

Package ‘hmm.discnp’

November 26, 2018

Version 2.1-5

Date 2018-11-26

Title Hidden Markov Models with Discrete Non-Parametric Observation Distributions

Author Rolf Turner

Maintainer Rolf Turner <r.turner@auckland.ac.nz>

Depends R (>= 2.10)

Imports nnet (>= 7.3.12)

Description Fits hidden Markov models with discrete non-parametric observation distributions to data sets. The observations may be univariate or bivariate. Simulates data from such models. Finds most probable underlying hidden states, the most probable sequences of such states, and the log likelihood of a collection of observations given the parameters of the model. Auxiliary predictors are accommodated in the univariate setting.

LazyData true

ByteCompile true

License GPL (>= 2)

NeedsCompilation yes

Repository CRAN

Date/Publication 2018-11-26 05:50:03 UTC

R topics documented:

anova.hmm.discnp	2
ccprSim	3
cnvrtRho	4
fitted.hmm.discnp	5
hmm	7
hydroDat	16

lesionCount	17
logLikHmm	18
misstify	20
mps	22
nafracCalc	23
pr	24
predict.hmm.discnp	26
rhmm	27
scovmat	29
sp	30
squantCI	32
SydColDisc	34
update.hmm.discnp	35
viterbi	37
weissData	39

Index	42
--------------	-----------

anova.hmm.discnp	<i>Anova for hmm.discnp models</i>
------------------	------------------------------------

Description

Performs a likelihood ratio test to compare two discrete non-parametric hidden Markov models.

Usage

```
## S3 method for class 'hmm.discnp'
anova(object, ...)
```

Arguments

object	An object of class “hmm.discnp” as returned by the function <code>hmm()</code> .
...	A second object of class “hmm.discnp”. There must be only <i>one</i> such object.

Value

A list with entries

stat	The likelihood ratio statistic.
df	The degrees of freedom.
pvalue	The p-value of the test.

This list has an attribute “details” which is a vector consisting of the first and second log likelihoods and the associated numbers of parameters, in order of these numbers of parameters. (See **Warning**.)

Warning

Hidden Markov models can be numerically delicate and the fitting algorithm can converge to a local maximum of the likelihood surface which is not the global maximum. Thus it is entirely possible for the log likelihood of the model with the greater number of parameters to be *smaller* than that of the model with the lesser number of parameters.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

[hmm\(\)](#)

Examples

```
xxx <- with(SydColDisc, split(y, f=list(locn, depth)))
fit1 <- hmm(xxx, K=1, itmax=10)
fit2 <- hmm(xxx, K=2, itmax=10)
anova(fit1, fit2)
```

ccprSim

Simulated monocyte counts and psychosis symptoms.

Description

Discretised values of monocyte counts, and ratings of level of psychosis simulated from a model fitted to a data set consisting of observations made on a number of patients from the Northern District Health Board system. The real data must be kept confidential due to ethics constraints.

Usage

```
data("ccprSim")
```

Format

The object `ccprSim` is a list of length 1258. Each entry of this list is to be considered to correspond to an individual subject. The entries consist of matrices having two columns named `cellCount` and `psychosisRating`. The number of rows of these matrices varies from entry to entry of the list (i.e. from subject to subject).

Most of the entries of these matrices are NA. The entries are temporally ordered and correspond to the number of weeks from the start of observation. Observations in the real data set were made only when the patient in question visited a physician and so weeks in which no visit was made resulted in an "observation" of NA. The object `ccprSim` was simulated in such a way as to imitate this characteristic. The fraction of missing observations in each variate (i.e. `cellCount` and `psychosisRating`) is roughly commensurate with the corresponding fractions in the real data.

The values in the first column of each matrix (the cellCount column consist of integers from 1 to 5 and are to be interpreted as indicators of cell counts in units of 10^9 cells per litre, discretised according to the following scale:

- $0.0 \leq c \leq 0.3 \leftrightarrow 1$
- $0.3 < c \leq 0.5 \leftrightarrow 2$
- $0.5 < c \leq 0.7 \leftrightarrow 3$
- $0.7 < c \leq 1.0 \leftrightarrow 4$
- $1.0 < c \leq 2.0 \leftrightarrow 5$

where c represents “count”.

The values in the second column of each matrix (the psychosisRating column consist of integers from 0 to 4 and are to be interpreted as indicators of a physician’s assessment of the level of psychosis of the patient. A value of 0 corresponds to “no symptoms”; a value of 4 corresponds to “severe”.

The question of essential interest in respect of the real data was “Is there any association between the cell count values and the psychosis ratings?” More specifically it was “Can the level of psychosis be *predicted* from the cell counts?”

Source

The real data, on the basis of which these data were simulated, were supplied by Dr. Jonathan Williams, Northern District Health Board.

Examples

```
## Not run:
# Takes a long time:
  fit <- hmm(ccprSim,K=2,indep=FALSE,itmax=500,verbose=TRUE)
# Throws an error; inadequate data:
  subsetFit <- hmm(ccprsim[1021:1025],K=2,indep=FALSE,itmax=5)

## End(Not run)
```

cnvrtRho

Convert Rho between forms.

Description

Converts the “old style” (matrix) specification of the emission probabilities to the “new style” (data frame) specification, or vice versa.

Usage

```
cnvrtRho(Rho)
```

Arguments

Rho A specification of the emission probabilities of a discrete valued hidden Markov model. It may be either a matrix of these probabilities, in which case it is converted to a three column data frame, or it may be a three column data frame, in which case it is converted to a matrix of probabilities. See `hmm()` for more details about the structure of Rho, in either form.

Details

The “new style” specification of Rho allows for predictor variables x . If Rho is of the “new style” and is designed to allow for predictor variables, then it will have more than three columns and cannot be converted to the “old style”. In such cases `cnvrtRho` will throw an error.

Value

A data frame if the argument Rho is a matrix (“old style” specification), or a matrix if the argument Rho is a data frame (“new style” specification).

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

[hmm\(\)](#)

Examples

```
Yval <- LETTERS[1:10]
Tpm <- matrix(c(0.75,0.25,0.25,0.75),ncol=2,byrow=TRUE)
Rho <- cbind(c(rep(1,5),rep(0,5)),c(rep(0,5),rep(1,5)))/5
rownames(Rho) <- Yval
newRho <- cnvrtRho(Rho)
oldRho <- cnvrtRho(newRho)
```

fitted.hmm.discnp *Fitted values of a discrete non-parametric hidden Markov model.*

Description

Calculates the fitted values of a discrete non-parametric hidden Markov model as the conditional expectations of the observations, given the entire observation sequence (and the estimated parameters of the model).

Usage

```
## S3 method for class 'hmm.discnp'
fitted(object, ...)
```

Arguments

object An object of class `hmm.discnp` as returned by `hmm()`.
 ... Not used.

Details

This function is essentially a wrapper for `sp()`. This latter function gets called with `mean=TRUE`. The fitted values are the conditional expected values of the observations, given the entire observation sequence. Note that the “conditional expected value” concept makes sense *only* if the observations are *numeric*. If this is not the case, an error is thrown. The observation sequence must be present in object (which will be the case if object was returned by `hmm()`). If it is not present an error is thrown.

However, if such an error is thrown, do not despair! You *do not* have to start from scratch when fitting your model with `keep.y=TRUE`. If `fit` is your fitted model that you obtained *without* setting `keep.y=TRUE`, then you can just re-fit the model using `fit` as the starting values:

```
fit2 <- hmm(<whatever>,par0=fit,keep.y=TRUE)
```

This will of course converge instantaneously. You could also do:

```
fit2 <- update(fit,data=<whatever>,keep.y=TRUE)
```

Value

If the observations (stored in object) consist of a single sequence, then the object returned is a matrix. If the model is univariate (see `hmm()`) then matrix has a single column constituting the corresponding sequence of fitted values. If the model is bivariate (either independent or dependent) then the matrix has two columns corresponding respectively to the two variables in the bivariate model. If the observations consist of a list of observation sequences, then the object returned is a list of corresponding matrices whose columns are the fitted values.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

`sp()` `link{predict.hmm.discnp}()`

Examples

```
P <- matrix(c(0.7,0.3,0.1,0.9),2,2,byrow=TRUE)
R <- matrix(c(0.5,0,0.1,0.1,0.3,
             0.1,0.1,0,0.3,0.5),5,2)
set.seed(42)
lll <- sample(250:350,20,TRUE)
```

```

y <- rhmm(ylengths=111,nsim=1,drop=TRUE,tpm=P,Rho=R)
fit <- hmm(y,K=2,verb=TRUE,keep.y=TRUE,itmax=10)
fv <- fitted(fit)

```

hmm

Fit a hidden Markov model to discrete data.

Description

Effects a maximum likelihood fit of a hidden Markov model to discrete data where the observations come from one of a number of finite discrete distributions, depending on the (hidden) state of the Markov chain. These distributions (the “emission probabilities”) are specified non-parametrically. The observations may be univariate, independent bivariate, or dependent bivariate. By default this function uses the EM algorithm. In the univariate setting it may alternatively use a “brute force” method.

Usage

```

hmm(y, yval=NULL, par0=NULL, K=NULL, rand.start=NULL,
     method=c("EM", "bf", "LM", "SD"), hglmethod=c("fortran", "oraw", "raw"),
     optimiser=c("nlm", "optim"), optimMethod=NULL, stationary=cis,
     mixture=FALSE, cis=TRUE, indep=NULL, tolerance=1e-4, digits=NULL,
     verbose=FALSE, itmax=200, crit=c("PCLL", "L2", "Linf", "ABSGRD"),
     keep.y=FALSE, keep.X=keep.y, newstyle=TRUE, X=NULL,
     addIntercept=TRUE, lmc=10, hessian=FALSE,...)

```

Arguments

- | | |
|------|---|
| y | A one or two column matrix of discrete data or a list of such matrices; missing values are allowed. If y is a vector, or list of vectors (of discrete data) these vectors are coerced to one column matrices. |
| yval | A vector (of length m, say) of possible values for the data or a list of two such vectors (of lengths m1 and m2, say, one for each of the two variates in the bivariate settings). These vectors default to the sorted unique values of the respective variates as provided in y. If yval is supplied and any value of y does not match some value of yval, then an error is thrown. |
| par0 | An optional (<i>named</i>) list of starting values for the parameters of the model, with components tpm (transition probability matrix), <i>optionally</i> ispd (initial state probability distribution) and Rho. The object Rho specifies the probability that the observations take on each of the possible values of the variate or variates, given the state of the hidden Markov chain. See Details . Note that in the case of independent bivariate data Rho is a list of two matrices. These matrices may (and in general will) have different row dimensions, but must have identical column dimensions (equal to K, the number of states; see below).

