

Package: graphicalEvidence (via r-universe)

October 26, 2024

Type Package

Title Graphical Evidence

Version 1.0

Date 2024-10-07

Maintainer David Rowe <david@rowe-stats.com>

Description Computes marginal likelihood in Gaussian graphical models through a novel telescoping block decomposition of the precision matrix which allows estimation of model evidence. The top level function used to estimate marginal likelihood is called `evidence`, which expects the prior name, data, and relevant prior specific parameters. This package also provides an MCMC prior sampler using the same underlying approach, implemented in `prior_sampling`, which expects a prior name and prior specific parameters. Both functions also expect the number of burn-in iterations and the number of sampling iterations for the underlying MCMC sampler.

License GPL-3

Imports Rcpp, parallel, doParallel, foreach, mvtnorm

LinkingTo Rcpp, RcppArmadillo

StagedInstall true

RoxygenNote 7.3.1

NeedsCompilation yes

Author David Rowe [aut, cre]

Repository CRAN

Date/Publication 2024-10-25 07:30:06 UTC

Contents

graphicalEvidence-package	2
evidence	3
gen_params_evidence	4

graphical_evidence_G_Wishart	5
graphical_evidence_rmatrix	6
prior_sampling	7
set_seed_evidence	8
test_evidence	9
Index	10

graphicalEvidence-package

Compute Marginal Likelihood in Gaussian Graphical Models

Description

This package allows estimation of marginal likelihood in Gaussian graphical models through a novel telescoping block decomposition of the precision matrix which allows estimation of model evidence via an application of Chib's method. The currently implemented priors are: Bayesian graphical lasso (BGL), Graphical horseshoe (GHS), Wishart, and G-Wishart. The top level function used to estimate marginal likelihood is `evidence` which expects the prior name, data, and relevant prior specific parameters. This package also provides an MCMC prior sampler for the priors of BGL, GHS, and G-Wishart, implemented in `prior_sampling`, which expects a prior name and prior specific parameters. Both functions also expect the number of burnin iterations and the number of sampling iterations for the underlying MCMC sampler.

Bhadra, A., Sagar, K., Rowe, D., Banerjee, S., & Datta, J. (2022) "Evidence Estimation in Gaussian Graphical Models Using a Telescoping Block Decomposition of the Precision Matrix." <<https://arxiv.org/abs/2205.01016>>

Chib, S. "Marginal likelihood from the Gibbs output." (1995) <<https://www.jstor.org/stable/2291521>>

Details

This package implements marginal estimation for four priors, "Wishart", Bayesian Graphical Lasso ("BGL"), graphical horseshoe ("GHS"), and "G-Wishart". An MCMC prior sampler is also provided for "BGL", "GHS", and "G-Wishart".

For more information and a faster, less portable implementation, visit the package repository on GitHub: <https://github.com/dp-rho/graphicalEvidence>

Author(s)

Maintainer: David Rowe <david@rowe-stats.com>

References

Bhadra, A., Sagar, K., Rowe, D., Banerjee, S., & Datta, J. (2022) "Evidence Estimation in Gaussian Graphical Models Using a Telescoping Block Decomposition of the Precision Matrix." <<https://arxiv.org/abs/2205.01016>>

Chib, S. "Marginal likelihood from the Gibbs output." (1995) <<https://www.jstor.org/stable/2291521>>

See Also

[test_evidence](#): For basic example of functionality

[evidence](#): For top level estimation function

[prior_sampling](#): For the prior sampler function

Examples

```
test_results <- test_evidence(num_runs=3, prior_name='G-Wishart')
```

evidence

Compute Marginal Likelihood using Graphical Evidence

Description

Computes the marginal likelihood of input data `xx` under one of the following priors: Wishart, Bayesian Graphical Lasso (BGL), Graphical Horseshoe (GHS), and G-Wishart, specified under `prior_name`. The number of runs is specified by `num_runs`, where each run is by default using a random permutation of the columns of `xx`, as marginal likelihood should be independent of column permutation.

