

Package ‘bgsmttr’

October 12, 2022

Title Bayesian Group Sparse Multi-Task Regression

Version 0.7

Description Implementation of Bayesian multi-task regression models and was developed within the context of imaging genetics. The package can currently fit two models. The Bayesian group sparse multi-task regression model of Greenlaw et al. (2017) <[doi:10.1093/bioinformatics/btx215](https://doi.org/10.1093/bioinformatics/btx215)> can be fit with implementation using Gibbs sampling. An extension of this model developed by Song, Ge et al. to accommodate both spatial correlation as well as correlation across brain hemispheres can also be fit using either mean-field variational Bayes or Gibbs sampling. The model can also be used more generally for multivariate (non-imaging) phenotypes with spatial correlation.

Author Yin Song, Shufei Ge, Liangliang Wang, Jiguo Cao, Keelin Greenlaw, Mary Lesperance, Farouk S. Nathoo

Maintainer Yin Song <yinsong@uvic.ca>

Depends R (>= 3.5.0), Matrix (>= 1.2.6), mvtnorm (>= 1.0.5), matrixcalc (>= 1.0.3), miscTools (>= 0.6.22)

Imports coda (>= 0.18.1), EDISON (>= 1.1.1), statmod (>= 1.4.26), methods (>= 3.3.3), sparseMVN (> 0.2.0), inline (>= 0.3.15), LaplacesDemon (>= 16.1.0), glmnet (>= 2.0.13), CholWishart (>= 0.9.3), mnormt (>= 1.5.4), Rcpp (>= 0.12.14),

License GPL-2

Encoding UTF-8

LazyData true

RoxygenNote 7.0.0

NeedsCompilation no

biocViews

Repository CRAN

Date/Publication 2019-12-13 05:20:03 UTC

R topics documented:

bgsmttr	2
bgsmttr_example_data	5
sp_bgsmttr	5
sp_bgsmttr_example_data	8
sp_bgsmttr_path	9

Index	11
--------------	-----------

bgsmttr	<i>Bayesian Group Sparse Multi-Task Regression for Imaging Genetics</i>
---------	---

Description

Runs the the Gibbs sampling algorithm to fit a Bayesian group sparse multi-task regression model. Tuning parameters can be chosen using either the MCMC samples and the WAIC (multiple runs) or using an approximation to the posterior mode and five-fold cross-validation (single run).

Usage

```
bgsmttr(
  X,
  Y,
  group,
  tuning = "CV.mode",
  lam_1_fixed = NULL,
  lam_2_fixed = NULL,
  iter_num = 10000,
  burn_in = 5001
)
```

Arguments

- | | |
|-------|---|
| X | A d-by-n matrix; d is the number of SNPs and n is the number of subjects. Each row of X should correspond to a particular SNP and each column should correspond to a particular subject. Each element of X should give the number of minor alleles for the corresponding SNP and subject. The function will center each row of X to have mean zero prior to running the Gibbs sampling algorithm. |
| Y | A c-by-n matrix; c is the number of phenotypes (brain imaging measures) and n is the number of subjects. Each row of Y should correspond to a particular phenotype and each column should correspond to a particular subject. Each element of Y should give the measured value for the corresponding phenotype and subject. The function will center and scale each row of Y to have mean zero and unit variance prior to running the Gibbs sampling algorithm. |
| group | A vector of length d; d is the number of SNPs. Each element of this vector is a string representing a gene or group label associated with each SNP. The SNPs represented by this vector should be ordered according to the rows of X. |

