



hhsmm: an R package for hidden hybrid Markov/semi-Markov models

Morteza Amini¹ · Afarin Bayat¹ · Reza Salehian¹

Received: 25 September 2021 / Accepted: 3 June 2022 / Published online: 8 July 2022
© The Author(s), under exclusive licence to Springer-Verlag GmbH Germany, part of Springer Nature 2022

Abstract

This paper introduces the **hhsmm** R package, which involves functions for initializing, fitting, and predication of hidden hybrid Markov/semi-Markov models. These models are flexible models with both Markovian and semi-Markovian states, which are applied to situations where the model involves absorbing or macro-states. The left-to-right models and the models with series/parallel networks of states are two models with Markovian and semi-Markovian states. The **hhsmm** also includes Markov/semi-Markov switching regression model as well as the auto-regressive HHSMM, the nonparametric estimation of the emission distribution using penalized B-splines, prediction of future states and the residual useful lifetime estimation in the `predict` function. The commercial modular aero-propulsion system simulation (C-MAPSS) data-set is also included in the package, which is used for illustration of the application of the package features. The application of the **hhsmm** package to the analysis and prediction of the Spain's energy demand is also presented.

Keywords Continuous time sojourn · EM algorithm · Forward-backward · Mixture of multivariate normals · Viterbi algorithm · R

1 Introduction

The package **hhsmm**, developed in the R language (R Development Core Team 2010), involves new tools for modeling multivariate and multi-sample time series by hidden hybrid Markov/semi-Markov models, introduced by Guédon (2005). A

✉ Morteza Amini
morteza.amini@ut.ac.ir

Afarin Bayat
aftbayat@gmail.com

Reza Salehian
reza.salehian@ut.ac.ir

¹ Department of Statistics, School of Mathematics, Statistics and Computer Science, University of Tehran, Tehran, Iran

hidden hybrid Markov/semi-Markov Model (HHSMM) is a model with both Markovian and semi-Markovian states. This package is available from the Comprehensive R Archive Network (CRAN) at <https://cran.r-project.org/package=hhsmm>.

The hidden hybrid Markov/semi-Markov models have many applications for the situations in which there are absorbing or macro-states in the model. These flexible models decrease the time complexity of the hidden semi-Markov models, preserving their prediction power. Another important application of the hidden hybrid Markov/semi-Markov models is in the genetics, where we aim to analysis of the DNA sequences including long interionic zones.

Several packages are available for modeling hidden Markov and semi-Markov models. Some of the packages developed in the R language are **depmixS4** (Visser and Speekenbrink 2010), **HiddenMarkov** (Harte 2006), **msm** (Jackson 2007), **hsmm** (Bulla et al. 2010) and **mhsmm** (ÓConnel and Højsgaard 2011). The packages **depmixS4**, **HiddenMarkov** and **msm** only consider hidden Markov models (HMM), while the two packages **hsmm** and **mhsmm** focus on hidden Markov and hidden semi-Markov (HSMM) models from single and multiple sequences, respectively. These packages do not include hidden hybrid Markov/semi-Markov models, which are included in the **hhsmm** package. The **mhsmm** package has some tools for fitting the HMM and HSMM models to the multiple sequences, while the **hsmm** package has not such a capability. Such a capability is preserved in the **hhsmm** package. Furthermore, the **mhsmm** package is equipped with the ability to define new emission distributions, by using the `mstep` functions, which is also preserved for the **hhsmm** package. In addition to all these differences, the **hhsmm** package is distinguished from **hsmm** and **mhsmm** packages in the following features:

- Some initialization tools are developed for initial clustering, parameter estimation, and model initialization;
- The left-to-right models, which are the models in which the process goes from a state to the next state and never comes back to the previous state, such as the health states of a system and the states of a phoneme in speech recognition, are considered;
- The ability to initialize, fit and predict the models based on data sets containing missing values, using the EM algorithm and imputation methods, is involved;
- The regime Markov/semi-Markov switching linear and additive regression model as well as the auto-regressive HHSMM are involved;
- The nonparametric estimation of the emission distribution using penalized B-splines is added;
- The prediction of the future states is involved;
- The estimation of the residual useful lifetime (RUL), which is the remaining time to the failure of a system (the last state of the left-to-right model, considered for the health of a system), is developed for the left-to-right models used in the reliability theory applications;
- The continuous sojourn time distributions are considered in their correct form;
- The Commercial Modular Aero-Propulsion System Simulation (CMAPSS) data set is included in the package.

There are also tools for modeling HMM in other languages. For instance, the **hmmlearn** library in Python or `hmmtrain` and `hmmestimate` functions in

Statistics and Machine Learning Toolbox of Matlab are available for modeling HMM, while none of them are not suitable for modeling HSMM or HHSMM.

The remainder of the paper is organized as follows. In Sect. 2, we introduce the hidden hybrid Markov/semi-Markov models (HHSMM), proposed by Guédon (2005). Section 3 presents a simple example of the HHSMM model and the hhsmm package. Section 4 presents special features of the hhsmm package including tools for handling missing values, initialization tools, the nonparametric mixture of B-splines for estimation of the emission distribution, regime (Markov/semi-Markov) switching regression, and auto-regressive HHSMM, prediction of the future state sequence, residual useful lifetime (RUL) estimation for reliability applications, continuous-time sojourn distributions, and some other features of the hhsmm package. Finally, the analysis of two real data sets is considered in Sect. 5, to illustrate the application of the hhsmm package.

2 Hidden hybrid Markov/semi-Markov models

Consider a sequence of observations $\{X_t\}$, which is observed for $t = 1, \dots, T$. Assume that the distribution of X_t depends on an un-observed (latent) variable S_t , called *state*. If the sequence $\{S_t\}$ is a Markov chain of order 1, and for any $t \geq 1$, X_t and X_{t+1} are conditionally independent, given S_t , then the sequence $\{(S_t, X_t)\}$ forms a hidden Markov model (HMM). A graphical representation of the dependence structure of the HMM is shown in Fig. 1.


The parameters of an HMM are the *initial probabilities* of the states, the *transition probability* matrix of states, and the parameters of the conditional distribution of observations given states, which is called the *emission* distribution.

The time spent in a state is called the sojourn time. In the HMM model, the sojourn time distribution is simply proved to be geometric distribution. The hidden semi-Markov model (HSMM) is similar to HMM, while the sojourn time distribution can be any other distribution, including discrete and continuous distributions with positive support, such as Poisson, negative binomial, logarithmic, nonparametric, gamma, Weibull, log-normal, etc.

The hidden hybrid Markov/semi-Markov model (HHSMM), introduced by Guédon (2005), is a combination of the HMM and HSMM models. It is defined, for $t = 0, \dots, \tau - 1$ and $j = 1, \dots, J$, by the following parameters:

1. Initial probabilities $\pi_j = P(S_0 = j)$, $\sum_j \pi_j = 1$,
2. Transition probabilities, which are

Fig. 1 A graphical representation of the dependence structure of the HMM



- for a semi-Markovian state j ,

$$p_{jk} = P(S_{t+1} = k | S_{t+1} \neq j, S_t = j), \forall k \neq j; \sum_{k \neq j} p_{jk} = 1; p_{jj} = 0$$

- for a Markovian state j ,

$$\tilde{p}_{jk} = P(S_{t+1} = k | S_t = j); \sum_k \tilde{p}_{jk} = 1$$

By the above definition, any *absorbing state*, which is a state j with $p_{jj} = 1$, is Markovian. This means that if we want to conclude an absorbing state along with some semi-Markovian states in the model, we need to use the HHSMM model.

3. Emission distribution parameters, θ , for the following distribution function

$$f_j(x_t) = f(x_t | S_t = j; \theta)$$

4. The sojourn time distribution is defined for semi-Markovian state j , as follows

$$d_j(u) = P(S_{t+u+1} \neq j, S_{t+u-v} = j, v = 0, \dots, u - 2 | S_{t+1} = j, S_t \neq j), \\ u = 1, \dots, M_j,$$

where M_j stands for an upper bound to the time spent in state j . Also, the survival function of the sojourn time distribution is defined as $D_j(u) = \sum_{v \geq u} d_j(v)$. For a Markovian state j , the sojourn time distribution is the geometric distribution with the following probability mass function

$$d_j(u) = (1 - \tilde{p}_{jj})\tilde{p}_{jj}^{u-1}, \quad u = 1, 2, \dots$$

The parameter estimation of the model is performed via the *EM algorithm* (Dempster et al. 1977). The EM algorithm consists of two steps. In the first step, which is called the *E-step*, the conditional expectation of the unobserved variables (states) is computed given the observations, which are called the *E-step probabilities*. This step utilizes the *forward-backward* algorithm to calculate the E-step probabilities. The second step is the maximization step (*M-step*). In this step the parameters of the model are updated by maximizing the expectation of the logarithm of the joint probability density/mass function of the observed and unobserved data. A brief description of the EM and forward-backward algorithms, as well as the Viterbi algorithm, is given in the Appendix. The Viterbi algorithm obtains the most likely state sequence, for the HHSMM model.

2.1 Examples of hidden hybrid Markov/semi-Markov models

Some examples of HHSMM models are as follows:

- **Models with macro-states:** The macro-states are series or parallel networks of states with common emission distribution. A semi-Markovian model can not be used for macro-states and a hybrid Markov/semi-Markov model is a good choice

in such situations (see Cook and Russell 1986; Durbin et al. 1998; Guédon 2005).

- **Models with absorbing states:** An absorbing state is Markovian by definition. Thus, a model with an absorbing state can not be fully semi-Markovian.
- **Left to right models:** The left-to-right models are useful tools in the reliability analysis of failure systems. Another application of these models is in speech recognition, where the feature sequence extracted from a voice is modeled by a left to right model of states. The transition matrix of a left to right model is an upper triangle matrix with its final diagonal element equal to one, since the last state of a left-to-right model is absorbing. Thus, a hidden hybrid Markov/semi-Markov model might be used in such cases, instead of a hidden fully semi-Markov model.
- **Analysis of DNA sequences:** It is observed that the length of some interionic zones in DNA sequences are approximately, geometrically distributed, while the length of other zones might deviate from the geometric distribution (Guédon 2005).

3 A simple example

To illustrate the application of the **hhsmm** package for initializing, fitting, and prediction of a hidden hybrid Markov/semi-Markov model, we first propose a simple example. We emphasize that the aim of this example is not comparing different models, while this is an example to show how can we use the flexible options of the package **hhsmm** for initializing and fitting different models.

To do this, we define a model, with two Markovian and one semi-Markovian state, and 2, 3, and 2 mixture components in states 1–3, respectively, as follows. The sojourn time distribution for the semi-Markovian state is considered to be the gamma distribution (see Sect. 4.8). The Boolean vector `semi` is used to define the Markovian and semi-Markovian states. Also, the mixture component proportions are defined using the parameter list `mix.p`.

```
R> J <- 3
R> initial <- c(1, 0, 0)
R> semi <- c(FALSE, TRUE, FALSE)
R> P <- matrix(c(0.8, 0.1, 0.1, 0.5, 0, 0.5, 0.1, 0.2, 0.7),
+   nrow = J, byrow=TRUE)
R> par <- list(mu = list(list(7, 8), list(10, 9, 11),
+   list(12, 14)), sigma = list(list(3.8, 4.9),
+   list(4.3, 4.2, 5.4), list(4.5, 6.1))),
+   mix.p = list(c(0.3, 0.7), c(0.2, 0.3, 0.5), c(0.5, 0.5)))
R> sojourn <- list(shape = c(0, 3, 0), scale = c(0, 10, 0),
+   type = "gamma")
R> model <- hhsmmSpec(init = initial, transition = P,
+   parms.emis = par, dens.emis = dmixmvnorm,
+   sojourn = sojourn, semi = semi)
```

Now, we simulate `train` and `test` data sets, using the `simulate` function. The `remission` argument is considered to be `rmixmvnorm`, which is a function for random sample generation from mixture of multivariate normal distributions. The data sets are plotted using the `plot` function. The plots of the `train` and `test` data sets are presented in Figs. 2 and 3, respectively. Different states are distinguished with different colors in the horizontal axis.

```
R> train <- simulate(model, nsim = c(50, 40, 30, 70), seed = 1234,
+   remission = rmixmvnorm)
R> test <- simulate(model, nsim = c(80, 45, 20, 35), seed = 1234,
+   remission = rmixmvnorm)
R> plot(train)
R> plot(test)
```

In order to initialize the parameters of the HHSMM model, we first obtain an initial clustering of the `train` data set, using the `initial_cluster` function. The `nstate` argument is set to 3, and the number of mixture components in the three states is set to `c(2, 2, 2)`. The `ltr` and `final.absorb` arguments is set to `FALSE`, which means that the model is not left-to-right and the final element of each sequence is not in an absorbing state. Thus, the `kmeans` algorithm (Lloyd 1982) is used for the initial clustering.

```
R> clus = initial_cluster(train, nstate = 3, nmix = c(2, 2, 2),
+   ltr = FALSE, final.absorb = FALSE, verbose = FALSE)
```

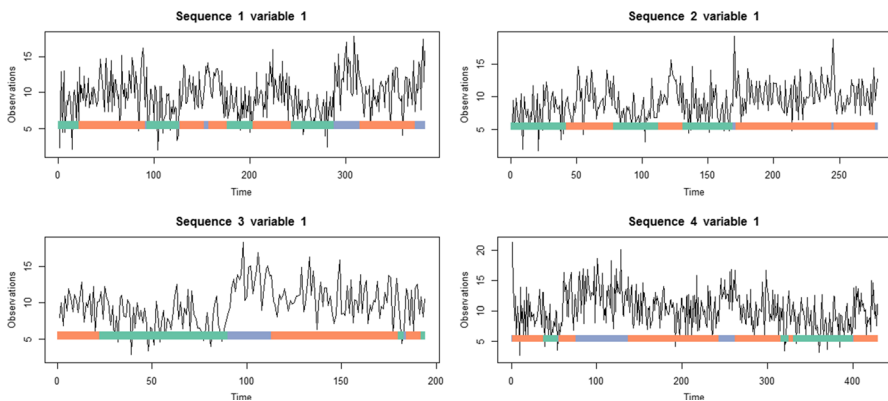


Fig. 2 The plots for 4 sequences of `train` data set

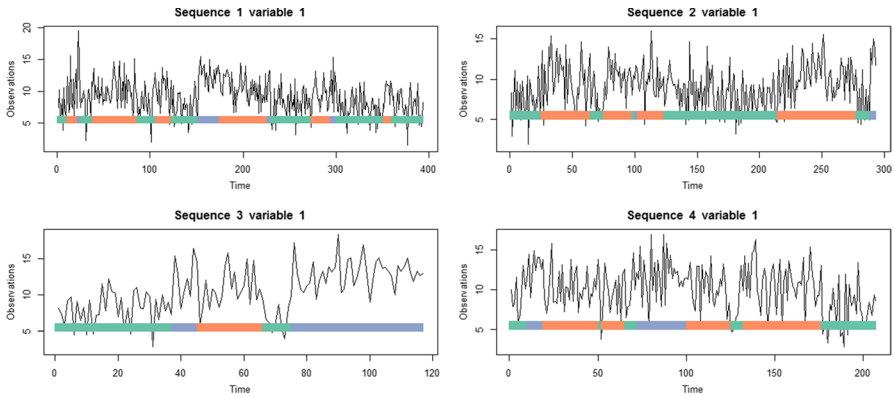


Fig. 3 The plots for 4 sequences of test data set

Now, we initialize the model parameters using the `initialize_model` function. The initial clustering output `clus` is used for estimation of the parameters. The sojourn time distribution is set to "gamma" distribution. First, we use the true value of the `semi` vector for modeling. Thus, the initialized model is a hidden hybrid Markov/semi-Markov model.

```
R> semi <- c(FALSE, TRUE, FALSE)
R> initmodel1 = initialize_model(clus = clus, sojourn = "gamma",
+   M = max(train$N), semi = semi)
```

