

Package ‘mcmc’

October 9, 2009

Version 0.7-3

Date 2009-10-08

Title Markov Chain Monte Carlo

Author Charles J. Geyer <charlie@stat.umn.edu>.

Maintainer Charles J. Geyer <charlie@stat.umn.edu>

Depends R (>= 2.6.0)

Suggests xtable, Iso

Description functions for Markov chain Monte Carlo (MCMC).

License X11

URL <http://www.stat.umn.edu/geyer/mcmc/>

Repository CRAN

Date/Publication 2009-10-09 06:55:25

R topics documented:

foo	2
initseq	2
logit	4
metrop	5
olbm	7
temper	8

Index	13
--------------	-----------

foo

Simulated logistic regression data.

Description

Like it says

Usage

```
data(foo)
```

Format

A data frame with variables

x1 quantitative predictor.

x2 quantitative predictor.

x3 quantitative predictor.

y Bernoulli response.

Examples

```
library(mcmc)
data(foo)
out <- glm(y ~ x1 + x2 + x3, family = binomial, data = foo)
summary(out)
```

initseq

Initial Sequence Estimators

Description

Variance of sample mean of functional of reversible Markov chain using methods of Geyer (1992).

Usage

```
initseq(x)
```

Arguments

x a numeric vector that is a scalar-valued functional of a reversible Markov chain.

Details

Let

$$\gamma_k = \text{cov}(X_i, X_{i+k})$$

considered as a function of the lag k be the autocovariance function of the input time series. Define

$$\Gamma_k = \gamma_{2k} + \gamma_{2k+1}$$

the sum of consecutive pairs of autocovariances. Then Theorem 3.1 in Geyer (1992) says that Γ_k considered as a function of k is strictly positive, strictly decreasing, and strictly convex, assuming the input time series is a scalar-valued functional of a reversible Markov chain. All of the MCMC done by this package is reversible. This R function estimates the “big gamma” function, Γ_k considered as a function of k , subject to three different constraints, (1) nonnegative, (2) nonnegative and nonincreasing, and (3) nonnegative, nonincreasing, and convex. It also estimates the variance in the Markov chain central limit theorem (CLT)

$$\gamma_0 + 2 \sum_{k=1}^{\infty} \gamma_k = -\gamma_0 + 2 \sum_{k=0}^{\infty} \Gamma_k$$

Note: The batch means provided by `metrop` are also scalar functionals of a reversible Markov chain. Thus these initial sequence estimators applied to the batch means give valid standard errors for the mean of the batch means even when the batch length is too short to provide a valid estimate of asymptotic variance. One does, of course, have to multiply the asymptotic variance of the batch means by the batch length to get the asymptotic variance for the unbached chain.

Value

a list containing the following components:

<code>gamma0</code>	the scalar γ_0 , the marginal variance of x .
<code>Gamma.pos</code>	the vector Γ , estimated so as to be nonnegative, where, as always, R uses one-origin indexing so <code>Gamma.pos[1]</code> is Γ_0 .
<code>Gamma.dec</code>	the vector Γ , estimated so as to be nonnegative and nonincreasing, where, as always, R uses one-origin indexing so <code>Gamma.dec[1]</code> is Γ_0 .
<code>Gamma.con</code>	the vector Γ , estimated so as to be nonnegative and nonincreasing and convex, where, as always, R uses one-origin indexing so <code>Gamma.con[1]</code> is Γ_0 .
<code>var.pos</code>	the scalar $-\text{gamma0} + 2 * \text{sum}(\text{Gamma.pos})$, which is the asymptotic variance in the Markov chain CLT. Divide by <code>length(x)</code> to get the approximate variance of the sample mean of x .
<code>var.dec</code>	the scalar $-\text{gamma0} + 2 * \text{sum}(\text{Gamma.dec})$, which is the asymptotic variance in the Markov chain CLT. Divide by <code>length(x)</code> to get the approximate variance of the sample mean of x .
<code>var.con</code>	the scalar $-\text{gamma0} + 2 * \text{sum}(\text{Gamma.con})$, which is the asymptotic variance in the Markov chain CLT. Divide by <code>length(x)</code> to get the approximate variance of the sample mean of x .

Bugs

Not precisely a bug, but `var.pos`, `var.dec`, and `var.con` can be negative. This happens only when the chain is way too short to estimate the variance, and even then rarely. But it does happen.

References

Geyer, C. J. (1992) Practical Markov Chain Monte Carlo. *Statistical Science* **7** 473–483.

