concern the transition matrix, the null asymptotic distribution does not necessarily follow an asymptotic chi-squared distribution, but a distribution of chi-bar-squared type (Bartolucci, 2006). We also illustrate extensions relaxing local independence. This is made in two ways: by including among the covariates the lagged response variable, and by allowing the response variables at the same time occasion to be dependent even conditionally on the corresponding latent variable. We review methods for obtaining standard errors for the model parameters. We illustrate the main approaches for selecting the number of latent states and also discuss the problem of path prediction through the Viterbi algorithm (Viterbi, 1967; Juang and Rabiner, 1991).

The paper is organized as follows. In Section 2 we illustrate the main cases of application of the LM model and we briefly outline the model history. Section 3 illustrates different types of LM model through various examples involving longitudinal categorical data, summarizing the results from other papers. In Section 4 we outline the general LM model and constrained versions

of the model based on parsimonious and interpretable parameterizations. In Section 5 we show how to obtain the manifest distribution of the response variables and the posterior distribution of the latent states by the Baum-Welch recursions (Baum et al, 1970; Welch, 2003). Then, we discuss maximum likelihood estimation for the general LM model based on the EM algorithm in Section 6, where we also deal with model identifiability and standard errors for the parameter estimates. Section 7 illustrates methods for selecting the number of states and path prediction. Section 8 briefly outlines Bayesian estimation as an alternative to maximum likelihood estimation. The paper ends with a section where we draw main conclusions and discuss further developments of the present framework.

2 Model motivations and historical review

The use of the LM models finds justification in different types of analysis. We illustrate in the following three main cases where it is sensible to apply LM models and then we provide a historical overview.

1. *Accounting for measurement errors.* In such cases the adopted LM model is seen as an extension of a Markov chain model (Anderson, 1951, 1954), which represents a fundamental model for stochastic processes. As an example, consider the case in which the response variables correspond to a self-rated evaluation of the quality-of-life of an elderly subject at different occasions. The different configurations of each latent variable are interpreted as different levels of the “true” quality-of-life, an individual characteristic which is observable with error due to several factors. In this case, it is reasonable to assume that the latent variables follow a Markov chain, so that the “true” level of quality-of-life is only affected by the previous level of this characteristic. Moreover, it is reasonable to assume that, given the latent variables, the response variables are conditionally independent since the observed quality-of-life only depends on the corresponding “true” level of this characteristic.
2. *Accounting for unobserved heterogeneity between subjects.* The latent variable may have the role of accounting for the unobserved heterogeneity between subjects. This aspect is in connection with the inclusion of the observed covariates in the measurement model which do not fully explain the heterogeneity between the responses provided by different individuals. The advantage with respect to a standard random effect or latent class model with covariates is that we admit that the effect of unobservable covariates has its own dynamics and is not constrained to be time constant. As an example, consider the case in which we observe, at repeated occasions, a binary response variable indicating if an individual has a job position. Then, through an LM model we can take into account the effect of unobservable factors, such as motivations or intelligence, which affect this variable.

3. *Finding latent subpopulations in the population of interest.* The latent states are identified as different subpopulations, with units in the same subpopulation having a common distribution for the response variables. In this context, an LM model may be seen as an extension of the latent class (LC) model (Lazarsfeld, 1950; Lazarsfeld and Henry, 1968) in which subjects are allowed to move between the latent classes during the period of observation. If available, covariates are included in the latent model and then may affect the initial and transition probabilities of the Markov chain. As an example, consider the case of the analysis of criminal data in which for each age band (period of time) we know which types of crime are committed by subjects in a certain cohort. Then, we can first classify subjects in different clusters and then study the transition between different clusters. Each cluster corresponds to a latent state, with subjects in the same state having the same criminal behavior.

The LM approach for longitudinal data dates back to Wiggins' Ph.D. thesis, Wiggins (1955), who introduced a version of this model based on a homogeneous Markov chain, a single outcome at each occasion, and without individual covariates. Wiggins (1955) formulated the model so that the manifest transition is a mixture of a true change and a spurious change due to measurement errors in the observed states. See also Wiggins (1973) for a deep illustration of this model and some simple generalizations.

The most natural extension of the basic LM model of Wiggins is for the inclusion of individual covariates. The formulation proposed by Vermunt et al (1999) is based on the parameterization of initial and transition probabilities of the latent Markov process through a series of multinomial logit regression models depending on time constant and time-varying covariates. Bartolucci et al (2007) extended this approach to the case of more than one response variable. Bartolucci and Pennoni (2007) also allowed transition probabilities to depend on lagged response variables. The formulation proposed by Bartolucci and Farcomeni (2009) includes the covariates in the measurement part of the model in the presence of multivariate responses. They propose to reparametrize the conditional distribution of the response variables given the latent process through a multivariate logistic transformation (McCullagh and Nelder, 1989; Glonek and McCullagh, 1995).

The basic LM approach has been generalized into the mixed LM approach by van de Pol and Langeheine (1990), who also transposed the general approach of Markov chains into the framework of mixture distribution models; see also Rost (2002). This extended approach has also been considered by Altman (2007). In the latter, both fixed and continuous random effects are introduced in the measurement or in the latent part of the model. An interesting review concerning the mixed hidden Markov model in which random effects are assumed to have a discrete distribution may be found in Maruotti (2011). A method based on fixed effects to represent the factors common to all units in the same cluster was proposed by Bartolucci et al (2009). A formulation based on random parameters having a discrete distribution was instead

proposed by Bartolucci et al (2011), in which the context of application is the evaluation of students' math achievement. A related approach is based on the *latent transition analysis* (Bye and Schechter, 1986; Langeheine, 1988; Collins and Wugalter, 1992; Kaplan, 2008) in which the parameters of the LM models are allowed to vary in different latent subpopulations.

It is important to also mention alternative, but related, approaches such as the *latent growth approach* and the *mixture Markov approach*. The latent growth approach is based on a parameterization of the response variables on individual-specific parameters and covariates associated to the time occasions. The individual-specific parameters have a Gaussian distribution with a mean which may also depend on the observable covariates. Models formulated on the basis of this approach have different names, such as *latent growth models* or *latent curve models*. For a review see Muthén (2004), Nagin (1999), and Vermunt (2010); see also Bollen and Curran (2006). The mixture Markov approach is based on separate Markov chain models for latent subpopulations. Each subpopulation has specific initial and transition probabilities of the Markov chain and the probability of belonging to a certain subpopulation typically depends on individual covariates. This approach was developed by Dias and Vermunt (2007) for an application in market segmentation. A particular case is the mover-stayer model of Goodman (1961).

3 Empirical illustrations

In order to illustrate the approaches reviewed in this paper, we provide a synthetic overview of some applications appeared in the literature. Here, we describe different datasets on which these applications are based, the aspects involved in practically fitting and using the LM models, and the interpretation of the results.

3.1 Marijuana Consumption dataset

The univariate version of the LM model without covariates was applied by Bartolucci (2006) to analyze a marijuana consumption dataset based on five annual waves of the "National Youth Survey" (Elliot et al, 1989). The dataset concerns $n = 237$ respondents who were aged 13 years in 1976. The use of marijuana was measured through five ordinal variables, one for each annual wave, with three categories corresponding to: "never in the past", "no more than once in a month in the past year", and "once a month in the past year". The substantive research question is whether there is an increase of marijuana use with age. The LM model, compared for instance to a multinomial logit model, allows to take account measurement errors and is flexible with respect to the distributional assumptions. More explicitly, it allows overdispersion with respect to the multinomial distribution.

The LM model selected in Bartolucci (2006) for the analysis of the marijuana consumption dataset is based on three latent states, homogeneous transition probabilities, and a parsimonious parameterization for the measurement model based on global logits as the one we will illustrate in Example 4 (see Section 4.1). This parametrization is based on one parameter for each latent state, which may be interpreted as the tendency to use marijuana for a subject in this state, and one cutpoint for each response category. Then, the latent states may be ordered representing subjects with “no tendency to use marijuana”, “incidental users of marijuana”, and “subjects with high tendency to use marijuana”. Also note that the cutpoints are common to all the response variables, since these variables correspond to repeated measurements of the same phenomenon under the same circumstances. This is because the dynamics of the marijuana consumption are only ascribed to the evolution of the underlying tendency to use it.

Bartolucci (2006) tested different hypotheses on the transition matrix of the latent process. In particular, he found that the hypothesis that the transition matrix has a tridiagonal structure, as will be illustrated in Example 6, cannot be rejected. The results under the selected LM model say that, at the beginning of the period of observation, 89.6% of the sample is in the first class (lowest tendency to marijuana consumption) and 1.5% is in the third class (highest tendency to marijuana consumption). An interesting interpretation of the pattern of consumption emerges from the estimated transition matrix. A large percentage of subjects remain in the same latent class, but almost the 25% of accidental users switch to the class of high frequency users. From the estimated marginal probabilities of the latent classes results that the tendency to use marijuana increases with age, since the probability of the third class increases across time.

3.2 Educational dataset

An interesting illustration of an LM model with individual covariates was given by Vermunt et al (1999). They used data from an educational panel study conducted by the “Institute for Science Education” in Kiel (Germany); see Hoffmann et al (1985). A cohort of secondary school pupils was interviewed once a year from grade 7 to grade 9 about their interests in physics as a school subject. Response variables have been dichotomized with categories “low” and “high” to avoid sparseness of the observed frequency table. Based on these data, the LM model was used to draw conclusions on whether interest in physics depends on the interest in the previous period of observation and on two available covariates: sex and grade in physics at the present time. The LM model in this application provides a dimension reduction and shows how interest evolves across time.

Vermunt et al (1999) estimated a univariate LM model with both initial and transition probabilities of the latent process depending on the available covariates according to a multinomial logit parametrization. In this model,

the measurement error was constrained to be the same for all time points, meaning that the conditional distribution of the response variables given the latent state is the same at every occasion. Then, Vermunt et al (1999) relaxed some of the basic assumptions of the LM model, such as the assumption that the Markov chain is of first-order.

According to the estimates of the parameters of the selected model, there is a significant effect of sex and grade on the interest in physics. Pupils with higher grades are more interested in physics than pupils with lower grades and girls are less interested in physics than boys. Moreover, the interest has a positive effect on the grade at the next time occasion. For the boys, the probability of switching from “low” to “high interest” is larger than that for girls, as well as to keep their interest high.

3.3 Criminal dataset

The multivariate version of the LM model where both the initial and transition probabilities of the latent process depend on time-constant covariates was illustrated by Bartolucci et al (2007); for a related study see Roeder et al (1999). Bartolucci et al (2007) analyzed the conviction histories of a cohort of offenders who were born in England and Wales in 1953. The offenders were followed from the age of criminal responsibility, 10 years, until the end of 1993. They were grouped in the following 10 major categories: violence against the person, sexual offenses, burglary, robbery, theft and handling stolen goods, fraud and forgery, criminal damage, drug offenses, motoring offenses, and other offenses. Gender was included in the model as an explanatory variable. The analysis was based on $T = 6$ age bands: 10-15, 16-20, 21-25, 26-30, 31-35, and 36-40 years. The adopted multivariate LM model allows us to estimate trajectories for behavioral types which are determined by the criminal conviction grouping. It also allows us to estimate different transition probabilities depending on gender. The approach uses a general population sample by augmenting the observed sample with not-convicted subjects.

According to Bartolucci et al (2007), the fit of the model is considerably improved by relaxing the assumption of homogeneity of the latent Markov chain, but retaining the constraint that males and females have the same transition probabilities. In particular, they selected a model based on partial homogeneity, in which there are two transition probability matrices: the first for transitions up to time \bar{t} and the second from time \bar{t} and beyond. The choice of $\bar{t} = 2$ has been made on the basis of a suitable statistical selection criterion.

