

Package ‘bsts’

May 8, 2026

Version 0.9.11

Date 2025-09-04

Title Bayesian Structural Time Series

Description Time series regression using dynamic linear models fit using MCMC. See Scott and Varian (2014) <[DOI:10.1504/IJMMNO.2014.059942](https://doi.org/10.1504/IJMMNO.2014.059942)>, among many other sources.

Depends BoomSpikeSlab (>= 1.2.7), zoo (>= 1.8), xts, Boom (>= 0.9.16), R(>= 4.5.0)

Suggests testthat

LinkingTo Boom (>= 0.9.16)

License LGPL-2.1 | MIT + file LICENSE

Encoding UTF-8

NeedsCompilation yes

Author Steven L. Scott [aut, cre]

Maintainer Steven L. Scott <steve.the.bayesian@gmail.com>

Repository CRAN

Date/Publication 2025-09-05 16:40:32 UTC

Contents

bsts-package	3
add.ar	4
add.dynamic.regression	5
add.local.level	9
add.local.linear.trend	10
add.monthly.annual.cycle	12
add.random.walk.holiday	14
add.seasonal	16
add.semilocal.linear.trend	18
add.shared.local.level	20
add.static.intercept	22

add.student.local.linear.trend	23
add.trig	25
aggregate.time.series	27
aggregate.weeks.to.months	28
auto.ar	30
bsts	31
bsts.options.Rd	37
compare.bsts.models	38
date.range	40
descriptive-plots	41
diagnostic-plots	42
dirm	43
dirm-model-optoins	46
estimate.time.scale	47
extend.time	48
format.timestamps	49
gdp	50
geometric.sequence	50
get.fraction	51
goog	52
HarveyCumulator	53
holiday	54
iclaims	56
last.day.in.month	57
MATCH.NumericTimestamps	58
match.week.to.month	58
max.window.width	59
mbsts	60
mixed.frequency	65
month.distance	67
named.holidays	68
new.home.sales	69
one.step.prediction.errors	69
plot.bsts	71
plot.bsts.mixed	75
plot.bsts.prediction	77
plot.bsts.predictors	79
plot.holiday	80
plot.mbsts	82
plot.mbsts.prediction	83
predict.bsts	84
predict.mbsts	88
quarter	90
regression.holiday	90
regularize.timestamps	93
residuals.bsts	94
rsxfs	95
shark	96

shorten	96
simulate.fake.mixed.frequency.data	97
spike.slabs.ar.prior	99
state.sizes	100
StateSpecification	101
SuggestBurn	102
summary.bsts	102
to.posixt	104
turkish	104
week.ends	105
weekday.names	106
wide.to.long	106

Index 108

<i>bsts-package</i>	<i>bsts</i>
---------------------	-------------

Description

Time series regression using dynamic linear models fit using MCMC. See Scott and Varian (2014) <DOI:10.1504/IJMMNO.2014.059942>, among many other sources.

Details

Installation note for Linux users: If you are installing *bsts* using `install.packages` on a Linux machine (and thus compiling yourself) you will almost certainly want to set the `Ncpus` argument to a large number. Windows and Mac users can ignore this advice.

Author(s)

Author: Steven L. Scott <steve.the.bayesian@gmail.com> Maintainer: Steven L. Scott <steve.the.bayesian@gmail.com>

References

Please see the references in the help page for the `bsts` function.

See Also

See the examples in the `bsts` function.

add.ar *AR(p) state component*

Description

Add an AR(p) state component to the state specification.

Usage

```
AddAr(state.specification,
      y,
      lags = 1,
      sigma.prior,
      initial.state.prior = NULL,
      sdy)
```

Arguments

<code>state.specification</code>	A list of state components. If omitted, an empty list is assumed.
<code>y</code>	A numeric vector. The time series to be modeled.
<code>lags</code>	The number of lags ("p") in the AR(p) process.
<code>sigma.prior</code>	An object created by SdPrior. The prior for the standard deviation of the process increments.
<code>initial.state.prior</code>	An object of class MvnPrior describing the values of the state at time 0. This argument can be NULL, in which case the stationary distribution of the AR(p) process will be used as the initial state distribution.
<code>sdy</code>	The sample standard deviation of the time series to be modeled. Used to scale the prior distribution. This can be omitted if <code>y</code> is supplied.

Details

The model is

$$\alpha_t = \phi_1 \alpha_{i,t-1} + \dots + \phi_p \alpha_{t-p} + \epsilon_{t-1} \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2)$$

The state consists of the last p lags of α . The state transition matrix has ϕ in its first row, ones along its first subdiagonal, and zeros elsewhere. The state variance matrix has σ^2 in its upper left corner and is zero elsewhere. The observation matrix has 1 in its first element and is zero otherwise.

Value

Returns `state.specification` with an AR(p) state component added to the end.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior](#)

Examples

```
n <- 100
residual.sd <- .001

# Actual values of the AR coefficients
true.phi <- c(-.7, .3, .15)
ar <- arima.sim(model = list(ar = true.phi),
                n = n,
                sd = 3)

## Layer some noise on top of the AR process.
y <- ar + rnorm(n, 0, residual.sd)
ss <- AddAr(list(), lags = 3, sigma.prior = SdPrior(3.0, 1.0))

# Fit the model with knowledge with residual.sd essentially fixed at the
# true value.
model <- bsts(y, state.specification=ss, niter = 500, prior = SdPrior(residual.sd, 100000))

# Now compare the empirical ACF to the true ACF.
acf(y, lag.max = 30)
points(0:30, ARMAacf(ar = true.phi, lag.max = 30), pch = "+")
points(0:30, ARMAacf(ar = colMeans(model$AR3.coefficients), lag.max = 30))
legend("topright", leg = c("empirical", "truth", "MCMC"), pch = c(NA, "+", "o"))
```

add.dynamic.regression

Dynamic Regression State Component

Description

Add a dynamic regression component to the state specification of a bsts model. A dynamic regression is a regression model where the coefficients change over time according to a random walk.

Usage

```

AddDynamicRegression(
  state.specification,
  formula,
  data,
  model.options = NULL,
  sigma.mean.prior.DEPRECATED = NULL,
  shrinkage.parameter.prior.DEPRECATED = GammaPrior(a = 10, b = 1),
  sigma.max.DEPRECATED = NULL,
  contrasts = NULL,
  na.action = na.pass)

DynamicRegressionRandomWalkOptions(
  sigma.prior = NULL,
  sdy = NULL,
  sdx = NULL)

DynamicRegressionHierarchicalRandomWalkOptions(
  sdy = NULL,
  sigma.mean.prior = NULL,
  shrinkage.parameter.prior = GammaPrior(a = 10, b = 1),
  sigma.max = NULL)

DynamicRegressionArOptions(lags = 1, sigma.prior = SdPrior(1, 1))

```

Arguments

<code>state.specification</code>	A list of state components that you wish to add to. If omitted, an empty list will be assumed.
<code>formula</code>	A formula describing the regression portion of the relationship between y and X. If no regressors are desired then the formula can be replaced by a numeric vector giving the time series to be modeled.
<code>data</code>	An optional data frame, list or environment (or object coercible by as.data.frame to a data frame) containing the variables in the model. If not found in data, the variables are taken from <code>'environment(formula)'</code> , typically the environment from which <code>AddDynamicRegression</code> is called.
<code>model.options</code>	An object inheriting from DynamicRegressionOptions giving the specific transition model for the dynamic regression coefficients, and the prior distribution for any hyperparameters associated with the transition model.
<code>sigma.mean.prior</code>	An object created by GammaPrior describing the prior distribution of b/a (see details below).
<code>sigma.mean.prior.DEPRECATED</code>	This option should be set using <code>model.options</code> . It will be removed in a future release.

<code>shrinkage.parameter.prior</code>	An object of class <code>GammaPrior</code> describing the shrinkage parameter, a (see details below).
<code>shrinkage.parameter.prior.DEPRECATED</code>	This option should be set using <code>model.options</code> . It will be removed in a future release.
<code>sigma.max</code>	The largest supported value of each <code>sigma[i]</code> . Truncating the support of <code>sigma</code> can keep ill-conditioned models from crashing. This must be a positive number (Inf is okay), or NULL. A NULL value will set <code>sigma.max = sd(y)</code> , which is a substantially larger value than one would expect, so in well behaved models this constraint will not affect the analysis.
<code>sigma.max.DEPRECATED</code>	This option should be set using <code>model.options</code> . It will be removed in a future release.
<code>contrasts</code>	An optional list. See the <code>contrasts.arg</code> of <code>model.matrix.default</code> . This argument is only used if a model formula is specified. It can usually be ignored even then.
<code>na.action</code>	What to do about missing values. The default is to allow missing responses, but no missing predictors. Set this to <code>na.omit</code> or <code>na.exclude</code> if you want to omit missing responses altogether.
<code>sd</code>	The standard deviation of the response variable. This is used to scale default priors and <code>sigma.max</code> if other arguments are left NULL. If all other arguments are non-NULL then <code>sd</code> is not used.
<code>sdx</code>	The vector of standard deviations of each predictor variable in the dynamic regression. Used only to scale the default prior. This argument is not used if a prior is specified directly.
<code>lags</code>	The number of lags in the autoregressive process for the coefficients.
<code>sigma.prior</code>	Either an object of class <code>SdPrior</code> or a list of such objects. If a single <code>SdPrior</code> is given then it specifies the prior on the innovation variance for all the coefficients. If a list of <code>SdPrior</code> objects is given, then each element gives the prior distribution for the corresponding regression coefficient. The length of such a list must match the number of predictors in the dynamic regression part of the model.

Details

For the standard "random walk" coefficient model, the model is

$$\beta_{i,t+1} = \beta_{i,t} + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma_i^2 / \text{variance}_{xi})$$

$$\frac{1}{\sigma_i^2} \sim \text{Ga}(a, b)$$

$$\sqrt{b/a} \sim \text{sigma.mean.prior}$$

$$a \sim \text{shrinkage.parameter.prior}$$

That is, each coefficient evolves independently, with its own variance term which is scaled by the variance of the i 'th column of X . The parameters of the hyperprior are interpretable as: $\sqrt{b/a}$ typical amount that a coefficient might change in a single time period, and 'a' is the 'sample size' or 'shrinkage parameter' measuring the degree of similarity in $\text{sigma}[i]$ among the arms.

In most cases we hope b/a is small, so that $\text{sigma}[i]$'s will be small and the series will be forecastable. We also hope that 'a' is large because it means that the $\text{sigma}[i]$'s will be similar to one another.

The default prior distribution is a pair of independent Gamma priors for $\sqrt{b/a}$ and a . The mean of $\text{sigma}[i]$ is set to $.01 * \text{sd}(y)$ with shape parameter equal to 1. The mean of the shrinkage parameter is set to 10, but with shape parameter equal to 1.

If the coefficients have AR dynamics, then the model is that each coefficient independently follows an AR(p) process, where p is given by the lags argument. Independent priors are assumed for each coefficient's model, with a uniform prior on AR coefficients (with support restricted to the finite region where the process is stationary), while the `sigma.prior` argument gives the prior for each coefficient's innovation variance.

Value

Returns a list with the elements necessary to specify a dynamic regression model.

Author(s)

Steven L. Scott

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior NormalPrior](#)

Examples

```
## Setting the seed to avoid small sample effects resulting from small
## number of iterations.
set.seed(8675309)
n <- 1000
x <- matrix(rnorm(n))

# beta follows a random walk with sd = .1 starting at -12.
beta <- cumsum(rnorm(n, 0, .1)) - 12

# level is a local level model with sd = 1 starting at 18.
```

```

level <- cumsum(rnorm(n)) + 18

# sigma.obs is .1
error <- rnorm(n, 0, .1)

y <- level + x * beta + error
par(mfrow = c(1, 3))
plot(y, main = "Raw Data")
plot(x, y - level, main = "True Regression Effect")
plot(y - x * beta, main = "Local Level Effect")

ss <- list()
ss <- AddLocalLevel(ss, y)
ss <- AddDynamicRegression(ss, y ~ x)
## In a real application you'd probably want more than 100
## iterations. See comment above about the random seed.
model <- bst(y, state.specification = ss, niter = 100, seed = 8675309)
plot(model, "dynamic", burn = 10)

xx <- rnorm(10)
pred <- predict(model, newdata = xx)
plot(pred)

```

add.local.level	<i>Local level trend state component</i>
-----------------	--

Description

Add a local level model to a state specification. The local level model assumes the trend is a random walk:

$$\alpha_{t+1} = \alpha_t + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma).$$

The prior is on the σ parameter.

Usage

```

AddLocalLevel(
  state.specification,
  y,
  sigma.prior,
  initial.state.prior,
  sdy,
  initial.y)

```

Arguments

state.specification

A list of state components that you wish to add to. If omitted, an empty list will be assumed.

<code>y</code>	The time series to be modeled, as a numeric vector.
<code>sigma.prior</code>	An object created by <code>SdPrior</code> describing the prior distribution for the standard deviation of the random walk increments.
<code>initial.state.prior</code>	An object created using <code>NormalPrior</code> , describing the prior distribution of the initial state vector (at time 1).
<code>sd</code>	The standard deviation of the series to be modeled. This will be ignored if <code>y</code> is provided, or if all the required prior distributions are supplied directly.
<code>initial.y</code>	The initial value of the series being modeled. This will be ignored if <code>y</code> is provided, or if the priors for the initial state are all provided directly.

Value

Returns a list with the elements necessary to specify a local linear trend state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior](#) [NormalPrior](#)

`add.local.linear.trend`

Local linear trend state component

Description

Add a local linear trend model to a state specification. The local linear trend model assumes that both the mean and the slope of the trend follow random walks. The equation for the mean is

$$\mu_{t+1} = \mu_t + \delta_t + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\mu).$$

The equation for the slope is

$$\delta_{t+1} = \delta_t + \eta_t \quad \eta_t \sim \mathcal{N}(0, \sigma_\delta).$$

The prior distribution is on the level standard deviation σ_μ and the slope standard deviation σ_δ .

