

Package ‘mombf’

April 1, 2022

Type Package

Version 3.1.3

Date 2022-04-01

Title Bayesian Model Selection and Averaging for Non-Local and Local Priors

Author David Rossell, John D. Cook, Donatello Telesca, P. Roebuck, Oriol Abril, Miquel Torrens

Maintainer David Rossell <rosselldavid@gmail.com>

Depends R (>= 2.14.0), methods, mvtnorm, ncvreg, mgcv

Imports Rcpp (>= 0.12.16), glmnet, mclust, sgd, survival, pracma

Suggests parallel, testthat, patrick

LinkingTo Rcpp, RcppArmadillo

Description Bayesian model selection and averaging for regression and mixtures for non-local and selected local priors.

License GPL (>= 2) | file LICENSE

URL <https://github.com/davidrusi/mombf>

BugReports <https://github.com/davidrusi/mombf/issues>

LazyLoad yes

Collate AllClasses.R AllGenerics.R alapl.R bms_ortho.R cil.R cox.R derivatives_nlps.R distribs.R dmom.R emomLM.R gam.R greedyGLM.R infocriteria.R initParameters.R msPriorSpec.R modelsearch.R modelSelection.R modelSelectionGLM.R mombf.R normaliwish.R normmix.R nlpMarginal.R pmomLM.R pmomPM.R postMode.R pplProbit.R pplmodel.R rmom.R RcppExports.R testfunction.R zellnerLM.R

NeedsCompilation yes

Repository CRAN

Date/Publication 2022-04-01 09:10:03 UTC

R topics documented:

bbPrior	2
bfnormmix	4
cil	6
dalapl	9
ddir	10
diwish	11
dmom	12
dpostNIW	14
eprod	15
getBIC	16
hald	17
marginalNIW	18
mixturebf-class	20
modelSelection	21
mombf	26
momknown	28
msfit-class	30
msPriorSpec-class	32
nlpmarginals	35
pmomLM	37
postModeOrtho	41
postProb	44
postSamples	45
priorp2g	45
rmlp	47
Index	49

 bbPrior

Priors on model space for variable selection problems

Description

unifPrior implements a uniform prior (equal a priori probability for all models). binomPrior implements a Binomial prior. bbPrior implements a Beta-Binomial prior.

Usage

```
unifPrior(sel, logscale=TRUE, groups=1:length(sel),
constraints=lapply(1:length(unique(groups)), function(z) integer(0)))
```

```
binomPrior(sel, prob=.5, logscale=TRUE, probconstr=prob, groups=1:length(sel),
constraints=lapply(1:length(unique(groups)), function(z) integer(0)))
```

```
bbPrior(sel, alpha=1, beta=1, logscale=TRUE, alphaconstr=alpha,
betaconstr=beta, groups=1:length(sel),
constraints=lapply(1:length(unique(groups)), function(z) integer(0)))
```

Arguments

sel	Logical vector indicating which variables are included in the model
logscale	Set to TRUE to return the log-prior probability.
groups	Group that each variable belongs to (e.g. dummy indicators for categorical variables with >2 categories). The idea is that all variables in a group are jointly added/removed from the model. By default all variables are assumed to be in separate groups
constraints	List with length equal to the number of groups (distinct elements in groups). Element j in the list should indicate any hierarchical constraints on the group, for instance constraints[[3]]=c(1,2) indicates that group 3 can only be included in the model if groups 1 and 2 are also in the model. This can be used to enforce that an interaction can only be in the model if the main effects are also in the model.
prob	Success probability for the Binomial prior
probconstr	Success probability for the Binomial prior for groups that are subject to constraints
alpha	First parameter of the Beta-Binomial prior, which is equivalent to specifying a Beta(alpha,beta) prior on prob.
beta	First parameter of the Beta-Binomial prior, which is equivalent to specifying a Beta(alpha,beta) prior on prob.
alphaconstr	Same as alpha for the groups that are subject to constraints
betaconstr	Same as beta for the groups that are subject to constraints

Value

Prior probability of the specified model

Author(s)

David Rossell

Examples

```
library(mombf)
sel <- c(TRUE, TRUE, FALSE, FALSE)
unifPrior(sel, logscale=FALSE)
binomPrior(sel, prob=.5, logscale=FALSE)
bbPrior(sel, alpha=1, beta=1, logscale=FALSE)
```

bfnormmix	<i>Number of Normal mixture components under Normal-IW and Non-local priors</i>
-----------	---------------------------------------------------------------------------------

Description

Posterior sampling and Bayesian model selection to choose the number of components k in multivariate Normal mixtures.

bfnormmix computes posterior probabilities under non-local MOM-IW-Dir(q) priors, and also for local Normal-IW-Dir(q ,niw) priors. It also computes posterior probabilities on cluster occupancy and posterior samples on the model parameters for several k .

Usage

```
bfnormmix(x, k=1:2, mu0=rep(0,ncol(x)), g, nu0, S0, q=3, q.niw=1,
          B=10^4, burnin= round(B/10), logscale=TRUE, returndraws=TRUE, verbose=TRUE)
```

Arguments

<code>x</code>	$n \times p$ input data matrix
<code>k</code>	Number of components
<code>mu0</code>	Prior on $\mu[j]$ is $N(\mu_0, g \Sigma[j])$
<code>g</code>	Prior on $\mu[j]$ is $N(\mu_0, g \Sigma[j])$. This is a critical MOM-IW prior parameter that specifies the separation between components deemed practically relevant. It defaults to assigning 0.95 prior probability to any pair of μ 's giving a bimodal mixture, see details
<code>S0</code>	Prior on $\Sigma[j]$ is $IW(\Sigma_j; \nu_0, S_0)$
<code>nu0</code>	Prior on $\Sigma[j]$ is $IW(\Sigma_j; \nu_0, S_0)$
<code>q</code>	Prior parameter in MOM-IW-Dir(q) prior
<code>q.niw</code>	Prior parameter in Normal-IW-Dir(q ,niw) prior
<code>B</code>	Number of MCMC iterations
<code>burnin</code>	Number of burn-in iterations
<code>logscale</code>	If set to TRUE then log-Bayes factors are returned
<code>returndraws</code>	If set to TRUE the MCMC posterior draws under the Normal-IW-Dir prior are returned for all k
<code>verbose</code>	Set to TRUE to print iteration progress

Details

The likelihood is

$$p(x[i,] | \mu, \Sigma, \eta) = \sum_j \eta_j N(x[i,]; \mu_j, \Sigma_j)$$

The Normal-IW-Dir prior is

$\text{Dir}(\eta; q, \text{niw}) \prod_j N(\mu_j; \mu_0, g \Sigma) \text{IW}(\Sigma_j; \nu_0, S_0)$

The MOM-IW-Dir prior is

$d(\mu, \Sigma) \text{Dir}(\eta; q) \prod_j N(\mu_j; \mu_0, g \Sigma) \text{IW}(\Sigma_j; \nu_0, S_0)$

where

$d(\mu, \Sigma) = [\prod_j <1 (\mu_j - \mu_l)' A (\mu_j - \mu_l)]$

and A is the average of $\Sigma_1^{-1}, \dots, \Sigma_k^{-1}$. Note that one must have $q > 1$ for the MOM-IW-Dir to define a non-local prior.

By default the prior parameter g is set such that

$P((\mu[j] - \mu[l])' A (\mu[j] - \mu[l]) < 4) = 0.05$.

The rationale when $\Sigma[j] = \Sigma[l]$ and $\eta[j] = \eta[l]$ then $(\mu[j] - \mu[l])' A (\mu[j] - \mu[l]) > 4$ corresponds to a bimodal density. That is, the default g focuses 0.95 prior prob on a degree of separation between components giving rise to a bimodal mixture density.

`bfnormmix` computes posterior model probabilities under the MOM-IW-Dir and Normal-IW-Dir priors using MCMC output. As described in Fuquene, Steel and Rossell (2018) the estimate is based on the posterior probability that one cluster is empty under each possible k .

Value

A list with elements

<code>k</code>	Number of components
<code>pp.momiw</code>	Posterior probability of k components under a MOM-IW-Dir(q) prior
<code>pp.niw</code>	Posterior probability of k components under a Normal-IW-Dir(q, niw) prior
<code>probempty</code>	Posterior probability that any one cluster is empty under a MOM-IW-Dir(q, niw) prior
<code>bf.momiw</code>	Bayes factor comparing 1 vs k components under a MOM-IW-Dir(q) prior
<code>logpen</code>	log of the posterior mean of the MOM-IW-Dir(q) penalty term
<code>logbf.niw</code>	Bayes factor comparing 1 vs k components under a Normal-IW-Dir(q, niw) prior

Author(s)

David Rossell

References

Fuquene J., Steel M.F.J., Rossell D. On choosing mixture components via non-local priors. 2018. arXiv

Examples

```
x <- matrix(rnorm(100*2), ncol=2)
```

```
bfnormmix(x=x, k=1:3)
```

`cil` *Treatment effect estimation for linear models via Confounder Importance Learning using non-local priors.*

Description

Treatment effect estimation for linear models in the presence of multiple treatments and a potentially high-dimensional number of controls, i.e. $p \gg n$ can be handled.

Confounder Importance Learning (CIL) proposes an estimation framework where the importance of the relationship between treatments and controls is factored in into the establishment of prior inclusion probabilities for each of these controls on the response model. This is combined with the use of non-local priors to obtain BMA estimates and posterior model probabilities.

`cil` is built on `modelSelection` and produces objects of type `cilfit`. Use `coef` and `postProb` to obtain treatment effect point estimates and posterior model probabilities, respectively, on this object class.

Usage

```
cil(y, D, X, I = NULL, R = 1e4, th.search = 'EB', mod1 = 'lasso_bic',
    th.prior = 'unif', beta.prior = 'nlp', rho.min = NULL, rho.max = NULL,
    th.range = NULL, tau = 0.348, max.mod = Inf, lpen = 'lambda.1se',
    eps = 1e-10, bvs.fit0 = NULL, th.EP = NULL)
```

Arguments

<code>y</code>	one-column matrix containing the observed responses. The response must be continuous (currently the only type supported)
<code>D</code>	treatment matrix with numeric columns, continuous or discrete. Any finite number of treatments are supported. If only one treatment is provided, supply this object in the same format used for <code>y</code>
<code>X</code>	matrix of controls with numeric columns, continuous or discrete. If only one treatment is provided, supply this object in the same format used for <code>y</code>
<code>I</code>	matrix with the desired interaction terms between <code>D</code> and <code>X</code> . If not informed, i.e. supplied as the default <code>NULL</code> , this term will not be included into the response model
<code>R</code>	integer scalar indicating the number of Gibbs sampling iterations to be run by <code>modelSelection</code> on each stage of CIL (see argument <code>niter</code> therein)
<code>th.search</code>	method to estimate theta values in the marginal prior inclusion probabilities of the CIL model. Options are: EB (Empirical Bayes, based on maximum marginal likelihood) and EP (Expectation propagation approximation)
<code>mod1</code>	method to estimate the feature parameters corresponding to the influence of the controls on the treatments. Supported values for this argument are 'ginv' (generalised pseudo-inverse), lasso (see argument <code>lpen</code>), lasso_bic (default), and ridge

<code>th.prior</code>	prior associated to the thetas for the Empirical Bayes estimation. Currently only <code>unif</code> (Uniform prior) is supported, effectively making the EB approach the maximisation of the marginal likelihood
<code>beta.prior</code>	prior on the response model parameters, for both treatments and covariates, used to compute posterior model probabilities and BMA estimates of the CIL model. Options are: <code>nlp</code> (pMOM non-local prior, by default, see also <code>tau</code> argument) and <code>zup</code> (Zellner's prior with $\tau = n$, equivalent to the Unit Information prior)
<code>rho.min</code>	value of ρ in $(0, 1/2)$ employed in the prior probability model of CIL. If left uninformed, i.e. supplied as the default <code>NULL</code> , it will be set to $1/p^2$, where p is the dimension of the response model.
<code>rho.max</code>	this argument is deprecated and should not be informed.
<code>th.range</code>	sequence of values to be considered in the grid when searching for points to initialise the search for the optimal theta parameters. If left uninformed, the function will determine a computationally suitable grid depending on the number of parameters to be estimated
<code>tau</code>	scalar hyper-parameter supplied to the non-local prior (see argument <code>beta.prior</code>). Set to 0.348 by default
<code>max.mod</code>	maximum number of models considered when computing the marginal quantities required to find the optimal values for theta. By default, this argument is equal to <code>Inf</code> (i.e. all visited models by the Gibbs sampler are considered), but it might be computationally desirable to restrict this number when the dimensionality of D and/or X is very large
<code>lpen</code>	penalty type supplied to <code>glmnet</code> if <code>mod1</code> is set to <code>lasso</code> . Default is <code>lambda.1se</code> (see documentation corresponding to <code>glmnet</code> for options on how to set this parameter)
<code>eps</code>	small scalar used to avoid round-offs to absolute zeroes or ones in marginal prior inclusion probabilities.
<code>bvs.fit0</code>	object returned by <code>modelSelection</code> under $\theta = 0$, used as a model exploration tool to compute EB approximation on the thetas. This argument is only supposed to be used in case of a second computation the model on the same data where <code>th.search</code> has ben changed to <code>EB</code> , in order to avoid repeating the computation of the initial <code>modelSelection</code> fit. To use this argument, supply the object residing in the slot <code>init.msfit</code> of a <code>cilfit</code> -class object.
<code>th.EP</code>	Optimal theta values under the EP approximation, obtained in a previous CIL run. This argument is only supposed to be used in case of a second computation the model on the same data where <code>th.search</code> has ben changed to <code>EB</code> , in order to save the cost of the EP search to initialise the optimisation algorithm. To use this argument, supply the object residing int the slot <code>th.hat</code> of a <code>cilfit</code> -class object.