The model is then fitted using the `hhsmmfit` function as follows. The initialized model `initmodel1` is used as the start value.

```
R> fit1 = hhsmmfit(x = train, model = initmodel1, M = max(train$N),
+   par = list(verbose = FALSE))
```

The log-likelihood trend can also be extracted and plotted as follows. This plot is presented in Fig. 4.

```
R> plot(fit1$loglik[-1], type = "b", ylab = "Log-likelihood",
+   xlab = "Iteration")
```

One can observe, for instance, the estimated initial probabilities, transition matrix, and the estimated parameters of the sojourn distribution as follows.

```

R> fit1$model$init
[1] 0.6261528 0.1144543 0.2593930
R> fit1$model$transition
      [,1]      [,2]      [,3]
[1,] 9.638477e-01 0.02294345 0.01320884
[2,] 4.571514e-01 0.00000000 0.54284857
[3,] 3.460872e-10 0.09880954 0.90119046
R> fit1$model$sojourn
$shape
[1] 0.0000000 0.9714843 0.0000000
$scale
[1] 0.0000 21.2017 0.0000
$type
[1] "gamma"

```

The state sequence is now predicted using the default method "viterbi" of the predict function for the test data set. Because of the displacement property of the states, the homogeneity of the predicted states is computed using the homogeneity function for three states. Since the states are indeed clusters, the homogeneity measures, which are used for clustering, are useful for measuring the homogeneity of two sequences of state. The homogeneity of a specified cluster (state) in two sequences, is defined as the percentage of members of both sequences that are in the same cluster (state) in both sequences. The output of the homogeneity function shows the homogeneity percent of two sequences of states.




Fig. 4 The log-likelihood trend during the model fitting

```
R> yhat1 <- predict(fit1, test)
R> homogeneity(yhat1$s , test$s)
[1] 0.9191686 0.8564920 0.7553957
```

Now, we initialize and fit a fully Markovian model (HMM) by setting `semi` to `c(FALSE, FALSE, FALSE)`. The same clustering output `clus` can be used here.

```
R> semi <- c(FALSE, FALSE, FALSE)
R> initmodel2 = initialize_model(clus = clus, M = max(train$N),
+   semi = semi)
```

The model is again fitted using `hhsmmfit` function.

```
R> fit2 = hhsmmfit(x = train, model = initmodel2, M = max(train$N),
+   par = list(lock.init = TRUE, verbose = FALSE))
```

We can compare some of the estimated parameters of this model with those of the previous one.

```
R> fit2$model$init
[1] 0.3333333 0.3333333 0.3333333
R> fit2$model$transition
      [,1]      [,2]      [,3]
[1,] 9.681322e-01 0.01670787 0.01515991
[2,] 1.695894e-02 0.96554352 0.01749754
[3,] 4.716631e-16 0.08631537 0.91368463
```

Now, we predict the state sequence of the fitted model and compute its homogeneity with the true state sequence.

```
R> yhat2 <- predict(fit2, test)
R> homogeneity(yhat2$s , test$s)
[1] 0.9237875 0.8609272 0.8400000
```

Finally, we initialize and fit a full semi-Markov model (HSMM) to the `train` data set, by setting `semi` to `c(TRUE, TRUE, TRUE)`. The "gamma" distribution is considered as the sojourn time distribution for all states.

```

R> semi <- c(TRUE, TRUE, TRUE)
R> initmodel3 = initialize_model(clus=clus, sojourn = "gamma",
+   M = max(train$N), semi = semi)
R> fit3 = hhsmmfit(x = train, model = initmodel3, M = max(train$N),
+   par = list(verbose = FALSE))
R> fit3$model$transition
      [,1]      [,2]      [,3]
[1,] 0.000000e+00 0.02704853 0.9729515
[2,] 2.242357e-01 0.00000000 0.7757643
[3,] 4.037597e-06 0.99999596 0.0000000
R> fit3$model$sojourn
$shape
[1] 4.2375333 3.4556658 0.1567049
$scale
[1] 8.802259 4.408482 31.021809
$type
[1] "gamma"

```

The prediction and homogeneity computation for this model is done as follows.

```

R> yhat3 <- predict(fit3, test)
R> homogeneity(yhat3$s, test$s)
[1] 0.9232737 0.8069414 0.7358491

```

4 Special features of the package

The **hhsmm** package has several special features, which are described in the following subsections.

4.1 Handling missing values

The **hhsmm** package is equipped with tools for handling data sets with missing values. A special imputation algorithm is used in the `initial_cluster` function. This algorithm, imputes a completely missed row of the data with the average of its previous and next rows, while if some columns are missed, the predictive mean matching method of the function `mice` from package **mice** (Van Buuren and Groothuis-Oudshoorn 2011), with $m = 1$, is used to initially impute the missing values. After performing the initial clustering and initial estimation of the parameters of the model, the `miss_mixmvnorm_mstep` function is considered, as the M-step function of the EM algorithm, for initializing and fitting the model. The function `miss_mixmvnorm_mstep` includes computation of the conditional means and conditional second moments of the missing values

given observed values in each iteration of the EM algorithm and updating the parameters of the Gaussian mixture emission distribution, using the `cov.miss.mix.wt` function. Furthermore, an approximation of the mixture component weights using the observed values and conditional means of the missing values given observed values is used in each iteration. The values of the emission density function, used in the E-step of the EM algorithm are computed by replacing the missing values with their conditional means given the observed values.

Here, we provide a simple example to examine the performance of the aforementioned method. First, we define a model with three states and two variables.

```
R> J <- 3
R> initial <- c(1, 0, 0)
R> semi <- c(FALSE, TRUE, FALSE)
R> P <- matrix(c(0.8, 0.1, 0.1, 0.5, 0, 0.5, 0.1, 0.2, 0.7),
+           nrow = J, byrow = TRUE)
R> par <- list(mu = list(list(c(7, 17), c(8, 18)),
+           list(c(15, 25), c(14, 24), c(16, 16)),
+           list(c(0, 10), c(2, 12))),
+           sigma = list(list(diag(c(2.8, 4.8)), diag(c(3.9, 5.9))),
+           list(diag(c(3.3, 5.3)), diag(c(3.2, 5.2)),
+           diag(c(4.4, 6.4))), list(diag(c(3.5, 5.5)), diag(c(5.1, 7.1))))),
+           mix.p = list(c(0.3, 0.7), c(0.2, 0.3, 0.5), c(0.5, 0.5)))
R> sojourn <- list(shape = c(0, 3, 0), scale = c(0, 10, 0),
+           type = "gamma")
R> model <- hhsmspec(init = initial, transition = P,
+           parms.emis = par, dens.emis = dmixmvnorm,
+           sojourn = sojourn, semi = semi)
```

Now, we simulate the complete train and test data sets.

```
R> train <- simulate(model, nsim = c(10, 8, 8, 18), seed = 1234,
+           remission = rmixmvnorm)
R> test <- simulate(model, nsim = c(8, 6, 6, 15), seed = 1234,
+           remission = rmixmvnorm)
```

First, we initialize and fit the model with complete data sets. To do this, we first use the `initial_cluster` to provide an initial clustering of the train data set.

```
R> clus = initial_cluster(train, nstate = 3, nmix = c(2, 2, 2),
+           ltr = FALSE, final.absorb = FALSE, verbose = FALSE)
```

Now, we initialize and fit the model.

```
R> semi <- c(FALSE, TRUE, FALSE)
R> initmodel1 = initialize_model(clus = clus, sojourn = "gamma",
+   M = max(train$N), semi = semi)
R> fit1 = hhsmmfit(x = train, model = initmodel1, M = max(train$N),
+   par = list(verbose = FALSE))
```

Finally, we predict the state sequence of the test data set, using the `predict.hhsmm` function and the default "viterbi" method.

```
R> yhat1 <- predict(fit1, test)
```

To examine the tools for modeling the data sets with missing values, we randomly select some elements of the `train` and `test` data sets and replace them with NA, as follows.

```
R> p = ncol(train$x)
R> n = nrow(train$x)
R> sammissless = sample(1:n, trunc(n / 10))
R> sammissall = sample(1:n, trunc(n / 20))
R> misrat = matrix(rbinom(trunc(n / 10) * p, 1, 0.2),
+   trunc(n / 10), p)
R> train$x[sammissall, ] <- NA
R> for(i in 1:trunc(n / 10))
+   train$x[sammissless[i], misrat[i,] == 1] <- NA
R> nt = nrow(test$x)
R> sammissless = sample(1:nt, trunc(nt / 12))
R> sammissall = sample(1:nt, trunc(nt / 25))
R> misrat = matrix(rbinom(trunc(nt / 12)*p, 1, 0.15),
+   trunc(nt / 12), p)
R> test$x[sammissall,] <- NA
R> for(i in 1:trunc(nt/12))
+   test$x[sammissless[i], misrat[i, ] == 1] <- NA
```

Now, we provide the initial clustering of the incomplete `train` data set using the `initial_cluster` function.

```
R> clus = initial_cluster(train, nstate = 3, nmix = c(2, 2, 2),
+   ltr = FALSE, final.absorb = FALSE, verbose = FALSE)
```

We can observe that the output of the `initial_cluster` function contains a flag that indicates the missingness in the data set.


```
R> clus$miss
TRUE
```

Now, we initialize and fit the model for the incomplete data set.

```
R> semi <- c(FALSE, TRUE, FALSE)
R> initmodel2 = initialize_model(clus = clus, sojourn = "gamma",
+   M = max(train$N), semi = semi)
R> fit2 = hhsmmfit(x = train, model = initmodel2,
+   M = max(train$N), par = list(lock.init = TRUE, verbose = FALSE))
```

Similarly, we predict the state sequence of the incomplete test data set, using the `predict.hhsmm` function.

```
R> yhat2 <- predict(fit2, test)
```

We can observe that the homogeneity of the predictions of the complete and incomplete data sets are very close to each other.

```
R> homogeneity(yhat1$s, test$s)
[1] 0.8487395 0.9793814 0.0000000
R> homogeneity(yhat2$s, test$s)
[1] 0.9830508 0.8595041 0.0000000
```

4.2 Tools and methods for initializing model

To initialize the HHSMM model, we need to obtain an initial clustering for the train data set. For a left to right model (option `ltr = TRUE` of the `initial_cluster` function), we propose Algorithm 2, which uses Algorithm 1, for a left to right initial clustering, which are included in the function `ltr_clus` of the **hhsmm** package. These algorithms use Hotelling's T-squared test statistic as the distance measure for clustering. The simulations and real data analysis show that the starting values obtained by the proposed algorithm perform well for a left to right model (see Sect. 5 for a real data application). If the model is not a left to right model, then the usual K-means algorithm is used for clustering. Furthermore, the K-means algorithm is used within each initial state to cluster data for mixture components. The number of mixture components can be determined automatically, using the option `nmix = "auto"`, by analysis of the within sum of squares obtained from the `kmeans` function. The number

of starting values of the `kmeans` is set to 10, for the stability of the results. The initial clustering is performed using the `initial_cluster` function.

Algorithm 1 The left to right clustering algorithm for two clusters.

0: For $s = 2, \dots, k - 2$ consider partitions $\{1, \dots, k\} = \{1, \dots, s\} \cup \{s + 1, \dots, k\}$ and compute the means

$$\bar{X}_{1s} = \frac{1}{s} \sum_{i=1}^s X_i, \quad \bar{X}_{2s} = \frac{1}{k-s} \sum_{i=s+1}^k X_i,$$

the variance-covariance matrices

$$\Sigma_{1s} = \frac{1}{s-1} \sum_{i=1}^s (X_i - \bar{X}_{1s})(X_i - \bar{X}_{1s})^\top, \Sigma_{2s} = \frac{1}{k-s-1} \sum_{i=s+1}^k (X_i - \bar{X}_{2s})(X_i - \bar{X}_{2s})^\top,$$

and the standardized distances (Hotelling's T-squared test statistic)

$$d_s = \frac{(s(k-s)/k)(k-p-1)}{(k-2)p} (\bar{X}_{1s} - \bar{X}_{2s})^\top \Sigma_{ps}^{-1} (\bar{X}_{1s} - \bar{X}_{2s}),$$

where

$$\Sigma_{ps} = \frac{(s-1)\Sigma_{1s} + (k-s-1)\Sigma_{2s}}{k-2}.$$

- 0: Let $s^* = \arg \max_s d_s$.
- 0: If $d_{s^*} > F_{(0.05;p,k-1-p)}$, the clusters would be $\{1, \dots, s^*\}$ and $\{s^* + 1, \dots, k\}$, otherwise, no clustering will be done, where $F_{(0.05;p,k-1-p)}$ stands for the 95th percentile of the F distribution with p and $k - 1 - p$ degrees of freedom.
-

After obtaining the initial clustering, the initial estimates of the parameters of the mixture of multivariate normal emission distribution are obtained. Furthermore, the parameters of the sojourn time distribution is obtained by running the method of moments estimation algorithms on the time duration observations of the initial clustering of each state. If we set `sojourn = "auto"` in the `initialize_model` function, the best sojourn time distribution is selected from the list of available sojourn time distributions, using the Chi-square goodness of fit testing on the initial cluster data of all states.

Algorithm 2 The left to right clustering algorithm for $K > 2$ clusters.