Usage

```
AddLocalLinearTrend(  
  state.specification = NULL,  
  y,  
  level.sigma.prior = NULL,  
  slope.sigma.prior = NULL,  
  initial.level.prior = NULL,  
  initial.slope.prior = NULL,  
  sdy,  
  initial.y)
```

Arguments

state.specification	A list of state components that you wish to add to. If omitted, an empty list will be assumed.
y	The time series to be modeled, as a numeric vector.
level.sigma.prior	An object created by SdPrior describing the prior distribution for the standard deviation of the level component.
slope.sigma.prior	An object created by SdPrior describing the prior distribution of the standard deviation of the slope component.
initial.level.prior	An object created by NormalPrior describing the initial distribution of the level portion of the initial state vector.
initial.slope.prior	An object created by NormalPrior describing the prior distribution for the slope portion of the initial state vector.
sd	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.
initial.y	The initial value of the series being modeled. This will be ignored if y is provided, or if the priors for the initial state are all provided directly.

Value

Returns a list with the elements necessary to specify a local linear trend state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

- Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.
- Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior NormalPrior](#)

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bsts(y, state.specification = ss, niter = 500)
pred <- predict(model, horizon = 12, burn = 100)
plot(pred)
```

add.monthly.annual.cycle

Monthly Annual Cycle State Component

Description

A seasonal state component for daily data, representing the contribution of each month to the annual seasonal cycle. I.e. this is the "January, February, March, ..." effect, with 12 seasons. There is a step change at the start of each month, and then the contribution of that month is constant over the course of the month.

Note that if you have anything other than daily data, then you're probably looking for [AddSeasonal](#) instead.

The state of this model is an 11-vector γ_t where the first element is the contribution to the mean for the current month, and the remaining elements are the values for the 10 most recent months. When t is the first day in the month then

$$\gamma_{t+1} = - \sum_{i=2}^{11} \gamma_{t,i} + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma)$$

And the remaining elements are γ_t shifted down one. When t is any other day then $\gamma_{t+1} = \gamma_t$.

Usage

```
AddMonthlyAnnualCycle(state.specification,
                        y,
                        date.of.first.observation = NULL,
                        sigma.prior = NULL,
                        initial.state.prior = NULL,
                        sdy)
```

Arguments

state.specification	A list of state components, to which the monthly annual cycle will be added. If omitted, an empty list will be assumed.
y	The time series to be modeled, as a numeric vector.
date.of.first.observation	The time stamp of the first observation in y, as a Date or POSIXt object. If y is a zoo object with appropriate time stamps then this argument will be deduced.
sigma.prior	An object created by SdPrior describing the prior distribution for the standard deviation of the random walk increments.
initial.state.prior	An object created using NormalPrior , describing the prior distribution of the the initial state vector (at time 1).
sd	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.

Examples

```
## Let's simulate some fake daily data with a monthly cycle.
## Not run:
residuals <- rnorm(365 * 5)

## End(Not run)

n <- length(residuals)
dates <- seq.Date(from = as.Date("2014-01-01"),
                 len = n,
                 by = 1)
monthly.cycle <- rnorm(12)
monthly.cycle <- monthly.cycle - mean(monthly.cycle)
timestamps <- as.POSIXlt(dates)
month <- timestamps$mon + 1

new.month <- c(TRUE, diff(timestamps$mon) != 0)
month.effect <- cumsum(new.month)
month.effect[month.effect == 0] <- 12

response <- monthly.cycle[month] + residuals
response <- zoo(response, timestamps)

## Now let's fit a bsts model to the daily data with a monthly annual
## cycle.
ss <- AddLocalLevel(list(), response)
ss <- AddMonthlyAnnualCycle(ss, response)

## In real life you'll probably want more iterations.
model <- bsts(response, state.specification = ss, niter = 200)
plot(model)
plot(model, "monthly")
```

 add.random.walk.holiday

Random Walk Holiday State Model

Description

Adds a random walk holiday state model to the state specification. This model says

$$y_t = \alpha_{d(t),t} + \epsilon_t$$

where there is one element in α_t for each day in the holiday influence window. The transition equation is

$$\alpha_{d(t+1),t+1} = \alpha_{d(t+1),t} + \epsilon_{t+1}$$

if $t+1$ occurs on day $d(t+1)$ of the influence window, and

$$\alpha_{d(t+1),t+1} = \alpha_{d(t+1),t}$$

otherwise.

Usage

```
AddRandomWalkHoliday(state.specification = NULL,
                      y,
                      holiday,
                      time0 = NULL,
                      sigma.prior = NULL,
                      initial.state.prior = NULL,
                      sdy = sd(as.numeric(y), na.rm = TRUE))
```

Arguments

state.specification	A list of state components that you wish augment. If omitted, an empty list will be assumed.
y	The time series to be modeled, as a numeric vector convertible to xts . This state model assumes y contains daily data.
holiday	An object of class Holiday describing the influence window of the holiday being modeled.
time0	An object convertible to Date containing the date of the initial observation in the training data. If omitted and y is a zoo or xts object, then time0 will be obtained from the index of <code>y[1]</code> .
sigma.prior	An object created by SdPrior describing the prior distribution for the standard deviation of the random walk increments.

initial.state.prior	An object created using NormalPrior , describing the prior distribution of the the initial state vector (at time 1).
sd	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.

Value

A list describing the specification of the random walk holiday state model, formatted as expected by the underlying C++ code.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.RegressionHolidayStateModel](#) [HierarchicalRegressionHolidayStateModel](#)

Examples

```
trend <- cumsum(rnorm(730, 0, .1))
dates <- seq.Date(from = as.Date("2014-01-01"), length = length(trend),
  by = "day")
y <- zoo(trend + rnorm(length(trend), 0, .2), dates)

AddHolidayEffect <- function(y, dates, effect) {
  ## Adds a holiday effect to simulated data.
  ## Args:
  ##   y: A zoo time series, with Dates for indices.
  ##   dates: The dates of the holidays.
  ##   effect: A vector of holiday effects of odd length. The central effect is
  ##     the main holiday, with a symmetric influence window on either side.
  ## Returns:
  ##   y, with the holiday effects added.
  time <- dates - (length(effect) - 1) / 2
  for (i in 1:length(effect)) {
    y[time] <- y[time] + effect[i]
    time <- time + 1
  }
  return(y)
}

## Define some holidays.
```

```

memorial.day <- NamedHoliday("MemorialDay")
memorial.day.effect <- c(.3, 3, .5)
memorial.day.dates <- as.Date(c("2014-05-26", "2015-05-25"))
y <- AddHolidayEffect(y, memorial.day.dates, memorial.day.effect)

presidents.day <- NamedHoliday("PresidentsDay")
presidents.day.effect <- c(.5, 2, .25)
presidents.day.dates <- as.Date(c("2014-02-17", "2015-02-16"))
y <- AddHolidayEffect(y, presidents.day.dates, presidents.day.effect)

labor.day <- NamedHoliday("LaborDay")
labor.day.effect <- c(1, 2, 1)
labor.day.dates <- as.Date(c("2014-09-01", "2015-09-07"))
y <- AddHolidayEffect(y, labor.day.dates, labor.day.effect)

## The holidays can be in any order.
holiday.list <- list(memorial.day, labor.day, presidents.day)
number.of.holidays <- length(holiday.list)

## In a real example you'd want more than 100 MCMC iterations.
niter <- 100
ss <- AddLocalLevel(list(), y)
ss <- AddRandomWalkHoliday(ss, y, memorial.day)
ss <- AddRandomWalkHoliday(ss, y, labor.day)
ss <- AddRandomWalkHoliday(ss, y, presidents.day)
model <- bst(y, state.specification = ss, niter = niter, seed = 8675309)

## Plot model components.
plot(model, "comp")

## Plot the effect of the specific state component.
plot(ss[[2]], model)

```

add.seasonal

Seasonal State Component

Description

Add a seasonal model to a state specification.

The seasonal model can be thought of as a regression on `nseasons` dummy variables with coefficients constrained to sum to 1 (in expectation). If there are S seasons then the state vector γ is of dimension $S-1$. The first element of the state vector obeys

$$\gamma_{t+1,1} = - \sum_{i=2}^S \gamma_{t,i} + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(0, \sigma)$$

Usage

```
AddSeasonal(
```

```
state.specification,  
y,  
nseasons,  
season.duration = 1,  
sigma.prior,  
initial.state.prior,  
sdy)
```

Arguments

state.specification	A list of state components that you wish to add to. If omitted, an empty list will be assumed.
y	The time series to be modeled, as a numeric vector.
nseasons	The number of seasons to be modeled.
season.duration	The number of time periods in each season.
sigma.prior	An object created by SdPrior describing the prior distribution for the standard deviation of the random walk increments.
initial.state.prior	An object created using NormalPrior , describing the prior distribution of the the initial state vector (at time 1).
sdy	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.

Value

Returns a list with the elements necessary to specify a seasonal state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts](#). [SdPrior](#) [NormalPrior](#)

Examples

```

data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 500)
pred <- predict(model, horizon = 12, burn = 100)
plot(pred)

```

```
add.semilocal.linear.trend
```

Semilocal Linear Trend

Description

The semi-local linear trend model is similar to the local linear trend, but more useful for long-term forecasting. It assumes the level component moves according to a random walk, but the slope component moves according to an AR1 process centered on a potentially nonzero value D . The equation for the level is

$$\mu_{t+1} = \mu_t + \delta_t + \epsilon_t \quad \epsilon_t \sim \mathcal{N}(t, \sigma_\mu).$$

The equation for the slope is

$$\delta_{t+1} = D + \phi(\delta_t - D) + \eta_t \quad \eta_t \sim \mathcal{N}(t, \sigma_\delta).$$

This model differs from the local linear trend model in that the latter assumes the slope δ_t follows a random walk. A stationary AR(1) process is less variable than a random walk when making projections far into the future, so this model often gives more reasonable uncertainty estimates when making long term forecasts.

The prior distribution for the semi-local linear trend has four independent components. These are:

- an inverse gamma prior on the level standard deviation σ_μ ,
- an inverse gamma prior on the slope standard deviation σ_δ ,
- a Gaussian prior on the long run slope parameter D ,
- and a potentially truncated Gaussian prior on the AR1 coefficient ϕ . If the prior on ϕ is truncated to $(-1, 1)$, then the slope will exhibit short term stationary variation around the long run slope D .

Usage

```

AddSemilocalLinearTrend(
  state.specification = list(),
  y = NULL,
  level.sigma.prior = NULL,
  slope.mean.prior = NULL,

```

```

slope.ar1.prior = NULL,
slope.sigma.prior = NULL,
initial.level.prior = NULL,
initial.slope.prior = NULL,
sdy = NULL,
initial.y = NULL)

```

Arguments

state.specification	A list of state components that you wish to add to. If omitted, an empty list will be assumed.
y	The time series to be modeled, as a numeric vector. This can be omitted if sdy and initial.y are supplied, or if all prior distributions are supplied directly.
level.sigma.prior	An object created by SdPrior describing the prior distribution for the standard deviation of the level component.
slope.mean.prior	An object created by NormalPrior giving the prior distribution for the mean parameter in the generalized local linear trend model (see below).
slope.ar1.prior	An object created by Ar1CoefficientPrior giving the prior distribution for the ar1 coefficient parameter in the generalized local linear trend model (see below).
slope.sigma.prior	An object created by SdPrior describing the prior distribution of the standard deviation of the slope component.
initial.level.prior	An object created by NormalPrior describing the initial distribution of the level portion of the initial state vector.
initial.slope.prior	An object created by NormalPrior describing the prior distribution for the slope portion of the initial state vector.
sdy	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.
initial.y	The initial value of the series being modeled. This will be ignored if y is provided, or if the priors for the initial state are all provided directly.

Value

Returns a list with the elements necessary to specify a generalized local linear trend state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior NormalPrior](#)

add.shared.local.level

Local level trend state component

Description

Add a shared local level model to a state specification. The shared local level model assumes the trend is a multivariate random walk:

$$\alpha_{t+1} = \alpha_t + \eta_t \quad \eta_{tj} \sim \mathcal{N}(0, \sigma_j).$$

The contribution to the mean of the observed series obeys

$$y_t = B\alpha_t + \epsilon_t.$$

plus observation error. Identifiability constraints imply that the observation coefficients B form a rectangular lower triangular matrix with diagonal 1.0.

If there are m time series and p factors, then B has m rows and p columns. Having B be lower triangular means that the first factor affects all series. The second affects all but the first, the third excludes the first two, etc.

Usage

```
AddSharedLocalLevel(
  state.specification,
  response,
  nfactors,
  coefficient.prior = NULL,
  initial.state.prior = NULL,
  timestamps = NULL,
  series.id = NULL,
  sdy,
  ...)
```

Arguments

state.specification	A pre-existing list of state components that you wish to add to. If omitted, an empty list will be assumed.
response	The time series to be modeled. This can either be a matrix with rows as time and columns as series, or it can be a numeric vector. If a vector is passed then timestamps and series.id are required. Otherwise they are unused.
nfactors	The number of latent factors to include in the model. This is the dimension of the state for this model component.
coefficient.prior	Prior distribution on the observation coefficients.
initial.state.prior	An object of class MvnPrior , describing the prior distribution of the initial state vector (at time 1).
timestamps	If response is in long format (i.e. a vector instead of a matrix) this argument is a vector of the same length indicating the time index to which each element of response belongs.
series.id	If response is in long format (i.e. a vector instead of a matrix) this argument is a vector of the same length indicating the time series to which each element of response belongs.
sdv	A vector giving the standard deviation of each series to be modeled. This argument is only necessary if response cannot be supplied directly.
...	Extra arguments passed to ConditionalZellnerPrior , used to create a default prior for the observation coefficients when coefficient.prior is left as NULL.

Value

Returns a list with the elements necessary to specify a local linear trend state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

- Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.
- Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior](#) [NormalPrior](#)

add.static.intercept *Static Intercept State Component*

Description

Adds a static intercept term to a state space model. If the model includes a traditional trend component (e.g. local level, local linear trend, etc) then a separate intercept is not needed (and will probably cause trouble, as it will be confounded with the initial state of the trend model). However, if there is no trend, or the trend is an AR process centered around zero, then adding a static intercept will shift the center to a data-determined value.

Usage

```
AddStaticIntercept(  
  state.specification,  
  y,  
  initial.state.prior = NormalPrior(y[1], sd(y, na.rm = TRUE)))
```

Arguments

`state.specification`
A list of state components that you wish to add to. If omitted, an empty list will be assumed.

`y`
The time series to be modeled, as a numeric vector.

`initial.state.prior`
An object created using [NormalPrior](#), describing the prior distribution of the intercept term.

Value

Returns a list with the information required to specify the state component. If `initial.state.prior` is specified then `y` is unused.

Author(s)

Steven L. Scott

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior](#) [NormalPrior](#)

