Package 'fitHeavyTail'

January 7, 2020

Title Mean and Covariance Matrix Estimation under Heavy Tails

Version 0.1.2

Date 2020-1-7

Description Robust estimation methods for the mean vector and covariance matrix from data (possibly containing NAs) under multivariate heavy-tailed distributions such as angular Gaussian (via Tyler's method), Cauchy, and Student's t.
Additionally, a factor model structure can be specified for the covariance matrix.
The package is based on the papers: Sun, Babu, and Palomar (2014), Sun, Babu, and Palomar (2015), Liu and Rubin (1995), and Zhou, Liu, Kumar, and Palomar (2019).

Maintainer Daniel P. Palomar <daniel.p.palomar@gmail.com>

URL https://github.com/dppalomar/fitHeavyTail

BugReports https://github.com/dppalomar/fitHeavyTail/issues

License GPL-3

Encoding UTF-8

LazyData true

RoxygenNote 7.0.2

Depends

Imports ICSNP, mvtnorm, stats

Suggests knitr, ggplot2, prettydoc, reshape2, rmarkdown, R.rsp, testthat

VignetteBuilder knitr, rmarkdown, R.rsp

NeedsCompilation no

Author Daniel P. Palomar [cre, aut], Rui Zhou [aut]

Repository CRAN

Date/Publication 2020-01-07 10:20:02 UTC

R topics documented:

`fitHeavyTail-package` *fitHeavyTail: Mean and Covariance Matrix Estimation under Heavy Tails*

Description

Robust estimation methods for the mean vector and covariance matrix from data (possibly containing NAs) under multivariate heavy-tailed distributions such as angular Gaussian (via Tyler's method), Cauchy, and Student's t. Additionally, a factor model structure can be specified for the covariance matrix.

Functions

`fit_Tyler`, `fit_Cauchy`, and `fit_mvt`

Help

For a quick help see the README file: GitHub-README.

For more details see the vignette: CRAN-vignette.

Author(s)

Daniel P. Palomar and Rui Zhou

References

Ying Sun, Prabhu Babu, and Daniel P. Palomar, "Regularized Tyler's Scatter Estimator: Existence, Uniqueness, and Algorithms," IEEE Trans. on Signal Processing, vol. 62, no. 19, pp. 5143-5156, Oct. 2014. <https://doi.org/10.1109/TSP.2014.2348944>

Ying Sun, Prabhu Babu, and Daniel P. Palomar, "Regularized Robust Estimation of Mean and Covariance Matrix Under Heavy-Tailed Distributions," IEEE Trans. on Signal Processing, vol. 63, no. 12, pp. 3096-3109, June 2015. <https://doi.org/10.1109/TSP.2015.2417513>

Chuanhai Liu and Donald B. Rubin, "ML estimation of the t-distribution using EM and its extensions, ECM and ECME," Statistica Sinica (5), pp. 19-39, 1995.

Rui Zhou, Junyan Liu, Sandeep Kumar, and Daniel P. Palomar, "Robust factor analysis parameter estimation," Lecture Notes in Computer Science (LNCS), 2019. <https://arxiv.org/abs/1909.12530>

`fit_Cauchy` *Estimate parameters of a multivariate elliptical distribution to fit data under a Cauchy distribution*

Description

Estimate parameters of a multivariate elliptical distribution, namely, the mean vector and the covariance matrix, to fit data. Any data sample with NAs will be simply dropped. The estimation is based on the maximum likelihood estimation (MLE) under a Cauchy distribution and the algorithm is obtained from the majorization-minimization (MM) optimization framework. The Cauchy distribution does not have second-order moments and the algorithm actually estimates the scatter matrix. Nevertheless, assuming that the observed data has second-order moments, the covariance matrix is returned by computing the missing scaling factor with a very effective method.

Usage

```
fit_Cauchy(
  X,
  initial = NULL,
  max_iter = 100,
  ptol = 0.001,
  ftol = Inf,
  return_iterates = FALSE,
  verbose = FALSE
)
```

Arguments

`X`	Data matrix containing the multivariate time series (each column is one time series).
`initial`	List of initial values of the parameters for the iterative estimation method. Possible elements include: • `mu`: default is the data sample mean, • `cov`: default is the data sample covariance matrix, • `scatter`: default follows from the scaled sample covariance matrix.
`max_iter`	Integer indicating the maximum number of iterations for the iterative estimation method (default is `100`).
`ptol`	Positive number indicating the relative tolerance for the change of the variables to determine convergence of the iterative method (default is `1e-3`).
`ftol`	Positive number indicating the relative tolerance for the change of the log-likelihood value to determine convergence of the iterative method (default is `Inf`, so it is not active). Note that using this argument might have a computational cost as a convergence criterion due to the computation of the log-likelihood (especially when `X` is high-dimensional).

`return_iterates`
: Logical value indicating whether to record the values of the parameters (and possibly the log-likelihood if `ftol < Inf`) at each iteration (default is `FALSE`).