In summary, the selected model is based on a partially homogeneous Markov chain with five latent states, different initial and equal transition probabilities for males and females. From the estimated conditional probabilities of conviction for any offense group and any latent state it was possible to determine classes of criminal activity. In accordance to the typologies found in the criminological latent class literature, these classes are interpreted as:

“non-offenders”, “incidental offenders”, “violent offenders”, “theft and fraud offenders”, and “high frequency and varied offenders”.

According to the estimated initial probabilities in the first age band, the percentage of males who are incidental offenders is higher than that of females. The common estimated transition probabilities for males and females from age band 10-15 to age band 16-20, and from one age band to the others for offenders over 16, show that the first transition occurs at an early age, 16 years, which in western society represents the peak of the age-crime curve.

At the first time occasions, “incidental offenders” have a quite high probability of persistence when moving from the age band 10-15 to age band 16-20. Moreover, “theft and fraud offenders” are mainly females and they have a high chance of moving to the class of non-offenders. The “high frequency and varied offenders” are mainly males and they have a high persistence.

According to the estimated transition probabilities, from age band 16-20 to the subsequent age bands the subjects belonging to the latent state of “non-offenders” have a very low chance of becoming offenders; “theft and fraud offenders” and “violent offenders” have a high probability of dropping out of crime. From the estimated proportion of males and females in each latent state at every time occasion resulted that 7% of males are “violent offenders” at age 16-20 years and 32% are “incidental offenders” at the same age. Only 3% of females are “theft and fraud offenders” at age 16-20 years.

3.4 Dataset on financial products preferences

An interesting analysis of data obtained from face to face interviews of the household ownership of 12 financial products is offered in Paas et al (2009). The panel was conducted by a market research company among 7676 Dutch households in 1996, 1998, 2000, and 2002. To have an accurate representation of the products portfolio the households were asked to retrieve their bank and insurance records. Households that dropped out were replaced to ensure the representativeness of the sample with respect to demographic variables, such as age, income, and marital status. The aim of the study was to get insights on the developments of the individual household product portfolio and on the effect of demographic covariates on such development. It also focused on predicting future behaviors of acquisition.

Paas et al (2009) proposed to use a time homogeneous multivariate LM model, with time-varying covariates affecting the latent process. They included additional assumptions, such as constant conditional probabilities of the response variables given the latent process. This is done to avoid manifest changes, so that the product penetration levels are consistent in latent states over measurement occasions. Moreover, they formulated the model with a time-constant effect of the covariates on the transition probabilities. Due to this parameterization of the latent model, the states of the chain have an interpretation according to different risk levels of investment.

The model selected on the basis of the Bayesian Information Criterion is based on nine latent states. These states can be ordered according to increasing penetration levels across the analyzed products, which range from bonds, the most common owned product, to saving accounts. The results highlight some divergences from common order of acquisitions, such as the acquisition of a mortgage before owing a credit card or vice-versa. Loans and unemployment insurance are the most often acquired products. According to the estimated transition matrix there is a high persistence in the same latent state: only 14% of the households changed latent state during the study period. The most common switch is from latent class 7 to 8, where latent state 7 is characterized by the acquisition of mortgage, life insurance, pension fund, car insurance, and saving accounts, whereas latent state 8 is for all the previous products plus the credit card. Another common switch is from latent state 4 to 7, where the first is characterized by the acquisition of life insurance, pension fund, car insurance, and savings accounts. This means that multiple products were acquired between consecutive measurement occasions.

Income, age of the head of the household, and household size have a significant effect on the initial and transition probabilities according to the Wald test. The covariate values implying a larger probability of belonging to an initial latent class also imply a greater probability of switching into the same latent class. For example, larger households are relatively often found in latent states where overall product ownership probabilities are relatively low.

The prediction of future purchase of a financial product was performed on the basis of the posterior latent state membership probabilities for each household at the last occasion, given all other observed information. To assess the accuracy of the forecasting, Paas et al (2009) used the Gini coefficient. Considering the empirical results in the last wave refereed to year 2002, which was not considered when estimating the model, the authors showed that the prediction equations are effective for forecasting household acquisition for most products.

3.5 Job position dataset

The multivariate LM model with covariates affecting the manifest probabilities proposed by Bartolucci and Farcomeni (2009) was applied by these authors to data extracted from the “Panel Study of Income Dynamics” database (University of Michigan). These data concern $n = 1446$ women who were followed from 1987 to 1993. The binary response variables are fertility, indicating whether a woman had given birth to a child in a certain year, and employment, indicating whether she was employed. The covariates are: race (dummy variable equal to 1 for a black woman), age (in 1986), education (year of schooling), child 1-2 (number of children in the family aged between 1 and 2 years, referred to the previous year), child 3-5, child 6-13, child 14-, income of the husband (in dollars, referred to the previous year).

The main issue concerns the direct effect of fertility on employment. Also of interest are the strength of the state dependence effect for both response variables and how these variables depend on the covariates. By a suitable formulation of the LM model it is possible to take into account that the unobserved covariates may have a time varying effect. Moreover, the contemporary dependence between two response variable may be disentangled from the serial dependence.

Bartolucci and Farcomeni (2009) used an LM model with covariates affecting the manifest probabilities since they were interested in separately estimating the effect of each covariate on each outcome. The proposed LM model allows to separate these effects from the unobserved heterogeneity, by modeling the latter with a latent Markov process. In this way, unobserved heterogeneity effects on the response variables are allowed to be time-varying; this is not allowed neither within a latent class model with covariates nor in the most common random-effects models.

The model selected using Akaike and Bayesian Information criteria is based on three latent states. Under this model, race has a significant effect on fertility, but not on employment according to the estimates of the parameters affecting the marginal logits of fertility and employment and the log-odds ratio between these variables. Age has a stronger effect on fertility than on employment. Education has a significant effect on both fertility and employment, whereas the number of children in the family strongly affects only the first response variable and income of the husband strongly affects only the second one.

The log-odds ratio between the two response variables, given the latent state, is negative and highly significant, meaning that fertility and employment are negatively associated when referred to the same year. On the other hand, lagged fertility has a significant negative effect on both response variables and lagged employment has a significant effect, which is positive, on both response variables. Therefore, fertility has a negative effect on the probability of having a job position in the same year of the birth, and the following one. Employment is serially positively correlated (as consequence of the state dependence effect) and fertility is negatively serially correlated.

From the estimates of the support points for each latent state it may be deduced that the latent states correspond to different levels of propensity to give birth to a child and to have a job position. The first latent state corresponds to subjects with the highest propensity to fertility and the lowest propensity to have a job position. On the contrary, the third latent state corresponds to subjects with the lowest propensity to fertility and the highest propensity to have a job position. Finally, the second state is associated to intermediate levels of both propensities. The two latent propensities are negatively correlated.

Overall, it results that the 78.5% of women started and persisted in the same latent state for the entire period, whereas for 21.5% of women had one or more transitions between states. The presence of these transitions is in accordance to the rejection of the hypothesis that a latent class model is suitable for these data.

3.6 Dataset on anorectic patients

An interesting extension of the LM model to account for a hierarchical structure of the data has been recently proposed by Rijmen et al (2007). They illustrated the model with an application based on a dataset from an ecological momentary assessment study (Vansteelandt et al, 2007) on the course of emotions among anorectic patients. At nine occasions for each of the seven days of observations, 32 females with eating disorders received a signal and were asked to rate themselves on a 7-point scale with respect to the intensity with which they experienced 12 emotional states. These were taken from the following emotional categories: “anger and irritation”, “shame and guilt”, “anxiety and tension”, “sadness and loneliness”, “happiness and joy”, and “love and appreciation”. The responses were dichotomized (0-2 vs. 3-6) and the signal was considered equally spaced. The aim of the study was to detect the course of emotion among the patients. The repeated measurements are useful to increase power mainly when there is high correlation among them. The LM model is useful to model the influence provided by each group of the hierarchy on the responses provided by subjects who are included in the group.

As a preliminary analysis, Rijmen et al (2007) used the univariate version of the LM model without covariates. They treated each person by day combination as a separate case assuming that the data stemming from different days are independent and that the parameters are constant over days. On the basis of such model the authors selected four latent states. The first state is interpreted as positive mood, the third as negative mood, the second as low intensity for all emotions except tension, and the fourth as neutral to moderately positive mood. According to the estimated transition matrix there is a high persistence in the same state. The probability of moving from state 1 to state 2 is 0.14, from state 3 to 4 is 0.18, and from state 4 to 3 is 0.14. They noted that there is an indirect transition from state 3 to state 1 via the emotionally more neutral state 4. Over the days there is an increase of the marginal probabilities of states 1 and 4 indicating that the mood of patients tends to become better later on in the day.

Then the authors also used a hierarchical LM model by introducing a latent variable at day level to account for the fact that data stemming from different days are not independent. They modeled the transition between latent states at day and signal levels by a first-order time homogeneous Markov chain. They estimated a model with two states at the day level and two signal states within each day-state. For the first day state, the signal state 1 is characterized by high probabilities of experiencing positive emotions and low probabilities of negative emotions. Therefore this state is interpreted as positive mood. Instead the second day state is interpreted as negative mood. In the signal state 2, positive emotions are not well separated from negative emotions and the state is considered as an emotionally neutral to moderately positive state. A tendency to experience more positive emotions emerges from the estimated initial and conditional probabilities of the chain over day and signal.

4 The General latent Markov model framework

In the following, we illustrate the general LM model framework and some constrained versions of the model.

4.1 Model formulation

Let T denote the number of occasions of observation and suppose that, at each time occasion, we observe r response variables denoted by $Y_j^{(t)}$, with $j = 1, \dots, r$ and $t = 1, \dots, T$. These response variables are collected, for each time occasion, in the random vector $\mathbf{Y}^{(t)}$, $t = 1, \dots, T$; we also denote by $\tilde{\mathbf{Y}}$ the overall vector of the rT response variables. When available, we also denote by $\mathbf{X}^{(t)}$ a vector of covariates corresponding to $\mathbf{Y}^{(t)}$ and by $\tilde{\mathbf{X}}$ we denote the vector of all the individual covariates which is obtained by stacking the vectors $\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(T)}$.

In order to model the distribution of $\tilde{\mathbf{Y}}$ given $\tilde{\mathbf{X}}$, the LM approach assumes the existence of a latent process $\mathbf{U} = (U^{(1)}, \dots, U^{(T)})$, which is assumed to follow a first-order Markov chain with state space $\{1, \dots, k\}$, where k is the number of latent states. Under the assumption of local independence, the random vectors $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(T)}$ are conditionally independent given the latent process. This assumption leads to a strong simplification of the model, but it can be relaxed as we illustrate in the following. Moreover, it is assumed that the distribution of each response vector $\mathbf{Y}^{(t)}$ only depends on the corresponding latent variable $U^{(t)}$ and the possible available covariates in $\mathbf{X}^{(t)}$.

In order to clarify the above assumptions and provide further developments, we introduce a general notation according to which, for a random variable A , by $f_A(a)$ we denote the probability mass (or density) function for the distribution of A and, given two random variables A and B , by $f_{B|A}(b|a)$ we denote that for the conditional distribution of B given A . Then, with reference to the LM framework introduced above, we let

$$\begin{aligned} f^{(t)}(\mathbf{y}|u, \mathbf{x}) &= f_{\mathbf{Y}^{(t)}|U^{(t)}, \mathbf{X}^{(t)}}(\mathbf{y}|u, \mathbf{x}), \quad t = 1, \dots, T, \quad u = 1, \dots, k, \\ p(u|\mathbf{x}) &= f_{U^{(1)}|\mathbf{X}^{(1)}}(u|\mathbf{x}), \quad u = 1, \dots, k, \\ p^{(t)}(u|\bar{u}, \mathbf{x}) &= f_{U^{(t)}|U^{(t-1)}, \mathbf{X}^{(t)}}(u|\bar{u}, \mathbf{x}), \quad t = 2, \dots, T, \quad \bar{u}, u = 1, \dots, k. \end{aligned}$$

In the above expressions, by u we indicate the current latent state, by \bar{u} the previous latent state, by \mathbf{x} a realization of $\mathbf{X}^{(t)}$, and by \mathbf{y} a realization of $\mathbf{Y}^{(t)}$. When we use an LM model for a specific application, we formulate suitable assumptions on the latent and the measurement models. Formulating assumptions on the measurement model amounts to express $f^{(t)}(\mathbf{y}|u, \mathbf{x})$ as a function of suitable parameters to be estimated. Accordingly, formulating assumptions on the latent model amounts to suitably parametrize the initial probability function $p(u|\mathbf{x})$ and the transition probability function $p^{(t)}(u|\bar{u}, \mathbf{x})$. How to express these parametrizations is clarified in the following.