```
add.student.local.linear.trend
      Robust local linear trend
```

Description

Add a local level model to a state specification. The local linear trend model assumes that both the mean and the slope of the trend follow random walks. The equation for the mean is

$$\mu_{t+1} = \mu_t + \delta_t + \epsilon_t \quad \epsilon_t \sim \mathcal{T}_{\nu_\mu}(0, \sigma_\mu).$$

The equation for the slope is

$$\delta_{t+1} = \delta_t + \eta_t \quad \eta_t \sim \mathcal{T}_{\nu_\delta}(0, \sigma_\delta).$$

Independent prior distributions are assumed on the level standard deviation, σ_μ the slope standard deviation σ_δ , the level tail thickness ν_μ , and the slope tail thickness ν_δ .

Usage

```
AddStudentLocalLinearTrend(
  state.specification = NULL,
  y,
  save.weights = FALSE,
  level.sigma.prior = NULL,
  level.nu.prior = NULL,
  slope.sigma.prior = NULL,
  slope.nu.prior = NULL,
  initial.level.prior = NULL,
  initial.slope.prior = NULL,
  sdy,
  initial.y)
```

Arguments

- | | |
|----------------------------------|--|
| <code>state.specification</code> | A list of state components that you wish to add to. If omitted, an empty list will be assumed. |
| <code>y</code> | The time series to be modeled, as a numeric vector. |
| <code>save.weights</code> | A logical value indicating whether to save the draws of the weights from the normal mixture representation. |
| <code>level.sigma.prior</code> | An object created by SdPrior describing the prior distribution for the standard deviation of the level component. |
| <code>level.nu.prior</code> | An object inheriting from the class DoubleModel , representing the prior distribution on the nu tail thickness parameter of the T distribution for errors in the evolution equation for the level component. |

<code>slope.sigma.prior</code>	An object created by SdPrior describing the prior distribution of the standard deviation of the slope component.
<code>slope.nu.prior</code>	An object inheriting from the class DoubleModel , representing the prior distribution on the nu tail thickness parameter of the T distribution for errors in the evolution equation for the slope component.
<code>initial.level.prior</code>	An object created by NormalPrior describing the initial distribution of the level portion of the initial state vector.
<code>initial.slope.prior</code>	An object created by NormalPrior describing the prior distribution for the slope portion of the initial state vector.
<code>sd</code>	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.
<code>initial.y</code>	The initial value of the series being modeled. This will be ignored if y is provided, or if the priors for the initial state are all provided directly.

Value

Returns a list with the elements necessary to specify a local linear trend state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior](#) [NormalPrior](#)

Examples

```
data(rsxfs)
ss <- AddStudentLocalLinearTrend(list(), rsxfs)
model <- bsts(rsxfs, state.specification = ss, niter = 500)
pred <- predict(model, horizon = 12, burn = 100)
plot(pred)
```

add.trig *Trigonometric Seasonal State Component*

Description

Add a trigonometric seasonal model to a state specification.

Usage

```
AddTrig(
  state.specification = NULL,
  y,
  period,
  frequencies,
  sigma.prior = NULL,
  initial.state.prior = NULL,
  sdy = sd(y, na.rm = TRUE),
  method = c("harmonic", "direct"))
```

Arguments

state.specification	A list of state components that you wish to add to. If omitted, an empty list will be assumed.
y	The time series to be modeled, as a numeric vector.
period	A positive scalar giving the number of time steps required for the longest cycle to repeat.
frequencies	A vector of positive real numbers giving the number of times each cyclic component repeats in a period. One sine and one cosine term will be added for each frequency.
sigma.prior	An object created by SdPrior describing the prior distribution for the standard deviation of the increments for the harmonic coefficients.
initial.state.prior	An object created using NormalPrior , describing the prior distribution of the the initial state vector (at time 1).
sdy	The standard deviation of the series to be modeled. This will be ignored if y is provided, or if all the required prior distributions are supplied directly.
method	The method of including the sinusoids. The "harmonic" method is strongly preferred, with "direct" offered mainly for teaching purposes.

Details

Harmonic Method:

Each frequency $\lambda_j = 2\pi j/S$ where S is the period (number of time points in a full cycle) is associated with two time-varying random components: γ_{jt} , and γ_{jt}^* . They evolve through time as

$$\begin{aligned}\gamma_{j,t+1} &= \gamma_{jt} \cos(\lambda_j) + \gamma_{j,t}^* \sin(\lambda_j) + \epsilon_{0t} \\ \gamma_{j,t+1}^* &= \gamma_{jt}^* \cos(\lambda_j) - \gamma_{jt} \sin(\lambda_j) + \epsilon_{1t}\end{aligned}$$

where ϵ_0 and ϵ_1 are independent with the same variance. This is the real-valued version of a harmonic function: $\gamma \exp(i\theta)$.

The transition matrix multiplies the function by $\exp(i\lambda_j)$, so that after 't' steps the harmonic's value is $\gamma \exp(i\lambda_j t)$.

The model dynamics allows gamma to drift over time in a random walk.

The state of the model is $(\gamma_{jt}, \gamma_{jt}^*)$, for $j = 1, \dots$ number of frequencies.

The state transition matrix is a block diagonal matrix, where block 'j' is

$$\begin{pmatrix} \cos(\lambda_j) & \sin(\lambda_j) \\ -\sin(\lambda_j) & \cos(\lambda_j) \end{pmatrix}$$

The error variance matrix is $\sigma^2 * I$. There is a common σ^2 parameter shared by all frequencies.

The model is full rank, so the state error expander matrix R_t is the identity.

The observation_matrix is $(1, 0, 1, 0, \dots)$, where the 1's pick out the 'real' part of the state contributions.

Direct Method: Under the 'direct' method the trig component adds a collection of sine and cosine terms with randomly varying coefficients to the state model. The coefficients are the states, while the sine and cosine values are part of the "observation matrix".

This state component adds the sum of its terms to the observation equation.

$$y_t = \sum_j \beta_{jt} \sin(f_j t) + \gamma_{jt} \cos(f_j t)$$

The evolution equation is that each of the sinusoid coefficients follows a random walk with standard deviation $\sigma[j]$.

$$\beta_{jt} = \beta_{j,t-1} + N(0, \sigma_{s_j}^2) \gamma_{jt} = \gamma_{j-1} + N(0, \sigma_{c_j}^2)$$

The direct method is generally inferior to the harmonic method. It may be removed in the future.

Value

Returns a list with the elements necessary to specify a seasonal state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

- Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.
- Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also[bsts.SdPrior MvnPrior](#)**Examples**

```

data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddTrig(ss, y, period = 12, frequencies = 1:3)
model <- bsts(y, state.specification = ss, niter = 200)
plot(model)

## The "harmonic" method is much more stable than the "direct" method.
ss <- AddLocalLinearTrend(list(), y)
ss <- AddTrig(ss, y, period = 12, frequencies = 1:3, method = "direct")
model2 <- bsts(y, state.specification = ss, niter = 200)
plot(model2)

```

aggregate.time.series *Aggregate a fine time series to a coarse summary*

Description

Aggregate measurements from a fine scaled time series into a coarse time series. This is similar to functions from the `xts` package, but it can handle aggregation from weeks to months.

Usage

```

AggregateTimeSeries(fine.series,
                    contains.end,
                    membership.fraction,
                    trim.left = any(membership.fraction < 1),
                    trim.right = NULL,
                    byrow = TRUE)

```

Arguments

<code>fine.series</code>	A numeric vector or matrix giving the fine scale time series to be aggregated.
<code>contains.end</code>	A logical vector corresponding to <code>fine.series</code> indicating whether each fine time interval contains the end of a coarse time interval.
<code>membership.fraction</code>	A numeric vector corresponding to <code>fine.series</code> , giving the fraction of each time interval's observation attributable to the coarse interval containing the fine interval's first day. This will usually be a vector of 1's, unless <code>fine.series</code> is weekly.
<code>trim.left</code>	Logical indicating whether the first observation in the coarse aggregate should be removed.

<code>trim.right</code>	Logical indicating whether the final observation in the coarse aggregate should be removed.
<code>byrow</code>	Logical. If <code>fine.series</code> is a matrix, this argument indicates whether rows (TRUE) or columns (FALSE) correspond to time points.

Value

A matrix (if `fine.series` is a matrix) or vector (otherwise) containing the aggregated values of `fine.series`.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
week.ending <- as.Date(c("2011-11-05",
                        "2011-11-12",
                        "2011-11-19",
                        "2011-11-26",
                        "2011-12-03",
                        "2011-12-10",
                        "2011-12-17",
                        "2011-12-24",
                        "2011-12-31"))
membership.fraction <- GetFractionOfDaysInInitialMonth(week.ending)
which.month <- MatchWeekToMonth(week.ending, as.Date("2011-11-01"))
contains.end <- WeekEndsMonth(week.ending)

weekly.values <- rnorm(length(week.ending))
monthly.values <- AggregateTimeSeries(weekly.values, contains.end, membership.fraction)
```

`aggregate.weeks.to.months`

Aggregate a weekly time series to monthly

Description

Aggregate measurements from a weekly time series into a monthly time series.

Usage

```
AggregateWeeksToMonths(weekly.series,
                       membership.fraction = NULL,
                       trim.left = TRUE,
                       trim.right = NULL)
```

Arguments

- `weekly.series` A numeric vector or matrix of class `zoo` giving the weekly time series to be aggregated. The index must be convertible to class `Date`.
- `membership.fraction` An optional numeric vector corresponding to `weekly.series`, giving the fraction of each week's observation attributable to the month containing the week's first day. If missing, then weeks will be split across months in proportion to the number of days in each month.
- `trim.left` Logical indicating whether the first observation in the monthly aggregate should be removed.
- `trim.right` Logical indicating whether the final observation in the monthly aggregate should be removed.

Value

A zoo-matrix (if `weekly.series` is a matrix) or vector (otherwise) containing the aggregated values of `weekly.series`.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[AggregateTimeSeries](#)

Examples

```
week.ending <- as.Date(c("2011-11-05",
                        "2011-11-12",
                        "2011-11-19",
                        "2011-11-26",
                        "2011-12-03",
                        "2011-12-10",
                        "2011-12-17",
                        "2011-12-24",
                        "2011-12-31"))

weekly.values <- zoo(rnorm(length(week.ending)), week.ending)
monthly.values <- AggregateWeeksToMonths(weekly.values)
```

auto.ar

*Sparse AR(p)***Description**

Add a sparse AR(p) process to the state distribution. A sparse AR(p) is an AR(p) process with a spike and slab prior on the autoregression coefficients.

Usage

```
AddAutoAr(state.specification,
            y,
            lags = 1,
            prior = NULL,
            sdy = NULL,
            ...)
```

Arguments

state.specification	A list of state components. If omitted, an empty list is assumed.
y	A numeric vector. The time series to be modeled. This can be omitted if sdy is supplied.
lags	The maximum number of lags ("p") to be considered in the AR(p) process.
prior	An object inheriting from SpikeSlabArPrior , or NULL. If the latter, then a default SpikeSlabArPrior will be created.
sdy	The sample standard deviation of the time series to be modeled. Used to scale the prior distribution. This can be omitted if y is supplied.
...	Extra arguments passed to SpikeSlabArPrior .

Details

The model contributes α_t to the expected value of y_t , where the transition equation is

$$\alpha_t = \phi_1 \alpha_{t-1} + \dots + \phi_p \alpha_{t-p} + \epsilon_{t-1} \quad \epsilon_t \sim \mathcal{N}(0, \sigma^2)$$

The state consists of the last p lags of α . The state transition matrix has ϕ_1 in its first row, ones along its first subdiagonal, and zeros elsewhere. The state variance matrix has σ^2 in its upper left corner and is zero elsewhere. The observation matrix has 1 in its first element and is zero otherwise.

This model differs from the one in [AddAr](#) only in that some of its coefficients may be set to zero.

Value

Returns `state.specification` with an AR(p) state component added to the end.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.SdPrior](#)

Examples

```
n <- 100
residual.sd <- .001

# Actual values of the AR coefficients
true.phi <- c(-.7, .3, .15)
ar <- arima.sim(model = list(ar = true.phi),
                n = n,
                sd = 3)

## Layer some noise on top of the AR process.
y <- ar + rnorm(n, 0, residual.sd)
ss <- AddAutoAr(list(), y, lags = 6)

# Fit the model with knowledge with residual.sd essentially fixed at the
# true value.
model <- bsts(y, state.specification=ss, niter = 500, prior = SdPrior(residual.sd, 100000))

# Now compare the empirical ACF to the true ACF.
acf(y, lag.max = 30)
points(0:30, ARMAacf(ar = true.phi, lag.max = 30), pch = "+")
points(0:30, ARMAacf(ar = colMeans(model$AR6.coefficients), lag.max = 30))
legend("topright", leg = c("empirical", "truth", "MCMC"), pch = c(NA, "+", "o"))
```

Description

Uses MCMC to sample from the posterior distribution of a Bayesian structural time series model. This function can be used either with or without contemporaneous predictor variables (in a time series regression).

If predictor variables are present, the regression coefficients are fixed (as opposed to time varying, though time varying coefficients might be added as state component). The predictors and response in the formula are contemporaneous, so if you want lags and differences you need to put them in the predictor matrix yourself.

If no predictor variables are used, then the model is an ordinary state space time series model.

The model allows for several useful extensions beyond standard Bayesian dynamic linear models.

- A spike-and-slab prior is used for the (static) regression component of models that include predictor variables. This is especially useful with large numbers of regressor series.
- Both the spike-and-slab component (for static regressors) and the Kalman filter (for components of time series state) require observations and state variables to be Gaussian. The `bsts` package allows for non-Gaussian error families in the observation equation (as well as some state components) by using data augmentation to express these families as conditionally Gaussian.
- As of version 0.7.0, `bsts` supports having multiple observations at the same time point. In this case the basic model is taken to be

$$y_{t,j} = Z_t^T \alpha_t + \beta^T x_{t,j} + \epsilon_{t,j}.$$

This is a regression model where all observations with the same time point share a common time series effect.