If the model is stationary (i.e. if stationary is TRUE) then you should almost surely not specify the ispd component of par0. If you do specify it, it really only makes sense to specify it to be the stationary distribution determined by |

	<p>tpm and this is a waste of time since this is what the code will take ispd to be if you leave it unspecified.</p> <p>If par0 is not specified, starting values are created by the function <code>init.all()</code>.</p>
K	<p>The number of states in the hidden Markov chain; if par0 is not specified K <i>MUST</i> be; if par0 is specified, K is ignored.</p> <p>Note that K=1 is acceptable; if K is 1 then all observations are treated as being independent and the non-parametric estimate of the distribution of the observations is calculated in the dQuoteobvious way.</p>
rand.start	<p>A list consisting of two logical scalars which must be named tpm and Rho, if tpm is TRUE then the function <code>init.all()</code> chooses entries for then starting value of tpm at random; likewise for Rho. This argument defaults to <code>list(tpm=FALSE,Rho=FALSE)</code>.</p>
method	<p>Character string, either "bf", "EM", "LM" or "SD" (i.e. use numerical maximisation via either <code>nlm()</code> or <code>optim()</code>, the EM algorithm, the Levenberg-Marquardt algorithm, or the method of steepest descent). May be abbreviated. Currently the "bf", "LM" and "SD" methods can be used only in the univariate setting, handle only stationary models (see below) and do not do mixtures.</p>
hglmethod	<p>Character string; one of "fortran", "oraw" or "raw". May be abbreviated. This argument determines the procedure by which the hessian, gradient and log likelihood of the model and data are calculated. If this is argument is equal to "fortran" (the default) then (obviously!) dynamically loaded fortran subroutines are used. The other two possibilities effect the calculations in raw R; "oraw" ("o" for "original" uses code that is essentially a direct transcription of the fortran code, do-loops being replaced by for-loops. With method "raw" the for-loops are eliminated and matrix-vector calculations are applied. The "oraw" method is about 25 times slower than the "fortran" method and the "raw" method is more than 30 times slower. The "raw" methods are present mainly for debugging purposes and would not usually be used in practice. This argument is used only if the method is "LM" or "SD" (and is involved only peripherally in the latter instance). It is ignored otherwise.</p>
optimiser	<p>Character string specifying the optimiser to use when the ""bf"" method of optimisation is chosen. It should be one of "nlm" or "optim", and may be abbreviated. Ignored unless method="bf".</p>
optimMethod	<p>Character string specifying the optimisation method to be used by <code>optim()</code>. Should be one of "Nelder-Mead", "BFGS", "CG", "L-BFGS-B", "SANN", or "Brent". Ignored if the method is not "bf" or if the optimiser is not "optim".</p>
stationary	<p>Logical scalar. If TRUE then the model is fitted under the stationarity assumption, i.e. that the Markov chain was in steady state at the time that observations commenced. In this case the initial state probability distribution is estimated as the stationary distribution determined by the (estimated) transition probability matrix. Otherwise if cis (see below) is TRUE the initial state probability distribution is estimated as the mean of the vectors of conditional probabilities of the states, given the observation sequences, at time t=1. If stationary is TRUE and cis is FALSE an error is thrown. Currently if the method is "bf", "LM" or "SD", and stationary is FALSE, then an error is thrown.</p>
mixture	<p>A logical scalar; if TRUE then a mixture model (all rows of the transition probability matrix are identical) is fitted rather than a general hidden Markov model.</p>

	Currently an error is thrown if <code>mixture=TRUE</code> and the method is "bf", "LM" or "SD".
<code>cis</code>	A logical scalar specifying whether there should be a constant initial state probability distribution. If <code>stationary</code> is FALSE and <code>cis</code> is FALSE then the initial state probability distribution for a given observation sequence is equal to 1 where the (first) maximum of the vector of conditional probabilities of the states, given the observation sequences, at time $t=1$, occurs, and is 0 elsewhere. If <code>stationary</code> is TRUE and <code>cis</code> is FALSE an error is given.
<code>indep</code>	Logical scalar. Should the bivariate model be fitted under the assumption that the two variables are (conditionally) independent given the state? If this argument is left as NULL its value is inferred from the structure of <code>Rho</code> in <code>par0</code> if the latter is supplied. If the data are bivariate and neither <code>indep</code> nor <code>par0</code> is supplied, then an error is given. If the data are bivariate and if the value of <code>indep</code> is inconsistent with the structure of <code>par0</code> \$ <code>Rho</code> then an error is given. If the data are univariate then <code>indep</code> is ignored.
<code>tolerance</code>	If the value of the quantity used for the stopping criterion is less than <code>tolerance</code> then the algorithm is considered to have converged. Ignored if <code>method="bf"</code> . Defaults to $1e-4$.
<code>digits</code>	Integer scalar. The number of digits to which to print out "progress reports" (when <code>verbose</code> is TRUE). There is a "sensible" default (calculated from <code>tolerance</code>). Not used if the method is "bf".
<code>verbose</code>	A logical scalar determining whether to print out details of the progress of the algorithm. If the method is "EM", "LM" or "SD" then when <code>verbose</code> is TRUE information about the convergence criteria is printed out at every step that the algorithm takes. If <code>method="bf"</code> then the value of <code>verbose</code> determines the value of the argument <code>print.level</code> of <code>nlm()</code> or the value of the argument <code>trace</code> of <code>optim()</code> . In the first case, if <code>verbose</code> is TRUE then <code>print.level</code> is set to 2, otherwise it is set to 0. In the second case, if <code>verbose</code> is TRUE then <code>trace</code> is set to 6, otherwise it is set to 0.
<code>itmax</code>	When the method is "EM", "LM" or "SD" this is the maximum number of steps that the algorithm takes. If the convergence criterion has not been met by the time <code>itmax</code> steps have been performed, a warning message is printed out, and the function stops. A value is returned by the function anyway, with the logical component "converged" set to FALSE. When <code>method="bf"</code> the <code>itmax</code> argument is passed to <code>nlm()</code> as the value of <code>iterlim</code> or to <code>optim()</code> as the value of <code>maxit</code> . If the (somewhat obscure) convergence criteria of <code>nlm()</code> or <code>optim()</code> have not been met by the time <code>itmax</code> "iterations" have been performed, the algorithm ceases. In this case, if <code>nlm()</code> is used. the value of <code>code</code> in the object returned set equal to 4 and if <code>optim()</code> is used then the value of <code>convergence</code> returned is set equal to 1. Note that the value of <code>code</code> , respectively <code>convergence</code> is returned as the <code>converged</code> component of the object returned by <code>hmm()</code> . A value of 1 indicates successful completion of the <code>nlm()</code> procedure. A value of 0 indicates successful completion of the <code>optim()</code> procedure.
<code>crit</code>	The name of the stopping criterion used. When <code>method="EM"</code> it must be one of "PCLL" (percent change in log-likelihood; the default), "L2" (L-2 norm, i.e. square root of sum of squares of change in coefficients), or "Linf" (L-infinity norm, i.e. maximum absolute value of change in coefficients). When

	<p>method="LM" or method="SD" there is a fourth possibility, namely "ABSGRD" the (maximum) absolute value of the gradient. It may not be advisable to use this criterion in the current context (i.e. that of discrete non-parametric distributions). See Warnings. This argument defaults to "PCLL". It is ignored if method="bf". (The <code>nlm()</code> and <code>optim()</code> functions have their own obscure stopping criteria.)</p>
keep.y	Logical scalar; should the observations <code>y</code> be returned as a component of the value of this function?
keep.X	Logical scalar; should the predictors <code>X</code> be returned as a component of the value of this function? Note that the value of <code>keep.X</code> will be silently set equal to FALSE unless it actually "makes sense" to keep <code>X</code> . I.e. unless the observations are <i>univariate</i> , <code>newstyle</code> is TRUE, and <code>X</code> is actually supplied, i.e. is not NULL.
newstyle	Logical scalar; if TRUE then the "new style" of parameterising <code>Rho</code> (see Details) is used, in the univariate setting. Note that if <code>X</code> is not NULL or if the method is not "EM" then <code>newstyle</code> <i>cannot</i> be set to FALSE; if it is, then an error is thrown.
X	<p>An optional <i>numeric</i> matrix, or a list of such matrices, of <i>predictors</i>. The use of such predictors is (currently, at least) applicable only in the univariate emissions setting, and then only if <code>Rho</code> is supplied in the "newstyle" format. If <code>X</code> is a list it must be of the same length as <code>y</code> and all entries of this list must have the same number of columns. If the columns of any entry of the list are named, then they must be named for <i>all</i> entries, and the column names must be the <i>same</i> for all entries. The number of rows of each entry must be equal to the length of the corresponding entry of <code>y</code>. If <code>X</code> is a matrix then <code>y</code> should be a vector or one-column matrix (or a list with a single entry equal to such).</p> <p>There may be at most one constant column in <code>X</code> or the components thereof. If there are <i>any</i> constant columns there must be precisely one (in all components of <code>X</code>), it must be the first column and all of its entries must be equal to 1. If the columns have names, the names of this first column must be "Intercept".</p> <p>Note that <code>X</code> (or its entries) must be a <i>numeric</i> matrix (or must be matrices) — no data frames! Factor predictors are not permitted. It may be possible to use factor predictors by supplying <code>X</code> or its entries as the output of <code>model.matrix()</code>; this will depend on circumstances.</p>
addIntercept	Logical scalar. Should a column of ones, corresponding to an intercept term, be prepended to each of the matrices in the list <code>X</code> ? (The user should remember that this argument defaults to TRUE.)
lmc	Numeric scalar. The (initial) "Levenberg-Marquardt correction" parameter. Used only if method="LM", otherwise ignored.
hessian	Logical scalar. Should the hessian matrix be returned? This argument is relevant only if method="bf" (in which case it is passed along to <code>hmmNumOpt()</code>) and is ignored otherwise. This argument should be set to TRUE only if you <i>really</i> want the hessian matrix. Setting it to TRUE causes a substantial delay between the time when <code>hmm()</code> finishes its iterations and when it actually returns a value.
...	Additional arguments passed to <code>hmmNumOpt()</code> . There is one noteworthy argument <code>useAnalGrad</code> which is used "directly" by <code>hmmNumOpt()</code> . This argument is a logical scalar and if it is TRUE then calls to <code>nlm()</code> or <code>optim()</code> are structured so that an analytic calculation of the gradient vector (implemented by the internal

function `get.gl()` is applied. If it is FALSE then finite difference methods are used to calculate the gradient vector. If this argument is not specified it defaults to FALSE. Note that the name of this argument **cannot be abbreviated**.