0: Let `Nclust = 1`. While `Nclust < K` and the clusters change, do

- for all clusters run Algorithm 1 to obtain two clusters

0: If `Nclust > K`, while `Nclust = K` do

- merge clusters with minimum d_{s^*} values its closest neighbour on its right or left.
-

4.3 Nonparametric mixture of B-splines emission

Usually, the emission distribution belongs to a parametric family of distributions. Although the mixture of multivariate normals is shown to be a good choice in many practical situations, there are also examples in which this class of emission distribution fails to model the skewness and tail weight of the data set. Furthermore, the choice of the number of components of the mixture distribution in each state is a challenge of using mixture of multivariate normals as the emission distribution. As an alternative to parametric emission distribution, HMMs and HSMMs with non-parametric estimates of state-dependent distributions are shown to be more parsimonious in terms of the number of states, easier to interpret, and well fitted to the data (Langrock et al. 2015; Adam et al. 2019). The proposed nonparametric estimation approach of Langrock et al. (2015) and Adam et al. (2019) is based on the idea of representing the densities of the emission distributions as linear combinations of B-spline basis functions, by adding a smoothing penalty term to the quasi-log-likelihood function. In this model, the emission distribution is defined as follows

$$f_j(x) = \sum_{k=-K}^K a_{j,k} \phi_k(x), \quad j = 1, \dots, J, \tag{1}$$

where $\{\phi_{-K}(\cdot), \dots, \phi_K(\cdot)\}$ is a sequences of B-splines and $\{a_{j,k}\}$ is the sequences of unknown coefficients to be estimated. These parameters are estimated in the M-step of the EM algorithm, by maximizing the following penalized quasi-log-likelihood function

$$\ell_p^{\text{HHSMM}}(\theta, \lambda) = \log(L^{\text{HHSMM}}(\theta)) - \frac{1}{2} \sum_{j=1}^J \lambda_j \sum_{k=-K+2}^K (\Delta^2 a_{j,k})^2, \tag{2}$$

in which $L^{\text{HHSMM}}(\theta)$ the quasi-likelihood of the HHSMM model, θ is the parameters of the model, $\Delta a_k = a_k - a_{k-1}$, $\Delta^2 a_k = \Delta(\Delta a_k)$, and $\lambda_1, \dots, \lambda_J$ are the smoothing parameters, which are estimated as follows (Schellhase and Kauermann 2012)

$$\hat{\lambda}_j = \frac{\text{df}(\hat{\lambda}_j) - p}{\sum_{k=-K+2}^K (\Delta^2 \hat{a}_{j,k})^2},$$

where p is the dimension of the data,

$$\text{df}(\hat{\lambda}_j) = \text{tr}(H^{-1}(\hat{a}_j; \lambda_j = \hat{\lambda}_j)H(\hat{a}_j; \lambda_j = 0)),$$

and $H(\hat{a}; \lambda)$ is the hessian matrix of the log-quasi-likelihood at \hat{a} with the specified value of λ .

To illustrate the application of the **hhsmm** package with a nonparametric mixture of B-splines emission distribution, we present a simple simulated data example. To do this, we first simulate data from an HHSMM model with a mixture of multivariate normals as the emission distribution, as follows.

```

R> J <- 3
R> initial <- c(1,0,0)
R> semi <- c(FALSE,TRUE,FALSE)
R> P <- matrix(c(0.8, 0.1, 0.1, 0.5, 0, 0.5, 0.1, 0.2, 0.7),
+   nrow = J, byrow = TRUE)
R> par <- list(mu = list(
+   list(c(7, 17), c(8, 18)),
+   list(c(15, 25), c(14, 24), c(16, 16)),
+   list(c(0, 10), c(2, 12))),
+   sigma = list(list(diag(c(2.8, 4.8)), diag(c(3.9, 5.9))),
+   list(diag(c(3.3, 5.3)), diag(c(3.2, 5.2)), diag(c(4.4, 6.4))),
+   list(diag(c(3.5, 5.5)), diag(c(5.1, 7.1)))),
+   mix.p = list(c(0.3, 0.7),
+   c(0.2, 0.3, 0.5),c(0.5, 0.5)))
R> sojourn <- list(shape = c(0,3,0), scale = c(0,10,0),
+   type = "gamma")
R> model <- hhsmspec(ini = initial, transition = P,
+   parms.emis = par, dens.emis = dmixmvnorm,
+   sojourn = sojourn, semi = semi)
R> train <- simulate(model, nsim = c(10, 8, 8, 18), seed = 1234,
+   remission = rmixmvnorm)
R> test <- simulate(model, nsim = c(8, 6, 6, 15), seed = 1234,
+   remission = rmixmvnorm)

```

Now, we obtain an initial clustering of the data set using the `initial_cluster` function. Note that for a nonparametric emission distribution, we have no mixture components and we should use the option `nmix = NULL`.

```

R> clus = initial_cluster(train, nstate = 3, nmix = NULL,
+   ltr = FALSE, final.absorb = FALSE, verbose = FALSE)

```

In order to initialize a HHSMM with non-parametric estimates of the emission distribution, we use the `initialize_model` function with the arguments `mstep = nonpar_mstep` and `dens.emission = dnonpar`, as follows.

```

R> semi <- c(FALSE, TRUE, FALSE)
R> initmodel1 = initialize_model(clus = clus, mstep = nonpar_mstep,
+   dens.emission = dnonpar, sojourn = "gamma", M = max(train$N),
+   semi = semi)

```

Now, we can use the `hhsmmfit` function to fit the model.

```

R> fit1 = hhsmmfit(x = train, model = initmodel1, M = max(train$N),
+   par = list(verbose = FALSE))

```

Finally, we predict the state sequence of the test data and compute the homogeneity of the predicted sequence and the reals sequence as follows.

```
R> yhat1 <- predict(fit1, test)
R> homogeneity(yhat1$s, test$s)
[1] 0.9210526 0.8508772 0.8750000
```

As one can see from the output of the homogeneity function, the fitted model has a high precision for the prediction of the state sequence of the new data set.

4.4 Regime (Markov/semi-Markov) switching regression model

Kim et al. (2008) considered the following Gaussian regime-switching model


$$y_t = x_t^T \beta_{s_t} + \sigma_{s_t} \epsilon_t, \tag{3}$$

where y_t is the response variable, x_t is a vector of covariates, which may include lagged values of y_t (auto-regressive HHSMM, see the next subsection), s_t is the state, and ϵ_t is the regression error, which is assumed to be distributed as standard normal distribution, for $t = 1, \dots, T$. Model (3) can easily be extended to the case of multivariate responses and also to the case of mixture of multivariate normals. The difference between the regime-switching model (3) and the HHSMM model is that, instead of using the density of y_t given s_t in the likelihood function, we use the conditional density of y_t given x_t and s_t . A graphical representation of the regime-switching model is presented in Fig. 5. The parameters of the regime switching regression model can be estimated using the EM algorithm.

Langrock et al. (2018) considered an extension of the model (3) to the following additive regime-switching model

$$y_t = \mu_{s_t} + \sum_{j=1}^p f_{j,s_t}(x_{j,t}) + \sigma_{s_t} \epsilon_t, \tag{4}$$

Fig. 5 Graphical representation of the regime-switching model



where $f_{j,s_t}(\cdot)$, $j = 1, \dots, p$ are unknown regression functions for p covariates. They utilized the penalized B-splines for estimation of the regression functions.

The estimation of extensions of models (3) and (4) is considered in **hhsmm** package, using the `mixlm_mstep` and `additive_reg_mstep` functions, respectively, as the M-step estimation and `dmixlm` and `dnorm_additive_reg` functions, respectively, which define mixture of multivariate normals and multivariate normal densities, respectively, as the conditional density of the responses given the covariates. The response variables are determined using the argument `resp.ind` in all of these functions, with its default value equal to one, which means that the first column of the input matrix \mathbf{x} , is the univariate response variable.

To illustrate usage of these functions in the **hhsmm** package, we present the following simple simulated data example. First, we simulate data using the function `simulate.hhsmm`, using the argument `remission = rmixlm, covar = list(mean = 0, cov = 1)`. The argument `covar` is indeed an argument of the `rmixlm` function, which is either a function which generates the covariate vector or a list containing the mean vector and the variance-covariance matrix of covariates to be generated from multivariate normal distribution. The `rmixlm` is a function for generation of the data from mixture of linear models, for each state. The list of parameters of this emission distribution consist of the following items:

- `intercept`, a list of the intercepts of the regression models for each state and each mixture component,
- `coefficient`, a list of the coefficient vectors/matrices of the regression models for each state and each mixture component,
- `csigma`, a list of the conditional variances/variance-covariance matrices of the response for each state and each mixture component,
- `mix.p`, a list of mixture component probabilities for each state.

First, we define the model parameters and simulate the data as follows.

```
R> J <- 3
R> initial <- c(1, 0, 0)
R> semi <- rep(FALSE, 3)
R> P <- matrix(c(0.5, 0.2, 0.3, 0.2, 0.5, 0.3, 0.1, 0.4, 0.5),
+           nrow = J, byrow = TRUE)
R> par <- list(intercept = list(3, list(-10, -1), 14),
+           coefficient = list(-1, list(1, 5), -7),
+           csigma = list(1.2, list(2.3, 3.4), 1.1),
+           mix.p = list(1, c(0.4, 0.6), 1))
R> model <- hhsmmSpec(init = initial, transition = P,
+           parms.emis = par, dens.emis = dmixlm, semi = semi)
R> train <- simulate(model, nsim = c(20, 30, 42, 50), seed = 1234,
+           remission = rmixlm, covar = list(mean = 0, cov = 1))
```

Now, we obtain an initial clustering of the data using `initial_cluster` function, with the argument `regress = TRUE`, which is essential for estimation the parameters of the regime switching regression models. By letting `regress =`

TRUE and `ltr = FALSE` the `initial_cluster` function uses an algorithm similar to that of Lloyd (1982) for the K-means method, by fitting linear regression models instead of computing simple means, in each iteration of an algorithm. When using `regress = TRUE` and `ltr = TRUE`, an algorithm similar to that, described in Sect. 4.2 is used for left-to-right clustering, by using regression coefficients instead of mean vectors, and the associated Hotelling's T-square statistic.

```
R> clus = initial_cluster(train = train, nstate = 3,
+   nmix = 2, ltr = FALSE, final.absorb = FALSE,
+   verbose = FALSE, regress = TRUE)
```

We initialize the model, using the `initialize_model` function, with arguments `mstep = mixlm_mstep`, which is a function for M-step estimation of the EM algorithm in the regime switching regression model, and `dens.emission = dmixlm`, which is a function for computation of the probability density function of the mixture Gaussian linear model, for a specified observation vector, a specified state and a specified model's parameters. Next, we fit the model, by using the `hhsmmfit` function.

```
R> initmodel = initialize_model(clus = clus, mstep = mixlm_mstep,
+   dens.emission = dmixlm, sojourn = NULL, semi = rep(FALSE, 3),
+   M = max(train$N), verbose = FALSE)
R> fit1 = hhsmmfit(x = train, model = initmodel, mstep = mixlm_mstep,
+   M = max(train$N), par = list(lock.init = TRUE, verbose = FALSE))
```

The plots of the clustered data as well as the estimated regime-switching regression model lines are then plotted as follows. The resulting plot is shown in Fig. 6.

```
R> plot(train$x[,1] ~ train$x[,2], col = train$s, pch = 16,
+   xlab = "x", ylab = "y")
R> abline(fit1$model$parms.emission$intercept[[1]][[1]],
+   fit1$model$parms.emission$coefficient[[1]][[1]], col = 1)
R> abline(fit1$model$parms.emission$intercept[[1]][[2]],
+   fit1$model$parms.emission$coefficient[[1]][[2]], col = 1)
R> abline(fit1$model$parms.emission$intercept[[2]][[1]],
+   fit1$model$parms.emission$coefficient[[2]][[1]], col = 2)
R> abline(fit1$model$parms.emission$intercept[[2]][[2]],
+   fit1$model$parms.emission$coefficient[[2]][[2]], col = 2)
R> abline(fit1$model$parms.emission$intercept[[3]][[1]],
+   fit1$model$parms.emission$coefficient[[3]][[1]], col = 3)
R> abline(fit1$model$parms.emission$intercept[[3]][[2]],
+   fit1$model$parms.emission$coefficient[[3]][[2]], col = 3)
```

To fit the regime-switching additive regression model to the `train` data, we make an initial clustering of the data, using the `initial_cluster` function, by letting `nstate = 3`, `nmix = NULL` and `regress = TRUE`.




Fig. 6 The regime Markov switching example

Using the argument `nmix = NULL` is essential in this case, since the parameters of the regime switching additive regression model does not involve mixture components.

```
R> clus = initial_cluster(train = train, nstate = 3, nmix = NULL,
+ verbose = FALSE, regress = TRUE)
```

Now, we initialize the model using the function `initialize_model`, with arguments `mstep = additive_reg_mstep` and `dens.emission = dnorm_additive_reg`. Note that here, we only consider a full-Markovian model and thus we let `semi = rep(FALSE, 3)` and `sojourn = NULL`, while one can also consider HSM or HHSMM models by considering different `semi` and `sojourn` arguments.

```
R> initmodel = initialize_model(clus = clus, mstep = additive_reg_mstep,
+ dens.emission = dnorm_additive_reg, sojourn = NULL, semi = rep(FALSE, 3),
+ M = max(train$N), verbose = FALSE)
```

Next, we fit the model by calling the `hhsmmfit` function.

```
R> fit1 = hhsmmfit(x = train, model = initmodel, mstep = additive_reg_mstep,
+ M = max(train$N), par = list(verbose = FALSE))
```


Again, we plot the data to add the fitted lines. The colors of the points show the true states, while the characters present the predicted states.

```
R>
+   xlab = "x", ylab = "y")
R> text(0, 30, "colors_are_real_states", col="red")
R> text(0, 28, "characters_are_predicted_states")
```

To obtain the predicted values of the response variable, we use the `addreg_hhsmm_predict` function as follows.

```
R> pred <- addreg_hhsmm_predict(fit1, train$x[, 2], 5)
R> yhat1 <- pred[[1]]
R> yhat2 <- pred[[2]]
R> yhat3 <- pred[[3]]
```

We add the predicted curves to the plot. The resulting plot is shown in Fig. 7.

```
R> lines(yhat1[order(train$x[, 2])]~sort(train$x[, 2]), col = 2)
R> lines(yhat2[order(train$x[, 2])]~sort(train$x[, 2]), col = 1)
R> lines(yhat3[order(train$x[, 2])]~sort(train$x[, 2]), col = 3)
```

As one can see from Fig. 7, the curves have a proper fit to the data points.




