

Mixture Hidden Markov Models for Sequence Data: The seqHMM Package in R

Satu Helske

Linköping University, Sweden
University of Oxford, UK
University of Jyväskylä, Finland

Jouni Helske

Linköping University, Sweden
University of Jyväskylä, Finland

Abstract

Sequence analysis is being more and more widely used for the analysis of social sequences and other multivariate categorical time series data. However, it is often complex to describe, visualize, and compare large sequence data, especially when there are multiple parallel sequences per subject. Hidden (latent) Markov models (HMMs) are able to detect underlying latent structures and they can be used in various longitudinal settings: to account for measurement error, to detect unobservable states, or to compress information across several types of observations. Extending to mixture hidden Markov models (MHMMs) allows clustering data into homogeneous subsets, with or without external covariates.

The **seqHMM** package in R is designed for the efficient modeling of sequences and other categorical time series data containing one or multiple subjects with one or multiple interdependent sequences using HMMs and MHMMs. Also other restricted variants of the MHMM can be fitted, e.g., latent class models, Markov models, mixture Markov models, or even ordinary multinomial regression models with suitable parameterization of the HMM.

Good graphical presentations of data and models are useful during the whole analysis process from the first glimpse at the data to model fitting and presentation of results. The package provides easy options for plotting parallel sequence data, and proposes visualizing HMMs as directed graphs.

Keywords: multichannel sequences, categorical time series, visualizing sequence data, visualizing models, latent Markov models, latent class models, R.

Vignette based on the corresponding paper at *Journal of Statistical Software* (Helske and Helske 2019).

1. Introduction

Social sequence analysis is being more and more widely used for the analysis of longitudinal data consisting of multiple independent subjects with one or multiple interdependent sequences (channels). Sequence analysis is used for computing the (dis)similarities of sequences, and often the goal is to find patterns in data using cluster analysis. However, describing, visualizing, and comparing large sequence data is often complex, especially in the case of multiple channels. Hidden (latent) Markov models (HMMs) can be used to compress and visualize information in such data. These models are able to detect underlying latent

structures. Extending to mixture hidden Markov models (MHMMs) allows clustering via latent classes, possibly with additional covariate information. One of the major benefits of using hidden Markov modeling is that all stages of analysis are performed, evaluated, and compared in a probabilistic framework.

The **seqHMM** (Helske and Helske 2019) package for R (R Core Team 2018) is designed for modeling sequence data and other categorical time series with one or multiple subjects and one or multiple channels using HMMs and MHMMs. The package provides functions for the estimation and inference of models, as well as functions for the easy visualization of multichannel sequences and HMMs. Even though the package was originally developed for researchers familiar with social sequence analysis and the examples are related to life course, knowledge on sequence analysis or social sciences is not necessary for the usage of **seqHMM**. The package is available on Comprehensive R Archive Repository (CRAN) and easily installed via `install.packages("seqHMM")`. Development versions can be obtained from GitHub¹.

There are also other R packages in CRAN for HMM analysis of categorical data. The **HMM** package (Himmelmann 2010) is a compact package designed for fitting an HMM for a single observation sequence. The **hmm.discnp** package (Turner and Liu 2014) can handle multiple observation sequences with possibly varying lengths. For modeling continuous-time processes as hidden Markov models, the **msm** package (Jackson 2011) is available. Both **hmm.discnp** and **msm** support only single-channel observations. The **depmixS4** package (Visser and Speekenbrink 2010) is able to fit HMMs for multiple interdependent time series (with continuous or categorical values), but for one subject only. In the **msm** and **depmixS4** packages, covariates can be added for initial and transition probabilities. The **mhsmm** package (O’Connell and Højsgaard 2011) allows modeling of multiple sequences using hidden Markov and semi-Markov models. There are no ready-made options for modeling categorical data, but users can write their own extensions for arbitrary distributions. The **LMest** package (Bartolucci and Pandolfi 2015) is aimed to panel data with a large number of subjects and a small number of time points. It can be used for hidden Markov modeling of multivariate and multichannel categorical data, using covariates in emission and transition processes. **LMest** also supports mixed latent Markov models, where the latent process is allowed to vary in different latent subpopulations. This differs from mixture hidden Markov models used in **seqHMM**, where also the emission probabilities vary between groups. The **seqHMM** package also supports covariates in explaining group memberships. A drawback in the **LMest** package is that the user cannot define initial values or zero constraints for model parameters, and thus important special cases such as left-to-right models cannot be used.

We start with describing data and methods: a short introduction to sequence data and sequence analysis, then the theory of hidden Markov models for such data, an expansion to mixture hidden Markov models and a glance at some special cases, and then some propositions on visualizing multichannel sequence data and hidden Markov models. After the theoretic part we take a look at features of the **seqHMM** package and at the end show an example on using the package for the analysis of life course data. The appendix shows the list of notations.

¹<https://github.com/helske/seqHMM>

2. Methods

2.1. Sequences and sequence analysis

By the term *sequence* we refer to an ordered set of categorical states. It can be a time series, such as a career trajectory or residential history, or any other series with ordered categorical observations, e.g., a DNA sequence or a structure of a story. Typically, *sequence data* consist of multiple independent subjects (*multivariate* data). Sometimes there are also multiple interdependent sequences per subject, often referred to as *multichannel* or multidimensional sequence data.

As an example we use the `biofam` data available in the **TraMineR** package (Gabadinho, Ritschard, Müller, and Studer 2011). It is a sample of 2000 individuals born in 1909–1972, constructed from the Swiss Household Panel survey in 2002 (Müller, Studer, and Ritschard 2007). The data set contains sequences of annual family life statuses from age 15 to 30. Eight observed states are defined from the combination of five basic states: living with parents, left home, married, having children, and divorced. To show a more complex example, we split the original data into three separate *channels* representing different life domains: marriage, parenthood, and residence. The data for each individual now includes three parallel sequences constituting of two or three *states* each: single/married/divorced, childless/parent, and living with parents / having left home.

Sequence analysis (SA), as defined in the social science framework, is a model-free data-driven approach to the analysis of successions of states. The approach has roots in bioinformatics and computer science (see e.g. Durbin, Eddy, Krogh, and Mitchison 1998), but during the past few decades SA has also become more common in other disciplines for the analysis of longitudinal data. In social sciences SA has been used increasingly often and is now “central to the life-course perspective” (Blanchard, Bühlmann, and Gauthier 2014).

SA is used for computing (dis)similarities of sequences. The most well-known method is optimal matching (McVicar and Anyadike-Danes 2002), but several alternatives exist (see e.g. Aisenbrey and Fasang 2010; Elzinga and Studer 2014; Gauthier, Widmer, Bucher, and Notredame 2009; Halpin 2010; Hollister 2009; Lesnard 2010). Also a method for analyzing multichannel data has been developed (Gauthier, Widmer, Bucher, and Notredame 2010). Often the goal in SA is to find typical and atypical patterns in trajectories using cluster analysis, but any approach suitable for compressing information on the dissimilarities can be used. The data are usually presented also graphically in some way. So far the **TraMineR** package has been the most extensive and frequently used software for social sequence analysis.

2.2. Hidden Markov models

In the context of hidden Markov models, sequence data consists of *observed states*, which are regarded as probabilistic functions of *hidden states*. Hidden states cannot be observed directly, but only through the sequence(s) of observations, since they emit the observations on varying probabilities. A discrete first order hidden Markov model for a single sequence is characterized by the following:

- *Observed state sequence* $\mathbf{y} = (y_1, y_2, \dots, y_T)$ with observed states $m \in \{1, \dots, M\}$.
- *Hidden state sequence* $\mathbf{z} = (z_1, z_2, \dots, z_T)$ with hidden states $s \in \{1, \dots, S\}$.

- *Initial probability vector* $\pi = \{\pi_s\}$ of length S , where π_s is the probability of starting from the hidden state s :

$$\pi_s = P(z_1 = s); \quad s \in \{1, \dots, S\}.$$

- *Transition matrix* $A = \{a_{sr}\}$ of size $S \times S$, where a_{sr} is the probability of moving from the hidden state s at time $t - 1$ to the hidden state r at time t :

$$a_{sr} = P(z_t = r | z_{t-1} = s); \quad s, r \in \{1, \dots, S\}.$$

We only consider homogeneous HMMs, where the transition probabilities a_{sr} are constant over time.

- *Emission matrix* $B = \{b_s(m)\}$ of size $S \times M$, where $b_s(m)$ is the probability of the hidden state s emitting the observed state m :

$$b_s(m) = P(y_t = m | z_t = s); \quad s \in \{1, \dots, S\}, m \in \{1, \dots, M\}.$$

The (first order) Markov assumption states that the hidden state transition probability at time t only depends on the hidden state at the previous time point $t - 1$:

$$P(z_t | z_{t-1}, \dots, z_1) = P(z_t | z_{t-1}). \quad (1)$$

Also, the observation at time t is only dependent on the current hidden state, not on previous hidden states or observations:

$$P(y_t | y_{t-1}, \dots, y_1, z_t, \dots, z_1) = P(y_t | z_t). \quad (2)$$

For a more detailed description of hidden Markov models, see e.g., [Rabiner \(1989\)](#), [MacDonald and Zucchini \(1997\)](#), and [Durbin *et al.* \(1998\)](#).

HMM for multiple sequences

We can also fit the same HMM for multiple subjects; instead of one observed sequence \mathbf{y} we have N sequences as $Y = (\mathbf{y}_1, \dots, \mathbf{y}_N)^\top$, where the observations $\mathbf{y}_i = (y_{i1}, \dots, y_{iT})$ of each subject i take values in the observed state space. Observed sequences are assumed to be mutually independent given the hidden states. The observations are assumed to be generated by the same model, but each subject has its own hidden state sequence.

