

Package ‘spcov’

July 23, 2025

Type Package

Title Sparse Estimation of a Covariance Matrix

Version 1.3

Date 2022-09-21

Author Jacob Bien and Rob Tibshirani

Maintainer Jacob Bien <jbien@usc.edu>

Description Provides a covariance estimator for multivariate normal data that is sparse and positive definite. Implements the majorize-minimize algorithm described in Bien, J., and Tibshirani, R. (2011), ``Sparse Estimation of a Covariance Matrix," Biometrika. 98(4). 807--820.

License GPL-2

LazyLoad yes

Encoding UTF-8

RoxygenNote 7.2.1

NeedsCompilation no

Repository CRAN

Date/Publication 2022-09-23 20:20:02 UTC

Contents

spcov-package	2
GenerateCliquesCovariance	3
ProxADMM	4
spcov	5

Index	9
--------------	----------

spcov-package

Sparse Estimation of a Covariance Matrix

Description

Performs the majorize-minimize algorithm described in

Details

Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," *Biometrika*. 98(4). 807–820.

Generalized gradient descent is used to minimize each majorizer.

Package:	spcov
Type:	Package
Version:	1.01
Date:	2012-03-04
License:	GPL-2
LazyLoad:	yes

See the function `spcov`.

Author(s)

Jacob Bien and Rob Tibshirani

Maintainer: Jacob Bien <jbien@stanford.edu>

References

Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," *Biometrika*. 98(4). 807–820.

See Also

[spcov](#)

`GenerateCliquesCovariance`*Generate a block diagonal covariance matrix*

Description

This function is included in the package so that it can be used in the example code provided in [spcov](#).

Usage

```
GenerateCliquesCovariance(ncliques, cliquesize, theta)
```

Arguments

<code>ncliques</code>	number of blocks
<code>cliquesize</code>	size of each block
<code>theta</code>	magnitude of non-zeros

Details

This function generates a block diagonal positive definite matrix with randomly-signed, non-zero elements. A shift is added to the diagonal of the matrix so that its condition number equals p , the number of variables.

Value

<code>Sigma</code>	the covariance matrix
<code>A</code>	symmetric square root of <code>Sigma</code>
<code>shift</code>	how much the eigenvalues were shifted. See details.

Author(s)

Jacob Bien and Rob Tibshirani

References

Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," accepted for publication in *Biometrika*.

ProxADMM

*Solving penalized Frobenius problem.***Description**

This function solves the optimization problem

Usage

```
ProxADMM(A, del, lam, P, rho = 0.1, tol = 1e-06, maxiters = 100, verb = FALSE)
```

Arguments

A	A symmetric matrix.
del	A non-negative scalar. Lower bound on eigenvalues.
lam	A non-negative scalar. L1 penalty parameter.
P	Matrix with non-negative elements and dimension of A. Allows for differing L1 penalty parameters.
rho	ADMM parameter. Can affect rate of convergence a lot.
tol	Convergence threshold.
maxiters	Maximum number of iterations.
verb	Controls whether to be verbose.

Details

Minimize $X (1/2)\|X - A\|_F^2 + \text{lam}\|P*X\|_1$ s.t. $X \geq \text{del} * I$.

This is the prox function for the generalized gradient descent of Bien & Tibshirani 2011 (see full reference below).

This is the R implementation of the algorithm in Appendix 3 of Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," *Biometrika*. 98(4). 807–820. It uses an ADMM approach to solve the problem

Minimize $X (1/2)\|X - A\|_F^2 + \text{lam}\|P*X\|_1$ s.t. $X \geq \text{del} * I$.

Here, the multiplication between P and X is elementwise. The inequality in the constraint is a lower bound on the minimum eigenvalue of the matrix X.

Note that there are two variables X and Z that are outputted. Both are estimates of the optimal X. However, Z has exact zeros whereas X has eigenvalues at least del. Running the ADMM algorithm long enough, these two are guaranteed to converge.

Value

X	Estimate of optimal X.
Z	Estimate of optimal X.
obj	Objective values.

Author(s)

Jacob Bien and Rob Tibshirani

References

Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," *Biometrika*. 98(4). 807–820.

See Also

spcov

Examples

```
set.seed(1)
n <- 100
p <- 200
# generate a covariance matrix:
model <- GenerateCliquesCovariance(ncliques=4, cliquesize=p / 4, 1)

# generate data matrix with  $x[i, ] \sim N(0, \text{model}\$Sigma)$ :
x <- matrix(rnorm(n * p), ncol=p) %*% model$A
S <- var(x)

# compute sparse, positive covariance estimator:
P <- matrix(1, p, p)
diag(P) <- 0
lam <- 0.1
aa <- ProxADMM(S, 0.01, lam, P)
```

spcov

Sparse Covariance Estimation

Description

Provides a sparse and positive definite estimate of a covariance matrix. This function performs the majorize-minimize algorithm described in Bien & Tibshirani 2011 (see full reference below).