Fig. 7 The Markov regime-switching additive regression fit

4.5 Auto-regressive HHSMM

A special case of the regime-switching regression models (3) and (4) is the auto-regressive HHSMM model, in which we take $x_t = (y_{t-1}, \dots, y_{t-\ell})$, for a specified lag $\ell \geq 1$. Here, we present a simulated data example, to illustrate this special case.

The model specification of the auto-regressive HHSMM is similar to that of the regime-switching regression model, noting that the dimension of x_t is always ℓ times the dimension of y_t . So, we specify the model as follows.

```
R> J <- 2
R> initial <- c(1, 0)
R> semi <- rep(FALSE, 2)
R> P <- matrix(c(0.2, 0.8, 0.1, 0.9),
+             nrow = J, byrow = TRUE)
R> par <- list(intercept = list(0.5, -0.8),
+             coefficient = list(-0.8, 0.7),
+             csigma = list(0.5, 0.2), mix.p = list(1, 1))
R> model <- hhsmspec(init = initial, transition = P,
+                  parms.emis = par, dens.emis = dmixlm, semi = semi)
```

To simulate the data using the `simulate.hhsmm` function, we have to use the argument `emission.control = list(autoregress = TRUE)`. We, then, plot the simulated data by using `plot.hhsmmdata` as follows. The resulting plot is shown in Fig. 8.

```
R> train <- simulate(model, nsim = c(50, 60, 84, 100),
+                  seed = 1234, emission.control = list(autoregress = TRUE))
R> plot(train)
```

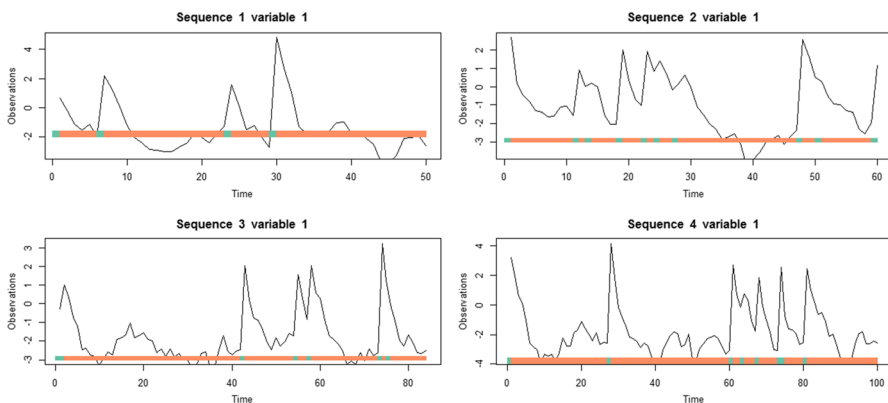


Fig. 8 The ARHMM example simulated train data set

To prepare the data for fitting the regime-switching regression model, we should first construct the lagged data matrix by using the function `lagdata` as follows. The default of the parameter `lags` of this function is equal to 1, which is the number of lags to be calculated.

```
R> train2 = lagdata(train)
```

The resulting lagged data set is then used for obtaining the initial clustering, using the argument `regress = TRUE` and `resp.ind = 2` in the `initial_cluster` function as follows.

```
R> clus = initial_cluster(train = train2, nstate = 2, nmix = 1,
+   ltr = FALSE, final.absorb = FALSE, verbose = FALSE,
+   regress = TRUE, resp.ind = 2)
```

Now, we initialize and fit the model as before. Note that we should use the argument `resp.ind = 2` in place of `...` in both functions (see the manual of the `hhsmm` package <https://cran.r-project.org/web/packages/hhsmm/hhsmm.pdf>)

```
R> initmodel = initialize_model(clus = clus, mstep = mixlm_mstep,
+   dens.emission = dmixlm, sojourn = NULL, semi=rep(FALSE, 2),
+   M = max(train$N), verbose = FALSE, resp.ind = 2)
R> fit1 = hhsmmfit(x = train2, model = initmodel, mstep = mixlm_mstep,
+   resp.ind = 2, M = max(train$N), par = list(verbose = FALSE))
```

To test the performance of the fitted model for prediction of the future time series, we need to simulate a test data set and then right-trim the test data set, using the `train_test_split`, and by setting `train.ratio = 1`, `trim = TRUE` and `trim.ratio = 0.9`, as follows. As one can see, the length of the sequences in `trimmed_test$trimmed` is 90% of the associated lengths in test data set.

```
R> test <- simulate(model, nsim = c(100, 95), seed = 1234,
+   emission.control = list(autoregress = TRUE))
R> trimmed_test = train_test_split(test, train.ratio = 1,
+   trim = TRUE, trim.ratio = 0.9)
R> test$N
[1] 100 95
R> trimmed_test$trimmed$N
[1] 90 85
R> trimmed = trimmed_test$trimmed
R> tc = trimmed_test$trimmed.count
```

The option `train.ratio = 1` means that we do not wish to split the test samples into new train and test subsets and we only need to right trim the

sequences. Now, we have both trimmed sequences in `trimmed` object and the complete test samples in `test` data set, so that we can compare the true and predicted states. The object `tc` contains the number of trimmed items in each sequence, which has to be predicted.

Now, we use the estimated parameters of the ARHMM to predict the future values of the sequence. To do this, we predict the state sequence of the lagged trimmed test data set using the `predict.hhsmm` function and then we obtain the linear predictors for the future values as follows.

```
R> lag_trimmed = lagdata(trimmed)
R> n <- length(tc)
R> Nc1 <- cumsum(c(0, lag_trimmed$N))
R> Nc2 <- cumsum(c(0, test$N))
R> Nc <- cumsum(c(0, trimmed$N))
R> par(mfrow = c(2,1))
R> for (i in 1:n) {
+   new_lag_data = list(x = lag_trimmed$x[(Nc1[i] + 1):Nc1[i + 1]],
+   N = lag_trimmed$N[i])
+   new_data = list(x = as.matrix(trimmed$x[(Nc[i] + 1):Nc[i + 1]]),
+   N = trimmed$N[i])
+   yhat1 <- predict(fit1, new_lag_data, future = tc[i])
+   fstates <- yhat1$state[(test$N[i] - tc[i]):(test$N[i] - 1)]
+   intercept = coefficients = csigma = c()
+   xcurrent = as.vector(new_data$x[new_data$N, ])
+   pred <- xcurrent
+   for(j in 1:tc[i]){
+     intercept[j] <-
+       as.vector(fit1$model$parms.emission$intercept[[fstates[j]]])
+     coefficients[j] <-
+       as.vector(fit1$model$parms.emission$coefficients[[fstates[j]]])
+     predicted <- intercept[j] + xcurrent * coefficients[j]
+     xcurrent <- predicted
+     pred <- c(pred, predicted)
+   }
+   tr_time = ((test$N[i] - tc[i] - 1):(test$N[i] - 1) + 1
+   plot(test$x[(Nc2[i] + 1):Nc2[i + 1]], type = "l", xlab = "time",
+     ylab = "x", main = paste("sequence", i))
+   lines(pred ~ tr_time, lwd = 3, col = 2)
+ }
```

The resulting plot is presented in Fig. 9. The colored lines are the predicted values.

We try to apply the regime-switching additive regression model to fit the ARHMM model. Again, the only difference of the initial clustering using the `initial_cluster` function is to set `nmix = NULL`.

```
R> clus = initial_cluster(train = train2, nstate = 2, nmix = NULL,
+   verbose = FALSE, regress = TRUE, resp.ind = 2)
```

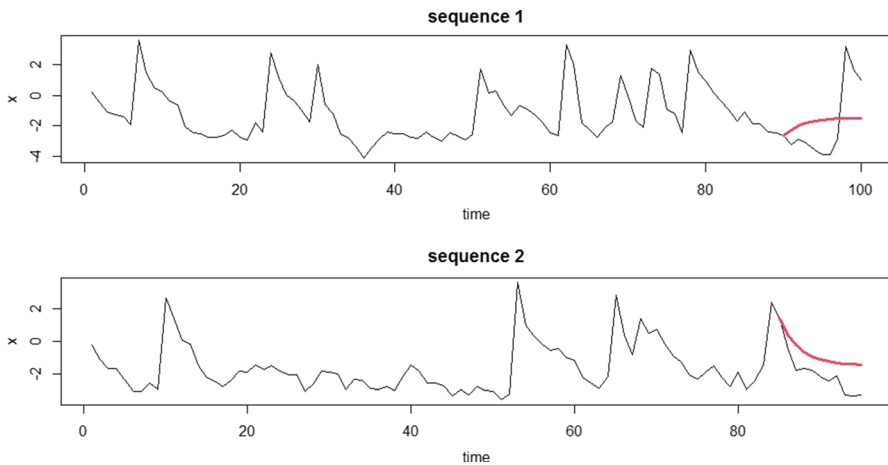


Fig. 9 Trimmed test data set and the predicted values

We initialize the model using the function `initialize_model` by setting `mstep = additive_reg_mstep` and `dens.emission = dnorm_additive_reg`. The difference here is that we pass the parameters of these function to the `initialize_model` function through the argument `control`. In the following, we pass the response indicator and the degree of the B-splines by setting the argument `control = list(resp.ind = 2, K = 7)`.

```
R> initmodel = initialize_model(clus = clus, mstep =
+   additive_reg_mstep, dens.emission = dnorm_additive_reg,
+   sojourn = NULL, semi = rep(FALSE, 2), M = max(train$N),
+   verbose = FALSE, control = list(resp.ind = 2, K = 7))
```

Now, we fit the model by calling the `hhsmmfit` function, as follows.

```
R> fit1 = hhsmmfit(x = train2, model = initmodel,
+   mstep = additive_reg_mstep, M = max(train$N),
+   control = list(resp.ind = 2, K = 7),
+   par = list(verbose = FALSE))
```

Finally, we provide the prediction plots, through the following codes.

```

R> par(mfrow = c(2,1))
R> for (i in 1:n) {
+   new_lag_data = list(x = lag_trimmed$x[(Nc1[i] + 1):Nc1[i + 1]], ),
+   N = lag_trimmed$N[i])
+   new_data = list(x = as.matrix(trimmed$x[(Nc[i] + 1):Nc[i + 1]], ),
+   N = trimmed$N[i])
+   yhat1 <- predict(fit1, new_lag_data, future = tc[i])
+   fstates <- yhat1$s[:(test$N[i] - tc[i]):(test$N[i] - 1)]
+   tr_time = ((test$N[i] - tc[i] - 1):(test$N[i] - 1) + 1
+   intercept = c()
+   coefficients = list()
+   xcurrent = as.vector(new_data$x[new_data$N, ])
+   pred <- xcurrent
+   for(j in 1:tc[i]){
+     preds <- addreg_hhsmm_predict(fit1, c(xcurrent,train$x), 7)
+     predicted <- preds[[fstates[j]]][1]
+     xcurrent <- predicted
+     pred <- c(pred, predicted)
+   }
+   plot(test$x[(Nc2[i] + 1):Nc2[i + 1]], type = "l", xlab = "time",
+   ylab = "x", main = paste("sequence", i))
+   lines(pred ~ tr_time, lwd = 3, col = 2)
+ }

```

The resulting plots are presented in Fig. 10. The colored lines are the predicted values. By comparing the Figs. 9 and 10 one can see that the regime-switching additive regression model results in more accurate prediction especially for the second sequence of the test data set.




Fig. 10 The predicted values using AR-HMM model, using the regime switching additive regression

4.6 Prediction of the future state sequence

To predict the future state sequence at times $T + 1, \dots, T + h$, first, we use viterbi (smoothing) algorithm (see the Appendix) to estimate the probabilities of the most likely path $\alpha_j(t)$ ($L_j(t)$) for $j = 1, \dots, J$ and $t = 0, \dots, \tau - 1$, as well as the current most likely state $\hat{s}_t^* = \arg \max_{1 \leq j \leq J} \alpha_j(t)$ ($\hat{s}_t^* = \arg \max_{1 \leq j \leq J} L_j(t)$). Also, we might compute the probabilities

$$\bar{\delta}_t(j) = \frac{\alpha_j(t)}{\sum_{k=1}^J \alpha_k(t)} \quad (\bar{\delta}_t(j) = \frac{L_j(t)}{\sum_{k=1}^J L_k(t)}). \quad (5)$$

Next, for $j = 1, \dots, h$, we compute the probability of the next state, by multiplying the transition matrix by the current state probability as follows

$$\bar{\delta}_{next} = \left(P \right)^T \bar{\delta}_{current} \quad (6)$$

Then, the j th future state is predicted as

$$\hat{s}_{next}^* = \arg \max_{1 \leq j \leq J} \bar{\delta}_{next}(j) \quad (7)$$

This process continues until the required time $T + h$. The prediction of the future state sequence is done using the function `predict.hhsmm` function in the **hhsmm** package, by determining the argument `future`, which is equal to zero by default. To examine this ability, we study a simple example as follows. First, we define a simple model, just like the model in section 3, and simulate train and test samples from this model, as follows.

```
R> J <- 3
R> initial <- c(1, 0, 0)
R> semi <- c(FALSE, TRUE, FALSE)
R> P <- matrix(c(0.8, 0.1, 0.1, 0.5, 0, 0.5, 0.1, 0.2, 0.7),
+   nrow = J, byrow = TRUE)
R> par <- list(mu = list(list(7, 8), list(10, 9, 11), list(12, 14)),
+   sigma = list(list(3.8, 4.9), list(4.3, 4.2, 5.4), list(4.5, 6.1)),
+   mix.p = list(c(0.3, 0.7), c(0.2, 0.3, 0.5), c(0.5, 0.5)))
R> sojourn <- list(shape = c(0, 3, 0), scale = c(0, 10, 0),
+   type = "gamma")
R> model <- hhsmmSpec(init = initial, transition = P,
+   parms.emis = par, dens.emis = dmixmvnorm,
+   sojourn = sojourn, semi = semi)
R> train <- simulate(model, nsim = c(50, 40, 30, 70), seed = 1234,
+   remission = rmixmvnorm)
R> test <- simulate(model, nsim = c(80, 45, 20, 35), seed = 1234,
+   remission = rmixmvnorm)
```

To examine the prediction performance of the model, we split the test sample from the right, using `train_test_split` function and a trim ratio equal to 0.9, as follows.

```
R> tt = train_test_split(test, train.ratio = 1, trim = TRUE,
+   trim.ratio = 0.9)
R> trimmed = tt$trimmed
R> tc = tt$trimmed.count
```

As in Sect. 3, we initialize and fit an HHSMM model to the `train` data set, as follows.

```
R> clus = initial_cluster(train, nstate=3, nmix=c(2, 2, 2),
+   ltr = FALSE, final.absorb = FALSE, verbose = FALSE)
R> semi <- c(FALSE, TRUE, FALSE)
R> initmodel1 = initialize_model(clus = clus, sojourn = "gamma",
+   M = max(train$N), semi = semi, verbose = FALSE)
R> fit1 = hhsmmfit(x = train, model = initmodel1, M = max(train$N),
+   par = list(verbose = FALSE))
```

Now, we predict the future states of each sequence of the test data set, separately, using the option `future = tc[i]`. Then, we print the homogeneity of real and predicted state sequences, by using the homogeneity function, as follows.

```
R> n <- length(tc)
R> Nc <- cumsum(c(0, trimmed$N))
R> Nc2 <- cumsum(c(0, test$N))
R> for(i in 1:n){
+   newdata = list(x = trimmed$x[(Nc[i] + 1):Nc[i + 1]], ,
+     N = trimmed$N[i])
+   yhat1 <- predict(fit1, newdata, future = tc[i])
+   yhat2 <- predict(fit1, newdata, future = 0)
+   cat("homogeneity without future sequence")
+   print(homogeneity(yhat1$s, test$s[(Nc2[i] + 1):Nc2[i + 1]]))
+   cat("homogeneity without future sequence")
+   print(homogeneity(yhat2$s, trimmed$s[(Nc[i] + 1):Nc[i + 1]]))
+ }
```

```
homogeneity with future sequence[1] 0.8965517 0.8066667 0.6097561
homogeneity without future sequence[1] 0.9079755 0.8066667 0.6097561
homogeneity with future sequence[1] 0.8205128 0.9000000 0.0000000
homogeneity without future sequence[1] 0.9140625 0.9000000 0.0000000
homogeneity with future sequence[1] 1 1 1
homogeneity without future sequence[1] 1 1 1
homogeneity with future sequence[1] 0.9333333 0.8306452 0.6578947
homogeneity without future sequence[1] 0.8750000 0.8306452 0.6578947
```

As, one can see from the above homogeneities, the predictions are quite good.

