

Package ‘bsplinePsd’

February 23, 2018

Type Package

Title Bayesian Nonparametric Spectral Density Estimation Using
B-Spline Priors

Version 0.2.0

Date 2018-02-23

Author Matthew C. Edwards [aut, cre],
Renate Meyer [aut],
Nelson Christensen [aut]

Maintainer Matthew C. Edwards <matt.edwards@auckland.ac.nz>

Description Implementation of a Metropolis-within-Gibbs MCMC algorithm to flexibly estimate the spectral density of a stationary time series. The algorithm updates a nonparametric B-spline prior using the Whittle likelihood to produce pseudo-posterior samples and is based on the work presented in Edwards, M.C., Meyer, R. and Christensen, N., Statistics and Computing (2018). <doi.org/10.1007/s11222-017-9796-9>.

License GPL (>= 3)

Imports Rcpp (>= 0.12.5), splines (>= 3.2.3)

LinkingTo Rcpp

RoxygenNote 6.0.1

NeedsCompilation yes

Repository CRAN

Date/Publication 2018-02-23 12:55:00 UTC

R topics documented:

bsplinePsd-package	2
db spline	2
gibbs_bspline	3

Index	6
--------------	----------

dbsplinePsd-package *Bayesian Nonparametric Spectral Density Estimation Using B-Spline Priors*

Description

Implementation of a Metropolis-within-Gibbs MCMC algorithm to flexibly estimate the spectral density of a stationary time series. The algorithm updates a nonparametric B-spline prior using the Whittle likelihood to produce pseudo-posterior samples.

Details

The function `gibbs_bspline` is an implementation of the (serial version of the) MCMC algorithm presented in Edwards et al. (2017). This algorithm uses a nonparametric B-spline prior to estimate the spectral density of a stationary time series and can be considered a generalisation of the algorithm of Choudhuri et al. (2004), which used the Bernstein polynomial prior. A Dirichlet process prior is used to find the weights for the B-spline densities used in the finite mixture and a separate and independent Dirichlet process prior used to place knots. The algorithm therefore allows for a data-driven choice of the number of knots/mixtures and their locations.

Author(s)

Matthew C. Edwards, Renate Meyer, Nelson Christensen

Maintainer: Matthew C. Edwards <matt.edwards@auckland.ac.nz>

References

Edwards, M. C., Meyer, R., and Christensen, N. (2017), Bayesian nonparametric spectral density estimation using B-spline priors, <arXiv:1707.04878>.

Choudhuri, N., Ghosal, S., and Roy, A. (2004), Bayesian estimation of the spectral density of a time series, *Journal of the American Statistical Association*, 99(468):1050–1059.

dbspline *Generate a cubic B-spline density basis*

Description

This function generates a cubic B-spline density basis.

Usage

```
dbspline(x, knots)
```

Arguments

x numeric vector for which the B-spline densities are to be generated
 knots knots used to generate the cubic B-spline densities

Details

[splineDesign](#) is used to generate a cubic B-spline basis. Each B-spline is then normalised to become a B-spline density using analytical integration. Note that the two end knots are each coincident four times.

Value

matrix of the cubic B-spline density basis

See Also

[splineDesign](#)

Examples

```
## Not run:

# Generate basis functions
x = seq(0, 1, length = 256)
knots = sort(c(0, runif(10), 1))
basis = dbspline(x, knots)

# Plot basis functions
plot(x, basis[1, ], type = "l", ylim = c(min(basis), max(basis)))
for (i in 2:nrow(basis)) lines(x, basis[i, ], col = i)

## End(Not run)
```

gibbs_bspline	<i>Metropolis-within-Gibbs sampler for spectral inference of a stationary time series using a B-spline prior</i>
---------------	--

Description

This function updates the (cubic) B-spline prior using the Whittle likelihood and obtains samples from the pseudo-posterior to infer the spectral density of a stationary time series.