Other “additional arguments” may be supplied for the control of `nlm()` and are passed on appropriately to `nlm()`. These are used only if `method="bf"` and if `optimiser="nlm"`. These “...” arguments might typically include `gradtol`, `stepmax` and `steptol`. They should **NOT** include `print.level` or `iterlim`. The former argument is automatically passed to `nlm()` as `0` if `verbose` is FALSE and as `2` if `verbose` is TRUE. The latter argument is automatically passed to `nlm()` with the value of `itmax`.

Details

In the univariate case the emission probabilities are, in the case in which `newstyle` is TRUE (the default), specified by means of a data frame `Rho`. The first column of `Rho`, named `y`, is a factor consisting of the possible values of the emissions, repeated `K` times (where `K` is the number of states). The second column, named `states`, is a factor consisting of integer values `1, 2, ..., K`. Each of these values is repeated `m` times where `m` is the length of `yval`. Further columns of `Rho` are numeric and consist of coefficients of the linear predictor of the probabilities of the various values of `y`. If `X` is NULL then `Rho` has only one further column named `Intercept`.

If `X` is not NULL then the `Intercept` column is present only if `addIntercept` is TRUE. There are many (other, in addition to the possible `Intercept` column) numeric columns as there are columns in `X` or in the matrices in the list `X`. The names of these columns are taken to be the column names of `X` or the *first* entry of `X` if such column names are present. Otherwise the names default to `V1, V2, ...`

The probabilities of the emissions taking on their various possible values are given by

$$\Pr(Y = y_i | \mathbf{x}, \text{state} = S) = \ell_i / \sum_{j=1}^m \ell_j$$

where ℓ_j is the j th entry of $\beta^\top \mathbf{x}$ and where in turn \mathbf{x} is the vector of predictors and β is the coefficient vector in the linear predictor that corresponds to y_i and the hidden state S . For identifiability the vectors β corresponding to the first value of Y (the first level of `Rho$y`) are set equal to the zero vector for all values of the state S .

If `newstyle` is FALSE `Rho` is specified by a matrix `Rho = [\rho_{ij}]` where $\rho_{ij} = \Pr(Y = y_i | S = j)$. If `newstyle` is FALSE and if `X` is not NULL then an error is thrown.

In the independent bivariate case the emission probabilities are specified by a list of two such matrices. In this setting $P(Y_1, Y_2) = (y_{i1}, y_{i2}) | S = j = \rho_{i_1, j}^{(1)} \rho_{i_2, j}^{(2)}$ where $R^{(k)} = [\rho_{ij}^{(k)}]$ ($k = 1, 2$) are the two emission probability matrices.

In the independent bivariate case the emission probabilities are specified by a list of two such matrices. In this setting $P(Y_1, Y_2) = (y_{i1}, y_{i2}) | S = j = \rho_{i_1, j}^{(1)} \rho_{i_2, j}^{(2)}$ where $R^{(k)} = [\rho_{ij}^{(k)}]$ ($k = 1, 2$) are the two emission probability matrices. In the dependent bivariate case the emission probabilities are specified by a three dimensional array. In this setting $P((Y_1, Y_2) = (y_{i1}, y_{i2}) | S = j) = \rho_{i_1, i_2, j}$ where $R_k = [\rho_{ij}^{(k)}]$ is the emission probability array.

The hard work of calculating the recursive probabilities used to fit the model is done by a Fortran subroutine "recurse" (actually coded in Ratfor) which is dynamically loaded. In the univariate case,

when X is provided, the estimation of the “linear predictor” vectors β is handled by the function `multinom()` from the `nnet` package. Note that this is a “Recommended” package and is thereby automatically available (i.e. does not have to be installed).

Value

A list with components:

<code>Rho</code>	The fitted value of the matrix, data frame, list of two matrices, or array <code>Rho</code> specifying the distributions of the observations (the “emission” probabilities).
<code>tpm</code>	The fitted value of the transition probability matrix <code>tpm</code> .
<code>ispd</code>	The fitted initial state probability distribution, or a matrix of (trivial or “deterministic”) initial state probability distributions, one (column) of <code>ispd</code> for each observation sequence. If <code>stationary</code> is <code>TRUE</code> then <code>ispd</code> is assumed to be the (unique) stationary distribution for the chain, and thereby determined by the transition probability matrix <code>tpm</code> . If <code>stationary</code> is <code>FALSE</code> and <code>cis</code> is <code>TRUE</code> then <code>ispd</code> is estimated as the mean of the vectors of conditional probabilities of the states, given the observation sequences, at time $t=1$. If <code>cis</code> is <code>FALSE</code> then <code>ispd</code> is a matrix whose columns are trivial probability vectors, as described above.
<code>log.like</code>	The final (maximal, we hope!) value of the log likelihood, as determined by the maximisation procedure.
<code>grad</code>	The gradient of the log likelihood. Present only if the method is “LM” or “bf” and in the latter case then only if the optimiser is <code>nlm()</code> .
<code>hessian</code>	The hessian of the log likelihood. Present only if the method is “LM” or “bf”.
<code>stopCrit</code>	A vector of the (final) values of the stopping criteria, with names “PCLL”, “L2”, “Linf” unless the method is “LM” or “SD” in which case this vector has a fourth entry named “ABSGRD”.
<code>par0</code>	The starting values used by the algorithms. Either the argument <code>par0</code> , or a similar object with either or both components (<code>tpm</code> and <code>Rho</code>) being created by <code>rand.start()</code> .
<code>npar</code>	The number of parameters in the fitted model. Equal to <code>nispar + ntpmpar + nrhopar</code> where (1) <code>nispar</code> is 0 if <code>stationary</code> is <code>TRUE</code> and is $K-1$ otherwise; (2) <code>ntpmpar</code> is $K*(K-1)$ (3) <code>nrhopar</code> is <ul style="list-style-type: none"> • $K*(nrow(Rho)-1)$ for univariate models with <code>newstyle</code> equal to <code>FALSE</code> • $(nrow(Rho) - K)*(ncol(Rho)-2)$ for univariate models with <code>newstyle</code> equal to <code>TRUE</code> • $K*(sum(sapply(Rho, nrow))-K)$ for bivariate independent models • $prod(dim(Rho))-K$ for bivariate dependent models.
<code>bicm</code>	Numeric scalar. The number by which <code>npar</code> is multiplied to form the BIC criterion. It is essentially the log of the number of observations. See the code of <code>hmm()</code> for details.

converged	<p>A logical scalar indicating whether the algorithm converged. If the EM, LM or steepest descent algorithm was used it simply indicates whether the stopping criterion was met before the maximum number (<code>itmax</code>) of steps was exceeded. If <code>method="bf"</code> then <code>converged</code> is based on the code component of the object returned by the optimiser when <code>nlm()</code> was used, or on the convergence component when <code>optim()</code> was used. In these cases <code>converged</code> has an <i>attribute</i> (code or convergence respectively) giving the (integer) value of the relevant component.</p> <p>Note that in the <code>nlm()</code> case a value of code equal to 2 indicates “probable” convergence, and a value of 3 indicates “possible” convergence. However in this context <code>converged</code> is set equal to TRUE <i>only</i> if code is 1.</p>
nstep	The number of steps performed by the algorithm if the method was "EM", "LM" or "SD". The value of <code>nstep</code> is set equal to the <code>iterations</code> component of the value returned by <code>nlm()</code> if <code>method="bf"</code> .
prior.emsteps	The number of EM steps that were taken before the method was switched from "EM" to "bf" or to "LM". Present only in values returned under the "bf" or "LM" methods after a switch from "EM" and is equal to 0 if either of these methods was specified in the initial call (rather than arising as the result of a switch).
ylengths	Integer vector of the lengths of the observation sequences (number of rows if the observations are in the form of one or two column matrices).
nafrac	A real number between 0 and 1 or a pair (two dimensional vector) of such numbers. Each number is the the fraction of missing values if the corresponding components of the observations.
y	A (slightly tidied up) version of the observations (argument <code>y</code>). Present only if <code>keep.y</code> is TRUE.
X	A (slightly tidied up) version of the predictor matrix or list of predictor matrices (argument <code>X</code>). Present only if <code>X</code> is supplied, is an appropriate argument, and if <code>keep.X</code> is TRUE.
parity	Character string; "univar" if the data were univariate, "bivar" if they were bivariate.
numeric	Logical scalar; TRUE if the (original) data were numeric, FALSE otherwise.
AIC	The value of $AIC = -2 \cdot \log.likelihood + 2 \cdot npar$ for the fitted model.
BIC	The value of $BIC = -2 \cdot \log.likelihood + \log(nobs) \cdot npar$ for the fitted model. In the forgoing <code>nobs</code> is the number of observations. This is the number of <i>non-missing</i> values in <code>unlist(y)</code> in the univariate setting and one half of this number in the bivariate setting.
args	<p>A list of argument values supplied. This component is returned in the interest of making results reproducible. It is also needed to facilitate the updating of a model via the update method for the class <code>hmm.discnp</code>, <code>update.hmm.discnp()</code>.</p> <p>It has components:</p> <ul style="list-style-type: none"> • <code>newstyle</code> • <code>method</code> • <code>optimiser</code> • <code>optimMethod</code>

- stationary
- mixture
- cis
- tolerance
- itmax
- crit
- addIntercept

Warnings

The ordering of the (hidden) states can be arbitrary. What the estimation procedure decides to call “state 1” may not be what *you* think of as being state number 1. The ordering of the states will be affected by the starting values used.

Some experiences with using the “ABSGRD” stopping criterion indicate that it may be problematic in the context of discrete non-parametric distributions. For example a value of 1854.955 was returned after 200 LM steps in one (non-convergent, of course!) attempt at fitting a model. The stopping criterion “PCLL” in this example was took the “reasonable” value of 0.03193748 when iterations ceased.

Notes

The package *used* to require the argument *y* to be a *matrix* in the case of multiple observed sequences. If the series were of unequal length the user was expected to pad them out with NAs to equalize the lengths.