See Also

[metrop](#)

Examples

```
n <- 2e4
rho <- 0.99
x <- arima.sim(model = list(ar = rho), n = n)
out <- initseq(x)
## Not run:
plot(seq(along = out$Gamma.pos) - 1, out$Gamma.pos,
      xlab = "k", ylab = expression(Gamma[k]), type = "l")
lines(seq(along = out$Gamma.dec) - 1, out$Gamma.dec, col = "red")
lines(seq(along = out$Gamma.con) - 1, out$Gamma.con, col = "blue")
## End(Not run)
# asymptotic 95% confidence interval for mean of x
mean(x) + c(-1, 1) * qnorm(0.975) * sqrt(out$var.con / length(x))
# estimated asymptotic variance
out$var.con
# theoretical asymptotic variance
(1 + rho) / (1 - rho) * 1 / (1 - rho^2)
# illustrating use with batch means
bm <- apply(matrix(x, nrow = 5), 2, mean)
initseq(bm)$var.con * 5
```

logit

Simulated logistic regression data.

Description

Like it says

Usage

```
data(logit)
```

Format

A data frame with variables

x1 quantitative predictor.

x2 quantitative predictor.

x3 quantitative predictor.

x4 quantitative predictor.

y Bernoulli response.

Examples

```
library(mcmc)
data(logit)
out <- glm(y ~ x1 + x2 + x3 + x4, family = binomial, data = logit)
summary(out)
```

metrop

Metropolis Algorithm

Description

Markov chain Monte Carlo for continuous random vector using a Metropolis algorithm.

Usage

```
metrop(obj, initial, nbatch, blen = 1, nspac = 1, scale = 1, outfun,
       debug = FALSE, ...)
```

Arguments

<code>obj</code>	an R function that evaluates the log unnormalized probability density of the desired equilibrium distribution of the Markov chain. First argument is the state vector of the Markov chain. Other arguments arbitrary and taken from the <code>...</code> arguments of this function. Should return <code>- Inf</code> for points of the state space having probability zero under the desired equilibrium distribution. Alternatively, an object of class <code>"metropolis"</code> from a previous run can be supplied, in which case any missing arguments (including the log unnormalized density function) are taken from this object (up until version 0.7-2 this was incorrect with respect to the <code>debug</code> argument, now it applies to it too).
<code>initial</code>	a real vector, the initial state of the Markov chain.
<code>nbatch</code>	the number of batches.
<code>blen</code>	the length of batches.
<code>nspac</code>	the spacing of iterations that contribute to batches.

scale	controls the proposal step size. If scalar or vector, the proposal is $x + \text{scale} * z$ where x is the current state and z is a standard normal random vector. If matrix, the proposal is $x + \text{scale} \%*\% z$.
outfun	controls the output. If a function, then the batch means of <code>outfun(state, ...)</code> are returned. If a numeric or logical vector, then the batch means of <code>state[outfun]</code> (if this makes sense). If missing, the the batch means of <code>state</code> are returned.
debug	if TRUE extra output useful for testing.
...	additional arguments for <code>obj</code> or <code>outfun</code> .

Details

Runs a “random-walk” Metropolis algorithm with multivariate normal proposal producing a Markov chain with equilibrium distribution having a specified unnormalized density. Distribution must be continuous. Support of the distribution is the support of the density specified by argument `obj`. The initial state must satisfy `obj(state, ...) > 0`. Description of a complete MCMC analysis (Bayesian logistic regression) using this function can be found in the vignette [demo](#) (`./doc/demo.pdf`).

Value

an object of class "mcmc", subclass "metropolis", which is a list containing at least the following components:

accept	fraction of Metropolis proposals accepted.
batch	<code>nbatch</code> by <code>p</code> matrix, the batch means, where <code>p</code> is the dimension of the result of <code>outfun</code> if <code>outfun</code> is a function, otherwise the dimension of <code>state[outfun]</code> if that makes sense, and the dimension of <code>state</code> when <code>outfun</code> is missing.
initial	value of argument <code>initial</code> .
final	final state of Markov chain.
initial.seed	value of <code>.Random.seed</code> before the run.
final.seed	value of <code>.Random.seed</code> after the run.
time	running time of Markov chain from <code>system.time()</code> .
lud	the function used to calculate log unnormalized density, either <code>obj</code> or <code>obj\$lud</code> from a previous run.
nbatch	the argument <code>nbatch</code> or <code>obj\$nbatch</code> .
blen	the argument <code>blen</code> or <code>obj\$blen</code> .
nspac	the argument <code>nspac</code> or <code>obj\$nspac</code> .
outfun	the argument <code>outfun</code> or <code>obj\$outfun</code> .