Details

We estimate treatment effects for the features present in the treatment matrix D . Features in X , which may or may not be causal factors of the treatments of interest, only act as controls and, therefore, are not used as inferential subjects.

Confounder importance learning is a flexible treatment effect estimation framework that essentially determines how the role of the influence of X on D should affect their relationship with the response, through establishing prior inclusion probabilities on the response model for y according to said role. This is regulated through a hyper-parameter θ that is set according to the method supplied to `th.search`. While the EB option obtains a more precise estimate a priori, the EP alternative achieves a reasonable approximation at a fraction of the computational cost.

See references for further details on implementation and computation.

Value

Object of class `cilfit`, which extends a list with elements

<code>cil.teff</code>	numeric vector containing the BMA point estimates for the treatment effects featured in D
<code>cil.teff.postdist</code>	numeric vector featuring quantiles of the posterior distribution for each of the treatment effects, obtained via Gibbs sampling
<code>cil.bma.mcmc</code>	numeric vector containing the BMA point estimate at each MCMC run on the model search Gibbs sampler
<code>model.postprobs</code>	matrix returning the posterior model probabilities computed in the CIL model
<code>marg.postprobs</code>	numeric vector containing the estimated marginal posterior inclusion probabilities of the featured treatments and controls
<code>theta.hat</code>	Values used for the hyper-parameter θ , estimated according to the argument <code>th.search</code> specified
<code>treat.coefs</code>	Estimated weights of the effect of the control variables on each of the treatments, as estimated with the method specified in argument <code>mod1</code>
<code>msfit</code>	Object returned by <code>modelSelection</code> (of class <code>msfit</code>) of the final model estimated by CIL.
<code>theta.EP</code>	Estimated values of θ using the EP algorithm. It coincides with <code>theta.hat</code> if the argument <code>th.search</code> is set to EB
<code>init.msfit</code>	Initial <code>msfit</code> object used to estimate the initial model where all elements in θ are set to zero (used in the optimisation process of this hyper-parameter)

Author(s)

Miquel Torrens

References

Torrens i Dinares M., Papaspiliopoulos O., Rossell D. Confounder importance learning for treatment effect inference. <https://arxiv.org/abs/2110.00314>, 2021, 1–48.

See Also

[postProb](#) to obtain posterior model probabilities.
[coef](#) for inference on the treatment parameters.

Examples

```

# Simulate data
set.seed(1)
X <- matrix(rnorm(100 * 50), nrow = 100, ncol = 50)
beta_y <- matrix(c(rep(1, 6), rep(0, 44)), ncol = 1)
beta_d <- matrix(c(rep(1, 6), rep(0, 44)), ncol = 1)
alpha <- 1
d <- X %*% beta_d + rnorm(100)
y <- d * alpha + X %*% beta_y + rnorm(100)

# Confounder Importance Learning
fit1 <- cil(y = y, D = d, X = X, th.search = 'EP')

# Posterior model probabilities (comparison)
postProb(fit1, nmax = 5)

# Treatment effects
coef(fit1)

```

dalapl

*Density and random draws from the asymmetric Laplace distribution***Description**

dalapl evaluates the probability density function, palapl the cumulative probability function and ralapl generates random draws.

Usage

```

dalapl(x, th=0, scale=1, alpha=0, logscale=FALSE)

palapl(x, th=0, scale=1, alpha=0)

ralapl(n, th=0, scale=1, alpha=0)

```

Arguments

x	Vector of values at which to evaluate the pdf/cdf
n	Number of random draws
th	Location parameter (mode)
scale	Scale parameter (proportional to variance)
alpha	Asymmetry parameter, must be between -1 and 1
logscale	If TRUE the log-pdf is returned

Details

For $x \leq th$ the asymmetric Laplace pdf is
 $0.5 * \exp(-\text{abs}(th-x) / (\text{sqrt}(\text{scale}) * (1+\alpha))) / \text{sqrt}(\text{scale})$
 and for $x > th$ it is
 $0.5 * \exp(-\text{abs}(th-x) / (\text{sqrt}(\text{scale}) * (1-\alpha))) / \text{sqrt}(\text{scale})$

Value

dalapl returns the density function, palapl the cumulative probability, ralapl random draws.

Author(s)

David Rossell

Examples

```
library(mombf)
e <- ralapl(n=10^4, th=1, scale=2, alpha=0.5)
thseq <- seq(min(e), max(e), length=1000)
hist(e, main='', breaks=30, prob=TRUE)
lines(thseq, dalapl(thseq, th=1, scale=2, alpha=0.5), col=2)
```

 ddir

Dirichlet density

Description

Evaluate the density of a Dirichlet distribution

Usage

```
ddir(x, q, logscale=TRUE)
```

Arguments

x	Vector or matrix containing the value at which to evaluate the density. If a matrix, the density is evaluated for each row. Rows are renormalized to ensure they add up to 1
q	Dirichlet parameters. Must have the same length as ncol(x), or length 1 (in which case a symmetric Dirichlet density is valued)
logscale	For logscale==TRUE, dimom returns the natural log of the prior density

Value

Density of a Dirichlet(q) distribution evaluated at each row of x

Author(s)

David Rossell

Examples

```
library(mombf)
x= matrix(c(1/3,2/3, .5, .5),nrow=2,byrow=TRUE)
ddir(x,q=2)
```

diwish

Density for Inverse Wishart distribution

Description

diwish returns the density for the inverse Wishart(ν, S) evaluated at Sigma.

Usage

```
diwish(Sigma, nu, S, logscale=FALSE)
```

Arguments

Sigma	Positive-definite matrix
nu	Degrees of freedom of the inverse Wishart
S	Scale matrix of the inverse Wishart
logscale	If logscale==TRUE the log-density is returned

Value

Inverse Wishart(ν, S) density evaluated at Sigma

Author(s)

David Rossell

See Also

[dpostNIW](#) for the Normal-IW posterior density

Examples

```
library(mombf)
Sigma= matrix(c(2,1,1,2),nrow=2)
diwish(Sigma,nu=4,S=diag(2))
```

dmom

*Non-local prior density, cdf and quantile functions.***Description**

dmom, dimom and demom return the density for the moment, inverse moment and exponential moment priors. pmom, pimom and pemom return the distribution function for the univariate moment, inverse moment and exponential moment priors (respectively). qmom and qimom return the quantiles for the univariate moment and inverse moment priors. dmomigmarg returns the marginal density implied by a $MOM(x;\tau*\phi)*Invgamma(\phi;a/2,b/2)$, pmomigmarg its cdf. Analogously demomigmarg and demomigmarg for $eMOM(x;\tau*\phi)*Invgamma(\phi;a/2,b/2)$

Usage

```
dmom(x, tau, a.tau, b.tau, phi=1, r=1, V1, baseDensity='normal', nu=3,
logscale=FALSE, penalty='product')
dimom(x, tau=1, phi=1, V1, logscale=FALSE, penalty='product')
demom(x, tau, a.tau, b.tau, phi=1, logscale=FALSE)

pmom(q, V1 = 1, tau = 1)
pimom(q, V1 = 1, tau = 1, nu = 1)
pemom(q, tau, a.tau, b.tau)

qmom(p, V1 = 1, tau = 1)
qimom(p, V1 = 1, tau = 1, nu = 1)

dmomigmarg(x, tau, a, b, logscale=FALSE)
pmomigmarg(x, tau, a, b)

demomigmarg(x, tau, a, b, logscale=FALSE)
pemomigmarg(x, tau, a, b)
```

Arguments

x	In the univariate setting, x is a vector with the values at which to evaluate the density. In the multivariate setting it is a matrix with an observation in each row.
q	Vector of quantiles.
p	Vector of probabilities.
V1	Scale matrix (ignored if penalty=='product'). Defaults to 1 in univariate setting and the identity matrix in the multivariate setting.
tau	Prior dispersion parameter is $\tau*\phi$. See details.
a.tau	If tau is left missing, an Inverse Gamma($a.tau/2, b.tau/2$) is placed on tau. In this case dmom and demom return the density marginalized with respect to tau.
b.tau	See a.tau.
phi	Prior dispersion parameter is $\tau*\phi$. See details.

r	Prior power parameter for MOM prior is $2*r$
baseDensity	For baseDensity=='normal' a Normal MOM prior is used, for baseDensity=='laplace' a Laplace MOM prior, for baseDensity=='t' a T MOM prior with nu degrees of freedom is used.
nu	Prior parameter indicating the degrees of freedom for the quadratic T MOM and iMOM prior densities. The tails of the inverse moment prior are proportional to the tails of a multivariate T with nu degrees of freedom.
penalty	penalty=='product' indicates that product MOM/iMOM should be used. penalty=='quadratic' indicates quadratic iMOM. See Details.
logscale	For logscale==TRUE, dmom returns the natural log of the prior density.
a	The marginal prior on phi is $IG(a/2, b/2)$
b	The marginal prior on phi is $IG(a/2, b/2)$

Details

For type=='quadratic' the density is as follows. Define the quadratic form $q(\theta) = (\theta - \theta_0)' * \text{solve}(V_1) * (\theta - \theta_0) / (\tau * \phi)$. The normal moment prior density is proportional to $q(\theta) * \text{dmvnorm}(\theta, \theta_0, \tau * \phi * V_1)$. The T moment prior is proportional to $q(\theta) * \text{dmvt}(\theta, \theta_0, \tau * \phi * V_1, \nu)$. The inverse moment prior density is proportional to $q(\theta)^{-(\nu+d)/2} * \exp(-1/q(\theta))$.

pmom, pimom and qimom use closed-form expressions, while qmom uses nlminb to find quantiles numerically. Only the univariate version is implemented. In this case the product MOM is equivalent to the quadratic MOM. The same happens for the iMOM.

dmomigmarg returns the marginal density

$p(x) = \int \text{MOM}(x; 0, \tau * \phi) \text{IG}(\phi; a/2, b/2) d\phi$

Value

Prior density, cumulative distribution function or quantile.

Author(s)

David Rossell

References

Johnson V.E., Rossell D. Non-Local Prior Densities for Default Bayesian Hypothesis Tests. Journal of the Royal Statistical Society B, 2010, 72, 143-170.

Johnson V.E., Rossell D. Bayesian model selection in high-dimensional settings. Journal of the American Statistical Association, 2012, 107, 649-660

See <http://rosselldavid.googlepages.com> for technical reports.