Usage

```
spcov(
  Sigma,
  S,
  lambda,
  step.size,
  nesterov = TRUE,
  n.outer.steps = 10000,
```

```

n.inner.steps = 10000,
tol.outer = 1e-04,
thr.inner = 0.01,
backtracking = 0.2,
trace = 0
)

```

Arguments

<code>Sigma</code>	an initial guess for Sigma (suggestions: <code>S</code> or <code>diag(diag(S))</code>).
<code>S</code>	the empirical covariance matrix of the data. Must be positive definite (if it is not, add a small constant to the diagonal).
<code>lambda</code>	penalty parameter. Either a scalar or a matrix of the same dimension as Sigma. This latter choice should be used to penalize only off-diagonal elements. All elements of <code>lambda</code> must be non-negative.
<code>step.size</code>	the step size to use in generalized gradient descent. Affects speed of algorithm.
<code>nesterov</code>	indicates whether to use Nesterov's modification of generalized gradient descent. Default: <code>TRUE</code> .
<code>n.outer.steps</code>	maximum number of majorize-minimize steps to take (recall that MM is the outer loop).
<code>n.inner.steps</code>	maximum number of generalized gradient steps to take (recall that generalized gradient descent is the inner loop).
<code>tol.outer</code>	convergence threshold for outer (MM) loop. Stops when drop in objective between steps is less than <code>tol.outer</code> .
<code>thr.inner</code>	convergence threshold for inner (i.e. generalized gradient) loop. Stops when mean absolute change in Sigma is less than <code>thr.inner * mean(abs(S))</code> .
<code>backtracking</code>	if <code>FALSE</code> , then fixed step size used. If numeric and in $(0,1)$, this is the parameter of backtracking that multiplies <code>step.size</code> on each step. Usually, in range of $(0.1, 0.8)$. Default: <code>0.2</code> .
<code>trace</code>	controls how verbose output should be.

Details

This is the R implementation of Algorithm 1 in Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," *Biometrika*. 98(4). 807–820. The goal is to approximately minimize (over Sigma) the following non-convex optimization problem:

minimize $\log\det(\text{Sigma}) + \text{trace}(S \text{Sigma}^{-1}) + \|\text{lambda} * \text{Sigma}\|_1$ subject to Sigma positive definite.

Here, the L1 norm and matrix multiplication between `lambda` and Sigma are elementwise. The empirical covariance matrix must be positive definite for the optimization problem to have bounded objective (see Section 3.3 of paper). We suggest adding a small constant to the diagonal of `S` if it is not. Since the above problem is not convex, the returned matrix is not guaranteed to be a global minimum of the problem.

In Section 3.2 of the paper, we mention a simple modification of gradient descent due to Nesterov. The argument `nesterov` controls whether to use this modification (we suggest that it be used). We

also strongly recommend using backtracking. This allows the algorithm to begin by taking large steps (the initial size is determined by the argument `step.size`) and then to gradually reduce the size of steps.

At the start of the algorithm, a lower bound (`delta` in the paper) on the eigenvalues of the solution is calculated. As shown in Equation (3) of the paper, the prox function for our generalized gradient descent amounts to minimizing (over a matrix X) a problem of the form

minimize $(1/2)\|X-A\|_F^2 + \|\lambda X\|_1$ subject to $X \succeq \delta I$

This is implemented using an alternating direction method of multipliers approach given in Appendix 3.

Value

<code>Sigma</code>	the sparse covariance estimate
<code>n.iter</code>	a vector giving the number of generalized gradient steps taken on each step of the MM algorithm
<code>obj</code>	a vector giving the objective values after each step of the MM algorithm

Author(s)

Jacob Bien and Rob Tibshirani

References

Bien, J., and Tibshirani, R. (2011), "Sparse Estimation of a Covariance Matrix," *Biometrika*. 98(4). 807–820.

See Also

ProxADMM

Examples

```
set.seed(1)
n <- 100
p <- 20
# generate a covariance matrix:
model <- GenerateCliquesCovariance(ncliques=4, cliquesize=p / 4, 1)

# generate data matrix with x[i, ] ~ N(0, model$Sigma):
x <- matrix(rnorm(n * p), ncol=p) %*% model$A
S <- var(x)

# compute sparse, positive covariance estimator:
step.size <- 100
tol <- 1e-3
P <- matrix(1, p, p)
diag(P) <- 0
lam <- 0.06
mm <- spcov(Sigma=S, S=S, lambda=lam * P,
            step.size=step.size, n.inner.steps=200,
```

```
        thr.inner=0, tol.outer=tol, trace=1)
sqrt(mean((mm$Sigma - model$Sigma)^2))
sqrt(mean((S - model$Sigma)^2))
## Not run: image(mm$Sigma!=0)
```


Index

* **multivariate**

GenerateCliquesCovariance, [3](#)

ProxADMM, [4](#)

spcov, [5](#)

spcov-package, [2](#)

GenerateCliquesCovariance, [3](#)

ProxADMM, [4](#)

spcov, [2](#), [3](#), [5](#)

spcov-package, [2](#)