`verbose`
: Logical value indicating whether to allow the function to print messages (default is `FALSE`).

Value

A list containing possibly the following elements:

`mu`
: Mean vector estimate.

`cov`
: Covariance matrix estimate.

`scatter`
: Scatter matrix estimate.

`converged`
: Boolean denoting whether the algorithm has converged (`TRUE`) or the maximum number of iterations `max_iter` has reached (`FALSE`).

`num_iterations`
: Number of iterations executed.

`cpu_time`
: Elapsed CPU time.

`log_likelihood`
: Value of log-likelihood after converge of the estimation algorithm (if `ftol < Inf`).

`iterates_record`
: Iterates of the parameters (`mu`, `scatter`, and possibly `log_likelihood` (if `ftol < Inf`)) along the iterations (if `return_iterates = TRUE`).

Author(s)

Daniel P. Palomar

References

Ying Sun, Prabhu Babu, and Daniel P. Palomar, "Regularized Robust Estimation of Mean and Covariance Matrix Under Heavy-Tailed Distributions," IEEE Trans. on Signal Processing, vol. 63, no. 12, pp. 3096-3109, June 2015.

See Also

`fit_Tyler` and `fit_mvt`

Examples

```
library(mvtnorm)       # to generate heavy-tailed data
library(fitHeavyTail)

X <- rmvt(n = 1000, df = 6)  # generate Student's t data
fit_Cauchy(X)
```

`fit_mvt` *Estimate parameters of a multivariate Student's t distribution to fit data*

Description

Estimate parameters of a multivariate Student's t distribution to fit data, namely, the mean vector, the covariance matrix, the scatter matrix, and the degrees of freedom. The data can contain missing values denoted by NAs. It can also consider a factor model structure on the covariance matrix. The estimation is based on the maximum likelihood estimation (MLE) and the algorithm is obtained from the expectation-maximization (EM) method.

Usage

```
fit_mvt(
  X,
  na_rm = TRUE,
  nu = c("kurtosis", "MLE-diag", "MLE-diag-resampled", "iterative"),
  nu_iterative_method = c("ECME-diag", "ECME", "ECM", "ECME-cov"),
  initial = NULL,
  factors = ncol(X),
  max_iter = 100,
  ptol = 0.001,
  ftol = Inf,
  return_iterates = FALSE,
  verbose = FALSE
)
```

Arguments

`X` Data matrix containing the multivariate time series (each column is one time series).

`na_rm` Logical value indicating whether to remove observations with some NAs (default) or not, in which case they will be imputed at a higher computational cost.

`nu` Degrees of freedom of the t distribution. Either a number (`>2`) or a string indicating the method to compute it:

- `"kurtosis"`: based on the kurtosis obtained from the sampled moments;
- `"MLE-diag"`: based on the MLE assuming a diagonal sample covariance;
- `"MLE-diag-resampled"`: method "MLE-diag" resampled for better stability;
- `"iterative"`: iterative estimation with the rest of the parameters via the EM algorithm.

`nu_iterative_method` String indicating the method for iteratively estimating `nu` (in case `nu = "iterative"`):

- `"ECM"`: maximization of the Q function;

- "ECME": maximization of the log-likelihood function;
- "ECME-diag": maximization of the log-likelihood function assuming a diagonal scatter matrix (default method).

This argument is used only when there are no NAs in the data and no factor model is chosen.

initial — List of initial values of the parameters for the iterative EM estimation method (in case nu = "iterative"). Possible elements include:

- mu: default is the data sample mean,
- cov: default is the data sample covariance matrix,
- scatter: default follows from the scaled sample covariance matrix,
- nu: can take the same values as argument nu, default is 4,
- B: default is the top eigenvectors of initial$cov multiplied by the sqrt of the eigenvalues,
- psi: default is diag(initial$cov -initial$B %*% t(initial$B)).

factors — Integer indicating number of factors (default is ncol(X), so no factor model assumption).

max_iter — Integer indicating the maximum number of iterations for the iterative estimation method (default is 100).

ptol — Positive number indicating the relative tolerance for the change of the variables to determine convergence of the iterative method (default is 1e-3).

ftol — Positive number indicating the relative tolerance for the change of the log-likelihood value to determine convergence of the iterative method (default is Inf, so it is not active). Note that using this argument might have a computational cost as a convergence criterion due to the computation of the log-likelihood (especially when X is high-dimensional).

return_iterates — Logical value indicating whether to record the values of the parameters (and possibly the log-likelihood if ftol < Inf) at each iteration (default is FALSE).

verbose — Logical value indicating whether to allow the function to print messages (default is FALSE).

Details

This function estimates the parameters of a multivariate Student's t distribution (mu, cov, scatter, and nu) to fit the data via the expectation–maximization (EM) algorithm. The data matrix X can contain missing values denoted by NAs. The estimation of nu if very flexible: it can be directly passed as an argument (without being estimated), it can be estimated with several one-shot methods (namely, "kurtosis", "MLE-diag", "MLE-diag-resampled"), and it can also be iteratively estimated with the other parameters via the EM algorithm.