4.1.1 Formulation of the measurement model

In order to parametrize the probability (or density) function $f^{(t)}(\mathbf{y}|u, \mathbf{x})$, it is necessary to disentangle the univariate case, in which only one response variable is observed at each time occasion, and then $\mathbf{Y}^{(t)} = Y^{(t)}$ and we use the notation $f^{(t)}(y|u, \mathbf{x})$, from the multivariate case in which we observe more response variables at each occasion.

In the univariate case, we have first to formulate a specific assumption on the distribution of $Y^{(t)}$ given $U^{(t)}$ and $\mathbf{X}^{(t)}$, which may belong to an arbitrary family depending on a set of parameters $\boldsymbol{\eta}_{u\mathbf{x}}^{(t)}$. See also Farcomeni (2012) for the assumptions needed to model quantiles of a continuous outcome. Then, we assume that

$$\boldsymbol{\eta}_{u\mathbf{x}}^{(t)} = \mathbf{W}_{u\mathbf{x}}^{(t)}\boldsymbol{\beta}, \quad u = 1, \dots, k, \quad (1)$$

where $\mathbf{W}_{u\mathbf{x}}^{(t)}$ is a design matrix depending on u and \mathbf{x} and $\boldsymbol{\beta}$ is a reduced vector of parameters. When $\boldsymbol{\eta}_{u\mathbf{x}}^{(t)}$ is obtained by a suitable link function of the mean of this distribution, it reduces to a scalar which is denoted by $\eta_{u\mathbf{x}}^{(t)}$. Link functions commonly used are found within the Generalized Linear Model (GLM) literature (McCullagh and Nelder, 1989). When $Y^{(t)}$ is a categorical variable with $c > 2$ categories, indexed from 0 to $c - 1$, more complex link functions have to be used, which also may take into account the ordinal nature of the response variables. Following Colombi and Forcina (2001), these link functions may be expressed as

$$\boldsymbol{\eta}_{u\mathbf{x}}^{(t)} = \mathbf{C} \log[\mathbf{M}\mathbf{f}^{(t)}(u, \mathbf{x})], \quad (2)$$

where \mathbf{C} is a suitable matrix of constraints, \mathbf{M} is a marginalization matrix with elements 0 and 1, and $\mathbf{f}^{(t)}(u, \mathbf{x})$ is a c -dimensional column vector with elements $f^{(t)}(y|u, \mathbf{x})$ for all possible values of y . In this case we denote by $\eta_{y|u\mathbf{x}}^{(t)}$, $y = 1, \dots, c - 1$, each element of $\boldsymbol{\eta}_{u\mathbf{x}}^{(t)}$.

The two following examples clarify the possible formulations.

Example 1 In the case of continuous response variables, we can simply assume that, given $U^{(t)}$ and $\mathbf{X}^{(t)}$, $Y^{(t)}$ has normal distribution with mean $\eta_{u\mathbf{x}}^{(t)}$. This amounts to use an identity link function in the GLM terminology. The resulting model then assumes that

$$Y^{(t)} = \beta_{1u} + \mathbf{x}'\boldsymbol{\beta}_2 + \varepsilon_t, \quad t = 1, \dots, T,$$

with $\varepsilon_t \sim N(0, \sigma^2)$, where σ^2 is a variance parameter to be estimated. This model is then an extension of the linear regression model, with time-varying discrete random effects.

Example 2 In the case of binary response variables, it is natural to assume that, given $U^{(t)}$ and $\mathbf{X}^{(t)}$, $Y^{(t)}$ has a Bernoulli distribution with a certain

“success” probability, the logit of which corresponds to $\eta_{u\mathbf{x}}^{(t)}$. The resulting model then assumes that

$$\log \frac{f^{(t)}(1|u, \mathbf{x})}{f^{(t)}(0|u, \mathbf{x})} = \beta_{1u} + \mathbf{x}'\boldsymbol{\beta}_2, \quad t = 1, \dots, T.$$

In this case, we have an extension of the standard logit model with random effects.

Example 3 Suppose that $Y^{(t)}$ is a nominal categorical variable with $c > 2$ categories, which are labelled from 0 to $c-1$. Then, the conditional distribution of this variable, given $U^{(t)}$ and $\mathbf{X}^{(t)}$, is assumed to be multinomial. This distribution may be parametrized by (local) logits, which are computed as

$$\eta_{y|u\mathbf{x}} = \log \frac{f(y|u, \mathbf{x})}{f(0|u, \mathbf{x})}, \quad y = 1, \dots, c-1.$$

Once these logits are collected in the $c-1$ dimensional vector $\boldsymbol{\eta}_{u\mathbf{x}}$, we can formulate different assumptions. A common approach is to formulate category specific parameters as in

$$\eta_{y|u\mathbf{x}} = \beta_{1y} + \beta_{2u} + \mathbf{x}'\boldsymbol{\beta}_{3y}, \quad y = 1, \dots, c-1, u = 1, \dots, k.$$

This model can also be reduced to

$$\eta_{y|u\mathbf{x}} = \beta_{1y} + \beta_{2u} + \mathbf{x}'\boldsymbol{\beta}_3, \quad y = 1, \dots, c-1, u = 1, \dots, k.$$

Example 4 Suppose that $Y^{(t)}$ is an ordinal variable with c categories. Its distribution may be parametrized by the so-called global logits, which are computed as

$$\eta_{y|u\mathbf{x}} = \log \frac{f(y|u, \mathbf{x}) + \dots + f(c-1|u, \mathbf{x})}{f(0|u, \mathbf{x}) + \dots + f(y-1|u, \mathbf{x})}, \quad y = 1, \dots, c-1. \quad (3)$$

Once these logits are collected in the $c-1$ dimensional vector $\boldsymbol{\eta}_{u\mathbf{x}}$, we can formulate an assumption of type

$$\eta_{y|u\mathbf{x}} = \beta_{1y} + \beta_{2u} + \mathbf{x}'\boldsymbol{\beta}_3, \quad y = 1, \dots, c-1, u = 1, \dots, k,$$

where β_{1y} are cutpoints, β_{2u} are intercepts specific of the corresponding latent state, and $\boldsymbol{\beta}_3$ is a vector of parameters for the covariates. In this way, we are generalizing the proportional odds model of Glonek and McCullagh (1995) by including time-varying discrete random effects.

It is interesting to note that the link function may be expressed as in (2). With $c = 4$, for instance, we obtain global logits by letting

$$\mathbf{C} = \begin{pmatrix} -1 & 0 & 0 & 1 & 0 & 0 \\ 0 & -1 & 0 & 0 & 1 & 0 \\ 0 & 0 & -1 & 0 & 0 & 1 \end{pmatrix} \quad \mathbf{M} = \begin{pmatrix} 1 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

In the case of multivariate responses, it is common to formulate an extended version of the assumption of local independence, according to which the single response variables $Y_j^{(t)}$ in every vector $\mathbf{Y}^{(t)}$ are conditionally independent given $U^{(t)}$ and $\mathbf{X}^{(t)}$. Adopting an obvious notation, we have that

$$f^{(t)}(\mathbf{y}|u, \mathbf{x}) = \prod_j f_j^{(t)}(y_j|u, \mathbf{x}). \quad (4)$$

Each univariate distribution $f_j^{(t)}(y|u, \mathbf{x})$ may be parametrized as above depending on a specific parameters $\boldsymbol{\eta}_{j|u\mathbf{x}}^{(t)}$. Then, we assume that

$$\boldsymbol{\eta}_{j|u\mathbf{x}}^{(t)} = \mathbf{W}_{j|u\mathbf{x}}^{(t)}\boldsymbol{\beta}, \quad j = 1, \dots, r, u = 1, \dots, k, \quad (5)$$

where $\mathbf{W}_{j|u\mathbf{x}}^{(t)}$ is a suitable design matrix.

In the case of multivariate data, it may be also reasonable to assume that the response variables in $\mathbf{Y}^{(t)}$ are not conditionally independent given $U^{(t)}$ and $\mathbf{X}^{(t)}$, allowing then for a form of *contemporary dependence*. In this case, we adopt a single link function for all the multivariate distributions $f^{(t)}(\mathbf{y}|u, \mathbf{x})$ depending on a single vector of parameters $\boldsymbol{\eta}_{u\mathbf{x}}$ which is parametrized as in (1). Such an approach has been proposed for categorical response variables by Bartolucci and Farcomeni (2009). However, the above approach is difficult to adopt when outcomes are of mixed nature, or when all outcomes are continuous but not conditionally Gaussian. For this reason, it is in general assumed that (4) holds.

In certain applications we may want to make the conditional distribution of $\mathbf{Y}^{(t)}$ given $U^{(t)}$ independent of the covariates $\mathbf{X}^{(t)}$. These covariates may simply not be available. In these cases, every function $f^{(t)}(\mathbf{y}|u, \mathbf{x})$ depends on a vector of parameters of type $\boldsymbol{\eta}_u$ or on vectors of parameters $\boldsymbol{\eta}_{j|u}$ in the multivariate case under assumption (4). These parameters may be considered as unconstrained or may be constrained by a formulation similar to those in (1) and (5), in which the design matrix does not depend on the covariates.

This is clarified in the following example.

Example 5 As in Example 4, suppose that $Y^{(t)}$ is an ordinal variable with c categories. Then, we may directly parametrize the conditional distribution of each variable $Y^{(t)}$, given $U^{(t)}$ and $\mathbf{X}^{(t)}$, by letting $\boldsymbol{\eta}_u$ equal to the vector of global logits defined in (3), which are denoted by $\eta_{y|u}$. Otherwise, we may assume that

$$\eta_{y|u} = \beta_{1y} + \beta_{2u}, \quad y = 1, \dots, c-1, u = 1, \dots, k,$$

where the parameter β_{2u} is a measure of the effect of the latent state on the probability of responding by a higher category, so that the latent states are easily interpretable in terms of tendency or propensity towards a certain behavior. An example of parametrizations of this type is for modeling data about drug consumption measured by an ordinal variables where the parameter β_{2u} directly measures the tendency towards this behavior for subject in latent state u ; see Section 3.1 for an description of an application of this type.

Finally, it is interesting to note that by including in each vector of covariates $\mathbf{X}^{(t)}$ the lagged response variables, that is the vector $\mathbf{Y}^{(t-1)}$ or a transformation of this vector, we can easily relax the hypothesis of local independence by allowing *serial dependence*; see Bartolucci and Farcomeni (2009) for details.

4.1.2 Formulation of the latent model

Here we illustrate different parameterizations of the initial and transition probabilities of the latent Markov chain.