Usage

```
bsts(formula,
     state.specification,
     family = c("gaussian", "logit", "poisson", "student"),
     data,
     prior,
     contrasts = NULL,
     na.action = na.pass,
     niter,
     ping = niter / 10,
     model.options = BstsOptions(),
     timestamps = NULL,
     seed = NULL,
     ...)
```

Arguments

formula	A formula describing the regression portion of the relationship between y and X . If no regressors are desired then the formula can be replaced by a numeric vector giving the time series to be modeled. Missing values are not allowed in predictors, but they are allowed in the response variable. If the response variable is of class <code>zoo</code> , <code>xts</code> , or <code>ts</code> , then the time series information it contains will be used in many of the plotting methods called from <code>plot.bsts</code> .
---------	---

<code>state.specification</code>	A list with elements created by AddLocalLinearTrend , AddSeasonal , and similar functions for adding components of state. See the help page for state.specification .
<code>family</code>	The model family for the observation equation. Non-Gaussian model families use data augmentation to recover a conditionally Gaussian model.
<code>data</code>	An optional data frame, list or environment (or object coercible by as.data.frame to a data frame) containing the variables in the model. If not found in <code>data</code> , the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>bsts</code> is called.
<code>prior</code>	If regressors are supplied in the model formula, then this is a prior distribution for the regression component of the model, as created by SpikeSlabPrior . The prior for the time series component of the model will be specified during the creation of <code>state.specification</code> . This argument is only used if a formula is specified. If the model contains no regressors, then this is simply the prior on the residual standard deviation, expressed as an object created by SdPrior .
<code>contrasts</code>	An optional list containing the names of contrast functions to use when converting factors numeric variables in a regression formula. This argument works exactly as it does in lm . The names of the list elements correspond to factor variables in your model formula. The list elements themselves are the names of contrast functions (see <code>help(contr.treatment)</code> and the <code>contrasts.arg</code> argument to model.matrix.default). This argument is only used if a model formula is specified, and even then the default is probably what you want.
<code>na.action</code>	What to do about missing values. The default is to allow missing responses, but no missing predictors. Set this to <code>na.omit</code> or <code>na.exclude</code> if you want to omit missing responses altogether.
<code>niter</code>	A positive integer giving the desired number of MCMC draws.
<code>ping</code>	A scalar giving the desired frequency of status messages. If <code>ping > 0</code> then the program will print a status message to the screen every <code>ping</code> MCMC iterations.
<code>model.options</code>	An object (list) returned by BstsOptions . See that function for details.
<code>timestamps</code>	The timestamp associated with each value of the response. This argument is primarily useful in cases where the response has missing gaps, or where there are multiple observations per time point. If the response is a "regular" time series with a single observation per time point then you can leave this argument as <code>NULL</code> . In that case, if either the response or the <code>data</code> argument is a type convertible to zoo then timestamps will be inferred.
<code>seed</code>	An integer to use as the random seed for the underlying C++ code. If <code>NULL</code> then the seed will be set using the clock.
<code>...</code>	Extra arguments to be passed to SpikeSlabPrior (see the entry for the <code>prior</code> argument, above).

Details

If the model family is logit, then there are two ways one can format the response variable. If the response is 0/1, TRUE/FALSE, or 1/-1, then the response variable can be passed as with any other model family. If the response is a set of counts out of a specified number of trials then it can be

passed as a two-column matrix, where the first column contains the counts of successes and the second contains the count of failures.

Likewise, if the model family is Poisson, the response can be passed as a single vector of counts, under the assumption that each observation has unit exposure. If the exposures differ across observations, then the response can be a two column matrix, with the first column containing the event counts and the second containing exposure times.

Value

An object of class `bsts` which is a list with the following components

<code>coefficients</code>	A <code>niter</code> by <code>ncol(X)</code> matrix of MCMC draws of the regression coefficients, where <code>X</code> is the design matrix implied by formula. This is only present if a model formula was supplied.
<code>sigma.obs</code>	A vector of length <code>niter</code> containing MCMC draws of the residual standard deviation.

The returned object will also contain named elements holding the MCMC draws of model parameters belonging to the state models. The names of each component are supplied by the entries in `state.specification`. If a model parameter is a scalar, then the list element is a vector with `niter` elements. If the parameter is a vector then the list element is a matrix with `niter` rows. If the parameter is a matrix then the list element is a 3-way array with first dimension `niter`.

Finally, if a model formula was supplied, then the returned object will contain the information necessary for the `predict` method to build the design matrix when a new prediction is made.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Scott and Varian (2014) "Predicting the Present with Bayesian Structural Time Series", *International Journal of Mathematical Modelling and Numerical Optimization*. 4–23.

Scott and Varian (2015) "Bayesian Variable Selection for Nowcasting Economic Time Series", *Economic Analysis of the Digital Economy*, pp 119-135.

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

George and McCulloch (1997) "Approaches for Bayesian variable selection", *Statistica Sinica* pp 339–374.

See Also

`bsts`, `AddLocalLevel`, `AddLocalLinearTrend`, `AddSemilocalLinearTrend`, `AddSeasonal`, `AddDynamicRegression`, `SpikeSlabPrior`, `SdPrior`.

Examples

```

## Example 1: Time series (ts) data
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 500)
pred <- predict(model, horizon = 12, burn = 100)
par(mfrow = c(1,2))
plot(model)
plot(pred)

## Not run:

MakePlots <- function(model, ask = TRUE) {
  ## Make all the plots callable by plot.bsts.
  opar <- par(ask = ask)
  on.exit(par(opar))
  plot.types <- c("state", "components", "residuals",
                 "prediction.errors", "forecast.distribution")
  for (plot.type in plot.types) {
    plot(model, plot.type)
  }
  if (model$has.regression) {
    regression.plot.types <- c("coefficients", "predictors", "size")
    for (plot.type in regression.plot.types) {
      plot(model, plot.type)
    }
  }
}

## Example 2: GOOG is the Google stock price, an xts series of daily
## data.
data(goog)
ss <- AddSemilocalLinearTrend(list(), goog)
model <- bst(goog, state.specification = ss, niter = 500)
MakePlots(model)

## Example 3: Change GOOG to be zoo, and not xts.
goog <- zoo(goog, index(goog))
ss <- AddSemilocalLinearTrend(list(), goog)
model <- bst(goog, state.specification = ss, niter = 500)
MakePlots(model)

## Example 4: Naked numeric data works too
y <- rnorm(100)
ss <- AddLocalLinearTrend(list(), y)
model <- bst(y, state.specification = ss, niter = 500)
MakePlots(model)

## Example 5: zoo data with intra-day measurements
y <- zoo(rnorm(100),

```

```

        seq(from = as.POSIXct("2012-01-01 7:00 EST"), len = 100, by = 100))
ss <- AddLocalLinearTrend(list(), y)
model <- bst(y, state.specification = ss, niter = 500)
MakePlots(model)

\dontrun{
  ## Example 6: Including regressors
  data(iclaims)
  ss <- AddLocalLinearTrend(list(), initial.claims$iclaimsNSA)
  ss <- AddSeasonal(ss, initial.claims$iclaimsNSA, nseasons = 52)
  model <- bst(iclaimsNSA ~ ., state.specification = ss, data =
    initial.claims, niter = 1000)

  plot(model)
  plot(model, "components")
  plot(model, "coefficients")
  plot(model, "predictors")
}

## End(Not run)

## Not run:
## Example 7: Regressors with multiple time stamps.
number.of.time.points <- 50
sample.size.per.time.point <- 10
total.sample.size <- number.of.time.points * sample.size.per.time.point
sigma.level <- .1
sigma.obs <- 1

## Simulate some fake data with a local level state component.
trend <- cumsum(rnorm(number.of.time.points, 0, sigma.level))
predictors <- matrix(rnorm(total.sample.size * 2), ncol = 2)
colnames(predictors) <- c("X1", "X2")
coefficients <- c(-10, 10)
regression <- as.numeric(predictors %*% coefficients)
y.hat <- rep(trend, sample.size.per.time.point) + regression
y <- rnorm(length(y.hat), y.hat, sigma.obs)

## Create some time stamps, with multiple observations per time stamp.
first <- as.POSIXct("2013-03-24")
dates <- seq(from = first, length = number.of.time.points, by = "month")
timestamps <- rep(dates, sample.size.per.time.point)

## Run the model with a local level trend, and an unnecessary seasonal component.
ss <- AddLocalLevel(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 7)
model <- bst(y ~ predictors, ss, niter = 250, timestamps = timestamps,
  seed = 8675309)

plot(model)
plot(model, "components")

## End(Not run)

## Example 8: Non-Gaussian data

```

```
## Poisson counts of shark attacks in Florida.
data(shark)
logshark <- log1p(shark$Attacks)
ss.level <- AddLocalLevel(list(), y = logshark)
model <- bsts(shark$Attacks, ss.level, niter = 500,
             ping = 250, family = "poisson", seed = 8675309)

## Poisson data can have an 'exposure' as the second column of a
## two-column matrix.
model <- bsts(cbind(shark$Attacks, shark$Population / 1000),
             state.specification = ss.level, niter = 500,
             family = "poisson", ping = 250, seed = 8675309)
```

bsts.options.Rd

Bsts Model Options

Description

Rarely used modeling options for bsts models.

Usage

```
BstsOptions(save.state.contributions = TRUE,
            save.prediction.errors = TRUE,
            bma.method = c("SSVS", "ODA"),
            oda.options = list(
              fallback.probability = 0.0,
              eigenvalue.fudge.factor = 0.01),
            timeout.seconds = Inf,
            save.full.state = FALSE)
```

Arguments

`save.state.contributions`

Logical. If TRUE then a 3-way array named `state.contributions` will be stored in the returned object. The indices correspond to MCMC iteration, state model number, and time. Setting `save.state.contributions` to FALSE yields a smaller object, but `plot` will not be able to plot the "state", "components", or "residuals" for the fitted model.

`save.prediction.errors`

Logical. If TRUE then a matrix named `one.step.prediction.errors` will be saved as part of the model object. The rows of the matrix represent MCMC iterations, and the columns represent time. The matrix entries are the one-step-ahead prediction errors from the Kalman filter.

bma.method	If the model contains a regression component, this argument specifies the method to use for Bayesian model averaging. "SSVS" is stochastic search variable selection, which is the classic approach from George and McCulloch (1997). "ODA" is orthogonal data augmentation, from Ghosh and Clyde (2011). It adds a set of latent observations that make the $X^T X$ matrix diagonal, vastly simplifying complete data MCMC for model selection.
oda.options	If bma.method == "ODA" then these are some options for fine tuning the ODA algorithm. <ul style="list-style-type: none"> • fallback.probability: Each MCMC iteration will use SSVS instead of ODA with this probability. In cases where the latent data have high leverage, ODA mixing can suffer. Mixing in a few SSVS steps can help keep an errant algorithm on track. • eigenvalue.fudge.factor: The latent X's will be chosen so that the complete data $X^T X$ matrix (after scaling) is a constant diagonal matrix equal to the largest eigenvalue of the observed (scaled) $X^T X$ times (1 + eigenvalue.fudge.factor). This should be a small positive number.
timeout.seconds	The number of seconds that sampler will be allowed to run. If the timeout is exceeded the returned object will be truncated to the final draw that took place before the timeout occurred, as if that had been the requested number of iterations.
save.full.state	Logical. If TRUE then the full distribution of the state vector will be preserved. It will be stored in the model under the name full.state, which is a 3-way array with dimensions corresponding to MCMC iteration, state dimension, and time.

Value

The arguments are checked to make sure they have legal types and values, then a list is returned containing the arguments.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

compare.bsts.models *Compare bsts models*

Description

Produce a set of line plots showing the cumulative absolute one step ahead prediction errors for different models. This plot not only shows which model is doing the best job predicting the data, it highlights regions of the data where the predictions are particularly good or bad.

Usage

```
CompareBstsModels(model.list,
                  burn = SuggestBurn(.1, model.list[[1]]),
                  filename = "",
                  colors = NULL,
                  lwd = 2,
                  xlab = "Time",
                  main = "",
                  grid = TRUE,
                  cutpoint = NULL)
```

Arguments

<code>model.list</code>	A list of bsts models.
<code>burn</code>	The number of initial MCMC iterations to remove from each model as burn-in.
<code>filename</code>	A string. If non-empty string then a pdf of the plot will be saved in the specified file.
<code>colors</code>	A vector of colors to use for the different lines in the plot. If NULL then the rainbow palette will be used.
<code>lwd</code>	The width of the lines to be drawn.
<code>xlab</code>	Label for the horizontal axis.
<code>main</code>	Main title for the plot.
<code>grid</code>	Logical. Should gridlines be drawn in the background?
<code>cutpoint</code>	Either NULL, or an integer giving the observation number used to define a holdout sample. Prediction errors occurring after the cutpoint will be true out of sample errors. If NULL then all prediction errors are "in sample". See the discussion in bsts.prediction.errors .

Value

Invisibly returns the matrix of cumulative one-step ahead prediction errors (the lines in the top panel of the plot). Each row in the matrix corresponds to a model in `model.list`.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
trend.only <- bsts(y, ss, niter = 250)

ss <- AddSeasonal(ss, y, nseasons = 12)
trend.and.seasonal <- bsts(y, ss, niter = 250)
```

```
CompareBstsModels(list(trend = trend.only,
                       "trend and seasonal" = trend.and.seasonal))
```

```
CompareBstsModels(list(trend = trend.only,
                       "trend and seasonal" = trend.and.seasonal),
                    cutpoint = 100)
```

date.range

Date Range

Description

Returns the first and last dates of the influence window for the given holiday, among the given timestamps.

Usage

```
DateRange(holiday, timestamps)
```

Arguments

holiday	An object of class Holiday .
timestamps	A vector of timestamps of class Date or class POSIXt . This function assumes daily data. Use with care in other settings.

Value

Returns a two-column data frame giving the first and last dates of the influence window for the holiday in the period covered by timestamps.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
holiday <- NamedHoliday("MemorialDay", days.before = 2, days.after = 2)
timestamps <- seq.Date(from = as.Date("2001-01-01"), by = "day",
                       length.out = 365 * 10)
influence <- DateRange(holiday, timestamps)
```

Description

Plots for describing time series data.

