

Package ‘msm’

May 11, 2012

Version 1.1.1

Date 2012-05-11

Title Multi-state Markov and hidden Markov models in continuous time

Author Christopher Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

Maintainer Christopher Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

Description Functions for fitting general continuous-time Markov and hidden Markov multi-state models to longitudinal data. Both Markov transition rates and the hidden Markov output process can be modelled in terms of covariates, which may be constant or piecewise-constant in time. A variety of observation schemes are supported, including processes observed at arbitrary times (panel data), continuously-observed processes, and censored states.

License GPL (>= 2)

Imports survival,mvtnorm

Suggests mstate,diagram

Repository CRAN

Date/Publication 2012-05-11 15:27:50

R topics documented:

aneur	2
boot.msm	3
bos	6
cav	7
coef.msm	8
crudeinits.msm	9
deltamethod	10
ematrix.msm	12
fev	13

hazard.msm	14
hmm-dists	15
logLik.msm	18
lrtest.msm	19
MatrixExp	19
medists	21
msm	23
msm.summary	35
odds.msm	36
pearson.msm	37
pexp	40
plot.msm	42
plot.prevalence.msm	43
plot.surfit.msm	44
plotprog.msm	46
pmatrix.msm	47
pmatrix.piecewise.msm	49
pnext.msm	51
prevalence.msm	53
psor	55
qmatrix.msm	56
qratio.msm	58
scorer resid.msm	59
sim.msm	60
simmulti.msm	62
sojourn.msm	64
statetable.msm	66
surface.msm	67
tnorm	68
totlos.msm	70
transient.msm	71
viterbi.msm	72

Index	74
--------------	-----------

aneur

Aortic aneurysm progression data

Description

This dataset contains longitudinal measurements of grades of aortic aneurysms, measured by ultrasound examination of the diameter of the aorta.

Usage

data(aneur)

Format

A data frame containing 4337 rows, with each row corresponding to an ultrasound scan from one of 838 men over 65 years of age.

ptnum	(numeric)	Patient identification number
age	(numeric)	Recipient age at examination (years)
diam	(numeric)	Aortic diameter
state	(numeric)	State of aneurysm.

The states represent successive degrees of aneurysm severity, as indicated by the aortic diameter.

State 1	Aneurysm-free	< 30 cm
State 2	Mild aneurysm	30-44 cm
State 3	Moderate aneurysm	45-54 cm
State 4	Severe aneurysm	> 55 cm

683 of these men were aneurysm-free at age 65 and were re-screened every two years. The remaining men were aneurysmal at entry and had successive screens with frequency depending on the state of the aneurysm. Severe aneurysms are repaired by surgery.

Source

The Chichester, U.K. randomised controlled trial of screening for abdominal aortic aneurysms by ultrasonography.

References

Jackson, C.H., Sharples, L.D., Thompson, S.G. and Duffy, S.W. and Couto, E. Multi-state Markov models for disease progression with classification error. *The Statistician*, 52(2): 193–209 (2003)

Couto, E. and Duffy, S. W. and Ashton, H. A. and Walker, N. M. and Myles, J. P. and Scott, R. A. P. and Thompson, S. G. (2002) *Probabilities of progression of aortic aneurysms: estimates and implications for screening policy* Journal of Medical Screening 9(1):40–42

 boot.msm

Bootstrap resampling for multi-state models

Description

Draw a number of bootstrap resamples, refit a `msm` model to the resamples, and calculate statistics on the refitted models.

Usage

```
boot.msm(x, stat=pmatix.msm, B=1000, file=NULL)
```

Arguments

x	A fitted msm model, as output by <code>msm</code> .
stat	A function to call on each refitted msm model. By default this is <code>pmatrix.msm</code> , returning the transition probability matrix in one time unit. If NULL then no function is computed.
B	Number of bootstrap resamples.
file	Name of a file in which to save partial results after each replicate. This is saved using <code>save</code> and can be restored using <code>load</code> , producing an object called <code>boot.list</code> containing the partial results. Not currently supported when using parallel processing.

Details

The bootstrap datasets are computed by resampling independent transitions between pairs of states (for non-hidden models without censoring), or independent individual series (for hidden models or models with censoring). Therefore this approach doesn't work if, for example, the data for a HMM consist of a series of observations from just one individual, and is inaccurate for small numbers of independent transitions or individuals.

Confidence intervals or standard errors for the corresponding statistic can be calculated by summarising the returned list of B replicated outputs. This is currently implemented for most the output functions `qmatrix.msm`, `ematrix.msm`, `qratio.msm`, `pmatrix.msm`, `pmatrix.piecewise.msm`, `totlos.msm` and `prevalence.msm`. For other outputs, users will have to write their own code to summarise the output of `boot.msm`.

Most of `msm`'s output functions present confidence intervals based on asymptotic standard errors calculated from the Hessian. These are expected to be underestimates of the true standard errors (Cramer-Rao lower bound). Some of these functions use a further approximation, the delta method (see `deltamethod`) to obtain standard errors of transformed parameters. Bootstrapping should give a more accurate estimate of the uncertainty.

An alternative method which is less accurate though faster than bootstrapping, but more accurate than the delta method, is to draw a sample from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix), and summarise the transformed estimates. See `pmatrix.msm`.

All objects used in the original call to `msm` which produced `x`, such as the `qmatrix`, should be in the working environment, or else `boot.msm` will produce an "object not found" error. This enables `boot.msm` to refit the original model to the replicate datasets. However there is currently a limitation. In the original call to `msm`, the "formula" argument should be specified directly, as, for example,

```
msm(state ~ time, data = ...)
```

and not, for example,

```
form = data$state ~ data$time
```

```
msm(formula=form, data = ...)
```

otherwise `boot.msm` will be unable to draw the replicate datasets.

If `stat` is NULL, then B different `msm` model objects will be stored in memory. This is inadvisable, as `msm` objects tend to be large, since they contain the original data used for the `msm` fit, so this will be wasteful of memory.

To specify more than one statistic, write a function consisting of a list of different function calls, for example,

```
stat = function(x) list (pmatrix.msm(x, t=1), pmatrix.msm(x, t=2))
```

Value

A list with B components, containing the result of calling function `stat` on each of the refitted models. If `stat` is NULL, then each component just contains the refitted model. If one of the B model fits was unsuccessful and resulted in an error, then the corresponding list component will contain the error message.

Author(s)

C.H.Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

Efron, B. and Tibshirani, R.J. (1993) *An Introduction to the Bootstrap*, Chapman and Hall.

See Also

[qmatrix.msm](#), [qratio.msm](#), [sojourn.msm](#), [ematrix.msm](#), [pmatrix.msm](#), [pmatrix.piecewise.msm](#), [totlos.msm](#), [prevalence.msm](#).

Examples

```
## Not run:
## Psoriatic arthritis example
data(psor)
psor.q <- rbind(c(0,0.1,0,0),c(0,0,0.1,0),c(0,0,0,0.1),c(0,0,0,0))
psor.msm <- msm(state ~ months, subject=ptnum, data=psor, qmatrix =
  psor.q, covariates = ~ollwsdrt+hieffusn,
  constraint = list(hieffusn=c(1,1,1),ollwsdrt=c(1,1,2)),
  control = list(REPORT=1,trace=2), method="BFGS")
## Bootstrap the baseline transition intensity matrix. This will take a long time.
q.list <- boot.msm(psor.msm, function(x)x$Qmatrices$baseline)
## Manipulate the resulting list of matrices to calculate bootstrap standard errors.
apply(array(unlist(q.list), dim=c(4,4,5)), c(1,2), sd)
## Similarly calculate a bootstrap 95% confidence interval
apply(array(unlist(q.list), dim=c(4,4,5)), c(1,2),
  function(x)quantile(x, c(0.025, 0.975)))
## Bootstrap standard errors are larger than the asymptotic standard
## errors calculated from the Hessian
psor.msm$QmatricesSE$baseline

## End(Not run)
```

bos

Bronchiolitis obliterans syndrome after lung transplants

Description

A dataset containing histories of bronchiolitis obliterans syndrome (BOS) from lung transplant recipients. BOS is a chronic decline in lung function, often observed after lung transplantation. The condition is classified into four stages of severity: none, mild, moderate and severe.

Usage

```
data(bos)
```

Format

A data frame containing 638 rows, grouped by patient, including histories of 204 patients. The first observation for each patient is defined to be stage 1, no BOS, at six months after transplant. Subsequent observations denote the entry times into stages 2, 3, 4, representing mild, moderate and severe BOS respectively, and stage 5, representing death.

ptnum	(numeric)	Patient identification number
time	(numeric)	Months after transplant
state	(numeric)	BOS state entered at this time

Details

The entry time of each patient into each stage of BOS was estimated by clinicians, based on their history of lung function measurements and acute rejection and infection episodes. BOS is only assumed to occur beyond six months after transplant. In the first six months the function of each patient's new lung stabilises. Subsequently BOS is diagnosed by comparing the lung function against the "baseline" value.

Source

Papworth Hospital, U.K.

References

Heng, D. et al. (1998). Bronchiolitis Obliterans Syndrome: Incidence, Natural History, Prognosis, and Risk Factors. *Journal of Heart and Lung Transplantation* 17(12)1255–1263.

cav *Heart transplant monitoring data*

Description

A series of approximately yearly angiographic examinations of heart transplant recipients. The state at each time is a grade of cardiac allograft vasculopathy (CAV), a deterioration of the arterial walls.

Usage

`data(cav)`

Format

A data frame containing 2846 rows. There are 622 patients, the rows are grouped by patient number and ordered by years after transplant, with each row representing an examination and containing additional covariates.

PTNUM	(numeric)	Patient identification number
age	(numeric)	Recipient age at examination (years)
years	(numeric)	Examination time (years after transplant)
dage	(numeric)	Age of heart donor (years)
sex	(numeric)	sex (0=male, 1=female)
pdiag	(factor)	Primary diagnosis (reason for transplant) IHD=ischaemic heart disease, IDC=idiopathic dilated cardiomyopathy.
cumrej	(numeric)	Cumulative number of acute rejection episodes
state	(numeric)	State at the examination. State 1 represents no CAV, state 2 is mild/moderate CAV and state 3 is severe CAV. State 4 indicates death.
firstobs	(numeric)	0 = record represents an angiogram or date of death. 1 = record represents transplant (patient's first observation)

Source

Papworth Hospital, U.K.

References

Sharples, L.D. and Jackson, C.H. and Parameshwar, J. and Wallwork, J. and Large, S.R. (2003). Diagnostic accuracy of coronary angiopathy and risk factors for post-heart-transplant cardiac allograft vasculopathy. *Transplantation* 76(4):679-82

coef.msm	<i>Extract model coefficients</i>
----------	-----------------------------------

Description

Extract the estimated log transition intensities and the corresponding linear effects of each covariate.

Usage

```
## S3 method for class 'msm'  
coef(object, ...)
```

Arguments

object	A fitted multi-state model object, as returned by msm .
...	(unused) further arguments passed to or from other methods.

Value

If there is no misclassification, `coef.msm` returns a list of matrices. The first component, labelled `logbaseline`, is a matrix containing the estimated transition intensities on the log scale with any covariates fixed at their means in the data. Each remaining component is a matrix giving the linear effects of the labelled covariate on the matrix of log intensities.

For misclassification models, `coef.msm` returns a list of lists. The first component, `Qmatrices`, is a list of matrices as described in the previous paragraph. The additional component `Ematrices` is a list of similar format containing the logit-misclassification probabilities and any estimated covariate effects.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#)

crudeinits.msm	<i>Calculate crude initial values for transition intensities</i>
----------------	--

Description

Calculates crude initial values for transition intensities by assuming that the data represent the exact transition times of the Markov process.

Usage

```
crudeinits.msm(formula, subject, qmatrix, data=NULL, censor=NULL, censor.states=NULL)
```

Arguments

formula	A formula giving the vectors containing the observed states and the corresponding observation times. For example, <code>state ~ time</code> Observed states should be in the set $1, \dots, n$, where n is the number of states.
subject	Vector of subject identification numbers for the data specified by formula. If missing, then all observations are assumed to be on the same subject. These must be sorted so that all observations on the same subject are adjacent.
qmatrix	Matrix of indicators for the allowed transitions. An initial value will be estimated for each value of qmatrix that is greater than zero. Transitions are taken as disallowed for each entry of qmatrix that is 0.
data	An optional data frame in which the variables represented by subject and state can be found.
censor	A state, or vector of states, which indicates censoring. See msm .
censor.states	Specifies the underlying states which censored observations can represent. See msm .

Details

Suppose we want a crude estimate of the transition intensity q_{rs} from state r to state s . If we observe n_{rs} transitions from state r to state s , and a total of n_r transitions from state r , then q_{rs}/q_{rr} can be estimated by n_{rs}/n_r . Then, given a total of T_r years spent in state r , the mean sojourn time $1/q_{rr}$ can be estimated as T_r/n_r . Thus, n_{rs}/T_r is a crude estimate of q_{rs} .

If the data do represent the exact transition times of the Markov process, then these are the exact maximum likelihood estimates.

Observed transitions which are incompatible with the given qmatrix are ignored. Censored states are ignored.

Value

The estimated transition intensity matrix. This can be used as the qmatrix argument to [msm](#).

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[statetable.msm](#)

Examples

```
data(cav)
twoway4.q <- rbind(c(-0.5, 0.25, 0, 0.25), c(0.166, -0.498, 0.166, 0.166),
c(0, 0.25, -0.5, 0.25), c(0, 0, 0, 0))
statetable.msm(state, PTNUM, data=cav)
crudeinits.msm(state ~ years, PTNUM, data=cav, qmatrix=twoway4.q)
```

deltamethod

The delta method

Description

Delta method for approximating the standard error of a transformation $g(X)$ of a random variable $X = (x_1, x_2, \dots)$, given estimates of the mean and covariance matrix of X .

Usage

```
deltamethod(g, mean, cov, ses=TRUE)
```

Arguments

g	A formula representing the transformation. The variables must be labelled x_1, x_2, \dots . For example, $\sim 1 / (x_1 + x_2)$ If the transformation returns a vector, then a list of formulae representing (g_1, g_2, \dots) can be provided, for example <code>list($\sim x_1 + x_2$, $\sim x_1 / (x_1 + x_2)$)</code>
mean	The estimated mean of X
cov	The estimated covariance matrix of X
ses	If TRUE, then the standard errors of $g_1(X), g_2(X), \dots$ are returned. Otherwise the covariance matrix of $g(X)$ is returned.

Details

The delta method expands a differentiable function of a random variable about its mean, usually with a first-order Taylor approximation, and then takes the variance. For example, an approximation to the covariance matrix of $g(X)$ is given by

$$\text{Cov}(g(X)) = g'(\mu)\text{Cov}(X)[g'(\mu)]^T$$

where μ is an estimate of the mean of X .

A limitation of this function is that variables created by the user are not visible within the formula g . To work around this, it is necessary to build the formula as a string, using functions such as `sprintf`, then to convert the string to a formula using `as.formula`. See the example below.

If you can spare the computational time, bootstrapping is a more accurate method of calculating confidence intervals or standard errors for transformations of parameters. See [boot.msm](#).

Value

A vector containing the standard errors of $g_1(X), g_2(X), \dots$ or a matrix containing the covariance of $g(X)$.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

Oehlert, G. W. *A note on the delta method*. *American Statistician* 46(1), 1992

Examples

```
## Simple linear regression, E(y) = alpha + beta x
x <- 1:100
y <- rnorm(100, 4*x, 5)
toy.lm <- lm(y ~ x)
estmean <- coef(toy.lm)
estvar <- summary(toy.lm)$cov.unscaled * summary(toy.lm)$sigma^2

## Estimate of (1 / (alphahat + betahat))
1 / (estmean[1] + estmean[2])
## Approximate standard error
deltamethod (~ 1 / (x1 + x2), estmean, estvar)

## We have a variable z we would like to use within the formula.
z <- 1
## deltamethod (~ z / (x1 + x2), estmean, estvar) will not work.
## Instead, build up the formula as a string, and convert to a formula.
form <- sprintf("~ %f / (x1 + x2)", z)
form
deltamethod(as.formula(form), estmean, estvar)
```

ematrix.msm

*Misclassification probability matrix***Description**

Extract the estimated misclassification probability matrix, and corresponding confidence intervals, from a fitted multi-state model at a given set of covariate values.