tuning	A string, either 'WAIC' or 'CV.mode'. If 'WAIC', the Gibbs sampler is run with fixed values of the tuning parameters specified by the arguments <i>lam_1_fixed</i> and <i>lam_2_fixed</i> and the WAIC is computed based on the sampling output. This can then be used to choose optimal values for <i>lam_1_fixed</i> and <i>lam_2_fixed</i> based on multiple runs with each run using different values of <i>lam_1_fixed</i> and <i>lam_2_fixed</i> . This option is best suited for either comparing a small set of tuning parameter values or for computation on a high performance computing cluster where different nodes can be used to run the function with different values of <i>lam_1_fixed</i> and <i>lam_2_fixed</i> . Posterior inference is then based on the run that produces the lowest value for the WAIC. The option 'CV.mode', which is the default, is best suited for computation using just a single processor. In this case the tuning parameters are chosen based on five-fold cross-validation over a grid of possible values with out-of-sample prediction based on an approximate posterior mode. The Gibbs sampler is then run using the chosen values of the tuning parameters. When tuning = 'CV.mode' the values for the arguments <i>lam_1_fixed</i> and <i>lam_2_fixed</i> are not required.
lam_1_fixed	Only required if tuning = 'WAIC'. A positive number giving the value for the gene-specific tuning parameter. Larger values lead to a larger degree of shrinkage to zero of estimated regression coefficients at the gene level (across all SNPs and phenotypes).
lam_2_fixed	Only required if tuning = 'WAIC'. A positive number giving the value for the SNP-specific tuning parameter. Larger values lead to a larger degree of shrinkage to zero of estimated regression coefficients at the SNP level (across all phenotypes).
iter_num	Positive integer representing the total number of iterations to run the Gibbs sampler. Defaults to 10,000.
burn_in	Nonnegative integer representing the number of MCMC samples to discard as burn-in. Defaults to 5001.

Value

A list with the elements

WAIC	If tuning = 'WAIC' this is the value of the WAIC computed from the MCMC output. If tuning = 'CV.mode' this component is excluded.
Gibbs_setup	A list providing values for the input parameters of the function.
Gibbs_W_summaries	A list with five components, each component being a d-by-c matrix giving some posterior summary of the regression parameter matrix W, where the ij-th element of W represents the association between the i-th SNP and j-th phenotype. - <i>Gibbs_W_summaries\$W_post_mean</i> is a d-by-c matrix giving the posterior mean of W. - <i>Gibbs_W_summaries\$W_post_mode</i> is a d-by-c matrix giving the posterior mode of W. - <i>Gibbs_W_summaries\$W_post_sd</i> is a d-by-c matrix giving the posterior standard deviation for each element of W.

-Gibbs_W_summaries\$W_2.5_quantile is a d-by-c matrix giving the posterior 2.5 percent quantile for each element of W.

-Gibbs_W_summaries\$W_97.5_quantile is a d-by-c matrix giving the posterior 97.5 percent quantile for each element of W.'

Author(s)

Farouk S. Nathoo, <nathoo@uvic.ca>

Keelin Greenlaw, <keelinggreenlaw@gmail.com>

Mary Lesperance, <mlespera@uvic.ca>

References

Greenlaw, Keelin, Elena Szefer, Jinko Graham, Mary Lesperance, and Farouk S. Nathoo. "A Bayesian Group Sparse Multi-Task Regression Model for Imaging Genetics." arXiv preprint arXiv:1605.02234 (2016).

Nathoo, Farouk S., Keelin Greenlaw, and Mary Lesperance. "Regularization Parameter Selection for a Bayesian Multi-Level Group Lasso Regression Model with Application to Imaging Genomics." arXiv preprint arXiv:1603.08163 (2016).