Usage

```
DayPlot(y, colors = NULL, ylab = NULL, ...)  
MonthPlot(y, seasonal.identifier = months, colors = NULL, ylab = NULL, ...)  
YearPlot(y, colors = NULL, ylab = NULL, ylim = NULL, legend = TRUE, ...)
```

Arguments

<code>y</code>	A time series to plot. Must be of class <code>ts</code> , or <code>zoo</code> . If a <code>zoo</code> object then the timestamps must be of type <code>Date</code> , <code>yearmon</code> , or <code>POSIXt</code> .
<code>seasonal.identifier</code>	A function that takes a vector of class <code>POSIXt</code> (date/time) and returns a character vector indicating the season to which each element belongs. Each unique element returned by this function returns a "season" to be plotted. See <code>weekdays</code> , <code>months</code> , and <code>quarters</code> for examples of how this should work.
<code>colors</code>	A vector of colors to use for the lines.
<code>legend</code>	Logical. If <code>TRUE</code> then a legend is added to the plot.
<code>ylab</code>	Label for the vertical axis.
<code>ylim</code>	Limits for the vertical axis. (a 2-vector)
<code>...</code>	Extra arguments passed to <code>plot</code> or <code>lines</code> .

Details

`DayPlot` and `MonthPlot` plot the time series one season at a time, on the same set of axes. The intent is to use `DayPlot` for daily data and `MonthPlot` for monthly or quarterly data.

`YearPlot` plots each year of the time series as a separate line on the same set of axes.

Both sets of plots help visualize seasonal patterns.

Value

Returns `invisible{NULL}`.

See Also

`monthplot` is a base R function for plotting time series of type `ts`.

Examples

```
## Plot a 'ts' time series.
data(AirPassengers)
par(mfrow = c(1,2))
MonthPlot(AirPassengers)
YearPlot(AirPassengers)

## Plot a 'zoo' time series.
data(turkish)
par(mfrow = c(1,2))
YearPlot(turkish)
DayPlot(turkish)
```

diagnostic-plots

Diagnostic Plots

Description

Diagnostic plots for distributions of residuals.

Usage

```
qqdist(draws, ...)
AcfDist(draws, lag.max = NULL, xlab = "Lag", ylab = "Autocorrelation", ...)
```

Arguments

draws	A matrix of Monte Carlo draws of residual errors. Each row is a Monte Carlo draw, and each column is an observation. In the case of <code>AcfDist</code> successive observations are assumed to be sequential in time.
lag.max	The number of lags to plot in the autocorrelation function. See acf .
xlab	Label for the horizontal axis.
ylab	Label for the vertical axis.
...	Extra arguments passed to either boxplot (for <code>AcfDist</code>) or PlotDynamicDistribution (for <code>qqdist</code>).

Details

`qqdist` sorts the columns of `draws` by their mean, and plots the resulting set of curves against the quantiles of the standard normal distribution. A reference line is added, and the mean of each column of `draws` is represented by a blue dot. The dots and the line are the transpose of what you get with [qqnorm](#) and [qqline](#).

`AcfDist` plots the posterior distribution of the autocorrelation function using a set of side-by-side boxplots.

Examples

```
data(AirPassengers)
y <- log(AirPassengers)

ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, ss, niter = 500)

r <- residuals(model)
par(mfrow = c(1,2))
qqdist(r) ## A bit of departure in the upper tail
AcfDist(r)
```

dirm

Dynamic intercept regression model

Description

A dynamic intercept regression is a regression model where the intercept term is a state space model. This model differs from [bsts](#) in that there can be multiple observations per time point.

Usage

```
dirm(formula,
     state.specification,
     data,
     prior = NULL,
     contrasts = NULL,
     na.action = na.pass,
     niter,
     ping = niter / 10,
     model.options = DirmModelOptions(),
     timestamps = NULL,
     seed = NULL,
     ...)
```

Arguments

formula A formula, as you would supply to [lm](#) describing the regression portion of the relationship between y and X.

state.specification A list with elements created by [AddLocalLinearTrend](#), [AddSeasonal](#), and similar functions for adding components of state. See the help page for [state.specification](#). The state specification describes the dynamic intercept term in the regression model.

<code>data</code>	An optional data frame, list or environment (or object coercible by <code>as.data.frame</code> to a data frame) containing the variables in the model. If not found in <code>data</code> , the variables are taken from <code>environment(formula)</code> , typically the environment from which <code>dirm</code> is called.
<code>prior</code>	A prior distribution for the regression component of the model, as created by <code>SpikeSlabPrior</code> . The prior for the time series component of the model will be specified during the creation of <code>state.specification</code> .
<code>contrasts</code>	An optional list containing the names of contrast functions to use when converting factors numeric variables in a regression formula. This argument works exactly as it does in <code>lm</code> . The names of the list elements correspond to factor variables in your model formula. The list elements themselves are the names of contrast functions (see <code>help(contr.treatment)</code> and the <code>contrasts.arg</code> argument to <code>model.matrix.default</code>). This argument can usually be omitted.
<code>na.action</code>	What to do about missing values. The default is to allow missing responses, but no missing predictors. Set this to <code>na.omit</code> or <code>na.exclude</code> if you want to omit missing responses altogether.
<code>niter</code>	A positive integer giving the desired number of MCMC draws.
<code>ping</code>	A scalar giving the desired frequency of status messages. If <code>ping > 0</code> then the program will print a status message to the screen every <code>ping</code> MCMC iterations.
<code>model.options</code>	An object created by <code>DirmModelOptions</code> specifying the desired model options.
<code>timestamps</code>	The timestamp associated with each value of the response. This is most likely a <code>Date</code> or <code>POSIXt</code> . It is expected that there will be multiple observations per time point (otherwise <code>'bsts'</code> should be used instead of <code>'dirm'</code>), and thus the <code>'timestamps'</code> argument will contain many duplicate values.
<code>seed</code>	An integer to use as the random seed for the underlying C++ code. If <code>NULL</code> then the seed will be set using the clock.
<code>...</code>	Extra arguments to be passed to <code>SpikeSlabPrior</code> (see the entry for the <code>prior</code> argument, above).

Details

The fitted model is a regression model with an intercept term given by a structural time series model. This is similar to the model fit by `bsts`, but it allows for multiple observations per time period.

Currently `dirm` only supports Gaussian observation errors, but look for that to change in future releases.

Value

An object of class `bsts` which is a list with the following components

<code>coefficients</code>	A <code>niter</code> by <code>ncol(X)</code> matrix of MCMC draws of the regression coefficients, where <code>X</code> is the design matrix implied by <code>formula</code> . This is only present if a model formula was supplied.
<code>sigma.obs</code>	A vector of length <code>niter</code> containing MCMC draws of the residual standard deviation.

The returned object will also contain named elements holding the MCMC draws of model parameters belonging to the state models. The names of each component are supplied by the entries in `state.specification`. If a model parameter is a scalar, then the list element is a vector with `niter` elements. If the parameter is a vector then the list element is a matrix with `niter` rows. If the parameter is a matrix then the list element is a 3-way array with first dimension `niter`.

Finally, if a model formula was supplied, then the returned object will contain the information necessary for the `predict` method to build the design matrix when a new prediction is made.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

George and McCulloch (1997) "Approaches for Bayesian variable selection", *Statistica Sinica* pp 339–374.

See Also

[bsts](#), [AddLocalLevel](#), [AddLocalLinearTrend](#), [AddSemilocalLinearTrend](#), [AddSeasonal](#) [AddDynamicRegression](#) [SpikeSlabPrior](#), [SdPrior](#).

Examples

```
SimulateDirmData <- function(observation.sd = 1, trend.sd = .1,
                             time.dimension = 100, nobs.per.period = 3,
                             xdim = 4) {
  trend <- cumsum(rnorm(time.dimension, 0, trend.sd))
  total.sample.size <- nobs.per.period * time.dimension
  predictors <- matrix(rnorm(total.sample.size * xdim),
                       nrow = total.sample.size)
  coefficients <- rnorm(xdim)
  expanded.trend <- rep(trend, each = nobs.per.period)
  response <- expanded.trend + predictors %*% coefficients + rnorm(
    total.sample.size, 0, observation.sd)
  timestamps <- seq.Date(from = as.Date("2008-01-01"),
                         len = time.dimension, by = "day")
  extended.timestamps <- rep(timestamps, each = nobs.per.period)
  return(list(response = response,
              predictors = predictors,
              timestamps = extended.timestamps,
              trend = trend,
              coefficients = coefficients))
}
```

```
data <- SimulateDirmData(time.dimension = 20)
ss <- AddLocalLevel(list(), data$response)

# In real life you'd want more than 50 MCMC iterations.
model <- dirm(data$response ~ data$predictors, ss, niter = 50,
  timestamps = data$timestamps)
```

dirm-model-opts *Specify Options for a Dynamic Intercept Regression Model*

Description

Specify modeling options for a dynamic intercept regression model.

Usage

```
DirmModelOptions(timeout.seconds = Inf,
  high dimensional.threshold.factor = 1.0)
```

Arguments

`timeout.seconds`

The overall time budget for model fitting. If the MCMC algorithm takes longer than this number, the current iteration will complete, and then the fitting algorithm will return with however many MCMC iterations were managed during the allotted time.

`high dimensional.threshold.factor`

When doing Kalman filter updates for the model, Sherman-Morrisson-Woodbury style updates are applied for high dimensional data, while direct linear algebra is used for low dimensional data. The definition of "high dimensional" is relative to the dimension of the state. An observation is considered high dimensional if its dimension exceeds the state dimension times this factor.

Value

An object of class `DirmModelOptions`, which is simply a list containing values of the function arguments.

The value of using this function instead of making a list "by hand" is that argument types are properly checked, and list names are sure to be correct.

estimate.time.scale *Intervals between dates*

Description

Estimate the time scale used in time series data.

Usage

```
EstimateTimeScale(dates)
```

Arguments

dates A sorted vector of class `Date`.

Value

A character string. Either "daily", "weekly", "yearly", "monthly", "quarterly", or "other". The value is determined based on counting the number of days between successive observations in dates.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
weekly.data <- as.Date(c("2011-10-01",
                        "2011-10-08",
                        "2011-10-15",
                        "2011-10-22",
                        "2011-10-29",
                        "2011-11-05"))

EstimateTimeScale(weekly.data) # "weekly"

almost.weekly.data <- as.Date(c("2011-10-01",
                                "2011-10-08",
                                "2011-10-15",
                                "2011-10-22",
                                "2011-10-29",
                                "2011-11-06")) # last day is one later

EstimateTimeScale(weekly.data) # "other"
```

extend.time *Extends a vector of dates to a given length*

Description

Pads a vector of dates to a specified length.

Usage

```
ExtendTime(dates, number.of.periods, dt = NULL)
```

Arguments

dates An ordered vector of class [Date](#).
number.of.periods The desired length of the output.
dt A character string describing the frequency of the dates in dates. Possible values are "daily", "weekly", "monthly", "quarterly", "yearly", or "other". An attempt to deduce dt will be made if it is missing.

Value

If number.of.periods is longer than length(dates), then dates will be padded to the desired length. Extra dates are added at time intervals matching the average interval in dates. Thus they may not be

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[bsts.mixed](#).

Examples

```
origin.month <- as.Date("2011-09-01")
week.ending <- as.Date(c("2011-10-01", ## 1
                        "2011-10-08", ## 2
                        "2011-12-03", ## 3
                        "2011-12-31")) ## 4
MatchWeekToMonth(week.ending, origin.month) == 1:4
```

Description

Tools for checking if a series of timestamps is 'regular' meaning that it has no duplicates, and no gaps. Checking for regularity can be tricky. For example, if you have monthly observations with [Date](#) or [POSIXt](#) timestamps then gaps between timestamps can be 28, 29, 30, or 31 days, but the series is still "regular".

Usage

```
NoDuplicates(timestamps)
NoGaps(timestamps)
IsRegular(timestamps)

HasDuplicateTimestamps(bsts.object)
```

Arguments

<code>timestamps</code>	A set of (possibly irregular or non-unique) timestamps. This could be a set of integers (like 1, 2, , 3...), a set of numeric like (1945, 1945.083, 1945.167, ...) indicating years and fractions of years, a Date object, or a POSIXt object.
<code>bsts.object</code>	A <code>bsts</code> model object.

Value

All four functions return scalar logical values. `NoDuplicates` returns `TRUE` if all elements of `timestamps` are unique.

`NoGaps` examines the smallest nonzero gap between time points. As long as no gaps between time points are more than twice as wide as the smallest gap, it returns `TRUE`, indicating that there are no missing timestamps. Otherwise it returns `FALSE`.

`IsRegular` returns `TRUE` if `NoDuplicates` and `NoGaps` both return `TRUE`.

`HasDuplicateTimestamps` returns `FALSE` if the data used to fit `bsts.model` either has `NULL` timestamps, or if the timestamps contain no duplicate values.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
first <- as.POSIXct("2015-04-19 08:00:04")
monthly <- seq(from = first, length.out = 24, by = "month")
IsRegular(monthly) ## TRUE

skip.one <- monthly[-8]
```

```

IsRegular(skip.one) ## FALSE

has.duplicates <- monthly
has.duplicates[1] <- has.duplicates[2]
IsRegular(has.duplicates) ## FALSE

```

gdp

Gross Domestic Product for 57 Countries

Description

Annual gross domestic product for 57 countries, as produced by the OECD.

Fields:

- LOCATION: Three letter country code.
- MEASURE: MLN_USD signifies a total GDP number in millions of US dollars. USD_CAP is per capita GDP in US dollars.
- TIME: The year of the measurement.
- Value: The measured value.
- Flag.Codes: P for provisional data, B for a break in the series, and E for an estimated value.

Usage

```
data(gdp)
```

Format

data frame

Source

OECD website: See <https://data.oecd.org/gdp/gross-domestic-product-gdp.htm>

geometric.sequence

Create a Geometric Sequence

Description

Create a geometric sequence.

Usage

```
GeometricSequence(length, initial.value = 1, discount.factor = .5)
```

Arguments

length A positive integer giving the length of the desired sequence.
 initial.value The first term in the sequence. Cannot be zero.
 discount.factor The ratio between a sequence term and the preceding term. Cannot be zero.