Usage

```
ematrix.msm(x, covariates="mean", ci=c("delta","normal","bootstrap","none"),
            cl=0.95, B=1000)
```

Arguments

x	A fitted multi-state model, as returned by msm .
covariates	<p>The covariate values for which to estimate the misclassification probability matrix. This can either be:</p> <p>the string "mean", denoting the means of the covariates in the data (this is the default),</p> <p>the number 0, indicating that all the covariates should be set to zero,</p> <p>or a list of values, with optional names. For example</p> <pre>list (60, 1)</pre> <p>where the order of the list follows the order of the covariates originally given in the model formula, or a named list,</p> <pre>list (age = 60, sex = 1)</pre>
ci	<p>If "delta" (the default) then confidence intervals are calculated by the delta method, or by simple transformation of the Hessian in the very simplest cases.</p> <p>If "normal", then calculate a confidence interval by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the multinomial-logit-transformed misclassification probabilities and covariate effects, then transforming back.</p> <p>If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm.</p>
cl	Width of the symmetric confidence interval to present. Defaults to 0.95.
B	Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs

Details

Misclassification probabilities and covariate effects are estimated on the multinomial-logit scale by [msm](#). A covariance matrix is estimated from the Hessian of the maximised log-likelihood. From these, the delta method can be used to obtain standard errors of the probabilities on the natural scale at arbitrary covariate values. Confidence intervals are estimated by assuming normality on the multinomial-logit scale.

Value

A list with components:

estimate	Estimated misclassification probability matrix. The rows correspond to true states, and columns observed states.
SE	Corresponding approximate standard errors.
L	Lower confidence limits.
U	Upper confidence limits.

Or if `ci="none"`, then `ematrix.msm` just returns the estimated misclassification probability matrix. The default print method for objects returned by `ematrix.msm` presents estimates and confidence limits. To present estimates and standard errors, do something like `ematrix.msm(x)[c("estimates", "SE")]`

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[qmatrix.msm](#)

 fev

FEV1 measurements from lung transplant recipients

Description

A series of measurements of the forced expiratory volume in one second (FEV1) from lung transplant recipients, from six months onwards after their transplant.

Usage

```
data(fev)
```

Format

A data frame containing 5896 rows. There are 204 patients, the rows are grouped by patient number and ordered by days after transplant. Each row represents an examination and containing an additional covariate.

ptnum	(numeric)	Patient identification number.
days	(numeric)	Examination time (days after transplant).
fev	(numeric)	Percentage of baseline FEV1. A code of 999 indicates the patient's date of death.
acute	(numeric)	0/1 indicator for whether the patient suffered an acute infection or rejection within 14 days of the visit.

Details

A baseline "normal" FEV1 for each individual is calculated using measurements from the first six months after transplant. After six months, as presented in this dataset, FEV1 is expressed as a percentage of the baseline value.

FEV1 is monitored to diagnose bronchiolitis obliterans syndrome (BOS), a long-term lung function decline, thought to be a form of chronic rejection. Acute rejections and infections also affect the lung function in the short term.

Source

Papworth Hospital, U.K.

References

Jackson, C.H. and Sharples, L.D. Hidden Markov models for the onset and progression of bronchiolitis obliterans syndrome in lung transplant recipients *Statistics in Medicine*, 21(1): 113–128 (2002).

hazard.msm	<i>Calculate tables of hazard ratios for covariates on transition intensities</i>
------------	---

Description

Hazard ratios are computed by exponentiating the estimated covariate effects on the log-transition intensities. This function is called by [summary.msm](#).

Usage

```
hazard.msm(x, hazard.scale = 1, c1 = 0.95)
```

Arguments

x	Output from msm representing a fitted multi-state model.
hazard.scale	Vector with same elements as number of covariates on transition rates. Corresponds to the increase in each covariate used to calculate its hazard ratio. Defaults to all 1.
c1	Width of the symmetric confidence interval to present. Defaults to 0.95.

Value

A list of tables containing hazard ratio estimates, one table for each covariate. Each table has three columns, containing the hazard ratio, and an approximate upper and lower confidence limit respectively (assuming normality on the log scale), for each Markov chain transition intensity.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#), [summary.msm](#), [odds.msm](#)

hmm-dists

Hidden Markov model constructors

Description

These functions are used to specify the distribution of the response conditionally on the underlying state in a hidden Markov model. A list of these function calls, with one component for each state, should be used for the `hmodel` argument to `msm`. The initial values for the parameters of the distribution should be given as arguments.

Usage

```
hmmCat(prob, basecat)
hmmIdent(x)
hmmUnif(lower, upper)
hmmNorm(mean, sd)
hmmLNorm(meanlog, sdlog)
hmmExp(rate)
hmmGamma(shape, rate)
hmmWeibull(shape, scale)
hmmPois(rate)
hmmBinom(size, prob)
hmmTNorm(mean, sd, lower, upper)
hmmMETNorm(mean, sd, lower, upper, sderr, meanerr=0)
hmmMEUnif(lower, upper, sderr, meanerr=0)
hmmNBinom(dis, prob)
hmmBeta(shape1, shape2)
```

Arguments

`hmmCat` represents a categorical response distribution on the set $1, 2, \dots, \text{length}(\text{prob})$. The Markov model with misclassification is an example of this type of model. The categories in this case are (some subset of) the underlying states.

The `hmmIdent` distribution is used for underlying states which are observed exactly without error.

`hmmUnif`, `hmmNorm`, `hmmLNorm`, `hmmExp`, `hmmGamma`, `hmmWeibull`, `hmmPois`, `hmmBinom`, `hmmTNorm`, `hmmNBinom` and `hmmBeta` represent Uniform, Normal, log-Normal, exponential, Gamma, Weibull, Poisson, Binomial, truncated Normal, negative binomial and beta distributions, respectively, with parameterisations the same as the default parameterisations in the corresponding base R distribution functions.

The `hmmMETNorm` and `hmmMEUnif` distributions are truncated Normal and Uniform distributions, but with additional Normal measurement error on the response. These are generalisations of the distributions proposed by Satten and Longini (1996) for modelling the progression of CD4 cell counts in monitoring HIV disease. See [medists](#) for density, distribution, quantile and random generation functions for these distributions. See also [tnorm](#) for density, distribution, quantile and random generation functions for the truncated Normal distribution.

(`hmmCat`) Vector of probabilities of observing category 1, 2, ..., `length(prob)` respectively. Or the probability governing a binomial or negative binomial distribution.

<code>probcat</code>	(<code>hmmCat</code>) Category which is considered to be the "baseline", so that during estimation, the probabilities are parameterised as probabilities relative to this baseline category. By default, the category with the greatest probability is used as the baseline.
<code>x</code>	(<code>hmmIdent</code>) Code in the data which denotes the exactly-observed state.
<code>mean</code>	(<code>hmmNorm</code> , <code>hmmLNorm</code> , <code>hmmTNorm</code>) Mean defining a Normal, or truncated Normal distribution.
<code>sd</code>	(<code>hmmNorm</code> , <code>hmmLNorm</code> , <code>hmmTNorm</code>) Standard deviation defining a Normal, or truncated Normal distribution.
<code>meanlog</code>	(<code>hmmNorm</code> , <code>hmmLNorm</code> , <code>hmmTNorm</code>) Mean on the log scale, for a log Normal distribution.
<code>sdlog</code>	(<code>hmmNorm</code> , <code>hmmLNorm</code> , <code>hmmTNorm</code>) Standard deviation on the log scale, for a log Normal distribution.
<code>rate</code>	(<code>hmmPois</code> , <code>hmmExp</code> , <code>hmmGamma</code>) Rate of a Poisson, Exponential or Gamma distribution (see dpois , dexp , dgamma).
<code>shape</code>	(<code>hmmPois</code> , <code>hmmExp</code> , <code>hmmGamma</code>) Shape parameter of a Gamma or Weibull distribution (see dgamma , dweibull).
<code>shape1</code> , <code>shape2</code>	First and second parameters of a beta distribution (see dbeta).
<code>scale</code>	(<code>hmmGamma</code>) Scale parameter of a Gamma distribution (see dgamma).
<code>size</code>	Order of a Binomial distribution (see dbinom).
<code>disp</code>	Dispersion parameter of a negative binomial distribution, also called <code>size</code> or <code>order</code> . (see dnbinom).
<code>lower</code>	(<code>hmmUnif</code> , <code>hmmTNorm</code> , <code>hmmMEUnif</code>) Lower limit for an Uniform or truncated Normal distribution.
<code>upper</code>	(<code>hmmUnif</code> , <code>hmmTNorm</code> , <code>hmmMEUnif</code>) Upper limit for an Uniform or truncated Normal distribution.

sderr	(hmmMETNorm, hmmUnif) Standard deviation of the Normal measurement error distribution.
meanerr	(hmmMETNorm, hmmUnif) Additional shift in the measurement error, fixed to 0 by default. This may be modelled in terms of covariates.

Details

See the PDF manual ‘msm-manual.pdf’ in the ‘doc’ subdirectory for algebraic definitions of all these distributions. New hidden Markov model response distributions can be added to **msm** by following the instructions in Section 2.17.1.

Parameters which can be modelled in terms of covariates, on the scale of a link function, are as follows.

PARAMETER NAME	LINK FUNCTION
mean	identity
meanlog	identity
rate	log
scale	log
meanerr	identity
prob	(multinomial logistic regression)

Parameters basecat, lower, upper, size, meanerr are fixed at their initial values. All other parameters are estimated while fitting the hidden Markov model, unless the appropriate `fixedpars` argument is supplied to `msm`.

For categorical response distributions (`hmmCat`) the outcome probabilities initialized to zero are fixed at zero, and the probability corresponding to basecat is fixed to one minus the sum of the remaining probabilities. These remaining probabilities are estimated, and can be modelled in terms of covariates via multinomial logistic regression (relative to basecat).

Value

Each function returns an object of class `hmodel`, which is a list containing information about the model. The only component which may be useful to end users is `r`, a function of one argument `n` which returns a random sample of size `n` from the given distribution.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

- Satten, G.A. and Longini, I.M. Markov chains with measurement error: estimating the ‘true’ course of a marker of the progression of human immunodeficiency virus disease (with discussion) *Applied Statistics* 45(3): 275-309 (1996).
- Jackson, C.H. and Sharples, L.D. Hidden Markov models for the onset and progression of bronchiolitis obliterans syndrome in lung transplant recipients *Statistics in Medicine*, 21(1): 113–128 (2002).

Jackson, C.H., Sharples, L.D., Thompson, S.G. and Duffy, S.W. and Couto, E. Multi-state Markov models for disease progression with classification error. *The Statistician*, 52(2): 193–209 (2003).

See Also

[msm](#)

logLik.msm

Extract model log-likelihood

Description

Extract the log-likelihood and the number of parameters of a model fitted with [msm](#).

Usage

```
## S3 method for class 'msm'  
logLik(object, ...)
```

Arguments

`object` A fitted multi-state model object, as returned by [msm](#).
`...` (unused) further arguments passed to or from other methods.

Value

The log-likelihood of the model represented by 'object' evaluated at the maximum likelihood estimates.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#)

lrtest.msm	<i>Likelihood ratio test</i>
------------	------------------------------

Description

Likelihood ratio test between two or more fitted multi-state models

Usage

```
lrtest.msm(...)
```

Arguments

... Two or more fitted multi-state models, as returned by [msm](#), supplied in increasing order of complexity.

Value

A matrix with three columns, giving the likelihood ratio statistic, difference in degrees of freedom and the chi-squared p-value for a comparison of the first model supplied with each subsequent model.

Warning

The comparison between models will only be valid if they are fitted to the same dataset. This may be a problem if there are missing values and R's default of 'na.action = na.omit' is used.

The likelihood ratio statistic only has the indicated chi-squared distribution if the models are nested. An alternative for comparing non-nested models is Akaike's information criterion.

See Also

[logLik.msm,msm](#)

MatrixExp	<i>Matrix exponential</i>
-----------	---------------------------

Description

Calculates the exponential of a square matrix.

Usage

```
MatrixExp(mat, t = 1, n = 20, k = 3, method="pade")
```

Arguments

mat	A square matrix
t	An optional scaling factor, or a vector or scaling factors, for the eigenvalues of mat
n	Number of terms in the series approximation to the exponential
k	Underflow correction factor, for the series approximation
method	"pade" for the Pade approximation, or "series" for the power series approximation. Ignored unless mat has repeated eigenvalues.

Details

The exponential E of a square matrix M is calculated as

$$E = U \exp(D) U^{-1}$$

where D is a diagonal matrix with the eigenvalues of M on the diagonal, $\exp(D)$ is a diagonal matrix with the exponentiated eigenvalues of M on the diagonal, and U is a matrix whose columns are the eigenvectors of M .

This method of calculation is used if M has distinct eigenvalues. If M has repeated eigenvalues, then its eigenvector matrix may be non-invertible. In this case, the matrix exponential is calculated using the Pade approximation defined by Moler and van Loan (2003), or the less robust power series approximation,

$$\exp(M) = I + M + M^2/2 + M^3/3! + M^4/4! + \dots$$

For a continuous-time homogeneous Markov process with transition intensity matrix Q , the probability of occupying state s at time $u + t$ conditional on occupying state r at time u is given by the (r, s) entry of the matrix $\exp(tQ)$.

The implementation of the Pade approximation was taken from JAGS by Martyn Plummer (<http://www-fis.iarc.fr/~martyn/software/jags>).

The series approximation method was adapted from the corresponding function in Jim Lindsey's R package rmutil (<http://popgen.unimaas.nl/~jlindsey/rcode.html>).

Value

The exponentiated matrix $\exp(mat)$. Or, if t is a vector of length 2 or more, an array of exponentiated matrices.

References

- Cox, D. R. and Miller, H. D. *The theory of stochastic processes*, Chapman and Hall, London (1965)
- Moler, C and van Loan, C (2003). Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. *SIAM Review* **45**, 3–49.
- At <http://epubs.siam.org/sam-bin/dbq/article/41801>

Description

Truncated Normal and Uniform distributions, where the response is also subject to a Normally distributed measurement error.

Usage

```
dmenorm(x, mean=0, sd=1, lower=-Inf, upper=Inf, sderr=0, meanerr=0,
        log = FALSE)
pmenorm(q, mean=0, sd=1, lower=-Inf, upper=Inf, sderr=0, meanerr=0,
        lower.tail = TRUE, log.p = FALSE)
qmenorm(p, mean=0, sd=1, lower=-Inf, upper=Inf, sderr=0, meanerr=0,
        lower.tail = TRUE, log.p = FALSE)
rmenorm(n, mean=0, sd=1, lower=-Inf, upper=Inf, sderr=0, meanerr=0)
dmeunif(x, lower=0, upper=1, sderr=0, meanerr=0, log = FALSE)
pmeunif(q, lower=0, upper=1, sderr=0, meanerr=0, lower.tail = TRUE,
        log.p = FALSE)
qmeunif(p, lower=0, upper=1, sderr=0, meanerr=0, lower.tail = TRUE,
        log.p = FALSE)
rmeunif(n, lower=0, upper=1, sderr=0, meanerr=0)
```

Arguments

x, q	vector of quantiles.
p	vector of probabilities.
n	number of observations. If <code>length(n) > 1</code> , the length is taken to be the number required.
mean	vector of means.
sd	vector of standard deviations.
lower	lower truncation point.
upper	upper truncation point.
sderr	Standard deviation of measurement error distribution.
meanerr	Optional shift for the measurement error distribution.
log, log.p	logical; if TRUE, probabilities p are given as $\log(p)$.
lower.tail	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise, $P[X > x]$.

Details

The normal distribution with measurement error has density

$$\frac{\Phi(u, \mu_2, \sigma_3) - \Phi(l, \mu_2, \sigma_3)}{\Phi(u, \mu_0, \sigma_0) - \Phi(l, \mu_0, \sigma_0)} \phi(x, \mu_0 + \mu_\epsilon, \sigma_2)$$

where

$$\begin{aligned}\sigma_2^2 &= \sigma_0^2 + \sigma_\epsilon^2, \\ \sigma_3 &= \sigma_0 \sigma_\epsilon / \sigma_2, \\ \mu_2 &= (x - \mu_\epsilon) \sigma_0^2 + \mu_0 \sigma_\epsilon^2,\end{aligned}$$

μ_0 is the mean of the original Normal distribution before truncation,

σ_0 is the corresponding standard deviation,

u is the upper truncation point,

l is the lower truncation point,

σ_ϵ is the standard deviation of the additional measurement error,

μ_ϵ is the mean of the measurement error (usually 0).

$\phi(x)$ is the density of the corresponding normal distribution, and

$\Phi(x)$ is the distribution function of the corresponding normal distribution.