Examples

```
data(bgsmttr_example_data)
names(bgsmttr_example_data)

## Not run:
## test run the sampler for 100 iterations with fixed tuning parameters and compute WAIC
## we recommend at least 5,000 iterations for actual use
fit = bgsmttr(X = bgsmttr_example_data$SNP_data, Y = bgsmttr_example_data$BrainMeasures,
group = bgsmttr_example_data$SNP_groups, tuning = 'WAIC', lam_1_fixed = 2, lam_2_fixed = 2,
iter_num = 100, burn_in = 50)
## posterior mean for regression parameter relating 100th SNP to 14th phenotype
fit$Gibbs_W_summaries$W_post_mean[100,14]
## posterior mode for regression parameter relating 100th SNP to 14th phenotype
fit$Gibbs_W_summaries$W_post_mode[100,14]
## posterior standard deviation for regression parameter relating 100th SNP to 14th phenotype
fit$Gibbs_W_summaries$W_post_sd[100,14]
## 95
c(fit$Gibbs_W_summaries$W_2.5_quantile[100,14],fit$Gibbs_W_summaries$W_97.5_quantile[100,14])

## End(Not run)

## Not run:
## run the sampler for 10,000 iterations with tuning parameters set using cross-validation
## On a standard computer with a small number of cores this is the recommended option
fit = bgsmttr(X = bgsmttr_example_data$SNP_data, Y = bgsmttr_example_data$BrainMeasures,
group = bgsmttr_example_data$SNP_groups, tuning = 'CV.mode', iter_num = 10000, burn_in = 5000)

## End(Not run)
```

bgsmttr_example_data *Example Structural Neuroimaging and Genetic Data*

Description

Simulated data with 632 subjects, 486 SNPs from 33 genes, 15 structural neuroimaging measures.

Usage

```
data(bgsmttr_example_data)
```

Format

A list with three components: "SNP_data", "SNP_groups", "BrainMeasures". SNP_data is a 486-by-632 matrix containing minor allele counts for 632 subjects and 486 SNPs. SNP_groups is a vector of length 486 with labels partitioning the 486 SNPs into 33 genes. BrainMeasures is a 15-by-632 matrix containing simulated volumetric and cortical thickness measures for 15 regions of interest.

Examples

```
data(bgsmttr_example_data)
names(bgsmttr_example_data)
dim(bgsmttr_example_data$SNP_data)
dim(bgsmttr_example_data$BrainMeasures)
unique(bgsmttr_example_data$SNP_groups)
```

sp_bgsmttr *A Bayesian Spatial Model for Imaging Genetics*

Description

Fits a Bayesian spatial model that allows for two types of correlation typically seen in structural brain imaging data. First, the spatial correlation in the imaging phenotypes obtained from neighbouring regions of the brain. Second, the correlation between corresponding measures on opposite hemispheres.

Usage

```
sp_bgsmttr(
  X,
  Y,
  method = "MCMC",
  rho = NULL,
  lambdasq = NULL,
  alpha = NULL,
```

```

A = NULL,
c.star = NULL,
FDR_opt = TRUE,
WAIC_opt = TRUE,
iter_num = 10000,
burn_in = 5001
)

```

Arguments

X	A d-by-n matrix; d is the number of SNPs and n is the number of subjects. Each row of X should correspond to a particular SNP and each column should correspond to a particular subject. Each element of X should give the number of minor alleles for the corresponding SNP and subject. The function will center each row of X to have mean zero prior to running the Gibbs sampling algorithm.
Y	A c-by-n matrix; c is the number of phenotypes (brain imaging measures) and n is the number of subjects. Each row of Y should correspond to a particular phenotype and each column should correspond to a particular subject. Each element of Y should give the measured value for the corresponding phenotype and subject. The function will center and scale each row of Y to have mean zero and unit variance prior to running the Gibbs sampling algorithm.
method	A string, either 'MCMC' or 'MFVB'. If 'MCMC', the Gibbs sampling method will be used. If 'MFVB', mean field variational Bayes method will be used.
rho	Spatial cohesion parameter. If no value is assigned, it takes 0.95 by default.
lambdasq	A tuning parameter. If no value is assigned, the algorithm will estimate and assign a value for it based on a moment estimate.
alpha	Bayesian False Discovery Rate (FDR) level. Default level is 0.05.
A	A c/2 by c/2 neighborhood structure matrix for different brain regions.
c.star	The threshold for computing posterior tail probabilities p_{ij} for Bayesian FDR as defined in Section 3.2 of Song, Ge et al.(2019). If not specified the default is to set this threshold as the minimum posterior standard deviation, where the minimum is taken over all regression coefficients in the model.
FDR_opt	A logical operator for computing Bayesian FDR. By default, it's TRUE.
WAIC_opt	A logical operator for computing WAIC from MCMC method. By default, it's TRUE.
iter_num	Positive integer representing the total number of iterations to run the Gibbs sampler. Defaults to 10,000.
burn_in	Nonnegative integer representing the number of MCMC samples to discard as burn-in. Defaults to 5001.