HMM for multichannel sequences

In the case of multichannel sequence data, such as the example described in Section 2.1, for each subject i there are C parallel sequences. Observations are now of the form y_{itc} , $i = 1, \dots, N$; $t = 1, \dots, T$; $c = 1, \dots, C$, so that our complete data is $Y = \{Y^1, \dots, Y^C\}$. In **seqHMM**, multichannel data are handled as a list of C data frames of size $N \times T$. We also define Y_i as all the observations corresponding to subject i .

We apply the same latent structure for all channels. In such a case the model has one transition matrix A but several emission matrices B_1, \dots, B_C , one for each channel. We assume that the observed states in different channels at a given time point t are independent of each other given the hidden state at t , i.e., $P(\mathbf{y}_{it} | z_{it}) = P(y_{it1} | z_{it}) \cdots P(y_{itC} | z_{it})$.

Sometimes the independence assumption does not seem theoretically plausible. For example, even conditioning on a hidden state representing a general life stage, are marital status and parenthood truly independent? On the other hand, given a person's religious views, could their opinions on abortion and gay marriage be thought as independent?

If the goal is to use hidden Markov models for prediction or simulating new sequence data, the analyst should carefully check the validity of independence assumptions. However, if the goal is merely to describe structures and compress information, it can be useful to accept the independence assumption even though it is not completely reasonable in a theoretical sense. When using multichannel sequences, the number of observed states is smaller, which leads to a more parsimonious representation of the model and easier inference of the phenomenon. Also due to the decreased number of observed states, the number of parameters of the model is decreased leading to the improved computational efficiency of model estimation.

The multichannel approach is particularly useful if some of the channels are only partially observed; combining missing and non-missing information into one observation is usually problematic. One would have to decide whether such observations are coded completely missing, which is simple but loses information, or whether all possible combinations of missing and non-missing states are included, which grows the state space larger and makes the interpretation of the model more difficult. In the multichannel approach the data can be used as it is.

Missing data

Missing observations are handled straightforwardly in the context of HMMs. When observation y_{itc} is missing, we gain no additional information regarding hidden states. In such a case, we set the emission probability $b_s(y_{itc}) = 1$ for all $s \in 1, \dots, S$. Sequences with varying lengths are handled by setting missing values before and/or after the observed states.

Log-likelihood and parameter estimation

The unknown transition, emission and initial probabilities are commonly estimated via maximum likelihood. The log-likelihood of the parameters $\mathcal{M} = \{\pi, A, B_1, \dots, B_C\}$ for the HMM is written as

$$\log L = \sum_{i=1}^N \log P(Y_i | \mathcal{M}), \quad (3)$$

where Y_i are the observed sequences in channels $c = 1, \dots, C$ for subject i . The probability of the observation sequence of subject i given the model parameters is

$$\begin{aligned} P(Y_i | \mathcal{M}) &= \sum_{\text{all } z} P(Y_i | z, \mathcal{M}) P(z | \mathcal{M}) \\ &= \sum_{\text{all } z} P(z_1 | \mathcal{M}) P(\mathbf{y}_{i1} | z_1, \mathcal{M}) \prod_{t=2}^T P(z_t | z_{t-1}, \mathcal{M}) P(\mathbf{y}_{it} | z_t, \mathcal{M}) \\ &= \sum_{\text{all } z} \pi_{z_1} b_{z_1}(y_{i11}) \cdots b_{z_1}(y_{i1C}) \prod_{t=2}^T [a_{z_{t-1}z_t} b_{z_t}(y_{it1}) \cdots b_{z_t}(y_{itC})], \end{aligned} \quad (4)$$

where the hidden state sequences $z = (z_1, \dots, z_T)$ take all possible combinations of values in the hidden state space $\{1, \dots, S\}$ and where \mathbf{y}_{it} are the observations of subject i at t in

channels $1, \dots, C$; π_{z_1} is the initial probability of the hidden state at time $t = 1$ in sequence z ; $a_{z_{t-1}z_t}$ is the transition probability from the hidden state at time $t - 1$ to the hidden state at t ; and $b_{z_t}(y_{itc})$ is the probability that the hidden state of subject i at time t emits the observed state at t in channel c .

For direct numerical maximization (DNM) of the log-likelihood, any general-purpose optimization routines such as BFGS or Nelder–Mead can be used (with suitable reparameterizations). Another common estimation method is the expectation–maximization (EM) algorithm, also known as the Baum–Welch algorithm in the HMM context. The EM algorithm rapidly converges close to a local optimum, but compared to DNM, the converge speed is often slow near the optimum.

The probability (4) is efficiently calculated using the forward part of the *forward–backward algorithm* (Baum and Petrie 1966; Rabiner 1989). The backward part of the algorithm is needed for the EM algorithm, as well as for the computation of analytic gradients for derivative based optimization routines. For more information on the algorithms, see a supplementary vignette on CRAN (Helske 2017a).

The estimation process starts by giving initial values to the estimates. Good starting values are needed for finding the optimal solution in a reasonable time. In order to reduce the risk of being trapped in a poor local maximum, a large number of initial values should be tested.

Inference on hidden states

Given our model and observed sequences, we can make several interesting inferences regarding the hidden states. Forward probabilities $\alpha_{it}(s)$ (Rabiner 1989) are defined as the joint probability of hidden state s at time t and the observation sequences $\mathbf{y}_{i1}, \dots, \mathbf{y}_{it}$ given the model \mathcal{M} , whereas backward probabilities $\beta_{it}(s)$ are defined as the joint probability of hidden state s at time t and the observation sequences $\mathbf{y}_{i(t+1)}, \dots, \mathbf{y}_{iT}$ given the model \mathcal{M} .

From forward and backward probabilities we can compute the *posterior probabilities* of states, which give the probability of being in each hidden state at each time point, given the observed sequences of subject i . These are defined as

$$P(z_{it} = s | Y_i, \mathcal{M}) = \frac{\alpha_{it}\beta_{it}}{P(Y_i | \mathcal{M})}. \quad (5)$$

Posterior probabilities can be used to find the locally most probable hidden state at each time point, but the resulting sequence is not necessarily globally optimal. To find the single best hidden state sequence $\hat{z}_i(Y_i) = \hat{z}_{i1}, \hat{z}_{i2}, \dots, \hat{z}_{iT}$ for subject i , we maximize $P(z | Y_i, \mathcal{M})$ or, equivalently, $P(z, Y_i | \mathcal{M})$. A dynamic programming method, the *Viterbi algorithm* (Rabiner 1989), is used for solving the problem.

Model comparison

Models with the same number of parameters can be compared with the log-likelihood. For choosing between models with a different number of hidden states, we need to take account of the number of parameters. We define the Bayesian information criterion (BIC) as

$$BIC = -2 \log(L_d) + p \log \left(\sum_{i=1}^N \sum_{t=1}^T \frac{1}{C} \sum_{c=1}^C \mathbf{I}(y_{itc} \text{ observed}) \right), \quad (6)$$

where L_d is computed using Equation 3, p is the number of estimated parameters, I is the indicator function, and the summation in the logarithm is the size of the data. If data are completely observed, the summation is simplified to $N \times T$. Missing observations in multichannel data may lead to non-integer data size.

2.3. Clustering by mixture hidden Markov models

There are many approaches for finding and describing clusters or latent classes when working with HMMs. A simple option is to group sequences beforehand (e.g., using sequence analysis and a clustering method), after which one HMM is fitted for each cluster. This approach is simple in terms of HMMs. Models with a different number of hidden states and initial values are explored and compared one cluster at a time. HMMs are used for compressing information and comparing different clustering solutions, e.g., finding the best number of clusters. The problem with this solution is that it is, of course, very sensitive to the original clustering and the estimated HMMs might not be well suited for borderline cases.

Instead of fixing sequences into clusters, it is possible to fit one model for the whole data and determine clustering during modeling. Now sequences are not in fixed clusters but get assigned to clusters with certain probabilities during the modeling process. In this section we expand the idea of HMMs to mixture hidden Markov models (MHMMs). This approach was formulated by [van de Pol and Langeheine \(1990\)](#) as a mixed Markov latent class model and later generalized to include time-constant and time-varying covariates by [Vermunt, Tran, and Magidson \(2008\)](#) (who named the resulting model as mixture latent Markov model, MLMM). The MHMM presented here is a variant of MLMM where only time-constant covariates are allowed. Time-constant covariates deal with unobserved heterogeneity and they are used for predicting cluster memberships of subjects.

Mixture hidden Markov model

Assume that we have a set of HMMs $\mathcal{M} = \{\mathcal{M}^1, \dots, \mathcal{M}^K\}$, where $\mathcal{M}^k = \{\pi^k, A^k, B_1^k, \dots, B_C^k\}$ for submodels $k = 1, \dots, K$. For each subject Y_i , denote $P(\mathcal{M}^k) = w_k$ as the prior probability that the observation sequences of a subject follow the submodel \mathcal{M}^k . Now the log-likelihood of the parameters of the MHMM is extended from Equation 3 as

$$\begin{aligned} \log L &= \sum_{i=1}^N \log P(Y_i | \mathcal{M}) \\ &= \sum_{i=1}^N \log \left[\sum_{k=1}^K P(\mathcal{M}^k) \sum_{\text{all } z} P(Y_i | z, \mathcal{M}^k) P(z | \mathcal{M}^k) \right] \\ &= \sum_{i=1}^N \log \left[\sum_{k=1}^K w_k \sum_{\text{all } z} \pi_{z_1}^k b_{z_1}^k(y_{i11}) \cdots b_{z_1}^k(y_{i1C}) \prod_{t=2}^T \left[a_{z_{t-1}z_t}^k b_{z_t}^k(y_{it1}) \cdots b_{z_t}^k(y_{itC}) \right] \right]. \end{aligned} \quad (7)$$

Compared to the usual hidden Markov model, there is an additional summation over the clusters in Equation 7, which seems to make the computations less straightforward than in the non-mixture case. Fortunately, by redefining MHMM as a special type HMM allows us to use standard HMM algorithms without major modifications. We combine the K submodels into one large hidden Markov model consisting of $\sum_{k=1}^K S_k$ states, where the initial state

vector contains elements of the form $w_k\pi^k$. Now the transition matrix is block diagonal

$$A = \begin{pmatrix} A^1 & 0 & \cdots & 0 \\ 0 & A^2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & A^K \end{pmatrix}, \quad (8)$$

where the diagonal blocks $A^k, k = 1, \dots, K$, are square matrices containing the transition probabilities of one cluster. The off-diagonal blocks are zero matrices, so transitions between clusters are not allowed. Similarly, the emission matrices for each channel contain stacked emission matrices B^k .

