Description

gibbs.lm: This function produces one or multiple Gibbs chains for normal linear regression models. The function works for models with zero, one or two explanatory variables:

$$M_0 : y = a + \varepsilon$$
$$M_1 : y = a + bx_1 + \varepsilon$$
$$M_2 : y = a + b_1 x_1 + b_2 x_2 + \varepsilon$$

with $\varepsilon \sim N(0, \sigma^2)$ and prior specification under each model

M_0: $a \sim N(0, k\left(\mathbf{X}_0\mathbf{X}_0^T\right)^{-1}\sigma^2)$, $\sigma^2 \sim IG(A, B)$

M_1: $\left(a, b_1\right)^T \sim N_2(\mathbf{0}, k\left(\mathbf{X}_1\mathbf{X}_1^T\right)^{-1}\sigma^2)$, $\sigma^2 \sim IG(A, B)$

M_2: $\left(a, b_1, b_2\right)^T \sim N_3(\mathbf{0}, k\left(\mathbf{X}_2\mathbf{X}_2^T\right)^{-1}\sigma^2)$, $\sigma^2 \sim IG(A, B)$

$\mathbf{X}_j$ is the corresponding design matrix for each model for $j = 0,1,2$.

Usage

```
gibbs.lm(y, x, n, l=1, discard=n/2, initial.matrix, hyper.par=c(0.001,0.001),
        prior.var='square', plot=TRUE)
```

Arguments

y: the dependent variable vector.

x: the independent variable vector or matrix. This argument can be a vector (model with 1 explanatory variable), a matrix (model with 2 explanatory variables each corresponding to a column) or a character set equal to ‘none’ (for the null model).

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. The dimension of the matrix depends on the number of Gibbs chains and on the number of parameters. The rows must equal the number of chains (argument *l*) while the columns must equal the number of parameters. For example, for a regression model with parameters (a, b_1, σ) and 5 parallel Gibbs chains, the dimension must be 5×3.

hyper.par: a vector of size 2 corresponding to the values of the hyper parameters A and B. The default values are $A=B=0.001$.

prior.var: A character related to the multiplying parameter *k*. The options are 'square', 'simple' or 'regressors'.

If set to 'square' then $k=(sample.size)^2$

If set to 'simple' then $k=(sample.size)$

If set to 'regressors' then $k=(number.of.regressor.parameters)^2$

The last case refers *only* to regressor parameters and *not* to the constant term, therefore *k* cannot be set to 'regressors' for a model with no independent variables (null model). The default choice is 'square'.

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

gibbs.lm 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.