Examples

```
#evaluate and plot the moment and inverse moment priors
library(mombf)
tau <- 1
thseq <- seq(-3,3,length=1000)
plot(thseq,dmom(thseq,tau=tau),type='l',ylab='Prior density')
lines(thseq,dimom(thseq,tau=tau),lty=2,col=2)
```

dpostNIW

*Posterior Normal-IWishart density***Description**

dpostNIW evaluates the posterior Normal-IWishart density at (μ, Σ) . rpostNIW draws independent samples. This posterior corresponds to a Normal model for the data

$$x[i,] \sim N(\mu, \Sigma) \text{ iid } i=1, \dots, n$$

under conjugate priors

$$\mu \mid \Sigma \sim N(\mu_0, g \Sigma) \quad \Sigma \sim IW(\nu_0, S_0)$$
Usage

```
dpostNIW(mu, Sigma, x, g=1, mu0=rep(0,length(mu)), nu0=nrow(Sigma)+1, S0,
  logscale=FALSE)
```

```
rpostNIW(n, x, g=1, mu0=0, nu0, S0, precision=FALSE)
```

Arguments

mu	Vector of length p
Sigma	p x p positive-definite covariance matrix
x	n x p data matrix (individuals in rows, variables in columns)
g	Prior dispersion parameter for mu
mu0	Prior mean for mu
nu0	Prior degrees of freedom for Sigma
S0	Prior scale matrix for Sigma, by default set to I/nu0
logscale	set to TRUE to get the log-posterior density
n	Number of samples to draw
precision	If set to TRUE, samples from the precision matrix (inverse of Sigma) are returned instead

Value

dpostNIW returns the Normal-IW posterior density evaluated at (mu,Sigma).

rpostNIW returns a list with two elements. The first element are posterior draws for the mean. The second element are posterior draws for the covariance (or its inverse if precision==TRUE). Only lower-diagonal elements are returned (Sigma[lower.tri(Sigma,diag=TRUE)]).

Author(s)

David Rossell

See Also

[diwish](#) for the inverse Wishart prior density, [marginalNIW](#) for the integrated likelihood under a Normal-IW prior

Examples

```
#Simulate data
x= matrix(rnorm(100),ncol=2)
#Evaluate posterior at data-generating truth
mu= c(0,0)
Sigma= diag(2)
dpostNIW(mu,Sigma,x=x,g=1,nu0=4,log=FALSE)
```

eprod

Expectation of a product of powers of Normal or T random variables

Description

Compute the mean of $\text{prod}(x)^{\text{power}}$ when x follows $T_{\text{dof}}(\mu,\text{sigma})$ distribution (dof= -1 for multivariate Normal).

Usage

```
eprod(m, S, power = 1, dof = -1)
```

Arguments

m	Location parameter
S	Scale matrix. For multivariate T with dof>2 the covariance is $S*\text{dof}/(\text{dof}-2)$. For the multivariate Normal the covariance is S.
power	Power that the product is raised to
dof	Degrees of freedom of the multivariate T. Set to -1 for the multivariate Normal.

Details

The calculation is based on the computationally efficient approach by Kan (2008).

Value

Expectation of the above-mentioned product

Author(s)

John Cook

References

Kan R. From moments of sum to moments of product. *Journal of Multivariate Analysis* 99 (2008), 542-554.

Examples

```
#Check easy independence case
m <- c(0,3); S <- matrix(c(2,0,0,1),ncol=2)

eprod(m, S, power=2)

(m[1]^2+S[1][1])*(m[2]^2+S[2][2])
```

getBIC

Obtain BIC and EBIC

Description

Obtain Bayesian Information Criterion and Extended Bayesian Information Criterion

Usage

```
getBIC(object)
```

```
getEBIC(object)
```

Arguments

object Object of class msfit returned by modelSelection

Details

Let p be the total number of parameters and n the sample size. The BIC of a model k with p_k parameters is

$$-2 L_k + p_k \log(n)$$

and the EBIC is

$$-2 L_k + p_k \log(n) + 2 \log(\binom{p}{p_k})$$

Function modelSelection returns the log posterior probability of a model, $\text{postProb} = \log(m_k) + \log(\text{prior } k)$, where m_k is the marginal likelihood of the model and prior k its prior probability.

When running function `modelSelection` with `priorCoef=bicprior()` and `priorDelta=modelunifprior()`, the BIC approximation is used for m_k , that is

$$\log(m_k) = L_k - 0.5 * p_k \log(n)$$

and all models are equally likely a priori, $\log(\text{prior } k) = p \log(1/2)$. Then the BIC can be easily recovered

$$\text{BIC}_k = -2 * [\text{postProb} + p \log(2)]$$

When using `priorCoef=bicprior()` and `priorDelta=modelbbprior()`, $\log(\text{prior } k) = -\log(p+1) - \log(p \text{ choose } p_k)$, hence

$$\text{EBIC}_k = -2 * [\text{postProb} + \log(p+1)].$$

Value

BIC or EBIC values for all models enumerated / visited by `modelSelection`

Author(s)

David Rossell

See Also

[modelSelection](#) to perform model selection

Examples

```
x <- matrix(rnorm(100*3),nrow=100,ncol=3)
theta <- matrix(c(1,1,0),ncol=1)
y <- x %*% theta + rnorm(100)
ybin <- y>0

#Obtain BIC
ms= modelSelection(ybin, x=x, priorCoef=bicprior(),
priorDelta=modelunifprior(), family='binomial')
getBIC(ms)

#Obtain EBIC
ms2= modelSelection(ybin, x=x, priorCoef=bicprior(),
priorDelta=modelbbprior(), family='binomial')
getEBIC(ms2)
```

Description

Montgomery and Peck (1982) illustrated variable selection techniques on the Hald cement data and gave several references to other analysis. The response variable y is the *heat evolved* in a cement mix. The four explanatory variables are ingredients of the mix, i.e., x_1 : *tricalcium aluminate*, x_2 : *tricalcium silicate*, x_3 : *tetracalcium alumino ferrite*, x_4 : *dicalcium silicate*. An important feature of these data is that the variables x_1 and x_3 are highly correlated ($\text{corr}(x_1, x_3) = -0.824$), as well as the variables x_2 and x_4 (with $\text{corr}(x_2, x_4) = -0.975$). Thus we should expect any subset of (x_1, x_2, x_3, x_4) that includes one variable from highly correlated pair to do as any subset that also includes the other member.

Usage

```
data(hald)
```

Format

hald is a matrix with 13 observations (rows) and 5 variables (columns), the first column is the dependent variable. y .hald and x .hald are also availables.

Source

Montgomery, D.C., Peck, E.A. (1982) *Introduction to linear regression analysis*, John Wiley, New York.

marginalNIW

Marginal likelihood under a multivariate Normal likelihood and a conjugate Normal-inverse Wishart prior.

Description

The argument z can be used to specify cluster allocations. If left missing then the usual marginal likelihood is computed, else it is computed conditional on the clusters (this is equivalent to the product of marginal likelihoods across clusters)

Usage

```
marginalNIW(x, xbar, samplecov, n, z, g, mu0=rep(0, ncol(x)),
nu0=ncol(x)+4, S0, logscale=TRUE)
```

Arguments

x	Data matrix (individuals in rows, variables in columns). Alternatively you can leave missing and specify $xbar$, $samplecov$ and n instead
$xbar$	Either a vector with column means of x or a list where each element corresponds to the column means for each cluster
$samplecov$	Either the sample covariance matrix $cov(x)$ or a list where each element contains the covariance for each cluster

n	Either an integer indicating the sample size <code>nrow(x)</code> or a vector indicating the cluster counts <code>table(z)</code>
z	Optional argument specifying cluster allocations
g	Prior dispersion parameter for mu
mu0	Prior mean for mu
nu0	Prior degrees of freedom for Sigma
S0	Prior scale matrix for Sigma, by default set to I/ν_0
logscale	set to TRUE to get the log-posterior density

Details

The function computes

$$p(x) = \int p(x | \mu, \Sigma) p(\mu, \Sigma) d\mu d\Sigma$$

where $p(x[i]) = N(x[i]; \mu, \Sigma)$ iid $i=1, \dots, n$

$$p(\mu | \Sigma) = N(\mu; \mu_0, g \Sigma) \quad p(\Sigma) = IW(\Sigma; \nu_0, S_0)$$

Value

If z is missing the integrated likelihood under a Normal-IW prior. If z was specified then the product of integrated likelihoods across clusters

Author(s)

David Rossell

See Also

[dpostNIW](#) for the posterior Normal-IW density.

Examples

```
#Simulate data
x= matrix(rnorm(100),ncol=2)

#Integrated likelihood under correct model
marginalNIW(x,g=1,nu0=4,log=FALSE)

#Integrated likelihood under random cluster allocations
z= rep(1:2,each=25)
marginalNIW(x,z=z,g=1,nu0=4,log=FALSE)
```

mixturebf-class *Class "mixturebf"*

Description

Stores the output of Bayesian model selection for mixture models, e.g. as produced by function `bfnormmix`.

Methods are provided for retrieving the posterior probability of a given number of mixture components, posterior means and posterior samples of the mixture model parameters.

Objects from the Class

Typically objects are automatically created by a call to `bfnormmix`.

Slots

The class has the following slots:

postprob data.frame containing posterior probabilities for different numbers of components (k) and log-posterior probability of a component being empty (contain no individuals)

p Number of variables in the data to which the model was fit

n Number of observations in the data to which the model was fit

priorpars Prior parameters used when fitting the model

postpars Posterior parameters for a 1-component mixture, e.g. for a Normal mixture the posterior is $N(\mu_1, \text{Sigma}/\text{prec}) \text{IW}(\nu_1, S_1)$

mcmc For each considered value of k, posterior samples for the parameters of the k-component model are stored

Methods

coef Computes posterior means for all parameters

show `signature(object = "mixturebf")`: Displays general information about the object.

postProb `signature(object = "mixturebf")`: Extracts posterior model probabilities, Bayes factors and posterior probability of a cluster being empty

postSamples `signature(object = "mixturebf")`: Extracts posterior samples

Author(s)

David Rossell

References

Fuquene J., Steel M.F.J., Rossell D. On choosing mixture components via non-local priors. 2018. arXiv

See Also

See also [bfnormmix](#)

Examples

```
showClass("mixturebf")
```

modelSelection	<i>Bayesian variable selection for linear models via non-local priors.</i>
----------------	----------------------------------------------------------------------------

Description

Bayesian model selection for linear, asymmetric linear, median and quantile regression under non-local or Zellner priors. $p \gg n$ can be handled.

modelSelection enumerates all models when feasible and uses a Gibbs scheme otherwise. See coef and coefByModel for estimates and posterior intervals of regression coefficients, and rnlp for posterior samples.

modelsearchBlockDiag seeks the highest posterior probability model using an iterative block search.

Usage

```
modelSelection(y, x, data, smoothterms, nknots=9,
  groups=1:ncol(x), constraints, center=TRUE, scale=TRUE,
  enumerate, includevars=rep(FALSE,ncol(x)), models,
  maxvars, niter=5000, thinning=1,
  burnin=round(niter/10), family='normal', priorCoef,
  priorGroup, priorDelta=modelbbprior(1,1),
  priorConstraints,
  priorVar=igprior(.01,.01),
  priorSkew=momprior(tau=0.348), phi, deltaini=rep(FALSE,ncol(x)),
  initSearch='greedy', method='auto', adj.overdisp='intercept',
  hess='asympt', optimMethod, optim_maxit, initpar='none', B=10^5,
  XtXprecomp= ifelse(ncol(x)<10^4,TRUE,FALSE), verbose=TRUE)
```

```
modelsearchBlockDiag(y, x, priorCoef=momprior(tau=0.348),
  priorDelta=modelbbprior(1,1), priorVar=igprior(0.01,0.01),
  blocksize=10, maxiter=10, maxvars=100, maxlogmargdrop=20,
  maxenum=10, verbose=TRUE)
```