Covariates and cluster probabilities

Covariates can be added to MHMM to explain cluster memberships as in latent class analysis. The prior cluster probabilities now depend on the subject's covariate values \mathbf{x}_i and are defined as multinomial distribution:

$$P(\mathcal{M}^k|\mathbf{x}_i) = w_{ik} = \frac{e^{\mathbf{x}_i^\top \gamma_k}}{1 + \sum_{j=2}^K e^{\mathbf{x}_i^\top \gamma_j}}. \quad (9)$$

The first submodel is set as the reference by fixing $\gamma_1 = (0, \dots, 0)^\top$.

As in MHMM without covariates, we can still use standard HMM algorithms with a slight modification; we now allow initial state probabilities π to vary between subjects, i.e., for subject i we have $\pi_i = (w_{i1}\pi^1, \dots, w_{iK}\pi^K)^\top$. Of course, we also need to estimate the coefficients γ . For direct numerical maximization the modifications are straightforward. In the EM algorithm, regarding the M-step for γ , **seqHMM** uses iterative Newton's method with analytic gradients and Hessian which are straightforward to compute given all other model parameters. This Hessian can also be used for computing the conditional standard errors of coefficients. For unconditional standard errors, which take account of possible correlation between the estimates of γ and other model parameters, the Hessian is computed using finite difference approximation of the Jacobian of the analytic gradients.

The posterior cluster probabilities $P(\mathcal{M}^k|Y_i, \mathbf{x}_i)$ are obtained as

$$\begin{aligned} P(\mathcal{M}^k|Y_i, \mathbf{x}_i) &= \frac{P(Y_i|\mathcal{M}^k, \mathbf{x}_i)P(\mathcal{M}^k|\mathbf{x}_i)}{P(Y_i|\mathbf{x}_i)} \\ &= \frac{P(Y_i|\mathcal{M}^k, \mathbf{x}_i)P(\mathcal{M}^k|\mathbf{x}_i)}{\sum_{j=1}^K P(Y_i|\mathcal{M}^j, \mathbf{x}_i)P(\mathcal{M}^j|\mathbf{x}_i)} = \frac{L_k^i}{L^i}, \end{aligned} \quad (10)$$

where L^i is the likelihood of the complete MHMM for subject i , and L_k^i is the likelihood of cluster k for subject i . These are straightforwardly computed from forward probabilities. Posterior cluster probabilities are used e.g., for computing classification tables.

2.4. Important special cases

The hidden Markov model is not the only important special case of the mixture hidden Markov model. Here we cover some of the most important special cases that are included in the **seqHMM** package.

Markov model

The Markov model (MM) is a special case of the HMM, where there is no hidden structure. It can be regarded as an HMM where the hidden states correspond to the observed states perfectly. Now the number of hidden states matches the number of the observed states. The emission probability $P(y_{it}) = 1$ if $z_t = y_{it}$ and 0 otherwise, i.e., the emission matrices are identity matrices. Note that for building Markov models the data must be in a single-channel format.

Mixture Markov model

Like MM, the mixture Markov model (MMM) is a special case of the MHMM, where there is no hidden structure. The likelihood of the model is now of the form

$$\begin{aligned} \log L &= \sum_{i=1}^N \log P(\mathbf{y}_i | \mathbf{x}_i, \mathcal{M}^k) = \sum_{i=1}^N \log \sum_{k=1}^K P(\mathcal{M}^k | \mathbf{x}_i) P(\mathbf{y}_i | \mathbf{x}_i, \mathcal{M}^k) \\ &= \sum_{i=1}^N \log \sum_{k=1}^K P(\mathcal{M}^k | \mathbf{x}_i) P(y_{i1} | \mathbf{x}_i, \mathcal{M}^k) \prod_{t=2}^T P(y_{it} | y_{i(t-1)}, \mathbf{x}_i, \mathcal{M}^k). \end{aligned} \quad (11)$$

Again, the data must be in a single-channel format.

Latent class model

Latent class models (LCM) are another class of models that are often used for longitudinal research. Such models have been called, e.g., (latent) growth models, latent trajectory models, or longitudinal latent class models (Vermunt *et al.* 2008; Collins and Wugalter 1992). These models assume that dependencies between observations can be captured by a latent class, i.e., a time-constant variable which we call cluster in this paper.

The **seqHMM** includes a function for fitting an LCM as a special case of MHMM where there is only one hidden state for each cluster. The transition matrix of each cluster is now reduced to a scalar 1 and the likelihood is of the form

$$\begin{aligned} \log L &= \sum_{i=1}^N \log P(Y_i | \mathbf{x}_i, \mathcal{M}^k) = \sum_{i=1}^N \log \sum_{k=1}^K P(\mathcal{M}^k | \mathbf{x}_i) P(Y_i | \mathbf{x}_i, \mathcal{M}^k) \\ &= \sum_{i=1}^N \log \sum_{k=1}^K P(\mathcal{M}^k | \mathbf{x}_i) \prod_{t=1}^T P(y_{it} | \mathbf{x}_i, \mathcal{M}^k). \end{aligned} \quad (12)$$

For LCMs, the data can consist of multiple channels, i.e., the data for each subject consists of multiple parallel sequences. It is also possible to use **seqHMM** for estimating LCMs for non-longitudinal data with only one time point, e.g., to study multiple questions in a survey.

3. Package features

The purpose of the **seqHMM** package is to offer tools for the whole HMM analysis process from sequence data manipulation and description to model building, evaluation, and visualization. Naturally, **seqHMM** builds on other packages, especially the **TraMineR** package designed for sequence analysis. For constructing, summarizing, and visualizing sequence data, **TraMineR**

Usage	Functions/methods
Model construction	<code>build_hmm</code> , <code>build_mhmm</code> , <code>build_mm</code> , <code>build_mmm</code> , <code>build_lcm</code> , <code>simulate_initial_probs</code> , <code>simulate_transition_probs</code> , <code>simulate_emission_probs</code>
Model estimation	<code>fit_model</code>
Model visualization	<code>plot</code> , <code>ssplot</code> , <code>mssplot</code>
Model inference	<code>logLik</code> , <code>BIC</code> , <code>summary</code>
State inference	<code>hidden_paths</code> , <code>posterior_probs</code> , <code>forward_backward</code>
Data visualization	<code>ssplot</code> , <code>ssp + plot</code> , <code>ssp + gridplot</code>
Data and model manipulation	<code>mc_to_sc</code> , <code>mc_to_sc_data</code> , <code>trim_model</code> , <code>separate_mhmm</code>
Data simulation	<code>simulate_hmm</code> , <code>simulate_mhmm</code>

Table 1: Functions and methods in the **seqHMM** package.

provides many useful features. First of all, we use the **TraMineR**'s `stslist` class as the sequence data structure of **seqHMM**. These state sequence objects have attributes such as color palette and alphabet, and they have specific methods for plotting, summarizing, and printing. Many other **TraMineR**'s features for plotting or data manipulation are also used in **seqHMM**.

On the other hand, **seqHMM** extends the functionalities of **TraMineR**, e.g., by providing easy-to-use plotting functions for multichannel data and a simple function for converting such data into a single-channel representation.

Other significant packages used by **seqHMM** include the **igraph** package (Csardi and Nepusz 2006), which is used for drawing graphs of HMMs, and the **nloptr** package (Ypma, Borchers, and Eddelbuettel 2014; Johnson 2014), which is used in direct numerical optimization of model parameters. The computationally intensive parts of the package are written in C++ with the help of the **Rcpp** (Eddelbuettel and François 2011; Eddelbuettel 2013) and **RcppArmadillo** (Eddelbuettel and Sanderson 2014) packages. In addition to using C++ for major algorithms, **seqHMM** also supports parallel computation via the OpenMP interface (Dagum and Enon 1998) by dividing computations for subjects between threads.

Table 1 shows the functions and methods available in the **seqHMM** package. The package includes functions for estimating and evaluating HMMs and MHMMs as well as visualizing data and models. There are some functions for manipulating data and models, and for simulating model parameters or sequence data given a model. In the next sections we discuss the usage of these functions more thoroughly.

As the straightforward implementation of the forward–backward algorithm poses a great risk of under- and overflow, typically forward probabilities are scaled so that there should be no underflow. **seqHMM** uses the scaling as in Rabiner (1989), which is typically sufficient for numerical stability. In case of MHMM though, we have sometimes observed numerical issues in the forward algorithm even with proper scaling. Fortunately this usually means that the backward algorithm fails completely, giving a clear signal that something is wrong. This is especially true in the case of global optimization algorithms which can search unfeasible

areas of the parameter space, or when using bad initial values often with large number of zero-constraints. Thus, **seqHMM** also supports computation on the logarithmic scale in most of the algorithms, which further reduces the numerical unstabilities. On the other hand, as there is a need to back-transform to the natural scale during the algorithms, the log-space approach is somewhat slower than the scaling approach. Therefore, the default option is to use the scaling approach, which can be changed to the log-space approach by setting the `log_space` argument to `TRUE` in, e.g., `fit_model`.