The uniform distribution with measurement error has density

$$(\Phi(x, \mu_\epsilon + l, \sigma_\epsilon) - \Phi(x, \mu_\epsilon + u, \sigma_\epsilon)) / (u - l)$$

These are calculated from the original truncated Normal or Uniform density functions $f(\cdot | \mu, \sigma, l, u)$ as

$$\int f(y | \mu, \sigma, l, u) \phi(x, y + \mu_\epsilon, \sigma_\epsilon) dy$$

If `sderr` and `meanerr` are not specified they assume the default values of 0, representing no measurement error variance, and no constant shift in the measurement error, respectively.

Therefore, for example with no other arguments, `dmenorm(x)`, is simply equivalent to `dtnorm(x)`, which in turn is equivalent to `dnorm(x)`.

These distributions were used by Satten and Longini (1996) for CD4 cell counts conditionally on hidden Markov states of HIV infection, and later by Jackson and Sharples (2002) for FEV1 measurements conditionally on states of chronic lung transplant rejection.

These distribution functions are just provided for convenience, and are not optimised for numerical accuracy or speed. To fit a hidden Markov model with these response distributions, use a `hmmMETNorm` or `hmmMEUnif` constructor. See the [hmm-dists](#) help page for further details.

Value

`dmenorm`, `dmeunif` give the density, `pmenorm`, `pmeunif` give the distribution function, `qmenorm`, `qmeunif` give the quantile function, and `rmenorm`, `rmeunif` generate random deviates, for the Normal and Uniform versions respectively.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

Satten, G.A. and Longini, I.M. Markov chains with measurement error: estimating the 'true' course of a marker of the progression of human immunodeficiency virus disease (with discussion) *Applied Statistics* 45(3): 275-309 (1996)

Jackson, C.H. and Sharples, L.D. Hidden Markov models for the onset and progression of bronchiolitis obliterans syndrome in lung transplant recipients *Statistics in Medicine*, 21(1): 113–128 (2002).

See Also

[dnorm](#), [dunif](#), [dtnorm](#)

Examples

```
## what does the distribution look like?
x <- seq(50, 90, by=1)
plot(x, dnorm(x, 70, 10), type="l", ylim=c(0,0.06)) ## standard Normal
lines(x, dtnorm(x, 70, 10, 60, 80), type="l")      ## truncated Normal
## truncated Normal with small measurement error
lines(x, dmenorm(x, 70, 10, 60, 80, sderr=3), type="l")
```

msm

Multi-state Markov and hidden Markov models in continuous time

Description

Fit a continuous-time Markov or hidden Markov multi-state model by maximum likelihood. Observations of the process can be made at arbitrary times, or the exact times of transition between states can be known. Covariates can be fitted to the Markov chain transition intensities or to the hidden Markov observation process.

Usage

```
msm ( formula, subject=NULL, data = list(), qmatrix, gen.inits = FALSE,
      ematrix=NULL, hmodel=NULL, obstype=NULL, obstrue=NULL,
      covariates = NULL, covinits = NULL, constraint = NULL,
      misccovariates = NULL, misccovinits = NULL, miscconstraint = NULL,
      hcovariates = NULL, hcovinits = NULL, hconstraint = NULL,
      qconstraint=NULL, econstraint=NULL, initprobs = NULL,
      est.initprobs=FALSE, initcovariates = NULL, initcovinits = NULL,
      death = FALSE, exacttimes = FALSE, censor=NULL,
      censor.states=NULL, pci=NULL, cl = 0.95, fixedpars = NULL, center=TRUE,
      opt.method=c("optim","nlm"), hessian=TRUE, use.deriv=TRUE,
      analyticp=TRUE, ... )
```

Arguments

formula	<p>A formula giving the vectors containing the observed states and the corresponding observation times. For example,</p> <pre>state ~ time</pre> <p>Observed states should be in the set $1, \dots, n$, where n is the number of states. The times can indicate different types of observation scheme, so be careful to choose the correct <code>obstype</code>.</p> <p>For hidden Markov models, <code>state</code> refers to the outcome variable, which need not be a discrete state.</p>
subject	<p>Vector of subject identification numbers for the data specified by formula. If missing, then all observations are assumed to be on the same subject. These must be sorted so that all observations on the same subject are adjacent.</p>
data	<p>Optional data frame in which to interpret the variables supplied in formula, <code>subject</code>, <code>covariates</code>, <code>miscovariates</code>, <code>hcovariates</code>, <code>obstype</code> and <code>obstrue</code>.</p>
qmatrix	<p>Initial transition intensity matrix of the Markov chain. If an instantaneous transition is not allowed from state r to state s, then <code>qmatrix</code> should have (r, s) entry 0, otherwise it should be non-zero. Any diagonal entry of <code>qmatrix</code> is ignored, as it is constrained to be equal to minus the sum of the rest of the row. For example,</p> <pre>rbind(c(0, 0.1, 0.01), c(0.1, 0, 0.2), c(0, 0, 0))</pre> <p>represents a 'health - disease - death' model, with transition intensities 0.1 from health to disease, 0.01 from health to death, 0.1 from disease to health, and 0.2 from disease to death. The initial intensities given here are with any covariates set to their means in the data (or set to zero, if <code>center = FALSE</code>). If any intensities are constrained to be equal using <code>qconstraint</code>, then the initial value is taken from the first of these (reading across rows).</p>
gen.inits	<p>If TRUE, then initial values for the transition intensities are generated automatically using the method in <code>crudeinits.msm</code>. The non-zero entries of the supplied <code>qmatrix</code> are assumed to indicate the allowed transitions of the model.</p>
ematrix	<p>If misclassification between states is to be modelled, this should be a matrix of initial values for the misclassification probabilities. The rows represent underlying states, and the columns represent observed states. If an observation of state s is not possible when the subject occupies underlying state r, then <code>ematrix</code> should have (r, s) entry 0. Otherwise <code>ematrix</code> should have (r, s) entry corresponding to the probability of observing s conditionally on occupying true state r. The diagonal of <code>ematrix</code> is ignored, as rows are constrained to sum to 1. For example,</p> <pre>rbind(c(0, 0.1, 0), c(0.1, 0, 0.1), c(0, 0.1, 0))</pre> <p>represents a model in which misclassifications are only permitted between adjacent states.</p>

If any probabilities are constrained to be equal using `econstraint`, then the initial value is taken from the first of these (reading across rows).

For an alternative way of specifying misclassification models, see `hmodel`.

`hmodel`

Specification of the hidden Markov model. This should be a list of return values from the constructor functions described in the [hmm-dists](#) help page. Each element of the list corresponds to the outcome model conditionally on the corresponding underlying state.

For example, consider a three-state hidden Markov model. Suppose the observations in underlying state 1 are generated from a Normal distribution with mean 100 and standard deviation 16, while observations in underlying state 2 are Normal with mean 54 and standard deviation 18. Observations in state 3, representing death, are exactly observed, and coded as 999 in the data. This model is specified as

```
hmodel = list(hmmNorm(mean=100, sd=16), hmmNorm(mean=54, sd=18), hmmIdent(999))
```

The mean and standard deviation parameters are estimated starting from these initial values. If multiple parameters are constrained to be equal using `hconstraint`, then the initial value is taken from the value given on the first occasion that parameter appears in `hmodel`.

See the [hmm-dists](#) help page for details of the constructor functions for each available distribution.

A misclassification model, that is, a hidden Markov model where the outcomes are misclassified observations of the underlying states, can either be specified using a list of [hmmCat](#) objects, or by using an `ematrix`.

For example,

```
ematrix = rbind( c( 0, 0.1, 0, 0 ), c( 0.1, 0, 0.1, 0 ), c( 0, 0.1, 0, 0 ), c(
)
```

is equivalent to

```
hmodel = list( hmmCat(prob=c(0.9, 0.1, 0, 0)), hmmCat(prob=c(0.1, 0.8, 0.1, 0)),
```

`obstype`

A vector specifying the observation scheme for each row of the data. This can be included in the data frame `data` along with the state, time, subject IDs and covariates. Its elements should be either 1, 2 or 3, meaning as follows:

- 1 An observation of the process at an arbitrary time (a "snapshot" of the process, or "panel-observed" data). The states are unknown between observation times.
- 2 An exact transition time, with the state at the previous observation retained until the current observation.
- 3 An exact transition time, but the state at the instant before entering this state is unknown. A common example is death times in studies of chronic diseases.

If `obstype` is not specified, this defaults to all 1. If `obstype` is a single number, all observations are assumed to be of this type. The `obstype` value for the first observation from each subject is not used.

- This is a generalisation of the `death` and `exacttimes` arguments to allow different schemes per observation. `obstype` overrides both `death` and `exacttimes`. `exacttimes=TRUE` specifies that all observations are of `obstype` 2.
- `death = death.states` specifies that all observations of `death.states` are of type 3. `death = TRUE` specifies that all observations in the final absorbing state are of type 3.
- obstrue** A vector of logicals (TRUE or FALSE) or numerics (1 or 0) specifying which observations (TRUE, 1) are observations of the underlying state without error, and which (FALSE, 0) are realisations of a hidden Markov model. Only used in misclassification or hidden Markov models.
- In HMMs where there are also censored states, `obstrue` should be set to 1 for observed states which are *censored* but not *misclassified*.
- covariates** Formula representing the covariates on the transition intensities via a log-linear model. For example,
- ```
~ age + sex + treatment
```
- covinits** Initial values for log-linear effects of covariates on the transition intensities. This should be a named list with each element corresponding to a covariate. A single element contains the initial values for that covariate on each transition intensity, reading across the rows in order. For a pair of effects constrained to be equal, the initial value for the first of the two effects is used.
- For example, for a model with the above `qmatrix` and `age` and `sex` covariates, the following initialises all covariate effects to zero apart from the `age` effect on the 2-1 transition, and the `sex` effect on the 1-3 transition. `covinits = list(sex=c(0, 0, 0.1, 0), age=c(0, 0, 0, 0.1, 0))`
- For factor covariates, name each level by concatenating the name of the covariate with the level name, quoting if necessary. For example, for a covariate `agegroup` with three levels 0-15, 15-60, 60-, use something like
- ```
covinits = list("agegroup15-60"=c(0, 0.1, 0, 0), "agegroup60-"=c(0.1, 0.1, 0, 0))
```
- If not specified or wrongly specified, initial values are assumed to be zero.
- constraint** A list of one numeric vector for each named covariate. The vector indicates which covariate effects on intensities are constrained to be equal. Take, for example, a model with five transition intensities and two covariates. Specifying
- ```
constraint = list (age = c(1,1,1,2,2), treatment = c(1,2,3,4,5))
```
- constrains the effect of `age` to be equal for the first three intensities, and equal for the fourth and fifth. The effect of `treatment` is assumed to be different for each intensity. Any vector of increasing numbers can be used as indicators. The intensity parameters are assumed to be ordered by reading across the rows of the transition matrix, starting at the first row, ignoring the diagonals.
- Negative elements of the vector can be used to indicate that particular covariate effects are constrained to be equal to minus some other effects. For example:
- ```
constraint = list (age = c(-1,1,1,2,-2), treatment = c(1,2,3,4,5))
```
- constrains the second and third `age` effects to be equal, the first effect to be minus the second, and the fifth `age` effect to be minus the fourth. For example, it may be

realistic that the effect of a covariate on the "reverse" transition rate from state 2 to state 1 is minus the effect on the "forward" transition rate, state 1 to state 2. Note that it is not possible to specify exactly which of the covariate effects are constrained to be positive and which negative. The maximum likelihood estimation chooses the combination of signs which has the higher likelihood.

For categorical covariates, defined using `factor(covname)`, specify constraints as follows:

```
list(..., covnameVALUE1 = c(...), covnameVALUE2 = c(...), ...)
```

where `covname` is the name of the factor, and `VALUE1`, `VALUE2`, ... are the levels of the factor. Make sure the `contrasts` option is set appropriately, for example, the default

```
options(contrasts=c(contr.treatment, contr.poly))
```

sets the first (baseline) level of unordered factors to zero.

To assume no covariate effect on a certain transition, use the `fixedpars` argument to fix it at its initial value (which is zero by default) during the optimisation.

- | | |
|-----------------------------|--|
| <code>misccovariates</code> | A formula representing the covariates on the misclassification probabilities, analogously to <code>covariates</code> , via multinomial logistic regression. Only used if the model is specified using <code>ematrix</code> , rather than <code>hmodel</code> . |
| <code>misccovinits</code> | Initial values for the covariates on the misclassification probabilities, defined in the same way as <code>covinits</code> . Only used if the model is specified using <code>ematrix</code> . |
| <code>misconstraint</code> | A list of one vector for each named covariate on misclassification probabilities. The vector indicates which covariate effects on misclassification probabilities are constrained to be equal, analogously to <code>constraint</code> . Only used if the model is specified using <code>ematrix</code> . |
| <code>hcovariates</code> | List of formulae the same length as <code>hmodel</code> , defining any covariates governing the hidden Markov outcome models. The covariates operate on a suitably link-transformed linear scale, for example, log scale for a Poisson outcome model. If there are no covariates for a certain hidden state, then insert a <code>NULL</code> in the corresponding place in the list. For example, <code>hcovariates = list(~acute + age, ~acute, NULL)</code> . |
| <code>hcovinits</code> | Initial values for the hidden Markov model covariate effects. A list of the same length as <code>hcovariates</code> . Each element is a vector with initial values for the effect of each covariate on that state. For example, the above <code>hcovariates</code> can be initialised with <code>hcovariates = list(c(-8, 0), -8, NULL)</code> . Initial values must be given for all or no covariates, if none are given these are all set to zero. The initial value given in the <code>hmodel</code> constructor function for the corresponding baseline parameter is interpreted as the value of that parameter with any covariates fixed to their means in the data. If multiple effects are constrained to be equal using <code>hconstraint</code> , then the initial value is taken from the first of the multiple initial values supplied. |
| <code>hconstraint</code> | A named list. Each element is a vector of constraints on the named hidden Markov model parameter. The vector has length equal to the number of times that class of parameter appears in the whole model.

For example consider the three-state hidden Markov model described above, with normally-distributed outcomes for states 1 and 2. To constrain the outcome |

	variance to be equal for states 1 and 2, and to also constrain the effect of acute on the outcome mean to be equal for states 1 and 2, specify <code>hconstraint = list(sd = c(1,1), acute=c(1,1))</code>
<code>qconstraint</code>	A vector of indicators specifying which baseline transition intensities are equal. For example, <code>qconstraint = c(1,2,3,3)</code> constrains the third and fourth intensities to be equal, in a model with four allowed instantaneous transitions. When there are covariates on the intensities and <code>center=TRUE</code> (the default), <code>qconstraint</code> is applied to the intensities with covariates taking the values of the means in the data. When <code>center=FALSE</code> , <code>qconstraint</code> is applied to the intensities with covariates set to zero.
<code>econstraint</code>	A similar vector of indicators specifying which baseline misclassification probabilities are constrained to be equal. Only used if the model is specified using <code>ematrix</code> , rather than <code>hmodel</code> .
<code>initprobs</code>	Currently only used in hidden Markov models. Vector of assumed underlying state occupancy probabilities at each individual's first observation. If these are estimated (see <code>est.initprobs</code>), then this defaults to equal probability for each state. Otherwise this defaults to <code>c(1, rep(0, nstates-1))</code> , that is, in state 1 with a probability of 1. Scaled to sum to 1 if necessary. The state 1 occupancy probability should be non-zero.
<code>est.initprobs</code>	Only used in hidden Markov models. If <code>TRUE</code> , then the underlying state occupancy probabilities at the first observation will be estimated, starting from initial values taken from the <code>initprobs</code> argument. Structural zeroes are allowed: if any of these initial values are zero they will be fixed at zero during optimisation, and no covariate effects on them are estimated. The exception is state 1, which should have non-zero occupancy probability. Be warned that if any of these initial values are 1, then <code>optim</code> will give a "non-finite value" error, since these are logit-transformed during estimation. If you wish to fix any initial probabilities to 1, then that implies all the others are fixed to zero, and there is no need to use <code>est.initprobs</code> at all. Note that the free parameters during this estimation exclude the state 1 occupancy probability, which is fixed at one minus the sum of the other probabilities.
<code>initcovariates</code>	Formula representing covariates on the initial state occupancy probabilities, via multinomial logistic regression. The linear effects of these covariates, observed at the individual's first observation time, operate on the log ratio of the state r occupancy probability to the state 1 occupancy probability, for each $r = 2$ to the number of states. Thus the state 1 occupancy probability should be non-zero. If <code>est.initprobs</code> is <code>TRUE</code> , these effects are estimated starting from their initial values. If <code>est.initprobs</code> is <code>FALSE</code> , these effects are fixed at their initial values.
<code>initcovinits</code>	Initial values for the covariate effects <code>initcovariates</code> . A named list with each element corresponding to a covariate, as in <code>covinits</code> . Each element is a vector with $(1 - \text{number of states})$ elements, containing the initial values for the linear effect of that covariate on the log odds of that state relative to state 1, from state 2 to the final state. If <code>initcovinits</code> is not specified, all covariate effects are initialised to zero.