For the initial probabilities, in particular, we assume

$$\boldsymbol{\lambda}_{\mathbf{x}} = \mathbf{Z}_{\mathbf{x}}\boldsymbol{\gamma},$$

where $\boldsymbol{\lambda}_{\mathbf{x}}$ is a vector obtained by a particular link function of the probabilities $(p(1|\mathbf{x}), \dots, p(k|\mathbf{x}))$, $\mathbf{Z}_{\mathbf{x}}$ is a design matrix depending on the covariates in \mathbf{x} , and $\boldsymbol{\gamma}$ is the corresponding parameter vector. Similarly, for the transition probabilities we have

$$\boldsymbol{\lambda}_{\bar{u}\mathbf{x}}^{(t)} = \mathbf{Z}_{\bar{u}\mathbf{x}}^{(t)}\boldsymbol{\delta},$$

where $\boldsymbol{\lambda}_{\bar{u}\mathbf{x}}^{(t)}$ is a link function of $(p^{(t)}(1|\bar{u}), \dots, p^{(t)}(k|\bar{u}))$, $\mathbf{Z}_{\bar{u}\mathbf{x}}^{(t)}$ is a suitable design matrix depending on the previous latent state \bar{u} , and $\boldsymbol{\delta}$ is the corresponding vector of parameters. The above link functions may be formulated on the basis of different types of logit. Typically, we use multinomial logits or global logits. As reference category, the multinomial logits have the first category when modeling the initial probabilities and category u when modeling the transition probabilities. Global logits are used when the latent states are ordered on the basis of a suitable parametrization of the conditional distribution of the response variables given the latent process. In any case, these link functions may be expressed as in (2). In certain cases it is reasonable to parametrize only some of the transition probabilities, while leaving the others constrained to 0, so that a more parsimonious model results.

The following example helps to clarify the parametrizations described above.

Example 6 - Tridiagonal transition matrix with logit parametrization.

Assuming that the latent states are ordered, we may require that the initial probabilities depend on the covariates in $\mathbf{X}^{(1)}$ through a global logits link function as follows

$$\lambda_{u|\mathbf{x}} = \log \frac{p(u+1|\mathbf{x}) + \dots + p(k|\mathbf{x})}{p(1|\mathbf{x}) + \dots + p(u|\mathbf{x})} = \gamma_{1u} + \mathbf{x}'\boldsymbol{\gamma}_2, \quad u = 1, \dots, k-1.$$

Moreover, it is reasonable to assume a particular structure for the transition matrices in order to achieve a reduction of the number of model parameters. An example is when they are tridiagonal, so that transition from state \bar{u} is

only allowed to state $u = \bar{u} - 1, \bar{u} + 1$; with $k = 4$, for instance, we have the transition matrix

$$\mathbf{P}_{\mathbf{x}}^{(t)} = \begin{pmatrix} p^{(t)}(1|1, \mathbf{x}) & p^{(t)}(2|1, \mathbf{x}) & 0 & 0 \\ p^{(t)}(1|2, \mathbf{x}) & p^{(t)}(2|2, \mathbf{x}) & p^{(t)}(3|2, \mathbf{x}) & 0 \\ 0 & p^{(t)}(2|3, \mathbf{x}) & p^{(t)}(3|3, \mathbf{x}) & p^{(t)}(4|3, \mathbf{x}) \\ 0 & 0 & p^{(t)}(3|4, \mathbf{x}) & p^{(t)}(4|4, \mathbf{x}) \end{pmatrix}. \quad (6)$$

This constraint makes sense only if the latent states are suitably ordered. In this case, we may also assume a parameterization based on logits of the following type

$$\lambda_{u|\bar{u}\mathbf{x}}^{(t)} = \log \frac{p^{(t)}(u|\bar{u}, \mathbf{x})}{p^{(t)}(\bar{u}|\bar{u}, \mathbf{x})} = \delta_{1u} + \delta_{2, (3+u-\bar{u})/2} + \mathbf{x}' \boldsymbol{\delta}_3, \quad t = 2, \dots, T, \bar{u} = 1, \dots, k,$$

where $u = 2$ for $\bar{u} = 1$, $u = k - 1$ for $\bar{u} = k$, and $u = \bar{u} - 1, \bar{u} + 1$ for $\bar{u} = 2, \dots, k - 1$.

Finally, when covariates are assumed to enter in the latent model, they are typically excluded from the measurement model and vice-versa. This simplifies the interpretation of the model and reduce the possibilities that the resulting model is not identifiable. See also Section 2 for further comments about where to include the individual covariates depending on the aim of the application of interest and the type of data.

5 Manifest and posterior distributions

Regardless of the specific model formulation, the assumption of local independence between the response vectors $\mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(T)}$ implies that for the conditional distribution of $\tilde{\mathbf{Y}}$, given $\tilde{\mathbf{X}}$ and \mathbf{U} , we have

$$f_{\tilde{\mathbf{Y}}|\mathbf{U}, \tilde{\mathbf{X}}}(\tilde{\mathbf{y}}|\mathbf{u}, \tilde{\mathbf{x}}) = \prod_t f^{(t)}(\mathbf{y}^{(t)}|u^{(t)}, \mathbf{x}^{(t)}),$$

for any realization $\tilde{\mathbf{y}}$ of $\tilde{\mathbf{Y}}$, $\tilde{\mathbf{x}}$ of $\tilde{\mathbf{X}}$, \mathbf{u} of \mathbf{U} . Moreover, since we assume that the latent process follows a first-order Markov chain, we have that

$$f_{\mathbf{U}|\tilde{\mathbf{X}}}(\mathbf{u}|\tilde{\mathbf{x}}) = p(u^{(1)}|\mathbf{x}^{(1)}) \prod_{t=2}^T p^{(t)}(u^{(t)}|u^{(t-1)}, \mathbf{x}^{(t)}).$$

Now let

$$f(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = f_{\tilde{\mathbf{Y}}|\tilde{\mathbf{X}}}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})$$

denote the probability mass (or density) function for the *manifest distribution* of $\tilde{\mathbf{Y}}$ given $\tilde{\mathbf{X}}$. We have that

$$f(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = \sum_{\mathbf{u}} f_{\tilde{\mathbf{Y}}|\mathbf{U}, \tilde{\mathbf{X}}}(\tilde{\mathbf{y}}|\mathbf{u}, \tilde{\mathbf{x}}) f_{\mathbf{U}|\tilde{\mathbf{X}}}(\mathbf{u}|\tilde{\mathbf{x}}), \quad (7)$$

where the sum $\sum_{\mathbf{u}}$ is extended to all possible configurations $\mathbf{u} = (u^{(1)}, \dots, u^{(T)})$ of the latent process.

Another distribution of fundamental importance is the *posterior distribution* of the latent process, which corresponds to the conditional distribution of U given $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$. In this case, we adopt the notation

$$q(\mathbf{u}|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = f_{U|\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}}(\mathbf{u}|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}),$$

which corresponds to the probability mass function

$$q(\mathbf{u}|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = \frac{f_{\tilde{\mathbf{Y}}|U, \tilde{\mathbf{X}}}(\tilde{\mathbf{y}}|\mathbf{u}, \tilde{\mathbf{x}})f_{U|\tilde{\mathbf{X}}}(\mathbf{u}|\tilde{\mathbf{x}})}{f(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})}. \quad (8)$$

This distribution will be used for parameter estimation and for predicting the latent state at each time occasion.

It has to be clear that the sum in (7) may be extended over a huge number of configurations \mathbf{u} . The number of these configurations is k^T . Therefore, computing the manifest distribution of $\tilde{\mathbf{Y}}$, as expressed in (7), may be computationally demanding in real applications. If for instance we fix $k = 4$ and have $T = 10$, the number of terms involved is about one million. The same also happens for the posterior distribution in (8), since it involves the manifest distribution at the denominator. However, in order to deal with these distributions, we can employ certain recursions which have been developed, in a series of fundamental papers in the HM literature, by Baum and Welch and their coauthors; see Baum et al (1970) and Welch (2003). For a review see Khreich et al (2010) and Bartolucci and Pandolfi (2013).

The first recursion of Baum and Welch is a forward recursion which allows us to compute the manifest distribution of $\tilde{\mathbf{Y}}$. Let

$$l^{(t)}(u, \tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = f_{U^{(t)}, \mathbf{Y}^{(1)}, \dots, \mathbf{Y}^{(t)}|\mathbf{X}^{(1)}, \dots, \mathbf{X}^{(t)}}(u, \mathbf{y}^{(1)}, \dots, \mathbf{y}^{(t)}|\mathbf{x}^{(1)}, \dots, \mathbf{x}^{(t)}),$$

for $t = 1, \dots, T$ and $u = 1, \dots, k$. This recursion is initialized with

$$l^{(1)}(u, \tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = p(u|\mathbf{x}^{(1)})f^{(1)}(\mathbf{y}^{(1)}|u, \mathbf{x}^{(1)}), \quad u = 1, \dots, k,$$

and then is based on the following step

$$l^{(t)}(u, \tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = \sum_{\bar{u}} l^{(t-1)}(\bar{u}, \tilde{\mathbf{y}}|\tilde{\mathbf{x}})p^{(t)}(u|\bar{u}, \mathbf{x}^{(t)})f^{(t)}(\mathbf{y}^{(t)}|u, \mathbf{x}^{(t)}), \quad u = 1, \dots, k,$$

to be performed for $t = 2, \dots, T$. At the end of the forward recursion, the manifest distribution is simply obtained as

$$f(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}) = \sum_u l^{(T)}(u, \tilde{\mathbf{y}}|\tilde{\mathbf{x}}).$$

The other recursion introduced by Baum and Welch is a backward recursion, which allows us to obtain the posterior distribution of every latent state and of every pair of consecutive latent states. Let

$$m^{(t)}(\tilde{\mathbf{y}}|\bar{u}, \tilde{\mathbf{x}}) = f_{\mathbf{Y}^{(t+1)}, \dots, \mathbf{Y}^{(T)}|U^{(t)}, \mathbf{X}^{(t+1)}, \dots, \mathbf{X}^{(T)}}(\mathbf{y}^{(t+1)}, \dots, \mathbf{y}^{(T)}|\bar{u}, \mathbf{x}^{(t+1)}, \dots, \mathbf{x}^{(T)}),$$

for $t = 1, \dots, T - 1$ and $\bar{u} = 1, \dots, k$. This recursion is initialized with

$$m^{(T)}(\tilde{\mathbf{y}}|\bar{u}, \tilde{\mathbf{x}}) = 1, \quad \bar{u} = 1, \dots, k,$$

and is then based on the following steps

$$m^{(t)}(\tilde{\mathbf{y}}|\bar{u}, \tilde{\mathbf{x}}) = \sum_u m^{(t+1)}(\tilde{\mathbf{y}}|u, \tilde{\mathbf{x}}) p^{(t+1)}(u|\bar{u}, \mathbf{x}^{(t+1)}) f^{(t+1)}(\mathbf{y}^{(t+1)}|u), \quad (9)$$

where $\bar{u} = 1, \dots, k$, to be performed in reverse order, that is from $t = T - 1$ to $t = 1$.

From the results of above recursion, and also using the results of the forward recursion, we obtain the following posterior distributions

$$q^{(t)}(u|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) = f_{U^{(t)}|\tilde{\mathbf{Y}}}(u|\tilde{\mathbf{y}}) = \frac{l^{(t)}(u, \tilde{\mathbf{y}}|\tilde{\mathbf{x}}) m^{(t)}(\tilde{\mathbf{y}}|u, \tilde{\mathbf{x}})}{f(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})}, \quad u = 1, \dots, k, \quad (10)$$

for $t = 1, \dots, T$. We also obtain

$$\begin{aligned} q^{(t)}(\bar{u}, u|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) &= f_{U^{(t-1)}, U^{(t)}|\tilde{\mathbf{Y}}}(\bar{u}, u|\tilde{\mathbf{y}}) = \\ &= \frac{l^{(t-1)}(\bar{u}, \tilde{\mathbf{y}}|\tilde{\mathbf{x}}) p^{(t)}(u|\bar{u}, \mathbf{x}^{(t)}) f^{(t)}(\mathbf{y}^{(t)}|u, \mathbf{x}^{(t)}) m^{(t)}(\tilde{\mathbf{y}}|u, \tilde{\mathbf{x}})}{f(\tilde{\mathbf{y}}|\tilde{\mathbf{x}})}, \end{aligned} \quad (11)$$

with $\bar{u}, u = 1, \dots, k$ and for $t = 2, \dots, T$.

