Description

linear.mh: Function which produces one or multiple Metropolis-Hastings chains for normal linear regression cases of the type

$$y = a + bx + \varepsilon \text{, with } \varepsilon \sim N(0, \sigma^2)$$

with prior specification $a \mid \sigma^2 \sim N(0, k\sigma^2)$, $b \mid \sigma^2 \sim N(0, k\sigma^2)$ and $\sigma^{-2} \sim G(a_0, b_0)$ and random walk proposals of the form

$$a^t \sim N(a^{t-1}, \sigma_a^2)$$

$$b^t \mid a^t \sim N\left(b^{t-1} + \rho \frac{\sigma_b}{\sigma_a}(a^t - a^{t-1}), \sigma_b^2(1-\rho^2) \right)$$

$$\log \sigma^t \sim N\left(\log \sigma^{t-1}, \sigma_\sigma^2\right) \text{ or } \sigma^t \sim G(\sigma^{(t-1)} b_{prop}, b_{prop})$$

Usage

linear.mh(y, x, n, l=1, discard=n/2, initial.matrix, select='normal', sd, cor=0,
gamma.parameter=5, hyper.par=c(0.001,0.001,1000), deflate=1, plot=TRUE)

Arguments

y: the dependent variable vector.

x: the independent variable vector.

n: the number of Metropolis-Hastings iterations.

l: the number of Metropolis-Hastings chains to be produced, the default is 1 chain.

discard: the number of discarded iterations for the "burn-in" period. The default value n/2 discards the first half of the chain. This arguments must *always* be smaller than *n*.

initial.matrix: A matrix which contains the initial values for the simulation. This argument must be given in matrix form. If l=1 (one chain) the initial.matrix must be of dimension 1×3 with column one containing the initial value for parameter *a*, column two the initial value for parameter *b* and column three the initial value for parameter σ (not σ^2). If l>1 i.e. l=4 then the initial matrix must be of dimension 4×3. In this case the 1st column must contain 4 initial values for parameter *a*, the 2nd column must contain 4 initial values for parameter *b* and the 3rd column must contain 4 initial values for parameter σ.

select: A character related to the type of proposal distribution for parameter σ. If select='normal' then a normal proposal distribution is used, if select='gamma' then a gamma proposal distribution is used. The default is 'normal'.

sd: A vector with 3 elements corresponding to the proposal's *standard deviations* σ_a, σ_b and σ_σ for parameters *a* (1st element), *b* (2nd element) and σ (3rd element). If select='gamma' then the third element is ignored.

cor: The correlation ρ between parameters *a*, *b*. The default value is 0.

gamma.pararameter: Parameter b_{prop} of the gamma proposal distribution for parameter σ if select='gamma'. The default value is 5.

hyper.par: a vector of size 3 corresponding to the values the values of the hyper parameters a_0, b_0 and *k*. The default values are $a_0 = b_0 = 0.001$ and $k = 1000$.

deflate: Provides the option to increase or decrease (usually decrease) the scale of the proposals when acceptance ratio is lower than 20% at iteration n/2. If for example deflate=10 then a normal proposal's scale will by divided with 10 and parameter b_{prop} of the gamma proposal will be multiplied with 10. The default value is 1.

plot: A logical argument. If TRUE and l=1 (one chain) then time series plots, autocorrelation plots and histograms for the draws of parameters *a*, *b* and σ are returned. If TRUE and l>1 (multilple chains) then ergodic mean plots and histograms for the draws of parameters *a*, *b* and σ are returned.

Components

linear.mh returns the following components:

parameters: The draws of parameters from the posterior distribution.

acceptance_ratio: The acceptance ratio of each chain.

R_root: The calculated R reduction measure (returned only if *l*>1).

lengths: The size of the MCMC sample kept for inference.

means: The posterior means of the parameters.

standard.deviations: The posterior standard deviations of the parameters.

correlations: The posterior correlation matrix of the parameters.

quantiles: The posterior quantiles of the parameters.