The old matrix format for multiple observation sequences was permitted for a while (and the matrix was internally changed into a list) but this is no longer allowed. If *y* is indeed given as a matrix then this corresponds to a single observation sequence and it must have one (univariate setting) or two (bivariate setting) columns which constitute the observations of the respective variates.

If $K=1$ then *tpm*, *ispd*, *converged*, and *nstep* are all set equal to NA in the list returned by this function.

The estimate of *ispd* in the non-stationary setting is inevitably very poor, unless the number of sequences of observations (the length of the list *y*) is very large. We have in effect “less than one” relevant observation for each such sequence.

The returned values of *tpm* and *Rho* (or the entries of *Rho* when *Rho* is a list) have dimension names. These are the same as the dimension names of the starting values of these objects (as provided in *par0*) if these exist. Otherwise they are formed from the sorted unique values of the observations in *y* or are taken to be the appropriate sequences of integers. Likewise the returned value of *ispd* is a named vector, the names being the same as the row (and column) names of *tpm* and the corresponding dimension names of *Rho* or of its entries.

If *method* is equal to “EM” it *may* get switched to “bf” at some EM step if there is a decrease in the log likelihood. This is “theoretically impossible” but can occur in practice due to an intricacy in the way that the EM algorithm treats *ispd* when *stationary* is TRUE. It turns out to be effectively impossible to maximise the expected log likelihood unless the term in that quantity corresponding to *ispd* is ignored (whence it *is* ignored). Ignoring this term is “asymptotically negligible” but can have the unfortunate effect of occasionally leading to a decrease in the log likelihood.

If such a decrease is detected, then the algorithm switches over to using the "bf" method of maximisation (which is not beset by this difficulty). The current estimates of ispd, tpm and Rho are used as starting values for the "bf" method.

It seems to me that it *should* be the case that such switching can occur only if stationary is TRUE. However I have encountered instances in which the switch occurs when stationary is FALSE. I have yet to figure out/track down what is going on here.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

References

Rabiner, L. R., "A tutorial on hidden Markov models and selected applications in speech recognition," Proc. IEEE vol. 77, pp. 257 – 286, 1989.

Zucchini, W. and Guttorp, P., "A hidden Markov model for space-time precipitation," Water Resources Research vol. 27, pp. 1917-1923, 1991.

MacDonald, I. L., and Zucchini, W., "Hidden Markov and Other Models for Discrete-valued Time Series", Chapman & Hall, London, 1997.

Liu, Limin, "Hidden Markov Models for Precipitation in a Region of Atlantic Canada", Master's Report, University of New Brunswick, 1997.

See Also

[rhmm\(\)](#), [mps\(\)](#), [viterbi\(\)](#)

Examples

```
# TO DO: Create one or more bivariate examples.
#
# The value of itmax in the following examples is so much
# too small as to be risible. This is just to speed up the
# R CMD check process.
# 1.
Yval <- LETTERS[1:10]
Tpm <- matrix(c(0.75,0.25,0.25,0.75),ncol=2,byrow=TRUE)
Rho <- cbind(c(rep(1,5),rep(0,5)),c(rep(0,5),rep(1,5)))/5
rownames(Rho) <- Yval
set.seed(42)
xxx <- rhmm(ylengths=rep(1000,5),nsim=1,tpm=Tpm,Rho=Rho,yval=Yval,drop=TRUE)
fit1 <- hmm(xxx,par0=list(tpm=Tpm,Rho=Rho),newstyle=FALSE,itmax=10)
print(fit1$Rho)
newRho <- cnvrtRho(Rho)
fit2 <- hmm(xxx,par0=list(tpm=Tpm,Rho=newRho),itmax=10)
print(round(cnvrtRho(fit2$Rho),6))
# Agrees with fit1$Rho to 5 decimal places; the difference
# is due to the fitting procedure used with the "new style"
# Rho, which calls upon multinom() from the nnet package.
```

```

# 2.
# See the help for rhmm() for how to generate y.num.
## Not run:
  fit.num      <- hmm(y.num,K=2,verb=TRUE,itmax=10)
  fit.num.mix  <- hmm(y.num,K=2,verb=TRUE,mixture=TRUE,itmax=10)
  print(fit.num[c("tpm","Rho")])

## End(Not run)
# Note that states 1 and 2 get swapped.

# 3.
xxx <- with(SydColDisc,split(y,f=list(locn,depth)))
fitSydCol <- hmm(xxx,K=2,itmax=10) # Two states: above and below the thermocline.

# 4.
fitLesCount <- hmm(lesionCount,K=2,itmax=10) # Two states: relapse and remission.

```

hydroDat

Canadian hydrological data sets.

Description

Five data sets obtained from the “HYDAT” database, Environment and Climate Change Canada’s database of historical hydrometric data. The data were obtained using the `tidyhydroDat` package. The data have been trimmed so that there are no gaps in the observation dates and are presented in “raw” form and in discretised form as deciles of the residuals (difference between raw values and the daily mean over years).

Usage

```

data("linLandFlows")
data("ftLiardFlows")
data("portMannFlows")
data("portMannSedLoads")
data("portMannSedCon")

```

Format

Data frames with observations on the following 3 variables.

`Date` Dates on which observations were made.

`Value` Numeric vector of observation values.

`mean` The mean over years of `Value`.

`resid` The difference `Value - mean`.

`deciles` A factor with levels `d1, ..., d10`, which are the deciles of the variable `resid`

Details

The variable mean was calculated as follows:

```
yday <- as.POSIXlt(X$Date)$yday
mn <- tapply(X$Value,yday,mean,na.rm=TRUE)
mean <- mn[as.character(yday)]
```

where X is the data set being processed.

The data set linLandFlows originally consisted of 2008 observations; there were 1980 observations after “trimming”. The data set ftLiardFlows originally consisted of 22364 observations; there were 11932 observations after “trimming”. The data set portMannFlows originally consisted of 6455 observations; there were 3653 observations after “trimming”. The data set portMannSedLoads consists of 2771 observations; no observations were trimmed. The data set portMannSedCon consists of 4597 observations; no observations were trimmed.

The units of the “Flows” variables are cubic metres per second (m^3/s); the units of “portMannSedLoads” are tonnes; the units of “portMannSedCon” are milligrams per litre (mg/l).

The “linLandFlows” data were obtained at the Lindberg Landing hydrometric station on the Liard River in the Northwest Territories of Canada. The “ftLiardFlows” data were obtained at the Fort Liard hydrometric station on the Liard River in the Northwest Territories of Canada. The “portMann” data were obtained at the hydrometric station located at the Port Mann pumping station on the Fraser River in the Province of British Columbia in Canada.

Source

Environment and Climate Change Canada’s database “HYDAT”, a database of historical hydrometric data. The data were obtained via the tidyhydat package, which is available from “CRAN”, <https://cran.r-project.org>

Examples

```
fit <- hmm(linLandFlows$deciles,K=4,itmax=10)
```

lesionCount

Multiple sclerosis lesion counts for three patients.

Description

Lesion counts for three multiple sclerosis patients. The counts were obtained by magnetic resonance imaging, and were observed at monthly intervals.

Usage

```
lesionCount
```

Format

A list with three components each component being the sequence of counts for a given patient and consisting of a vector with non-negative integer entries.

Modelling

The hidden Markov models applied to these data by Albert et al. and by MacKay and Petkau were much more complex and elaborate than those fitted in the examples in this package. See the references for details.

Source

The data were originally studied by Albert et al., (1994). They were also analyzed by Altman and Petkau (2005). The data were kindly provided by Prof. Altman.

References

Albert, P. S., McFarland, H. F., Smith, M. E., and Frank, J. A. Time series for modelling counts from a relapsing-remitting disease: application to modelling disease activity in multiple sclerosis. *Statistics in Medicine* **13** (1994) 453–466.

Altman, Rachel MacKay, and Petkau, A. John. Application of hidden Markov models to multiple sclerosis lesion count data. *Statistics in Medicine* **24** (2005) 2335–2344.

 logLikHmm

Log likelihood of a hidden Markov model

Description

Calculate the log likelihood of a hidden Markov model with discrete non-parametric observation distributions.

Usage

```
logLikHmm(y, model=NULL, tpm=NULL, ispd=NULL, Rho=NULL, X=NULL,
          addIntercept=NULL, warn=TRUE)
```

Arguments

<code>y</code>	A vector, or list of vectors, or a two column matrix or a list of such matrices, whose entries consist of observations from a hidden Markov model with discrete non-parametric observation distributions.
<code>model</code>	An object specifying a hidden Markov model, usually returned by <code>hmm()</code> .
<code>tpm</code>	The transition probability matrix of the Markov chain. Ignored (and extracted from <code>model</code>) if <code>model</code> is non-NULL.

ispd	The vector of probabilities specifying the initial state probability distribution, or a matrix each of whose columns is a trivial (“delta function”) vector specifying the “most probable” initial state for each observation sequence. If ispd is missing then ispd is calculated as the stationary distribution determined by tpm. Ignored (and extracted from model) if model is non-NULL.
Rho	An object specifying the “emission” probabilities of the observations. (See the Details in the help for <code>hmm()</code> .) Ignored (and extracted from model) if model is non-NULL.
X	An optional <i>numeric</i> matrix, or a list of such matrices, of <i>predictors</i> . The use of such predictors is (currently, at least) applicable only in the univariate emissions setting, and then only if Rho is supplied in the “newstyle” format. If X is a list it must be of the same length as y and all entries of this list must have the same number of columns. The number of rows of each entry must be equal to the length of the corresponding entry of y. If X is a matrix then y should be a vector or one-column matrix (or a list with a single entry equal to such).
addIntercept	Logical scalar. See the documentation of <code>hmm()</code> . If this argument is not specified, and if model is NULL then an error is thrown.
warn	Logical scalar; should a warning be issued if Rho hasn’t got relevant dimension names? (Note that if this is so, then the corresponding dimension names are formed from the sorted unique values of y or of the appropriate column(s) of y. And if <i>this</i> is so, then the user should be sure that the ordering of the entries of Rho corresponds properly to the the sorted unique values of y.) This argument is passed to the utility function <code>check.yval()</code> which actually issues the warning if <code>warn=TRUE</code> .

Details

If y is not provided the function simply returns the `log.like` component of model (which could be NULL if model was not produced by `hmm()`).

The observation values (the entries of the vector or matrix y, or of the vectors or matrices which constitute the entries of y if y is a list) must be consistent with the appropriate dimension names of Rho or of its entries when Rho is a list. More specifically, if Rho has dimension names (or its entries have dimension names) then the observation values must all be found as entries of the appropriate dimension name vector. If a vector of dimension names is NULL then the corresponding dimension must be equal to the number of unique observations of the appropriate variate. integers between 1 and `nrow(Rho)`.