4.7 Residual useful lifetime (RUL) estimation, for reliability applications

The residual useful lifetime (RUL) is defined as the remaining lifetime of a system at a specified time point. If we analyse a reliable system with a hidden Markov or semi-Markov model, a suitable choice would be a left to right model, with the final state as the failure state. The RUL of such model is defined at time t as

$$RUL_t = \tilde{D} : S_{t+\tilde{D}} = J, S_{t+\tilde{D}-1} = i; \quad 1 \leq i < k \leq J. \tag{8}$$

We describe a method of RUL estimation (see Cartella et al. 2015), which is used in **hhsmm** package. First, we should compute the probabilities in (5), using the Viterbi or smoothing algorithm.

Two different methods are used to obtain point and interval estimates of the RUL in the **hhsmm** package. The first method (the option `confidence = "mean"` in the `predict` function of the **hhsmm** package) is based on the method described in Cartella et al. (2015). This method computes the average time in the current state as follows

$$\tilde{d}_{avg}(\hat{s}_t^*) = \sum_{j=1}^J \left(\mu_{d_j} - \hat{d}_t(j) \right) \bar{\delta}_t(j), \tag{9}$$

where $\mu_{d_j} = \sum_{u=1}^{M_j} u d_j(u)$ is the expected value of the duration variable in state j , and $\hat{d}_t(j)$ is the estimated states duration, computed as follows (Azimi 2004)

$$\hat{d}_t(j) = \hat{d}_{t-1}(j) \bar{\delta}_t(j), \quad t = 2, \dots, M_j, \quad \hat{d}_1(j) = 1, \quad j = 1, \dots, J.$$

In order to obtain a confidence interval for the RUL, Cartella et al. (2015) also computed the standard deviation of the duration variable in state j , σ_{d_j} , and

$$\tilde{d}_{low}(\hat{s}_t^*) = \sum_{j=1}^J \left(\mu_{d_j} - \hat{d}_t(j) - \sigma_{d_j} \right) \bar{\delta}_t(j), \tag{10}$$

$$\tilde{d}_{up}(\hat{s}_t^*) = \sum_{j=1}^J \left(\mu_{d_j} - \hat{d}_t(j) + \sigma_{d_j} \right) \bar{\delta}_t(j) \tag{11}$$

However, to obtain a confidence interval of the specified level $1 - \gamma \in (0, 1)$, we have corrected Eqs. (10) and (11) in the **hhsmm** package as follows

$$\tilde{d}_{low}(\hat{s}_t^*) = \sum_{j=1}^J \left(\mu_{d_j} - \hat{d}_t(j) - z_{1-\gamma/2} \sigma_{d_j} \right) \bar{\delta}_t(j), \tag{12}$$

$$\tilde{d}_{up}(\hat{s}_t^*) = \sum_{j=1}^J \left(\mu_{d_j} - \hat{d}_t(j) + z_{1-\gamma/2} \sigma_{d_j} \right) \bar{\delta}_t(j), \tag{13}$$

where $z_{1-\gamma/2}$ is the $1 - \gamma/2$ quantile of the standard normal distribution.

The probability of the next state is obtained by multiplying the transition matrix by the current state probability as follows

$$\bar{\delta}_{next} = \left[\bar{\delta}_{t+\tilde{d}}(j) \right]_{1 \leq j \leq J} = \left(P \right)^T \bar{\delta}_t \tag{14}$$

while the maximum a posteriori estimate of the next state is calculated as

$$\hat{s}_{next}^* = \hat{s}_{t+\tilde{d}}^* = \arg \max_{1 \leq j \leq J} \bar{\delta}_{t+\tilde{d}}(j) \tag{15}$$

If $\hat{S}_{t+\tilde{d}}(j)$ coincides with the failure state J , the failure will happen after the remaining time at the current state is over and the average estimation of the failure time is $\tilde{D}_{avg} = \tilde{d}_{avg}(\hat{s}_t^*)$, with the lower and upper bounds $\tilde{D}_{low} = \tilde{d}_{low}(\hat{s}_t^*)$ and $\tilde{D}_{up} = \tilde{d}_{up}(\hat{s}_t^*)$, respectively, otherwise, the sojourn time of the next state is calculated as

$$\tilde{d}_{avg}(\hat{S}_{t+\tilde{d}}^*) = \sum_{j=1}^J \mu_{d_j} \bar{\delta}_{t+\tilde{d}}(j) \tag{16}$$

$$\tilde{d}_{low}(\hat{S}_{t+\tilde{d}}^*) = \sum_{j=1}^J \left(\mu_{d_j} - z_{1-\gamma/2} \sigma_{d_j} \right) \bar{\delta}_{t+\tilde{d}}(j) \tag{17}$$

$$\tilde{d}_{up}(\hat{S}_{t+\tilde{d}}^*) = \sum_{j=1}^J \left(\mu_{d_j} + z_{1-\gamma/2} \sigma_{d_j} \right) \bar{\delta}_{t+\tilde{d}}(j) \tag{18}$$

This procedure is iterated until the failure state is encountered in the prediction of the next state. The estimate of the RUL is then calculated by summing all the aforementioned estimated remaining times, as follows

$$\tilde{D}_{avg} = \sum \tilde{d}_{avg}, \quad \tilde{D}_{low} = \sum \tilde{d}_{low}, \quad \tilde{D}_{up} = \sum \tilde{d}_{up} \tag{19}$$

In the second method (the option `confidence = "max"` in the `predict` function of the `hhsmm` package), we relax the normal assumption and use the mode and quantiles of the sojourn time distribution, by replacing the mean μ_{d_j} with the mode $m_{d_j} = \arg \max_{1 \leq u \leq M_j} d_j(u)$ and replacing $-z_{1-\gamma/2} \sigma_{d_j}$ with $\min\{v; \sum_{u=1}^v d_j(u) \leq \gamma/2\}$ and $+z_{1-\gamma/2} \sigma_{d_j}$ with $M_j - \min\{v; \sum_{u=v}^{M_j} d_j(u) \leq \gamma/2\}$ in Eqs. (10), (12), (13), (16), (17) and (18).

4.8 Continuous time sojourn distributions

Since the measurements of the observations are always performed on discrete time units (assumed to be positive integers), the sojourn time probabilities of the

sojourn time distribution with probability density function g_j , in state j , is obtained as follows

$$\begin{aligned}
 d_j(u) &= P(S_{t+u+1} \neq j, S_{t+u-v} = j, v = 0, \dots, u - 2 | S_{t+1} = j, S_t \neq j) \\
 &= \int_{u-1}^u g_j(y) dy / \int_0^{M_j} g_j(y) dy, \quad j = 1, \dots, J, \quad u = 1, \dots, M_j.
 \end{aligned}
 \tag{20}$$

Almost all flexible continuous distributions with positive domain, which are used as the lifetime distribution, including gamma, weibull, log-normal, Birnbaum-Saunders, inverse-gamma, Fréchet, Gumbel and many other distributions, might be used as the continuous-time sojourn distribution. Some of the continuous sojourn time distributions, included in the **hhsmm** package, are as follows:

- **Gamma sojourn:** The gamma sojourn time density functions are

$$g_j(y) = \frac{y^{\alpha_j-1} e^{-\frac{y}{\beta_j}}}{\beta_j^{\alpha_j} \Gamma(\alpha_j)}, \quad j = 1, \dots, J,$$

which result in

$$d_j(u) = \int_{u-1}^u y^{\alpha_j-1} e^{-\frac{y}{\beta_j}} dy / \int_0^{M_j} y^{\alpha_j-1} e^{-\frac{y}{\beta_j}} dy$$

- **Weibull sojourn:** The Weibull sojourn time density functions are

$$g_j(y) = \frac{\alpha_j}{\beta_j} \left(\frac{y}{\beta_j}\right)^{\alpha_j-1} \exp\left\{-\left(\frac{y}{\beta_j}\right)^{\alpha_j}\right\}, \quad j = 1, \dots, J,$$

which result in

$$d_j(u) = \int_{u-1}^u y^{\alpha_j-1} \exp\left\{-\left(\frac{y}{\beta_j}\right)^{\alpha_j}\right\} dy / \int_0^{M_j} y^{\alpha_j-1} \exp\left\{-\left(\frac{y}{\beta_j}\right)^{\alpha_j}\right\} dy$$

- **log-normal sojourn:** The log-normal sojourn time density functions are

$$g_j(y) = \frac{1}{\sqrt{2\pi}\sigma_j} \exp\left\{\frac{-1}{2\sigma_j^2}(\log y - \mu_j)^2\right\}, \quad j = 1, \dots, J,$$

which result in

$$d_j(u) = \int_{u-1}^u \exp\left\{\frac{-1}{2\sigma_j^2}(\log y - \mu_j)^2\right\} dy / \int_0^{M_j} \exp\left\{\frac{-1}{2\sigma_j^2}(\log y - \mu_j)^2\right\} dy$$

4.9 Other features of the package

There are some other features included in the **hhsmm** package, which are listed below:

- `dmixmvnorm`: Computes the probability density function of a mixture of multivariate normals for a specified observation vector, a specified state, and a specified model's parameters
- `mixmvnorm_mstep`: The M step function of the EM algorithm for the mixture of multivariate normals as the emission distribution using the observation matrix and the estimated weight vectors
- `rmixmvnorm`: Generates a vector of observations from mixture multivariate normal distribution in a specified state and using the parameters of a specified model
- `train_test_split`: Splits the data sets to train and test subsets with an option to right trim the sequences
- `lagdata`: Creates lagged time series of a data
- `score`: Computes the score (log-likelihood) of new observations using a trained model
- `homogeneity`: Computes maximum homogeneity of two state sequences
- `hhsmmdata`: Converts a matrix of data and its associated vector of sequence lengths to a data list of class "hhsmmdata"

5 Real data analysis

To examine the performance of the **hhsmm** package, we consider the analysis of two real data sets. The first data set is the Spain energy market data set from **MSwM** package and the second one is the Commercial Modular Aero-Propulsion System Simulation (CMAPSS) data set from the **CMAPSS** data package.

5.1 Spain energy market data set

The Spain energy market data set (Fontdecaba et al. 2009) contains the price of the energy in Spain with other economic data. The daily data is from January 1, 2002 to October 31, 2008, during working days (Monday to Friday). This data set is available in **MSwM** package (<https://cran.r-project.org/package=MSwM>), in a data-frame named `energy` and contains 1785 observations on 7 variables: `Price` (Average price of energy in Cent/kwh), `Oil` (Oil price in Euro/barril), `Gas` (Gas price in Euro/MWh), `Coal` (Coal price in Euro/T), `EurDol` (Exchange rate between Dolar-Euro in USD-Euro), `Ibex35` (Ibex 35 index divided by one thousand) and `Demand` (Daily demand of energy in GWh). This data-set is also

analysed in Fontdecaba et al. (2009), using Markov switching regression model. The objective of the analysis is to predict the response variable `Price`, based on the information in other variables (covariates).

In order to analyze the `energy` data set, we load it from `MSwM` package, and transform it into a "hhsmmdata" using `hhsmmdata` function as follows.

```
R> library(MSwM)
R> data(energy)
R> energy.hmm = hhsmmdata(energy)
R> p = ncol(energy.hmm$x) - 1
```

We consider a two-state model. Here, we consider a fully Markovian model and thus, we let `semi <- rep(FALSE, J)`. Although an optimal value of K might be obtained by minimizing the AIC, BIC or even a cross-validation error, we set the degree of the B-splines to $K = 20$ for this analysis, for the sake of brevity.

```
R> K <- 20
R> J <- 2
R> initial <- rep(1/J, J)
R> semi <- rep(FALSE, J)
```

First, we make an initial clustering for the data set. Again, we point out that we should consider `nmix = NULL` and `regress = TRUE` in the `initial_cluster` function.

```
R> clus = initial_cluster(train = energy.hmm, nstate = J,
+   nmix = NULL, ltr = FALSE, final.absorb = FALSE,
+   verbose = TRUE, regress = TRUE)
```

To initialize the model, we use the `initialize_model` function, with arguments `mstep = additive_reg_mstep`, `dens.emission = dnorm_additive_reg` and `control = list(K = K)`.

```
R> initmodel = initialize_model(clus = clus, mstep =
+   additive_reg_mstep, dens.emission = dnorm_additive_reg,
+   sojourn = NULL, semi = semi, M = max(energy.hmm$N),
+   verbose = FALSE, control = list(K = K))
```

Next, we fit the model by calling the `hhsmmfit` function as follows.

```
R> fit1 = hhsmmfit(x = energy.hmm, model = initmodel,
+   mstep = additive_reg_mstep, M = max(energy.hmm$N),
+   control = list(K = K))
```

Now, we can obtain the response predictions using the `addreg_hhsmm_predict` function as follows.

```
R> pred2 <- addreg_hhsmm_predict(fit1, energy.hmm$x[, 2:(p+1)], K)
```

To visualize the results, first, we add the predicted states to a `hhsmmdata` set made by the response variable. Then, we plot it using the `plot.hhsmm` function, as follows.

```
R> s = fit1$yhat
R> newdata = hhsmmdata(x = as.matrix(energy.hmm$x[, 1]),
+   N = energy.hmm$N)
R> newdata$s = s
R> plot(newdata)
```

The resulting plot is presented in Fig. 11. The predicted states are shown by two different colors on the horizontal axis.