Description of additional output when `debug = TRUE` can be found in the vignette [debug](#) (`./doc/debug.pdf`).

Warning

If `outfun` is missing or not a function, then the log unnormalized density can be defined without a `...` argument and that works fine. One can define it starting `ludfun <- function(state)` and that works or `ludfun <- function(state, foo, bar)`, where `foo` and `bar` are supplied as additional arguments to `metrop`.

If `outfun` is a function, then both it and the log unnormalized density function can be defined without `...` arguments *if they have exactly the same arguments list* and that works fine. Otherwise it doesn't work. Start the definitions `ludfun <- function(state, foo)` and `outfun <- function(state, bar)` and you get an error about unused arguments. Instead start the definitions `ludfun <- function(state, foo, ...)` and `outfun <- function(state, bar, ...)`, supply `foo` and `bar` as additional arguments to `metrop`, and that works fine.

In short, the log unnormalized density function and `outfun` need to have `...` in their arguments list to be safe. Sometimes it works when `...` is left out and sometimes it doesn't.

Of course, one can avoid this whole issue by always defining the log unnormalized density function and `outfun` to have only one argument `state` and use global variables (objects in the R global environment) to specify any other information these functions need to use. That too follows the R way. But some people consider that bad programming practice.

Examples

```
h <- function(x) if (all(x >= 0) && sum(x) <= 1) return(1) else return(-Inf)
out <- metrop(h, rep(0, 5), 1000)
out$accept
# acceptance rate too low
out <- metrop(out, scale = 0.1)
out$accept
# acceptance rate o. k. (about 25 percent)
plot(out$batch[, 1])
# but run length too short (few excursions from end to end of range)
out <- metrop(out, nbatch = 1e4)
out$accept
plot(out$batch[, 1])
hist(out$batch[, 1])
```

Description

Variance of sample mean of time series calculated using overlapping batch means.

Usage

```
olbm(x, batch.length, demean = TRUE)
```

Arguments

`x` a matrix or time series object. Each column of `x` is treated as a scalar time series.

`batch.length` length of batches.

`demean` when `demean = TRUE` (the default) the sample mean is subtracted from each batch mean when estimating the variance. Using `demean = FALSE` would essentially assume the true mean is known to be zero, which might be useful in a toy problem where the answer is known.

Value

The estimated variance of the sample mean.

See Also

[ts](#)

Examples

```
h <- function(x) if (all(x >= 0) && sum(x) <= 1) return(1) else return(-Inf)
out <- metrop(h, rep(0, 5), 1000)
out <- metrop(out, scale = 0.1)
out <- metrop(out, nbatch = 1e4)
olbm(out$batch, 150)
# monte carlo estimates (true means are same by symmetry)
apply(out$batch, 1, mean)
# monte carlo standard errors (true s. d. are same by symmetry)
sqrt(diag(olbm(out$batch, 150)))
# check that batch length is reasonable
acf(out$batch, lag.max = 200)
```

temper

Simulated Tempering and Umbrella Sampling

Description

Markov chain Monte Carlo for continuous random vectors using parallel or serial simulated tempering, also called umbrella sampling. For serial tempering the state of the Markov chain is a pair (i, x) , where i is an integer between 1 and k and x is a vector of length p . This pair is represented as a single real vector $c(i, x)$. For parallel tempering the state of the Markov chain is vector of vectors (x_1, \dots, x_k) , where each x is of length p . This vector of vectors is represented as a $k \times p$ matrix.

Usage

```
temper(obj, initial, neighbors, nbatch, blen = 1, nspac = 1, scale = 1,
       outfun, debug = FALSE, parallel = FALSE, ...)
```