Arguments

y Either a formula with the regression equation or a vector with observed responses. The response can be either continuous or of class Surv (survival outcome). If y is a formula then x, groups and constraints are automatically created

x	Design matrix with linear covariates for which we want to assess if they have a linear effect on the response. Ignored if y is a formula
data	If y is a formula then data should be a data frame containing the variables in the model
smoothterms	Formula for non-linear covariates (cubic splines), modelSelection assesses if the variable has no effect, linear or non-linear effect. smoothterms can also be a design matrix or data.frame containing linear terms, for each column modelSelection creates a spline basis and tests no/linear/non-linear effects
nknots	Number of spline knots. For cubic splines the non-linear basis adds knots-4 coefficients for each linear term, we recommend setting nknots to a small/moderate value
groups	If variables in x such be added/dropped in groups, groups indicates the group that each variable corresponds to (by default each variable goes in a separate group)
constraints	Constraints on the model space. List with length equal to the number of groups; if group[[i]]=c(j,k) then group i can only be in the model if groups j and k are also in the model
center	If TRUE, y and x are centered to have zero mean. Dummy variables corresponding to factors are NOT centered
scale	If TRUE, y and columns in x are scaled to have variance=1. Dummy variables corresponding to factors are NOT scaled
enumerate	Default is TRUE if there's less than 15 variable groups. If TRUE all models with up to maxvars are enumerated, else Gibbs sampling is used to explore the model space
includevars	Logical vector of length ncol(x) indicating variables that should always be included in the model, i.e. variable selection is not performed for these variables
models	Optional logical matrix indicating the models to be enumerated with rows equal to the number of desired models and columns to the number of variables in x.
maxvars	When enumerate==TRUE only models with up to maxvars variables enumerated (defaults to all variables). In modelsearchBlockDiag a sequence of models is defined from 1 up to maxvars
niter	Number of Gibbs sampling iterations
thinning	MCMC thinning factor, i.e. only one out of each thinning iterations are reported. Defaults to thinning=1, i.e. no thinning
burnin	Number of burn-in MCMC iterations. Defaults to .1*niter. Set to 0 for no burn-in
family	Family of parametric distribution. Possible values are 'normal' for Normal errors, 'twopiecenormal' for two-piece Normal, 'laplace' for Laplace errors, 'twopiecelaplace' for double exponential, 'binomial' for logistic regression, 'poisson' for Poisson regression. For 'auto' the error distribution is inferred from the data among 'normal', 'laplace', 'twopiecenormal' and 'twopiecelaplace'. 'laplace' corresponds to median regression and 'twopiecelaplace' to quantile regression. See argument priorSkew

priorCoef	Prior on coefficients, created by momprior, imomprior, emomprior or zellnerprior. Prior dispersion is on coefficients/sqrt(scale) for Normal and two-piece Normal, and on coefficients/sqrt(2*scale) for Laplace and two-piece Laplace.
priorGroup	Prior on grouped coefficients (e.g. categorical predictors with >2 categories, splines). Created by groupmomprior, groupemomprior, groupimomprior or groupzellnerprior
priorDelta	Prior on model space. Use modelbbprior() for Beta-Binomial prior, modelbinomprior(p) for Binomial prior with prior inclusion probability p, modelcomplexprior for Complexity prior, or modelunifprior() for Uniform prior
priorConstraints	Prior distribution on the number of terms subject to hierarchical constraints that are included in the model
priorVar	Inverse gamma prior on scale parameter. For Normal outcomes variance=scale, for Laplace outcomes variance=2*scale
priorSkew	Either a fixed value for tanh(alpha) where alpha is the asymmetry parameter or a prior on tanh(alpha). For family=='twopiecelaplace' setting alpha=a is equivalent to performing quantile regression for the quantile (1+a)/2. Ignored if family is 'normal' or 'laplace'.
phi	The error variance in Gaussian models, typically this is unknown and is left missing
deltaini	Logical vector of length ncol(x) indicating which coefficients should be initialized to be non-zero. Defaults to all variables being excluded from the model
initSearch	Algorithm to refine deltaini. initSearch=='greedy' uses a greedy Gibbs sampling search. initSearch=='SCAD' sets deltaini to the non-zero elements in a SCAD fit with cross-validated regularization parameter. initSearch=='none' leaves deltaini unmodified
method	Method to compute marginal likelihood. method=='Laplace' for Laplace approx, method=='ALA' for approximate Laplace approximation. method=='MC' for Importance Sampling, method=='Hybrid' for Hybrid Laplace-IS (only available for piMOM prior). See Details. method=='auto' attempts to use exact calculations when possible, otherwise ALA if available, otherwise Laplace approx.
adj.overdisp	Only used when method=='ALA'. Over-dispersion adjustment in models with fixed dispersion parameter, as in logistic and Poisson regression. adj.overdisp='none' for no adjustment (not recommended, particularly for Poisson models). adj.overdisp='intercept' to estimate over-dispersion from the intercept-only model, and adj.overdisp='residuals' from the Pearson residuals of each model
hess	Method to estimate the hessian in the Laplace approximation to the integrated likelihood under Laplace or asymmetric Laplace errors. When hess=='asymp' the asymptotic hessian is used, hess=='asympDiagAdj' a diagonal adjustment is applied (see Rossell and Rubio for details).
optimMethod	Algorithm to maximize objective function when method=='Laplace'. Leave unspecified or set optimMethod=='auto' for an automatic choice. optimMethod=='LMA' uses modified Newton-Raphson algorithm, 'CDA' coordinate descent algorithm
optim_maxit	Maximum number of iterations when method=='Laplace'

initpar	Initial regression parameter values when finding the posterior mode to approximate the integrated likelihood. 'none', 'MLE', 'MLE-aisgd', 'L1', 'L2-aisgd', or a numeric vector with initial values. If $p < n/2$ MLE is used, else L1 (regularization parameter set via BIC). 'auto': if $n > 10,000$ or $p > 200$, then MLE-aisgd or L2-aisgd are used. aisgd stands for averaged intrinsic stochastic gradient descent (see function sgd in package sgd)
B	Number of samples to use in Importance Sampling scheme. Ignored if method=='Laplace'
XtXprecomp	Set to TRUE to pre-compute the Gram matrix $x'x$ upfront (saves time), to FALSE to compute and store elements only as needed (saves memory)
verbose	Set verbose==TRUE to print iteration progress
blocksize	Maximum number of variables in a block. Careful, the cost of the algorithm is of order $2^{\text{blocksize}}$
maxiter	Maximum number of iterations, each iteration includes a screening pass to add and subtract variables
maxlogmargdrop	Stop the sequence of models when the drop in $\log p(y \text{model})$ is greater than maxlogmargdrop. This option avoids spending unnecessary time exploring overly large models
maxenum	If the posterior mode found has less than maxenum variables then do a full enumeration of all its submodels

Details

Let δ be the vector indicating inclusion/exclusion of each column of x in the model. The Gibbs algorithm sequentially samples from the posterior of each element in δ conditional on all the remaining elements in δ and the data. To do this it is necessary to evaluate the marginal likelihood for any given model. These have closed-form expression for the MOM prior, but for models with >15 variables these are expensive to compute and Laplace approximations are used instead (for the residual variance a log change of variables is used, which improves the approximation). For other priors closed forms are not available, so by default Laplace approximations are used. For the iMOM prior we also implement a Hybrid Laplace-IS which uses a Laplace approximation to evaluate the integral wrt β and integrates wrt ϕ (residual variance) numerically.

It should be noted that Laplace approximations tend to under-estimate the marginal densities when the MLE for some parameter is very close to 0. That is, it tends to be conservative in the sense of excluding more variables from the model than an exact calculation would.

Finally, method=='plugin' provides a BIC-type approximation that is faster than exact or Laplace methods, at the expense of some accuracy. In non-sparse situations where models with many variables have large posterior probability method=='plugin' can be substantially faster.

For more details on the methods used to compute marginal densities see Johnson & Rossell (2012).

modelsearchBlockDiag uses the block search method described in Papaspiliopoulos & Rossell. Briefly, spectral clustering is run on $X'X$ to cluster variables into blocks of blocksize and subsequently the Coolblock algorithm is used to define a sequence of models of increasing size. The exact integrated likelihood is evaluated for all models in this path, the best model chosen, and the scheme iteratively repeated to add and drop variables until convergence.

Value

Object of class `msfit`, which extends a list with elements

<code>postSample</code>	matrix with posterior samples for the model indicator. <code>postSample[i, j]==1</code> indicates that variable <code>j</code> was included in the model in the MCMC iteration <code>i</code>
<code>postOther</code>	<code>postOther</code> returns posterior samples for parameters other than the model indicator, i.e. basically hyper-parameters. If hyper-parameters were fixed in the model specification, <code>postOther</code> will be empty.
<code>margpp</code>	Marginal posterior probability for inclusion of each covariate. This is computed by averaging marginal post prob for inclusion in each Gibbs iteration, which is much more accurate than simply taking <code>colMeans(postSample)</code>
.	.
<code>postMode</code>	Model with highest posterior probability amongst all those visited
<code>postModeProb</code>	Unnormalized posterior prob of posterior mode (log scale)
<code>postProb</code>	Unnormalized posterior prob of each visited model (log scale)
<code>priors</code>	List with priors specified when calling <code>modelSelection</code>

Author(s)

David Rossell

References

- Johnson V.E., Rossell D. Non-Local Prior Densities for Default Bayesian Hypothesis Tests. *Journal of the Royal Statistical Society B*, 2010, 72, 143-170.
- Johnson V.E., Rossell D. Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, 2012, 107, 649-660.
- Papaspiliopoulos O., Rossell, D. Scalable Bayesian variable selection and model averaging under block orthogonal design. 2016
- Rossell D., Rubio F.J. Tractable Bayesian variable selection: beyond normality. 2016

See Also

[msfit-class](#) for details on the output. [postProb](#) to obtain posterior model probabilities. `coef.msfit` for Bayesian model averaging estimates and intervals. `predict.msfit` for BMA estimates and intervals for user-supplied covariate values. [rnlp](#) to obtain posterior samples for the coefficients. [nlpMarginal](#) to compute marginal densities for a given model.

Examples

```
#Simulate data
x <- matrix(rnorm(100*3),nrow=100,ncol=3)
theta <- matrix(c(1,1,0),ncol=1)
y <- x %*% theta + rnorm(100)

#Specify prior parameters
```

```

priorCoef <- momprior(tau=0.348)
priorDelta <- modelunifprior()

#Alternative model space prior: 0.5 prior prob for including any covariate
priorDelta <- modelbinomprior(p=0.5)

#Alternative: Beta-Binomial prior for model space
priorDelta <- modelbbprior(alpha.p=1,beta.p=1)

#Model selection
fit1 <- modelSelection(y=y, x=x, center=FALSE, scale=FALSE,
priorCoef=priorCoef, priorDelta=priorDelta)
postProb(fit1) #posterior model probabilities

fit1$margpp #posterior marginal inclusion prob

coef(fit1) #BMA estimates, 95% intervals, marginal post prob

```

mombf

Moment and inverse moment Bayes factors for linear models.

Description

mombf computes moment Bayes factors to test whether a subset of regression coefficients are equal to some user-specified value. imombf computes inverse moment Bayes factors. zellnerbf computes Bayes factors based on the Zellner-Siow prior (used to build the moment prior).

Usage

```

mombf(lm1, coef, g, prior.mode, baseDensity='normal', nu=3, theta0,
logbf=FALSE, B=10^5)
imombf(lm1, coef, g, prior.mode, nu = 1, theta0, method='adapt',
nquant=100, B = 10^5)

```

Arguments

lm1	Linear model fit, as returned by lm1.
coef	Vector with indexes of coefficients to be tested. e.g. coef==c(2,3) and theta0==c(0,0) tests coef(lm1)[2]=coef(lm1)[3]=0.
g	Vector with prior parameter values. See dmom and dimom for details.
prior.mode	If specified, g is set such that the prior mode is prior.mode
baseDensity	Density upon which the Mom prior is based. baseDensity=='normal' results in the normal Mom prior, baseDensity=='t' in the t Mom prior with nu degrees of freedom.
nu	For mombf, nu specifies the degrees of freedom of the t Mom prior. It is ignored unless baseDensity=='t'. nu defaults to 3. For imombf, nu specifies the degrees of freedom for the inverse moment prior (see dimom for details). Defaults to nu=1, which Cauchy-like tails.

theta0	Null value for the regression coefficients. Defaults to 0.
logbf	If logbf==TRUE the natural logarithm of the Bayes factor is returned.
method	Numerical integration method to compute the bivariate integral (only used by imombf). For method=='adapt', the inner integral is evaluated (via integrate) at a series of nquant quantiles of the residual variance posterior distribution, and then averaged as described in Johnson (1992). Set method=='MC' to use Monte Carlo integration.
nquant	Number of quantiles at which to evaluate the integral for known sigma. Only used if method=='adapt'.
B	Number of Monte Carlo samples to estimate the T Mom and the inverse moment Bayes factor. Only used in mombf if baseDensity=='t'. Only used in imombf if method=='MC'.

Details

These functions actually call momunknown and imomunknown, but they have a simpler interface. See dmom and dimom for details on the moment and inverse moment priors. The Zellner-Siow g-prior is given by $dmvnorm(\theta, \theta_0, n * \sigma * V_1)$.

Value

mombf returns the moment Bayes factor to compare the model where $\theta \neq \theta_0$ with the null model where $\theta = \theta_0$. Large values favor the alternative model; small values favor the null. imombf returns inverse moment Bayes factors. zellnerbf returns Bayes factors based on the Zellner-Siow g-prior.