Usage

```
gibbs_bspline(data, Ntotal, burnin, thin = 1, k.theta = 0.01, MG = 1,
  G0.alpha = 1, G0.beta = 1, LG = 20, MH = 1, H0.alpha = 1,
  H0.beta = 1, LH = 20, tau.alpha = 0.001, tau.beta = 0.001,
  kmax = 100, k1 = 20)
```

Arguments

data	numeric vector
Ntotal	total number of iterations to run the Markov chain
burnin	number of initial iterations to be discarded
thin	thinning number (post-processing)
k.theta	prior parameter for number of B-spline densities k (proportional to $\exp(-k.\text{theta}*k^2)$) in mixture
MG	Dirichlet process base measure constant for weights of B-spline densities in mixture (> 0)
G0.alpha, G0.beta	parameters of Beta base measure of Dirichlet process for weights of B-spline densities in mixture (default is Uniform[0, 1])
LG	truncation parameter of Dirichlet process in stick breaking representation for weights of B-spline densities
MH	Dirichlet process base measure constant for knot placements of B-spline densities (> 0)
H0.alpha, H0.beta	parameters of Beta base measure of Dirichlet process for knot placements of B-spline densities (default is Uniform[0, 1])
LH	truncation parameter of Dirichlet process in stick breaking representation for knot placements of B-spline densities
tau.alpha, tau.beta	prior parameters for tau (Inverse Gamma)
kmax	upper bound for number of B-spline densities in mixture
k1	starting value for k. If k1 = NA then a random starting value between 5 and kmax is selected.

Details

The function `gibbs_bspline` is an implementation of the (serial version of the) MCMC algorithm presented in Edwards et al. (2018). This algorithm uses a nonparametric B-spline prior to estimate the spectral density of a stationary time series and can be considered a generalisation of the algorithm of Choudhuri et al. (2004), which used the Bernstein polynomial prior. A Dirichlet process prior is used to find the weights for the B-spline densities used in the finite mixture and a separate and independent Dirichlet process prior used to place knots. The algorithm therefore allows for a data-driven choice of the number of knots/mixtures and their locations.

Value

A list containing the following components:

psd.median, psd.mean	psd estimates: (pointwise) posterior median and mean
psd.p05, psd.p95	pointwise credibility interval

```

psd.u05, psd.u95
                    uniform credibility interval
k, tau, V, Z, U, X  posterior traces of model parameters
knots.trace        trace of knot placements
ll.trace           trace of log likelihood

```

References

Edwards, M. C., Meyer, R., and Christensen, N. (2018), Bayesian nonparametric spectral density estimation using B-spline priors, *Statistics and Computing*, <https://doi.org/10.1007/s11222-017-9796-9>.

Choudhuri, N., Ghosal, S., and Roy, A. (2004), Bayesian estimation of the spectral density of a time series, *Journal of the American Statistical Association*, 99(468):1050–1059.

Examples

```

## Not run:

# Generate AR(1) data with rho = 0.9
n = 128
data = arima.sim(n, model = list(ar = 0.9))
data = data - mean(data)

# Run MCMC (may take some time)
mcmc = gibbs_bspline(data, Ntotal = 4000, burnin = 2000, thin = 1)

# Compare estimate with true PSD
require(beyondWhittle) # For psd_arma() function
freq = 2 * pi / n * (1:(n / 2 + 1) - 1)[-c(1, n / 2 + 1)]
psd.true <- psd_arma(freq, ar = 0.9, ma = numeric(0), sigma2 = 1)
plot(x = freq, y = psd.true, col = 2, type = "l", xlab = "Frequency", ylab = "PSD")
lines(x = freq, y = mcmc$psd.median, type = "l")
lines(x = freq, y = mcmc$psd.p05, type = "l", lty = 2)
lines(x = freq, y = mcmc$psd.p95, type = "l", lty = 2)
legend(x = "topright", legend = c("true psd", "pointwise median", "pointwise CI"),
lty = c(1, 1, 2), col = c(2, 1, 1))

## End(Not run)

```

Index

*Topic **package**

bsplinePsd-package, [2](#)

bsplinePsd (bsplinePsd-package), [2](#)

bsplinePsd-package, [2](#)

db spline, [2](#)

gibbs_bspline, [3](#)

splineDesign, [3](#)