Usage

```
evidence(
  xx,
  burnin,
  nmc,
  prior_name = c("Wishart", "BGL", "GHS", "G-Wishart"),
  runs = 1,
  print_progress = FALSE,
  permute_columns = TRUE,
  alpha = NULL,
  lambda = NULL,
  V = NULL,
  G = NULL
)
```

Arguments

<code>xx</code>	The input data specified by a user for which the marginal likelihood is to be calculated. This should be input as a matrix like object with each individual sample of <code>xx</code> representing one row
<code>burnin</code>	The number of iterations the MCMC sampler should iterate through and discard before beginning to save results
<code>nmc</code>	The number of samples that the MCMC sampler should use to estimate marginal likelihood

prior_name	The name of the prior for which the marginal should be calculated, this is one of 'Wishart', 'BGL', 'GHS', 'G_Wishart'
runs	The number of complete runs of the graphical evidence method that will be executed. Specifying multiple runs allows estimation of the variance of the estimator and by default will permute the columns of xx such that each run uses a random column ordering, as marginal likelihood should be independent of column permutations
print_progress	A boolean which indicates whether progress should be displayed on the console as each row of the telescoping sum is computed and each run is completed
permute_columns	A boolean which indicates whether columns of xx for runs beyond the first should be randomly permuted to ensure that marginal calculation is consistent across different column permutations
alpha	A number specifying alpha for the priors of 'Wishart' and 'G_Wishart'
lambda	A number specifying lambda for the priors of 'BGL' and 'GHS' prior
V	The scale matrix when specifying 'Wishart' or 'G_Wishart' prior
G	The adjacency matrix when specifying 'G_Wishart' prior

Value

A list of results which contains the mean marginal likelihood, the standard deviation of the estimator, and the raw results in a vector

Examples

```
# Compute the marginal 10 times with random column permutations of xx at each
# individual run for G-Wishart prior using 2,000 burnin and 10,000 sampled
# values at each call to the MCMC sampler
g_params <- gen_params_evidence('G_Wishart')
marginal_results <- evidence(
  g_params$x_mat, 2e3, 1e4, 'G_Wishart', 3, alpha=2,
  V=g_params$scale_mat, G=g_params$g_mat
)
```

gen_params_evidence *Generate Test Parameters*

Description

Generates predetermined parameters for testing the functionality of the graphical evidence method

Usage

```
gen_params_evidence(prior_name = c("Wishart", "BGL", "GHS", "G_Wishart"))
```

Arguments

prior_name The name of the prior for being tested with preexisting test parameters, this is one of 'Wishart', 'BGL', 'GHS', 'G_Wishart'

Value

A list of matrices representing test parameters dependent on the prior specified in prior_name

Examples

```
# Generate test parameter matrices for G-Wishart prior
gen_params_evidence('G_Wishart')
```

```
graphical_evidence_G_Wishart
      Compute Marginal Likelihood using Graphical Evidence under G
      Wishart
```

Description

Computes the marginal likelihood of input data xx under G-Wishart prior using graphical evidence.

Usage

```
graphical_evidence_G_Wishart(
  xx,
  burnin,
  nmc,
  alpha,
  V,
  G,
  print_progress = FALSE
)
```

Arguments

xx The input data specified by a user for which the marginal likelihood is to be calculated. This should be input as a matrix like object with each individual sample of xx representing one row

burnin The number of iterations the MCMC sampler should iterate through and discard before beginning to save results

nmc The number of samples that the MCMC sampler should use to estimate marginal likelihood

alpha A number specifying alpha for G-Wishart prior

V The scale matrix of G-Wishart prior

G The adjacency matrix of G-Wishart prior

print_progress A boolean which indicates whether progress should be displayed on the console as each row of the telescoping sum is computed

Value

An estimate for the marginal likelihood under G-Wishart prior with the specified parameters

Examples

```
# Compute the marginal likelihood of xx for G-Wishart prior using
# 2,000 burnin and 10,000 sampled values at each call to the MCMC sampler
g_params <- gen_params_evidence('G_Wishart')
marginal_results <- graphical_evidence_G_Wishart(
  g_params$x_mat, 2e3, 1e4, 2, g_params$scale_mat, g_params$g_mat
)
```

graphical_evidence_rmatrix

Compute Marginal Likelihood using Graphical Evidence for Wishart, BGL, and GHS

Description

Computes the marginal likelihood of input data `xx` under one of the following priors: Wishart, Bayesian Graphical Lasso (BGL), and Graphical Horseshoe (GHS), specified under `prior_name`.

