Package ‘bigtime’

November 9, 2017

Type Package

Title Sparse Estimation of Large Time Series Models

Version 0.1.0

Author Ines Wilms [cre, aut],
David S. Matteson [aut],
Jacob Bien [aut],
Sumanta Basu [aut]

Maintainer Ines Wilms <ines.wilms@kuleuven.be>

Description Estimation of large Vector AutoRegressive (VAR), Vector AutoRegressive with Exogenous Variables X (VARX) and Vector AutoRegressive Moving Average (VARMA) Models with Structured Lasso Penalties, see Nicholson, Bien and Matteson (2017) <arXiv:1412.5250v2> and Wilms, Basu, Bien and Matteson (2017) <arXiv:1707.09208>.

Depends R (>= 3.1.0), methods

License GPL (>= 2)

Encoding UTF-8

LazyData true

RoxygenNote 6.0.1

Imports MASS, zoo, lattice, Rcpp, stats, utils, grDevices, graphics, corrplot

SystemRequirements C++11

LinkingTo Rcpp, RcppArmadillo, RcppEigen

URL http://github.com/ineswilms/bigtime

NeedsCompilation yes

Repository CRAN

Date/Publication 2017-11-09 18:45:41 UTC

R topics documented:

`bigtime` *bigtime: A package for obtaining sparse estimates of large time series models.*

Description

The bigtime package provides sparse estimators for three large time series models: Vector AutoRegressive Models, Vector AutoRegressive Models with Exogenous variables, and Vector AutoRegressive Moving Average Models. The univariate cases are also supported.

Details

To use the facilities of this package, start with a T by k time series matrix Y (for the VAR and VARMA), and an exogenous time series matrix X (for the VARX). Run sparseVAR, sparseVARX or sparseVARMA to get the estimated model. The function lagmatrix returns the lag matrix of estimated coefficients of the estimated model. The function directforecast gives h-step ahead forecasts based on the estimated model.

Author(s)

Ines Wilms <ines.wilms@kuleuven.be>, Jacob Bien, David S. Matteson, Sumanta Basu

References

Nicholson William B., Bien Jacob and Matteson David S. (2017), "High Dimensional Forecasting via Interpretable Vector Autoregression" arXiv preprint <arXiv:1412.5250v2>.

Wilms Ines, Sumanta Basu, Bien Jacob and Matteson David S. (2017), "Sparse Identification and Estimation of High-Dimensional Vector AutoRegressive Moving Averages" arXiv preprint <arXiv:1707.09208>.

Examples

```
# Fit a sparse VAR model
data(Y)
VARfit <- sparseVAR(Y) # sparse VAR
Lhat <- lagmatrix(fit=VARfit, model="VAR") # get estimated lagmatrix
VARforecast <- directforecast(fit=VARfit, model="VAR", h=1) # get one-step ahead forecasts
```

`directforecast` *Function to obtain h-step ahead direct forecast based on estimated VAR, VARX or VARMA model*

Description

Function to obtain h-step ahead direct forecast based on estimated VAR, VARX or VARMA model

Usage

```
directforecast(fit, model, h = 1)
```

Arguments

`fit`	Fitted sparse VAR, VARX or VARMA model.
`model`	Type of model that was estimated: VAR, VARX or VARMA.
`h`	Desired forecast horizon.

Value

Vector of length k containing the h-step ahead forecasts for the k time series.

Examples

```
data(Y)
VARfit <- sparseVAR(Y) # sparse VAR
VARforecast <- directforecast(fit=VARfit, model="VAR", h=1)
```

`lagmatrix` *Creates Lagmatrix of Estimated Coefficients*

Description

Creates Lagmatrix of Estimated Coefficients

Usage

```
lagmatrix(fit, model, returnplot = F)
```

Arguments

`fit`	Fitted VAR, VARX or VARMA model.
`model`	Type of model that was estimated: VAR, VARX or VARMA.
`returnplot`	TRUE or FALSE: return plot of lag matrix or not.

Value

A list with estimated lag matrix of the VAR model, or lag matrices of the VARX or VARMA model. The rows contain the responses, the columns contain the predictors.

Examples

```
data(Y)
data(X)
VARXfit <- sparseVARX(Y=Y, X=X) # sparse VARX
Lhats <- lagmatrix(fit=VARXfit, model="VARX")
```

`sparseVAR` *Sparse Estimation of the Vector AutoRegressive (VAR) Model*

Description

Sparse Estimation of the Vector AutoRegressive (VAR) Model

Usage

```
sparseVAR(Y, p = NULL, VARpen = "HLag", VARlseq = NULL, VARgran = NULL,
  VARalpha = 0, cvcut = 0.9, h = 1, eps = 0.001)
```