In order to efficiently implement the recursions it is useful to express them using matrix notation. In this regard we refer the reader to Bartolucci (2006), Bartolucci and Pennoni (2007), Bartolucci and Farcomeni (2009) and Zucchini and MacDonald (2009) for a detailed description.

6 Likelihood inference

In the presence of individual covariates, the observed data correspond to the vectors $\tilde{\mathbf{x}}_i$ and $\tilde{\mathbf{y}}_i$, for $i = 1, \dots, n$. In particular, $\tilde{\mathbf{x}}_i$ may be decomposed into time-specific subvectors of covariates $\mathbf{x}_i^{(1)}, \dots, \mathbf{x}_i^{(T)}$. Similarly, we recall that $\tilde{\mathbf{y}}_i$ is made of subvectors $\mathbf{y}_i^{(1)}, \dots, \mathbf{y}_i^{(T)}$, that become scalars in the univariate case.

Using the above notation, we have the following expression for the model log-likelihood:

$$\ell(\boldsymbol{\theta}) = \sum_i \log f_{\tilde{\mathbf{Y}}|\tilde{\mathbf{X}}}(\tilde{\mathbf{y}}_i|\tilde{\mathbf{x}}_i),$$

where $\boldsymbol{\theta}$ is the vector of parameters. An equivalent expression that is computationally more convenient is

$$\ell(\boldsymbol{\theta}) = \sum_{\tilde{\mathbf{x}}} \sum_{\tilde{\mathbf{y}}} n_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} \log f_{\tilde{\mathbf{Y}}|\tilde{\mathbf{X}}}(\tilde{\mathbf{y}}|\tilde{\mathbf{x}}),$$

where $n_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}}$ is the joint frequency of the covariate configuration $\tilde{\mathbf{x}}$ and the response configuration $\tilde{\mathbf{y}}$.

The likelihood function can be maximized by an EM algorithm.

6.1 Expectation-Maximization algorithm

The EM algorithm is based on the complete data log-likelihood that, in the univariate categorical case, has expression

$$\begin{aligned} \ell^*(\boldsymbol{\theta}) = & \sum_t \sum_u \sum_{\mathbf{x}} \sum_y a_{u\mathbf{x}y}^{(t)} \log f^{(t)}(y|u, \mathbf{x}) + \\ & + \sum_u \sum_{\mathbf{x}} b_{u\mathbf{x}}^{(1)} \log p(u|\mathbf{x}) + \sum_{t>2} \sum_{\bar{u}} \sum_u \sum_{\mathbf{x}} b_{\bar{u}u\mathbf{x}}^{(t)} \log p^{(t)}(u|\bar{u}, \mathbf{x}), \end{aligned} \quad (12)$$

where $b_{u\mathbf{x}}^{(t)}$ is a frequency of the latent state u and covariate configuration \mathbf{x} at occasion t ; with reference to the same occasion and covariate configuration, $b_{\bar{u}u\mathbf{x}}^{(t)}$ is the number of transitions from state \bar{u} to state u , whereas $a_{u\mathbf{x}y}^{(t)}$ is the number of subjects that are in latent state u and provide response y . Expressions similar to that in (12) result for the complete data log-likelihood in more general cases, such as that of multivariate responses.

Based on the complete data log-likelihood, the two steps of the EM algorithm are as follows:

- **E-step:** it computes the conditional expected value of each frequency involved in (12); these expected values are denoted by $\hat{a}_{u\mathbf{x}y}^{(t)}$, $\hat{b}_{u\mathbf{x}}^{(t)}$, and $\hat{b}_{\bar{u}u\mathbf{x}}^{(t)}$. In particular, we have

$$\begin{aligned} \hat{a}_{u\mathbf{x}y}^{(t)} &= \sum_{\tilde{\mathbf{x}}} \sum_{\tilde{\mathbf{y}}} n_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} q^{(t)}(u|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) I(\mathbf{x}^{(t)} = \mathbf{x}, y^{(t)} = y), \\ \hat{b}_{u\mathbf{x}}^{(t)} &= \sum_{\tilde{\mathbf{x}}} \sum_{\tilde{\mathbf{y}}} n_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} q^{(t)}(u|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) I(\mathbf{x}^{(t)} = \mathbf{x}), \\ \hat{b}_{\bar{u}u\mathbf{x}}^{(t)} &= \sum_{\tilde{\mathbf{x}}} \sum_{\tilde{\mathbf{y}}} n_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} q^{(t)}(\bar{u}, u|\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) I(\mathbf{x}^{(t)} = \mathbf{x}), \end{aligned}$$

where $I(\cdot)$ is the indicator function. These expressions are based on the posterior probabilities introduced in Section 5; see, in particular, expressions (10) and (11).

- **M-step:** it maximizes the complete data log-likelihood expressed as in (12), with each frequency substituted by the corresponding expected value. How to maximize this function depends on the specific formulation of the model and, in particular, on whether the covariates are included in the measurement model or in the latent model. We refer to Bartolucci (2006) and Bartolucci and Farcomeni (2009) for details.

6.1.1 Initialization of the algorithm.

As typically happens for latent variable and mixture models, the likelihood function may be multimodal. In particular, the EM algorithm could converge to a mode of the likelihood which does not correspond to the global maximum. In order to increase the chance of reaching the global maximum, the EM

shall be initialized properly. We give guidelines below on deterministic starting solutions and suggest to compare the value at convergence with values obtained starting from randomly chosen initial values. For a similar multi-start strategy for mixture models see Berchtold (2004).

We illustrate in detail, in the case of univariate categorical responses, the deterministic rule we suggest to use for initializing the EM algorithm. It consists of computing the global logits for the observed distribution of the response variables. In practice, this amounts to computing:

$$\eta_y = \log \frac{\sum_{i=1}^n \sum_{t=1}^T I(y_i^{(t)} \geq y)}{\sum_{i=1}^n \sum_{t=1}^T I(y_i^{(t)} < y)}, \quad y = 1, \dots, c-1.$$

Moreover, once a grid of points ν_1, \dots, ν_k is chosen (we suggest a grid of equally spaced points between $-k$ and k), the conditional response probabilities are initialized as follows:

$$f(y|u, \mathbf{x}) = \phi_{y|u} - \phi_{y+1|u}, \quad y = 0, \dots, c-1,$$

where

$$\phi_{y|u} = \begin{cases} 1 & y = 0, \\ \frac{\exp(\eta_{y-1} + \nu_u)}{1 + \exp(\eta_{y-1} + \nu_u)} & y = 1, \dots, c-1, \\ 0 & y = c, \end{cases}$$

for $u = 1, \dots, k$. Note that $\phi_{y|u}$ are the values of the corresponding survival function, that is $\phi_{y|u} = p(Y^{(t)} \geq y | U^{(t)} = u)$. This rule guaranties that the conditional probabilities $\phi_{y|u}$ sum to 1 and the resulting distribution of the response variables is statistically increasing with u . In the multivariate case, we suggest to apply this rule separately for each response variable, so as to initialize the conditional response probabilities $\phi_{jy|u}$, $j = 1, \dots, r$, $u = 1, \dots, k$, $y = 0, \dots, c_j - 1$.

Finally, regardless of the nature of the response variables, we suggest to use the following starting values for the initial probabilities:

$$p(u|\mathbf{x}) = \frac{1}{k}, \quad u = 1, \dots, k,$$

whereas, for the transition probabilities, we suggest to use

$$p^{(t)}(u|\bar{u}, \mathbf{x}) = \frac{1}{h+k} \begin{cases} h+1, & u = \bar{u}, \\ 1, & u \neq \bar{u}, \end{cases}$$

for $t = 2, \dots, T$, where h is a suitable constant (we generally use $h = 9$ in our applications).

The random starting rule that we propose is based on suitably normalized random numbers drawn from a uniform distribution from 0 and 1. In particular, in the univariate case we first draw each $\phi_{y|u}$, $u = 1, \dots, k$, $y = 0, \dots, c-1$, from this distribution and then we normalize these probabilities so that the constraint $\sum_{y=0}^{c-1} \phi_{y|u} = 1$ is satisfied for $u = 1, \dots, k$. In a similar way, we suggest to choose the initial probabilities $p(u|\mathbf{x})$, $u = 1, \dots, k$, as random

numbers drawn from uniform distribution between 0 and 1 which are suitably normalized. The same may be applied, for $t = 2, \dots, T$ and $\bar{u} = 1, \dots, k$, to draw the transition probabilities $p^{(t)}(u|\bar{u}, \mathbf{x})$, $u = 1, \dots, k$, which must sum to 1.

Inference is finally based on the solution corresponding to the largest value of the likelihood at convergence, which typically corresponds to the global optimum. The other local modes of the likelihood are not of interest for inference. In fact, the random rule allows us to adequately explore the parameter space, when its application is repeated a suitable number of times. Therefore, once an estimate is obtained starting with the deterministic rule, say $\hat{\theta}_0$, we suggest to perform again the algorithm starting from a suitable number R of randomly chosen points of the parameters space, obtaining the estimates $\hat{\theta}_1, \dots, \hat{\theta}_R$. Then, we compare $\ell(\hat{\theta}_0)$ with the maximum of $\ell(\hat{\theta}_1), \dots, \ell(\hat{\theta}_R)$. Provided that R is large enough, if $\ell(\hat{\theta}_0)$ is not smaller than this maximum (up to a negligible tolerance level), then we can be confident that the solution based on the deterministic starting rule corresponds to the global maximum of $\ell(\theta)$. Otherwise, this rule needs to be somehow improved. In any case, we can take, as final estimate of the parameters, denoted by $\hat{\theta}$, the one corresponding to the highest log-likelihood among $\hat{\theta}_0, \dots, \hat{\theta}_R$. Other possible strategies will be described in Section 9.

6.2 Information matrix, local identifiability, and standard errors

The EM algorithm uses neither the observed nor the expected information matrix which are suitable transformations of the second derivative matrix of the (incomplete) log-likelihood. From the inverse of either the observed or the expected information matrix we obtain standard errors for the parameter estimates. Then, from the output of the EM algorithm, we do not have an obvious mean of assessing the precision of the maximum likelihood estimates. This motivated Louis (1982) to develop a procedure for calculating the observed information matrix from the EM algorithm output. The proposed procedure involves expressing this matrix as the difference between two matrices corresponding to the total information, which we would have if we knew the latent states, and the missing information. However, the expression of the latter may be cumbersome to compute; see also Oakes (1999).

Several methods have been proposed to overcome this difficulty. Most of these methods have been developed within the literature on hidden Markov models; for a concise review see Lystig and Hughes (2002). The more interesting methods are based on the information matrix obtained from the EM algorithm by the technique of Louis (1982) or related techniques; see for instance Turner et al (1998) and Bartolucci and Farcomeni (2009). Other interesting methods obtain the information matrix on the basis of the second derivative of the manifest probability of the response variables by a recursion similar to (9); see Lystig and Hughes (2002) and Bartolucci (2006). Among the methods related to the EM algorithm, that proposed by Bartolucci and Farcomeni (2009) is very simple to implement and requires a small extra code over that required

for the maximum likelihood estimation. The method exploits a well-known result according to which the score of the complete data log-likelihood computed at the E-step corresponds to the score of the incomplete data log-likelihood. More precisely, we have

$$\mathbf{s}(\boldsymbol{\theta}) = \frac{\partial \ell(\boldsymbol{\theta})}{\partial \boldsymbol{\theta}} = \left. \frac{\partial E_{\bar{\boldsymbol{\theta}}}[\ell^*(\boldsymbol{\theta})|\mathbf{X}, \mathbf{Y}]}{\partial \boldsymbol{\theta}} \right|_{\bar{\boldsymbol{\theta}}=\boldsymbol{\theta}},$$

where $E_{\bar{\boldsymbol{\theta}}}[\ell^*(\boldsymbol{\theta})|\mathbf{X}, \mathbf{Y}]$ denotes the conditional expected value of the complete-data log-likelihood computed at the parameter value $\bar{\boldsymbol{\theta}}$. Then, the observed information matrix, denoted by $\mathbf{J}(\boldsymbol{\theta})$, is obtained as minus the numerical derivative of $\mathbf{s}(\boldsymbol{\theta})$. The standard error of each parameter estimate is then obtained as the square root of the corresponding diagonal element of $\mathbf{J}(\hat{\boldsymbol{\theta}})^{-1}$. These standard errors may be used for hypothesis testing and for obtaining confidence intervals in the usual way. The information matrix can also be used for checking local identifiability at $\hat{\boldsymbol{\theta}}$. We consider the model to be local identifiable if the matrix $\mathbf{J}(\hat{\boldsymbol{\theta}})$ is of full rank. See also McHugh (1956) and Goodman (1974).