Arguments

<code>obj</code>	either an R function or an object of class "tempering" from a previous run. If a function, should evaluate the log unnormalized density $\log h(i, x)$ of the desired equilibrium distribution of the Markov chain for serial tempering (the same function is used for both serial and parallel tempering, see details below for further explanation). If an object, the log unnormalized density function is <code>obj\$lod</code> , and missing arguments of <code>temper</code> are obtained from the corresponding elements of <code>obj</code> . The first argument of the log unnormalized density function is the state for simulated tempering (i, x) is supplied as an R vector <code>c(i, x)</code> ; other arguments are arbitrary and taken from the <code>...</code> arguments of <code>temper</code> . The log unnormalized density function should return <code>-Inf</code> for points of the state space having probability zero.
<code>initial</code>	for serial tempering, a real vector <code>c(i, x)</code> as described above. For parallel tempering, a real $k \times p$ matrix as described above. In either case, the initial state of the Markov chain.
<code>neighbors</code>	a logical symmetric matrix of dimension k by k . Elements that are <code>TRUE</code> indicate jumps or swaps attempted by the Markov chain.
<code>nbatch</code>	the number of batches.
<code>blen</code>	the length of batches.
<code>nspac</code>	the spacing of iterations that contribute to batches.
<code>scale</code>	controls the proposal step size for real elements of the state vector. For serial tempering, proposing a new value for the x part of the state (i, x). For parallel tempering, proposing a new value for the x_i part of the state (x_1, \dots, x_k). In either case, the proposal is a real vector of length p . If scalar or vector, the proposal is $x + \text{scale} * z$ where x is the part x or x_i of the state the proposal may replace. If matrix, the proposal is $x + \text{scale} \%*\% z$. If list, the length must be k , and each element must be scalar, vector, or matrix, and operate as described above. The i -th component of the list is used to update x when the state is (i, x) or x_i otherwise.
<code>outfun</code>	controls the output. If a function, then the batch means of <code>outfun(state, ...)</code> are returned. The argument <code>state</code> is like the argument <code>initial</code> of this function. If missing, the batch means of the real part of the state vector or matrix are returned, and for serial tempering the batch means of a multivariate Bernoulli indicating the current component are returned.
<code>debug</code>	if <code>TRUE</code> extra output useful for testing.
<code>parallel</code>	if <code>TRUE</code> does parallel tempering, if <code>FALSE</code> does serial tempering.
<code>...</code>	additional arguments for <code>obj</code> or <code>outfun</code> .

Details

Serial tempering simulates a mixture of distributions of a continuous random vector. The number of components of the mixture is k , and the dimension of the random vector is p . Denote the state (i, x) , where i is a positive integer between 1 and k , and let $h(i, x)$ denote the unnormalized joint density of their equilibrium distribution. The logarithm of this function is what `obj` or `obj$lod`

calculates. The mixture distribution is the marginal for x derived from the equilibrium distribution $h(i, x)$, that is,

$$h(x) = \sum_{i=1}^k h(i, x)$$

Parallel tempering simulates a product of distributions of a continuous random vector. Denote the state (x_1, \dots, x_k) , then the unnormalized joint density of the equilibrium distribution is

$$h(x_1, \dots, x_k) = \prod_{i=1}^k h(i, x_i)$$

The update mechanism of the Markov chain combines two kinds of elementary updates: jump/swap updates (jump for serial tempering, swap for parallel tempering) and within-component updates. Each iteration of the Markov chain one of these elementary updates is done. With probability 1/2 a jump/swap update is done, and with probability 1/2 a with-component update is done.

Within-component updates are the same for both serial and parallel tempering. They are “random-walk” Metropolis updates with multivariate normal proposal, the proposal distribution being determined by the argument `scale`. In serial tempering, the x part of the current state (i, x) is updated preserving $h(i, x)$. In parallel tempering, an index i is chosen at random and the part of the state x_i representing that component is updated, again preserving $h(i, x)$.

Jump updates choose uniformly at random a neighbor of the current component: if i indexes the current component, then it chooses uniformly at random a j such that `neighbors[i, j] == TRUE`. It then does a Metropolis-Hastings update for changing the current state from (i, x) to (j, x) .

Swap updates choose a component uniformly at random and a neighbor of that component uniformly at random: first an index i is chosen uniformly at random between 1 and k , then an index j is chosen uniformly at random such that `neighbors[i, j] == TRUE`. It then does a Metropolis-Hastings update for swapping the states of the two components: interchanging x_i and x_j while perserving $h(x_1, \dots, x_k)$.

The initial state must satisfy `lud(initial, ...) > - Inf` for serial tempering or must satisfy `lud(initial[i,], ...) > - Inf` for each i for parallel tempering, where `lud` is either `obj` or `obj$lud`. That is, the initial state must have positive probability.

Value

an object of class "mcmc", subclass "tempering", which is a list containing at least the following components:

<code>batch</code>	the batch means of the continuous part of the state. If <code>outfun</code> is missing, an <code>nbatch</code> by <code>k</code> by <code>p</code> array. Otherwise, an <code>nbatch</code> by <code>m</code> matrix, where <code>m</code> is the length of the result of <code>outfun</code> .
<code>ibatch</code>	(returned for serial tempering only) an <code>nbatch</code> by <code>k</code> matrix giving batch means for the multivariate Bernoulli random vector that is all zeros except for a 1 in the i -th place when the current state is (i, x) .
<code>acceptx</code>	fraction of Metropolis within-component proposals accepted. A vector of length <code>k</code> giving the acceptance rate for each component.