Arguments

`Y`	A T by k matrix of time series. If k=1, a univariate autoregressive model is estimated.
`p`	User-specified maximum autoregressive lag order of the VAR. Typical usage is to have the program compute its own maximum lag order based on the time series length.
`VARpen`	"HLag" (hierarchical sparse penalty) or "L1" (standard lasso penalty) penalization.
`VARlseq`	User-specified grid of values for regularization parameter corresponding to sparse penalty. Typical usage is to have the program compute its own grid. Supplying a grid of values overrides this. WARNING: use with care.
`VARgran`	User-specified vector of granularity specifications for the penalty parameter grid: First element specifies how deep the grid should be constructed. Second element specifies how many values the grid should contain.
`VARalpha`	a small positive regularization parameter value corresponding to squared Frobenius penalty. The default is zero.
`cvcut`	Proportion of observations used for model estimation in the time series cross-validation procedure. The remainder is used for forecast evaluation.
`h`	Desired forecast horizon in time-series cross-validation procedure.
`eps`	a small positive numeric value giving the tolerance for convergence in the proximal gradient algorithm.

Value

A list with the following components

Y	T by k matrix of time series.
k	Number of time series.
p	Maximum autoregressive lag order of the VAR.
Phihat	Matrix of estimated autoregressive coefficients of the VAR.
phi0hat	vector of VAR intercepts.

References

Nicholson William B., Bien Jacob and Matteson David S. (2017), "High Dimensional Forecasting via Interpretable Vector Autoregression" arXiv preprint <arXiv:1412.5250v2>.

See Also

lagmatrix and directforecast

Examples

```
data(Y)
VARfit <- sparseVAR(Y) # sparse VAR
y <- matrix(Y[,1], ncol=1)
ARfit <- sparseVAR(y) # sparse AR
```

---

sparseVARMA — *Sparse Estimation of the Vector AutoRegressive Moving Average (VARMA) Model*

Description

Sparse Estimation of the Vector AutoRegressive Moving Average (VARMA) Model

Usage

```
sparseVARMA(Y, U = NULL, VARp = NULL, VARpen = "HLag", VARlseq = NULL,
  VARgran = NULL, VARalpha = 0, VARMAp = NULL, VARMAq = NULL,
  VARMApen = "HLag", VARMAlPhiseq = NULL, VARMAPhigran = NULL,
  VARMAlThetaseq = NULL, VARMAThetagran = NULL, VARMAalpha = 0, h = 1,
  cvcut = 0.9, eps = 10^-3)
```

Arguments

`Y` A T by k matrix of time series. If k=1, a univariate autoregressive moving average model is estimated.

`U` A T by k matrix of (approximated) error terms. Typical usage is to have the program estimate a high-order VAR model (Phase I) to get approximated error terms U.

`VARp` User-specified maximum autoregressive lag order of the PhaseI VAR. Typical usage is to have the program compute its own maximum lag order based on the time series length.

`VARpen` "HLag" (hierarchical sparse penalty) or "L1" (standard lasso penalty) penalization in PhaseI VAR.

`VARlseq` User-specified grid of values for regularization parameter in the PhaseI VAR. Typical usage is to have the program compute its own grid. Supplying a grid of values overrides this. WARNING: use with care.

`VARgran` User-specified vector of granularity specifications for the penalty parameter grid of the PhaseI VAR: First element specifies how deep the grid should be constructed. Second element specifies how many values the grid should contain.

`VARalpha` a small positive regularization parameter value corresponding to squared Frobenius penalty in PhaseI VAR. The default is zero.

`VARMAp` User-specified maximum autoregressive lag order of the VARMA. Typical usage is to have the program compute its own maximum lag order based on the time series length.

`VARMAq` User-specified maximum moving average lag order of the VARMA. Typical usage is to have the program compute its own maximum lag order based on the time series length.

`VARMApen` "HLag" (hierarchical sparse penalty) or "L1" (standard lasso penalty) penalization in the VARMA.

`VARMAlPhiseq` User-specified grid of values for regularization parameter corresponding to the autoregressive coefficients in the VARMA. Typical usage is to have the program compute its own grid. Supplying a grid of values overrides this. WARNING: use with care.

`VARMAPhigran` User-specified vector of granularity specifications for the penalty parameter grid corresponding to the autoregressive coefficients in the VARMA: First element specifies how deep the grid should be constructed. Second element specifies how many values the grid should contain.

`VARMAlThetaseq` User-specified grid of values for regularization parameter corresponding to the moving average coefficients in the VARMA. Typical usage is to have the program compute its own grid. Supplying a grid of values overrides this. WARNING: use with care.

`VARMAThetagran` User-specified vector of granularity specifications for the penalty parameter grid corresponding to the moving average coefficients in the VARMA: First element specifies how deep the grid should be constructed. Second element specifies how many values the grid should contain.

VARMAalpha	a small positive regularization parameter value corresponding to squared Frobenius penalty in VARMA. The default is zero.
h	Desired forecast horizon in time-series cross-validation procedure.
cvcut	Proportion of observations used for model estimation in the time series cross-validation procedure. The remainder is used for forecast evaluation.
eps	a small positive numeric value giving the tolerance for convergence in the proximal gradient algorithms.