- death** Vector of indices of absorbing states whose time of entry is known exactly, but the individual is assumed to be in an unknown transient state ("alive") at the previous instant. This is the usual situation for times of death in chronic disease monitoring data. For example, if you specify `death = c(4, 5)` then states 4 and 5 are assumed to be exactly-observed death states.
- See the `obstype` argument. States of this kind correspond to `obstype=3`. `death = TRUE` indicates that the final absorbing state is of this kind, and `death = FALSE` (the default) indicates that there is no state of this kind.
- The `death` argument is overridden by `obstype` or `exacttimes`. Note that times of death do not necessarily have `obstype=3`. If the state is known between the time of death and the previous observation, then you should specify `obstype=2` for the death times, or `exacttimes=TRUE` if the state is known at all times, and the `death` argument is ignored.
- censor** A state, or vector of states, which indicates censoring. Censoring means that the observed state is known only to be one of a particular set of states. For example, `censor=999` indicates that all observations of 999 in the vector of observed states are censored states. By default, this means that the true state could have been anything other than an absorbing state. To specify corresponding true states explicitly, use a `censor.states` argument.
- Note that in contrast to the usual terminology of survival analysis, here it is the *state* which is considered to be censored, rather than the *event time*. If at the end of a study, an individual has not died, but their true state is *known*, then `censor` is unnecessary, since the standard multi-state model likelihood is applicable.
- Note in particular that general time-inhomogeneous Markov models with piecewise constant transition intensities can be constructed using the `censor` facility. If the true state is unknown on occasions when a piecewise constant covariate is known to change, then censored states can be inserted in the data on those occasions. The covariate may represent time itself, in which case the `pci` option to `msm` can be used to perform this trick automatically, or some other time-dependent variable.
- censor.states** Specifies the underlying states which censored observations can represent. If `censor` is a single number (the default) this can be a vector, or a list with one element. If `censor` is a vector with more than one element, this should be a list, with each element a vector corresponding to the equivalent element of `censor`. For example
- ```
censor = c(99, 999), censor.states = list(c(2,3), c(3,4))
```
- means that observations coded 99 represent either state 2 or state 3, while observations coded 999 are really either state 3 or state 4.
- pci** Model for piecewise-constant intensities. Vector of cut points defining the times, since the start of the process, at which intensities change. For example
- ```
pci = c(5, 10)
```
- specifies that the intensity changes at time points 5 and 10. This will automatically construct a model with a categorical (factor) covariate called `timeperiod`, with levels "`[-Inf, 5)`", "`[5, 10)`" and "`[10, Inf)`", where the first level is the baseline. This covariate defines the time period in which the observation was made. Initial values and constraints on covariate effects are specified the same way as for a model with a covariate of this name, for example,

```
covinits = list("timeperiod[5,10]"=c(0.1,0.1), "timeperiod[10,Inf]"=
```

Internally, this works by inserting censored observations in the data at times when the intensity changes but the state is not observed.

If the supplied times are outside the range of the time variable in the data, `pci` is ignored and a time-homogeneous model is fitted.

After fitting a time-inhomogeneous model, `qmatrix.msm` can be used to obtain the fitted intensity matrices for each time period, for example,

```
qmatrix.msm(example.msm, covariates=list(timeperiod="[5,Inf)"))
```

This facility does not support interactions between time and other covariates. Such models need to be specified "by hand", using a state variable with censored observations inserted. Note that the data component of the `msm` object returned from a call to `msm` with `pci` supplied contains the states with inserted censored observations and time period indicators. These can be used to construct such models.

<code>exacttimes</code>	By default, the transitions of the Markov process are assumed to take place at unknown occasions in between the observation times. If <code>exacttimes</code> is set to <code>TRUE</code> , then all observation times are assumed to represent the exact and complete times of transition of the process. This is equivalent to every row of the data having <code>obstype = 2</code> . See the <code>obstype</code> argument. If both <code>obstype</code> and <code>exacttimes</code> are specified then <code>exacttimes</code> is ignored.
<code>c1</code>	Width of symmetric confidence intervals for maximum likelihood estimates, by default 0.95.
<code>fixedpars</code>	<p>Vector of indices of parameters whose values will be fixed at their initial values during the optimisation. These are given in the order: transition intensities (reading across rows of the transition matrix), covariates on intensities (ordered by intensities within covariates), hidden Markov model parameters, including misclassification probabilities (ordered by parameters within states), hidden Markov model covariate parameters (ordered by covariates within parameters within states), initial state occupancy probabilities (excluding the first probability, which is fixed at one minus the sum of the others).</p> <p>If there are equality constraints on certain parameters, then <code>fixedpars</code> indexes the set of unique parameters, excluding those which are constrained to be equal to previous parameters.</p> <p>For covariates on misclassification probabilities, this is a change from version 0.4 in the parameter ordering. Previously these were ordered by misclassification probabilities within covariates.</p> <p>This can be useful for profiling likelihoods, and building complex models stage by stage. To fix all parameters, specify <code>fixedpars = TRUE</code>.</p>
<code>center</code>	If <code>TRUE</code> (the default, unless <code>fixedpars=TRUE</code>) then covariates are centered at their means during the maximum likelihood estimation. This usually improves stability of the numerical optimisation.
<code>opt.method</code>	Quoted name of the R function to perform minimisation of the minus twice log likelihood. Either "optim" or "nlm". <code>optim</code> is the default.
<code>hessian</code>	If <code>TRUE</code> (the default) then the Hessian matrix is computed at the maximum likelihood estimates, to obtain standard errors and confidence intervals.

<code>use.deriv</code>	If TRUE then analytic first derivatives are used in the optimisation of the likelihood, when an appropriate quasi-Newton optimisation method, such as BFGS, is being used. Note that the default for <code>optim</code> is a Nelder-Mead method which cannot use derivatives. However, these derivatives, if supplied, are always used to calculate the Hessian.
<code>analyticp</code>	By default, the likelihood for certain simpler 3, 4 and 5 state models is calculated using an analytic expression for the transition probability (P) matrix. To revert to the original method of using the matrix exponential, specify <code>analyticp=FALSE</code> . See the PDF manual for a list of the models for which analytic P matrices are implemented.
<code>...</code>	Optional arguments to the general-purpose R optimisation routines, <code>optim</code> or <code>nlm</code> . Useful options for <code>optim</code> include <code>method="BFGS"</code> for using a quasi-Newton optimisation algorithm, which can often be faster than the default Nelder-Mead. If the optimisation fails to converge, consider normalising the problem using, for example, <code>control=list(fnscale = 2500)</code> , for example, replacing 2500 by a number of the order of magnitude of the likelihood. If 'false' convergence is reported and the standard errors cannot be calculated due to a non-positive-definite Hessian, then consider tightening the tolerance criteria for convergence. If the optimisation takes a long time, intermediate steps can be printed using the <code>trace</code> argument of the control list. See <code>optim</code> for details.

Details

For full details about the methodology behind the **msm** package, refer to the PDF manual ‘`msm-manual.pdf`’ in the ‘`doc`’ subdirectory of the package. This includes a tutorial in the typical use of **msm**. The paper by Jackson (2011) in *Journal of Statistical Software* presents the material in this manual in a more concise form.

msm was designed for fitting *continuous-time* Markov models, processes where transitions can occur at any time. These models are defined by *intensities*, which govern both the time spent in the current state and the probabilities of the next state. In *discrete-time models*, transitions are known in advance to only occur at multiples of some time unit, and the model is purely governed by the probability distributions of the state at the next time point, conditionally on the state at the current time. These can also be fitted in **msm**, assuming that there is a continuous-time process underlying the data. Then the fitted transition probability matrix over one time period, as returned by `pmatrix.msm(..., t=1)` is equivalent to the matrix that governs the discrete-time model. However, these can be fitted more efficiently using multinomial logistic regression, for example, using `multinom` from the R package **nnet** (Venables and Ripley, 2002).

For simple continuous-time multi-state Markov models, the likelihood is calculated in terms of the transition intensity matrix Q . When the data consist of observations of the Markov process at arbitrary times, the exact transition times are not known. Then the likelihood is calculated using the transition probability matrix $P(t) = \exp(tQ)$, where `exp` is the matrix exponential. If state i is observed at time t and state j is observed at time u , then the contribution to the likelihood from this pair of observations is the i, j element of $P(u - t)$. See, for example, Kalbfleisch and Lawless (1985), Kay (1986), or Gentleman *et al.* (1994).

For hidden Markov models, the likelihood for an individual with k observations is calculated directly by summing over the unknown state at each time, producing a product of k matrices. The calculation is a generalisation of the method described by Satten and Longini (1996), and also by Jackson and Sharples (2002), and Jackson *et al.* (2003).

There must be enough information in the data on each state to estimate each transition rate, otherwise the likelihood will be flat and the maximum will not be found. It may be appropriate to reduce the number of states in the model, the number of allowed transitions, or the number of covariate effects, to ensure convergence. Hidden Markov models, and situations where the value of the process is only known at a series of snapshots, are particularly susceptible to non-identifiability, especially when combined with a complex transition matrix.

Choosing an appropriate set of initial values for the optimisation can also be important. For flat likelihoods, 'informative' initial values will often be required.

Value

To obtain summary information from models fitted by the `msm` function, it is recommended to use extractor functions such as `qmatrix.msm`, `pmatrix.msm`, `sojourn.msm`. These provide estimates and confidence intervals for quantities such as transition probabilities for given covariate values.

For advanced use, it may be necessary to directly use information stored in the object returned by the `msm` function. This is a list of class `msm`, with components:

<code>call</code>	The original call to <code>msm</code> .
<code>Qmatrices</code>	A list of matrices. The first component, labelled <code>logbaseline</code> , is a matrix containing the estimated transition intensities on the log scale with any covariates fixed at their means in the data (or at zero, if <code>center=FALSE</code>). The component labelled <code>baseline</code> is the equivalent on the untransformed scale. Each remaining component is a matrix giving the linear effects of the labelled covariate on the matrix of log intensities. To extract an estimated intensity matrix on the natural scale, at an arbitrary combination of covariate values, use the function <code>qmatrix.msm</code> .
<code>QmatricesSE</code>	The standard error matrices corresponding to <code>Qmatrices</code> .
<code>QmatricesL, QmatricesU</code>	Corresponding lower and upper symmetric confidence limits, of width 0.95 unless specified otherwise by the <code>cl</code> argument.
<code>Ematrices</code>	A list of matrices. The first component, labelled <code>logitbaseline</code> , is the estimated misclassification probability matrix (expressed as log odds relative to the probability of the true state) with any covariates fixed at their means in the data (or at zero, if <code>center=FALSE</code>). The component labelled <code>baseline</code> is the equivalent on the untransformed scale. Each remaining component is a matrix giving the linear effects of the labelled covariate on the matrix of logit misclassification probabilities. To extract an estimated misclassification probability matrix on the natural scale, at an arbitrary combination of covariate values, use the function <code>ematrix.msm</code> .
<code>EmatricesSE</code>	The standard error matrices corresponding to <code>Ematrices</code> .
<code>EmatricesL, EmatricesU</code>	Corresponding lower and upper symmetric confidence limits, of width 0.95 unless specified otherwise by the <code>cl</code> argument.
<code>sojourn</code>	A list with components: <code>mean</code> = estimated mean sojourn times in the transient states, with covariates fixed at their means (if <code>center=TRUE</code>) or at zero (if <code>center=FALSE</code>). <code>se</code> = corresponding standard errors.

<code>minus2loglik</code>	Minus twice the maximised log-likelihood.
<code>deriv</code>	Derivatives of the minus twice log-likelihood at its maximum.
<code>estimates</code>	Vector of untransformed maximum likelihood estimates returned from <code>optim</code> . Transition intensities are on the log scale and misclassification probabilities are given as log odds relative to the probability of the true state.
<code>estimates.t</code>	Vector of transformed maximum likelihood estimates with intensities and probabilities on their natural scales.
<code>fixedpars</code>	Indices of <code>estimates</code> which were fixed during the maximum likelihood estimation.
<code>center</code>	Indicator for whether the estimation was performed with covariates centered on their means in the data.
<code>covmat</code>	Covariance matrix corresponding to <code>estimates</code> .
<code>ci</code>	Matrix of confidence intervals corresponding to <code>estimates.t</code>
<code>opt</code>	Return value from <code>optim</code> or <code>nlm</code> , giving information about the results of the optimisation.
<code>foundse</code>	Logical value indicating whether the Hessian was positive-definite at the supposed maximum of the likelihood. If not, the covariance matrix of the parameters is unavailable. In these cases the optimisation has probably not converged to a maximum.
<code>data</code>	A list of constants and vectors giving the data, for use in post-processing.
<code>qmodel</code>	A list of objects specifying the model for transition intensities, for use in post-processing.
<code>emodel</code>	A list of objects specifying the model for misclassification.
<code>qcmmodel</code>	A list of objects specifying the model for covariates on the transition intensities.
<code>ecmodel</code>	A list of objects specifying the model for covariates on misclassification probabilities.
<code>hmodel</code>	A list of class "hmodel", containing objects specifying the hidden Markov model. Estimates of "baseline" location parameters are presented with any covariates fixed to their means in the data, or if <code>center=FALSE</code> , with covariates fixed to zero.
<code>cmodel</code>	A list of objects specifying any model for censoring.

Printing a `msm` object by typing the object's name at the command line implicitly invokes `print.msm`. This formats and prints the important information in the model fit. This includes the fitted transition intensity matrix, matrices containing covariate effects on intensities, and mean sojourn times from a fitted `msm` model. When there is a hidden Markov model, the chief information in the `hmodel` component is also formatted and printed. This includes estimates and confidence intervals for each parameter.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

- Jackson, C. H. (2011). Multi-State Models for Panel Data: The msm Package for R., *Journal of Statistical Software*, 38(8), 1-29. URL <http://www.jstatsoft.org/v38/i08/>.
- Kalbfleisch, J., Lawless, J.F., The analysis of panel data under a Markov assumption *Journal of the Americal Statistical Association* (1985) 80(392): 863–871.
- Kay, R. A Markov model for analysing cancer markers and disease states in survival studies. *Biometrics* (1986) 42: 855–865.
- Gentleman, R.C., Lawless, J.F., Lindsey, J.C. and Yan, P. Multi-state Markov models for analysing incomplete disease history data with illustrations for HIV disease. *Statistics in Medicine* (1994) 13(3): 805–821.
- Satten, G.A. and Longini, I.M. Markov chains with measurement error: estimating the 'true' course of a marker of the progression of human immunodeficiency virus disease (with discussion) *Applied Statistics* 45(3): 275-309 (1996)
- Jackson, C.H. and Sharples, L.D. Hidden Markov models for the onset and progression of bronchiolitis obliterans syndrome in lung transplant recipients *Statistics in Medicine*, 21(1): 113–128 (2002).
- Jackson, C.H., Sharples, L.D., Thompson, S.G. and Duffy, S.W. and Couto, E. Multi-state Markov models for disease progression with classification error. *The Statistician*, 52(2): 193–209 (2003)
- Venables, W.N. and Ripley, B.D. (2002) *Modern Applied Statistics with S*, second edition. Springer.

See Also

[simmulti.msm](#), [plot.msm](#), [summary.msm](#), [qmatrix.msm](#), [pmatrix.msm](#), [sojourn.msm](#).

Examples

```
### Heart transplant data
### For further details and background to this example, see
### the PDF manual in the doc directory.
data(cav)
print(cav[1:10,])
twoway4.q <- rbind(c(-0.5, 0.25, 0, 0.25), c(0.166, -0.498, 0.166, 0.166),
c(0, 0.25, -0.5, 0.25), c(0, 0, 0, 0))
statetable.msm(state, PTNUM, data=cav)
crudeinits.msm(state ~ years, PTNUM, data=cav, qmatrix=twoway4.q)
cav.msm <- msm( state ~ years, subject=PTNUM, data = cav,
               qmatrix = twoway4.q, death = 4,
               control = list ( trace = 2, REPORT = 1 ) )

cav.msm
qmatrix.msm(cav.msm)
pmatrix.msm(cav.msm, t=10)
sojourn.msm(cav.msm)
```

msm.summary	<i>Summarise a fitted multi-state model</i>
-------------	---

Description

Summary method for fitted [msm](#) models. Currently, this produces a table of observed and expected state prevalences for each time. For models with covariates, prints hazard ratios with confidence intervals for covariate effects.