Value

The loglikelihood of y given the parameter values specified in `par`.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

References

See `hmm()` for references.

See Also

`hmm()`, `pr()`, `sp()`

Examples

```
# TO DO: One or more bivariate examples.
P <- matrix(c(0.7,0.3,0.1,0.9),2,2,byrow=TRUE)
R <- matrix(c(0.5,0,0.1,0.1,0.3,
             0.1,0.1,0,0.3,0.5),5,2)

set.seed(42)
l11 <- sample(250:350,20,TRUE)
y.num <- rhmm(ylengths=l11,nsim=1,tpm=P,Rho=R,drop=TRUE)
y.let <- rhmm(ylengths=l11,nsim=1,tpm=P,Rho=R,yval=letters[1:5],drop=TRUE)
l11 <- logLikHmm(y.num,tpm=P,Rho=R,warn=FALSE)
row.names(R) <- letters[1:5]
l12 <- logLikHmm(y.let,tpm=P,Rho=R)
# Note: l12 differs from l11 since y.num and y.let were
# simulated from different random seeds.
l13 <- logLikHmm(y.let,tpm=P,Rho=R,ispd=c(0.5,0.5))
fit <- hmm(y.num,K=2,itmax=10)
l14 <- logLikHmm(y.num,fit) # Use the fitted rather than the "true" parameters.
```

misstify

Insert missing values.

Description

Insert missing values into data simulated by `rhmm`.

Usage

```
misstify(X, nafrac, fep = NULL)
```

Arguments

X	A data set (list of matrices with one or two columns, or list of such list) generated by <code>rhmm()</code> .
nafrac	A scalar probability or a vector of two such probabilities if the data are bivariate. The values should in all cases be numbers between zero and 1. These values represent the fraction of missing values in the data (or in the respective components of the data if the model is bivariate). If the data are bivariate <code>nafrac</code> is duplicated if it is of length 1. If an entry of <code>nafrac</code> is equal to 1, an error is thrown. (Having <i>everything</i> missing makes no sense!) The fraction(s) of missing values in a given data set may be determined by <code>nafracCalc()</code> .

`fep` “First entry present”. A list with one or two entries, the first being a logical scalar (which might be named “present”). If there is a second entry it should be a scalar probability (which might be named “p2”). In an application of interest, observation sequences always begin at an observed event, i.e. at a time point at which the “emission” has at least one non-missing value. If `fep[[1]]` is TRUE the NAs will be inserted in such a way that the resulting data have this characteristic. If `fep` is left NULL then its first (possibly only) entry is set to TRUE.

For *bivariate* data, `fep[[2]]` specifies the probability that *both* values of the initial pair of observations are non-missing. In this case one of the entries of the initial pair is chosen to be “potentially” missing, with probabilities `nafrac/sum(nafrac)`. This entry is left non-missing with probability `fep[[2]]`. (The other entry is always left non-missing.)

If the data are univariate or if `fep[[1]]` is FALSE, then `fep[[2]]` is ignored. If the data are bivariate and `fep[[2]]` is not specified, it defaults to the (estimated) conditional probability that both entries of the initial pair of observations are present given that at least one is present, under the assumption of independence of these events. I.e. it is set equal to `prod(1-nafrac)/(1-prod(1-nafrac))`.

Value

A list, like that returned by `rhmm()`. The latter function calls upon `misstify()` to effect the replacement of observations by the appropriate (or specified) fraction of missing values.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

`rhmm()` `nafracCalc()`

Examples

```
P <- matrix(c(0.7,0.3,0.1,0.9),2,2,byrow=TRUE)
R <- matrix(c(0.5,0,0.1,0.1,0.3,
             0.1,0.1,0,0.3,0.5),5,2)
set.seed(42)
l11 <- sample(250:350,20,TRUE)
y <- rhmm(ylengths=l11,nsim=1,tpm=P,Rho=R,drop=TRUE)
y <- misstify(y,nafrac=0.5,fep=list(present=TRUE))
```

mps *Most probable states.*

Description

Calculates the most probable hidden state underlying each observation.

Usage

```
mps(y, model = NULL, tpm, Rho, ispd=NULL, warn=TRUE)
```

Arguments

y	The observations for which the underlying most probable hidden states are required. May be a sequence of observations, or a list each component of which constitutes a (replicate) sequence of observations. If y is missing it is set equal to the y component of model, given that that object and that component exist. Otherwise an error is given.
model	An object describing a fitted hidden Markov model, as returned by <code>hmm()</code> . In order to make any kind of sense, model should bear some reasonable relationship to y.
tpm	The transition probability matrix for a hidden Markov model; ignored if model is non-null. Should bear some reasonable relationship to y.
Rho	An object specifying the probability distributions of the observations (“emission” probabilities) for a hidden Markov model. See <code>hmm()</code> . Ignored if model is non-null. Should bear some reasonable relationship to y.
ispd	A vector specifying the initial state probability distribution for a hidden Markov model, or a matrix each of whose columns are trivial (“delta function”) vectors specifying the “most probable” initial state for each observation sequence. This argument is ignored if model is non-null. It should bear some reasonable relationship to y. If both ispd and model are NULL then ispd is taken to be the stationary distribution of the chain, calculated from tpm.
warn	Logical scalar; should a warning be issued if Rho hasn’t got relevant dimension names? (Note that if this is so, then the corresponding dimension names are formed from the sorted unique values of y or of the appropriate column(s) of y. And if <i>this</i> is so, then the user should be sure that the ordering of the entries of Rho corresponds properly to the the sorted unique values of y.) This argument is passed to the utility function <code>check.yval()</code> which actually issues the warning if warn=TRUE.

Details

For each t the maximum value of $\gamma_t(i)$, i.e. of the (estimated) probability that the state at time t is equal to i , is calculated, and the value of the state with the corresponding index is returned.

Value

If y is a single observation sequence, then the value is a vector of corresponding most probable states.

If y is a list of replicate sequences, then the value is a list, the j -th entry of which constitutes the vector of most probable states underlying the j -th replicate sequence.

Warning

The *sequence of most probable states* as calculated by this function will not in general be the *most probable sequence of states*. It may not even be a *possible* sequence of states. This function looks at the state probabilities separately for each time t , and not at the states in their sequential context.

To obtain the most probable sequence of states use `viterbi()`.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

References

Rabiner, L. R., "A tutorial on hidden Markov models and selected applications in speech recognition," Proc. IEEE vol. 77, pp. 257 – 286, 1989.

See Also

`hmm()`, `rhmm()`, `viterbi()`

Examples

```
# See the help for rhmm() for how to generate y.num.
## Not run:
fit.num <- hmm(y.num,K=2,verb=TRUE)
s.1 <- mps(y.num,fit.num)
s.2 <- mps(y.num,tpm=P,ispd=c(0.25,0.75),Rho=R) # P and R as in the help
                                                # for rhmm().
# The order of the states has gotten swapped; (3 - s.1[,1]) is much
# more similar to s.2[,1] than is s.1[,1].

## End(Not run)
```

nafracCalc

Calculate fractions of missing values.

Description

Calculate the fraction (univariate case) or fractions (bivariate case) of missing values in the data or in each component of the data.

Usage

```
nafracCalc(y)
```

Arguments

`y` A vector or a one or two column matrix of discrete data or a list of such vectors or matrices.

Value

A scalar (between 0 and 1) if the data are univariate or a pair of such scalars if the data are bivariate. The values are equal to the ratios of the total count of missing values in the appropriate column to the total number of observations.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

```
rhmm() misstify()
```

Examples

```
xxx <- with(SydColDisc, split(y, f=list(locln, depth)))  
nafracCalc(xxx) # 0.7185199
```

pr

Probability of state sequences.

Description

Calculates the conditional probability of one or more state sequences, given the corresponding observations sequences (and the model parameters).

Usage

```
pr(s, y, model=NULL, tpm, Rho, ispd=NULL, warn=TRUE)
```


Arguments

<code>s</code>	A sequence of states of the underlying Markov chain, or a list of such sequences.
<code>y</code>	A one or two column matrix of observations from a hidden Markov model, corresponding to the state sequence <code>s</code> , or a list of such matrices corresponding to the state sequences in the list <code>s</code> . If <code>y</code> consists of a single matrix, it is assumed to correspond to each of the state sequences in <code>s</code> in turn. Otherwise the length of the list <code>y</code> must be the same as the length of the list <code>s</code> . (If not, then an error is given). If <code>y</code> is missing, it is extracted from <code>model</code> provided that <code>model</code> and its <code>y</code> component are not NULL. Otherwise an error is given.
<code>model</code>	An object of class <code>hmm.discnp</code> as returned by <code>hmm()</code> .
<code>tpm</code>	The transition probability matrix of the chain. Ignored (and extracted from <code>model</code> instead) if <code>model</code> is not NULL.
<code>Rho</code>	An object specifying the “emission” probabilities of observations, given the underlying state. See <code>hmm()</code> . Ignored (and extracted from <code>model</code> instead) if <code>model</code> is not NULL.
<code>ispd</code>	The vector specifying the initial state probability distribution of the Markov chain. Ignored (and extracted from <code>model</code> instead) if <code>model</code> is not NULL. If both <code>ispd</code> and <code>model</code> are NULL then <code>ispd</code> is taken to be the stationary distribution of the chain, calculated from <code>tpm</code> .
<code>warn</code>	Logical scalar; should a warning be issued if <code>Rho</code> hasn’t got relevant dimension names? (Note that if this is so, then the corresponding dimension names are formed from the sorted unique values of <code>y</code> or of the appropriate column(s) of <code>y</code> . And if <i>this</i> is so, then the user should be sure that the ordering of the entries of <code>Rho</code> corresponds properly to the the sorted unique values of <code>y</code> .) This argument is passed to the utility function <code>check.yval()</code> which actually issues the warning if <code>warn=TRUE</code> .

Value

The probability of `s` given `y`, or a vector of such probabilities if `s` and `y` are lists.