In the second plot, we want to show the separate predictions for the two states along with the true response values. To do this, we use the following lines of codes.

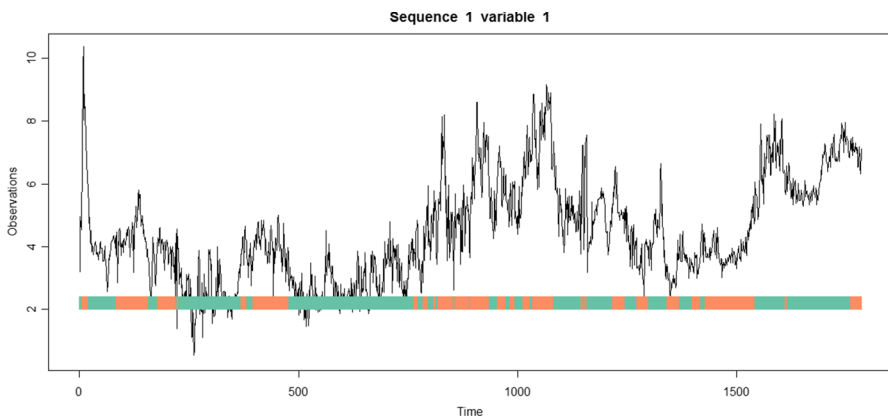


Fig. 11 Spain energy data and its estimated states

```

R> col.states = fit1$yhat
R> col.states[col.states == 1] = 'goldenrod2'
R> col.states[col.states == 2] = 'green4'
R> plot(energy.hmm$x[, 1], col = col.states,
+      pch = 16, xlab = "Time", ylab = "EnergyOilPrice")
R> time = 1:length(pred2[[1]])
R> lines(pred2[[1]] ~ time, col = "red", lwd = 0.25)
R> lines(pred2[[2]] ~ time, col = "blue", lwd = 0.25)

```

The resulting plot is presented in Fig. 12. The predictions associated with the two states' are shown by blue and red lines.

To visualize the nonparametric regression curves, we consider only the covariate `Oil Price`, which is the second column of the `energy` data set. In the following, we initialize and fit the model to the data set containing only the first column as the response and the second column as the covariate.

```

R> energy.hmm2 = hhsmmdata(energy[, 1:2])
R> clus = initial_cluster(train = energy.hmm2, nstate = J, nmix = NULL,
+ regress = TRUE)
R> initmodel = initialize_model(clus = clus, mstep = additive_reg_mstep,
+ dens.emission = dnorm_additive_reg, sojourn = NULL, semi = semi,
+ M = max(energy.hmm$N))
R> fit1 = hhsmmfit(x = energy.hmm2, model = initmodel,
+ mstep = additive_reg_mstep, M = max(energy.hmm2$N))

```

Again, we obtain the response predictions as follows.




Fig. 12 Two-state prediction of the energy price based on the other variables

```
R> pred <- addreg_hhsmm_predict(fit1, energy.hmm2$x[, 2], K)
```

Now, we can plot the corresponding graph as follows.

```
R> col.states = fit1$yhat
R> col.states[col.states == 1] = 'goldenrod2'
R> col.states[col.states == 2] = 'green4'
R> plot(energy.hmm$x[, 1] ~ energy.hmm$x[, 2], col = col.states,
+       pch = 16, xlab = "Oil Price", ylab = "Energy Price", lwd = 2,
+       main = "Case p=1")
R> lines(pred[[1]][order(energy.hmm$x[, 2])] ~ sort(energy.hmm$x[, 2]),
+        col = 1, lwd = 2)
R> lines(pred[[2]][order(energy.hmm$x[, 2])] ~ sort(energy.hmm$x[, 2]),
+        col = 1, lwd = 2)
R> text(30, 7, "State 1", cex = 1.5)
R> text(60, 3, "State 2", cex = 1.5)
```

The resulting plot is presented in Fig. 13. As one can see from Fig. 13, the two curves are well-fitted to the data points.

To compare the prediction performance of the above-mentioned model with the simple (single state) additive regression model, we can use the `additive_reg_mstep` function, with a weight matrix with a single column and all components equal to 1.




Fig. 13 Prediction curves for regime switching nonparametric regression model of energy price on oil price


```

R> n = energy.hmm$N
R> wt = matrix(rep(1, n), n, 1)
R> emission = additive_reg_mstep(energy.hmm$x, wt,
+   control = list(K = K))
R> tmpfit = list(model = list(J = 1, parms.emission = emission))
R> pred1 <- addreg_hhsmm_predict(tmpfit, energy.hmm$x[,2:(p+1)], K)

```

We compute the sum of squared errors (SSE) of the two competitive models as follows.

```

R> sse1 = sum((pred1 - energy.hmm$x[, 1])^2)
R> sse2 = sum(sapply(1:J, function(j)
+   sum((pred2[[j]][s == j] - energy.hmm$x[s == j, 1]) ^ 2)))

```

We plot the predictions of two competitive models by adding their SSE values to plots, as follows.




Fig. 14 Comparison with simple additive regression

Table 1 Sensor description of the C-MAPSS data set (see Li et al. 2019)

No.	Symbol	Description	Units	Included in the package?
1	T2	Total Temperature at fan inlet	$^{\circ}\text{R}$	×
2	T24	Total temperature at LPC outlet	$^{\circ}\text{R}$	✓
3	T30	Total temperature at HPC outlet	$^{\circ}\text{R}$	✓
4	T50	Total temperature LPT outlet	$^{\circ}\text{R}$	✓
5	P2	Pressure at fan inlet	psia	×
6	P15	Total pressure in bypass-duct	psia	×
7	P30	Total pressure at HPC outlet	psia	✓
8	Nf	Physical fan speed	rpm	✓
9	Nc	Physical core speed	rpm	✓
10	Epr	Engine pressure ratio	-	×
11	Ps30	Static pressure at HPC outlet	psia	✓
12	Phi	Ratio of fuel flow to Ps30	pps/psi	✓
13	NRf	Corrected fan speed	rpm	✓
14	NRc	Corrected core speed	rpm	✓
15	BPR	Bypass ratio	-	✓
16	farB	Burner fuel-air ratio	-	×
17	htBleed	Bleed enthalpy	-	✓
18	NF dmd	Demanded fan speed	rpm	×
19	PCNR dmd	Demanded corrected fan speed	rpm	×
20	W31	HPT coolant bleed	lbm/s	✓
21	W32	LPT coolant bleed	lbm/s	✓

Table 2 C-MAPSS data set overview

	FD001	FD002	FD003	FD004
Training Units	100	260	100	249
Testing Units	100	259	100	248

```
R> par(mfrow = c(1, 2))
R> plot(energy.hmm$x[, 1], type = "l", xlab = 'Time',
+      ylab = 'Energy Price', main = "Regime switching
+      additive regression", lwd = 2)
R> time = 1:length(pred2[[1]])
R> predict = (s == 1) * pred2[[1]] + (s == 2) * pred2[[2]]
R> lines(predict ~ time, col = "red", lwd = 2)
R> lines(1:energy.hmm$N, rep(0.5, energy.hmm$N),
+      col = s, lwd = 2, type = "h", pch = 16)
R> text(500, 9, paste("SSE=", round(sse2, 2)))
R> plot(energy.hmm$x[, 1], type = "l", xlab = 'Time',
+      ylab = 'Energy Price', main = "Simple additive
+      regression", lwd = 2)
R> time = 1:length(pred1)
R> lines(pred1 ~ time, col = "blue", lwd = 2)
R> text(500, 9, paste("SSE=", round(sse1, 2)))
```

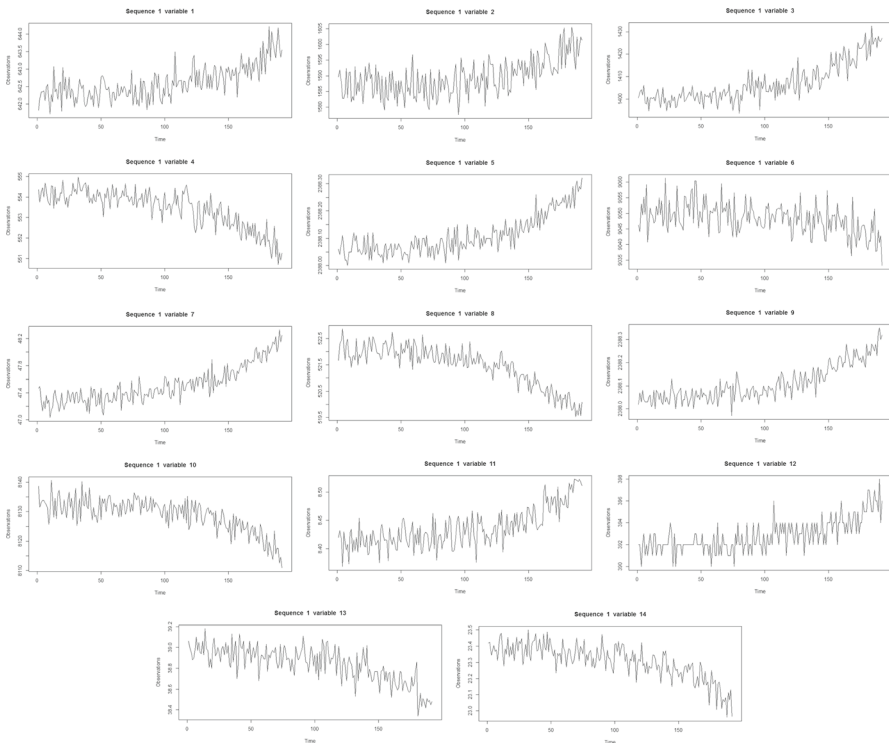


Fig. 15 The time series plot of 14 variables of the first sequence of the train set for the CMAPSS data set

The resulting plot is given in Fig. 14. As one can see from Fig. 14, the two-state regime switching additive regression model, performs much better than the simple (single state) additive regression model.

5.2 RUL estimation for the C-MAPSS data set

The turbofan engine data is from the Prognostic Center of Excellence (PCoE) of NASA Ames Research Center, which is simulated by the Commercial Modular Aero-Propulsion System Simulation (C-MAPSS). Only 14 out of 21 variables are selected, by a method mentioned by Li et al. (2019) and are included in the **hhsmm** package. A list of all 21 variable, as well as a description and the selected 14 variables are tabulated in Table 1. The `train` and `test` lists are of class "hhsmmdata". The original data set contains the subsets FD001-FD004, which are concatenated in the CMAPSS data set. These sets are described in Table 2. This table is presented in `CMAPSS$subsets` in the CMAPSS data set.

We load the data set and extract the `train` and `test` sets as follows.

```
R> data(CMAPSS)
R> train = CMAPSS$train
R> test = CMAPSS$test
```

To visualize the data set, we plot only the first sequence of the `train` set. To do this, this sequence is converted to a data set of class "hhsmmdata", using the function `hhsmmdata` as follows. The plots are presented in Fig. 15.

```
R> train1 = hhsmmdata(x = train$x[1:train$N[1]], N = train$N[1])
R> plot(train1)
```

Initial clustering of the states and mixture components is obtained by the `initial_cluster` function. Since, the suitable reliability model for the CMAPSS data set is a left to right model, the option `ltr = TRUE` is used. Also, since the engines are failed in the final time of each sequence, the final time of each sequence is considered the absorbing state (final state of the left to right model). This assumption is given to the model by the option `final.absorb = TRUE`. The number of states is assumed to be 5 states, which could be one healthy state, 3 levels of damage state, and one failure state in the reliability model. The number of mixture components is computed automatically using the option `nmix = "auto"`.

```
R> J = 5
R> clus = initial_cluster(train = train, nstate = J, nmix = "auto",
+   ltr = TRUE, final.absorb = TRUE, verbose = TRUE)
Within sequence clustering ...
clustering [=====] 100% in 3m
State 1
Between sequence clustering ...
Automatic determination of the number of mixture components ...
State 2
Between sequence clustering ...
Automatic determination of the number of mixture components ...
State 3
Between sequence clustering ...
Automatic determination of the number of mixture components ...
State 4
Between sequence clustering ...
Automatic determination of the number of mixture components ...
State 5
Between sequence clustering ...
Automatic determination of the number of mixture components ...
```

Now, we initialize the model using the `initialize_model` function. The sojourn time distribution is assumed to be "gamma" distribution.