Usage

```
## S3 method for class 'msm'
summary(object, times=NULL, timezero=NULL, initstates=NULL,
covariates="mean", misccovariates="mean", hazard.scale=1, ...)
```

Arguments

object	A fitted multi-state model object, as returned by msm .
times	A sequence of times at which to compare observed and expected prevalences of each state. Defaults to <code>seq(min(times), max(times), (max(times) - min(times))/10)</code> .
timezero	Initial time of the Markov process. Expected values are forecasted from here. Defaults to the minimum of the observation times given in the data.
initstates	Optional vector of the same length as the number of states. Gives the numbers of individuals occupying each state at <code>timezero</code> . The default is that all individuals are in state 1.
covariates	Covariate values for which to forecast expected state occupancy. See qmatrix.msm . Defaults to the mean values of the covariates in the data set.
misccovariates	(Misclassification models only) Values of covariates on the misclassification probability matrix for which to forecast expected state occupancy. Defaults to the mean values of the covariates in the data set.
hazard.scale	Vector with same elements as number of covariates on transition rates. Corresponds to the increase in each covariate used to calculate its hazard ratio. Defaults to all 1.
...	further arguments passed to or from other methods.

Value

A list of class `summary.msm`, with components:

prevalences	Output from prevalence.msm .
hazard	Output from hazard.msm
.	.
hazard.scale	Value of the <code>hazard.scale</code> argument

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#), [prevalence.msm](#), [hazard.msm](#)

odds.msm	<i>Calculate tables of odds ratios for covariates on misclassification probabilities</i>
----------	--

Description

Odds ratios are computed by exponentiating the estimated covariate effects on the logit-misclassification probabilities.

Usage

```
odds.msm(x, odds.scale = 1, cl = 0.95)
```

Arguments

x	Output from msm representing a fitted multi-state model.
odds.scale	Vector with same elements as number of covariates on misclassification probabilities. Corresponds to the increase in each covariate used to calculate its odds ratio. Defaults to all 1.
cl	Width of the symmetric confidence interval to present. Defaults to 0.95.

Value

A list of tables containing odds ratio estimates, one table for each covariate. Each table has three columns, containing the odds ratio, and an approximate upper 95% and lower 95% confidence limit respectively (assuming normality on the log scale), for each misclassification probability.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#), [hazard.msm](#)

pearson.msm	<i>Pearson-type goodness-of-fit test</i>
-------------	--

Description

Pearson-type goodness-of-fit test for multi-state models fitted to panel-observed data.

Usage

```
pearson.msm(x, transitions=NULL, timegroups=3, intervalgroups=3,
            covgroups=3, groups=NULL, boot=FALSE, B=500,
            next.obstime=NULL, N=100, indep.cens=TRUE, maxtimes=NULL)
```

Arguments

x	A fitted multi-state model, as returned by <code>msm</code> .
transitions	<p>This should be an integer vector indicating which interval transitions should be grouped together in the contingency table. Its length should be the number of allowed interval transitions, excluding transitions from absorbing states to absorbing states.</p> <p>The allowed interval transitions are the set of pairs of states (a, b) for which it is possible to observe a at one time and b at any later time. For example, in a "well-disease-death" model with allowed <i>instantaneous</i> 1-2, 2-3 transitions, there are 5 allowed <i>interval</i> transitions. In numerical order, these are 1-1, 1-2, 1-3, 2-2 and 2-3, excluding absorbing-absorbing transitions.</p> <p>Then, to group transitions 1-1,1-2 together, and transitions 2-2,2-3 together, specify</p> <pre>transitions = c(1,1,2,3,3).</pre> <p>Only transitions from the same state may be grouped. By default, each interval transition forms a separate group.</p>
timegroups	Number of groups based on quantiles of the time since the start of the process.
intervalgroups	Number of groups based on quantiles of the time interval between observations.
covgroups	<p>Number of groups based on quantiles of $\sum_r q_{irr}$, where q_{irr} are the diagonal entries of the transition intensity matrix for the ith transition. These are a function of the covariate effects and the covariate values at the ith transition: q_{irr} is minus the sum of the off-diagonal entries $q_{rs}^{(0)} \exp(\beta_{rs}^T z_i)$ on the rth row.</p> <p>Thus <code>covgroups</code> summarises the impact of covariates at each observation, by calculating the overall rate of progression through states at that observation.</p> <p>For time-inhomogeneous models specified using the <code>pci</code> argument to <code>msm</code>, if the only covariate is the time period, <code>covgroups</code> is set to 1, since <code>timegroups</code> ensures that transitions are grouped by time.</p>
groups	A vector of arbitrary groups in which to categorise each transition. This can be an integer vector or a factor. This can be used to diagnose specific areas of poor fit. For example, the contingency table might be grouped by arbitrary

combinations of covariates to detect types of individual for whom the model fits poorly.

The length of groups should be `x$data$n`, the number of observations used in the model fit, which is the number of observations in the original dataset with any missing values excluded. The value of groups at observation i is used to categorise the transition which *ends* at observation i . Values of groups at the first observation for each subject are ignored.

boot	Estimate an "exact" p-value using a parametric bootstrap. All objects used in the original call to <code>msm</code> which produced <code>x</code> , such as the <code>qmatrix</code> , should be in the working environment, or else an "object not found" error will be given. This enables the original model to be refitted to the replicate datasets.
B	Number of bootstrap replicates.
next.obstime	This is a vector of length <code>x\$data\$n</code> (the number of observations used in the model fit) giving the time to the next <i>scheduled</i> observation following each time point. This is only used when times to death are known exactly. For individuals who died (entered an absorbing state) before the next scheduled observation, and the time of death is known exactly, <code>next.obstime</code> would be <i>greater</i> than the observed death time. If the individual did not die, and a scheduled observation did follow that time point, <code>next.obstime</code> should just be the same as the time to that observation. <code>next.obstime</code> is used to determine a grouping of the time interval between observations, which should be based on scheduled observations. If exact times to death were used in the grouping, then shorter intervals would contain excess deaths, and the goodness-of-fit statistic would be biased. If <code>next.obstime</code> is unknown, it is multiply-imputed using a product-limit estimate based on the intervals to observations other than deaths. The resulting tables of transitions are averaged over these imputations. This may be slow.
N	Number of imputations for the estimation of the distribution of the next scheduled observation time, when there are exact death times.
indep.cens	If TRUE, then times to censoring are included in the estimation of the distribution to the next scheduled observation time. If FALSE, times to censoring are assumed to be systematically different from other observation times.
maxtimes	A vector of length <code>x\$data\$n</code> , or a common scalar, giving an upper bound for the next scheduled observation time. Used in the multiple imputation when times to death are known exactly. If a value greater than <code>maxtimes</code> is simulated, then the next scheduled observation is taken as censored. This should be supplied, if known. If not supplied, this is taken to be the maximum interval occurring in the data, plus one time unit. For observations which are not exact death times, this should be the time since the previous observation.

Details

This method (Aguirre-Hernandez and Farewell, 2002) is intended for data which represent observations of the process at arbitrary times ("snapshots", or "panel-observed" data). For data which represent the exact transition times of the process, `prevalence.msm` can be used to assess fit, though without a formal test.

When times of death are known exactly, states are misclassified, or an individual's final observation is a censored state, the modification by Titman and Sharples (2008) is used. The only form of censoring supported is a state at the end of an individual's series which represents an unknown transient state (i.e. the individual is only known to be alive at this time). Other types of censoring are omitted from the data before performing the test.

See these references for further details of the methods. The method used for censored states is a modification of the method in the appendix to Titman and Sharples (2008), described at <http://www.mrc-bsu.cam.ac.uk/personal/chris/papers/robustcensoring.pdf> (Titman, 2007).

Groupings of the time since initiation, the time interval and the impact of covariates are based on equally-spaced quantiles. The number of groups should be chosen that there are not many cells with small expected numbers of transitions, since the deviance statistic will be unstable for sparse contingency tables. Ideally, the expected numbers of transitions in each cell of the table should be no less than about 5. Conversely, the power of the test is reduced if there are too few groups. Therefore, some sensitivity analysis of the test results to the grouping is advisable. To ensure that the null distribution of the statistic is approximated reasonably by a chi-squared distribution with `df.lower` degrees of freedom, the table should always be grouped by any fitted covariates (Aguirre-Hernandez and Farewell, 2002).

Saved model objects fitted with previous versions of R (versions less than 0.8) will need to be refitted under the current R for use with `pearson.msm`.

Value

A list whose first two elements are contingency tables of observed transitions O and expected transitions E , respectively, for each combination of groups. The third element is a table of the deviances $(O - E)^2/E$ multiplied by the sign of $O - E$. If the expected number of transitions is zero then the deviance is zero. Entries in the third matrix will be bigger in magnitude for groups for which the model fits poorly.

- "test" , the fourth element of the list, is a data frame with one row containing the Pearson-type goodness-of-fit test statistic `stat`. The test statistic is the sum of the deviances. For panel-observed data without exact death times, `p.lower` and `p.upper` are theoretical lower and upper limits for the unknown p-value of the test statistic, based on chi-squared distributions with `df.lower` and `df.upper` degrees of freedom, respectively. `df.upper` is the number of independent cells in the contingency table, and `df.lower` is `df.upper` minus the number of estimated parameters in the model.
- "intervalq" (not printed by default) contains the definition of the grouping of the intervals between observations. These groups are defined by quantiles within the groups corresponding to the time since the start of the process.
- "sim" If there are exact death times, this contains simulations of the contingency tables and test statistics for each imputation of the next scheduled sampling time. These are averaged over to produce the presented tables and test statistic. This element is not printed by default.
- With exact death times, the null variance of the test statistic (formed by taking mean of simulated test statistics) is less than twice the mean (Titman, 2008), and the null distribution is not chi-squared. In this case, `p.upper` is an upper

limit for the true asymptotic p-value, but `p.lower` is not a lower limit, and is not presented.

"boot" If the bootstrap has been used, the element will contain the bootstrap replicates of the test statistics (not printed by default).

Author(s)

Andrew Titman (<a.titman@lancaster.ac.uk>), Chris Jackson (<chris.jackson@mrc-bsu.cam.ac.uk>)

References

Aguirre-Hernandez, R. and Farewell, V. (2002) A Pearson-type goodness-of-fit test for stationary and time-continuous Markov regression models. *Statistics in Medicine* 21:1899-1911.

Titman, A. and Sharples, L. (2008) A general goodness-of-fit test for Markov and hidden Markov models. *Statistics in Medicine* 27(12):2177-2195

Titman, A. (2008) Model diagnostics in multi-state models of biological systems. PhD thesis, University of Cambridge.

See Also

[msm](#), [prevalence.msm](#), [scorereresid.msm](#),

Examples

```
data(psor)
psor.q <- rbind(c(0,0.1,0,0),c(0,0,0.1,0),c(0,0,0,0.1),c(0,0,0,0))
psor.msm <- msm(state ~ months, subject=ptnum, data=psor,
               qmatrix = psor.q, covariates = ~ollwsdrt+hieffusn,
               constraint = list(hieffusn=c(1,1,1),ollwsdrt=c(1,1,2)),
               control = list(REPORT=1,trace=2), method="BFGS")
pearson.msm(psor.msm, timegroups=2, intervalgroups=2, covgroups=2)
# More 1-2, 1-3 and 1-4 observations than expected in shorter time
# intervals - the model fits poorly.
# A random effects model might accommodate such fast progressors.
```

pexp

Exponential distribution with piecewise-constant rate

Description

Density, distribution function, quantile function and random generation for a generalisation of the exponential distribution, in which the rate changes at a series of times.

Usage

```
dpexp(x, rate=1, t=0, log = FALSE)
ppexp(q, rate=1, t=0, lower.tail = TRUE, log.p = FALSE)
qpexp(p, rate=1, t=0, lower.tail = TRUE, log.p = FALSE)
rpep(n, rate=1, t=0)
```

Arguments

x, q	vector of quantiles.
p	vector of probabilities.
n	number of observations. If <code>length(n) > 1</code> , the length is taken to be the number required.
rate	vector of rates.
t	vector of the same length as <code>rate</code> , giving the times at which the rate changes. The first element of <code>t</code> should be 0, and <code>t</code> should be in increasing order.
log, log.p	logical; if TRUE, probabilities <code>p</code> are given as <code>log(p)</code> .
lower.tail	logical; if TRUE (default), probabilities are $P[X \leq x]$, otherwise, $P[X > x]$.

Details

Consider the exponential distribution with rates r_1, \dots, r_n changing at times t_1, \dots, t_n , with $t_1 = 0$. Suppose t_k is the maximum t_i such that $t_i < x$. The density of this distribution at $x > 0$ is $f(x)$ for $k = 1$, and

$$\prod_{i=1}^k (1 - F(t_i - t_{i-1}, r_i)) f(x - t_k, r_k)$$

for $k > 1$.

where $F()$ and $f()$ are the distribution and density functions of the standard exponential distribution.

If `rate` is of length 1, this is just the standard exponential distribution. Therefore, for example, `dpexp(x)`, with no other arguments, is simply equivalent to `dexp(x)`.

Only `rpexp` is used in the `msm` package, to simulate from Markov processes with piecewise-constant intensities depending on time-dependent covariates. These functions are merely provided for completion, and are not optimized for numerical stability or speed.

Value

`dpexp` gives the density, `ppexp` gives the distribution function, `qpexp` gives the quantile function, and `rpexp` generates random deviates.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[dexp](#), [sim.msm](#).

Examples

```
x <- seq(0.1, 50, by=0.1)
rate <- c(0.1, 0.2, 0.05, 0.3)
t <- c(0, 10, 20, 30)
## standard exponential distribution
plot(x, dexp(x, 0.1), type="l")
```

```
## distribution with piecewise constant rate
lines(x, dpexp(x, rate, t), type="l", lty=2)
## standard exponential distribution
plot(x, pexp(x, 0.1), type="l")
## distribution with piecewise constant rate
lines(x, ppexp(x, rate, t), type="l", lty=2)
```

plot.msm

Plots of multi-state models

Description

This produces a plot of the expected probability of survival against time, from each transient state. Survival is defined as not entering an absorbing state.

Usage

```
## S3 method for class 'msm'
plot(x, from, to, range, covariates, legend.pos, xlab="Time",
      ylab="Fitted survival probability", lwd=1, ...)
```

Arguments

x	Output from msm , representing a fitted multi-state model object.
from	States from which to consider survival. Defaults to the complete set of transient states.
to	Absorbing state to consider. Defaults to the highest-labelled absorbing state.
range	Vector of two elements, giving the range of times to plot for.
covariates	Covariate values for which to evaluate the expected probabilities. This can either be: <ul style="list-style-type: none"> the string "mean", denoting the means of the covariates in the data (this is the default), the number 0, indicating that all the covariates should be set to zero, or a list of values, with optional names. For example <pre>list (60, 1)</pre> where the order of the list follows the order of the covariates originally given in the model formula, or a named list, <pre>list (age = 60, sex = 1)</pre>
legend.pos	Vector of the x and y position, respectively, of the legend.
xlab	x axis label.
ylab	y axis label.
lwd	Line width. See par .
...	Other arguments to be passed to the generic plot and lines functions.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#)

plot.prevalence.msm *Plot of observed and expected prevalences*

Description

Provides a rough indication of goodness of fit of a multi-state model, by estimating the observed numbers of individuals occupying a state at a series of times, and plotting these against forecasts from the fitted model, for each state. Observed prevalences are indicated as solid lines, expected prevalences as dashed lines.

Usage

```
## S3 method for class 'prevalence.msm'
plot(x, mintime=NULL, maxtime=NULL, timezero=NULL,
     initstates=NULL, interp=c("start","midpoint"),
     covariates="mean", miscovariates="mean",
     piecewise.times=NULL, piecewise.covariates=NULL,
     xlab="Times",ylab="Prevalence (%)", lwd.obs=1,
     lwd.exp=1, lty.obs=1, lty.exp=2, col.obs="blue",
     col.exp="red", legend.pos=NULL,
     ...)
```

Arguments

x	A fitted multi-state model produced by msm .
mintime	Minimum time at which to compute the observed and expected prevalences of states.
maxtime	Maximum time at which to compute the observed and expected prevalences of states.
timezero	Initial time of the Markov process. Expected values are forecasted from here. Defaults to the minimum of the observation times given in the data.
initstates	Optional vector of the same length as the number of states. Gives the numbers of individuals occupying each state at the initial time, to be used for forecasting expected prevalences. The default is those observed in the data. These should add up to the actual number of people in the study at the start.
interp	Interpolation method for observed states, see prevalence.msm
covariates	Covariate values for which to forecast expected state occupancy. See qmatrix.msm . Defaults to the mean values of the covariates in the data set.

misccovariates	(Misclassification models only) Values of covariates on the misclassification probability matrix for which to forecast expected state occupancy. Defaults to the mean values of the covariates in the data set.
piecewise.times	Times at which piecewise-constant intensities change. See pmatrix.piecewise.msm for how to specify this.
piecewise.covariates	Covariates on which the piecewise-constant intensities depend. See pmatrix.piecewise.msm for how to specify this.
xlab	x axis label.
ylab	y axis label.
lwd.obs	Line width for observed prevalences. See par .
lwd.exp	Line width for expected prevalences. See par .
lty.obs	Line type for observed prevalences. See par .
lty.exp	Line type for expected prevalences. See par .
col.obs	Line colour for observed prevalences. See par .
col.exp	Line colour for expected prevalences. See par .
legend.pos	Vector of the x and y position, respectively, of the legend.
...	Further arguments to be passed to the generic plot function.