<code>accepti</code>	fraction of Metropolis jump/swap proposals accepted. A k by k matrix giving the acceptance rate for each allowed jump or swap component. NA for elements such that the corresponding elements of <code>neighbors</code> is FALSE.
<code>initial</code>	value of argument <code>initial</code> .
<code>final</code>	final state of Markov chain.
<code>initial.seed</code>	value of <code>.Random.seed</code> before the run.
<code>final.seed</code>	value of <code>.Random.seed</code> after the run.
<code>time</code>	running time of Markov chain from <code>system.time()</code> .
<code>lud</code>	the function used to calculate log unnormalized density, either <code>obj</code> or <code>obj\$lud</code> from a previous run.
<code>nbatch</code>	the argument <code>nbatch</code> or <code>obj\$nbatch</code> .
<code>blen</code>	the argument <code>blen</code> or <code>obj\$blen</code> .
<code>nspac</code>	the argument <code>nspac</code> or <code>obj\$nspac</code> .
<code>outfun</code>	the argument <code>outfun</code> or <code>obj\$outfun</code> .

Description of additional output when `debug = TRUE` can be found in the vignette [debug \(. . /doc/debug.pdf\)](#).

Warning

If `outfun` is missing, then the log unnormalized density function can be defined without a `...` argument and that works fine. One can define it starting `ludfun <- function(state)` and that works or `ludfun <- function(state, foo, bar)`, where `foo` and `bar` are supplied as additional arguments to `temper` and that works too.

If `outfun` is a function, then both it and the log unnormalized density function can be defined without `...` arguments *if they have exactly the same arguments list* and that works fine. Otherwise it doesn't work. Start the definitions `ludfun <- function(state, foo)` and `outfun <- function(state, bar)` and you get an error about unused arguments. Instead start the definitions `ludfun <- function(state, foo, ...)` and `outfun <- function(state, bar, ...)`, supply `foo` and `bar` as additional arguments to `temper`, and that works fine.

In short, the log unnormalized density function and `outfun` need to have `...` in their arguments list to be safe. Sometimes it works when `...` is left out and sometimes it doesn't.

Of course, one can avoid this whole issue by always defining the log unnormalized density function and `outfun` to have only one argument `state` and use global variables (objects in the R global environment) to specify any other information these functions need to use. That too follows the R way. But some people consider that bad programming practice.

Examples

```
d <- 9
witch.which <- c(0.1, 0.3, 0.5, 0.7, 1.0)
ncomp <- length(witch.which)

neighbors <- matrix(FALSE, ncomp, ncomp)
neighbors[row(neighbors) == col(neighbors) + 1] <- TRUE
neighbors[row(neighbors) == col(neighbors) - 1] <- TRUE
```

```
ludfun <- function(state, log.pseudo.prior = rep(0, ncomp)) {
  stopifnot(is.numeric(state))
  stopifnot(length(state) == d + 1)
  icoomp <- state[1]
  stopifnot(icoomp == as.integer(icoomp))
  stopifnot(1 <= icoomp && icoomp <= ncomp)
  stopifnot(is.numeric(log.pseudo.prior))
  stopifnot(length(log.pseudo.prior) == ncomp)
  theta <- state[-1]
  if (any(theta > 1.0)) return(-Inf)
  bnd <- witch.which[icoomp]
  lpp <- log.pseudo.prior[icoomp]
  if (any(theta > bnd)) return(lpp)
  return(- d * log(bnd) + lpp)
}

# parallel tempering
thetas <- matrix(0.5, ncomp, d)
out <- temper(ludfun, initial = thetas, neighbors = neighbors, nbatch = 20,
  blen = 10, nspac = 5, scale = 0.56789, parallel = TRUE, debug = TRUE)

# serial tempering
theta.initial <- c(1, rep(0.5, d))
# log pseudo prior found by trial and error
qux <- c(0, 9.179, 13.73, 16.71, 20.56)

out <- temper(ludfun, initial = theta.initial, neighbors = neighbors,
  nbatch = 50, blen = 30, nspac = 2, scale = 0.56789,
  parallel = FALSE, debug = FALSE, log.pseudo.prior = qux)
```

Index

*Topic **datasets**

foo, 1

logit, 4

*Topic **misc**

metrop, 5

temper, 8

*Topic **ts**

initseq, 2

olbm, 7

foo, 1

initseq, 2

logit, 4

metrop, 3, 5

olbm, 7

temper, 8

ts, 7