Author(s)

David Rossell

References

See <http://rosselldavid.googlepages.com> for technical reports. For details on the quantile integration, see Johnson, V.E. A Technique for Estimating Marginal Posterior Densities in Hierarchical Models Using Mixtures of Conditional Densities. Journal of the American Statistical Association, Vol. 87, No. 419. (Sep., 1992), pp. 852-860.

See Also

[momunknown](#), [imomunknown](#) and [zbfunknown](#) for another interface to compute Bayes factors. [momknown](#), [imomknown](#) and [zbfknown](#) to compute Bayes factors assuming that the dispersion parameter is known, and for approximate Bayes factors for GLMs

Examples

```
##compute Bayes factor for Hald's data
data(hald)
lm1 <- lm(hald[,1] ~ hald[,2] + hald[,3] + hald[,4] + hald[,5])
```

```
# Set g so that interval (-0.2,0.2) has 5% prior probability
# (in standardized effect size scale)
priorp <- .05; q <- .2
gmom <- priorp2g(priorp=priorp,q=q,prior='normalMom')
gimom <- priorp2g(priorp=priorp,q=q,prior='iMom')

mombf(lm1,coef=2,g=gmom) #moment BF
imombf(lm1,coef=2,g=gimom,B=10^5) #inverse moment BF
zellnerbf(lm1,coef=2,g=1) #BF based on Zellner's g-prior
```

momknown

Bayes factors for moment, inverse moment and Zellner-Siow g-prior.

Description

momknown and momunknown compute moment Bayes factors for linear models when σ^2 is known and unknown, respectively. The functions can also be used to compute approximate Bayes factors for generalized linear models and other settings. imomknown, imomunknown compute inverse moment Bayes factors. zbfknown, zbfunknown compute Bayes factors based on the Zellner-Siow g-prior.

Usage

```
momknown(theta1hat, V1, n, g = 1, theta0, sigma, logbf = FALSE)
momunknown(theta1hat, V1, n, nuisance.theta, g = 1, theta0, ssr, logbf =
FALSE)
imomknown(theta1hat, V1, n, nuisance.theta, g = 1, nu = 1, theta0,
sigma, method='adapt', B=10^5)
imomunknown(theta1hat, V1, n, nuisance.theta, g = 1, nu = 1, theta0,
ssr, method='adapt', nquant = 100, B = 10^5)
zbfknown(theta1hat, V1, n, g = 1, theta0, sigma, logbf = FALSE)
zbfunknown(theta1hat, V1, n, nuisance.theta, g = 1, theta0, ssr, logbf =
FALSE)
```

Arguments

theta1hat	Vector with regression coefficients estimates.
V1	Matrix proportional to the covariance of theta1hat. For linear models, the covariance is $\sigma^2 \cdot V1$.
n	Sample size.
nuisance.theta	Number of nuisance regression coefficients, i.e. coefficients that we do not wish to test for.
ssr	Sum of squared residuals from a linear model call.
g	Prior parameter. See dmom and dimom for details.
theta0	Null value for the regression coefficients. Defaults to 0.

sigma	Dispersion parameter is σ^2 .
logbf	If <code>logbf==TRUE</code> the natural logarithm of the Bayes factor is returned.
nu	Prior parameter for the inverse moment prior. See <code>dimom</code> for details. Defaults to <code>nu=1</code> , which Cauchy-like tails.
method	Numerical integration method (only used by <code>imomknown</code> and <code>imomunknown</code>). Set <code>method=='adapt'</code> in <code>imomknown</code> to integrate using adaptive quadrature of functions as implemented in the function <code>integrate</code> . In <code>imomunknown</code> the integral is evaluated as in <code>imomknown</code> at a series of <code>nquant</code> quantiles of the posterior for <code>sigma</code> , and then averaged as described in Johnson (1992). Set <code>method=='MC'</code> to use Monte Carlo integration.
nquant	Number of quantiles at which to evaluate the integral for known <code>sigma</code> .
B	Number of Monte Carlo samples to estimate the inverse moment Bayes factor. Ignored if <code>method!='MC'</code> .

Details

See `dmom` and `dimom` for details on the moment and inverse moment priors. The Zellner-Siow g-prior is given by `dmvnorm(theta,theta0,n*g*V1)`.

Value

`momknown` and `momunknown` return the moment Bayes factor to compare the model where `theta!=theta0` with the null model where `theta==theta0`. Large values favor the alternative model; small values favor the null. `imomknown` and `imomunknown` return inverse moment Bayes factors. `zbfknown` and `zbfunknown` return Bayes factors based on the Zellner-Siow g-prior.

Author(s)

David Rossell

References

See <http://rosselldavid.googlepages.com> for technical reports.

For details on the quantile integration, see Johnson, V.E. A Technique for Estimating Marginal Posterior Densities in Hierarchical Models Using Mixtures of Conditional Densities. Journal of the American Statistical Association, Vol. 87, No. 419. (Sep., 1992), pp. 852-860.

See Also

[mombf](#) and [imombf](#) for a simpler interface to compute Bayes factors in linear regression

Examples

```
#simulate data from probit regression
set.seed(4*2*2008)
n <- 50; theta <- c(log(2),0)
x <- matrix(NA,nrow=n,ncol=2)
x[,1] <- rnorm(n,0,1); x[,2] <- rnorm(n,.5*x[,1],1)
```

```

p <- pnorm(x[,1]*theta[1]+x[,2]+theta[2])
y <- rbinom(n,1,p)

#fit model
glm1 <- glm(y~x[,1]+x[,2],family=binomial(link = "probit"))
thetahat <- coef(glm1)
V <- summary(glm1)$cov.scaled

#compute Bayes factors to test whether x[,1] can be dropped from the model
g <- .5
bfmom.1 <- momknown(thetahat[2],V[2,2],n=n,g=g,sigma=1)
bfimom.1 <- imomknown(thetahat[2],V[2,2],n=n,nuisance.theta=2,g=g,sigma=1)
bfmom.1
bfimom.1

```

msfit-class

Class "msfit"

Description

Stores the output of Bayesian variable selection, as produced by function `modelSelection`. The class extends a list, so all usual methods for lists also work for `msfit` objects, e.g. accessing elements, retrieving names etc.

Methods are provided to compute posterior probabilities, obtaining regression coefficient estimates and posterior intervals (both via Bayesian model averaging and for individual models), and sampling from their posterior distribution, as indicated below.

Objects from the Class

Typically objects are automatically created by a call to `modelSelection`. Alternatively, objects can be created by calls of the form `new("msfit", x)` where `x` is a list with the adequate elements (see slots).

Slots

The class extends a list with elements:

postSample matrix with posterior samples for the model indicator. `postSample[i, j]==1` indicates that variable `j` was included in the model in the MCMC iteration `i`

postOther `postOther` returns posterior samples for parameters other than the model indicator, i.e. basically hyper-parameters. If hyper-parameters were fixed in the model specification, `postOther` will be empty.

margpp Marginal posterior probability for inclusion of each covariate. This is computed by averaging marginal post prob for inclusion in each Gibbs iteration, which is much more accurate than simply taking `colMeans(postSample)`.

postMode Model with highest posterior probability amongst all those visited

postModeProb Unnormalized posterior prob of posterior mode (log scale)
postProb Unnormalized posterior prob of each visited model (log scale)
family Residual distribution, i.e. argument family when calling `modelSelection`
p Number of variables
priors Priors specified when calling `modelSelection`
ystd For internal use. Stores the response variable, standardized if center or scale were set to TRUE
xstd For internal use. Stores the covariates, standardized if center or scale were set to TRUE
stdconstants For internal use. If center or scale were set to TRUE, stores the sample mean and standard deviation of the outcome and covariates
call Stores info about the call, the formula used (if any), splines used etc

Methods

coef Obtains posterior means and intervals via Bayesian model averaging
coefByModel Obtains posterior means and intervals for individual models
predict Obtains posterior means and intervals for given covariate values. These are posterior intervals for the mean, not posterior predictive intervals for the outcome
show `signature(object = "msfit")`: Displays general information about the object.
postProb `signature(object = "msfit")`: Extracts posterior model probabilities.
rnlp `signature(object = "msfit")`: Obtain posterior samples for regression coefficients.

Author(s)

David Rossell

References

Johnson VE, Rossell D. Non-Local Prior Densities for Default Bayesian Hypothesis Tests. *Journal of the Royal Statistical Society B*, 2010, 72, 143-170
 Johnson VE, Rossell D. Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, 107, 498:649-660.

See Also

See also [modelSelection](#) and [rnlp](#).

Examples

```
showClass("msfit")
```

msPriorSpec-class *Class "msPriorSpec"*

Description

Stores the prior distributions to be used for Bayesian variable selection in normal regression models. This class can be used to specify the prior on non-zero regression coefficients, the model indicator or the nuisance parameters.

Usage

```

bicprior()

momprior(taustd=1, tau, tau.adj=10^6, r=1)
imomprior(tau, tau.adj=10^6)
emomprior(tau, tau.adj=10^6)
zellnerprior(taustd=1, tau, tau.adj=10^6)
normalidprior(taustd=1, tau, tau.adj=10^6)

groupmomprior(taustd=1, tau, tau.adj=10^6)
groupimomprior(tau, tau.adj=10^6)
groupemomprior(tau, tau.adj=10^6)
groupzellnerprior(taustd=1, tau, tau.adj=10^6)

modelunifprior()
modelbinomprior(p=0.5)
modelbbprior(alpha.p=1, beta.p=1)
modelcomplexprior(c=1)

igprior(alpha=.01, lambda=.01)

```

Arguments

tau	Prior dispersion parameter for covariates undergoing selection
taustd	Prior dispersion parameter for covariates undergoing selection. It is calibrated so that 'taustd=1' equals the unit information prior.
tau.adj	Prior variance in Normal prior for covariates not undergoing selection
r	MOM prior parameter is $2*r$
p	Prior inclusion probability for binomial prior on model space
alpha.p	Beta-binomial prior on model space has parameters alpha.p, beta.p
beta.p	Beta-binomial prior on model space has parameters alpha.p, beta.p
c	Under the Complexity prior the prior probability of having k variables in the model is proportional to $1/p^{(ck)}$
alpha	Inverse gamma prior has parameters alpha/2, lambda/2
lambda	Inverse gamma prior has parameters alpha/2, lambda/2

Details

DISCUSSION OF PRIOR ON PARAMETERS

Let $\beta=(\beta_1,\dots,\beta_p)$ be the regression coefficients for individual variables and $\delta=(\delta_1,\dots,\delta_q)$ those for grouped variables (e.g. factors or smooth terms in modelSelection).

momprior, emomprior, imomprior, zellnerprior and normalid can be priors on both β or δ . For further information see the vignette.

groupzellnerprior is the prior density on δ

$$p_z(\delta; \tau) = \prod_j N(\delta_j; 0, (\tau/p_j)) (X_j' X_j)^{-1}$$

where X_j are the design matrix columns associated to δ_j and $p_j = \text{ncol}(X_j)$ is the number of covariates in the group (for groupmomprior, the term in the denominator is $(p_j + 2)$ instead of p_j). A default $\tau = n = \text{nrow}(X_j)$ mimics the unit information prior and implies that the ratio of variance explained by X_j / residual variance is expected to be 1 a priori. To set the dispersion in terms of unit information prior, taustd is also available.

groupmomprior adds a quadratic MOM penalty

$$p_m(\delta; \tau) = p_z(\delta; \tau * n) \prod_j \delta_j' X_j' X_j \delta_j \text{ncol}(X_j) / (\tau * n * p_j / (p_j + 2))$$

and analogously for eMOM and iMOM. Note that unlike groupzellnerprior, the $\text{nrow}(X_j)$ factor is already included in the code. This is done to give user introduced τ values a roughly similar meaning between momprior and groupmomprior.

DISCUSSION OF PRIOR ON MODELS

Under the uniform prior, the prior probability of any model is $1 / \text{number of models}$.

Under the Binomial, Beta-Binomial and Complexity priors a model with k out of K active variables has prior probability $P(Z=k) / (K \text{ choose } k)$, where $Z \sim \text{Binomial}(K, p)$, $Z \sim \text{BetaBinomial}(K, \alpha, p, \beta)$ or for the Complexity prior $P(Z=k)$ proportional to $1/K^c (c * k)$.

Objects from the Class

Objects can be created by calls of the form `new("msPriorSpec", ...)`, but it is easier to use creator functions.

For priors on regression coefficients use momprior, imomprior or emomprior. For prior on model space modelunifprior, modelbinomprior modelbbprior, or modelcomplexprior. For prior on residual variance use igprior.