Value

A numeric vector containing the desired sequence.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
GeometricSequence(4, .8, .6)
# [1] 0.8000 0.4800 0.2880 0.1728

GeometricSequence(5, 2, 3)
# [1] 2 6 18 54 162

## Not run:
GeometricSequence(0, -1, -2)
# Error: length > 0 is not TRUE

## End(Not run)
```

get.fraction	<i>Compute membership fractions</i>
--------------	-------------------------------------

Description

Returns the fraction of days in a week that occur in the ear

Usage

```
GetFractionOfDaysInInitialMonth(week.ending)
GetFractionOfDaysInInitialQuarter(week.ending)
```

Arguments

week.ending A vector of class [Date](#). Each entry contains the date of the last day in a week.

Value

Returns a numeric vector of the same length as week.ending. Each entry gives the fraction of days in the week that occur in the coarse time interval (month or quarter) containing the start of the week (i.e the date 6 days before).

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[bsts.mixed](#).

Examples

```
dates <- as.Date(c("2003-03-31",
                  "2003-04-01",
                  "2003-04-02",
                  "2003-04-03",
                  "2003-04-04",
                  "2003-04-05",
                  "2003-04-06",
                  "2003-04-07"))
fraction <- GetFractionOfDaysInInitialMonth(dates)
fraction == c(1, 6/7, 5/7, 4/7, 3/7, 2/7, 1/7, 1)
```

goog

Google stock price

Description

Daily closing price of Google stock.

Usage

```
data(goog)
```

Format

xts time series

Source

The Internets

HarveyCumulator	<i>HarveyCumulator</i>
-----------------	------------------------

Description

Given a state space model on a fine scale, the Harvey cumulator aggregates the model to a coarser scale (e.g. from days to weeks, or weeks to months).

Usage

```
HarveyCumulator(fine.series,  
                contains.end,  
                membership.fraction)
```

Arguments

<code>fine.series</code>	The fine-scale time series to be aggregated.
<code>contains.end</code>	A logical vector, with length matching <code>fine.series</code> indicating whether each fine scale time interval contains the end of a coarse time interval. For example, months don't contain a fixed number of weeks, so when cumulating a weekly time series into a monthly series, you need to know which weeks contain the end of a month.
<code>membership.fraction</code>	The fraction of each fine-scale time observation belonging to the coarse scale time observation at the beginning of the time interval. For example, if week <code>i</code> started in March and ended in April, <code>membership.fraction[i]</code> is the fraction of <code>fine.series[i]</code> that should be attributed to March. This should be 1 for most observations.

Value

Returns a vector containing the course scale partial aggregates of `fine.series`.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.mixed](#),

Examples

```

data(goog)
days <- factor(weekdays(index(goog)),
               levels = c("Monday", "Tuesday", "Wednesday",
                          "Thursday", "Friday"),
               ordered = TRUE)

## Because of holidays, etc the days do not always go in sequence.
## (Sorry, Rebecca Black! https://www.youtube.com/watch?v=kfVsf0SbJY0)
## diff.days[i] is the number of days between days[i-1] and days[i].
## We know that days[i] is the end of a week if diff.days[i] < 0.
diff.days <- tail(as.numeric(days), -1) - head(as.numeric(days), -1)
contains.end <- c(FALSE, diff.days < 0)

goog.weekly <- HarveyCumulator(goog, contains.end, 1)

```

 holiday

Specifying Holidays

Description

Specify holidays for use with holiday state models.

Usage

```

FixedDateHoliday(holiday.name,
                 month = base::month.name,
                 day,
                 days.before = 1,
                 days.after = 1)

NthWeekdayInMonthHoliday(holiday.name,
                         month = base::month.name,
                         day.of.week = weekday.names,
                         week.number = 1,
                         days.before = 1,
                         days.after = 1)

LastWeekdayInMonthHoliday(holiday.name,
                          month = base::month.name,
                          day.of.week = weekday.names,
                          days.before = 1,
                          days.after = 1)

NamedHoliday(holiday.name = named.holidays,
             days.before = 1,

```

```
days.after = 1)
```

```
DateRangeHoliday(holiday.name,
                  start.date,
                  end.date)
```

Arguments

holiday.name	A string that can be used to label the holiday in output.
month	A string naming the month in which the holiday occurs. Unambiguous partial matches are acceptable. Capitalize the first letter.
day	An integer specifying the day of the month on which the FixedDateHoliday occurs.
day.of.week	A string giving the day of the week on which the holiday occurs.
week.number	An integer specifying the week of the month on which the NthWeekdayInMonthHoliday occurs.
days.before	An integer giving the number of days of influence that the holiday exerts prior to the actual holiday.
days.after	An integer giving the number of days of influence that holiday exerts after the actual holiday.
named.holidays	A character vector containing one or more recognized holiday names.
start.date	A vector of starting dates for the holiday. Each instance of the holiday in the training data or the forecast period must be represented by an element in this vector. Thus if this is an annual holiday and, there are 10 years of training data, and a 1-year forecast is needed, then this will be a vector of length 11.
end.date	A vector of ending dates for the holiday. Each date must occur on or after the corresponding element of start.date, and end.date[i] must come before start.date[i+1].

Value

Each function returns a list containing the information from the function arguments, formatted as expected by the underlying C++ code. State models that focus on holidays, such as [AddRandomWalkHoliday](#), [AddRegressionHoliday](#), and [AddHierarchicalRegressionHoliday](#), will expect one or more holiday objects as arguments.

- `FixedDateHoliday` describes a holiday that occurs on the same date each year, like US independence day (July 4).
- `NthWeekdayInMonthHoliday` describes a holiday that occurs a particular weekday of a particular week of a particular month. For example, US Labor Day is the first Monday in September.
- `LastWeekdayInMonthHoliday` describes a holiday that occurs on the last instance of a particular weekday in a particular month. For example, US Memorial Day is the last Monday in May.
- `DateRangeHoliday` describes an irregular holiday that might not follow a particular pattern. You can handle this type of holiday by manually specifying a range of dates for each instance of the holiday in your data set. NOTE: If you plan on using the model to forecast, be sure to include date ranges in the forecast period as well as the period covered by the training data.

- NamedHoliday is a convenience class for describing several important holidays in the US.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[AddRandomWalkHoliday](#), [AddRegressionHoliday](#), [AddHierarchicalRegressionHoliday](#)

Examples

```
july4 <- FixedDateHoliday("July4", "July", 4)
memorial.day <- LastWeekdayInMonthHoliday("MemorialDay", "May", "Monday")
labor.day <- NthWeekdayInMonthHoliday("LaborDay", "September", "Monday", 1)
another.way.to.get.memorial.day <- NamedHoliday("MemorialDay")
easter <- NamedHoliday("Easter")
winter.olympics <- DateRangeHoliday("WinterOlympicsSince2000",
  start = as.Date(c("2002-02-08",
    "2006-02-10",
    "2010-02-12",
    "2014-02-07",
    "2018-02-07")),
  end = as.Date(c("2002-02-24",
    "2006-02-26",
    "2010-02-28",
    "2014-02-23",
    "2018-02-25")))
```

iclaims

Initial Claims Data

Description

A weekly time series of US initial claims for unemployment. The first column contains the initial claims numbers from FRED. The others contain a measure of the relative popularity of various search queries identified by Google Correlate.

Usage

```
data(iclaims)
```

Format

zoo time series

Source

FRED. <http://research.stlouisfed.org/fred2/series/ICNSA>,
 Google correlate. <http://www.google.com/trends/correlate>

See Also[bsts](#)**Examples**

```
data(iclaims)
plot(initial.claims)
```

last.day.in.month	<i>Find the last day in a month</i>
-------------------	-------------------------------------

Description

Finds the last day in the month containing a specified date.

Usage

```
LastDayInMonth(dates)
```

Arguments

dates A vector of class [Date](#).

Value

A vector of class [Date](#) where each entry is the last day in the month containing the corresponding entry in dates.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
inputs <- as.Date(c("2007-01-01",
                   "2007-01-31",
                   "2008-02-01",
                   "2008-02-29",
                   "2008-03-14",
                   "2008-12-01",
                   "2008-12-31"))
expected.outputs <- as.Date(c("2007-01-31",
                              "2007-01-31",
                              "2008-02-29",
                              "2008-02-29",
                              "2008-03-31",
                              "2008-12-31",
                              "2008-12-31"))
LastDayInMonth(inputs) == expected.outputs
```

 MATCH.NumericTimestamps

Match Numeric Timestamps

Description

S3 generic method for MATCH function supplied in the zoo package.

Usage

```
## S3 method for class 'NumericTimestamps'
MATCH(x, table, nomatch = NA, ...)
```

Arguments

x	A numeric set of timestamps.
table	A set of regular numeric timestamps to match against.
nomatch	The value to be returned in the case when no match is found. Note that it is coerced to integer.
...	Additional arguments passed to match .

Details

Numeric timestamps match if they agree to 8 significant digits.

Value

Returns the index of the entry in table matched by each argument in x. If an entry has no match then nomatch is returned at that position.

See Also

[MATCH](#)

 match.week.to.month *Find the month containing a week*

Description

Returns the index of a month, in a sequence of months, that contains a given week.

Usage

```
MatchWeekToMonth(week.ending, origin.month)
```

Arguments

week.ending A vector of class [Date](#). Each entry contains the date of the last day in a week.
 origin.month A [Date](#), giving any day in the month to use as the origin of the sequence (month 1).

Value

The index of the month matching the month containing the first day in week.ending. The origin is month 1. It is the caller's responsibility to ensure that these indices correspond to legal values in a particular vector of months.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[bsts.mixed](#).

Examples

```
origin.month <- as.Date("2011-09-01")
week.ending <- as.Date(c("2011-10-01", ## 1
                        "2011-10-08", ## 2
                        "2011-12-03", ## 3
                        "2011-12-31")) ## 4
MatchWeekToMonth(week.ending, origin.month) == 1:4
```

max.window.width	<i>Maximum Window Width for a Holiday</i>
------------------	---

Description

The maximum width of a holiday's influence window

Usage

```
## Default S3 method:
MaxWindowWidth(holiday, ...)
## S3 method for class 'DateRangeHoliday'
MaxWindowWidth(holiday, ...)
```

Arguments

holiday An object of class [Holiday](#).
 ... Other arguments (not used).

Value

Returns the number of days in a holiday's influence window.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[Holiday](#). [AddRegressionHoliday](#). [AddRandomWalkHoliday](#). [AddHierarchicalRegressionHoliday](#).

Examples

```
easter <- NamedHoliday("Easter", days.before = 2, days.after = 1)
if (MaxWindowWidth(easter) == 4) {
  print("That's the right answer!\n")
}

## This holiday lasts two days longer in 2005 than in 2004.
may18 <- DateRangeHoliday("May18",
  start = as.Date(c("2004-05-17",
    "2005-05-16")),
  end   = as.Date(c("2004-05-19",
    "2005-05-20")))

if (MaxWindowWidth(may18) == 5) {
  print("Right again!\n")
}
```

mbsts

Multivariate Bayesian Structural Time Series

Description

Fit a multivariate Bayesian structural time series model, also known as a "dynamic factor model."

**** NOTE **** This code is experimental. Please feel free to experiment with it and report any bugs to the maintainer. Expect it to improve substantially in the next release.

Usage

```
mbsts(formula,
  shared.state.specification,
  series.state.specification = NULL,
  data = NULL,
  timestamps = NULL,
  series.id = NULL,
  prior = NULL, # TODO
```

```

opts = NULL,
contrasts = NULL,
na.action = na.pass,
niter,
ping = niter / 10,
data.format = c("long", "wide"),
seed = NULL,
...)
```

Arguments

- formula** A formula describing the regression portion of the relationship between y and X.
If no regressors are desired then the formula can be replaced by a numeric matrix giving the multivariate time series to be modeled.
- shared.state.specification**
A list with elements created by [AddSharedLocalLevel](#), and similar functions for adding components of state.
This list defines the components of state which are shared across all time series. These are the "factors" in the dynamic factor model.
- series.state.specification**
This argument specifies state components needed by a particular series. Not all series need have the same state components (e.g. some series may require a seasonal component, while others do not). It can be NULL, indicating that there are no series-specific state components.
It can be a list of elements created by [AddLocalLevel](#), [AddSeasonal](#), and similar functions for adding state component to scalar bsts models. In this case the same, independent, individual components will be added to each series. For example, each series will get its own independent Seasonal state component if [AddSeasonal](#) was used to add a seasonal component to this argument.
In its most general form, this argument can be a list of lists, some of which can be NULL, but with non-NULL lists specifying state components for individual series, as above.
- data** An optional data frame, list or environment (or object coercible by [as.data.frame](#) to a data frame) containing the variables mentioned in the `formula` argument. If not found in `data`, the variables are taken from `environment(formula)`, typically the environment from which `bsts` is called.
- timestamps** A vector of timestamps indicating the time of each observation. If `data.format` is "long" then this argument is required. If "wide" data is passed then it is optional.
TODO: TEST THIS under wide and long formats in regression and non-regression settings.
- series.id** A factor (or object coercible to factor) indicating the series to which each observation in "long" format belongs. This argument is ignored for data in "wide" format.

prior	A list of <code>SpikeSlabPrior</code> objects, one for each time series. Or this argument can be NULL in which case a default prior will be used. Note that the prior is on both the regression coefficients and the residual sd for each time series.
opts	A list containing model options. This is currently only used for debugging, so leave this as NULL.
contrasts	An optional list containing the names of contrast functions to use when converting factors numeric variables in a regression formula. This argument works exactly as it does in <code>lm</code> . The names of the list elements correspond to factor variables in your model formula. The list elements themselves are the names of contrast functions (see <code>help(contr.treatment)</code> and the <code>contrasts.arg</code> argument to <code>model.matrix.default</code>). This argument is only used if a model formula is specified, and even then the default is probably what you want.
na.action	What to do about missing values. The default is to allow missing responses, but no missing predictors. Set this to <code>na.omit</code> or <code>na.exclude</code> if you want to omit missing responses altogether.
niter	A positive integer giving the desired number of MCMC draws.
ping	A scalar giving the desired frequency of status messages. If <code>ping > 0</code> then the program will print a status message to the screen every <code>ping</code> MCMC iterations.
data.format	Whether the data are store in wide (each row is a time point, and columns are values from different series) or long (each row is the value of a particular series at a particular point in time) format. For "long" see <code>timestamps</code> and <code>series.id</code> .
seed	An integer to use as the random seed for the underlying C++ code. If NULL then the seed will be set using the clock.
...	Extra arguments to be passed to <code>SpikeSlabPrior</code> (see the entry for the prior argument, above).

Value

An object of class `mbsts` which is a list with the following components

coefficients	A <code>niter</code> by <code>ncol(X)</code> matrix of MCMC draws of the regression coefficients, where <code>X</code> is the design matrix implied by formula. This is only present if a model formula was supplied.
sigma.obs	A vector of length <code>niter</code> containing MCMC draws of the residual standard deviation.