Usage

```
graphical_evidence_rmatrix(
  xx,
  burnin,
  nmc,
  prior_name = c("Wishart", "BGL", "GHS"),
  lambda = 0,
  alpha = 0,
  V = 0,
  print_progress = FALSE
)
```

Arguments

<code>xx</code>	The input data specified by a user for which the marginal likelihood is to be calculated. This should be input as a matrix like object with each individual sample of <code>xx</code> representing one row
<code>burnin</code>	The number of iterations the MCMC sampler should iterate through and discard before beginning to save results
<code>nmc</code>	The number of samples that the MCMC sampler should use to estimate marginal likelihood
<code>prior_name</code>	The name of the prior for which the marginal should be calculated, this is one of 'Wishart', 'BGL', 'GHS'

lambda	A number specifying lambda for the priors of 'BGL' and 'GHS' prior
alpha	A number specifying alpha for the priors of 'Wishart'
V	The scale matrix when specifying 'Wishart'
print_progress	A boolean which indicates whether progress should be displayed on the console as each row of the telescoping sum is computed

Value

An estimate for the marginal likelihood under specified prior with the specified parameters

Examples

```
# Compute the marginal likelihood of xx for GHS prior using 1,000
# burnin and 5,000 sampled values at each call to the MCMC sampler
g_params <- gen_params_evidence('GHS')
marginal_results <- graphical_evidence_rmatrix(
  g_params$x_mat, 1e3, 5e3, 'GHS', lambda=1
)
```

prior_sampling *Sample The Precision Matrix*

Description

Takes specified prior_name and relevant parameters to sample the precision matrix nmc times after discarding the first number of runs specified by burnin.

Usage

```
prior_sampling(
  p,
  burnin,
  nmc,
  prior_name = c("BGL", "GHS", "G_Wishart"),
  G = NULL,
  V = NULL,
  alpha = NULL,
  lambda = NULL
)
```

Arguments

p	The dimension of the precision matrix that will be sampled
burnin	The number of iterations the MCMC sampler should iterate through and discard before beginning to save results
nmc	The number of samples that will be drawn

prior_name	The name of the prior for which the marginal should be calculated, this is one of 'Wishart', 'BGL', 'GHS', 'G_Wishart'
G	The adjacency matrix when specifying 'G_Wishart' prior
V	The scale matrix when specifying 'Wishart' or 'G_Wishart' prior
alpha	A number specifying alpha for the priors of 'Wishart' and 'G_Wishart'
lambda	A number specifying lambda for the priors of 'BGL' and 'GHS' prior

Value

An array of dim nmc x p x p where each p x p slice is one sample of the precision matrix

Examples

```
# Draw 5000 samples of the precision matrix for GHS prior distribution with
# parameter lambda set to 1
prior_sampling(5, 1e3, 5e3, 'GHS', lambda=1)
```

set_seed_evidence *Set the Random Seed*

Description

Sets the random seed of both the R session (using set.seed) and the compiled sampler, as both samplers are used during any calls to evidence(...) or prior_sampling(...)

Usage

```
set_seed_evidence(seed)
```

Arguments

seed	a random seed that will be passed to the interpreted random number generator using set.seed, and will be passed to the compiled random number generator using private Rcpp package function set_seed
------	--

Value

No return value, this function is called to set the random seed.

Examples

```
set_seed_evidence(42)
```

test_evidence	<i>Test Graphical Evidence</i>
---------------	--------------------------------

Description

Tests any of the allowed priors with preexisting test inputs which should yield reproducible results, as the random seed is reset during this function call

Usage

```
test_evidence(  
  num_runs,  
  prior_name = c("Wishart", "BGL", "GHS", "G-Wishart"),  
  input_seed = NULL  
)
```

Arguments

num_runs	An integer number that specifies how many runs of graphical evidence will be executed on the test parameters, as multiple runs allows us to quantify uncertainty on the estimator.
prior_name	The name of the prior for being tested with preexisting test parameters, this is one of 'Wishart', 'BGL', 'GHS', 'G-Wishart'
input_seed	An integer value that will be used as a random seed to make outputs repeatable.

Value

A list of results which contains the mean marginal likelihood, the standard deviation of the estimator, and the raw results in a vector

Examples

```
# Compute the marginal 10 times with random column permutations of the  
# preexisting test parameters for G-Wishart prior  
test_evidence(num_runs=3, 'G-Wishart')
```

Index

* **package**

graphicalEvidence-package, 2

evidence, 2, 3, 3

gen_params_evidence, 4

graphical_evidence_G_Wishart, 5

graphical_evidence_rmatrix, 6

graphicalEvidence

(graphicalEvidence-package), 2

graphicalEvidence-package, 2

prior_sampling, 2, 3, 7

set_seed_evidence, 8

test_evidence, 3, 9