Slots

priorType: Object of class "character". "coefficients" indicates that the prior is for the non-zero regression coefficients. "modelIndicator" that it is for the model indicator, and "nuisancePars" that it is for the nuisance parameters. Several prior distributions are available for each choice of priorType, and these can be specified in the slot priorDist.

priorDistr: Object of class "character". If priorType=="coefficients", priorDistr can be equal to "pMOM", "piMOM", "peMOM", "zellner", "normalid", "groupMOM" or "groupzellner" (product moment, product inverse moment, product exponential moment, Zellner prior, normal prior with $\Sigma = I$, respectively). If priorType=="modelIndicator", priorDistr can be equal to "uniform" or "binomial" to specify a uniform prior (all models equally likely a

priori) or a binomial prior, or to "complexity" for the Complexity prior of Castillo et al 2015. For a binomial prior, the prior inclusion probability for any single variable must be specified in slot `priorPars['p']`. For a beta-binomial prior, the Beta hyper-prior parameters must be in `priorPars['alpha.p']` and `priorPars['beta.p']`. For the Complexity prior, the prior parameter must be in the slot `priorPars['c']`. If `priorType=="nuisancePars"`, `priorDistr` must be equal to "invgamma". This corresponds to an inverse gamma distribution for the residual variance, with parameters specified in the slot `priorPars`.

`priorPars`: Object of class "vector", where each element must be named. For `priorDistr=='pMOM'`, there must be an element "r" (MOM power is $2r$). For any `priorDistr` there must be either an element "tau" indicating the prior dispersion or elements "a.tau" and "b.tau" specifying an inverse gamma hyper-prior for "tau". Optionally, there may be an element "tau.adj" indicating the prior dispersion for the adjustment variables (i.e. not undergoing variable selection). If not defined, "tau.adj" is set to 0.001 by default. For `priorDistr=='binomial'`, there must be either an element "p" specifying the prior inclusion probability for any single covariate, or a vector with elements "alpha.p" and "beta.p" specifying a $\text{Beta}(\alpha.p, \beta.p)$ hyper-prior on p . For `priorDistr=='invgamma'` there must be elements "alpha" and "lambda". The prior for the residual variance is an inverse gamma with parameteres $.5*\alpha$ and $.5*\lambda$.

Methods

No methods defined with class "msPriorSpec" in the signature.

Note

When new instances of the class are created a series of check are performed to ensure that a valid prior specification is produced.

Author(s)

David Rossell

References

Johnson VE, Rossell D. Non-Local Prior Densities for Default Bayesian Hypothesis Tests. *Journal of the Royal Statistical Society B*, 2010, 72, 143-170

Johnson VE, Rossell D. Bayesian model selection in high-dimensional settings. *Journal of the American Statistical Association*, 107, 498:649-660.

See Also

See also [modelSelection](#) for an example of defining an instance of the class and perform Bayesian model selection.

Examples

```
showClass("msPriorSpec")
```

nlpmarginals	<i>Marginal density of the observed data for linear regression with Normal, two-piece Normal, Laplace or two-piece Laplace residuals under non-local and Zellner priors</i>
--------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Description

The marginal density of the data, i.e. the likelihood integrated with respect to the given prior distribution on the regression coefficients of the variables included in the model and an inverse gamma prior on the residual variance.

nlpMarginal is the general function, the remaining ones correspond to particular cases and are kept for backwards compatibility with old code, and will be deprecated in the future.

Usage

```
nlpMarginal(sel, y, x, data, smoothterms, nknots=9, groups=1:ncol(x),
family="normal", priorCoef, priorGroup,
priorVar=igprior(alpha=0.01,lambda=0.01), priorSkew=momprior(tau=0.348),
phi, method='auto', adj.overdisp='intercept', hess='asyp', optimMethod,
optim_maxit, initpar='none', B=10^5, logscale=TRUE, XtX, ytX)
```

```
pimomMarginalK(sel, y, x, phi, tau=1, method='Laplace', B=10^5, logscale=TRUE, XtX, ytX)
pimomMarginalU(sel, y, x, alpha=0.001, lambda=0.001, tau=1,
method='Laplace', B=10^5, logscale=TRUE, XtX, ytX)
pmomMarginalK(sel, y, x, phi, tau, r=1, method='auto', B=10^5,
logscale=TRUE, XtX, ytX)
pmomMarginalU(sel, y, x, alpha=0.001, lambda=0.001, tau=1,
r=1, method='auto', B=10^5, logscale=TRUE, XtX, ytX)
```

Arguments

sel	Vector with indexes of columns in x to be included in the model. Ignored if y is a formula
y	Either a formula with the regression equation or a vector with observed responses. The response can be either continuous or of class Surv (survival outcome). If y is a formula then x, groups and constraints are automatically created
x	Design matrix with linear covariates for which we want to assess if they have a linear effect on the response. Ignored if y is a formula
data	If y is a formula then data should be a data frame containing the variables in the model
smoothterms	Formula for non-linear covariates (cubic splines), modelSelection assesses if the variable has no effect, linear or non-linear effect. smoothterms can also be a design matrix or data.frame containing linear terms, for each column modelSelection creates a spline basis and tests no/linear/non-linear effects

nknots	Number of spline knots. For cubic splines the non-linear basis adds knots-4 coefficients for each linear term, we recommend setting nknots to a small/moderate value
groups	If variables in x such be added/dropped in groups, groups indicates the group that each variable corresponds to (by default each variable goes in a separate group)
family	Residual distribution. Possible values are 'normal','twopiecenormal','laplace', 'twopiecelaplace'
priorCoef	Prior on coefficients, created by momprior, imomprior, emomprior or zellnerprior. Prior dispersion is on coefficients/sqrt(scale) for Normal and two-piece Normal, and on coefficients/sqrt(2*scale) for Laplace and two-piece Laplace.
priorGroup	Prior on grouped coefficients (e.g. categorical predictors with >2 categories, splines). Created by groupmomprior, groupemomprior, groupimomprior or groupzellnerprior
priorVar	Inverse gamma prior on scale parameter, created by igprior(). For Normal variance=scale, for Laplace variance=2*scale.
priorSkew	Either a number fixing tanh(alpha) where alpha is the asymmetry parameter or a prior on residual skewness parameter, assumed to be of the same family as priorCoef. Ignored if family is 'normal' or 'laplace'.
method	Method to approximate the integral. See help(modelSelection).
adj.overdisp	Only used for method=='ALA'. Over-dispersion adjustment for models with fixed dispersion parameter such as logistic and Poisson regression
hess	Method to estimat the hessian in the Laplace approximation to the integrated likelihood under Laplace or asymmetric Laplace errors. When hess=='asymp' the asymptotic hessian is used, hess=='asympDiagAdj' a diagonal adjustment is applied (see Rossell and Rubio for details).
optimMethod	Algorithm to maximize objective function when method=='Laplace'. Leave unspecified or set optimMethod=='auto' for an automatic choice. optimMethod=='LMA' uses modified Newton-Raphson algorithm, 'CDA' coordinate descent algorithm
optim_maxit	Maximum number of iterations when method=='Laplace'
initpar	Initial regression parameter values when finding the posterior mode to approximate the integrated likelihood. See help(modelSelection)
B	Number of Monte Carlo samples to use (ignored unless method=='MC')
logscale	If logscale==TRUE the log marginal density is returned.
XtX	Optionally, specify the matrix X'X. Useful when the function must be called a large number of times.
ytX	Optionally, specify the vector y'X. Useful when the function must be called a large number of times.
phi	Dispersion parameter. See help(modelSelection)
alpha	Prior for phi is inverse gamma alpha/2, lambda/2
lambda	Prior for phi is inverse gamma alpha/2, lambda/2
tau	Prior dispersion parameter for MOM and iMOM priors (see details)
r	Prior power parameter for MOM prior is 2*r

Details

The marginal density of the data is equal to the integral of $N(y; x[,sel]^* \theta, \phi * I) * \pi(\theta | \phi, \tau) * IG(\phi; \alpha/2, \lambda/2)$ with respect to θ , where $\pi(\theta | \phi, \tau)$ is a non-local prior and IG denotes the density of an inverse gamma.

`pmomMarginalK` and `pimomMarginalK` assume that the residual variance is known and therefore the inverse-gamma term in the integrand can be omitted.

The product MOM and iMOM densities can be evaluated using the functions `dmom` and `dimom`.

Value

Marginal density of the observed data under the specified prior.

Author(s)

David Rossell

References

Johnson V.E., Rossell D. Non-Local Prior Densities for Default Bayesian Hypothesis Tests. *Journal of the Royal Statistical Society B*, 2010, 72, 143-170. See <http://rosselldavid.googlepages.com> for technical reports.

See Also

`modelSelection` to perform model selection based on product non-local priors. `momunknown`, `imomunknown`, `momknown`, `imomknown` to compute Bayes factors for additive MOM and iMOM priors

Examples

```
x <- matrix(rnorm(100*2), ncol=2)
y <- x %*% matrix(c(.5,1), ncol=1) + rnorm(nrow(x))
pmomMarginalK(sel=1, y=y, x=x, phi=1, tau=1, method='Laplace')
pimomMarginalK(sel=1:2, y=y, x=x, phi=1, tau=1, method='Laplace')
```

pmomLM

Bayesian variable selection and model averaging for linear and probit models via non-local priors.

Description

Variable selection for linear and probit models, providing a sample from the joint posterior of the model and regression coefficients. `pmomLM` and `pmomPM` implement product Normal MOM and heavy-tailed product MOM as prior distribution for linear and probit model coefficients (respectively). `emomLM` and `emomPM` set an eMOM prior.

pp1PM finds the value of the prior dispersion parameter tau minimizing posterior expected predictive loss (Gelfand and Ghosh, 1998) for the Probit model, i.e. can be used to automatically set up tau.

ppmodel returns the proportion of visits to each model.

Usage

```
pmomLM(y, x, xadj, center=FALSE, scale=FALSE, niter=10^4, thinning=1,
burnin=round(niter/10), priorCoef, priorDelta, priorVar, initSearch='greedy',
verbose=TRUE)
```

```
pmomPM(y, x, xadj, niter=10^4, thinning=1, burnin=round(niter/10),
priorCoef, priorDelta, initSearch='greedy', verbose=TRUE)
```

```
emomLM(y, x, xadj, center=FALSE, scale=FALSE, niter=10^4, thinning=1,
burnin=round(niter/10), priorCoef, priorDelta, priorVar, initSearch='greedy',
verbose=TRUE)
```

```
emomPM(y, x, xadj, niter=10^4, thinning=1, burnin =round(niter/10),
priorCoef, priorDelta, initSearch='greedy', verbose=TRUE)
```

```
pp1PM(tauseq=exp(seq(log(.01),log(2),length=20)), kPen=1, y, x, xadj, niter=10^4,
thinning=1, burnin=round(niter/10), priorCoef, priorDelta, priorVar,
initSearch='greedy', mc.cores=1)
```

```
ppmodel(nlpfit)
```

Arguments

y	Vector with observed responses. For pmomLM this must be a numeric vector. For pmomPM it can either be a logical vector, a factor with 2 levels or a numeric vector taking only two distinct values.
x	Design matrix with all potential predictors which are to undergo variable selection.
xadj	Design matrix for adjustment covariates, i.e. variables which are included in the model with probability 1. For instance, xadj can be used to force the inclusion of an intercept in the model.
center	If center==TRUE, y and x are centered to have zero mean, therefore eliminating the need to include an intercept term in x
scale	If scale==TRUE, y and columns in x are scaled to have standard deviation 1
niter	Number of MCMC sampling iterations
thinning	MCMC thinning factor, i.e. only one out of each thinning iterations are reported. Defaults to thinning=1, i.e. no thinning
burnin	Number of burn-in MCMC iterations. Defaults to .1*niter. Set to 0 for no burn-in
priorCoef	Prior distribution for the coefficients. Must be object of class msPriorSpec with slot priorType set to 'coefficients'. Possible values for slot priorDistr are 'pMOM', 'piMOM' and 'peMOM'

priorDelta	Prior on model indicator space. Must be object of class msPriorSpec with slot priorType set to 'modelIndicator'. Possible values for slot priorDistr are 'uniform' and 'binomial'. For 'binomial', you can either set the prior probability 'p' or specify a Beta-binomial prior by specifying the parameters 'alpha.p','beta.p'.
priorVar	Prior on residual variance. Must be object of class msPriorSpec with slot priorType set to 'nuisancePars'. Slot priorDistr must be equal to 'invgamma'
initSearch	Algorithm to refine deltaini. initSearch=='greedy' uses a greedy Gibbs sampling search. initSearch=='SCAD' sets deltaini to the non-zero elements in a SCAD fit with cross-validated regularization parameter. initSearch=='none' initializes to the null model with no variables in x included.
verbose	Set verbose==TRUE to print iteration progress
tauseq	Grid of tau values for which the posterior predictive loss should be evaluated.
kPen	Penalty term specifying the relative importance of deviations from the observed data vs deviation from the posterior predictive. kPen can be set either to a numeric value or to 'msize' to set penalty equal to the average model size. Loss is $Dev(y_p, \hat{y}) + kPen * Dev(\hat{y}, y_{obs})$, where y_p : draw from post predictive, y_{obs} : observed data and \hat{y} is $E(y y_{obs})$.
mc.cores	Allows for parallel computing. mc.cores is the number of processors to use. Setting mc.cores>1 requires the parallel package.
nlpfit	Non-local prior model fit, as returned by pmomLM, pmomPM, emomLM or emomPM.