3.1. Building and fitting models

A model is first constructed using an appropriate build function. As Table 1 illustrates, several such functions are available: `build_hmm` for hidden Markov models, `build_mhmm` for mixture hidden Markov models, `build_mm` for Markov models, `build_mmm` for mixture Markov models, and `build_lcm` for latent class models.

The user may give their own starting values for model parameters, which is typically advisable for improved efficiency, or use random starting values. Build functions check that the data and parameter matrices (when given) are of the right form and create an object of class `hmm` (for HMMs and MMs) or `mhmm` (for MHMMs, MMMs, and LCMs). For ordinary Markov models, the `build_mm` function automatically estimates the initial probabilities and the transition matrix based on the observations. For this type of model, starting values or further estimation are not needed. For mixture models, covariates can be omitted or added with the usual `formula` argument using symbolic formulas familiar from, e.g., the `lm` function. Even though missing observations are allowed in sequence data, covariates must be completely observed.

After a model is constructed, model parameters may be estimated with the `fit_model` function. MMs, MMMs, and LCMs are handled internally as their more general counterparts, except in the case of `print` methods, where some redundant parts of the model are not printed.

In all models, initial zero probabilities are regarded as structural zeroes and only positive probabilities are estimated. Thus it is easy to construct, e.g., a left-to-right model by defining the transition probability matrix as an upper triangular matrix.

The `fit_model` function provides three estimation steps: 1) EM algorithm, 2) global DNM, and 3) local DNM. The user can call for one method or any combination of these steps, but should note that they are performed in the above-mentioned order. At the first step, starting values are based on the model object given to `fit_model`. Results from a former step are then used as starting values in the latter. Exceptions to this rule include some global optimization algorithms, which do not use initial values (because of this, performing just the local DNM step can lead to a better solution than global DNM with a small number of iterations).

We have used our own implementation of the EM algorithm for MHMMs whereas the DNM steps (2 and 3) rely on the optimization routines provided by the `nloptr` package. The EM algorithm and computation of gradients were written in C++ with an option for parallel computation between subjects. The user can choose the number of parallel threads (typically, the number of cores) with the `threads` argument.

In order to reduce the risk of being trapped in a poor local optimum, a large number of initial values should be tested. The **seqHMM** package strives to automatize this. One option is to run the EM algorithm multiple times with more or less random starting values for transition or emission probabilities or both. These are called for in the `control_em` argument. Although

not done by default, this method seems to perform very well as the EM algorithm is relatively fast compared to DNM.

Another option is to use a global DNM approach such as the multilevel single-linkage method (MLSL) (Rinnooy Kan and Timmer 1987a,b). It draws multiple random starting values and performs local optimization from each starting point. The LDS modification uses low-discrepancy sequences instead of random numbers as starting points and should improve the convergence rate (Kucherenko and Sytsko 2005).

By default, the `fit_model` function uses the EM algorithm with a maximum of 1000 iterations and skips the local and global DNM steps. For the local step, the L-BFGS algorithm (Nocedal 1980; Liu and Nocedal 1989) is used by default. Setting `global_step = TRUE`, the function performs MSLS-LDS with the L-BFGS as the local optimizer. In order to reduce the computation time spent on non-global optima, the convergence tolerance of the local optimizer is set relatively large, so again local optimization should be performed at the final step.

Unfortunately, there is no universally best optimization method. For unconstrained problems, the computation time for a single EM or DNM rapidly increases as the model size increases and at the same time the risk of getting trapped in a local optimum or a saddle point also increases. As **seqHMM** provides functions for analytic gradients, the optimization routines of **nloptr** which make use of this information are likely preferable. In practice we have had most success with randomized EM, but it is advisable to try a couple of different settings; e.g., randomized EM, EM followed by global DNM, and only global DNM, perhaps with different optimization routines. Documentation of the `fit_model` function gives examples of different optimization strategies and how they can lead to different solutions.

For examples on model estimation and starting values, see a supplementary vignette on CRAN (Helske 2017b).

State and model inference

In **seqHMM**, forward and backward probabilities are computed using the `forward_backward` function, either on the logarithmic scale or in the form of scaled probabilities, depending on the argument `log_space`. Posterior probabilities are obtained from the `posterior_probs` function. In **seqHMM**, the most probable paths are computed with the `hidden_paths` function. For details of the Viterbi and the forward-backward algorithm, see e.g., Rabiner (1989).

The **seqHMM** package provides the `logLik` method for computing the log-likelihood of a model. The method returns an object of class `logLik` which is compatible with the generic information criterion functions `AIC` and `BIC` of R. When constructing the `hmm` and `mhmm` objects via model building functions, the number of observations and the number of parameters of the model are stored as attributes `nobs` and `df` which are extracted by the `logLik` method for the computation of information criteria. The number of model parameters defined from the initial model by taking account of the parameter redundancy constraints (stemming from sum-to-one constraints of transition, emission, and initial state probabilities) and by defining all zero probabilities as structural, fixed values.

The `summary` method automatically computes some features for the MHMM, MMM, and the latent class model, e.g., standard errors for covariates and prior and posterior cluster probabilities for subjects. A `print` method for this summary shows an output of the summaries: estimates and standard errors for covariates, log-likelihood and BIC, and information on most




Figure 1: Stacked sequence plot of the first ten individuals in the `biofam` data plotted with the `ssplot` function. The top plot shows the original sequences, and the three bottom plots show the sequences in the separate channels for the same individuals. The sequences are in the same order in each plot, i.e., the same row always matches the same individual.

probable clusters and prior probabilities.

3.2. Visualizing sequence data

Good graphical presentations of data and models are useful during the whole analysis process from the first glimpse into the data to the model fitting and presentation of results. The `TraMineR` package provides nice plotting options and summaries for simple sequence data, but at the moment there is no easy way of plotting multichannel data. We propose to use a so-called *stacked sequence plot* (`ssp`), where the channels are plotted on top of each other so that the same row in each figure matches the same subject. Figure 1 illustrates an example of a stacked sequence plot with the ten first sequences of the `biofam` data set. The code for creating the figure is shown in Section 4.1.

The `ssplot` function is the simplest way of plotting multichannel sequence data in `seqHMM`. It can be used to illustrate state distributions or sequence index plots. The former is the default option, since index plots can take a lot of time and memory if data are large. Figure 2 illustrates a default plot which the user can modify in many ways (see the code in Section 4.1). More examples are shown in the documentation pages of the `ssplot` function.

Another option is to define function arguments with the `ssp` function and then use previously saved arguments for plotting with a simple `plot` method. It is also possible to combine several `ssp` figures into one plot with the `gridplot` function. Figure 3 illustrates an example of such a plot showing sequence index plots for women and men (see the code in Section 4.1). Sequences are ordered in a more meaningful order using multidimensional scaling scores of




Figure 2: Stacked sequence plot of annual state distributions in the three-channel `biofam` data. This is the default output of the `ssplot` function. The labels for the channels are taken from the named list of state sequence objects, and the labels for the x axis ticks are taken from the column names of the first object.

observations (computed from sequence dissimilarities). After defining the plot for one group, a similar plot for other groups is easily defined using the `update` function.

The `gridplot` function is useful for showing different features for the same subjects or the same features for different groups. The user has a lot of control over the layout, e.g., dimensions of the grid, widths and heights of the cells, and positions of the legends.

We also provide a function `mc_to_sc_data` for the easy conversion of multichannel sequence data into a single channel representation. Plotting combined data is often useful in addition to (or instead of) showing separate channels.

3.3. Visualizing hidden Markov models

For the easy visualization of the model structure and parameters, we propose plotting HMMs as directed graphs. Such graphs are easily called with the `plot` method, with an object of class `hmm` as an argument. Figure 4 illustrates a five-state HMM. The code for producing the plot is shown in Section 4.4.

Hidden states are presented with pie charts as vertices (or nodes), and transition probabilities are shown as edges (arrows, arcs). By default, the higher the transition probability, the thicker the stroke of the edge. Emitted observed states are shown as slices in the pies. For gaining a simpler view, observations with small emission probabilities (less than 0.05 by default) can be combined into one category. Initial state probabilities are given below or next to the respective vertices. In the case of multichannel sequences, the data and the model are converted into a single-channel representation with the `mc_to_sc` function.

A simple default plot is easy to call, but the user has a lot of control over the layout. Figure 5 illustrates another possible visualization of the same model. The code is shown in Section



Figure 3: Showing state distribution plots for women and men in the `biofam` data. Two figures were defined with the `ssp` function and then combined into one figure with the `gridplot` function.




Figure 4: Illustrating a hidden Markov model as a directed graph. Pies represent five hidden states, with slices showing emission probabilities of combinations of observed states. States with emission probability less than 0.05 are combined into one slice. Edges show the transition probabilities. Initial probabilities of hidden states are given below the pies.




Figure 5: Another version of the hidden Markov model of Figure 4 with a different layout and modified labels, legends, and colors. All observed states are shown.

4.4.

The `ssplot` function (see Section 3.2) also accepts an object of class `hmm`. The user can easily choose to plot observations, most probable paths of hidden states, or both. The function automatically computes hidden paths if the user does not provide them.

Figure 6 shows observed sequences with the most probable paths of hidden states given the model. Sequences are sorted according to multidimensional scaling scores computed from hidden paths. The code for creating the plot is shown in Section 4.4.

The `plot` method works for `mhmm` objects as well. The user can choose between an interactive mode, where the model for each (chosen) cluster is plotted separately, and a combined plot with all models in one plot. The equivalent to the `ssplot` function for MHMMs is `mssplot`. It plots stacked sequence plots separately for each cluster. If the user asks to plot more than one cluster, the function is interactive by default.

4. Examples with life course data

In this section we show examples of using the `seqHMM` package. We start by constructing and visualizing sequence data, then show how HMMs are built and fitted for single-channel and multichannel data, then move on to clustering with MHMMs, and finally illustrate how to plot HMMs.