The returned object will also contain named elements holding the MCMC draws of model parameters belonging to the state models. The names of each component are supplied by the entries in `state.specification`. If a model parameter is a scalar, then the list element is a vector with `niter` elements. If the parameter is a vector then the list element is a matrix with `niter` rows. If the parameter is a matrix then the list element is a 3-way array with first dimension `niter`.

Finally, if a model formula was supplied, then the returned object will contain the information necessary for the `predict` method to build the design matrix when a new prediction is made.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

- Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.
- Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.
- George and McCulloch (1997) "Approaches for Bayesian variable selection", *Statistica Sinica* pp 339–374.

See Also

[bsts](#), [AddLocalLevel](#), [AddLocalLinearTrend](#), [AddSemilocalLinearTrend](#), [AddSeasonal](#), [AddDynamicRegression](#), [SpikeSlabPrior](#), [SdPrior](#).

Examples

```
## Not run:

# This example takes 12s on Windows, which is longer than CRAN's 10s
# limit. Marking code as 'dontrun' to prevent CRAN auto checks from
# timing out.

seed <- 8675309
set.seed(seed)

ntimes <- 250
nseries <- 20
nfactors <- 6
residual.sd <- 1.2
state.innovation.sd <- .75

##-----
## simulate latent state for fake data.
##-----
state <- matrix(rnorm(ntimes * nfactors, 0, state.innovation.sd), nrow = ntimes)
for (i in 1:ncol(state)) {
  state[, i] <- cumsum(state[, i])
}

##-----
## Simulate "observed" data from state.
##-----
observation.coefficients <- matrix(rnorm(nseries * nfactors), nrow = nseries)
diag(observation.coefficients) <- 1.0
observation.coefficients[upper.tri(observation.coefficients)] <- 0

errors <- matrix(rnorm(nseries * ntimes, 0, residual.sd), ncol = ntimes)
y <- t(observation.coefficients %*% t(state) + errors)

##-----
## Plot the data.
##-----
```

```

par(mfrow=c(1,2))
plot.ts(y, plot.type="single", col = rainbow(nseries), main = "observed data")
plot.ts(state, plot.type = "single", col = 1:nfactors, main = "latent state")

##-----
## Fit the model
##-----
ss <- AddSharedLocalLevel(list(), y, nfactors = nfactors)

opts <- list("fixed.state" = t(state),
            fixed.residual.sd = rep(residual.sd, nseries),
            fixed.regression.coefficients = matrix(rep(0, nseries), ncol = 1))

model <- mbsts(y, shared.state.specification = ss, niter = 100,
              data.format = "wide", seed = seed)

##-----
## Plot the state
##-----
par(mfrow=c(1, nfactors))
ylim <- range(model$shared.state, state)
for (j in 1:nfactors) {
  PlotDynamicDistribution(model$shared.state[, j, ], ylim=ylim)
  lines(state[, j], col = "blue")
}

##-----
## Plot the factor loadings.
##-----
opar <- par(mfrow=c(nfactors,1), mar=c(0, 4, 0, 4), omi=rep(.25, 4))
burn <- 10
for(j in 1:nfactors) {
  BoxplotTrue(model$shared.local.level.coefficients[-(1:burn), j, ],
              t(observation.coefficients)[, j], axes=F, truth.color="blue")
  abline(h=0, lty=3)
  box()
  axis(2)
}
axis(1)
par(opar)

##-----
## Plot the predicted values of the series.
##-----
index <- 1:12
nr <- floor(sqrt(length(index)))
nc <- ceiling(length(index) / nr)
opar <- par(mfrow = c(nr, nc), mar = c(2, 4, 1, 2))
for (i in index) {
  PlotDynamicDistribution(
    model$shared.state.contributions[, 1, i, ]
    + model$regression.coefficients[, i, 1]
    , ylim=range(y))
}

```

```

    points(y[, i], col="blue", pch = ".", cex = .2)
  }
  par(opar)
  # next line closes 'dontrun'

  ## End(Not run)
  # next line closes 'examples'

```

mixed.frequency *Models for mixed frequency time series*

Description

Fit a structured time series to mixed frequency data.

Usage

```

bsts.mixed(target.series,
           predictors,
           which.coarse.interval,
           membership.fraction,
           contains.end,
           state.specification,
           regression.prior,
           niter,
           ping = niter / 10,
           seed = NULL,
           truth = NULL,
           ...)

```

Arguments

- `target.series` A vector object of class `zoo` indexed by calendar dates. The date associated with each element is the LAST DAY in the time interval measured by the corresponding value. The value is what Harvey (1989) calls a 'flow' variable. It is a number that can be viewed as an accumulation over the measured time interval.
- `predictors` A matrix of class `zoo` indexed by calendar dates. The date associated with each row is the LAST DAY in the time interval encompassing the measurement. The dates are expected to be at a finer scale than the dates in `target.series`. Any predictors should be at sufficient lags to be able to predict the rest of the cycle.
- `which.coarse.interval`
 A numeric vector of length `nrow(predictors)` giving the index of the coarse interval corresponding to the end of each fine interval.
- `membership.fraction`
 A numeric vector of length `nrow(predictors)` giving the fraction of activity attributed to the coarse interval corresponding to the beginning of each fine interval. This is always positive, and will be 1 except when a fine interval spans the boundary between two coarse intervals.

<code>contains.end</code>	A logical vector of length <code>nrow(predictors)</code> indicating whether each fine interval contains the end of a coarse interval.
<code>state.specification</code>	A state specification like that required by <code>bsts</code> .
<code>regression.prior</code>	A prior distribution created by <code>SpikeSlabPrior</code> . A default prior will be generated if none is specified.
<code>niter</code>	The desired number of MCMC iterations.
<code>ping</code>	An integer indicating the frequency with which progress reports get printed. E.g. setting <code>ping = 100</code> will print a status message with a time and iteration stamp every 100 iterations. If you don't want these messages set <code>ping < 0</code> .
<code>seed</code>	An integer to use as the random seed for the underlying C++ code. If <code>NULL</code> then the seed will be set using the clock.
<code>truth</code>	For debugging purposes only. A list containing one or more of the following elements. If any are present then corresponding values will be held fixed in the MCMC algorithm. <ul style="list-style-type: none"> • A matrix named <code>state</code> containing the state of the coarse model from a fake-data simulation. • A vector named <code>beta</code> of regression coefficients. • A scalar named <code>sigma.obs</code>.
<code>...</code>	Extra arguments passed to <code>SpikeSlabPrior</code>

Value

An object of class `bsts.mixed`, which is a list with the following elements. Many of these are arrays, in which case the first index of the array corresponds to the MCMC iteration number.

<code>coefficients</code>	A matrix containing the MCMC draws of the regression coefficients. Rows correspond to MCMC draws, and columns correspond to variables.
<code>sigma.obs</code>	The standard deviation of the weekly latent observations.
<code>state.contributions</code>	A three-dimensional array containing the MCMC draws of each state model's contributions to the state of the weekly model. The three dimensions are MCMC iteration, state model, and week number.
<code>weekly</code>	A matrix of MCMC draws of the weekly latent observations. Rows are MCMC iterations, and columns are weekly time points.
<code>cumulator</code>	A matrix of MCMC draws of the cumulator variable.

The returned object also contains MCMC draws for the parameters of the state models supplied as part of `state.specification`, relevant information passed to the function call, and other supplemental information.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts](#), [AddLocalLevel](#), [AddLocalLinearTrend](#), [AddSemilocalLinearTrend](#), [SpikeSlabPrior](#), [SdPrior](#).

Examples

```
## Not run:
data <- SimulateFakeMixedFrequencyData(nweeks = 104, xdim = 20)

## Setting an upper limit on the standard deviations can help keep the
## MCMC from flying off to infinity.
sd.limit <- sd(data$coarse.target)
state.specification <-
  AddLocalLinearTrend(list(),
    data$coarse.target,
    level.sigma.prior = SdPrior(1.0, 5, upper.limit = sd.limit),
    slope.sigma.prior = SdPrior(.5, 5, upper.limit = sd.limit))
weeks <- index(data$predictor)
months <- index(data$coarse.target)
which.month <- MatchWeekToMonth(weeks, months[1])
membership.fraction <- GetFractionOfDaysInInitialMonth(weeks)
contains.end <- WeekEndsMonth(weeks)

model <- bsts.mixed(target.series = data$coarse.target,
  predictors = data$predictors,
  membership.fraction = membership.fraction,
  contains.end = contains.end,
  which.coarse = which.month,
  state.specification = state.specification,
  niter = 500,
  expected.r2 = .999,
  prior.df = 1)

plot(model, "state")
plot(model, "components")

## End(Not run)
```

Description

The (integer) number of months between dates.

Usage

```
MonthDistance(dates, origin)
```

Arguments

dates A vector of class `Date` to be measured.
origin A scalar of class `Date`.

Value

Returns a numeric vector giving the integer number of months that have elapsed between `origin` and each element in `dates`. The daily component of each date is ignored, so two dates that are in the same month will have the same measured distance. Distances are signed, so months that occur before `origin` will have negative values.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
dates <- as.Date(c("2008-04-17",
                  "2008-05-01",
                  "2008-05-31",
                  "2008-06-01"))
origin <- as.Date("2008-05-15")
MonthDistance(dates, origin) == c(-1, 0, 0, 1)
```

named.holidays	<i>Holidays Recognized by Name</i>
----------------	------------------------------------

Description

A character vector listing the names of pre-specified holidays.

Usage

```
named.holidays
```

Value

```
"NewYearsDay" "SuperBowlSunday" "MartinLutherKingDay" "PresidentsDay" "ValentinesDay"
"SaintPatricksDay" "USDaylightSavingsTimeBegins" "USDaylightSavingsTimeEnds" "EasterSun-
day" "USMothersDay" "IndependenceDay" "LaborDay" "ColumbusDay" "Halloween" "Thanks-
giving" "MemorialDay" "VeteransDay" "Christmas"
```

new.home.sales	<i>New home sales and Google trends</i>
----------------	---

Description

The first column, HSN1FNSA is a time series of new home sales in the US, obtained from the FRED online data base. The series has been manually deseasonalized. The remaining columns contain search terms from Google trends (obtained from <http://trends.google.com/correlate>). These show the relative popularity of each search term among all search terms typed into Google. All series in this data set have been standardized by subtracting off their mean and dividing by their standard deviation.

Usage

```
data(new.home.sales)
```

Format

zoo time series

Source

FRED and trends.google.com

one.step.prediction.errors	<i>Prediction Errors</i>
----------------------------	--------------------------

Description

Computes the one-step-ahead prediction errors for a [bsts](#) model.

Usage

```
bsts.prediction.errors(bsts.object,  
                       cutpoints = NULL,  
                       burn = SuggestBurn(.1, bsts.object),  
                       standardize = FALSE)
```

Arguments

<code>bsts.object</code>	An object of class <code>bsts</code> .
<code>cutpoints</code>	An increasing sequence of integers between 1 and the number of time points in the training data for <code>bsts.object</code> , or NULL. If NULL then the in-sample one-step prediction errors from the <code>bsts</code> object will be extracted and returned. Otherwise the model will be re-fit with a separate MCMC run for each entry in <code>'cutpoints'</code> . Data up to each cutpoint will be included in the fit, and one-step prediction errors for data after the cutpoint will be computed.
<code>burn</code>	An integer giving the number of MCMC iterations to discard as burn-in. If <code>burn <= 0</code> then no burn-in sample will be discarded.
<code>standardize</code>	Logical. If TRUE then the prediction errors are divided by the square root of the one-step-ahead forecast variance. If FALSE the raw errors are returned.

Details

Returns the posterior distribution of the one-step-ahead prediction errors from the `bsts.object`. The errors are computed using the Kalman filter, and are of two types.

Purely in-sample errors are computed as a by-product of the Kalman filter as a result of fitting the model. These are stored in the `bsts.object` assuming the `save.prediction.errors` option is TRUE, which is the default (See [BstsOptions](#)). The in-sample errors are 'in-sample' in the sense that the parameter values used to run the Kalman filter are drawn from their posterior distribution given complete data. Conditional on the parameters in that MCMC iteration, each 'error' is the difference between the observed $y[t]$ and its expectation given data to $t-1$.

Purely out-of-sample errors can be computed by specifying the `'cutpoints'` argument. If cutpoints are supplied then a separate MCMC is run using just data up to the cutpoint. The Kalman filter is then run on the remaining data, again finding the difference between $y[t]$ and its expectation given data to $t-1$, but conditional on parameters estimated using data up to the cutpoint.

Value

A matrix of draws of the one-step-ahead prediction errors. Rows of the matrix correspond to MCMC draws. Columns correspond to time.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts](#), [AddLocalLevel](#), [AddLocalLinearTrend](#), [AddSemilocalLinearTrend](#), [SpikeSlabPrior](#), [SdPrior](#).

Examples

```

data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)

## Not run:
model <- bst(y, state.specification = ss, niter = 500)

## End(Not run)

errors <- bst.prediction.errors(model, burn = 100)
PlotDynamicDistribution(errors$in.sample)

## Compute out of sample prediction errors beyond times 80 and 120.
errors <- bst.prediction.errors(model, cutpoints = c(80, 120))
standardized.errors <- bst.prediction.errors(
  model, cutpoints = c(80, 120), standardize = TRUE)
plot(model, "prediction.errors", cutpoints = c(80, 120))
str(errors)      ## three matrices, with 400 (= 500 - 100) rows
                 ## and length(y) columns

```

plot.bsts

*Plotting functions for Bayesian structural time series***Description**

Functions to plot the results of a model fit using [bsts](#).

Usage

```

## S3 method for class 'bsts'
plot(x, y = c("state", "components", "residuals",
             "coefficients", "prediction.errors",
             "forecast.distribution",
             "predictors", "size", "dynamic", "seasonal", "monthly",
             "help"),
     ...)

PlotBstsCoefficients(bsts.object, burn = SuggestBurn(.1, bsts.object),
                    inclusion.threshold = 0, number.of.variables = NULL, ...)