Details

See [prevalence.msm](#) for details of the assumptions underlying this method.

Observed prevalences are plotted with a solid line, and expected prevalences with a dotted line.

References

Gentleman, R.C., Lawless, J.F., Lindsey, J.C. and Yan, P. Multi-state Markov models for analysing incomplete disease history data with illustrations for HIV disease. *Statistics in Medicine* (1994) 13(3): 805–821.

See Also

[prevalence.msm](#)

plot.survfit.msm

Plot empirical and fitted survival curves

Description

Plot a Kaplan-Meier estimate of the survival probability and compare it with the fitted survival probability from a `msm` model.

Usage

```
## S3 method for class 'survfit.msm'
plot(x, from=1, to=NULL, range=NULL, covariates="mean",
      interp=c("start", "midpoint"), ci=c("none", "normal", "bootstrap"), B=100,
      legend.pos=NULL, xlab="Time", ylab="Survival probability",
      lty=1, lwd=1, col="red", lty.ci=2, lwd.ci=1, col.ci="red",
      mark.time=TRUE, col.surv="blue", lty.surv=2, lwd.surv=1,
      ...)
```

Arguments

x	Output from msm , representing a fitted multi-state model object.
from	State from which to consider survival. Defaults to state 1.
to	Absorbing state to consider. Defaults to the highest-labelled absorbing state.
range	Vector of two elements, giving the range of times to plot for.
covariates	Covariate values for which to evaluate the expected probabilities. This can either be: <p>the string "mean", denoting the means of the covariates in the data (this is the default),</p> <p>the number 0, indicating that all the covariates should be set to zero,</p> <p>or a list of values, with optional names. For example</p> <pre>list (60, 1)</pre> <p>where the order of the list follows the order of the covariates originally given in the model formula, or a named list,</p> <pre>list (age = 60, sex = 1)</pre>
ci	If "none" (the default) no confidence intervals are plotted. If "normal" or "bootstrap", confidence intervals are plotted based on the respective method in pmatrix.msm . This is very computationally-intensive, since intervals must be computed at a series of times.
B	Number of bootstrap or normal replicates for the confidence interval. The default is 100 rather than the usual 1000, since these plots are for rough diagnostic purposes.
interp	If <code>interp="start"</code> (the default) then the entry time into the absorbing state is assumed to be the time it is first observed in the data. If <code>interp="midpoint"</code> then the entry time into the absorbing state is assumed to be halfway between the time it is first observed and the previous observation time. This is generally more reasonable for "progressive" models with observations at arbitrary times.
legend.pos	Vector of the <i>x</i> and <i>y</i> position, respectively, of the legend.
xlab	x axis label.
ylab	y axis label.

lty	Line type for the fitted curve. See par .
lwd	Line width for the fitted curve. See par .
col	Colour for the fitted curve. See par .
lty.ci	Line type for the fitted curve confidence limits. See par .
lwd.ci	Line width for the fitted curve confidence limits. See par .
col.ci	Colour for the fitted curve confidence limits. See par .
mark.time	Mark the empirical survival curve at each censoring point, see lines.survfit .
col.surv	Colour for the empirical survival curve, passed to lines.survfit . See par .
lty.surv	Line type for the empirical survival curve, passed to lines.survfit . See par .
lwd.surv	Line width for the empirical survival curve, passed to lines.survfit . See par .
...	Other arguments to be passed to the plot function which draws the fitted curve, or the lines.survfit function which draws the empirical curve.

Details

If the data represent observations of the process at arbitrary times, then the first occurrence of the absorbing state in the data will usually be greater than the actual first transition time to that state. Therefore the Kaplan-Meier estimate of the survival probability will be an overestimate.

This currently only handles time-homogeneous models.

See Also

[survfit](#), [plot.survfit](#), [plot.prevalence.msm](#)

plotprog.msm

Kaplan Meier estimates of incidence

Description

Compute and plot Kaplan-Meier estimates of the probability that each successive state has not occurred yet.

Usage

```
plotprog.msm(formula, subject, data, legend.pos=NULL, xlab="Time",
             ylab="1 - incidence probability", lwd=1, xlim=NULL,
             mark.time=TRUE, ...)
```

Arguments

formula	A formula giving the vectors containing the observed states and the corresponding observation times. For example, state ~ time Observed states should be in the set 1, . . . , n, where n is the number of states.
subject	Vector of subject identification numbers for the data specified by formula. If missing, then all observations are assumed to be on the same subject. These must be sorted so that all observations on the same subject are adjacent.
data	An optional data frame in which the variables represented by state, time and subject can be found.
legend.pos	Vector of the <i>x</i> and <i>y</i> position, respectively, of the legend.
xlab	x axis label.
ylab	y axis label.
lwd	Line width. See par .
xlim	x axis limits, e.g. c(0,10) for an axis ranging from 0 to 10. Default is the range of observation times.
mark.time	Mark the empirical survival curve at each censoring point, see lines.survfit .
...	Other arguments to be passed to the plot and lines.survfit functions.

Details

If the data represent observations of the process at arbitrary times, then the first occurrence of the state in the data will usually be greater than the actual first transition time to that state. Therefore the probabilities plotted by [plotprog.msm](#) will be overestimates.

See Also

[survfit](#), [plot.survfit](#)

pmatrix.msm

Transition probability matrix

Description

Extract the estimated transition probability matrix from a fitted multi-state model for a given time interval, at a given set of covariate values.

Usage

```
pmatrix.msm(x, t=1, t1=0, covariates="mean",
            ci=c("none", "normal", "bootstrap"), c1=0.95, B=1000,
            ...)
```

Arguments

x	A fitted multi-state model, as returned by msm .
t	The time interval to estimate the transition probabilities for, by default one unit.
t1	The starting time of the interval. Used for models x with piecewise-constant intensities fitted using the <code>pci</code> option to msm . The probabilities will be computed on the interval $[t1, t1+t]$.
covariates	The covariate values at which to estimate the transition probabilities. This can either be: <p>the string "mean", denoting the means of the covariates in the data (this is the default),</p> <p>the number 0, indicating that all the covariates should be set to zero,</p> <p>or a list of values, with optional names. For example</p> <pre>list (60, 1)</pre> <p>where the order of the list follows the order of the covariates originally given in the model formula, or a named list,</p> <pre>list (age = 60, sex = 1)</pre> <p>For time-inhomogeneous models fitted using the <code>pci</code> option to msm, "covariates" here include only those specified using the <code>covariates</code> argument to msm, and exclude the artificial covariates representing the time period.</p> <p>For time-inhomogeneous models fitted "by hand" by using a time-dependent covariate in the <code>covariates</code> argument to msm, the function pmatrix.piecewise.msm should be used to calculate transition probabilities.</p>
ci	If "normal", then calculate a confidence interval for the transition probabilities by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then calculating the resulting transition probability matrix for each replicate.
	If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm .
	If "none" (the default) then no confidence interval is calculated.
c1	Width of the symmetric confidence interval, relative to 1.
B	Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs
...	Optional arguments to be passed to MatrixExp to control the method of computing the matrix exponential.

Details

For a continuous-time homogeneous Markov process with transition intensity matrix Q , the probability of occupying state s at time $u + t$ conditionally on occupying state r at time u is given by the (r, s) entry of the matrix $P(t) = \exp(tQ)$, where $\exp()$ is the matrix exponential.

For non-homogeneous processes, where covariates and hence the transition intensity matrix Q are piecewise-constant in time, the transition probability matrix is calculated as a product of matrices over a series of intervals, as explained in [pmatrix.piecewise.msm](#).

The [pmatrix.piecewise.msm](#) function is only necessary for models fitted using a time-dependent covariate in the `covariates` argument to [msm](#). For time-inhomogeneous models fitted using "pci", [pmatrix.msm](#) can be used, with arguments `t` and `t1`, to calculate transition probabilities over any time period.

Value

The matrix of estimated transition probabilities $P(t)$ in the given time. Rows correspond to "from-state" and columns to "to-state".

Or if `ci="normal"` or `ci="bootstrap"`, [pmatrix.msm](#) returns a list with components `estimates` and `ci`, where `estimates` is the matrix of estimated transition probabilities, and `ci` is a list of two matrices containing the upper and lower confidence limits.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>.

See Also

[qmatrix.msm](#), [pmatrix.piecewise.msm](#), [boot.msm](#)

`pmatrix.piecewise.msm` *Transition probability matrix for processes with piecewise-constant intensities*

Description

Extract the estimated transition probability matrix from a fitted non-time-homogeneous multi-state model for a given time interval. This is a generalisation of [pmatrix.msm](#) to models with time-dependent covariates. Note that [pmatrix.msm](#) is sufficient to calculate transition probabilities for time-inhomogeneous models fitted using the `pci` argument to [msm](#).

Usage

```
pmatrix.piecewise.msm(x, t1, t2, times, covariates,
ci=c("none", "normal", "bootstrap"), cl=0.95, B=1000, ...)
```

Arguments

<code>x</code>	A fitted multi-state model, as returned by msm . This should be a non-homogeneous model, whose transition intensity matrix depends on a time-dependent covariate.
<code>t1</code>	The start of the time interval to estimate the transition probabilities for.
<code>t2</code>	The end of the time interval to estimate the transition probabilities for.

times	Cut points at which the transition intensity matrix changes.
covariates	A list with number of components one greater than the length of times. Each component of the list is specified in the same way as the covariates argument to <code>pmatrix.msm</code> . The components correspond to the covariate values in the intervals (t1, times[1]], (times[1], times[2]], ..., (times[length(times)], t2] (assuming that all elements of times are in the interval (t1, t2)).
ci	If "normal", then calculate a confidence interval for the transition probabilities by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then calculating the resulting transition probability matrix for each replicate. If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See <code>boot.msm</code> for more details of bootstrapping in <code>msm</code> . If "none" (the default) then no confidence interval is calculated.
c1	Width of the symmetric confidence interval, relative to 1.
B	Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs
...	Optional arguments to be passed to <code>MatrixExp</code> to control the method of computing the matrix exponential.

Details

Suppose a multi-state model has been fitted, in which the transition intensity matrix $Q(x(t))$ is modelled in terms of time-dependent covariates $x(t)$. The transition probability matrix $P(t_1, t_n)$ for the time interval (t_1, t_n) cannot be calculated from the estimated intensity matrix as $\exp((t_n - t_1)Q)$, because Q varies within the interval t_1, t_n . However, if the covariates are piecewise-constant, or can be approximated as piecewise-constant, then we can calculate $P(t_1, t_n)$ by multiplying together individual matrices $P(t_i, t_{i+1}) = \exp((t_{i+1} - t_i)Q)$, calculated over intervals where Q is constant:

$$P(t_1, t_n) = P(t_1, t_2)P(t_2, t_3) \dots P(t_{n-1}, t_n)$$

Value

The matrix of estimated transition probabilities $P(t)$ for the time interval $[t_1, t_n]$. That is, the probabilities of occupying state s at time t_n conditionally on occupying state r at time t_1 . Rows correspond to "from-state" and columns to "to-state".

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

`pmatrix.msm`

Examples

```
## Not run:
## In a clinical study, suppose patients are given a placebo in the
## first 5 weeks, then they begin treatment 1 at 5 weeks, and
## a combination of treatments 1 and 2 from 10 weeks.
## Suppose a multi-state model x has been fitted for the patients'
## progress, with treat1 and treat2 as time dependent covariates.

## Cut points for when treatment covariate changes
times <- c(0, 5, 10)

## Indicators for which treatments are active in the four intervals
## defined by the three cut points
covariates <- list( list (treat1=0, treat2=0), list (treat1=0, treat2=0), list(treat1=1, treat2=0),
list(treat1=1, treat2=1) )

## Calculate transition probabilities from the start of the study to 15 weeks
pmatrix.pieceswise.msm(x, 0, 15, times, covariates)

## End(Not run)
```

pnext.msm

*Probability of each state being next***Description**

Compute a matrix of the probability of each state s being the next state of the process after each state r . Together with the mean sojourn times in each state ([sojourn.msm](#)), these fully define a continuous-time Markov model.

Usage

```
pnext.msm(x, covariates = "mean",
          ci=c("delta","normal","bootstrap","none"), cl = 0.95,
          B=1000)
```

Arguments

x	A fitted multi-state model, as returned by msm .
covariates	The covariate values at which to estimate the intensities. This can either be: the string "mean", denoting the means of the covariates in the data (this is the default), the number 0, indicating that all the covariates should be set to zero, or a list of values, with optional names. For example

	list (60, 1)
	where the order of the list follows the order of the covariates originally given in the model formula, or a named list,
	list (age = 60, sex = 1)
ci	If "delta" (the default) then confidence intervals are calculated by the delta method. If "normal", then calculate a confidence interval by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then transforming. If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm .
c1	Width of the symmetric confidence interval to present. Defaults to 0.95.
B	Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs.

Details

For a continuous-time Markov process in state r , the probability that the next state is s is $-q_{rs}/q_{rr}$, where q_{rs} is the transition intensity ([qmatrix.msm](#)).

The model is fully parameterised by these probabilities together with the mean sojourn times $-1/q_{rr}$ in each state r . This gives a more intuitively meaningful description of a model than the intensity matrix.

Remember that **msm** deals with continuous-time not discrete-time models, so these are *not* the same as the probability of observing state s at a fixed time in the future. Those probabilities are given by [pmatrix.msm](#).

Value

The matrix of probabilities that the next move of a process in state r (rows) is to state s (columns).

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[qmatrix.msm](#), [pmatrix.msm](#), [qratio.msm](#)

```
prevalence.msm
```

Tables of observed and expected prevalences

Description

This provides a rough indication of the goodness of fit of a multi-state model, by estimating the observed numbers of individuals occupying each state at a series of times, and comparing these with forecasts from the fitted model.

Usage

```
prevalence.msm(x, times=NULL, timezero=NULL, initstates=NULL, covariates="mean",
  misccovariates="mean", piecewise.times=NULL, piecewise.covariates=NULL,
  ci=c("none", "normal", "bootstrap"), cl=0.95, B=1000,
  interp=c("start", "midpoint"), plot=FALSE, ...)
```

Arguments

<code>x</code>	A fitted multi-state model produced by msm .
<code>times</code>	Series of times at which to compute the observed and expected prevalences of states.
<code>timezero</code>	Initial time of the Markov process. Expected values are forecasted from here. Defaults to the minimum of the observation times given in the data.
<code>initstates</code>	Optional vector of the same length as the number of states. Gives the numbers of individuals occupying each state at the initial time, to be used for forecasting expected prevalences. The default is those observed in the data. These should add up to the actual number of people in the study at the start.
<code>covariates</code>	Covariate values for which to forecast expected state occupancy. See qmatrix.msm . Defaults to the mean values of the covariates in the data set.
<code>misccovariates</code>	(Misclassification models only) Values of covariates on the misclassification probability matrix for which to forecast expected state occupancy. Defaults to the mean values of the covariates in the data set.
<code>piecewise.times</code>	Times at which piecewise-constant intensities change. See pmatrix.piecewise.msm for how to specify this. This is only required for time-inhomogeneous models specified using explicit time-dependent covariates, and should not be used for models specified using "pci".
<code>piecewise.covariates</code>	Covariates on which the piecewise-constant intensities depend. See pmatrix.piecewise.msm for how to specify this.
<code>ci</code>	If "normal", then calculate a confidence interval for the expected prevalences by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then calculating the expected prevalences for each replicate.

	If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm .
	If "none" (the default) then no confidence interval is calculated.
c1	Width of the symmetric confidence interval, relative to 1
B	Number of bootstrap replicates
interp	Suppose an individual was observed in states S_{r-1} and S_r at two consecutive times t_{r-1} and t_r , and we want to estimate 'observed' prevalences at a time t between t_{r-1} and t_r . If <code>interp="start"</code> , then individuals are assumed to be in state S_{r-1} at time t , the same state as they were at t_{r-1} . If <code>interp="midpoint"</code> then if $t \leq (t_{r-1} + t_r)/2$, the midpoint of t_{r-1} and t_r , the state at t is assumed to be S_{r-1} , otherwise S_r . This is generally more reasonable for "progressive" models.
plot	Generate a plot of observed against expected prevalences. See plot.prevalence.msm
...	Further arguments to pass to plot.prevalence.msm .