Throughout the examples we use the same `biofam` data described in Section 2.1. We use both the original single-channel data and a three-channel modification named `biofam3c`, which is included in the `seqHMM` package. For more information on the conversion, see the documentation of the `biofam3c` data.




Figure 6: Using the `ssplot` function for an `hmm` object makes it easy to plot the observed sequences together with the most probable paths of hidden states given the model.

4.1. Sequence data

Before getting to the estimation, it is good to get to know the data. We start by loading the original `biofam` data as well as the three-channel version of the same data, `biofam3c`. We convert the data into the `stslist` form with the `seqdef` function. We set the starting age at 15 and set the order of the states with the `alphabet` argument (for plotting). Colors of the states can be modified and stored as an attribute in the `stslist` object – this way the user only needs to define them once.

```
library("seqHMM")

data("biofam", package = "TraMineR")
biofam_seq <- seqdef(biofam[, 10:25], start = 15, labels = c(
  "parent",
  "left", "married", "left+marr", "child", "left+child", "left+marr+ch",
  "divorced"
))

data("biofam3c")
marr_seq <- seqdef(biofam3c$marrried, start = 15, alphabet = c(
  "single",
  "married", "divorced"
))
child_seq <- seqdef(biofam3c$children,
  start = 15,
  alphabet = c("childless", "children")
)
left_seq <- seqdef(biofam3c$left, start = 15, alphabet = c(
  "with parents",
  "left home"
))

attr(marr_seq, "cpal") <- c("violetred2", "darkgoldenrod2", "darkmagenta")
attr(child_seq, "cpal") <- c("darkseagreen1", "coral3")
attr(left_seq, "cpal") <- c("lightblue", "red3")
```

Here we show codes for creating Figures 2, 1, and 3. Such plots give a good glimpse into multichannel data.

Figure 2: Plotting state distributions

We start by showing how to call the simple default plot of Figure 2 in Section 3.3. By default the function plots state distributions (`type = "d"`). Multichannel data are given as a list where each component is an `stslist` corresponding to one channel. If names are given, those will be used as labels in plotting.

```
ssplot(list(
  "Marriage" = marr_seq, "Parenthood" = child_seq,
  "Residence" = left_seq
))
```

Figure 1: Plotting sequences

Figure 1 with the whole sequences requires modifying more arguments. We call for sequence index plots (`type = "I"`) and sort sequences according to the first channel (the original sequences), starting from the beginning. We give labels to y and x axes and modify the positions of y labels. We give a title to the plot but omit the number of subjects, which by default is printed. We set the proportion of the plot given to legends and the number of columns in each legend.

```
ssplot(
  list(
    biofam_seq[1:10, ], marr_seq[1:10, ], child_seq[1:10, ],
    left_seq[1:10, ]
  ),
  sortv = "from.start", sort.channel = 1, type = "I",
  ylab = c("Original", "Marriage", "Parenthood", "Residence"),
  xtlab = 15:30, xlab = "Age", title = "Ten first sequences",
  title.n = FALSE, legend.prop = 0.63, ylab.pos = c(1, 1.5),
  ncol.legend = c(3, 1, 1, 1)
)
```

Figure 3: Plotting sequence data in a grid

For using the `gridplot` function, we first need to specify the `ssp` objects of the separate plots. Here we start by defining the first plot for women with the `ssp` function. It stores the features of the plot, but does not draw anything. We want to sort sequences according to multidimensional scaling scores. These are computed from optimal matching dissimilarities for observed sequences. Any dissimilarity method available in `TraMineR` can be used instead of the default (see the documentation of the `seqdef` function for more information). We want to use the same legends for the both plots, so we remove legends from the `ssp` objects.

Since we are going to plot to two similar figures, one for women and one for men, we can pass the first `ssp` object to the `update` function. This way we only need to define the changes and omit everything that is similar.

These two `ssp` objects are then passed on to the `gridplot` function. Here we make a 2×2 grid, of which the bottom row is for the legends, but the function can also automatically determine the number of rows and columns and the positions of the legends.

```
ssp_f <- ssp(
  list(
    marr_seq[biofam3c$covariates$sex == "woman", ],
```

```
  child_seq[biofam3c$covariates$sex == "woman", ],
  left_seq[biofam3c$covariates$sex == "woman", ]
),
type = "I", sortv = "mds.obs", with.legend = FALSE, title = "Women",
ylab.pos = c(1, 2, 1), xtlab = 15:30, ylab = c(
  "Married", "Children",
  "Residence"
)
)
)

ssp_m <- update(ssp_f,
  title = "Men",
  x = list(
    marr_seq[biofam3c$covariates$sex == "man", ],
    child_seq[biofam3c$covariates$sex == "man", ],
    left_seq[biofam3c$covariates$sex == "man", ]
  )
)

gridplot(list(ssp_f, ssp_m),
  ncol = 2, nrow = 2, byrow = TRUE,
  legend.pos = "bottom", legend.pos2 = "top", row.prop = c(0.65, 0.35)
)
```

For more examples on visualization, see a supplementary vignette on CRAN ([Helske 2017c](#)).

4.2. Hidden Markov models

We start by showing how to fit an HMM for single-channel `biofam` data. The model is initialized with the `build_hmm` function which creates an object of class `hmm`. The simplest way is to use automatic starting values by giving the number of hidden states.

```
sc_initmod_random <- build_hmm(observations = biofam_seq, n_states = 5)
```

It is, however, often advisable to set starting values for initial, transition, and emission probabilities manually. Here the hidden states are regarded as more general life stages, during which individuals are more likely to meet certain observable life events. We expect that the life stages are somehow related to age, so constructing starting values from the observed state frequencies by age group seems like an option worth a try (these are easily computed using the `seqstatf` function in **TraMineR**). We construct a model with four hidden states using age groups 15–18, 19–21, 22–24, 25–27 and 28–30.

The `fit_model` function uses the probabilities given by the initial model as starting values when estimating the parameters. Only positive probabilities are estimated; zero values are fixed to zero. Thus, the amount of 0.1 is added to each value in case of zero-frequencies in some categories (at this point we do not want to fix any parameters to zero). Each row is divided by its sum, so that the row sums equal to 1.

```
sc_init <- c(0.9, 0.06, 0.02, 0.01, 0.01)

sc_trans <- matrix(c(
  0.80, 0.10, 0.05, 0.03, 0.02, 0.02, 0.80, 0.10,
  0.05, 0.03, 0.02, 0.03, 0.80, 0.10, 0.05, 0.02, 0.03, 0.05, 0.80, 0.10,
  0.02, 0.03, 0.05, 0.05, 0.85
), nrow = 5, ncol = 5, byrow = TRUE)

sc_emiss <- matrix(NA, nrow = 5, ncol = 8)
sc_emiss[1, ] <- seqstatf(biofam_seq[, 1:4])[, 2] + 0.1
sc_emiss[2, ] <- seqstatf(biofam_seq[, 5:7])[, 2] + 0.1
sc_emiss[3, ] <- seqstatf(biofam_seq[, 8:10])[, 2] + 0.1
sc_emiss[4, ] <- seqstatf(biofam_seq[, 11:13])[, 2] + 0.1
sc_emiss[5, ] <- seqstatf(biofam_seq[, 14:16])[, 2] + 0.1
sc_emiss <- sc_emiss / rowSums(sc_emiss)

rownames(sc_trans) <- colnames(sc_trans) <- rownames(sc_emiss) <-
  paste("State", 1:5)

colnames(sc_emiss) <- attr(biofam_seq, "labels")

sc_trans

##           State 1 State 2 State 3 State 4 State 5
## State 1      0.80      0.10      0.05      0.03      0.02
## State 2      0.02      0.80      0.10      0.05      0.03
```

```
## State 3    0.02    0.03    0.80    0.10    0.05
## State 4    0.02    0.03    0.05    0.80    0.10
## State 5    0.02    0.03    0.05    0.05    0.85

round(sc_emiss, 3)

##          parent  left married left+marr  child left+child left+marr+ch
## State 1  0.928 0.063  0.002    0.002 0.001    0.001    0.002
## State 2  0.701 0.218  0.018    0.028 0.001    0.004    0.029
## State 3  0.417 0.290  0.050    0.114 0.001    0.006    0.117
## State 4  0.204 0.231  0.080    0.201 0.002    0.009    0.256
## State 5  0.101 0.157  0.097    0.196 0.002    0.013    0.400
##          divorced
## State 1   0.001
## State 2   0.001
## State 3   0.005
## State 4   0.018
## State 5   0.034
```

Now, the `build_hmm` checks that the data and matrices are of the right form.

```
sc_initmod <- build_hmm(
  observations = biofam_seq, initial_probs = sc_init,
  transition_probs = sc_trans, emission_probs = sc_emiss
)
```

We then use the `fit_model` function for parameter estimation. Here we estimate the model using the default options of the EM step.

```
sc_fit <- fit_model(sc_initmod)
```

The fitting function returns the estimated model, its log-likelihood, and information on the optimization steps.

```
sc_fit$logLik
## [1] -16781.99
```

```
sc_fit$model
## Initial probabilities :
## State 1 State 2 State 3 State 4 State 5
## 0.986 0.000 0.014 0.000 0.000
##
## Transition probabilities :
```

```
##           to
## from      State 1 State 2 State 3 State 4 State 5
## State 1   0.786  0.175  0.0391 0.00000 0.0000
## State 2   0.000  0.786  0.0751 0.07568 0.0631
## State 3   0.000  0.000  0.8898 0.08342 0.0267
## State 4   0.000  0.000  0.0000 0.78738 0.2126
## State 5   0.000  0.000  0.0000 0.00136 0.9986
##
## Emission probabilities :
##           symbol_names
## state_names 0 1      2      3      4      5      6      7
## State 1     1 0 0.00000 0.000 0.00000 0.0000 0.000 0.0000
## State 2     1 0 0.00000 0.000 0.00000 0.0000 0.000 0.0000
## State 3     0 1 0.00000 0.000 0.00000 0.0000 0.000 0.0000
## State 4     0 0 0.00195 0.992 0.00581 0.0000 0.000 0.0000
## State 5     0 0 0.21508 0.000 0.00000 0.0246 0.713 0.0474
```

As a multichannel example we fit a 5-state model for the 3-channel data. Emission probabilities are now given as a list of three emission matrices, one for each channel. The `alphabet` function from the **TraMineR** package can be used to check the order of the observed states – the same order is used in the `build` functions. Here we construct a left-to-right model where transitions to earlier states are not allowed, so the transition matrix is upper-triangular. This seems like a valid option from a life-course perspective. Also, in the previous single-channel model of the same data the transition matrix was estimated almost upper triangular. We also give names for channels – these are used when printing and plotting the model.