Value

A list containing possibly the following elements:

mu — Mean vector estimate.

cov — Covariance matrix estimate.

scatter	Scatter matrix estimate.
nu	Degrees of freedom estimate.
converged	Boolean denoting whether the algorithm has converged (TRUE) or the maximum number of iterations max_iter has been reached (FALSE).
num_iterations	Number of iterations executed.
cpu_time	Elapsed CPU time.
B	Factor model loading matrix estimate according to cov = (B %*% t(B) + diag(psi) (only if factor model requested).
psi	Factor model idiosynchratic variances estimates according to cov = (B %*% t(B) + diag(psi) (only if factor model requested).
log_likelihood	Value of log-likelihood after converge of the estimation algorithm (if ftol < Inf).
iterates_record	Iterates of the parameters (mu, scatter, nu, and possibly log_likelihood (if ftol < Inf)) along the iterations (if return_iterates = TRUE).

Author(s)

Daniel P. Palomar and Rui Zhou

References

Chuanhai Liu and Donald B. Rubin, "ML estimation of the t-distribution using EM and its extensions, ECM and ECME," Statistica Sinica (5), pp. 19-39, 1995.

Rui Zhou, Junyan Liu, Sandeep Kumar, and Daniel P. Palomar, "Robust factor analysis parameter estimation," Lecture Notes in Computer Science (LNCS), 2019. <https://arxiv.org/abs/1909.12530>

See Also

fit_Tyler and fit_Cauchy

Examples

```
library(mvtnorm)       # to generate heavy-tailed data
library(fitHeavyTail)

X <- rmvt(n = 1000, df = 6)  # generate Student's t data
fit_mvt(X)
```

`fit_Tyler` *Estimate parameters of a multivariate elliptical distribution to fit data via Tyler's method*

Description

Estimate parameters of a multivariate elliptical distribution, namely, the mean vector and the covariance matrix, to fit data. Any data sample with NAs will be simply dropped. The algorithm is based on Tyler's method, which normalizes the centered samples to get rid of the shape of the distribution tail. The data is first demeaned (with the geometric mean by default) and normalized. Then the estimation is based on the maximum likelihood estimation (MLE) and the algorithm is obtained from the majorization-minimization (MM) optimization framework. Since Tyler's method can only estimate the covariance matrix up to a scaling factor, a very effective method is employed to recover the scaling factor.

Usage

```
fit_Tyler(
  X,
  initial = NULL,
  max_iter = 100,
  ptol = 0.001,
  ftol = Inf,
  return_iterates = FALSE,
  verbose = FALSE
)
```

Arguments

`X`	Data matrix containing the multivariate time series (each column is one time series).
`initial`	List of initial values of the parameters for the iterative estimation method. Possible elements include: • `mu`: default is the data sample mean, • `cov`: default is the data sample covariance matrix.
`max_iter`	Integer indicating the maximum number of iterations for the iterative estimation method (default is `100`).
`ptol`	Positive number indicating the relative tolerance for the change of the variables to determine convergence of the iterative method (default is `1e-3`).
`ftol`	Positive number indicating the relative tolerance for the change of the log-likelihood value to determine convergence of the iterative method (default is `Inf`, so it is not active). Note that using this argument might have a computational cost as a convergence criterion due to the computation of the log-likelihood (especially when `X` is high-dimensional).

`return_iterates`
: Logical value indicating whether to record the values of the parameters (and possibly the log-likelihood if `ftol < Inf`) at each iteration (default is `FALSE`).

`verbose`
: Logical value indicating whether to allow the function to print messages (default is `FALSE`).

Value

A list containing possibly the following elements:

`mu`
: Mean vector estimate.

`cov`
: Covariance matrix estimate.

`converged`
: Boolean denoting whether the algorithm has converged (`TRUE`) or the maximum number of iterations `max_iter` has reached (`FALSE`).

`num_iterations`
: Number of iterations executed.

`cpu_time`
: Elapsed CPU time.

`log_likelihood`
: Value of log-likelihood after converge of the estimation algorithm (if `ftol < Inf`).

`iterates_record`
: Iterates of the parameters (`mu`, `scatter`, and possibly `log_likelihood` (if `ftol < Inf`)) along the iterations (if `return_iterates = TRUE`).

Author(s)

Daniel P. Palomar

References

Ying Sun, Prabhu Babu, and Daniel P. Palomar, "Regularized Tyler's Scatter Estimator: Existence, Uniqueness, and Algorithms," IEEE Trans. on Signal Processing, vol. 62, no. 19, pp. 5143-5156, Oct. 2014.

See Also

`fit_Cauchy` and `fit_mvt`

Examples

```
library(mvtnorm)       # to generate heavy-tailed data
library(fitHeavyTail)

X <- rmvt(n = 1000, df = 6)  # generate Student's t data
fit_Tyler(X)
```

Index