Details

The fitted transition probability matrix is used to forecast expected prevalences from the state occupancy at the initial time. To produce the expected number in state j at time t after the start, the number of individuals under observation at time t (including those who have died, but not those lost to follow-up) is multiplied by the product of the proportion of individuals in each state at the initial time and the transition probability matrix in the time interval t . The proportion of individuals in each state at the "initial" time is estimated, if necessary, in the same way as the observed prevalences.

For misclassification models (fitted using an `ematrix`), this aims to assess the fit of the full model for the *observed* states. That is, the combined Markov progression model for the true states and the misclassification model. Thus, expected prevalences of *true* states are estimated from the assumed proportion occupying each state at the initial time using the fitted transition probability matrix. The vector of expected prevalences of true states is then multiplied by the fitted misclassification probability matrix to obtain the expected prevalences of *observed* states.

For general hidden Markov models, the observed state is taken to be the predicted underlying state from the Viterbi algorithm ([viterbi.msm](#)). The goodness of fit of these states to the underlying Markov model is tested.

Note that this function assumes intensities are the same for all individuals. By default they are taken from the mean values of the covariates in the population. Piecewise-constant intensities may be assumed, through the arguments `piecewise.times` and `piecewise.covariates`.

For an example of this approach, see Gentleman *et al.* (1994).

Value

A list of matrices, with components:

Observed Table of observed numbers of individuals in each state at each time

Observed percentages

Corresponding percentage of the individuals at risk at each time.

Expected Table of corresponding expected numbers.

Expected percentages

Corresponding percentage of the individuals at risk at each time.

Or if `ci.boot = TRUE`, the component `Expected` is a list with components `estimates` and `ci`. `estimates` is a matrix of the expected prevalences, and `ci` is a list of two matrices, containing the confidence limits. The component `Expected percentages` has a similar format.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

Gentleman, R.C., Lawless, J.F., Lindsey, J.C. and Yan, P. Multi-state Markov models for analysing incomplete disease history data with illustrations for HIV disease. *Statistics in Medicine* (1994) 13(3): 805–821.

See Also

[msm](#), [summary.msm](#)

psor

Psoriatic arthritis data

Description

A series of observations of grades of psoriatic arthritis, as indicated by numbers of damaged joints.

Usage

`data(psor)`

Format

A data frame containing 806 observations, representing visits to a psoriatic arthritis (PsA) clinic from 305 patients. The rows are grouped by patient number and ordered by examination time. Each row represents an examination and contains additional covariates.

<code>ptnum</code>	(numeric)	Patient identification number
<code>months</code>	(numeric)	Examination time in months
<code>state</code>	(numeric)	Clinical state of PsA. Patients in states 1, 2, 3 and 4 have 0, 1 to 4, 5 to 9 and 10 or more damaged joints, respectively.
<code>hieffusn</code>	(numeric)	Presence of five or more effusions
<code>ollwsdr</code>	(character)	Erythrocyte sedimentation rate of less than 15 mm/h

References

Gladman, D. D. and Farewell, V.T. (1999) Progression in psoriatic arthritis: role of time-varying clinical indicators. *J. Rheumatol.* 26(11):2409-13

Examples

```
## Four-state progression-only model with high effusion and low
## sedimentation rate as covariates on the progression rates. High
## effusion is assumed to have the same effect on the 1-2, 2-3, and 3-4
## progression rates, while low sedimentation rate has the same effect
## on the 1-2 and 2-3 intensities, but a different effect on the 3-4.

data(psor)
psor.q <- rbind(c(0,0.1,0,0),c(0,0,0.1,0),c(0,0,0,0.1),c(0,0,0,0))
psor.msm <- msm(state ~ months, subject=ptnum, data=psor,
               qmatrix = psor.q, covariates = ~ollwsdrt+hieffusn,
               constraint = list(hieffusn=c(1,1,1),ollwsdrt=c(1,1,2)),
               fixedpars=FALSE, control = list(REPORT=1,trace=2), method="BFGS")
qmatrix.msm(psor.msm)
sojourn.msm(psor.msm)
hazard.msm(psor.msm)
```

qmatrix.msm

Transition intensity matrix

Description

Extract the estimated transition intensity matrix, and the corresponding standard errors, from a fitted multi-state model at a given set of covariate values.

Usage

```
qmatrix.msm(x, covariates="mean", sojourn=FALSE,
           ci=c("delta","normal","bootstrap","none"), cl=0.95,
           B=1000)
```

Arguments

x	A fitted multi-state model, as returned by <code>msm</code> .
covariates	The covariate values at which to estimate the intensity matrix. This can either be: the string "mean", denoting the means of the covariates in the data (this is the default), the number 0, indicating that all the covariates should be set to zero, or a list of values, with optional names. For example

	list (60, 1)
	where the order of the list follows the order of the covariates originally given in the model formula, or a named list,
	list (age = 60, sex = 1)
sojourn	Set to TRUE if the estimated sojourn times and their standard errors should also be returned.
ci	If "delta" (the default) then confidence intervals are calculated by the delta method, or by simple transformation of the Hessian in the very simplest cases. Normality on the log scale is assumed. If "normal", then calculate a confidence interval by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then transforming. If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm .
c1	Width of the symmetric confidence interval to present. Defaults to 0.95.
B	Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs.

Details

Transition intensities and covariate effects are estimated on the log scale by [msm](#). A covariance matrix is estimated from the Hessian of the maximised log-likelihood.

A more practically meaningful parameterisation of a continuous-time Markov model with transition intensities q_{rs} is in terms of the mean sojourn times $-1/q_{rr}$ in each state r and the probabilities that the next move of the process when in state r is to state s , $-q_{rs}/q_{rr}$.

Value

A list with components:

estimate	Estimated transition intensity matrix.
SE	Corresponding approximate standard errors.
L	Lower confidence limits
U	Upper confidence limits

Or if ci="none", then `qmatrix.msm` just returns the estimated transition intensity matrix.

If sojourn is TRUE, extra components called sojourn, sojournSE, sojournL and sojournU are included, containing the estimates, standard errors and confidence limits, respectively, of the mean sojourn times in each transient state.

The default print method for objects returned by `qmatrix.msm` presents estimates and confidence limits. To present estimates and standard errors, do something like

```
qmatrix.msm(x)[c("estimates", "SE")]
```

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[pmatix.msm](#), [sojourn.msm](#), [deltamethod](#), [ematrix.msm](#)

qratio.msm

Estimated ratio of transition intensities

Description

Compute the estimate and approximate standard error of the ratio of two estimated transition intensities from a fitted multi-state model at a given set of covariate values.

Usage

```
qratio.msm(x, ind1, ind2, covariates = "mean",
           ci=c("delta","normal","bootstrap","none"), c1 = 0.95,
           B=1000)
```

Arguments

x	A fitted multi-state model, as returned by msm .
ind1	Pair of numbers giving the indices in the intensity matrix of the numerator of the ratio, for example, c(1, 2).
ind2	Pair of numbers giving the indices in the intensity matrix of the denominator of the ratio, for example, c(2, 1).
covariates	The covariate values at which to estimate the intensities. This can either be: <ul style="list-style-type: none"> the string "mean", denoting the means of the covariates in the data (this is the default), the number 0, indicating that all the covariates should be set to zero, or a list of values, with optional names. For example <pre>list (60, 1)</pre> where the order of the list follows the order of the covariates originally given in the model formula, or a named list, <pre>list (age = 60, sex = 1)</pre>
ci	If "delta" (the default) then confidence intervals are calculated by the delta method. If "normal", then calculate a confidence interval by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum

likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then transforming.

If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See [boot.msm](#) for more details of bootstrapping in **msm**.

- c1 Width of the symmetric confidence interval to present. Defaults to 0.95.
- B Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs

Details

For example, we might want to compute the ratio of the progression rate and recovery rate for a fitted model `disease.msm` with a health state (state 1) and a disease state (state 2). In this case, the progression rate is the (1,2) entry of the intensity matrix, and the recovery rate is the (2,1) entry. Thus to compute this ratio with covariates set to their means, we call

```
qratio.msm(disease.msm, c(1,2), c(2,1)) .
```

Standard errors are estimated by the delta method. Confidence limits are estimated by assuming normality on the log scale.

Value

A named vector with elements `estimate`, `se`, `L` and `U` containing the estimate, standard error, lower and upper confidence limits, respectively, of the ratio of intensities.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[qmatrix.msm](#)

scorer resid.msm

Score residuals

Description

Score residuals for detecting outlying subjects.

Usage

```
scorer resid.msm(x, plot=FALSE)
```

Arguments

x	A fitted multi-state model, as returned by <code>msm</code> .
plot	If TRUE, display a simple plot of the residuals in subject order, labelled by subject identifiers

Details

The score residual for a single subject is

$$U(\theta)^T I(\theta)^{-1} U(\theta)$$

where $U(\theta)$ is the vector of first derivatives of the log-likelihood for that subject at maximum likelihood estimates θ , and $I(\theta)$ is the observed Fisher information matrix, that is, the matrix of second derivatives of minus the log-likelihood for that subject at theta.

Subjects with a higher influence on the maximum likelihood estimates will have higher score residuals.

Value

Vector of the residuals, named by subject identifiers.

Author(s)

Andrew Titman <a.titman@lancaster.ac.uk> (theory), Chris Jackson <chris.jackson@mrc-bsu.cam.ac.uk> (code)

sim.msm

Simulate one individual trajectory from a continuous-time Markov model

Description

Simulate one realisation from a continuous-time Markov process up to a given time.

Usage

```
sim.msm(qmatrix, maxtime, covs=NULL, beta=NULL, obstimes=0, start=1,
mintime=0)
```

Arguments

qmatrix	The transition intensity matrix of the Markov process. The diagonal of qmatrix is ignored, and computed as appropriate so that the rows sum to zero. For example, a possible qmatrix for a three state illness-death model with recovery is: <pre> rbind(c(0, 0.1, 0.02), c(0.1, 0, 0.01), c(0, 0, 0)) </pre>
maxtime	Maximum time for the simulated process.
covs	Matrix of time-dependent covariates, with one row for each observation time and one column for each covariate.
beta	Matrix of linear covariate effects on log transition intensities. The rows correspond to different covariates, and the columns to the transition intensities. The intensities are ordered by reading across rows of the intensity matrix, starting with the first, counting the positive off-diagonal elements of the matrix.
obstimes	Vector of times at which the covariates are observed.
start	Starting state of the process. Defaults to 1.
mintime	Starting time of the process. Defaults to 0.

Details

The effect of time-dependent covariates on the transition intensity matrix for an individual is determined by assuming that the covariate is a step function which remains constant in between the individual's observation times.

Value

A list with components,

states	Simulated states through which the process moves. This ends with either an absorption before <code>obstime</code> , or a transient state at <code>obstime</code> .
times	Exact times at which the process changes to the corresponding states
qmatrix	The given transition intensity matrix

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[simmulti.msm](#)

Examples

```

qmatrix <- rbind(
  c(-0.2,  0.1,  0.1 ),
  c(0.5,  -0.6,  0.1 ),

```

```

                                c(0, 0, 0)
                                )
sim.msm(qmatrix, 30)

```

simmulti.msm	<i>Simulate multiple trajectories from a multi-state Markov model with arbitrary observation times</i>
--------------	--

Description

Simulate a number of individual realisations from a multi-state Markov process. Observations of the process are made at specified arbitrary times for each individual.

Usage

```

simmulti.msm(data, qmatrix, covariates=NULL, death = FALSE, start,
              ematrix=NULL, misccovariates=NULL, hmodel=NULL, hcovariates=NULL,
              censor.states=NULL)

```

Arguments

data	A data frame with a mandatory column named <code>time</code> , representing observation times. The optional column named <code>subject</code> , corresponds to subject identification numbers. If not given, all observations are assumed to be on the same individual. Observation times should be sorted within individuals. The optional column named <code>cens</code> indicates the times at which simulated states should be censored. If <code>cens==0</code> then the state is not censored, and if <code>cens==k</code> , say, then all simulated states at that time which are in the set <code>censor.states</code> are replaced by <code>k</code> . Other named columns of the data frame represent any covariates.
qmatrix	The transition intensity matrix of the Markov process, with any covariates set to zero. The diagonal of <code>qmatrix</code> is ignored, and computed as appropriate so that the rows sum to zero. For example, a possible <code>qmatrix</code> for a three state illness-death model with recovery is: <pre> rbind(c(0, 0.1, 0.02), c(0.1, 0, 0.01), c(0, 0, 0)) </pre>
covariates	List of covariate effects on log transition intensities. Each element is a vector of the effects of one covariate on all the transition intensities. The intensities are ordered by reading across rows of the intensity matrix, starting with the first, counting the positive off-diagonal elements of the matrix. For example, for a multi-state model with three transition intensities, and two covariates <code>x</code> and <code>y</code> on each intensity, <pre> covariates=list(x = c(-0.3,-0.3,-0.3), y=c(0.1, 0.1, 0.1)) </pre>
death	Vector of indices of the death states. A death state is an absorbing state whose time of entry is known exactly, but the individual is assumed to be in an unknown transient state ("alive") at the previous instant. This is the usual situation for

times of death in chronic disease monitoring data. For example, if you specify `death = c(4, 5)` then states 4 and 5 are assumed to be death states. `death = TRUE` indicates that the final state is a death state, and `death = FALSE` (the default) indicates that there is no death state.

<code>start</code>	A vector with the same number of elements as there are distinct subjects in the data, giving the states in which each corresponding individual begins. Defaults to state 1 for each subject.
<code>ematrix</code>	An optional misclassification matrix for generating observed states conditionally on the simulated true states. As defined in msm .
<code>misccovariates</code>	Covariate effects on misclassification probabilities via multinomial logistic regression. Linear effects operate on the log of each probability relative to the probability of classification in the correct state. In same format as <code>covariates</code> .
<code>hmodel</code>	An optional hidden Markov model for generating observed outcomes conditionally on the simulated true states. As defined in msm .
<code>hcovariates</code>	List of the same length as <code>hmodel</code> , defining any covariates governing the hidden Markov outcome models. Unlike in the <code>msm</code> function, this should also define the values of the covariate effects. Each element of the list is a named vector of the initial values for each set of covariates for that state. For example, for a three-state hidden Markov model with two, one and no covariates on the state 1, 2 and 3 outcome models respectively, <pre>hcovariates = list (c(acute=-8, age=0), c(acute=-8), NULL)</pre>
<code>sensor.states</code>	Set of simulated states which should be replaced by a censoring indicator at censoring times. By default this is all transient states (representing alive, with unknown state).

Details

`sim.msm` is called repeatedly to produce a simulated trajectory for each individual. The state at each specified observation time is then taken to produce a new column state. The effect of time-dependent covariates on the transition intensity matrix for an individual is determined by assuming that the covariate is a step function which remains constant in between the individual's observation times. If the subject enters an absorbing state, then only the first observation in that state is kept in the data frame. Rows corresponding to future observations are deleted. The entry times into states given in `death` are assumed to be known exactly.

Value

A data frame with columns,

<code>subject</code>	Subject identification indicators
<code>time</code>	Observation times
<code>state</code>	Simulated (true) state at the corresponding time
<code>obs</code>	Observed outcome at the corresponding time, if <code>ematrix</code> or <code>hmodel</code> was supplied

plus any supplied covariates.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[sim.msm](#)

Examples

```
### Simulate 100 individuals with common observation times
sim.df <- data.frame(subject = rep(1:100, rep(13,100)), time = rep(seq(0, 24, 2), 100))
qmatrix <- rbind(c(-0.11, 0.1, 0.01 ),
                c(0.05, -0.15, 0.1 ),
                c(0.02, 0.07, -0.09))
simmulti.msm(sim.df, qmatrix)
```

sojourn.msm

Mean sojourn times from a multi-state model

Description

Estimate the mean sojourn times in the transient states of a multi-state model and their confidence limits.