Warning

The conditional probabilities will be tiny if the sequences involved are of any substantial length. Underflow may be a problem. The implementation of the calculations is not sophisticated.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

[hmm\(\)](#), [mps\(\)](#), [viterbi\(\)](#), [sp\(\)](#), [fitted.hmm.discnp\(\)](#)

Examples

```
# See the help for rhmm() for how to generate y.num.
## Not run:
fit.num <- hmm(y.num,K=2,verb=TRUE)
# Using fitted parameters.
s.vit.1 <- viterbi(y.num,fit.num)
pr.vit.1 <- pr(s.vit.1,model=fit.num)
# Using true parameters from which y.num was generated.
s.vit.2 <- viterbi(y.num,tpm=P,Rho=R)
pr.vit.2 <- pr(s.vit.2,y.num,tpm=P,Rho=R)

## End(Not run)
```

predict.hmm.discnp *Predicted values of a discrete non-parametric hidden Markov model.*

Description

Calculates the predicted values of a discrete non-parametric hidden Markov model as the conditional expectations of the observations, given the entire observation sequence (and the [estimated] parameters of the model).

Usage

```
## S3 method for class 'hmm.discnp'
predict(object, y=NULL, ...)
```

Arguments

object	An object of class <code>hmm.discnp</code> as returned by <code>hmm()</code> .
y	A data structure from which the fitted model object <i>could</i> have been fitted.
...	Not used.

Details

This function is essentially the same as `fitted.hmm.discnp()`, except that it allows the calculation of fitted/predicted values for a data object possibly different from that to which the model was fitted. Like `fitted.hmm.discnp()` this function is essentially a wrapper for `sp()`. Note that the predicted values are calculated as “conditional expected values” which makes sense *only* if the observations are *numeric*. If this is not the case, an error is thrown. If the observation list `y` is not supplied then it is taken to be the `y` component of `object`. If there is no such component (i.e. if `object` is the value returned from `hmm()` with the argument `keep.y` set equal to `FALSE`) then an error is thrown.

Value

If the observations (stored in `y`) consist of a single sequence, then the object returned is a matrix. If the model is univariate (see `hmm()`) then matrix has a single column constituting the corresponding sequence of fitted values. If the model is bivariate (either independent or dependent) then the matrix has two columns corresponding respectively to the two variables in the bivariate model. If the observations consist of a list of observation sequences, then the object returned is a list of corresponding matrices whose columns are the fitted values.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

`sp()` `link{fitted.hmm.discnp}()`

Examples

```
P <- matrix(c(0.7,0.3,0.1,0.9),2,2,byrow=TRUE)
R <- matrix(c(0.5,0,0.1,0.1,0.3,
             0.1,0.1,0,0.3,0.5),5,2)
set.seed(42)
l11 <- sample(250:350,20,TRUE)
y1 <- rhmm(ylengths=l11,nsim=1,tpm=P,Rho=R,drop=TRUE)
fit <- hmm(y1,K=2,verb=TRUE,keep.y=TRUE,itmax=10)
fv <- fitted(fit)
set.seed(176)
l12 <- sample(250:350,20,TRUE)
y2 <- rhmm(ylengths=l12,nsim=1,tpm=P,Rho=R,drop=TRUE)
pv <- predict(fit,y=y2)
```

rhmm

Simulate discrete data from a non-parametric hidden Markov model.

Description

Simulates one or more replicates of discrete data from a model such as is fitted by the function `hmm()`.

Usage

```
rhmm(model,...,nsim,verbose=FALSE)
## Default S3 method:
rhmm(model,...,nsim=1,verbose=FALSE,ylengths,
      nafrac=NULL,fep=NULL,tpm,Rho,ispd=NULL,yval=NULL,
      drop=TRUE)
## S3 method for class 'hmm.discnp'
```

```
rhmm(model, ..., nsim=1, verbose=FALSE, inMiss=TRUE,
      fep=NULL, drop=TRUE)
```

Arguments

model	An object of class <code>hmm.discnp</code> . This will have the form of a list specifying a hidden Markov model with discrete emissions and emission probabilities specified non-parametrically, i.e. by means of a table or tables. Usually this will be an object returned by <code>hmm()</code> . This argument is ignored by the default method.
...	Not used.
nsim	Integer scalar; the number of data sets to be simulated.
verbose	Logical scalar. If TRUE then the overall index of the simulated value that has been reached is printed out every 1000 iterations. Useful for reassurance when very “large” simulations are undertaken.
ylengths	Integer values vector specify the lengths (or number of rows in the bivariate setting) of the individual observation sequences constituting a data set.
nafrac	See <code>misstify()</code> for an explanation of this argument. If specified a fraction <code>nafrac[[j]]</code> of column <code>j</code> of the data will be randomly set equal to NA.
fep	“First entry present”. See <code>misstify()</code> for an explanation of this argument.
tpm	The transition probability matrix for the underlying hidden Markov chain(s). Note that the rows of <code>tpm</code> must sum to 1. Ignored if <code>ncol(Rho)==1</code> . Ignored by the <code>hmm.discnp</code> method and extracted from <code>model</code> .
Rho	An object specifying the probability distribution of the observations, given the state of the underlying hidden Markov chain. (I.e. the “emission” probabilities.) See <code>hmm()</code> . Note that <code>Rho</code> can be such that the number of states is 1, in which case the simulated data are i.i.d. from the single distribution specified by <code>Rho</code> . Ignored by the <code>hmm.discnp</code> method and extracted from <code>model</code> .
ispd	A vector specifying the initial state probability distribution of the chain. If this is not specified it is taken to be the stationary distribution of the chain, calculated from <code>tpm</code> . Ignored by the <code>hmm.discnp</code> method and extracted from <code>model</code> .
yval	Vector of possible values of the observations, or (in the bivariate setting) a list of two such vectors. If not supplied it is formed from the appropriate dimension names of <code>Rho</code> or of the entries of <code>Rho</code> if <code>Rho</code> is a list. Ignored by the <code>hmm.discnp</code> method.
drop	Logical scalar; if <code>nsim</code> is 1 and if <code>drop</code> is TRUE then the list to be returned by this function is replaced by its first and only entry. (See below.)
inMiss	Logical scalar; if TRUE then missing values will be randomly inserted into the data in the fraction <code>nafrac</code> determined from object.

Value

If `nsim>1` or `drop` is FALSE then the value returned is a list of length `nsim`. Each entry of this list is in turn a list of the same length as `ylengths`, each component of which is an independent vector or matrix of simulated observations. The length or number of rows of component `i` of this list is equal to `ylengths[i]`. The values of the observations are entries of `yval` or of `its` entries when `yval` is a list.

If `nsim=1` and `drop` is `TRUE` then the (“outer”) list described above is replaced by its first and only entry

Note

You may find it useful to avail yourself of the function `nafracCalc()` to determine the fraction of missing values in a given existing (presumably “real”) data set.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

`hmm()` `nafracCalc()` `misstify()`

Examples

```
# To do: one or more bivariate examples.
## Not run:
  y <- list(linLandFlows$deciles,ftLiardFlows$deciles)
  fit <- hmm(y,K=3)
  simX <- rhmm(fit)

## End(Not run)
```

scovmat

Simulation based covariance matrix.

Description

Produces an estimate of the covariance matrix of the parameter estimates in a model fitted by `hmm.discnp`. Uses a method based on simulation (or “parametric bootstrapping”).

Usage

```
scovmat(object, seed = NULL, nsim=100, verbose = TRUE)
```

Arguments

<code>object</code>	An object of class <code>hmm.discnp</code> as returned by <code>hmm()</code> .
<code>seed</code>	Integer scalar serving as a seed for the random number generator. If left <code>NULL</code> the seed itself is chosen randomly from the set of integers between 1 and 10^5 .
<code>nsim</code>	A positive integer. The number of simulations upon which the covariance matrix estimate will be based.
<code>verbose</code>	Logical scalar; if <code>TRUE</code> , iteration counts will be printed out during each of the simulation and model-fitting stages.

Details

This function is currently applicable only to models fitted to univariate data. The covariance matrix produced is for the “raw” parameters (entries of `tpm` with the last column dropped — since the rows sum to 1, and the entries of `Rho` with the last row dropped — since the columns sum to 1).

Value

A (positive definite) matrix which is an estimate of the covariance of the parameter estimates from the fitted model specified by `object`. It has row and column labels which indicate the parameters to which its entries pertain, in a reasonably perspicuous manner.

This matrix has an attribute `seed` (the random number generation seed that was used) so that the calculations can be reproduced.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

[squandCI\(\)](#) `link{rhmm}()` `link{hmm}()`

Examples

```
## Not run:
y <- list(lindLandFlows$deciles, ftLiardFlows$deciles)
fit <- hmm(y, K=3)
ccc <- scovmat(fit, nsim=100)

## End(Not run)
```

sp

Calculate the conditional state probabilities.

Description

Returns the probabilities that the underlying hidden state is equal to each of the possible state values, at each time point, given the observation sequence. Also can return the fitted conditional means, if requested, given that the observations are numeric.

Usage

```
sp(y, model = NULL, tpm=NULL, Rho=NULL, ispd=NULL, X=NULL,
  addIntercept=NULL, means = FALSE, warn=TRUE)
```

Arguments

y	The observations on the basis of which the probabilities of the underlying hidden states are to be calculated. May be a (one or two column) matrix of observations, or a list each component of which is such a matrix. If y is missing it is set equal to the y component of model, given that that argument is non-NULL and that that component exist. Otherwise an error is given.
model	An object of class <code>hmm.discnp</code> as returned by <code>hmm()</code> .
tpm	The transition probability matrix for the underlying hidden Markov chain. Ignored if model is not NULL. Ignored if model is not NULL (in which case tpm is extracted from model).
Rho	An object specifying the distribution of the observations, given the underlying state. I.e. the “emission” probabilities. See <code>hmm()</code> . Ignored if model is not NULL (in which case Rho is extracted from model).
ispd	Vector specifying the initial state probability distribution of the underlying hidden Markov chain. Ignored if model is not NULL (in which case ispd is extracted from model). If both model and ispd are NULL then ispd is calculated to be the stationary distribution of the chain as determined by tpm.
X	An optional <i>numeric</i> matrix, or a list of such matrices, of <i>predictors</i> . The use of such predictors is (currently, at least) applicable only in the univariate emissions setting, and then only if Rho is supplied in the “newstyle” format. If X is a list it must be of the same length as y and all entries of this list must have the same number of columns. The number of rows of each entry must be equal to the length of the corresponding entry of y. If X is a matrix then y should be a vector or one-column matrix (or a list with a single entry equal to such).
addIntercept	Logical scalar. See the documentation of <code>hmm()</code> .
means	A logical scalar; if means is TRUE then the conditional expected value of (each column of) the observations (given the observation sequence) is calculated at each time point. If means is TRUE and the observation values are not numeric, then an error is given.
warn	Logical scalar; should a warning be issued if Rho hasn’t got relevant dimension names? (Note that if this is so, then the corresponding dimension names are formed from the sorted unique values of y or of the appropriate column(s) of y. And if <i>this</i> is so, then the user should be sure that the ordering of the entries of Rho corresponds properly to the the sorted unique values of y.) This argument is passed to the utility function <code>check.yval()</code> which actually issues the warning if <code>warn=TRUE</code> .