Another possibility is to use the parametric bootstrap procedure (Davison and Hinkley, 1997). Bootstrap naturally takes into account the fact that the parameter space of the LM may be bounded. It is performed by repeatedly generating data from the estimated model and then fitting the model on them. A detailed illustration may be found in Bartolucci et al (2013), Chapter 7.

7 Model choice and path prediction

A fundamental issue concerning the LM model is the choice of the number of latent states, denoted by k . In certain applications, this number is a priori defined by the nature of the problem or the particular interest of the user. In most cases, however, it is selected on the basis of the observed data. For this aim, two main approaches are commonly used, which are based on likelihood ratio testing and on information criteria. The other aspect of interest is the prediction of the sequence of latent states for a given subject on the basis of his/her observable covariates and response variables. We discuss below how to find the maximum a posteriori sequence of latent states for a given subject.

7.1 Model selection

A first possibility is based on performing a likelihood ratio test between the model with k classes and that with $k + 1$ classes for increasing values of k , until the test is not rejected. The adopted test statistic may be expressed as

$$LR = -2(\hat{\ell}_0 - \hat{\ell}_1),$$

where $\hat{\ell}_0$ is the maximum log-likelihood of the smaller model and $\hat{\ell}_1$ is that of the larger model. The problem of this approach is that, in order to obtain

a p -value for LR , we need to use a parametric bootstrap procedure (Feng and McCulloch, 1996) based on a suitable number of samples simulated from the estimated model with k classes. This is because the standard regularity conditions, required to validly use the chi-squared distribution to compute these p -values, are not met in this case. Cheng and Liu (2001) propose a new consistency statistics to be evaluated with the bootstrap procedure for finding the best parsimonious model.

Another approach in order to choose the number of latent states is based on information criteria. The two most common criteria are the Akaike Information Criterion (AIC), see Akaike (1973), and the Bayesian Information Criterion (BIC), see Schwarz (1978). According to first criterion, we choose the number of states corresponding to the minimum of $AIC = -2\ell(\hat{\theta}) + 2g$, where g is the number of free parameters; according to the second, we choose the model with the smallest value of $BIC = -2\ell(\hat{\theta}) + g \log(n)$.

The performance of the two approaches above have been deeply studied in the literature on mixture models; see McLachlan and Peel (2000), Chapter 6. These criteria have also been studied in the hidden Markov literature for time-series, where the two indices above are penalized with a term depending on the number of time occasions; see Boucheron and Gassiat (2007). From these studies, it emerges that BIC is usually preferable to AIC, as the latter tends to overestimate the number of latent states. The use of BIC is also criticized anyway for being too strict in certain cases; see among others Rusakov and Geiger (2005). Theoretical properties of AIC and BIC applied to the LM models are less studied. However, BIC is a commonly accepted model choice criterion even for these models. It has been applied by many authors, such as Langeheine (1994), Langeheine and Van de Pol (1994), and Magidson and Vermunt (2001). In particular, Bartolucci et al (2009) suggest the use of this criterion together with diagnostic statistics measuring the goodness-of-fit and goodness-of-classification, whereas a simulation study may be found in Bartolucci and Farcomeni (2009). A recent study aimed at comparing the performance of some likelihood based criteria and classification based criteria, such as an entropy measure, for selecting the number of latent states of a multivariate LM model may be found in Bacci et al (2013).

Other indices have been proposed in the literature to assess the quality of the classification. Bartolucci et al (2009) proposed to use

$$S = \frac{\sum_{\tilde{\mathbf{x}}} \sum_{\tilde{\mathbf{y}}} n_{\tilde{\mathbf{x}}\tilde{\mathbf{y}}} \sum_{t=1}^T [\hat{f}^{*(t)}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}}) - 1/k]}{(1 - 1/k)nT},$$

where $\hat{f}^{*(t)}(\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$ is the maximum, with respect to u , of the posterior probabilities $f_{U^{(t)}|\tilde{\mathbf{X}}, \tilde{\mathbf{Y}}}^{(t)}(u|\tilde{\mathbf{x}}, \tilde{\mathbf{y}})$. Index S is always between 0 and 1, with 1 corresponding to the situation of absence of uncertainty in the classification, since one of such posterior probabilities is equal to 1 for every $\tilde{\mathbf{y}}$ and t , with all the other probabilities equal to 0.

A problem related to selection of k is the evaluation of hypotheses on the initial and transition probability matrix. Two related issues arise, one linked to

goodness of fit and selection of a constrained formulation, and another linked to hypothesis testing. A key example of the latter is the assumption that the transition matrix is diagonal, which leads to a latent class model in which subject-specific heterogeneity terms are time-fixed. It is shown in Bartolucci (2006) that the likelihood ratio statistic can be used to test hypotheses on the hidden distribution. The likelihood ratio statistic (LR) is once again defined as minus twice the difference in log-likelihood at convergence between the constrained model and the unconstrained model. A boundary problem arises, but it is shown in Bartolucci (2006) that the LR distribution may be approximated in general with a chi-bar-squared, which can be used to compute p -values.

7.2 Path prediction

Once the model has been estimated, a relevant issue is that of path prediction, that is, finding the most likely sequence of latent states for a given subject on the basis of the responses he/she provided. For each subject in the sample, this is the sequence $\hat{\mathbf{u}} = (\hat{u}^{(1)}, \dots, \hat{u}^{(T)})$ maximizing

$$f_{\mathbf{U}|\tilde{\mathbf{X}},\tilde{\mathbf{Y}}}(\mathbf{u}|\tilde{\mathbf{x}},\tilde{\mathbf{y}}). \quad (13)$$

The problem above is different from that of finding the most likely state occupied by a subject at certain occasion, which is found maximizing $f_{U^{(t)}|\tilde{\mathbf{X}},\tilde{\mathbf{Y}}}(u|\tilde{\mathbf{x}},\tilde{\mathbf{y}}) \propto q^{(t)}(u,\tilde{\mathbf{y}})$. This problem may be simply solved on the basis of the posterior probabilities computed at the last iteration of the EM algorithm (see Section 6) but does not correspond to maximizing the joint conditional probability in (13). In order to do so, we employ an efficient algorithm which avoids the evaluation of such a posterior probability at every configuration \mathbf{u} of the latent process. This algorithm is known as the Viterbi algorithm (Viterbi, 1967; Juang and Rabiner, 1991), and is illustrated in the following for the case of univariate responses without covariates.

For a given subject with response configuration \mathbf{y} , let $\hat{r}^{(1)}(u) = f_{U^{(1)}|Y^{(1)}}(u^{(1)},y^{(1)})$ and, for $t = 2, \dots, T$, let

$$\hat{r}^{(t)}(u) = \max_{u^{(1)}, \dots, u^{(t-1)}} f_{U^{(1)}, \dots, U^{(t)}|Y^{(1)}, \dots, Y^{(t)}}(u^{(1)}, \dots, u^{(t-1)}, u, y^{(1)}, \dots, y^{(t)}).$$

A forward recursion can be used to compute the above quantities, and then a backward recursion based on these quantities can then be used for path prediction:

1. for $u = 1, \dots, k$ let $\hat{r}^{(1)}(u) = p(u)\phi_{y^{(1)}|u}$;
2. for $t = 2, \dots, T$ and $v = 1, \dots, k$ compute $\hat{r}^{(t)}(v)$ as

$$\phi_{y^{(t)}|v} \max_u [\hat{r}^{(t-1)}(u)p^{(t)}(u|\bar{u})];$$

3. find the optimal state $\hat{u}^{(T)}$ as $\hat{u}^{(T)} = \arg \max_u \hat{r}^{(T)}(u)$;
4. for $t = T-1, \dots, 1$, find the optimal state $\hat{u}^{(t)}$ as $\hat{u}^{(t)} = \arg \max_u \hat{r}^{(t)}(u)p^{(t+1)}(\hat{u}|\bar{u})$.

All the above quantities are computed on the basis of the ML estimate of the parameter $\boldsymbol{\theta}$ of the model of interest.

8 Bayesian estimation

In the Bayesian framework, parameters are random variables and information is summarized by approximating the posterior distribution, which is proportional to the prior distribution and the likelihood.

An advantage of Bayesian inference is that one can easily include prior information, if available. Furthermore, there are computational advantages since there usually are less problems from multimodality when approximating the posterior with respect to numerical maximization of the likelihood. For a general introduction on Bayesian inference see for instance Bernardo and Smith (1994). There are different approaches to Bayesian estimation for hidden Markov models (e.g., Robert et al, 2000; Cappé et al, 2005; Zucchini and MacDonald, 2009; Spezia, 2010). Many of these approaches treat also k as an unknown, and efficiently derive a posterior on this quantity through transdimensional algorithms, such as the reversible jump algorithm (Green, 1995). Simply removing the transdimensional steps yields algorithms for fixed k .

In the Bayesian framework a first issue regards the choice of prior inputs. When prior information is available, it should be summarized in the prior distributions. When prior information is not available, one can use a Dirichlet with parameter \mathbf{e} , where \mathbf{e} denotes a vector of ones, for the initial probabilities and independently for each row of the transition matrix. This prior has been suggested as a universal default prior for multinomial parameters by Tuyl et al (2009). When covariates are included in the model, a natural prior distribution for the regression coefficients (and latent intercepts) is a zero-centered multivariate Gaussian, with covariance matrix either diagonal or proportional to the hat matrix. Zero-centered multivariate Gaussians can be also used for the log-odds parameters (Leonard, 1975; Nazaret, 1987) in the multivariate case. The posterior distribution is seldom available in closed form, and shall be approximated through a Markov chain Monte Carlo (MCMC) sampling scheme (Robert and Casella, 2010). We outline here an efficient Gibbs sampler for the basic LM model, in which there are no covariates and the conditional response probabilities are time-homogeneous. Our approach is based on the augmented likelihood (i.e., the complete likelihood), is adapted from Chib (1996), and to the best of our knowledge has not been considered previously in the LM literature. While simpler approaches work with time-series in the HMM context, in the LM context where the number of occasions is much lower, efficient algorithms are necessary. For reasons of space we restrict to the basic LM model without covariates for fixed k , but the approach can be naturally extended to more general cases.