Value

A list with the following components

Y	T by k matrix of time series.
U	Matrix of (approximated) error terms.
k	Number of time series.
VARp	Maximum autoregressive lag order of the PhaseI VAR.
VARPhihat	Matrix of estimated autoregressive coefficients of the Phase I VAR.
VARphi0hat	Vector of Phase I VAR intercepts.
VARMAp	Maximum autoregressive lag order of the VARMA.
VARMAq	Maximum moving average lag order of the VARMA.
Phihat	Matrix of estimated autoregressive coefficients of the VARMA.
Thetahat	Matrix of estimated moving average coefficients of the VARMA.
phi0hat	Vector of VARMA intercepts.

References

Wilms Ines, Sumanta Basu, Bien Jacob and Matteson David S. (2017), "Sparse Identification and Estimation of High-Dimensional Vector AutoRegressive Moving Averages" arXiv preprint <arXiv:1707.09208>.

See Also

lagmatrix and directforecast

Examples

```
data(Y)
VARMAfit <- sparseVARMA(Y) # sparse VARMA
y <- matrix(Y[,1], ncol=1)
ARMAfit <- sparseVARMA(y) # sparse ARMA
```

sparseVARX — *Sparse Estimation of the Vector AutoRegressive with Exogenous Variables X (VARX) Model*

Description

Sparse Estimation of the Vector AutoRegressive with Exogenous Variables X (VARX) Model

Usage

```
sparseVARX(Y, X, p = NULL, s = NULL, VARXpen = "HLag",
  VARXlPhiseq = NULL, VARXPhigran = NULL, VARXlBseq = NULL,
  VARXBgran = NULL, VARXalpha = 0, h = 1, cvcut = 0.9, eps = 10^-3)
```

Arguments

Y	A T by k matrix of time series. If k=1, a univariate autoregressive model is estimated.
X	A T by m matrix of time series.
p	User-specified maximum endogenous autoregressive lag order. Typical usage is to have the program compute its own maximum lag order based on the time series length.
s	User-specified maximum exogenous autoregressive lag order. Typical usage is to have the program compute its own maximum lag order based on the time series length.
VARXpen	"HLag" (hierarchical sparse penalty) or "L1" (standard lasso penalty) penalization in VARX.
VARXlPhiseq	User-specified grid of values for regularization parameter corresponding to the endogenous autoregressive coefficients in the VARX. Typical usage is to have the program compute its own grid. Supplying a grid of values overrides this. WARNING: use with care.
VARXPhigran	User-specified vector of granularity specifications for the penalty parameter grid corresponding to the endogenous autoregressive coefficients in the VARX: First element specifies how deep the grid should be constructed. Second element specifies how many values the grid should contain.
VARXlBseq	User-specified grid of values for regularization parameter corresponding to the exogenous autoregressive coefficients in the VARX. Typical usage is to have the program compute its own grid. Supplying a grid of values overrides this. WARNING: use with care.
VARXBgran	User-specified vector of granularity specifications for the penalty parameter grid corresponding to the exogenous autoregressive coefficients in the VARX: First element specifies how deep the grid should be constructed. Second element specifies how many values the grid should contain.
VARXalpha	a small positive regularization parameter value corresponding to squared Frobenius penalty. The default is zero.

`h`	Desired forecast horizon in time-series cross-validation procedure.
`cvcut`	Proportion of observations used for model estimation in the time series cross-validation procedure. The remainder is used for forecast evaluation.
`eps`	a small positive numeric value giving the tolerance for convergence in the proximal gradient algorithm.

Value

A list with the following components

`Y`	T by k matrix of endogenous time series.
`X`	T by m matrix of exogenous time series.
`k`	Number of endogenous time series.
`m`	Number of exogenous time series.
`p`	Maximum endogenous autoregressive lag order of the VARX.
`s`	Maximum exogenouss autoregressive lag order of the VARX.
`Phihat`	Matrix of estimated endogenous autoregressive coefficients.
`Bhat`	Matrix of estimated exogenous autoregressive coefficients.
`phi0hat`	vector of VARX intercepts.

See Also

lagmatrix and directforecast

Examples

```
data(Y)
data(X)
VARXfit <- sparseVARX(Y=Y, X=X) # sparse VARX
y <- matrix(Y[,1], ncol=1)
ARXfit <- sparseVARX(Y=y, X=X) # sparse ARX
```

`X` *Multivariate Time Series Example*

Description

A matrix containing three time series in its columns that are each observed over 50 time points

Usage

```
X
```

Format

A matrix with 50 rows and 3 columns

`Y` *Multivariate Time Series Example*

Description

A matrix containing two time series in its columns that are each observed over 50 time points

Usage

```
Y
```

Format

A matrix with 50 rows and 2 columns

Index