Value

A list with following elements

WAIC	WAIC is computed from the MCMC output if "MCMC" is chosen for method.
lower_boud	Lower bound from MFVB output if "MFVB" is chosen for method.

- Gibbs_setup A list providing values for the input parameters of the function.
- lambdasq_est Estimated value for tuning parameter lambda-squared.
- Gibbs_W_summaries
 A list with five components, each component being a d-by-c matrix giving some posterior summary of the regression parameter matrix W, where the ij-th element of W represents the association between the i-th SNP and j-th phenotype.
 -Gibbs_W_summaries\$W_post_mean is a d-by-c matrix giving the posterior mean of W.
 -Gibbs_W_summaries\$W_post_mode is a d-by-c matrix giving the posterior mode of W.
 -Gibbs_W_summaries\$W_post_sd is a d-by-c matrix giving the posterior standard deviation for each element of W.
 -Gibbs_W_summaries\$W_2.5_quantile is a d-by-c matrix giving the posterior 2.5 percent quantile for each element of W.
 -Gibbs_W_summaries\$W_97.5_quantile is a d-by-c matrix giving the posterior 97.5 percent quantile for each element of W.
- FDR_summaries A list with three components providing the summaries for estimated Bayesian FDR results for both MCMC and MFVB methods. Details for Bayesian FDR computation could be found at Morris et al.(2008).
 -sensitivity_rate is the estimated sensitivity rate for each region.
 -specificity_rate is the estimated specificity rate for each region.
 -significant_snp_idx is the index of estimated significant/important SNPs for each region.
- MFVB_summaries A list with four components, each component is the mean field variational bayes approximation summary of model paramters.
 -Number of Iteration is how many iterations it takes for convergence.
 -W_post_mean is MFVB approximation of W.
 -Sigma_post_mean is MFVB approximation of Sigma.
 -omega_post_mean is MFVB approximation of Omega.

Author(s)

Yin Song, <yinsong@uvic.ca>
 Shufei Ge, <shufeig@sfu.ca>
 Farouk S. Nathoo, <nathoo@uvic.ca>
 Liangliang Wang, <lwa68@sfu.ca>
 Jiguo Cao, <jiguo_cao@sfu.ca>

References

Song, Y., Ge, S., Cao, J., Wang, L., Nathoo, F.S., A Bayesian Spatial Model for Imaging Genetics. arXiv:1901.00068.

Examples

```

data(sp_bgsmttr_example_data)
names(sp_bgsmttr_example_data)

## Not run:

# Run the example data with Gibbs sampling and compute Bayesian FDR as follow:

fit_mcmc = sp_bgsmttr(X = sp_bgsmttr_example_data$SNP_data,
Y = sp_bgsmttr_example_data$BrainMeasures, method = "MCMC",
A = sp_bgsmttr_example_data$neighborhood_structure, rho = 0.8,
FDR_opt = TRUE, WAIC_opt = TRUE, lambdasq = 1000, iter_num = 10000)

# MCMC estimation results for regression parameter W and estimated Bayesian FDR summaries

fit_mcmc$Gibbs_W_summaries
fit_mcmc$FDR_summaries

# The WAIC could be also obtained as:

fit_mcmc$WAIC

# Run the example data with mean field variational Bayes and compute Bayesian FDR as follow:

fit_mfvb = sp_bgsmttr(X = sp_bgsmttr_example_data$SNP_data,
Y = sp_bgsmttr_example_data$BrainMeasures, method = "MFVB",
A = sp_bgsmttr_example_data$neighborhood_structure, rho = 0.8, FDR_opt = FALSE,
lambdasq = 1000, iter_num = 10000)

# MFVB estimated results for regression parameter W and estimated Bayesian FDR summaries
fit_mfvb$MFVB_summaries
fit_mfvb$FDR_summaries

# The corresponding lower bound of MFVB method after convergence is obtained as:
fit_mfvb$lower_boud

## End(Not run)