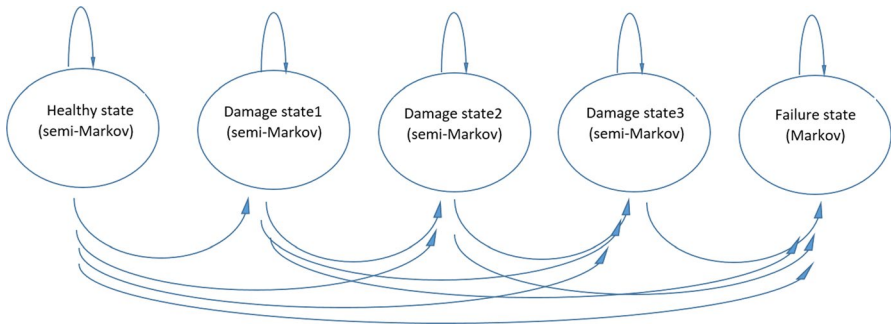


Fig. 16 Graphical representation of the reliability left to right model

```
R> initmodel = initialize_model(clus = clus, sojourn = "gamma",
+ M = max(train$N), verbose = TRUE)
Initial estimation ...
State 1 estimation
Mixture component 1 estimation
Mixture component 2 estimation
Mixture component 3 estimation
Mixture component 4 estimation
Mixture component 5 estimation
...
State 5 estimation
Mixture component 1 estimation
Mixture component 2 estimation
Mixture component 3 estimation
Mixture component 4 estimation
Mixture component 5 estimation
Initializing model ...
```

As a result, the initial estimates of the parameters of the sojourn time distribution and initial estimates of the transition probability matrix and the initial probability vector are obtained as follows.

```
R> initmodel$sojourn
$shape
[1] 10.753944 11.222101 5.617826 8.386559 0.000000
$scale
[1] 0.9138905 0.8885851 1.9869251 23.1578163 0.0000000
$type
[1] "gamma"
R> initmodel$transition
      [,1] [,2]      [,3]      [,4]      [,5]
[1,] 0 0.85 0.0500000 0.05000000 0.05000000
[2,] 0 0.00 0.8947368 0.05263158 0.05263158
[3,] 0 0.00 0.0000000 0.94444444 0.05555556
[4,] 0 0.00 0.0000000 0.00000000 1.00000000
[5,] 0 0.00 0.0000000 0.00000000 1.00000000
R> initmodel$init
[1] 1 0 0 0 0
```

Now, we fit the HHSMM model, using `hhsmmfit` function. The option `lock.init=TRUE` is a good option for a right to left model since the initial state is the first state (healthy system state in the reliability model) in such situations with probability 1. Graphical visualization of such a model is given in Fig. 16.

```

R> fit1 = hhsmmfit(x = train, model = initmodel, M = max(train$N),
+   par = list(lock.init = TRUE))
iteration: 1   log-likelihood = -1813700
iteration: 2   log-likelihood = -1381056
iteration: 3   log-likelihood = -1470325
iteration: 4   log-likelihood = -1471163
iteration: 5   log-likelihood = -1450113
iteration: 6   log-likelihood = -1429393
....
iteration: 79  log-likelihood = -1380078
iteration: 80  log-likelihood = -1379937
iteration: 81  log-likelihood = -1379810
AIC = 2770213
BIC = 2823105

```

The estimates of the transition probability matrix, the sojourn time probability matrix, the initial probability vector, and the AIC and BIC of the model, are extracted as follows.




Fig. 17 The estimated gamma sojourn time density functions

```

R> fit1$model$transition
      [,1]      [,2]      [,3]      [,4] [,5]
[1,]  0 0.1481077 0.7954866 0.05640566  0
[2,]  0 0.0000000 0.0000000 1.00000000  0
[3,]  0 0.0000000 0.0000000 0.00000000  1
[4,]  0 0.0000000 0.0000000 0.00000000  1
[5,]  0 0.0000000 0.0000000 0.00000000  1
R> head(fit1$model$d)
      [,1]      [,2]      [,3]      [,4] [,5]
[1,] 3.302753e-19 5.011609e-12 2.064632e-10 1.101749e-07 1e-100
[2,] 4.526557e-16 2.694189e-10 1.604557e-08 1.111672e-06 1e-100
[3,] 2.957092e-14 2.545406e-09 1.854185e-07 3.737021e-06 1e-100
[4,] 5.455948e-13 1.178541e-08 9.730204e-07 8.376700e-06 1e-100
[5,] 5.017141e-12 3.736808e-08 3.346571e-06 1.528679e-05 1e-100
[6,] 2.975970e-11 9.388985e-08 8.869356e-06 2.464430e-05 1e-100
R> fit1$model$init
[1] 1 0 0 0 0
R> fit1$AIC
[1] 2770213
R> fit1$BIC
[1] 2823105

```

We can plot the estimated gamma sojourn probability density functions as follows.

```

R> J <- 5
R> MM = max(train$N) * 1.2
R> f1 <- function(x) dgamma(x, shape = fit1$model$sojourn$shape[1],
+   scale = fit1$model$sojourn$scale[1])
R> plot(f1, 0, MM, type = "l", xlab = "Time", ylab =
+   "Sojourn_time_gamma_probability_density_function")
R> leg = "state_1"
R> for(j in 2:(J-1)){
+   f <- function(x) dgamma(x, shape = fit1$model$sojourn$shape[j],
+   scale = fit1$model$sojourn$scale[j])
+   xs <- seq(1, MM, 0.1)
+   ys <- sapply(xs, f)
+   lines(xs, ys, col = j)
+   leg = c(leg, paste("state", j))
+ }
R> legend(2*MM/3, max(ys), leg, lty = rep(1, J - 1), col = 1:(J - 1))

```

The resulted plot is shown in Fig. 17.

Now, we obtain the estimates of the RULs, as well as the confidence intervals, by four different methods as follows. These four methods are obtained by combination of two different methods "viterbi" and "smoothing" for the prediction and two different methods "mean" and "max" for RUL estimation and confidence interval computation. The option "viterbi" uses the Viterbi algorithm to find the most likely state sequence, while the option "smoothing" uses the estimates of the state probabilities, using the emission probabilities of the test data set. On the

other hand, in the "mean" method, the mean sojourn time and its standard deviation are used for estimation and confidence interval, while in the "max" method, the maximum probability sojourn time and its quantiles are used (see Sect. 4.7).

```
R> pp1 = predict(fit1, test, method = "viterbi",
+   RUL.estimate = TRUE, confidence = "mean")
R> pp2 = predict(fit1, test, method = "viterbi",
+   RUL.estimate = TRUE, confidence = "max")
R> pp3 = predict(fit1, test, method = "smoothing",
+   RUL.estimate = TRUE, confidence = "mean")
R> pp4 = predict(fit1, test, method = "smoothing",
+   RUL.estimate = TRUE, confidence = "max")
```

As a competitor, we fit the hidden Markov model (HMM) to the data set, which means that we consider all states to be Markovian. To do this, we try fitting the HMM to the `train` set, using the option `semi = rep(FALSE, J)` of the `hhsmmfit` function of the **hhsmm** package. We use the same initial values of the parameters, while we need to use `dens.emission = dmixmvnorm` in the `hhsmm-spec` function, and set the mixture components probabilities equal to 1 (for one mixture component in each state).

```
R> J <- 5
R> init0 <- c(1, 0, 0, 0, 0)
R> P0 <- initmodel$transition
R> b0 = list(mu = list(), sigma = list())
R> for(j in 1:J){
+   b0$mu[[j]] <- Reduce('+', initmodel$parms.emission$mu[[j]]) / J
+   b0$sigma[[j]] <- Reduce('+',
+     initmodel$parms.emission$sigma[[j]]) / J
+ }
R> for(j in 1:J) b0$mix.p[[j]] = 1
R> initmodel <- hhsmm-spec(init = init0, transition = P0,
+   parms.emission = b0, dens.emission = dmixmvnorm,
+   semi = rep(FALSE, J))
R> fit3 = hhsmmfit(train, initmodel, mstep = mixmvnorm_mstep)
iteration: 1  log-likelihood = -110772073
iteration: 2  log-likelihood = -1228418
iteration: 3  log-likelihood = -1246009
iteration: 4  log-likelihood = -1260494
iteration: 5  log-likelihood = -1284620
iteration: 6  log-likelihood = -1329087
iteration: 7  log-likelihood = -1506447
iteration: 8  log-likelihood = -2650675
iteration: 9  log-likelihood = -2652023
iteration: 10 log-likelihood = -2650571
iteration: 11 log-likelihood = -2650527
AIC = 5303215
BIC = 5313999
```

For the fitted HMM model, we estimate the RULs using the aforementioned options.

```
R> pp5 = predict(fit3, test, method = "viterbi",
+   RUL.estimate = TRUE, confidence = "mean")
R> pp6 = predict(fit3, test, method = "viterbi",
+   RUL.estimate = TRUE, confidence = "max")
R> pp7 = predict(fit3, test, method = "smoothing",
+   RUL.estimate = TRUE, confidence = "mean")
R> pp8 = predict(fit3, test, method = "smoothing",
+   RUL.estimate = TRUE, confidence = "max")
```

Now, we use the real values of the RULs, stored in `test$RUL` to compute the coverage probabilities of the confidence intervals of HHSMM and HMM models.

```
R> mean((test$RUL >= pp1$RUL.low) & (test$RUL <= pp1$RUL.up))
[1] 0.7963225
R> mean((test$RUL >= pp2$RUL.low) & (test$RUL <= pp2$RUL.up))
[1] 0.54314
R> mean((test$RUL >= pp3$RUL.low) & (test$RUL <= pp3$RUL.up))
[1] 0.864215
R> mean((test$RUL >= pp4$RUL.low) & (test$RUL <= pp4$RUL.up))
[1] 0.7213579
R> mean((test$RUL >= pp5$RUL.low) & (test$RUL <= pp5$RUL.up))
[1] 0
R> mean((test$RUL >= pp6$RUL.low) & (test$RUL <= pp6$RUL.up))
[1] 0
R> mean((test$RUL >= pp7$RUL.low) & (test$RUL <= pp7$RUL.up))
[1] 0
R> mean((test$RUL >= pp8$RUL.low) & (test$RUL <= pp8$RUL.up))
[1] 0
```

As one can see from the above results, the HHSMM model's coverage probabilities are much better than the HMM ones.

To visualize the results of RUL estimation, we plot the RUL estimates and RUL bounds as follows.

```

R> par(mfrow = c(2,2))
R> plot(test$RUL[order(pp1$RUL)],
+       ylim = c(min(pp1$RUL.low), max(pp1$RUL.up)),
+       pch=16, col = "green", xlab = "unit",
+       ylab = "RUL", main = "Viterbi-mean_method")
R> lines(pp1$RUL.low[order(pp1$RUL)], lty = 2, col = "red")
R> lines(pp1$RUL.up[order(pp1$RUL)], lty = 2,
+       col = "red")
R> lines(pp1$RUL[order(pp1$RUL)], col = "blue")
R> plot(test$RUL[order(pp2$RUL.low)],
+       ylim = c(min(pp2$RUL.low), max(pp2$RUL.up)),
+       pch=16, col = "green", xlab = "unit",
+       ylab = "RUL", main = "Viterbi-max_method")
R> lines(sort(pp2$RUL.low), lty = 2, col = "red")
R> lines(pp2$RUL.up[order(pp2$RUL.low)], lty = 2,
+       col = "red")
R> lines(pp2$RUL[order(pp2$RUL.low)], col = "blue")
R> plot(test$RUL[order(pp3$RUL.low)],
+       ylim = c(min(pp3$RUL.low), max(pp3$RUL.up)),
+       pch=16, col = "green", xlab = "unit",
+       ylab = "RUL", main = "Smoothing-mean_method")
R> lines(sort(pp3$RUL.low), lty = 2, col = "red")
R> lines(pp3$RUL.up[order(pp3$RUL.low)], lty = 2,
+       col = "red")
R> lines(pp3$RUL[order(pp3$RUL.low)], col = "blue")
R> plot(test$RUL[order(pp4$RUL.low)],
+       ylim = c(min(pp4$RUL.low), max(pp4$RUL.up)),
+       pch=16, col = "green", xlab = "unit",
+       ylab = "RUL", main = "Smoothing-max_method")
R> lines(sort(pp4$RUL.low), lty = 2, col = "red")
R> lines(pp4$RUL.up[order(pp4$RUL.low)], lty = 2,
+       col = "red")
R> lines(pp4$RUL[order(pp4$RUL.low)], col = "blue")

```

The resulting plots are presented in Fig. 18. From Fig. 18 and the above coverage probabilities, one can see that the “smoothing” and “max” methods perform better than other methods, in this example.

6 Concluding remarks

This paper presents several examples of the R package **hhsmm**. The scope of application of this package covers simulation, initialization, fitting, and prediction of HMM, HSMM, and HHSMM models, for different types of discrete and continuous sojourn distribution, including shifted Poisson, negative binomial, logarithmic, gamma, Weibull, and log-normal. This package contains density and M-step function for estimation of the emission distribution for different types of emission distribution, including the mixture of multivariate normals and penalized B-spline estimator of the emission distribution, the mixture of linear and additive regression (conditional multivariate normal distributions of the response

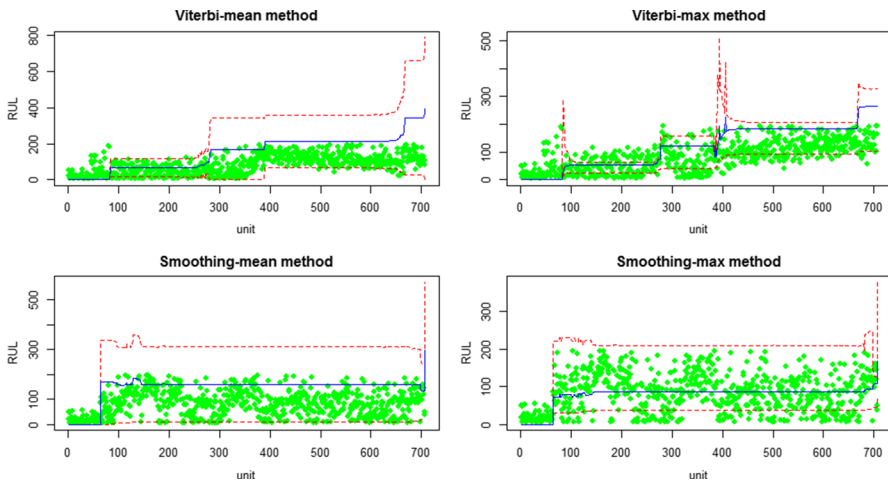


Fig. 18 RUL estimates (solid blue lines) and RUL bounds (dashed red lines) using four different methods for the CMAPSS test data set

given the covariates; regime-switching regression models) as well as the ability to define another emission distributions by the user. As a special case of the regime-switching regression models, the auto-regressive HHSMM models can be modeled by the **hhsmm** package. The left to right models are considered in the **hhsmm** package, especially in the initialization functions. The **hhsmm** package uses the EM algorithm to handle the missing values when the mixture of multivariate normals is considered as the emission distribution. The ability to predict the future states, residual useful lifetime estimation for the left to right models, computation of the score of new observations, computing the homogeneity of two sequences of states, and splitting the data to train and test sequences by the ability to right-trim the test sequences, are other useful features of the **hhsmm** package. The current version 0.3.2 of this package is now available on CRAN (<https://cran.r-project.org/package=hhsmm>), while the future improvements of this package are also considered by the authors. Any report of the possible bugs of the **hhsmm** package are welcome through <https://github.com/mortamini/hhsmm/issues> and we welcome the users' offers for any needed feature of the package in the future.

