Description

linear.gibbs: Function which produces one or multiple Gibbs chains for normal linear regression cases of the type

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

with prior specification

$$a \mid \sigma^2 \sim N(0, k\sigma^2), \quad b \mid \sigma^2 \sim N(0, k\sigma^2) \quad \text{and} \quad \sigma^{-2} \sim G(a_0, b_0)$$
$$\text{or}$$
$$a \sim N(0, k), \quad b \sim N(0, k) \quad \text{and} \quad \sigma^{-2} \sim G(a_0, b_0)$$

Usage

linear.gibbs(y, x, n, l=1, discard=n/2, initial.matrix, hyper.par=c(0.001,0.001,1000),
          model='conj', plot=TRUE)

Arguments

y: the dependent variable vector.

x: the independent variable vector.

n: the number of Gibbs iterations.

l: the number of Gibbs chains to be produced, the default choice 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.

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

model: A character related to the prior specification for parameters a and b. The options are 'conj' or 'semiconj'. If 'conj' then $a\,|\,\sigma^2 \sim N(0, k\sigma^2)$ and $b\,|\,\sigma^2 \sim N(0, k\sigma^2)$, if 'semiconj' then $a \sim N(0, k)$ and $b \sim N(0, k)$.

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.gibbs returns the following components:

parameters: The draws of parameters from the posterior distribution.

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.