We estimate model parameters using the local step with the default L-BFGS algorithm using parallel computation with 4 threads.

```
mc_init <- c(0.9, 0.05, 0.02, 0.02, 0.01)

mc_trans <- matrix(
  c(
    0.80, 0.10, 0.05, 0.03, 0.02, 0, 0.90, 0.05, 0.03,
    0.02, 0, 0, 0.90, 0.07, 0.03, 0, 0, 0, 0.90, 0.10, 0, 0, 0, 0, 1
  ),
  nrow = 5, ncol = 5, byrow = TRUE
)

mc_emiss_marr <- matrix(
  c(
    0.90, 0.05, 0.05, 0.90, 0.05, 0.05, 0.05, 0.90,
    0.05, 0.05, 0.90, 0.05, 0.30, 0.30, 0.40
  ),
  nrow = 5, ncol = 3,
  byrow = TRUE
)
```

```

mc_emiss_child <- matrix(c(
  0.9, 0.1, 0.9, 0.1, 0.1, 0.9, 0.1, 0.9, 0.5,
  0.5
), nrow = 5, ncol = 2, byrow = TRUE)

mc_emiss_left <- matrix(c(
  0.9, 0.1, 0.1, 0.9, 0.1, 0.9, 0.1, 0.9, 0.5,
  0.5
), nrow = 5, ncol = 2, byrow = TRUE)

mc_obs <- list(marr_seq, child_seq, left_seq)

mc_emiss <- list(mc_emiss_marr, mc_emiss_child, mc_emiss_left)

mc_initmod <- build_hmm(
  observations = mc_obs, initial_probs = mc_init,
  transition_probs = mc_trans, emission_probs = mc_emiss,
  channel_names = c("Marriage", "Parenthood", "Residence")
)

# For CRAN vignette: load the estimated model object for speed-up
data("hmm_biofam")
# mc_fit <- fit_model(mc_initmod, em_step = FALSE, local_step = TRUE,
# threads = 4)

```

We store the model as a separate object for the ease of use and then compute BIC.

```

# Vignette: already loaded hmm_biofam
# hmm_biofam <- mc_fit$model
BIC(hmm_biofam)

## [1] 28842.7

```

4.3. Clustering and mixture hidden Markov models

When fitting mixture hidden Markov models, the starting values are given as lists, with one component per cluster. For multichannel data, emission probabilities are given as a list of lists. Here we fit a model for two clusters with 5 and 4 hidden states. For the cluster with five states we use the same starting values as for the multichannel HMM described earlier. Covariates are defined with the usual formula and data arguments. Here we use sex and birth cohort to explain cluster memberships.

We fit a model using 100 random restarts of the EM algorithm followed by the local L-BFGS method. Again we use parallel computation.


```

mc_init2 <- c(0.9, 0.05, 0.03, 0.02)

mc_trans2 <- matrix(c(
  0.85, 0.05, 0.05, 0.05, 0, 0.90, 0.05, 0.05, 0, 0,
  0.95, 0.05, 0, 0, 0, 1
), nrow = 4, ncol = 4, byrow = TRUE)

mc_emiss_marr2 <- matrix(c(
  0.90, 0.05, 0.05, 0.90, 0.05, 0.05, 0.05,
  0.85, 0.10, 0.05, 0.80, 0.15
), nrow = 4, ncol = 3, byrow = TRUE)

mc_emiss_child2 <- matrix(c(0.9, 0.1, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5),
  nrow = 4, ncol = 2, byrow = TRUE
)

mc_emiss_left2 <- matrix(c(0.9, 0.1, 0.5, 0.5, 0.5, 0.5, 0.5, 0.5),
  nrow = 4, ncol = 2, byrow = TRUE
)

mhmm_init <- list(mc_init, mc_init2)

mhmm_trans <- list(mc_trans, mc_trans2)

mhmm_emiss <- list(
  list(mc_emiss_marr, mc_emiss_child, mc_emiss_left),
  list(mc_emiss_marr2, mc_emiss_child2, mc_emiss_left2)
)

biofam3c$covariates$cohort <- cut(
  biofam3c$covariates$birthyr,
  c(1908, 1935, 1945, 1957)
)

biofam3c$covariates$cohort <- factor(biofam3c$covariates$cohort,
  labels = c("1909-1935", "1936-1945", "1946-1957")
)

init_mhmm <- build_mhmm(
  observations = mc_obs, initial_probs = mhmm_init,
  transition_probs = mhmm_trans, emission_probs = mhmm_emiss,
  formula = ~ sex + cohort, data = biofam3c$covariates,
  channel_names = c("Marriage", "Parenthood", "Residence"),
  cluster_names = c("Cluster 1", "Cluster 2")
)

# vignette: less restarts and no parallelization

```

```

set.seed(1011)
mhmm_fit <- fit_model(init_mhmm,
  local_step = TRUE, threads = 2,
  control_em = list(restart = list(times = 10))
)
mhmm <- mhmm_fit$model

```

The `summary` method automatically computes some features for an MHMM, e.g., standard errors for covariates and prior and posterior cluster probabilities for subjects. A `print` method shows some summaries of these: estimates and standard errors for covariates (see Section 2.3), log-likelihood and BIC, and information on most probable clusters and prior probabilities. Parameter estimates for transitions, emissions, and initial probabilities are omitted by default. The classification table shows mean probabilities of belonging to each cluster by the most probable cluster (defined from posterior cluster probabilities). A good model should have values close to 1 on the diagonal.

```

summary(mhmm, conditional_se = FALSE)

## Covariate effects :
## Cluster 1 is the reference.
##
## Cluster 2 :
##           Estimate Std. error
## (Intercept)   -1.209    0.138
## sexwoman       0.213    0.141
## cohort1936-1945 -0.785    0.172
## cohort1946-1957 -1.238    0.165
##
## Log-likelihood: -12969.57   BIC: 26592.66
##
## Means of prior cluster probabilities :
## Cluster 1 Cluster 2
##    0.857    0.143
##
## Most probable clusters :
##           Cluster 1 Cluster 2
## count           1753    247
## proportion       0.876    0.124
##
## Classification table :
## Mean cluster probabilities (in columns) by the most probable cluster (rows)
##
##           Cluster 1 Cluster 2
## Cluster 1    0.9775    0.0225
## Cluster 2    0.0013    0.9987

```




Figure 7: A default plot of a hidden Markov model.

4.4. Visualizing hidden Markov models

The figures in Section 3.3 illustrate the five-state multichannel HMM fitted in Section 4.2.

A basic HMM graph is easily called with the `plot` method. Figure 7 illustrates the default plot.

```
plot(hmm_biofam)
```

A simple default plot is a convenient way of visualizing the models during the analysis process, but for publishing it is often better to modify the plot to get an output that best illustrates the structure of the model in hand. Figure 4 and Figure 5 show two variants of the same model.

Figure 4: HMM plot with modifications

In Figure 4 we draw larger vertices, control the distances of initial probabilities (vertex labels), set the curvatures of the edges, give a more descriptive label for the combined slices and give less space for the legend.

```
plot(hmm_biofam,
     vertex.size = 50, vertex.label.dist = 1.5,
     edge.curved = c(0, 0.6, -0.8, 0.6, 0, 0.6, 0), legend.prop = 0.3,
     combined.slice.label = "States with prob. < 0.05"
)
```

Figure 5: HMM plot with a different layout

Here we position the vertices using given coordinates. Coordinates are given in a two-column matrix, with x coordinates in the first column and y coordinates in the second. Arguments `xlim` and `ylim` set the lengths of the axes, and `rescale = FALSE` prevents rescaling the coordinates to the $[-1, 1] \times [-1, 1]$ interval (the default). We modify the positions of initial probabilities, fix edge widths to 1, reduce the size of the arrows in edges, position legend on top of the figure, and print labels in two columns in the legend. Parameter values are shown with one significant digit. All emission probabilities are shown regardless of their value (`combine.slices = 0`).

New colors are set from the ready-defined `colorpalette` data. The `seqHMM` package uses these palettes when determining colors automatically, e.g., in the `mc_to_sc` function. Since here there are 10 combined states, the default color palette is number 10. To get different colors, we choose the ten first colors from palette number 14.

```
vertex_layout <- matrix(c(1, 2, 2, 3, 1, 0, 0.5, -0.5, 0, -1),
  ncol = 2
)

plot(hmm_biofam,
  layout = vertex_layout, xlim = c(0.5, 3.5),
  ylim = c(-1.5, 1), rescale = FALSE, vertex.size = 50,
  vertex.label.pos = c("left", "top", "bottom", "right", "left"),
  edge.curved = FALSE, edge.width = 1, edge.arrow.size = 1,
  with.legend = "left", legend.prop = 0.4, label.signif = 1,
  combine.slices = 0, cpal = colorpalette[[30]][c(14:5)]
)
```

Figure 6: `ssplot` for an HMM object

Plotting observed and hidden state sequences is easy with the `ssplot` function: the function accepts an `hmm` object instead of (a list of) `stslists`. If hidden state paths are not provided, the function automatically computes them when needed.

```
ssplot(hmm_biofam,
  plots = "both", type = "I", sortv = "mds.hidden",
  title = "Observed and hidden state sequences", xtlab = 15:30,
  xlab = "Age"
)
```

4.5. Visualizing mixture hidden Markov models

Objects of class `mhmm` have similar plotting methods to `hmm` objects. The default way of visualizing a model is to plot in an interactive mode, where the model for each cluster is plotted separately. Another option is a combined plot with all models in one plot, although it can be difficult to fit several graphs and legends in one figure.