Appendix

Forward-backward algorithm for the HHSMM model

Denote the sequence $\{Y_t, \dots, Y_s\}$ by $Y_{t:s}$ and suppose that the observation sequence $X_{0:\tau-1}$ with the corresponding hidden state sequence $S_{0:\tau-1}$ is observed. The forward-backward algorithm is an algorithm to compute the probabilities

$$L_j(t) = P(S_t = j | X_{0:\tau-1} = x_{0:\tau-1})$$

within the E-step of the EM algorithm. The above probabilities are computed in the backward recursion of the forward-backward algorithm.

For a semi-Markovian state j , and $t = 0, \dots, \tau - 2$, the forward recursion computes the following probabilities (Guédon 2005),

$$\begin{aligned}
 F_j(t) &= P(S_{t+1} \neq j, S_t = j | X_{0:t} = x_{0:t}) \\
 &= \frac{f_j(x_t)}{N_t} \left[\sum_{u=1}^t \left\{ \prod_{v=1}^{u-1} \frac{f_j(x_{t-v})}{N_{t-v}} \right\} d_j(u) \sum_{i \neq j} p_{ij} F_i(t-u) \right. \\
 &\quad \left. + \left\{ \prod_{v=1}^t \frac{f_j(x_{t-v})}{N_{t-v}} \right\} d_j(t+1) \pi_j \right],
 \end{aligned}
 \tag{A.21}$$

and

$$\begin{aligned}
 F_j(\tau - 1) &= P(S_{\tau-1} = j | X_{0:\tau-1} = x_{0:\tau-1}) \\
 &= \frac{f_j(x_{\tau-1})}{N_{\tau-1}} \left[\sum_{u=1}^{\tau-1} \left\{ \prod_{v=1}^{u-1} \frac{f_j(x_{\tau-1-v})}{N_{\tau-1-v}} \right\} D_j(u) \sum_{i \neq j} p_{ij} F_i(\tau - 1 - u) \right. \\
 &\quad \left. + \left\{ \prod_{v=1}^{\tau-1} \frac{f_j(x_{\tau-1-v})}{N_{\tau-1-v}} \right\} D_j(\tau) \pi_j \right].
 \end{aligned}
 \tag{A.22}$$

where the normalizing factor N_t is computed as follows

$$N_t = P(X_t = x_t | X_{0:t-1} = x_{0:t-1}) = \sum_j P(S_t = j, X_t = x_t | X_{0:t-1} = x_{0:t-1})$$

For a Markovian state j , and for $t = 0$,

$$\tilde{F}_j(0) = P(S_0 = j | X_0 = x_0) = \frac{f_j(x_0)}{N_0} \pi_j$$

and for $t = 1, \dots, \tau - 1$

$$\tilde{F}_j(t) = P(S_t = j | X_{0:t} = x_{0:t}) = \frac{f_j(x_t)}{N_t} \sum_i \tilde{p}_{ij} \tilde{F}_i(t-1) \tag{A.23}$$

The log-likelihood of the model is then

$$\log P(X_{0:\tau-1} = x_{0:\tau-1}; \theta) = \sum_{t=0}^{\tau-1} \log N_t,$$

which is used as a criteria for convergence of the EM algorithm and the evaluate the quality of the model. The backward recursion is initialized by

$$L_j(\tau - 1) = P(S_{\tau-1} = j | X_{0:\tau-1} = x_{0:\tau-1}) = F_j(\tau - 1) = \tilde{F}_j(\tau - 1).$$

For a semi-Markovian state j , we have

$$L_j(\tau - 1) = P(S_{\tau-1} = j | X_{0:\tau-1} = x_{0:\tau-1}) = F_j(\tau - 1),$$

and for $t = \tau - 2, \dots, 0$,

$$L_j(t) = L1_j(t) + L_j(t + 1) - G_j(t + 1) \sum_{i \neq j} p_{ij} \tilde{F}_i(t), \tag{A.24}$$

where

$$L1_j(t) = \sum_{k \neq j} G_k(t + 1) p_{jk} F_j(t), \tag{A.25}$$

and

$$G_j(t + 1) = \sum_{u=1}^{\tau-1-t} G_j(t + 1, u)$$

with

$$G_j(t + 1, u) = \left\{ \prod_{v=0}^{\tau-2-t} \frac{f_j(x_{\tau-1-v})}{N_{\tau-1-v}} \right\} D_j(\tau - 1 - t).$$

For a Markovian state j and for $t = \tau - 2, \dots, 0$

$$L_j(t) = L1_j(t) \tag{A.26}$$

Furthermore, if a mixture of multivariate normal distributions with probability density function

$$f_j(x) = \sum_{k=1}^{K_j} \lambda_{kj} \mathcal{N}_p(x; \mu_{kj}, \Sigma_{kj}), \quad j = 1, \dots, J, \tag{A.27}$$

is considered as the emission distribution, then the following probabilities of the mixture components are computed in the E-step of the $(s + 1)$ th iteration of the EM algorithm

$$\gamma_{kj}^{(s+1)}(t) = \frac{\lambda_{kj}^{(s)} \mathcal{N}_p(x_t; \mu_{kj}^{(s)}, \Sigma_{kj}^{(s)})}{f_j(x_t)}, \tag{A.28}$$

where $\lambda_{kj}^{(s)}$, $\mu_{kj}^{(s)}$ and $\Sigma_{kj}^{(s)}$ are the s th updates of the emission parameters in the M-step of the s th iteration of the EM algorithm.

The M-step of the EM algorithm

In the M-step of the EM algorithm, the initial probabilities are updated as follows

$$\pi_j^{(s+1)} = P(S_0 = j | X_{0:\tau-1} = x_{0:\tau-1}; \theta^{(s)}) = L_j(0),$$

For a semi-Markovian state i , the transition probabilities are updated as follows

$$P_{ij}^{(s+1)} = \frac{\sum_{t=0}^{\tau-2} G_j(t+1) p_{ij}^{(s)} F_i(t)}{\sum_{t=0}^{\tau-2} L_i(t)} \tag{A.29}$$

and for a Markovian state i

$$\tilde{p}_{ij}^{(s+1)} = \frac{\sum_{t=0}^{\tau-2} \tilde{G}_j(t+1) \tilde{p}_{ij}^{(s)} \tilde{F}_i(t)}{\sum_{t=0}^{\tau-2} L_i(t)} \tag{A.30}$$

If we consider the mixture of multivariate normals, with the probability density function (A.27), as the emission distribution, then its parameters are updated as follows

$$\begin{aligned} \lambda_{kj}^{(s+1)} &= \frac{\sum_{t=0}^{\tau-1} \gamma_{kj}^{(s)}(t) L_j^{(s)}(t)}{\sum_{m=1}^{K_j} \sum_{t=0}^{\tau-1} \gamma_{mj}^{(s)}(t) L_j^{(s)}(t)} = \frac{\sum_{t=0}^{\tau-1} \gamma_{kj}^{(s)}(t) L_j^{(s)}(t)}{\sum_{t=0}^{\tau-1} L_j^{(s)}(t)}, \\ \mu_{kj}^{(s+1)} &= \frac{\sum_{t=0}^{\tau-1} \gamma_{kj}^{(s)}(t) L_j^{(s)}(t) x_t}{\sum_{t=0}^{\tau-1} \gamma_{kj}^{(s)}(t) L_j^{(s)}(t)}, \\ \Sigma_{kj}^{(s+1)} &= \frac{\sum_{t=0}^{\tau-1} \gamma_{kj}^{(s)}(t) L_j^{(s)}(t) (x_t - \mu_{kj}^{(s+1)})(x_t - \mu_{kj}^{(s+1)})^T}{\sum_{t=0}^{\tau-1} \gamma_{kj}^{(s)}(t) L_j^{(s)}(t)} \end{aligned}$$

Also, the parameters of the sojourn time distribution are updated by maximization of the following quasi-log-likelihood function

$$\tilde{Q}_d(\{d_j(u)\} | \theta^{(s)}) = \sum_{u=1}^{M_j} \tilde{\eta}_{j,u}^{(s)} \log d_j(u),$$

where

$$\begin{aligned} \tilde{\eta}_{j,u} &= \sum_{t=0}^{\tau-2} G_j(t+1, u) \sum_{i \neq j} p_{ij} F_i(t) + A(u) d_j(u) \pi_j \prod_{v=1}^u \frac{f_j(x_{u_\tau-v})}{N_{u_\tau-v}}, \\ A(u) &= \begin{cases} \frac{L_{ij}(u-1)}{F_j(u-1)}, & u \leq \tau - 1 \\ 1, & u > \tau - 1, \end{cases} \end{aligned}$$

and $u_\tau = \min(u, \tau)$.

Viterbi algorithm and smoothing for the HHSMM model

The Viterbi algorithm (Viterbi 1967) is an algorithm to obtain the most likely state sequence, given the observations and the estimated parameters of the model.

For a semi-Markovian state j , and for $t = 0, \dots, \tau - 2$, the probability of the most probable state sequence is obtained by the Viterbi recursion as follows

$$\begin{aligned} \alpha_j(t) &= \max_{s_0, \dots, s_{t-1}} P(S_{t+1} \neq j, S_t = j, S_{0:t-1} = s_{0:t-1}, X_{0:t} = x_{0:t}) \\ &= f_j(x_t) \max \left[\max_{1 \leq u \leq t} \left[\left\{ \prod_{v=1}^{u-1} f_j(x_{t-v}) \right\} d_j(u) \max_{i \neq j} \{p_{ij} \alpha_i(t-u)\} \right], \right. \\ &\quad \left. \times \left\{ \prod_{v=1}^t f_j(x_{t-v}) \right\} d_j(t+1) \pi_j \right], \end{aligned} \quad (\text{A.31})$$

and

$$\begin{aligned} \alpha_j(\tau-1) &= \max_{s_0, \dots, s_{\tau-2}} P(S_{\tau-1} = j, S_{0:\tau-2} = s_{0:\tau-2}, X_{0:\tau-1} = x_{0:\tau-1}) \\ &= f_j(x_{\tau-1}) \max \left[\max_{1 \leq u \leq \tau-1} \left[\left\{ \prod_{v=1}^{u-1} f_j(x_{\tau-1-v}) \right\} D_j(u) \max_{i \neq j} \{p_{ij} \alpha_i(\tau-1-u)\} \right], \right. \\ &\quad \left. \times \left\{ \prod_{v=1}^{\tau-1} f_j(x_{\tau-1-v}) \right\} D_j(\tau) \pi_j \right], \end{aligned} \quad (\text{A.32})$$

For a Markovian state j , the Viterbi recursion is initialized by

$$\tilde{\alpha}_j(0) = P(S_0 = j, X_0 = x_0) = f_j(x_0) \pi_j$$

and for $t = 1, \dots, \tau - 1$,

$$\tilde{\alpha}_j(t) = \max_{s_0, \dots, s_{t-1}} P(S_t = j, S_0^{t-1} = s_0^{t-1}, X_{0:t} = x_{0:t}) \quad (\text{A.33})$$

$$= f_j(x_t) \max_i \{\tilde{p}_{ij} \tilde{\alpha}_i(t-1)\}. \quad (\text{A.34})$$

After obtaining the probability of the most probable state sequence, the current most likely state is obtained as $\hat{s}_t^* = \arg \max_{1 \leq j \leq J} \alpha_j(t)$.

Another approach for obtaining the state sequence is the smoothing method, which uses the backward probabilities $L_j(t)$ instead of $\alpha_j(t)$.

Acknowledgements The authors would like to thank the two anonymous referees and the associate editor for their useful comments and suggestions, which improved an earlier version of the **hhsmm** package and this paper.

References

- Adam T, Langrock R, Weiß CH (2019) Penalized estimation of flexible hidden Markov models for time series of counts. *Metron* 77(2):87–104
- Azimi M (2004) *Data transmission schemes for a new generation of interactive digital television* (Doctoral dissertation, University of British Columbia)
- Bulla J, Bulla I, Nenadic O (2010) hhsmm an R package for analyzing hidden semi-Markov models. *Comput Stat Data Anal* 54(3):611–619
- Cartella F, Lemeire J, Dimiccoli L, Sahli H (2015) Hidden semi-Markov models for predictive maintenance. *Math Probl Eng*
- Cook AE, Russell MJ (1986) Improved duration modeling in hidden Markov models using series-parallel configurations of states. *Proc Inst Acoust* 8:299–306
- Dempster AP, Laird NM, Rubin DB (1977) Maximum likelihood from incomplete data via the EM algorithm. *J R Stat Soc Ser B (Methodol)* 39(1):1–22
- Durbin R, Eddy SR, Krogh A, Mitchison G (1998) *Biological sequence analysis: probabilistic models of proteins and nucleic acids*. Cambridge University Press, Cambridge
- Fontdecaba S, Muñoz MP, Sánchez JA (2009) Estimating Markovian Switching Regression Models in An application to model energy price in Spain. In *The R User Conference*, France
- Guédon Yann (2005) Hidden hybrid Markov/semi-Markov chains. *Comput Stat Data Anal* 49(3):663–688
- Harte D (2006) Mathematical background notes for package “HiddenMarkov”. *Statistics Re.*
- Jackson C (2007) Multi-state modelling with R: the msm package. Cambridge, UK, pp 1–53
- Kim CJ, Piger J, Startz R (2008) Estimation of Markov regime-switching regression models with endogenous switching. *J Econom* 143(2):263–273
- Langrock R, Kneib T, Sohn A, DeRuiter SL (2015) Nonparametric inference in hidden Markov models using P-splines. *Biometrics* 71(2):520–528
- Langrock R, Adam T, Leos-Barajas V, Mews S, Miller DL, Papastamatiou YP (2018) Spline-based non-parametric inference in general state-switching models. *Statistica Neerlandica* 72(3):179–200
- Li J, Li X, He D (2019) A directed acyclic graph network combined with CNN and LSTM for remaining useful life prediction. *IEEE Access* 7:75464–75475
- Lloyd S (1982) Least squares quantization in PCM. *IEEE Trans Inf Theory* 28(2):129–137
- ÓConnell J, Højsgaard S (2011) Hidden semi Markov models for multiple observation sequences: the mhsmm package for R. *J Stat Softw* 39(4):1–22
- R Development Core Team (2010) *R: a Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria
- Schellhase C, Kauermann G (2012) Density estimation and comparison with a penalized mixture approach. *Comput Stat* 27:757–777
- Van Buuren S, Groothuis-Oudshoorn K (2011) mice: multivariate imputation by chained equations in R. *J Stat Softw* 45:1–67
- Visser I, Speekenbrink M (2010) depmixS4: an R package for hidden Markov models. *J Stat Softw* 36(7):1–21
- Viterbi A (1967) Error bounds for convolutional codes and an asymptotically optimum decoding algorithm. *IEEE Trans Inf Theory* 13(2):260–269

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.