Usage

```
sojourn.msm(x, covariates="mean", ci=c("delta", "normal", "bootstrap", "none"),
            cl=0.95, B=1000)
```

Arguments

x	A fitted multi-state model, as returned by msm .
covariates	The covariate values at which to estimate the mean sojourn times. This can either be: <ul style="list-style-type: none"> the string "mean", denoting the means of the covariates in the data (this is the default), the number 0, indicating that all the covariates should be set to zero, a list of values, with optional names. For example, <ul style="list-style-type: none"> list(60, 1), where the order of the list follows the order of the covariates originally given in the model formula, or a named list, e.g. list (age = 60, sex = 1)

ci	<p>If "delta" (the default) then confidence intervals are calculated by the delta method, or by simple transformation of the Hessian in the very simplest cases.</p> <p>If "normal", then calculate a confidence interval by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then transforming.</p> <p>If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm.</p>
c1	Width of the symmetric confidence interval to present. Defaults to 0.95.
B	Number of bootstrap replicates, or number of normal simulations from the distribution of the MLEs

Details

The mean sojourn time in a transient state r is estimated by $-1/q_{rr}$, where q_{rr} is the r th entry on the diagonal of the estimated transition intensity matrix.

A continuous-time Markov model is fully specified by the mean sojourn times and the probability that each state is next ([pnext.msm](#)). This is a more intuitively meaningful description of a model than the transition intensity matrix ([qmatrix.msm](#)).

Value

A data frame with components:

estimates	Estimated mean sojourn times in the transient states.
SE	Corresponding standard errors.
L	Lower confidence limits.
U	Upper confidence limits.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#), [qmatrix.msm](#), [deltamethod](#)

statetable.msm	<i>Table of transitions</i>
----------------	-----------------------------

Description

Calculates a frequency table counting the number of times each pair of states were observed in successive observation times. This can be a useful way of summarising multi-state data.

Usage

```
statetable.msm(state, subject, data=NULL)
```

Arguments

state	Observed states, assumed to be ordered by time within each subject.
subject	Subject identification numbers corresponding to state. If not given, all observations are assumed to be on the same subject.
data	An optional data frame in which the variables represented by subject and state can be found.

Value

A frequency table with starting states as rows and finishing states as columns.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[crudeinits.msm](#)

Examples

```
## Heart transplant data
data(cav)

## 148 deaths from state 1, 48 from state 2 and 55 from state 3.
statetable.msm(state, PTNUM, data=cav)
```

surface.msm *Explore the likelihood surface*

Description

Plot the log-likelihood surface with respect to two parameters.

Usage

```
surface.msm(x, params=c(1,2), np=10, type=c("contour","filled.contour","persp","image"),
            point=NULL, xrange=NULL, yrange=NULL,...)
```

Arguments

x	Output from msm , representing a fitted msm model.
params	Integer vector with two elements, giving the indices of the parameters to vary. All other parameters will be fixed. Defaults to <code>c(1,2)</code> , representing the first two log transition intensities. See the <code>fixedpars</code> argument to <code>msm</code> for a definition of these indices.
np	Number of grid points to use in each direction, by default 10. An <code>np × np</code> grid will be used to evaluate the likelihood surface. If 100 likelihood function evaluations is slow, then reduce this.
type	Character string specifying the type of plot to produce.
"contour"	Contour plot, using the R function contour .
"filled.contour"	Solid-color contour plot, using the R function filled.contour .
"persp"	Perspective plot, using the R function persp .
"image"	Grid color plot, using the R function image .
point	Vector of length <code>n</code> , where <code>n</code> is the number of parameters in the model, including the parameters that will be varied here. This specifies the point at which to fix the likelihood. By default, this is the maximum likelihood estimates stored in the fitted model <code>x</code> , <code>x\$estimates</code> .
xrange	Range to plot for the first varied parameter. Defaults to plus and minus two standard errors, obtained from the Hessian at the maximum likelihood estimate.
yrange	Range to plot for the second varied parameter. Defaults to plus and minus two standard errors, obtained from the Hessian at the maximum likelihood estimate.
...	Further arguments to be passed to the plotting function.

Details

Draws a contour or perspective plot. Useful for diagnosing irregularities in the likelihood surface. If you want to use these plots before running the maximum likelihood estimation, then just run `msm` with all estimates fixed at their initial values.

contour.msm just calls surface.msm with type = "contour".

persp.msm just calls surface.msm with type = "persp".

image.msm just calls surface.msm with type = "image".

As these three functions are methods of the generic functions contour, persp and image, they can be invoked as contour(x), persp(x) or image(x), where x is a fitted msm object.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[msm](#), [contour](#), [filled.contour](#), [persp](#), [image](#).

tnorm

Truncated Normal distribution

Description

Density, distribution function, quantile function and random generation for the truncated Normal distribution with mean equal to mean and standard deviation equal to sd before truncation, and truncated on the interval [lower, upper].

Usage

```
dtnorm(x, mean=0, sd=1, lower=-Inf, upper=Inf, log = FALSE)
ptnorm(q, mean=0, sd=1, lower=-Inf, upper=Inf,
       lower.tail = TRUE, log.p = FALSE)
qtnorm(p, mean=0, sd=1, lower=-Inf, upper=Inf,
       lower.tail = TRUE, log.p = FALSE)
rtnorm(n, mean=0, sd=1, lower=-Inf, upper=Inf)
```

Arguments

x, q	vector of quantiles.
p	vector of probabilities.
n	number of observations. If length(n) > 1, the length is taken to be the number required.
mean	vector of means.
sd	vector of standard deviations.
lower	lower truncation point.
upper	upper truncation point.
log, log.p	logical; if TRUE, probabilities p are given as log(p).
lower.tail	logical; if TRUE (default), probabilities are P[X <= x], otherwise, P[X > x].

Details

The truncated normal distribution has density

$$f(x, \mu, \sigma) = \phi(x, \mu, \sigma) / (\Phi(u, \mu, \sigma) - \Phi(l, \mu, \sigma))$$

for $l \leq x \leq u$, and 0 otherwise.

μ is the mean of the original Normal distribution before truncation,

σ is the corresponding standard deviation,

u is the upper truncation point,

l is the lower truncation point,

$\phi(x)$ is the density of the corresponding normal distribution, and

$\Phi(x)$ is the distribution function of the corresponding normal distribution.

If mean or sd are not specified they assume the default values of 0 and 1, respectively.

If lower or upper are not specified they assume the default values of `-Inf` and `Inf`, respectively, corresponding to no lower or no upper truncation.

Therefore, for example, `dtnorm(x)`, with no other arguments, is simply equivalent to `dnorm(x)`.

Only `rtnorm` is used in the `msm` package, to simulate from hidden Markov models with truncated normal distributions. This uses the rejection sampling algorithms described by Robert (1995).

These functions are merely provided for completion, and are not optimized for numerical stability or speed. To fit a hidden Markov model with a truncated Normal response distribution, use a [hmmTNorm](#) constructor. See the [hmm-dists](#) help page for further details.

Value

`dtnorm` gives the density, `ptnorm` gives the distribution function, `qtnorm` gives the quantile function, and `rtnorm` generates random deviates.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

Robert, C. P. Simulation of truncated normal variables. *Statistics and Computing* (1995) 5, 121–125

See Also

[dnorm](#)

Examples

```
x <- seq(50, 90, by=1)
plot(x, dnorm(x, 70, 10), type="l", ylim=c(0,0.06)) ## standard Normal distribution
lines(x, dtnorm(x, 70, 10, 60, 80), type="l")      ## truncated Normal distribution
```

totlos.msm	<i>Total length of stay</i>
------------	-----------------------------

Description

Estimate the expected total length of stay in each set of states, for a given period of evolution of a multi-state model.

Usage

```
totlos.msm(x, start=1, end=NULL, fromt=0, tot=Inf, covariates="mean",
           ci=c("none", "normal", "bootstrap"), cl=0.95, B=1000, ...)
```

Arguments

x	A fitted multi-state model, as returned by msm .
start	State at the beginning of the period.
end	States to forecast the total length of stay in. Defaults to all states.
fromt	Time from which to estimate total length of stay. Defaults to 0, the beginning of the process.
tot	Time up to which total length of stay is estimated. Defaults to infinity, giving the expected time spent in the state until absorption. For models without an absorbing state, t must be specified.
covariates	The covariate values to estimate for. This can either be: <p>the string "mean", denoting the means of the covariates in the data (this is the default),</p> <p>the number 0, indicating that all the covariates should be set to zero,</p> <p>or a list of values, with optional names. For example</p> <pre>list (60, 1)</pre> <p>where the order of the list follows the order of the covariates originally given in the model formula, or a named list,</p> <pre>list (age = 60, sex = 1)</pre>
ci	If "normal", then calculate a confidence interval by simulating B random vectors from the asymptotic multivariate normal distribution implied by the maximum likelihood estimates (and covariance matrix) of the log transition intensities and covariate effects, then calculating the total length of stay for each replicate. <p>If "bootstrap" then calculate a confidence interval by non-parametric bootstrap refitting. This is 1-2 orders of magnitude slower than the "normal" method, but is expected to be more accurate. See boot.msm for more details of bootstrapping in msm.</p> <p>If "none" (the default) then no confidence interval is calculated.</p>

c1	Width of the symmetric confidence interval, relative to 1
B	Number of bootstrap replicates
...	Further arguments to be passed to the integrate function to control the numerical integration.

Details

The expected total length of stay in state j between times t_1 and t_2 , from the point of view of an individual in state i at time 0, is defined by the integral from t_1 to t_2 of the i, j entry of the transition probability matrix $P(t)$. As the individual entries of $P(t) = \exp(tQ)$ are not available explicitly in terms of t for a general Markov model, this integral is calculated numerically, using the [integrate](#) function. This may take a long time for models with many states where $P(t)$ is expensive to calculate.

For a model where the individual has only one place to go from each state, and each state is visited only once, for example a progressive disease model with no recovery or death, these are equal to the mean sojourn time in each state. However, consider a three-state health-disease-death model with transitions from health to disease, health to death, and disease to death, where everybody starts healthy. In this case the mean sojourn time in the disease state will be greater than the expected length of stay in the disease state. This is because the mean sojourn time in a state is conditional on entering the state, whereas the expected total time diseased is a forecast for a healthy individual, who may die before getting the disease.

The only time-inhomogeneous models handled by this function are those specified using the `pci` option to [msm](#). For any other inhomogeneous models, the function assumes that covariates are constant on the forecasted interval.

Value

A vector of expected total lengths of stay for each transient state.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

See Also

[sojourn.msm](#), [pmatrix.msm](#), [integrate](#), [boot.msm](#).

transient.msm

Transient and absorbing states

Description

Returns the transient and absorbing states of either a fitted model or a transition intensity matrix.

Usage

```
transient.msm(x=NULL, qmatrix=NULL)
absorbing.msm(x=NULL, qmatrix=NULL)
```

Arguments

`x` A fitted multi-state model as returned by [msm](#).

`qmatrix` A transition intensity matrix. The diagonal is ignored and taken to be minus the sum of the rest of the row.

Value

A vector of the ordinal indices of the transient or absorbing states.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

viterbi.msm

Calculate the most likely path through underlying states

Description

For a fitted hidden Markov model, or a model with censored state observations, the Viterbi algorithm recursively constructs the path with the highest probability through the underlying states.

Usage

```
viterbi.msm(x)
```

Arguments

`x` A fitted hidden Markov multi-state model, or a model with censored state observations, as produced by [msm](#)

Value

A data frame with columns:

`subject` = subject identification numbers

`time` = times of observations

`observed` = corresponding observed states

`fitted` = corresponding fitted states found by Viterbi recursion. If the model is not a hidden Markov model and there are no censored state observations, this is just the observed states.

Author(s)

C. H. Jackson <chris.jackson@mrc-bsu.cam.ac.uk>

References

Durbin, R., Eddy, S., Krogh, A. and Mitchison, G. *Biological sequence analysis*, Cambridge University Press, 1998.

See Also

[msm](#)

Index

*Topic **datasets**

aneur, 2
bos, 6
cav, 7
fev, 13
psor, 55

*Topic **distribution**

hmm-dists, 15
medists, 21
pexp, 40
tnorm, 68

*Topic **math**

deltamethod, 10
MatrixExp, 19

*Topic **models**

boot.msm, 3
coef.msm, 8
crudeinits.msm, 9
ematrix.msm, 12
hazard.msm, 14
logLik.msm, 18
lrtest.msm, 19
msm, 23
msm.summary, 35
odds.msm, 36
pearson.msm, 37
plot.msm, 42
plot.prevalence.msm, 43
plot.survfit.msm, 44
plotprog.msm, 46
pmatrix.msm, 47
pmatrix.piecewise.msm, 49
pnext.msm, 51
prevalence.msm, 53
qmatrix.msm, 56
qratio.msm, 58
scorer resid.msm, 59
sim.msm, 60
simmulti.msm, 62

sojourn.msm, 64
statetable.msm, 66
surface.msm, 67
totlos.msm, 70
transient.msm, 71
viterbi.msm, 72

absorbing.msm (transient.msm), 71
aneur, 2

boot.msm, 3, 4, 11, 12, 48–50, 52, 54, 57, 59,
65, 70, 71

bos, 6

cav, 7

coef.msm, 8

contour, 67, 68

contour.msm (surface.msm), 67

crudeinits.msm, 9, 24, 66

dbeta, 16

dbinom, 16

deltamethod, 4, 10, 58, 65

dexp, 16, 41

dgamma, 16

dmenorm (medists), 21

dmeunif (medists), 21

dnbinom, 16

dnorm, 23, 69

dpexp (pexp), 40

dpois, 16

dtnorm, 23

dtnorm (tnorm), 68

dunif, 23

dweibull, 16

ematrix.msm, 4, 5, 12, 13, 32, 58

fev, 13

filled.contour, 67, 68

- hazard.msm, 14, 35, 36
- hmm-dists, 15
- hmmBeta (hmm-dists), 15
- hmmBinom (hmm-dists), 15
- hmmCat, 25
- hmmCat (hmm-dists), 15
- hmmExp (hmm-dists), 15
- hmmGamma (hmm-dists), 15
- hmmIdent (hmm-dists), 15
- hmmLNorm (hmm-dists), 15
- hmmMETNorm, 22
- hmmMETNorm (hmm-dists), 15
- hmmMEUnif, 22
- hmmMEUnif (hmm-dists), 15
- hmmNBinom (hmm-dists), 15
- hmmNorm (hmm-dists), 15
- hmmPois (hmm-dists), 15
- hmmTNorm, 69
- hmmTNorm (hmm-dists), 15
- hmmUnif (hmm-dists), 15
- hmmWeibull (hmm-dists), 15

- image, 67, 68
- image.msm (surface.msm), 67
- integrate, 71

- lines, 42
- lines.survfit, 46, 47
- load, 4
- logLik.msm, 18, 19
- lrtest.msm, 19

- MatrixExp, 19, 48, 50
- medists, 16, 21
- msm, 3, 4, 8, 9, 12–15, 18, 19, 23, 32, 33, 35–38, 40, 42, 43, 45, 48, 49, 51, 53, 55–58, 60, 63–65, 67, 68, 70–73
- msm.summary, 35

- nlm, 31, 33

- odds.msm, 15, 36
- optim, 28, 30, 31, 33

- par, 42, 44, 46, 47
- pearson.msm, 37
- persp, 67, 68
- persp.msm (surface.msm), 67
- pexp, 40
- plot, 42, 44, 46, 47
- plot.msm, 34, 42
- plot.prevalence.msm, 43, 46, 54
- plot.survfit, 46, 47
- plot.survfit.msm, 44
- plotprog.msm, 46, 47
- pmatrix.msm, 4, 5, 32, 34, 45, 47, 49, 50, 52, 58, 71
- pmatrix.piecewise.msm, 4, 5, 44, 48, 49, 49, 53
- pmenorm (medists), 21
- pmeunif (medists), 21
- pnext.msm, 51, 65
- ppexp (pexp), 40
- prevalence.msm, 4, 5, 35, 36, 38, 40, 43, 44, 53
- print.summary.msm (msm.summary), 35
- psor, 55
- ptnorm (tnorm), 68

- qmatrix.msm, 4, 5, 13, 30, 32, 34, 35, 43, 49, 52, 53, 56, 57, 59, 65
- qmenorm (medists), 21
- qmeunif (medists), 21
- qpexp (pexp), 40
- qratio.msm, 4, 5, 52, 58
- qtnorm (tnorm), 68

- rmenorm (medists), 21
- rmeunif (medists), 21
- rpexp (pexp), 40
- rtnorm (tnorm), 68

- save, 4
- scorer resid.msm, 40, 59
- sim.msm, 41, 60, 63, 64
- simmulti.msm, 34, 61, 62
- sojourn.msm, 5, 32, 34, 51, 58, 64, 71
- statetable.msm, 10, 66
- summary.msm, 14, 15, 34, 55
- summary.msm (msm.summary), 35
- surface.msm, 67
- survfit, 46, 47

- tnorm, 16, 68
- totlos.msm, 4, 5, 70
- transient.msm, 71

- viterbi.msm, 54, 72