Details

The implemented MCMC scheme makes proposals from the joint posterior of $(\delta[i], \theta[i])$ given all other parameters and the data, where $\delta[i]$ is the indicator for inclusion/exclusion of covariate i and $\theta[i]$ is the coefficient value. In contrast with some model fitting options implemented in modelSelection, here the scheme is exact. However, sampling the coefficients can adversely affect the mixing when covariates are very highly correlated. In practice, the mixing seems to be reasonably good for correlations up to 0.9.

pmomPM uses the scheme of Albert & Chib (1993) for probit models.

Value

pmomLM and pmomPM returns a list with elements

postModel	matrix with posterior samples for the model indicator. $postModel[i, j]==1$ indicates that variable j was included in the model in the MCMC iteration i
postCoef1	matrix with posterior samples for coefficients associated to x
postCoef2	matrix with posterior samples for coefficients associated to xadj
postPhi	vector with posterior samples for residual variance
postOther	postOther returns posterior samples for other parameters, i.e. basically hyper-parameters. Currently the prior precision parameter tau
margpp	Marginal posterior probability for inclusion of each covariate. This is computed by averaging marginal post prob for inclusion in each MCMC iteration, which is much more accurate than simply taking $colMeans(postModel)$

.

pplPM returns a list with elements

optfit Probit model fit using tauopt. It is the result of a call to pmomPM.
 PPL data.frame indicating for each value in tauseq the posterior predictive loss
 (PPL=G+P), the goodness-of-fit (G) and penalty terms (P)

, the average number of covariates in the model (msize) including xadj and the smoothed sPPL
 obtained via a gam fit.

tauo Value of tau minimizing the PPL

Author(s)

David Rossell, Donatello Telesca

References

Johnson V.E., Rossell D. Non-Local Prior Densities for Default Bayesian Hypothesis Tests. Journal of the Royal Statistical Society B, 2010, 72, 143-170.

Johnson V.E., Rossell D. Bayesian model selection in high-dimensional settings. Technical report. 2011 See <http://rosselldavid.googlepages.com> for technical reports.

Albert, J. and Chib, S. (1993) Bayesian analysis of binary and polychotomous response data. Journal of the American Statistical Association, 88, p669-679

Gelfand, A. and Ghosh, S. (1998) Model choice: A minimum posterior predictive loss approach. Biometrika, 85, p1-11.

See Also

For more details on the prior specification see [msPriorSpec-class](#) To compute marginal densities for a given model see [pmomMarginalK](#), [pmomMarginalU](#), [pimomMarginalK](#), [pimomMarginalU](#).

Examples

```
#Simulate data
x <- matrix(rnorm(100*3),nrow=100,ncol=3)
xadj <- rep(1,nrow(x))
theta <- matrix(c(1,1,0),ncol=1)
y <- 10*xadj + x %*% theta + rnorm(100)

#Beta-binomial prior on model space
priorDelta <- modelbbprior(alpha.p=1,beta.p=1)

#Non-informative prior on residual variance
priorVar <- igprior(alpha=.01,lambda=.01)

#Product MOM prior with tau=0.3 on x coefficients
#Non-informative prior on xadj coefficients
priorCoef <- momprior(tau=0.3, tau.adj=10^6)
```



```

mom0 <- pmomLM(y=y,x=x,xadj=xadj,center=FALSE,scale=FALSE,niter=1000,
priorCoef=priorCoef,priorDelta=priorDelta,priorVar=priorVar)
round(colMeans(mom0$postModel),2)
round(colMeans(mom0$postCoef1),2)
round(colMeans(mom0$postCoef2),2)

#Alternative prior: hyper-prior on tau
priorCoef <- new("msPriorSpec",priorType='coefficients',priorDistr='pMOM',
priorPars=c(a.tau=1,b.tau=.135,tau.adj=10^6,r=1)) #hyper-prior
mom1 <-
pmomLM(y=y,x=x,xadj=xadj,center=FALSE,scale=FALSE,niter=1000,
priorCoef=priorCoef,priorDelta=priorDelta,priorVar=priorVar)
mean(mom1$postOther) #posterior mean for tau

#Probit model
n <- 500; rho <- .25; niter <- 1000
theta <- c(.4,.6,0); theta.adj <- 0
V <- diag(length(theta)); V[upper.tri(V)] <- V[lower.tri(V)] <- rho
x <- rmvnorm(n,rep(0,length(theta)),V); xadj <- matrix(1,nrow=nrow(x),ncol=1)
lpred <- as.vector(x %*% matrix(theta,ncol=1) + xadj %*% matrix(theta.adj,ncol=1))
p <- pnorm(lpred)
y <- runif(n)<p

mom2 <- pmomPM(y=y,x=x,xadj=xadj,niter=1000,priorCoef=priorCoef,
priorDelta=priorDelta,initSearch='greedy')
colMeans(mom2$postCoef1)
coef(glm(y ~ x + xadj -1, family=binomial(link='probit'))))

```

postModeOrtho

Bayesian model selection and averaging under block-diagonal $X'X$ for linear models.

Description

postModeOrtho is for diagonal $X'X$, postModeBlockDiag for the more general block-diagonal $X'X$, where X is the matrix with predictors.

Both functions return the model of highest posterior probability of any given size using an efficient search algorithm. This sequence of models includes the highest posterior probability model (HPM). Posterior model probabilities, marginal variable inclusion probabilities and Bayesian model averaging estimates are also provided. The unknown residual variance is integrated out using an exact deterministic algorithm of low computational cost (see details in reference).

Usage

```
postModeOrtho(y, x, priorCoef=momprior(tau=0.348), priorDelta=modelbbprior(1,1),
priorVar=igprior(0.01,0.01), bma=FALSE, includeModels, maxvars=100)
```

```
postModeBlockDiag(y, x, blocks, priorCoef=zellnerprior(tau=nrow(x)),
priorDelta=modelbinomprior(p=1/ncol(x)),priorVar=igprior(0.01,0.01), bma=FALSE,
```

```
maxvars=100, momcoef)
```

Arguments

y	Outcome
x	Matrix with predictors. If an intercept is desired x should include a column of 1's.
blocks	Factor or integer vector of length ncol(x) indicating the block that each column in x belongs to.
priorCoef	Prior distribution for the coefficients. Object created with momprior, imomprior, emomprior or zellnerprior.
priorDelta	Prior on model space. Use modelbbprior() for Beta-Binomial prior, modelbinomprior(p) for Binomial prior with prior inclusion probability p, modelcomplexprior for Complexity prior, or modelunifprior() for Uniform prior
priorVar	Inverse gamma prior on residual variance, created with igprior()
bma	Set to TRUE to obtain marginal inclusion probabilities and Bayesian model averaging parameter estimates for each column of x.
includeModels	Models that should always be included when computing posterior model probabilities. It must be a list, each element in the list corresponds to a model and must be a logical or numeric vector indicating the variables in that model
maxvars	The search for the HPM is restricted to models with up to maxvars variables (note: posterior model probabilities and BMA are valid regardless of maxvars)
momcoef	optional argument containing pre-computed coefficients needed to obtain the marginal likelihood under the pMOM prior. A first call to postModeBlockDiag returns these coefficients, thus this argument is useful to speed up successive calls.

Details

The first step is to list a sequence of models with 0,...,maxvars variables which, under fairly general conditions listed in Papaspiliopoulos & Rossell (2016), is guaranteed to include the HPM. Then posterior model probabilities are computed for all these models to determine the HPM, evaluate the marginal posterior of the residual variance on a grid, and subsequently compute the marginal density $p(y)$ via adaptive quadrature. Finally this adaptive grid is used to compute marginal inclusion probabilities and Bayesian model averaging estimates. For more details see Papaspiliopoulos & Rossell (2016).

Value

List with elements

models	data.frame indicating the variables included in the sequence of models found during the search of the HPM, and their posterior probabilities. The model with highest posterior probability in this list is guaranteed to be the HPM.
phi	data.frame containing an adaptive grid of phi (residual variance) values and their marginal posterior density $p(\text{phily})$.

logpy	log-marginal density $p(y)$, i.e. normalization constant of $p(\text{phily})$.
bma	Marginal posterior inclusion probabilities and Bayesian model averaging estimates for each column in x .
postmean.model	Coefficient estimates conditional on each of the models in <code>models</code>
momcoef	If a MOM prior was specified in <code>priorCoef</code> , <code>momcoef</code> stores some coefficients needed to compute its marginal likelihood

Author(s)

David Rossell

References

Papaspiliopoulos O., Rossell D. Scalable Bayesian variable selection and model averaging under block-orthogonal design. 2016

Examples

```
#Simulate data
set.seed(1)
p <- 400; n <- 410
x <- scale(matrix(rnorm(n*p),nrow=n,ncol=p),center=TRUE,scale=TRUE)
S <- cov(x)
e <- eigen(cov(x))
x <- t(t(x %*% e$vectors)/sqrt(e$values))
th <- c(rep(0,p-3),c(.5,.75,1)); phi <- 1
y <- x %*% matrix(th,ncol=1) + rnorm(n,sd=sqrt(phi))

#Fit
priorCoef=zellnerprior(tau=n); priorDelta=modelbinomprior(p=1/p); priorVar=igprior(0.01,0.01)
pm.zell <- postModeOrtho(y,x=x,priorCoef=priorCoef,priorDelta=priorDelta,priorVar=priorVar,
bma=TRUE)

#Best models
head(pm.zell$models)

#Posterior probabilities for sequence of models
nvars <- sapply(strsplit(as.character(pm.zell$models$modelid),split=','),length)
plot(nvars,pm.zell$models$pp,ylab='post prob',xlab='number of vars',ylim=0:1,xlim=c(0,50))

#Marginal posterior of phi
plot(pm.zell$phi,type='l',xlab='phi',ylab='p(phi|y)')

#Marginal inclusion prob & BMA estimates
plot(pm.zell$bma$margpp,ylab='Marginal inclusion prob')
plot(pm.zell$bma$coef,ylab='BMA estimate')
```

postProb	<i>Obtain posterior model probabilities</i>
----------	---------------------------------------------

Description

Obtain posterior model probabilities after running Bayesian model selection

Usage

```
postProb(object, nmax, method='norm')
```

Arguments

object	Object of class <code>msfit</code> returned by <code>modelSelection</code> or class <code>mixturebf</code> , e.g. returned by <code>bfnormmix</code>
nmax	Maximum number of models to report (defaults to no max)
method	Only when <code>class(object)</code> is <code>msfit</code> . For 'norm' probabilities are obtained by renormalizing the stored integrated likelihoods, for 'exact' they are given by the proportion of MCMC visits to each model. 'norm' has less variability but can be biased if the chain has not converged.

Value

A `data.frame` with posterior model probabilities in column `pp`. Column `modelid` indicates the indexes of the selected covariates (empty for the null model with no covariates).

Author(s)

David Rossell

See Also

[modelSelection](#) to perform model selection

Examples

```
#See help(modelSelection)
```

postSamples	<i>Extract posterior samples from an object</i>
-------------	-------------------------------------------------

Description

Obtain posterior model probabilities after running Bayesian model selection

Usage

```
postSamples(object)
```

Arguments

object	Object containing posterior samples, e.g. of class mixture bf as returned by bfnormmix
--------	----------------------------------------------------------------------------------------

Value

For objects of class mixturebf, a list with one element for each considered number of mixture components.