Figure 8 illustrates the MHMM fitted in Section 4.3. By setting `interactive = FALSE` and `nrow = 2` we plot graphs in a grid with two rows. The rest of the arguments are similar to basic HMM plotting and apply for all the graphs.

```
plot(mhmm,
     interactive = FALSE, nrow = 2, legend.prop = 0.45,
     vertex.size = 50, vertex.label.cex = 1.3, cex.legend = 1.3,
     edge.curved = 0.65, edge.label.cex = 1.3, edge.arrow.size = 0.8
)
```

The equivalent of the `ssplot` function for `mhmm` objects is `mssplot`. It shows data and/or hidden paths one cluster at a time. The function is interactive if more than one cluster is plotted (thus omitted here). Subjects are allocated to clusters according to the most probable hidden state paths.

```
mssplot(mhmm, ask = TRUE)
```

If the user wants more control than the default `mhmm` plotting functions offer, they can use the `separate_mhmm` function to convert a `mhmm` object into a list of separate `hmm` objects. These can then be plotted as any `hmm` objects, e.g., use `ssp` and `gridplot` for plotting sequences and hidden paths of each cluster into the same figure.

5. Conclusion

Hidden Markov models are useful in various longitudinal settings with categorical observations. They can be used for accounting measurement error in the observations (e.g., drug use as in Vermunt *et al.* 2008), for detecting true unobservable states (e.g., different periods of the bipolar disorder as in Lopez 2008), and for compressing information across several types of observations (e.g., finding general life stages as in Helske, Helske, and Eerola 2018).

The **seqHMM** package is designed for analyzing categorical sequences with hidden Markov models and mixture hidden Markov models, as well as their restricted variants Markov models, mixture Markov models, and latent class models. It can handle many types of data from a single sequence to multiple multichannel sequences. Covariates can be included in MHMMs to explain cluster membership. The package also offers versatile plotting options for sequence data and HMMs, and can easily convert multichannel sequence data and models into single-channel representations.

Parameter estimation in (M)HMMs is often very sensitive to starting values. To deal with that, **seqHMM** offers several fitting options with global and local optimization using direct numerical estimation and the EM algorithm.

Almost all intensive computations are done in C++. The package also supports parallel computation.

Especially combined with the **TraMineR** package, **seqHMM** is designed to offer tools for the whole analysis process from data preparation and description to model fitting, evaluation, and visualization. In future we plan to develop MHMMs to deal with time-varying covariates in transition and emission matrices (Bartolucci, Farcomeni, and Pennoni 2012), and add

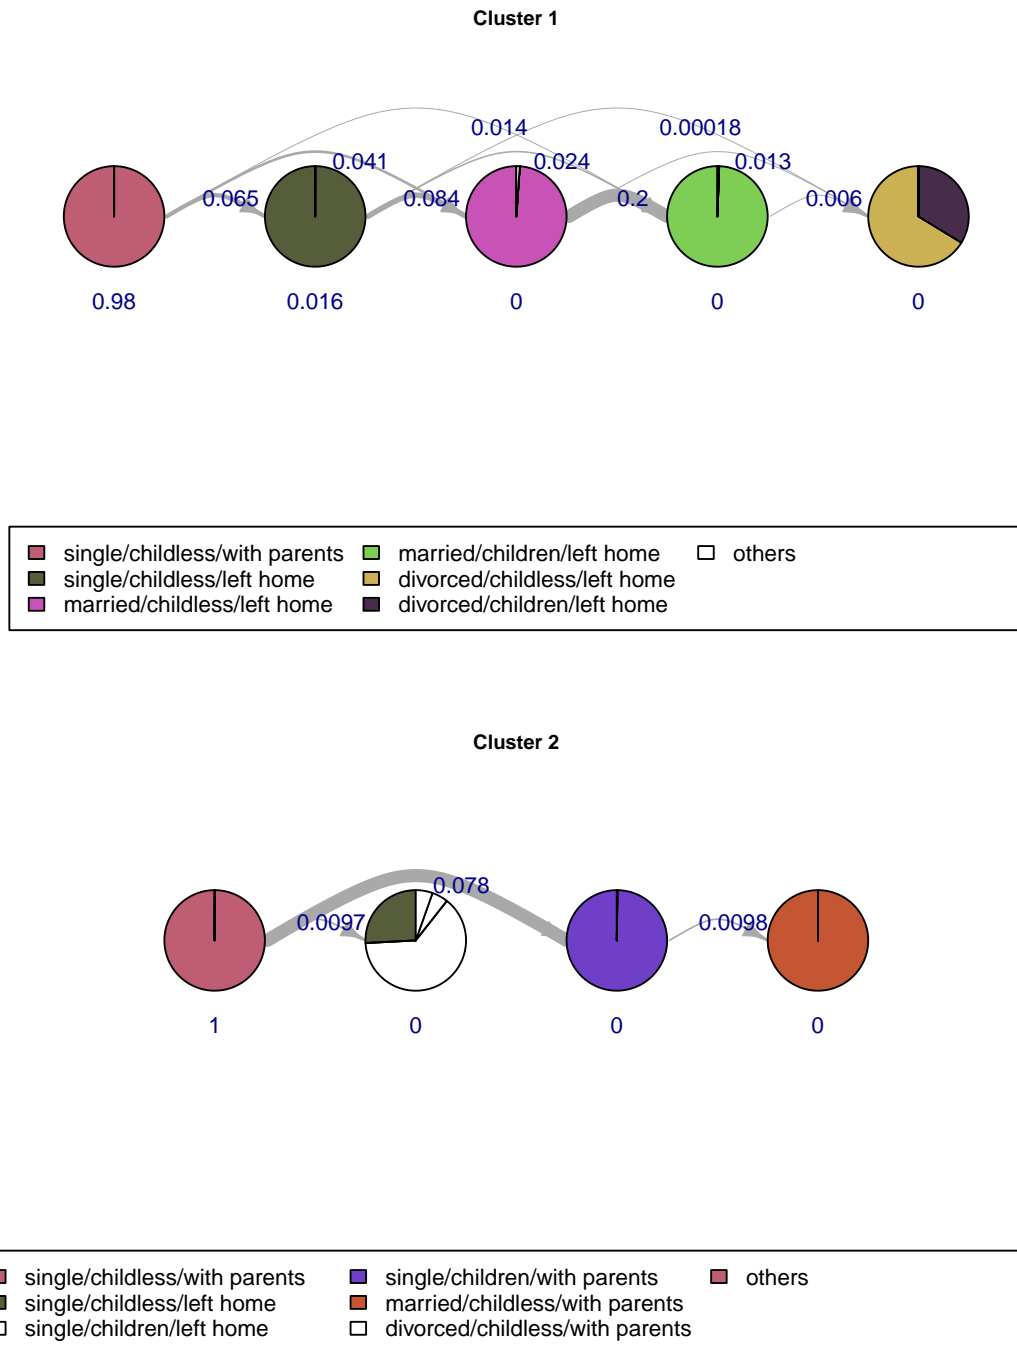


Figure 8: Plotting submodels of an MHMM with the `plot` method.

an option to incorporate sampling weights for model estimation. Also, the computational efficiency of the restricted variants of (M)HMMs, such as latent class models, could be improved by taking account of the restricted structure of those models in EM and log-likelihood computations.

Acknowledgements

Satu Helske is grateful for support for this research from the John Fell Oxford University Press (OUP) Research Fund and the Department of Mathematics and Statistics at the University of Jyväskylä, Finland, and Jouni Helske for the Emil Aaltonen Foundation and the Academy of Finland (research grants 284513 and 312605).

We also wish to thank Mervi Eerola and Jukka Nyblom as well as the editor and two anonymous referees for their helpful comments and suggestions. Comments, suggestions, and bug reports from various users of **seqHMM** have also been highly appreciated.

A. Notations

Symbol	Meaning
Y_i	Observation sequences of subject $i, i = 1 \dots, N$
\mathbf{y}_{it}	Observations of subject i at time $t, t = 1, \dots, T$
y_{itc}	Observation of subject i at time t in channel $c, c = 1, \dots, C$
$m_c \in \{1, \dots, M_c\}$	Observed state space for channel c
z_{it}	Hidden state at time t for subject i
$s \in \{1, \dots, S\}$	Hidden state space
$A = \{a_{sr}\}$	Transition matrix of size $S \times S$
$a_{sr} = P(z_t = r z_{t-1} = s)$	Transition probability between hidden states s and r
$B_c = \{b_s(m_c)\}$	Emission matrix of size $S \times M_c$ for channel c
$b_s(m_c) = P(y_{itc} = m_c z_{it} = s)$	Emission probability of observed state m_c in channel c given hidden state s
$b_s(\mathbf{y}_{it}) = b_s(y_{it1}) \dots b_s(y_{itC})$	Joint emission probability of observations at time t in channels $1, \dots, C$ given hidden state s
$\pi = (\pi_1, \dots, \pi_S)^\top$	Vector of initial probabilities
$\pi_s = P(z_1 = s)$	Initial probability of hidden state s
$\hat{z}_i(Y_i)$	The most probable hidden state sequence for subject i
\mathbf{x}_i	Covariates of subject i
$\mathcal{M}_k, k = 1, \dots, K$	Submodel for cluster k (latent class/cluster)
w_{ik}	Probability of cluster k for subject i
γ_k	Regression coefficients for cluster k
$\{\pi^k, A^k, B_1^k, \dots, B_C^k, \gamma_k\}$	Model parameters for cluster k

References

- Aisenbrey S, Fasang A (2010). “New Life for Old Ideas: The “Second Wave” of Sequence Analysis – Bringing the “Course” Back Into the Life Course.” *Sociological Methods & Research*, **38**(3), 420–462. doi:10.1177/0049124109357532.

