

Package: varlasso (via r-universe)

August 13, 2024

Type Package

Title Vector Autoregressive State Space Models With Shrinkage

Version 0.0.1

Maintainer Eric Ward <eric.ward@noaa.gov>

Description The varlasso package uses Stan (mc-stan.org) to fit VAR state space models with optional shrinkage priors on B matrix elements (autoregression coefficients).

License GPL (>=3)

Depends R (>= 3.4.0)

Imports MASS, methods, Rcpp (>= 0.12.0), RcppParallel (>= 5.0.1), rstan (>= 2.18.1), rstantools (>= 2.1.1), loo (>= 2.0.0), rlang (>= 0.3.1)

LinkingTo BH (>= 1.66.0), Rcpp (>= 0.12.0), RcppEigen (>= 0.3.3.3.0), RcppParallel (>= 5.0.1), rstan (>= 2.18.1), StanHeaders (>= 2.18.0)

Suggests testthat, knitr, rmarkdown

Encoding UTF-8

LazyData true

URL <https://nwfsc-timeseries.github.io/varlasso/>,
<https://github.com/nwfsc-timeseries/varlasso>

BugReports <https://github.com/nwfsc-timeseries/varlasso/issues>

Roxygen list(markdown = TRUE)

RoxygenNote 7.1.2

SystemRequirements GNU make

Biarch true

VignetteBuilder knitr

Repository <https://atsa-es.r-universe.dev>

RemoteUrl <https://github.com/atsa-es/varlasso>

RemoteRef HEAD

RemoteSha f5d2835faa546798cdb3e41833db91b3ddd9480e

Contents

varlasso-package	2
fit	2
Index	6

varlasso-package	<i>The 'varlasso' package.</i>
------------------	--------------------------------

Description

A DESCRIPTION OF THE PACKAGE

References

Stan Development Team (2020). RStan: the R interface to Stan. R package version 2.21.2. <https://mc-stan.org>

fit	<i>Fit a varlasso model to multivariate time series data</i>
-----	--

Description

fit is the primary function for fitting models with varlasso.

Usage

```
fit(
  data,
  shared_q = NULL,
  shared_r = NULL,
  fixed_r = NULL,
  est_trend = FALSE,
  varss = TRUE,
  b_diag = list(mu = 0.7, sd = 1),
  b_diag_min = 0,
  b_offdiag = list(mu = 0, sd = 1),
  off_diag_prior = c("normal", "student-t", "laplace", "hs", "normal_pp"),
  sigma_pro = list(mu = 0, sd = 1),
  sigma_obs = list(mu = 0, sd = 1),
  nu = NULL,
  sigma_scale = list(df = 3, sd = 2),
  hs = list(df = 1, df_global = 1, df_slab = 4, scale_global = 1, scale_slab = 2),
  unique_reg = FALSE,
  pars_to_monitor = NULL,
  iter = 1000,
```

```

warmup = floor(iter/2),
thin = 1,
chains = 3,
save_log_lik = FALSE,
estimation = c("sampling", "optimizing", "vb"),
est_hessian = FALSE,
prior_only = FALSE,
...
)

```

Arguments

data	The data, as a long format dataframe
shared_q	Optional vector with integers indicating which process variance parameters are shared; defaults to constant process variances shared across species.
shared_r	Optional vector with integers indicating which observation variance parameters are shared; defaults to shared observation variances across species.
fixed_r	Optional scalar or vector of values to be passed in for a model where observation error variances are not estimated. Defaults to NULL, and them being estimated. Cannot be set to zero.
est_trend	Whether or not to estimate a species-specific trend, defaults to FALSE
varss	Boolean, whether to fit a VARSS model (defaults to true) or VAR model
b_diag	A 2 element list specifying the mean and sd of a normal prior on the diagonal elements. These elements are mu (default= 0.7) and sd (default = 1)
b_diag_min	The minimum value of the B diagonal, defaults to 0. Setting this lower, e.g. -1, allows for oscillations
b_offdiag	A 2 element list specifying the mean and sd of a normal prior on the off-diagonal elements. These elements are mu (default= 0) and sd (default = 1)
off_diag_prior	Character string denoting which prior to use for off-diagonal elements. Can be "normal" (independent priors estimated for each) parameter, "normal_pp" (normal prior with partial pooling, where model is hierarchical with an estimated standard deviation), "student-t" (hierarchical student-t model used with estimated scale and df), "laplace" (double-exponential model used with estimated scale) or "hs" (regularized horseshoe prior used, following parameterization in rstanarm)
sigma_pro	A 2 element list specifying the mean and sd of a half Student-t prior on the process variance standard deviation. Elements of the list are mu (default= 0) and sd (default = 1), and the prior df is fixed at 3
sigma_obs	A 2 element list specifying the mean and sd of a half Student-t prior on the process variance standard deviation. Elements of the list are mu (default= 0) and sd (default = 1), and the prior df is fixed at 3
nu	Optional known parameter (df) for Student-t priors
sigma_scale	A 2 element list specifying the df and sd parameters for a half Student-t prior on the scale parameter. This is only used for the Laplace and Student-t priors; default values are 3 for df and 2 for sd