Each element in the list contains posterior samples on the mixture weights (eta) and other component-specific parameters such as means (mu) and Cholesky decomposition of the inverse covariance matrix (cholSigmainv)

Author(s)

David Rossell

Examples

```
#See help(bfnormmix)
```

priorp2g	<i>Moment and inverse moment prior elicitation</i>
----------	----------------------------------------------------

Description

priorp2g finds the g value giving priorp prior probability to the interval (-q,q).

Usage

```
priorp2g(priorp, q, nu=1, prior=c("iMom", "normalMom", "tMom"))
```

Arguments

prior	prior=='normalMom' does computations for the normal moment prior, prior=='tMom' for the T moment prior, prior=='iMom' does computations for the inverse moment prior. Currently prior=='tMom' is not implemented in priorp2g.
q	priorp2g returns g giving priorp prior probability to the interval (-q, q).
nu	Prior degrees of freedom for the T moment prior or the iMom prior (ignored if prior=='normalMom').
priorp	priorp2g returns g giving priorp prior probability to the interval (-q, q)

Details

See pmom and pimom for the MOM/iMOM cumulative distribution functions.

Value

priorp2g returns g giving priorp prior probability to the interval (-q, q).

Author(s)

David Rossell <rosselldavid@gmail.com>

References

See <http://rosselldavid.googlepages.com> for technical reports.

See Also

[pmom](#), [pimom](#)

Examples

```
data(hald)
lm1 <- lm(hald[, 1] ~ hald[, 2] + hald[, 3] + hald[, 4] + hald[, 5])

#find g value giving 0.05 probability to interval (-.2,.2)
priorp <- .05; q <- .2
gmom <- priorp2g(priorp=priorp, q=q, prior='normalMom')
gimom <- priorp2g(priorp=priorp, q=q, prior='iMom')
gmom
gimom
```

rnlp

*Posterior sampling for regression parameters***Description**

Gibbs sampler for linear, generalized linear and survival models under product non-local priors, Zellner's prior and a Normal approximation to the posterior. Both sampling conditional on a model and Bayesian model averaging are implemented (see Details).

If x and y not specified samples from non-local priors/posteriors with density proportional to $d(\theta) N(\theta; m, V)$ are produced, where $d(\theta)$ is the non-local penalty term.

Usage

```
rnlp(y, x, m, V, msfit, outcometype, family, priorCoef, priorGroup,
     priorVar, isgroup, niter=10^3, burnin=round(niter/10), thinning=1, pp='norm')
```

Arguments

<code>y</code>	Vector with observed responses. When <code>class(y)=='Surv'</code> sampling is based on the Cox partial likelihood, else a linear model is assumed.
<code>x</code>	Design matrix with all potential predictors
<code>m</code>	Mean for the Normal kernel
<code>V</code>	Covariance for the Normal kernel
<code>msfit</code>	Object of class <code>msfit</code> returned by <code>modelSelection</code> . If specified Bayesian model averaging posterior samples are returned, according to posterior model probabilities in <code>msfit</code> , and then arguments <code>y</code> , <code>x</code> , <code>m</code> , <code>V</code> etc. If <code>msfit</code> is missing then posterior samples under the full model $y \sim x$ are returned
<code>outcometype</code>	Type of outcome. Possible values are "Continuous", "glm" or "Survival"
<code>family</code>	Assumed family for the family. Some possible values are "normal", "binomial logit" and "Cox"
<code>priorCoef</code>	Prior distribution for the coefficients. Ignored if <code>msfit</code> is supplied. Must be object of class <code>msPriorSpec</code> , e.g. created by <code>momprior</code> , <code>emomprior</code> , <code>imomprior</code> , <code>zellnerprior</code>
<code>priorGroup</code>	Prior on grouped coefficients (e.g. categorical predictors with >2 categories, splines), as passed to <code>modelSelection</code>
<code>priorVar</code>	Prior on residual variance. Ignored if <code>msfit</code> supplied. Must be object of class <code>msPriorSpec</code> , e.g. created with <code>igprior</code>
<code>isgroup</code>	Logical vector where TRUE indicates that the variable is part of a group, e.g. one of several dummy indicators for a discrete covariate
<code>niter</code>	Number of MCMC iterations
<code>burnin</code>	Number of burn-in MCMC iterations. Defaults to <code>.1*niter</code> . Set to 0 for no burn-in

thinning	MCMC thinning factor, i.e. only one out of each thinning iterations are reported. Defaults to no thinning
pp	When <code>msfit</code> is provided this is the method to compute posterior model probabilities, which determine the sampled models. Can be 'norm' or 'exact', see <code>postProb</code> for details.

Details

The algorithm is implemented for product MOM (pMOM), product iMOM (piMOM) and product eMOM (peMOM) priors. The algorithm combines an orthogonalization that provides low serial correlation with a latent truncation representation that allows fast sampling.

When `y` and `x` are specified sampling is for the linear regression posterior. When argument `msfit` is left missing, posterior sampling is for the full model regressing `y` on all covariates in `x`. When `msfit` is specified each model is drawn with probability given by `postProb(msfit)`. In this case, a Bayesian Model Averaging estimate of the regression coefficients can be obtained by applying `colMeans` to the `rnlp` output matrix.

When `y` and `x` are left missing, sampling is from a density proportional to $d(\theta) N(\theta; m, V)$, where $d(\theta)$ is the non-local penalty (e.g. $d(\theta) = \text{prod}(\theta^{2r})$ for the pMOM prior).

Value

Matrix with posterior samples

Author(s)

David Rossell

References

D. Rossell and D. Telesca. Non-local priors for high-dimensional estimation, 2014. <http://arxiv.org/pdf/1402.5107v2.pdf>

See Also

[modelSelection](#) to perform model selection and compute posterior model probabilities. For more details on prior specification see [msPriorSpec-class](#).

Examples

```
#Simulate data
x <- matrix(rnorm(100*3),nrow=100,ncol=3)
theta <- matrix(c(1,1,0),ncol=1)
y <- x %*% theta + rnorm(100)
fit1 <- modelSelection(y=y, x=x, center=FALSE, scale=FALSE)

th <- rnlp(msfit=fit1, niter=100)
colMeans(th)
```


Index

- * **cil**
 - cil, 6
- * **classes**
 - mixturebf-class, 20
 - msfit-class, 30
 - msPriorSpec-class, 32
- * **datasets**
 - hald, 17
- * **distribution**
 - bbPrior, 2
 - dalapl, 9
 - ddir, 10
 - diwish, 11
 - dmom, 12
 - dpostNIW, 14
 - getBIC, 16
 - postProb, 44
 - postSamples, 45
 - priorp2g, 45
 - rnlp, 47
- * **distrib**
 - eprod, 15
- * **htest**
 - bfnormmix, 4
 - dmom, 12
 - marginalNIW, 18
 - modelSelection, 21
 - mombf, 26
 - momknown, 28
 - nlpmarginals, 35
 - pmomLM, 37
 - postModeOrtho, 41
 - priorp2g, 45
- * **models**
 - bfnormmix, 4
 - cil, 6
 - eprod, 15
 - getBIC, 16
 - marginalNIW, 18
 - modelSelection, 21
 - mombf, 26
 - momknown, 28
 - nlpmarginals, 35
 - pmomLM, 37
 - postModeOrtho, 41
 - postProb, 44
 - postSamples, 45
 - rnlp, 47
- bbPrior, 2
- bfnormmix, 4, 21
- bicprior (msPriorSpec-class), 32
- binomPrior (bbPrior), 2
- cil, 6
- coef.mixturebf (mixturebf-class), 20
- coefByModel (msfit-class), 30
- coefByModel, msfit-method (msfit-class), 30
- coefByModel-methods (msfit-class), 30
- dalapl, 9
- ddir, 10
- demom (dmom), 12
- demom, data.frame-method (dmom), 12
- demom, matrix-method (dmom), 12
- demom, vector-method (dmom), 12
- demom-methods (dmom), 12
- demomigmarg (dmom), 12
- dimom (dmom), 12
- diwish, 11, 15
- dmom, 12
- dmomigmarg (dmom), 12
- dpostNIW, 11, 14, 19
- emomLM (pmomLM), 37
- emomPM (pmomLM), 37
- emomprior (msPriorSpec-class), 32
- eprod, 15

- getBIC, 16
- getBIC,msfit-method (getBIC), 16
- getBIC-methods (getBIC), 16
- getEBIC (getBIC), 16
- getEBIC,msfit-method (getBIC), 16
- getEBIC-methods (getBIC), 16
- groupemomprior (msPriorSpec-class), 32
- groupimomprior (msPriorSpec-class), 32
- groupmomprior (msPriorSpec-class), 32
- groupzellnerprior (msPriorSpec-class), 32
- hald, 17
- igprior (msPriorSpec-class), 32
- imombf, 29
- imombf (mombf), 26
- imomknown, 27, 37
- imomknown (momknown), 28
- imomprior (msPriorSpec-class), 32
- imomunknown, 27, 37
- imomunknown (momknown), 28
- marginalNIW, 15, 18
- marginalNIW,matrix,missing,missing,missing,missing-method (marginalNIW), 18
- marginalNIW,matrix,missing,missing,missing,vector-method (marginalNIW), 18
- marginalNIW,missing,ANY,matrix,numeric,missing-method (marginalNIW), 18
- marginalNIW,missing,list,list,numeric,missing-method (marginalNIW), 18
- marginalNIW-methods (marginalNIW), 18
- mixturebf (mixturebf-class), 20
- mixturebf-class, 20
- modelbbprior (msPriorSpec-class), 32
- modelbinomprior (msPriorSpec-class), 32
- modelcomplexprior (msPriorSpec-class), 32
- modelsearchBlockDiag (modelSelection), 21
- modelSelection, 17, 21, 31, 34, 37, 44, 48
- modelunifprior (msPriorSpec-class), 32
- mombf, 26, 29
- momknown, 27, 28, 37
- momprior (msPriorSpec-class), 32
- momunknown, 27, 37
- momunknown (momknown), 28
- msfit (msfit-class), 30
- msfit-class, 30
- msfit.coef (msfit-class), 30
- msfit.predict (msfit-class), 30
- msPriorSpec (msPriorSpec-class), 32
- msPriorSpec-class, 32
- nlpMarginal, 25
- nlpMarginal (nlpmarginals), 35
- nlpmarginals, 35
- normalidprior (msPriorSpec-class), 32
- palapl (dalapl), 9
- pemom (dmom), 12
- pemomigmarg (dmom), 12
- pimom, 46
- pimom (dmom), 12
- pimomMarginalK, 40
- pimomMarginalK (nlpmarginals), 35
- pimomMarginalU, 40
- pimomMarginalU (nlpmarginals), 35
- pmom, 46
- pmom (dmom), 12
- pmomigmarg (dmom), 12
- pmomLM, 37
- pmomMarginalK, 40
- pmomMarginalK (nlpmarginals), 35
- pmomMarginalU, 40
- pmomMarginalU (nlpmarginals), 35
- pmomLM, 37
- postModeBlockDiag (postModeOrtho), 41
- postModeOrtho, 41
- postProb, 8, 25, 44
- postProb,cilfit-method (postProb), 44
- postProb,mixturebf-method (postProb), 44
- postProb,msfit-method (postProb), 44
- postProb-methods (postProb), 44
- postSamples, 45
- postSamples,mixturebf-method (postSamples), 45
- postSamples-methods (postSamples), 45
- ppIPM (pmomLM), 37
- ppmodel (pmomLM), 37
- priorp2g, 45
- qimom (dmom), 12
- qmom (dmom), 12
- ralapl (dalapl), 9
- rnlp, 25, 31, 47

rnlp,ANY,matrix,missing,missing,missing,character,character-method
(rnlp), 47

rnlp,ANY,matrix,missing,missing,missing,missing,missing-method
(rnlp), 47

rnlp,ANY,matrix,missing,missing,msfit,missing,missing-method
(rnlp), 47

rnlp,missing,missing,missing,missing,msfit,missing,missing-method
(rnlp), 47

rnlp,missing,missing,numeric,matrix,missing,missing,missing-method
(rnlp), 47

rnlp-methods (rnlp), 47

rpostNIW (dpostNIW), 14

show,mixturebf-method
(mixturebf-class), 20

show,msfit-method (msfit-class), 30

unifPrior (bbPrior), 2

x.hald (hald), 17

y.hald (hald), 17

zbfknown, 27

zbfknown (momknown), 28

zbfunknown, 27

zbfunknown (momknown), 28

zellnerbf (mombf), 26

zellnerprior (msPriorSpec-class), 32