Our Gibbs sampling scheme proceeds by iterating the following steps. First of all, the latent indicators \mathbf{u} are updated through an *ff-bs* algorithm (Chib, 1996). In practice the usual forward-backward recursions are used to compute $q^{(t)}(u, \tilde{\mathbf{y}})$. We then compute $f_{U^{(t)}|Y}(u|\mathbf{Y}) \propto q^{(t)}(u, \tilde{\mathbf{y}})$, and sample latent indicators from this distribution. The sampled latent indicators are used to compute $b_u^{(t)}$ and $b_{\bar{u}u}^{(t)}$ for the current iteration, that is, the frequency of each latent state at each time and the frequency of each of the possible k^2 tran-

sitions at each time, for $t = 2, \dots, T$. The initial probabilities shall then be sampled from a Dirichlet with parameter $\mathbf{e} + \mathbf{b}^{(1)}$. Assuming a time constant transition matrix, we then have that the \bar{u} -th row can be sampled from a Dirichlet with parameter $\mathbf{e} + \sum_{t=2}^T \mathbf{b}_{\bar{u}}^{(t)}$, with $\mathbf{b}_{\bar{u}}^{(t)} = (b_{\bar{u}1}^{(t)}, \dots, b_{\bar{u}k}^{(t)})$.

At the end of each iteration, label switching is tackled by sorting conditional probabilities of a baseline category in increasing order. By fixing an order *a priori*, we are implicitly choosing one of the $k!$ modes of the likelihood that could be visited by the algorithm, and excluding all others. Frühwirth-Schnatter (2001) discussed in detail the advantages and disadvantages of this strategy, and outlined alternative strategies which can be directly embedded in our MCMC algorithm.

Convergence of the MCMC sample sequence to the posterior distribution is guaranteed only when the number of iterations diverges. Therefore, one key issue regards convergence diagnosis, that is, a formal or informal assessment that the sequence sampled approximates the posterior in a satisfactory way, and therefore that sampling can be stopped. A general review of convergence diagnostics can be found in Cowles and Carlin (1996), while tools specifically derived for LM models can be found, for instance, in Robert et al (1999) and Scott (2002). The basic procedures often involve parallel sampling of multiple chains from different starting points and/or time-series analysis. Techniques of the first kind try to assess that the chains have converged to the same distribution, while time-series analysis involves checking that the sequence does not have trends and that the dependence is not strong. Autocorrelation Functions (ACF) are usually computed and visually inspected to this end. Robert et al (1999) and Scott (2002) suggested monitoring the log-posterior, using for instance the ACF. Formal tools in Robert et al (1999) involve non-parametric tests. For instance, subsamples of the MCMC chain can be compared to check that they have the same distribution using a Kolmogorov-Smirnov test, and convergence to normality of any pre-specified functional can be assessed similarly.

Model parameters are often estimated on the basis of posterior summaries, that is, by estimation of the posterior means. The posterior means are expected to be close to the maximum likelihood estimate when the sample size is large enough. A final issue regards the inference on k when a transdimensional algorithm is not used. The number of latent states can be estimated with the parallel sampling approach of Congdon (2006), which in practice involves fitting the model repeatedly for all values of k in a range of plausible values. A posterior distribution on k is then approximated by post-processing the MCMC output, given certain assumptions of prior independence. A simpler possibility is to use the Schwarz criterion (Schwarz (1978); Kass and Raftery (1995), see also Leroux and Puterman (1992)), which is analogous to BIC. The Schwarz criterion corresponds to a uniform prior on k and a multivariate normal approximation for the posterior of the remaining parameters, integrated over with a Laplace's approximation. A further approach is the Deviance Information Criterion (DIC), proposed in Spiegelhalter et al (2002). DIC is a Bayesian generalization of the Akaike information criterion. It can be directly

estimated from the MCMC output (see for instance Farcomeni and Arima, 2012) and can be seen as a Bayesian measure of fit and prediction accuracy.

9 Conclusions and further developments

In this paper we presented a review of the latent Markov model (LM). The most general formulation is given by a model for multivariate data with covariates. Then, we illustrated several constrained versions of the general LM model, and discussed maximum likelihood estimation and related inferential issues. We tried to keep the presentation simple and provided references to methods which are of simple implementation, taking ideas from methods widely applied in the literature on hidden Markov models for time-series data. We pointed out that longitudinal data are similar to time-series data but they are referred to several units observed at few time occasions. Longitudinal data are widespread in many research fields. Moreover, in order to keep the presentation simple, we dealt with the case of balanced data in which all subjects are observed at the same number of occasions, the same response variables are obtained at every occasion, and there are no missing responses.

The potential of the LM approach was illustrated by summarizing the results of several applications available in the literature. It is worth noting that the number of applications available in the literature is ever growing; recent applications also involve situations which are not typically those of longitudinal data analysis. Examples concern the estimation of closed-populations on the basis of capture-recapture data (Bartolucci and Pennoni, 2007) and the assessment of the peer review process of grant applications (Bornmann et al, 2008). Recent developments include also semiparametric (Dannemann, 2012) and nonparametric (Yau et al, 2011) approaches. These approaches focus on a continuous outcome, and estimate an LM model without distributional assumptions on the manifest distribution. Here, we summarize additional developments of the LM models that, in our opinion, deserve particular attention:

- *More flexible parametrizations for the conditional distribution of the response variables given the latent process:* We refer in particular to parametrizations which address the problem of multidimensionality when the model is applied, for instance, to data coming from the administration of test items assessing different types of ability. This is a problem well known in the Item Response Theory (Hambleton and Swaminathan, 1985), which may concern different fields of applications, such as psychology and assessment of the quality of life. Flexible parametrizations are also of interest when the response variables are of mixed nature; for instance, we may deal with a mixture of binary, ordinal, and continuous variables.
- *Presence of local maxima:* The likelihood may be multimodal. This is a well known problem also in the mixture models (McLachlan and Peel, 2000) and non-normal mixture models (Seidel and Ševčíková, 2004) literatures. This problem is particularly important for smaller sample sizes. We suggested

a multistart strategy, but formal assessment procedures are yet not available in the literature. Further, other estimation methods may be devised which at least increase the likelihood of obtaining the MLE and avoid spurious solutions. A possibility is to explore the robust estimation literature to this end. For instance, García-Escudero et al (2013) combine trimmed maximum likelihood methods and restrictions.

- *Higher-order Markov chains*: the assumptions that the latent Markov chain is of first-order may be restrictive when there is a “memory effect”, so that the latent state of a subject at a given occasion depends on two or more previous occasions. The extension to second or higher-order Markov chains may be of interest. This extension is rather simple to implement and has been already applied in certain contexts. However, to our knowledge, a systematic illustration of this higher-order approach does not seem to be available in the literature; and very few papers work on LM models with parsimonious representations for higher-order latent chains. Two exceptions are Bartolucci and Farcomeni (2010) and Farcomeni (2011).
- *Missing responses*: It is well known that longitudinal datasets typically suffer from the problem of missing responses due to several reasons. In applying an LM model, a trivial solution to this problem consists in eliminating from the sample subjects with at least one missing response. Obviously, this complete data solution may lead to a strong bias of the parameter estimates when the assumption that these responses are missing at random is not plausible. The interest here is in formulating an LM approach in which the event of missing response is explicitly modeled given the latent state and in which the evolution of the latent process takes into account this possibility by adopting a suitable parametrization of the transition matrices. In principle, another possibility is given by multiple imputation within the EM algorithm. In the Bayesian framework this is straightforward (Farcomeni and Arima, 2012).
- *Time-dependent covariates and causality*: We have made the assumption in this paper that covariates are strictly exogenous. There are cases in which time-dependent covariates may be influenced by other covariates not included and/or from the outcomes at previous occasions. This case of *time-dependent confounding* can possibly lead to biased estimates. A modification of the EM algorithm to account for the underlying causal structure would lead to unbiased estimates and is grounds for further work.

Acknowledgements F. Bartolucci and F. Pennoni acknowledge the financial support from the grant “Finite mixture and latent variable models for causal inference and analysis of socio-economic data” (FIRB - Futuro in ricerca) funded by the Italian Government (RBF12SHVV).

References

Akaike H (1973) Information theory as an extension of the maximum likelihood principle. In: Petrov BN, F C (eds) Second International symposium on

- information theory, Akademiai Kiado, Budapest, pp 267–281
- Altman RM (2007) Mixed hidden Markov models: an extension of the hidden Markov model to the longitudinal data setting. *Journal of the American Statistical Association* 102:201–210
- Anderson TW (1951) Probability models for analysing time changes in attitudes. In: Corporation TR (ed) *The use of mathematical models in the measurement of the attitudes*, Lazarsfeld P. F.
- Anderson TW (1954) Probability models for analysing time changes in attitudes. In: F LP (ed) *Mathematical Thinking in the Social Science*, The Free press
- Andersson S, Rydén T (2009) Subspace estimation and prediction methods for hidden Markov models. *The Annals of Statistics* 37:4131–4152
- Archer GEB, Titterton DM (2009) Parameter estimation for hidden Markov chains. *Journal of Statistical Planning and Inference* 108:365–390
- Bacci S, Pandolfi S, Pennoni F (2013) A comparison of some criteria for states selection in the latent Markov model for longitudinal data. *Advances in Data Analysis and Classification* p in press
- Bartolucci F (2006) Likelihood inference for a class of latent Markov models under linear hypotheses on the transition probabilities. *Journal of the Royal Statistical Society, Series B* 68:155–178
- Bartolucci F, Farcomeni A (2009) A multivariate extension of the dynamic logit model for longitudinal data based on a latent Markov heterogeneity structure. *Journal of the American Statistical Association* 104:816–831
- Bartolucci F, Farcomeni A (2010) A note on the mixture transition distribution and hidden Markov models. *Journal of Time Series Analysis* 31:132–138
- Bartolucci F, Pandolfi S (2013) A new constant memory recursion for hidden Markov models. *Journal of Computational Biology* in press
- Bartolucci F, Pennoni F (2007) A class of latent Markov models for capture-recapture data allowing for time, heterogeneity and behavior effects. *Biometrics* 63:568–578
- Bartolucci F, Pennoni F, Francis B (2007) A latent Markov model for detecting patterns of criminal activity. *Journal of the Royal Statistical Society, Series A* 170:151–132
- Bartolucci F, Lupparelli M, Montanari GE (2009) Latent Markov model for binary longitudinal data: an application to the performance evaluation of nursing homes. *Annals of Applied Statistics* 3:611–636
- Bartolucci F, Pennoni F, Vittadini G (2011) Assessment of school performance through a multilevel latent Markov Rasch model. *Journal of Educational and Behavioural Statistics* 36:491–522
- Bartolucci F, Farcomeni A, Pennoni F (2013) *Latent Markov Models for Longitudinal Data*. Chapman and Hall/CRC press, Boca Raton
- Baum L, Petrie T (1966) Statistical inference for probabilistic functions of finite state Markov chains. *Annals of Mathematical Statistics* 37:1554–1563
- Baum L, Petrie T, Soules G, Weiss N (1970) A maximization technique occurring in the statistical analysis of probabilistic functions of Markov chains. *Annals of Mathematical Statistics* 41:164–171