Details

The conditional mean value at time t is calculated as

$$\sum_k \gamma_t(k) \mu_k$$

where $\gamma_t(k)$ is the conditional probability (given the observations) that the hidden Markov chain is in state k at time t , and μ_k is the expected value of an observation given that the chain is in state k .

Value

If means is TRUE then the returned value is a list with components

probs The conditional probabilities of the states at each time point.
 means The conditional expectations of the observations at each time point.

Clearly this makes sense only if the observations are numeric.

Otherwise the returned value consists of probs as described above.

If there is a single matrix of observations y then probs is a matrix whose rows correspond to the states of the hidden Markov chain, and whose columns correspond to the observation times. If the observations consist of a list of observation vectors, then probs is a list of such matrices, one for each matrix of observations.

Likewise for the means component of the list returned when the argument means is TRUE.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

[hmm\(\)](#), [mps\(\)](#), [viterbi\(\)](#), [pr\(\)](#), [fitted.hmm.discnp\(\)](#)

Examples

```
P <- matrix(c(0.7,0.3,0.1,0.9),2,2,byrow=TRUE)
R <- matrix(c(0.5,0,0.1,0.1,0.3,
             0.1,0.1,0,0.3,0.5),5,2)
set.seed(42)
y.num <- rhmm(ylengths=rep(300,20),nsim=1,tpm=P,Rho=R,drop=TRUE)
fit.num <- hmm(y.num,K=2,verb=TRUE,keep.y=TRUE,itmax=10)
cpe1 <- sp(model=fit.num,means=TRUE) # Using the estimated parameters.
cpe2 <- sp(y.num,tpm=P,Rho=R,means=TRUE,
           warn=FALSE) # Using the ``true'' parameters.
# The foregoing would issue a warning that Rho had no row names
# were it not for the fact that "warn" has been set to FALSE.
```

squantCI

Simulation-quantile based confidence intervals.

Description

Calculates estimates of confidence intervals for the parameters of a model fitted by `hmm.discnp`. Uses a method based quantiles of estimates produced by simulation (or “parametric bootstrapping”).

Usage

```
squantCI(object, seed = NULL, alpha = 0.05, nsim=100, verbose = TRUE)
```


Arguments

object	An object of class <code>hmm.discnp</code> as returned by <code>hmm()</code> .
seed	Integer scalar serving as a seed for the random number generator. If left <code>NULL</code> the seed itself is chosen randomly from the set of integers between 1 and 10^5 .
alpha	Positive real number strictly between 0 and 1. A set of $100*(1-\alpha)\%$ confidence intervals will be produced.
nsim	A positive integer. The number of simulations upon which the confidence interval estimates will be based.
verbose	Logical scalar; if <code>TRUE</code> , iteration counts will be printed out during each of the simulation and model-fitting stages.

Details

This function is currently applicable only to models fitted to univariate data. The confidence intervals calculated are for the “raw” parameters (entries of `tpm` with the last column dropped — since the rows sum to 1, and the entries of `Rho` with the last row dropped — since the columns sum to 1.

Value

A 2-by-`npar` matrix (where `npar` is the number of “independent” parameters in the model) whose columns form the estimated confidence intervals. The column labels indicate the parameters to which each column pertains, in a reasonably perspicuous manner. The row labels indicate the relevant quantiles in percentages.

This matrix has an attribute `seed` (the random number generation seed that was used) so that the calculations can be reproduced.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

`scovmat()` `link{rhmm}()` `link{hmm}()`

Examples

```
## Not run:
y <- list(lindLandFlows$deciles, ftLiardFlows$deciles)
fit <- hmm(y, K=3)
CIs <- squantCI(fit, nsim=100)

## End(Not run)
```

 SydColDisc

Discretised version of coliform counts in sea-water samples

Description

Discretised version of counts of faecal coliform bacteria in sea water samples collected at seven locations near Sydney NSW, Australia. There were four “controls”: Longreef, Bondi East, Port Hacking “50”, and Port Hacking “100” and three “outfalls”: Bondi Offshore, Malabar Offshore and North Head Offshore. At each location measurements were made at four depths: 0, 20, 40, and 60 meters. A large fraction of the counts are missing values.

Usage

SydColDisc

Format

A data frame with 5432 observations on the following 6 variables.

`y` A factor consisting of a discretisation of counts of faecal coliform count bacteria in sea water samples. The original measures were obtained by a repeated dilution process. The data were discretised using the `cut()` function with breaks given by `c(0, 1, 5, 25, 200, Inf)` and labels equal to `c("lo", "mlo", "m", "mhi", "hi")`.

`locn` a factor with levels Longreef, Bondi East, Port Hacking 50, Port Hacking 100, Bondi Offshore, Malabar Offshore and North Head Offshore.

`depth` a factor with levels 0 (0 metres), 20 (20 metres), 40 (40 metres), 60 (60 metres).

`ma.com` A factor with levels no and yes, indicating whether the Malabar sewage outfall had been commissioned.

`nh.com` A factor with levels no and yes, indicating whether the North Head sewage outfall had been commissioned.

`bo.com` A factor with levels no and yes, indicating whether the Bondi Offshore sewage outfall had been commissioned.

Details

The observations corresponding to each location-depth combination constitute a (discrete valued) time series. The sampling interval is ostensibly 1 week; distinct time series are ostensibly synchronous. The measurements were made over a 194 week period. Due to exigencies of weather, the unreliability of boats and other factors the collection times were actually highly irregular and have been rounded to the nearest week. Often no sample was obtained at a given site within a week of the putative collection time, in which the observed count is given as a missing value. In fact over **75%** of the counts are missing. See Turner et al. (1998) for more detail.

Modelling

The hidden Markov models applied in the paper Turner et al. (1998) and in the paper Turner (2008) used a numeric version of the response in this data set. The numeric response was essentially a square root transformation of the original data, and the resulting values were modelled in terms of a Poisson distribution. See the references for details.

Source

The original data were kindly supplied by Geoff Coade, of the New South Wales Environment Protection Authority (Australia)

References

T. Rolf Turner, Murray A. Cameron, and Peter J. Thomson. Hidden Markov chains in generalized linear models. *Canadian J. Statist.* **26** (1998) 107 – 125.

Rolf Turner. Direct maximization of the likelihood of a hidden Markov model. *Comp. Statist. Data Anal.* **52** (2008) 4147–4160.

update.hmm.discnp	<i>Update a fitted hmm.discnp model.</i>
-------------------	--

Description

An update() method for objects of class hmm.discnp.

Usage

```
## S3 method for class 'hmm.discnp'
update(object,..., data, Kplus1=FALSE,
       tpm2=NULL, verbose=FALSE, method=NULL, optimiser=NULL,
       stationary=NULL, mixture=NULL, cis=NULL, tolerance=NULL,
       itmax=NULL, crit=NULL, X=NULL, addIntercept=NULL)
```

Arguments

object	An object of class hmm.discnp as returned by hmm() .
...	Not used.
data	The data set to which the (updated) model is to be fitted. See the description of the y argument of hmm() for more detail.
Kplus1	Logical scalar. Should the number of states be incremented by 1? If so then tpm (the transition probability matrix) is re-formed by <code>rbind()</code> -ing on a row all of whose entries are 1/K (where K is the “old” number of states) and then <code>cbind()</code> -ing on a column of zeroes. The emission probability matrix Rho is reformed by <code>cbind()</code> -ing on a column all of whose entries are 1/m where m is the number of discrete values of the emissions.

Note that the initial likelihood of the “new” model with $K+1$ states will (should?) be exactly the same as that of the “old” fitted K -state model.

The `Kplus1` argument is provided mainly so as to provide a set of starting values for the fitting process which will guarantee the log likelihood of a $K+1$ -state model will be at least as large as that of a K -state model fitted to the same data set.

Experience indicates that when `Kplus1=TRUE` is used, the fitting process does not “move very far” from the maximum log likelihood found for the K -state model. It is then advisable to try (many) random starting values so as to (try to) find the “true” maximum for the $K+1$ -state model.

<code>tpm2</code>	The transition probability matrix to use when updating a model fitted with $K=1$ and <code>Kplus1=TRUE</code> . This argument is ignored otherwise. The default value of this argument is <code>matrix(0.5,2,2)</code> . The value of <code>tpm2</code> makes no difference to the <i>initial</i> value of the likelihood of the $K=2$ model (which will be identical to the likelihood of the fitted $K=1$ model that is being updated). Any two-by-two transition probability matrix “will do”. However the value of <code>tpm2</code> could conceivably have an impact on the final likelihood of the $K=2$ model to which the fitting procedure converges. This is particularly true if the method is (or is switched to) “LM”.
<code>verbose</code>	See the help for <code>hmm()</code> .
<code>method</code>	See the help for <code>hmm()</code> .
<code>optimiser</code>	See the help for <code>hmm()</code> .
<code>stationary</code>	See the help for <code>hmm()</code> .
<code>mixture</code>	See the help for <code>hmm()</code> .
<code>cis</code>	See the help for <code>hmm()</code> .
<code>tolerance</code>	See the help for <code>hmm()</code> .
<code>itmax</code>	See the help for <code>hmm()</code> .
<code>crit</code>	See the help for <code>hmm()</code> .
<code>X</code>	See the help for <code>hmm()</code> .
<code>addIntercept</code>	See the help for <code>hmm()</code> .

Details

Except for argument `X`, any arguments that are left `NULL` have their values supplied from the `args` component of object.