```

sp_bgsmttr_example_data

Example Structural Neuroimaging and Genetic Data for Spatial Model.

Description

This example dataset has simulated dataset with 632 subjects, 486 SNPs from 24 structural neuroimaging measures. It also contains example datasets used for regularization path plotting function.

Usage

```
data(sp_bgsmttr_example_data)
```

Format

A list with four components: "SNP_data", "SNP_groups", "BrainMeasures", "path_data" SNP_data is a 486-by-632 matrix containing minor allele counts for 632 subjects and 486 SNPs. neighbourhood_structure is a 12 by 12 first order neighbourhood structure matrix. BrainMeasures is a 24-by-632 matrix containing simulated volumetric and cortical thickness measures for 24 regions of interest.

path_data is a list of two elements. The first element is "lambda_v", which is a vector of different lambda squared values. The second element is "W_est_list", which is a list containing estimated coefficients matrices W for each corresponding lambda squared value.

Examples

```
data(sp_bgsmttr_example_data)
names(sp_bgsmttr_example_data)
dim(sp_bgsmttr_example_data$SNP_data)
dim(sp_bgsmttr_example_data$BrainMeasures)
dim(sp_bgsmttr_example_data$neighbourhood_structure)
length(sp_bgsmttr_example_data$path_data$lambda_v)
length(sp_bgsmttr_example_data$path_data$W_est_list)
```

sp_bgsmttr_path

A Bayesian Spatial Model for Imaging Genetics

Description

A plotting function can be used to demonstrate the regularization paths for estimating parameters of each ROI when the spatial model is fitted with multiple values of tuning parameter lambda-squared.

Usage

```
sp_bgsmttr_path(lambda_v, W_est_list)
```

Arguments

lambda_v	A vector containing all the different tuning parameter lambda-squared values for model fitting.
W_est_list	A list containing all the estimated coefficients matrices W for each corresponding lambda squared value used in lambda_v for model fitting. Each element of this list is a d-by-c matrix.

Value

Regularization plots files in PDF format for each ROI.

Author(s)

Yin Song, <yinsong@uvic.ca>

Shufei Ge, <shufeig@sfu.ca>

Farouk S. Nathoo, <nathoo@uvic.ca>

Liangliang Wang, <lwa68@sfu.ca>

Jiguo Cao, <jiguo_cao@sfu.ca>

References

Song, Y., Ge, S., Cao, J., Wang, L., Nathoo, F.S., A Bayesian Spatial Model for Imaging Genetics. arXiv:1901.00068.

Examples

```
data(sp_bgsmttr_example_data$path_data)
```

```
## Not run:
```

```
# Creating the regulazaition path plots as follow:
```

```
sp_bgsmttr_path(lambda_v = sp_bgsmttr_example_data$path_data$lambda_v,  
W_est_list = sp_bgsmttr_example_data$path_data$W_est_list )
```

```
## End(Not run)
```

Index

* datasets

- bgsmtr_example_data, [5](#)
- sp_bgsmtr_example_data, [8](#)

- bgsmtr, [2](#)
- bgsmtr_example_data, [5](#)

- sp_bgsmtr, [5](#)
- sp_bgsmtr_example_data, [8](#)
- sp_bgsmtr_path, [9](#)