- Berchtold A (2004) Optimization of mixture models: Comparison of different strategies. *Computational statistics* 19:385–406
- Bernardo JM, Smith AFM (1994) *Bayesian Theory*. Wiley, Chichester
- Bickel PJ, Ritov Y, Rydén T (1998) Asymptotic normality of the maximum-likelihood estimator for general hidden Markov models. *Annals of Statistics* 26:1614–1635
- Bollen KA, Curran PJ (2006) *Latent curve models: A structural equation perspective*. Wiley, Hoboken, NJ
- Bornmann L, Mutz R, Daniel HD (2008) Latent Markov modeling applied to grand peer review. *Journal of Informetrics* 2:217–228
- Boucheron S, Gassiat E (2007) An information-theoretic perspective on order estimation. In: O Cappé TR E Moulines (ed) *Inference in Hidden Markov Models*, Springer, pp 565–602
- Bye BV, Schechter ES (1986) A latent Markov model approach to the estimation of response error in multiwave panel data. *Journal of the American Statistical Association* 81:375–380
- Cappé O, Moulines E, Rydén T (2005) *Inference in Hidden Markov Models*. Springer-Verlag, New York
- Cheng RCH, Liu WB (2001) The consistency of estimators in finite mixture models. *Scandinavian Journal of Statistics* 28:603–616
- Chib S (1996) Calculating posterior distributions and modal estimates in Markov mixture models. *Journal of Econometrics* 75:79–97
- Collins LM, Wugalter SE (1992) Latent class models for stage-sequential dynamic latent variables. *Multivariate Behavioral Research* 27:131–157
- Colombi R, Forcina A (2001) Marginal regression models for the analysis of positive association of ordinal response variables. *Biometrika* 88:1007–1019
- Congdon P (2006) Bayesian model choice based on Monte Carlo estimates of posterior model probabilities. *Computational Statistics & Data Analysis* 50:346–357
- Cowles MK, Carlin BP (1996) Markov chain Monte Carlo convergence diagnostics: a comparative review. *Journal of the American Statistical Association* 91:883–904
- Dannemann J (2012) Semiparametric hidden Markov models. *Journal of computational and graphical statistics* 21:677–692
- Davison AC, Hinkley DV (1997) *Bootstrap Methods and their Application*. Cambridge University Press, Cambridge, MA
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the em algorithm (with discussion). *Journal of the Royal Statistical Society, Series B* 39:1–38
- Dias JG, Vermunt JK (2007) Latent class modeling of website users' search patterns: Implications for online market segmentation. *Journal of Retailing and Consumer Services* 14:359–368
- Elliot DS, Huizinga D, Menard S (1989) *Multiple Problem Youth: Delinquency, Substance Use, and Mental Health Problems*. Springer-Verlag, New York
- Farcomeni A (2011) Hidden Markov partition models. *Statistics & Probability Letters* 81:1766–1770

- Farcomeni A (2012) Quantile regression for longitudinal data based on latent Markov subject-specific parameters. *Statistics and Computing* 22:141–152
- Farcomeni A, Arima S (2012) A Bayesian autoregressive three-state hidden Markov model for identifying switching monotonic regimes in Microarray time course data. *Statistical Applications in Genetics and Molecular Biology* 11(4):article 3
- Feng Z, McCulloch CE (1996) Using bootstrap likelihood ratios in finite mixture models. *J R Statist Soc* 58:609–617
- Fitzmaurice G, Davidian M, Verbeke G, G M (eds) (2009) *Longitudinal data analysis*. Chapman and Hall, CRC, London
- Frühwirth-Schnatter S (2001) Markov chain Monte Carlo estimation of classical and dynamic switching and mixture models. *Journal of the American Statistical Association* 96:194–209
- García-Escudero L, Gordaliza A, Mayo-Isacar A (2013) A constrained robust proposal for mixture modeling avoiding spurious solutions. *Advances in Data Analysis and Classification* pp 1–17
- Ghahramani Z, Jordan MI (1997) Factorial hidden Markov models. *Machine Learning* 29:245–273
- Glonek GFV, McCullagh P (1995) Multivariate logistic models. *Journal of the Royal Statistical Society B* 57:533–546
- Goodman LA (1961) Statistical methods for the mover-stayer model. *Journal of the American Statistical Association* 56:841–868
- Goodman LA (1974) Exploratory latent structure analysis using both identifiable and unidentifiable models. *Biometrika* 61:215–231
- Green PJ (1995) Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika* 82:711–732
- Hambleton RK, Swaminathan H (1985) *Item Response Theory: Principles and Applications*. Kluwer Nijhoff, Boston
- Hoffmann L, Lehrke M, Todt E (1985) Development and changes in pupils' interest in physics (grade 5 to 10): Design of a longitudinal study. In: Lehrke M, Hoffmann L, Gardner PL (eds) *Interest in Science and Technology Education*, IPN, Kiel, pp 71–80
- Juang B, Rabiner L (1991) Hidden Markov models for speech recognition. *Technometrics* 33:251–272
- Kaplan D (2008) An overview of Markov chain methods for the study of stage-sequential developmental processes. *Developmental Psychology* 44:457–467
- Kass RE, Raftery AE (1995) Bayes factors. *Journal of the American Statistical Association* 90:773–795
- Khreich W, Granger E, Miri A, Sabourin R (2010) On the memory complexity of the forward-backward algorithm. *Pattern Recognition Letters* 31:91–99
- Koski T (2001) *Hidden Markov Models for Bioinformatics*. Kluwer, Dordrecht
- Künsch HR (2005) State space and hidden Markov models. In: Barndorff-Nielsen OE, Cox DR, C Klüppelberg e (eds) *Complex Stochastic Systems*, Chapman and Hall/CRC, Boca Raton, FL, pp 109–173
- Langeheine R (1988) New development in latent class theory. In: Langeheine R, J R (eds) *Latent trait and latent class models*, New York: Plenum Press,

- pp 77–108
- Langeheine R (1994) Latent variables Markov models. In: von Eye A, Clogg C (eds) *Latent variables analysis: Applications for developmental research*, Sage, Thousand Oaks, CA, pp 373–395
- Langeheine R, Van de Pol F (1994) Discrete-time mixed Markov latent class models. In: Dale A, Davies R (eds) *Analyzing Social and Political Change: a Casebook of Methods*, Sage Publications, London, pp 171–197
- Lazarsfeld PF (1950) The logical and mathematical foundation of latent structure analysis. In: S A Stouffer EAS L Guttman (ed) *Measurement and Prediction*, Princeton University Press, New York
- Lazarsfeld PF, Henry NW (1968) *Latent Structure Analysis*. Houghton Mifflin, Boston
- Leonard T (1975) Bayesian estimation methods for two-way contingency tables. *Journal of the Royal Statistical Society, Series B* 37:23–37
- Leroux BG, Puterman ML (1992) Maximum-penalized-likelihood estimation for independent and Markov-dependent mixture models. *Biometrics* 48:545–558
- Levinson SE, Rabiner LR, Sondhi MM (1983) An introduction to the application of the theory of probabilistic functions of a Markov process to automatic speech recognition. *Bell System Technical Journal* 62:1035–1074
- Louis T (1982) Finding the observed information matrix when using the EM algorithm. *Journal of the Royal Statistical Society, Series B* 44:226–233
- Lystig TC, Hughes J (2002) Exact computation of the observed information matrix for hidden Markov models. *Journal of Computational and Graphical Statistics* 11:678–689
- MacDonald IL, Zucchini W (1997) *Hidden Markov and other Models for Discrete-Valued Time Series*. Chapman and Hall, London
- Magidson J, Vermunt JK (2001) Latent class factor and cluster models, biplots and related graphical displays. *Sociological Methodology* 31:223–264
- Maruotti A (2011) Mixed hidden Markov models for longitudinal data: An overview. *International Statistical Review* 79:427–454
- McCullagh P, Nelder JA (1989) *Generalized Linear Models*, 2nd Edition. Chapman and Hall, CRC, London
- McHugh RB (1956) Efficient estimation and local identification in latent class analysis. *Psychometrika* 21:331–347
- McLachlan G, Peel D (2000) *Finite Mixture Models*. Wiley
- Muthén B (2004) Latent variable analysis: Growth mixture modeling and related techniques for longitudinal data growth mixture modeling and related techniques for longitudinal data. In: Kaplan D (ed) *Handbook of quantitative methodology for the social sciences*, Sage Publications, Newbury Park, CA, pp 345–368
- Nagin D (1999) Analyzing developmental trajectories: A semi-parametric, group-based approach. *Psychological Methods* 4:139–157
- Nazaret W (1987) Bayesian log-linear estimates for three-way contingency tables. *Biometrika* 74:401–410

- Oakes D (1999) Direct calculation of the information matrix via the EM algorithm. *Journal of the Royal Statistical Society, Series B* 61:479–482
- Paas LJ, Vermunt JK, Bijmolt THA (2009) Discrete time, discrete state latent Markov modelling for assessing and predicting household acquisitions of financial products. *Journal of the Royal Statistical Society, Series A* 170:955–974
- van de Pol F, Langeheine R (1990) Mixed Markov latent class models. *Sociological Methodology* 20:213–247
- Rijmen F, Vansteelandt K, De Boeck P (2007) Latent class models for diary methods data: parameter estimation by local computations. *Psychometrika* 73:167–182
- Robert C, Ryden T, Titterton D (2000) Bayesian inference in hidden Markov models through the reversible jump Markov chain Monte Carlo method. *Journal of the Royal Statistical Society, Series B* 62:57–75
- Robert CP, Casella G (2010) *Monte Carlo Statistical Methods*, 2nd Edition. Springer-Verlag, New York
- Robert CP, Ryden T, Titterton D (1999) Convergence controls for MCMC algorithms, with applications to hidden Markov chains. *Journal of Statistical Computation and Simulation* 64:327–355
- Roeder K, Lynch KG, Nagin DS (1999) Modeling uncertainty in latent class membership: a case study in criminology. *Journal of the American Statistical Association* 94:766–776
- Rost J (2002) Mixed and latent Markov models as item response models. *Methods of Psychological Research Online*, Special Issue pp 53–70
- Rusakov D, Geiger D (2005) Asymptotic model selection for naive Bayesian networks. *Journal of Machine Learning Research* 6
- Schwarz G (1978) Estimating the dimension of a model. *Annals of Statistics* 6:461–464
- Scott SL (2002) Bayesian methods for hidden Markov models: Recursive computing in the 21st century. *Journal of the American Statistical Association* 97:337–351
- Seidel W, Ševčíková H (2004) Types of likelihood maxima in mixture models and their implication on the performance of tests. *Annals of the Institute of Statistical Mathematics* 56:631–654
- Spezia L (2010) Bayesian analysis of multivariate Gaussian hidden Markov models with an unknown number of regimes. *Journal of Time Series Analysis* 31:1–11
- Spiegelhalter DJ, Best NG, Carlin BP, van der Linde A (2002) Bayesian measures of model complexity and fit (with discussion). *Journal of the Royal Statistical Society, Series B* 64:583–639
- Turner R (2008) Direct maximization of the likelihood of a hidden Markov model. *Computational Statistics and Data Analysis* 52:4147–4160
- Turner TR, Cameron MA, Thomson PJ (1998) Hidden Markov chains in generalized linear models. *The Canadian Journal of Statistics* 26:107–125
- Tuyl F, Gerlach R, Mengersen K (2009) Posterior predictive arguments in favor of the Bayes-Laplace prior as the consensus prior for binomial and

- multinomial parameters. *Bayesian Analysis* 4:151–158
- Vansteelandt K, Rijmen F, Pieters G, Vanderlinden J (2007) Drive for thinness, affect regulation and physical activity in eating disorders: a daily life study. *Behaviour Research and Therapy* 45:1717–1734
- Vermunt J (2010) Longitudinal research with latent variables. In: van Montfort K, Oud J, Satorra A (eds) *Handbook of Advanced Multilevel Analysis*, Springer, Heidelberg, Germany, pp 119–152
- Vermunt JK, Langeheine R, Böckenholt U (1999) Discrete-time discrete-state latent Markov models with time-constant and time-varying covariates. *Journal of Educational and Behavioral Statistics* 24:179–207
- Viterbi A (1967) Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Transactions on Information Theory* 13:260–269
- Welch LR (2003) Hidden Markov models and the Baum-Welch algorithm. *IEEE Information Theory Society Newsletter* 53:1–13
- Wiggins L (1973) *Panel Analysis: Latent probability models for attitude and behaviour processes*. Elsevier, Amsterdam
- Wiggins LM (1955) *Mathematical models for the analysis of multi-wave panels*. PhD thesis, Columbia University, Ann Arbor
- Yau C, Papaspiliopoulos O, Roberts G, Holmes C (2011) Bayesian nonparametric hidden Markov models with application to the analysis of copy-number-variation in mammalian genomes. *Journal of the Royal Statistical Society, Series B* 73:37–57
- Zucchini W, MacDonald IL (2009) *Hidden Markov Models for time series: an introduction using R*. Springer-Verlag, New York