Value

An object of class `hmm.discnp` with an additional component `init.log.like` which is the initial log likelihood calculated at the starting values of the parameters (which may be modified from the parameters returned in the object being updated, if `Kplus1` is `TRUE`). The calculation is done by the function `logLikHmm()`. Barring the strange and unforeseen, `init.log.like` should be (reassuringly) equal to `object$log.like`. See `hmm()` for details of the other components of the returned value.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

See Also

[hmm\(\)](#) [rhmm.hmm.discnp\(\)](#)

Examples

```
set.seed(294)
fit <- hmm(WoodPeweeSong,K=2,rand.start=list(tpm=TRUE,Rho=TRUE),itmax=10)
xxx <- rhmm(fit,nsim=1)
sfit <- update(fit,data=xxx,itmax=10)
yyy <- with(SydColDisc,split(y,f=list(locn,depth)))
f1 <- hmm(yyy,K=1)
f2 <- update(f1,data=yyy,Kplus1=TRUE) # Big improvement, but ...
## Not run:
g2 <- hmm(yyy,K=2) # Substantially better than f2.

## End(Not run)
```

viterbi *Most probable state sequence.*

Description

Calculates “the” most probable state sequence underlying each of one or more replicate observation sequences.

Usage

```
viterbi(y, model = NULL, tpm, Rho, ispd=NULL,log=FALSE, warn=TRUE)
```

Arguments

y	The observations for which the underlying most probable hidden states are required. May be a sequence of observations, or a list each entry of which constitutes an independent sequence of observations. If y is missing (and if model is not NULL) then y is extracted from model, provided that the y component of model is present. Otherwise an error is given.
model	An object describing a hidden Markov model, as fitted to the data set y by hmm() .
tpm	The transition probability matrix for a hidden Markov model; ignored if model is non-null.

Rho	<p>An object specifying the probability distributions of the observations for a hidden Markov model. See <code>hmm()</code>. Ignored if <code>model</code> is non-null. Should bear some reasonable relationship to <code>y</code>.</p> <p>If Rho has dimension names (or if its entries have dimension names in the case where Rho is a list) then the appropriate dimension names must include all corresponding values of the observations. If a relevant vector of dimension names is NULL then it is formed as the sort unique values of the appropriate columns of the observation matrices. In this case the corresponding dimensions must match the number of unique values.</p>
<code>i.spd</code>	<p>The initial state probability distribution for a hidden Markov model; ignored if <code>model</code> is non-null. Should bear some reasonable relationship to <code>y</code>. If <code>model</code> and <code>i.spd</code> are both NULL then <code>i.spd</code> is set equal to the stationary distribution calculated from <code>tpm</code>.</p>
<code>log</code>	<p>Logical scalar. Should logarithms be used in the recursive calculations of the probabilities involved in the Viterbi algorithm, so as to avoid underflow? If <code>log</code> is FALSE then underflow is avoided instead by a normalization procedure. The quantity <code>delta</code> (see Rabiner 1989, page 264) is replaced by <code>delta/sum(delta)</code> at each step. It should actually make no difference whether <code>log</code> is set to TRUE. I just included the option because I could. Also the HMM package uses the logarithm approach so setting <code>log=TRUE</code> might be of interest if comparisons are to be made between the results of the two packages.</p>
<code>warn</code>	<p>Logical scalar; should a warning be issued if Rho hasn't got relevant dimension names? (Note that if this is so, then the corresponding dimension names are formed from the sorted unique values of <code>y</code> or of the appropriate column(s) of <code>y</code>. And if <i>this</i> is so, then the user should be sure that the ordering of the entries of Rho corresponds properly to the the sorted unique values of <code>y</code>.) This argument is passed to the utility function <code>check.yval()</code> which actually issues the warning if <code>warn=TRUE</code>.</p>

Details

Applies the Viterbi algorithm to calculate “the” most probable rotable state sequence underlying each observation sequences.

Value

If `y` consists of a single observation sequence, the value is the underlying most probable observation sequence, or a matrix whose columns consist of such sequences if there is more than one (equally) most probable sequence.

If `y` consists of a list of observation sequences, the value is a list each entry of which is of the form described above.

Warning

There *may* be more than one equally most probable state sequence underlying a given observation sequence. This phenomenon can occur but appears to be unlikely to do so in practice.

Thanks

The correction made to the code so as to avoid underflow problems was made due to an inquiry and suggestion from Owen Marshall.

Author(s)

Rolf Turner <r.turner@auckland.ac.nz>

References

Rabiner, L. R., "A tutorial on hidden Markov models and selected applications in speech recognition," Proc. IEEE vol. 77, pp. 257 – 286, 1989.

See Also

[hmm\(\)](#), [rhmm\(\)](#), [mps\(\)](#), [pr\(\)](#), [viterbi\(\)](#)

Examples

```
# See the help for rhmm() for how to generate y.num and y.let.
## Not run:
fit.num <- hmm(y.num,K=2,verb=TRUE)
v.1 <- viterbi(model=fit.num)
v.2 <- viterbi(y.num,tpm=P,Rho=R) # P and R as in the
                                # help for rhmm().
# The order of the states has gotten swapped; 3-v.1[[1]] is much
# more similar to v.2[[1]] than is v.1[[1]].

fit.let <- hmm(y.let,K=2,verb=TRUE)
v.3 <- viterbi(model=fit.let)    # Works.
v.4 <- viterbi(y.let,tpm=P,Rho=R) # Throws an error (R has no row names.)

## End(Not run)
```

weissData

Data from “An Introduction to Discrete-Valued Time Series”

Description

Data sets from the book “An Introduction to Discrete-Valued Time Series” by Christian H. Weiß.

Usage

```
data(Bovine)
data(Cryptosporidiosis)
data(Downloads)
data(EricssonB_Jul2)
```

```
data(FattyLiver)
data(FattyLiver2)
data(goldparticle380)
data(Hanta)
data(InfantEEGsleepstates)
data(IPs)
data(LegionnairesDisease)
data(OffshoreRigcountsAlaska)
data(PriceStability)
data(Strokes)
data(WoodPeweeSong)
```

Format

- Bovine A character vector of length 8419.
- Cryptosporidiosis A numeric (integer) vector of length 365.
- Downloads A numeric (integer) vector of length 267.
- EricssonB_Jul12 A numeric (integer) vector of length 460.
- FattyLiver2 A numeric (integer) vector of length 449.
- FattyLiver A numeric (integer) vector of length 928.
- goldparticle380 A numeric (integer) vector of length 380.
- Hanta A numeric (integer) vector of length 52.
- InfantEEGsleepstates A character vector of length 107.
- IPs A numeric (integer) vector of length 241.
- LegionnairesDisease A numeric (integer) vector of length 365.
- OffshoreRigcountsAlaska A numeric (integer) vector of length 417.
- PriceStability A numeric (integer) vector of length 152.
- Strokes A numeric (integer) vector of length 108.
- WoodPeweeSong A numeric (integer) vector of length 1327.

Details

For detailed information about each of these data sets, see the book cited in the **References**.

Note that the data sets Cryptosporidiosis and LegionnairesDisease are actually called Cryptosporidiosis_02-08 and LegionnairesDisease_02-08 in the given reference. The “suffixes” were removed since the minus sign causes problems in a variable name in R.

Source

These data sets were kindly provided by Prof. Christian H. Weiß. The package author is also pleased to acknowledge the kind permission granted by Prof. Kurt Brännäs (Professor Emeritus of Economics at Umeå University) to include the Ericsson time series data set (EricssonB_Jul12).

References

Christian H. Weiß (2018). *An Introduction to Discrete-Valued Time Series*. Chichester: John Wiley & Sons.

Examples

```
## Not run:
fit1 <- hmm(WoodPeweeSong,K=2,verbose=TRUE)
# EM converges in 6 steps --- suspicious.
set.seed(321)
fit2 <- hmm(WoodPeweeSong,K=2,verbose=TRUE,rand.start=list(tpm=TRUE,Rho=TRUE))
# 52 steps --- note the huge difference between fit1$log.like and fit2$log.like!
set.seed(321)
fit3 <- hmm(WoodPeweeSong,K=2,verbose=TRUE,method="bf",
            rand.start=list(tpm=TRUE,Rho=TRUE))
# log likelihood essentially the same as for fit2

## End(Not run)
```

Index

- *Topic **datagen**
 - misstify, 20
 - rhmm, 27
- *Topic **datasets**
 - ccprSim, 3
 - hydroDat, 16
 - lesionCount, 17
 - SydColDisc, 34
 - weissData, 39
- *Topic **methods**
 - anova.hmm.discnp, 2
 - update.hmm.discnp, 35
- *Topic **models**
 - anova.hmm.discnp, 2
 - fitted.hmm.discnp, 5
 - hmm, 7
 - logLikHmm, 18
 - mps, 22
 - pr, 24
 - predict.hmm.discnp, 26
 - sp, 30
 - update.hmm.discnp, 35
 - viterbi, 37
- *Topic **utilities**
 - cnvrtRho, 4
 - nafracCalc, 23
- *Topic **utility**
 - scovmat, 29
 - squantCI, 32
- anova.hmm.discnp, 2
- Bovine (weissData), 39
- ccprSim, 3
- cnvrtRho, 4
- Cryptosporidiosis (weissData), 39
- cut, 34
- Downloads (weissData), 39
- EricssonB_Jul2 (weissData), 39
- FattyLiver (weissData), 39
- FattyLiver2 (weissData), 39
- fitted.hmm.discnp, 5, 25, 26, 32
- ftLiardFlows (hydroDat), 16
- goldparticle380 (weissData), 39
- Hanta (weissData), 39
- hmm, 2, 3, 5, 6, 7, 18–20, 22, 23, 25–29, 31–33, 35–39
- hydroDat, 16
- InfantEEGsleepstates (weissData), 39
- IPs (weissData), 39
- LegionnairesDisease (weissData), 39
- lesionCount, 17
- linLandFlows (hydroDat), 16
- logLikHmm, 18
- misstify, 20, 24, 28, 29
- model.matrix, 10
- mps, 15, 22, 25, 32, 39
- nafracCalc, 20, 21, 23, 29
- nlm, 9, 11, 13
- OffshoreRigcountsAlaska (weissData), 39
- optim, 8, 9, 13
- portMannFlows (hydroDat), 16
- portMannSedCon (hydroDat), 16
- portMannSedLoads (hydroDat), 16
- pr, 20, 24, 32, 39
- predict.hmm.discnp, 26
- PriceStability (weissData), 39
- rhmm, 15, 20, 21, 23, 24, 27, 39
- rhmm.hmm.discnp, 37

scovmat, [29](#), [33](#)
sp, [6](#), [20](#), [25–27](#), [30](#)
squantCI, [30](#), [32](#)
Strikes (weissData), [39](#)
SydColDisc, [34](#)

update.hmm.discnp, [13](#), [35](#)

viterbi, [15](#), [23](#), [25](#), [32](#), [37](#), [39](#)

weissData, [39](#)
WoodPeweeSong (weissData), [39](#)