hs	A list containing the hyperparameters of the horseshoe prior. Default values are df (df of the Student-t prior for local shrinkage, defaults to 1), df_global (Student-t df of the global shrinkage, defaults to 1), df_slab (df of the Student-t prior on the slab regularization, defaults to 4), scale_global (Scale for Student-t prior on global shrinkage, defaults to 1), scale_slab (Scale of the Student-t prior on shrinkage parameter, defaults to 2).
unique_reg	Boolean, whether to estimate unique regularization parameters for the Student-t or Laplace models (defaults to false)
pars_to_monitor	Optional string vector of which parameters to monitor. By default this is NULL and only most important parameters are monitored (variances, B matrix, etc). This can be the word "all", in which case all parameters are saved (this may be good for loo::loo_moment_match() for example).
iter	Number of MCMC iterations, defaults to 1000
warmup	Warmup / burn in phase, defaults to 500
thin	MCMC thin, defaults to 1
chains	MCMC chains, defaults to 3
save_log_lik	Whether to save log_lik, defaults to FALSE for speed
estimation	Whether to do MCMC sampling ("sampling"; default), maximum posterior estimation ("optimizing") or Stan's variational algrotim ("vb")
est_hessian	Whether to estimate hessian matrix if doing MAP estimation, defaults to FALSE
prior_only	Whether to only generate samples from the prior distribution, defaults to FALSE
...	Extra arguments to pass to sampling

Value

an object of class 'stanfit'

Examples

```

iter <- 50
warmup <- 20
thin <- 1
chains <- 1

n_ts <- 4
n_time <- 70
n_burn <- 50
xx <- matrix(0, n_time, n_ts)
set.seed(123)
xx[1, ] <- rnorm(n_ts)
B <- matrix(rnorm(n = n_ts * n_ts, mean = 0, sd = 0.1), n_ts, n_ts)
diag(B) <- runif(n_ts, min = 0.7, max = 0.9)

for (i in 2:n_time) {
  xx[i, ] <- B %*% xx[i - 1, ] + rnorm(n_ts, mean = 0, sd = 0.04)
}

```

```
}
xx <- xx[-c(1:n_burn), ]

yy <- xx + rnorm(n = nrow(xx) * ncol(xx), mean = 0, sd = 0.02)
df <- data.frame(
  "time" = rep(1:nrow(yy), ncol(yy)),
  "species" = sort(rep(1:ncol(yy), nrow(yy))),
  "y" = c(yy)
)

# fit basic model with normal independent priors
m <- fit(data = df, chains = chains, iter = iter, warmup = warmup, thin = thin)
# fit same model with shrinkage using student-t priors
m_t <- fit(
  data = df, chains = chains, iter = iter, warmup = warmup, thin = thin,
  off_diag_prior = "student-t"
)
# fit same model using student-t priors and known df
m_t_known <- fit(
  data = df, chains = chains, iter = iter, warmup = warmup, thin = thin,
  off_diag_prior = "student-t", nu = 4
)
# fit same model using laplace priors
m_lp <- fit(
  data = df, chains = chains, iter = iter, warmup = warmup, thin = thin,
  off_diag_prior = "laplace"
)
# fit same model using horseshoe priors
m_hs <- fit(
  data = df, chains = chains, iter = iter, warmup = warmup, thin = thin,
  off_diag_prior = "hs"
)

# include example for just sampling from priors
m <- fit(
  data = df, chains = chains, iter = iter,
  warmup = warmup, thin = thin, prior_only = TRUE
)
```

Index

`fit`, [2](#)

`varlasso (varlasso-package)`, [2](#)

`varlasso-package`, [2](#)