- Bartolucci F, Farcomeni A, Pennoni F (2012). *Latent Markov models for longitudinal data*. CRC Press, Boca Raton, FL.
- Bartolucci F, Pandolfi S (2015). **LMest**: *Latent Markov Models with and without Covariates*. R Package Version 2.1, URL <https://CRAN.R-project.org/package=LMest>.
- Baum LE, Petrie T (1966). “Statistical Inference for Probabilistic Functions of Finite State Markov Chains.” *The Annals of Mathematical Statistics*, **67**(6), 1554–1563. doi:10.1214/aoms/1177699147.
- Blanchard P, Bühlmann F, Gauthier JA (eds.) (2014). *Advances in Sequence Analysis: Theory, Method, Applications*. Springer New York Heidelberg Dordrecht London. doi:10.1007/978-3-319-04969-4.
- Collins LM, Wugalter SE (1992). “Latent Class Models for Stage-Sequential Dynamic Latent Variables.” *Multivariate Behavioral Research*, **27**(1), 131–157. doi:10.1207/s15327906mbr2701_8.
- Csardi G, Nepusz T (2006). “The **igraph** Software Package for Complex Network Research.” *InterJournal Complex Systems*, **1695**. URL <https://igraph.org>.
- Dagum L, Enon R (1998). “OpenMP: An Industry Standard API for Shared-Memory Programming.” *Computational Science & Engineering, IEEE*, **5**(1), 46–55.
- Durbin R, Eddy S, Krogh A, Mitchison G (1998). *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, Cambridge, UK.
- Eddelbuettel D (2013). *Seamless R and C++ Integration with Rcpp*. Springer-Verlag, New York. ISBN 978-1-4614-6867-7.
- Eddelbuettel D, François R (2011). “**Rcpp**: Seamless R and C++ Integration.” *Journal of Statistical Software*, **40**(8), 1–18. doi:10.18637/jss.v040.i08.
- Eddelbuettel D, Sanderson C (2014). “RcppArmadillo: Accelerating R with High-Performance C++ Linear Algebra.” *Computational Statistics and Data Analysis*, **71**, 1054–1063. doi:10.1016/j.csda.2013.02.005.
- Elzinga CH, Studer M (2014). “Spell Sequences, State Proximities, and Distance Metrics.” *Sociological Methods & Research*, pp. 3–47. doi:10.1177/0049124114540707.
- Gabardinho A, Ritschard G, Müller NS, Studer M (2011). “Analyzing and Visualizing State Sequences in R with **TraMineR**.” *Journal of Statistical Software*, **40**(4), 1–37. doi:10.18637/jss.v040.i04.
- Gauthier JA, Widmer ED, Bucher P, Notredame C (2009). “How Much Does It Cost? Optimization of Costs in Sequence Analysis of Social Science Data.” *Sociological Methods & Research*, **38**(1), 197–231. doi:10.1177/0049124109342065.
- Gauthier JA, Widmer ED, Bucher P, Notredame C (2010). “Multichannel Sequence Analysis Applied to Social Science Data.” *Sociological Methodology*, **40**(1), 1–38. doi:10.1111/j.1467-9531.2010.01227.x.

- Halpin B (2010). “Optimal Matching Analysis and Life-Course Data: The Importance of Duration.” *Sociological Methods & Research*, **38**(3), 365–388. doi:10.1177/0049124110363590.
- Helske J (2017a). *The main algorithms used in the seqHMM package*. URL https://cran.r-project.org/web/packages/seqHMM/vignettes/seqHMM_algorithms.pdf.
- Helske S (2017b). *Examples and tips for estimating Markovian models with seqHMM*. URL https://cran.r-project.org/web/packages/seqHMM/vignettes/seqHMM_estimation.pdf.
- Helske S (2017c). *Visualization tools in the seqHMM package*. URL https://cran.r-project.org/web/packages/seqHMM/vignettes/seqHMM_visualization.pdf.
- Helske S, Helske J (2019). “Mixture Hidden Markov Models for Sequence Data: The seqHMM Package in R.” *Journal of Statistical Software*, **88**(3), 1–32. doi:10.18637/jss.v088.i03.
- Helske S, Helske J, Eerola M (2018). “Combining Sequence Analysis and Hidden Markov Models in the Analysis of Complex Life Sequence Data.” In *Life Course Research and Social Policies*, pp. 185–200. Springer International Publishing. doi:10.1007/978-3-319-95420-2_11. URL https://doi.org/10.1007/978-3-319-95420-2_11.
- Himmelman L (2010). *HMM – Hidden Markov Models*. R Package Version 1.0, URL <https://CRAN.R-project.org/package=HMM>.
- Hollister M (2009). “Is Optimal Matching Suboptimal?” *Sociological Methods & Research*, **38**(2), 235–264. doi:10.1177/0049124109346164.
- Jackson CH (2011). “Multi-State Models for Panel Data: The **msm** Package for R.” *Journal of Statistical Software*, **38**(8), 1–29. doi:10.18637/jss.v038.i08.
- Johnson SG (2014). *The NLOpt Nonlinear Optimization Package*. URL <http://ab-initio.mit.edu/nlopt>.
- Kucherenko S, Sytsko Y (2005). “Application of Deterministic Low-Discrepancy Sequences in Global Optimization.” *Computational Optimization and Applications*, **30**(3), 297–318. doi:10.1007/s10589-005-4615-1.
- Lesnard L (2010). “Setting Cost in Optimal Matching to Uncover Contemporaneous Socio-Temporal Patterns.” *Sociological Methods & Research*, **38**(3), 389–419. doi:10.1177/0049124110362526.
- Liu DC, Nocedal J (1989). “On the Limited Memory BFGS Method for Large Scale Optimization.” *Mathematical Programming*, **45**(1), 503–528. doi:10.1007/BF01589116.
- Lopez A (2008). *Markov Models for Longitudinal Course of Youth Bipolar Disorder*. ProQuest, Ann Arbor, MI. URL <https://d-scholarship.pitt.edu/6524/1/LopezAdrianaApril23.pdf>.
- MacDonald IL, Zucchini W (1997). *Hidden Markov and Other Models for Discrete-Valued Time Series*. CRC Press, Boca Raton, FL.

- McVicar D, Anyadike-Danes M (2002). “Predicting Successful and Unsuccessful Transitions from School to Work by Using Sequence Methods.” *Journal of the Royal Statistical Society: Series A (Statistics in Society)*, **165**(2), 317–334. doi:10.1111/1467-985X.00641.
- Müller NS, Studer M, Ritschard G (2007). “Classification de Parcours de Vie à l’Aide de l’Optimal Matching.” *XIVe Rencontre de la Société francophone de classification (SFC 2007)*, pp. 157–160.
- Nocedal J (1980). “Updating Quasi-Newton Matrices with Limited Storage.” *Mathematics of Computation*, **35**(151), 773–782. doi:10.1090/S0025-5718-1980-0572855-7.
- O’Connell J, Højsgaard S (2011). “Hidden Semi Markov Models for Multiple Observation Sequences: The **mhsmm** Package for R.” *Journal of Statistical Software*, **39**(4), 1–22. doi:10.18637/jss.v039.i04.
- R Core Team (2018). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. URL <https://www.R-project.org/>.
- Rabiner L (1989). “A Tutorial on Hidden Markov Models and Selected Applications in Speech Recognition.” *Proceedings of the IEEE*, **77**(2), 257–286. doi:10.1109/5.18626.
- Rinnooy Kan A, Timmer G (1987a). “Stochastic Global Optimization Methods Part I: Clustering Methods.” *Mathematical Programming*, **39**(1), 27–56. doi:10.1007/BF02592070.
- Rinnooy Kan A, Timmer G (1987b). “Stochastic Global Optimization Methods Part II: Multi-Level Methods.” *Mathematical Programming*, **39**(1), 57–78. doi:10.1007/BF02592071.
- Turner R, Liu L (2014). **hmm.discnp**: *Hidden Markov Models with Discrete Non-Parametric Observation Distributions*. R Package Version 0.2-3, URL <https://CRAN.R-project.org/package=hmm.discnp>.
- van de Pol F, Langeheine R (1990). “Mixed Markov Latent Class Models.” *Sociological Methodology*, **20**, 213–247. doi:10.2307/271087.
- Vermunt JK, Tran B, Magidson J (2008). “Latent Class Models in Longitudinal Research.” *Handbook of Longitudinal Research: Design, Measurement, and Analysis*, pp. 373–385. Elsevier, Burlington, MA.
- Visser I, Speekenbrink M (2010). “**depmixS4**: An R-package for Hidden Markov Models.” *Journal of Statistical Software*, **36**(7), 1–21. doi:10.18637/jss.v036.i07.
- Ypma J, Borchers HW, Eddelbuettel D (2014). **nloptr**: *R Interface to NLOptr*. R Package Version 1.0.4, URL <https://CRAN.R-project.org/package=nloptr>.

Affiliation:

Satu Helske
Department of Social Research
University of Turku
20500 Turku
Finland
E-mail: satu.helske@utu.fi

Jouni Helske
Department of Mathematics and Statistics
University of Jyväskylä
FI-40014
Finland
E-mail: Jouni.Helske@jyu.fi