PlotBstsComponents(bsts.object,
                  burn = SuggestBurn(.1, bsts.object),
                  time,
                  same.scale = TRUE,
                  layout = c("square", "horizontal", "vertical"),
                  style = c("dynamic", "boxplot"),

```

```
ylim = NULL,
components = 1:length(bsts.object$state.specification),
...)
```

```
PlotDynamicRegression(bsts.object,
  burn = SuggestBurn(.1, bsts.object),
  time = NULL,
  same.scale = FALSE,
  style = c("dynamic", "boxplot"),
  layout = c("square", "horizontal", "vertical"),
  ylim = NULL,
  zero.width = 2,
  zero.color = "green",
  ...)
```

```
PlotBstsState(bsts.object, burn = SuggestBurn(.1, bsts.object),
  time, show.actuals = TRUE,
  style = c("dynamic", "boxplot"),
  scale = c("linear", "mean"),
  ylim = NULL,
  ...)
```

```
PlotBstsResiduals(bsts.object, burn = SuggestBurn(.1, bsts.object),
  time, style = c("dynamic", "boxplot"), means =
  TRUE, ...)
```

```
PlotBstsPredictionErrors(bsts.object, cutpoints = NULL,
  burn = SuggestBurn(.1, bsts.object),
  style = c("dynamic", "boxplot"),
  xlab = "Time", ylab = "", main = "",
  ...)
```

```
PlotBstsForecastDistribution(bsts.object, cutpoints = NULL,
  burn = SuggestBurn(.1, bsts.object),
  style = c("dynamic", "boxplot"),
  xlab = "Time",
  ylab = "",
  main = "",
  show.actuals = TRUE,
  col.actuals = "blue",
  ...)
```

```
PlotBstsSize(bsts.object, burn = SuggestBurn(.1, bsts.object), style =
  c("histogram", "ts"), ...)
```

```
PlotSeasonalEffect(bsts.object, nseasons = 7, season.duration = 1,
  same.scale = TRUE, ylim = NULL, get.season.name = NULL,
  burn = SuggestBurn(.1, bsts.object), ...)
```

```
PlotMonthlyAnnualCycle(bsts.object, ylim = NULL, same.scale = TRUE,
                        burn = SuggestBurn(.1, bsts.object), ...)
```

Arguments

x	An object of class bsts .
bsts.object	An object of class bsts .
y	A character string indicating the aspect of the model that should be plotted.
burn	The number of MCMC iterations to discard as burn-in.
col.actuals	The color to use for the actual data when comparing actuals vs forecasts.
components	A numeric vector indicating which components to plot. Component indices correspond to elements of the state specification that was used to build the bsts model being plotted.
cutpoints	A numeric vector of integers, or NULL. For diagnostic plots of prediction errors or forecast distributions, the model will be re-fit with a separate MCMC run for each entry in 'cutpoints'. Data up to each cutpoint will be included in the fit, and one-step prediction errors for data after the cutpoint will be computed.
get.season.name	A function that can be used to infer the title of each seasonal plot. It should take a single POSIXt , Date , or similar object as an argument, and return a single string that can be used as a panel title. If <code>get.season.name</code> is NULL and <code>nseasons</code> is specified or inferred to be one of the following values, then the following functions will be used. <ul style="list-style-type: none"> • 4: quarters • 7: weekdays • 12: months
inclusion.threshold	An inclusion probability that individual coefficients must exceed in order to be displayed when <code>what == "coefficients"</code> . See the help file for plot.lm.spike .
layout	For controlling the layout of functions that generate multiple plots.
main	Main title for the plot.
means	Logical. If TRUE then the mean of each residual is plotted as a blue dot. If false only the distribution of the residuals is plotted.
nseasons	If there is only one seasonal component in the model, this argument is ignored. If there are multiple seasonal components then <code>nseasons</code> and <code>season.duration</code> are used to select the desired one.
number.of.variables	If non-NULL this specifies the number of coefficients to plot, taking precedence over <code>inclusion.threshold</code> . See plot.lm.spike .
same.scale	Logical. If TRUE then all the state components will be plotted with the same scale on the vertical axis. If FALSE then each component will get its own scale for the vertical axis.

scale	The scale on which to plot the state. If the error family is "logit" or "poisson" then the state can either be plotted on the scale of the linear predictor (e.g. trend + seasonal + regression) or the linear predictor can be passed through the link function so as to plot the distribution of the conditional mean.
season.duration	If there is only one seasonal component in the model, this argument is ignored. If there are multiple seasonal components then nseasons and season.duration are used to select the desired one.
show.actuals	Logical. If TRUE then actual values from the fitted series will be shown on the plot.
style	The desired plot style. Partial matching is allowed, so "dyn" would match "dynamic", for example.
time	An optional vector of values to plot against. If missing, the default is to diagnose the time scale of the original time series.
xlab	Label for the horizontal axis.
ylab	Label for the vertical axis.
ylim	Limits for the vertical axis. If NULL these will be inferred from the state components and the same.scale argument. Otherwise all plots will be created with the same ylim values.
zero.width	A numerical value for the width of the reference line at zero. If NULL then the line will be omitted.
zero.color	A color for the width of the reference line at zero. If NULL then the line will be omitted.
...	Additional arguments to be passed to PlotDynamicDistribution , or TimeSeriesBoxplot .

Details

[PlotBstsState](#), [PlotBstsComponents](#), and [PlotBstsResiduals](#) all produce dynamic distribution plots. [PlotBstsState](#) plots the aggregate state contribution (including regression effects) to the mean, while [PlotBstsComponents](#) plots the contribution of each state component. [PlotBstsResiduals](#) plots the posterior distribution of the residuals given complete data (i.e. looking forward and backward in time). [PlotBstsPredictionErrors](#) plots filtering errors (i.e. the one-step-ahead prediction errors given data up to the previous time point). [PlotBstsForecastDistribution](#) plots the one-step-ahead forecasts instead of the prediction errors.

[PlotBstsCoefficients](#) creates a significance plot for the predictors used in the state space regression model. It is obviously not useful for models with no regressors.

[PlotBstsSize](#) plots the distribution of the number of predictors included in the model.

[PlotSeasonalEffect](#) generates an array of plots showing how the distribution of the seasonal effect changes, for each season, for models that include a seasonal state component.

[PlotMonthlyAnnualCycle](#) produces an array of plots much like [PlotSeasonalEffect](#), for models that include a [MonthlyAnnualCycle](#) state component.

Value

These functions are called for their side effect, which is to produce a plot on the current graphics device.

PlotBstsState invisibly returns the state object being plotted.

See Also

[bsts](#) [PlotDynamicDistribution](#) [plot.lm.spike](#)

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 500)
plot(model, burn = 100)
plot(model, "residuals", burn = 100)
plot(model, "components", burn = 100)
plot(model, "forecast.distribution", burn = 100)
```

plot.bsts.mixed

Plotting functions for mixed frequency Bayesian structural time series

Description

Functions for plotting the output of a mixed frequency time series regression.

Usage

```
## S3 method for class 'bsts.mixed'
plot(x,
      y = c("state", "components",
            "coefficients", "predictors", "size"),
      ...)

PlotBstsMixedState(bsts.mixed.object,
                   burn = SuggestBurn(.1, bsts.mixed.object),
                   time = NULL,
                   fine.scale = FALSE,
                   style = c("dynamic", "boxplot"),
                   trim.left = NULL,
                   trim.right = NULL,
                   ...)

PlotBstsMixedComponents(bsts.mixed.object,
                        burn = SuggestBurn(.1, bsts.mixed.object),
```

```

time = NULL,
same.scale = TRUE,
fine.scale = FALSE,
style = c("dynamic", "boxplot"),
layout = c("square", "horizontal", "vertical"),
ylim = NULL,
trim.left = NULL,
trim.right = NULL,
...)

```

Arguments

x	An object of class <code>bsts.mixed</code> .
bsts.mixed.object	An object of class <code>bsts.mixed</code> .
y	A character string indicating the aspect of the model that should be plotted.
burn	The number of MCMC iterations to discard as burn-in.
time	An optional vector of values to plot against. If missing, the default is to obtain the time scale from the original time series.
fine.scale	Logical. If TRUE then the plots will be at the weekly level of granularity. If FALSE they will be at the monthly level.
same.scale	Logical. If TRUE then all the state components will be plotted with the same scale on the vertical axis. If FALSE then each component will get its own scale for the vertical axis.
style	character. If "dynamic" then a dynamic distribution plot will be shown. If "box" then boxplots will be shown.
layout	A character string indicating whether the plots showing components of state should be laid out in a square, horizontally, or vertically.
trim.left	A logical indicating whether the first (presumably partial) observation in the aggregated state time series should be removed.
trim.right	A logical indicating whether the last (presumably partial) observation in the aggregated state time series should be removed.
ylim	Limits for the vertical axis. Optional.
...	Additional arguments to be passed to <code>PlotDynamicDistribution</code> or <code>TimeSeriesBoxplot</code>

Details

`PlotBstsMixedState` plots the aggregate state contribution (including regression effects) to the mean, while `PlotBstsComponents` plots the contribution of each state component separately. `PlotBstsCoefficients` creates a significance plot for the predictors used in the state space regression model.

Value

These functions are called for their side effect, which is to produce a plot on the current graphics device.

See Also

[bsts.mixed](#) [PlotDynamicDistribution](#) [plot.lm.spike](#) [PlotBstsSize](#)

Examples

```
## Not run:
## This example is flaky and needs to be fixed
data <- SimulateFakeMixedFrequencyData(nweeks = 104, xdim = 20)
state.specification <- AddLocalLinearTrend(list(), data$coarse.target)
weeks <- index(data$predictor)
months <- index(data$coarse.target)
which.month <- MatchWeekToMonth(weeks, months[1])
membership.fraction <- GetFractionOfDaysInInitialMonth(weeks)
contains.end <- WeekEndsMonth(weeks)

model <- bsts.mixed(target.series = data$coarse.target,
                   predictors = data$predictors,
                   membership.fraction = membership.fraction,
                   contains.end = contains.end,
                   which.coarse = which.month,
                   state.specification = state.specification,
                   niter = 500)

plot(model, "state")
plot(model, "components")

## End(Not run)
```

plot.bsts.prediction *Plot predictions from Bayesian structural time series*

Description

Plot the posterior predictive distribution from a [bsts](#) prediction object.

Usage

```
## S3 method for class 'bsts.prediction'
plot(x,
     y = NULL,
     burn = 0,
     plot.original = TRUE,
     median.color = "blue",
     median.type = 1,
     median.width = 3,
     interval.quantiles = c(.025, .975),
     interval.color = "green",
     interval.type = 2,
```

```

interval.width = 2,
style = c("dynamic", "boxplot"),
ylim = NULL,
...)
```

Arguments

x	An object of class <code>bsts.prediction</code> created by calling <code>predict</code> on a <code>bsts</code> object.
y	A dummy argument necessary to match the signature of the <code>plot</code> generic function. This argument is unused.
plot.original	Logical or numeric. If TRUE then the prediction is plotted after a time series plot of the original series. If FALSE, the prediction fills the entire plot. If numeric, then it specifies the number of trailing observations of the original time series to plot in addition to the predictions.
burn	The number of observations you wish to discard as burn-in from the posterior predictive distribution. This is in addition to the burn-in discarded using <code>predict.bsts</code> .
median.color	The color to use for the posterior median of the prediction.
median.type	The type of line (lty) to use for the posterior median of the prediction.
median.width	The width of line (lwd) to use for the posterior median of the prediction.
interval.quantiles	The lower and upper limits of the credible interval to be plotted.
interval.color	The color to use for the upper and lower limits of the 95% credible interval for the prediction.
interval.type	The type of line (lty) to use for the upper and lower limits of the 95% credible interval for of the prediction.
interval.width	The width of line (lwd) to use for the upper and lower limits of the 95% credible interval for of the prediction.
style	Either "dynamic", for dynamic distribution plots, or "boxplot", for box plots. Partial matching is allowed, so "dyn" or "box" would work, for example.
ylim	Limits on the vertical axis.
...	Extra arguments to be passed to <code>PlotDynamicDistribution</code> and <code>lines</code> .

Details

Plots the posterior predictive distribution described by x using a dynamic distribution plot generated by `PlotDynamicDistribution`. Overlays the posterior median and 95% prediction limits for the predictive distribution.

Value

Returns NULL.

See Also

[bsts PlotDynamicDistribution](#) [plot.lm.spike](#)

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 500)
pred <- predict(model, horizon = 12, burn = 100)
plot(pred)
```

plot.bsts.predictors *Plot the most likely predictors*

Description

Creates a time series plot showing the most likely predictors of a time series used to fit a [bsts](#) object.

Usage

```
PlotBstsPredictors(bsts.object,
                   burn = SuggestBurn(.1, bsts.object),
                   inclusion.threshold = .1,
                   ylim = NULL,
                   flip.signs = TRUE,
                   show.legend = TRUE,
                   grayscale = TRUE,
                   short.names = TRUE,
                   ...)
```

Arguments

bsts.object	An object of class bsts .
burn	The number of observations you wish to discard as burn-in.
inclusion.threshold	Plot predictors with marginal inclusion probabilities above this threshold.
ylim	Scale for the vertical axis.
flip.signs	If true then a predictor with a negative sign will be flipped before being plotted, to better align visually with the target series.
show.legend	Should a legend be shown indicating which predictors are plotted?
grayscale	Logical. If TRUE then lines for different predictors grow progressively lighter as their inclusion probability decreases. If FALSE then lines are drawn in black.
short.names	Logical. If TRUE then a common prefix or suffix shared by all the variables will be discarded.
...	Extra arguments to be passed to plot .

See Also

[bsts PlotDynamicDistribution plot.lm.spike](#)

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
lag.y <- c(NA, head(y, -1))
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
## Call bsts with na.action = na.omit to omit the leading NA in lag.y
model <- bsts(y ~ lag.y, state.specification = ss, niter = 500,
             na.action = na.omit)
plot(model, "predictors")
```

plot.holiday

Plot Holiday Effects

Description

Plot the estimated effect of the given holiday.

Usage

```
PlotHoliday(holiday, model, show.raw.data = TRUE, ylim = NULL, ...)
```

Arguments

holiday	An object of class Holiday .
model	A model fit by bsts containing either a RegressionHolidayStateModel or HierarchicalRegressionHolidayStateModel that includes holiday.
show.raw.data	Logical indicating if the raw data corresponding to holiday should be superimposed on the plot. The 'raw data' are the actual values of the target series, minus the value of the target series the day before the holiday began, which is a (somewhat poor) proxy for remaining state elements. The raw data can appear artificially noisy if there are other strong state effects such as a day-of-week effect for holidays that don't always occur on the same day of the week.
ylim	Limits on the vertical axis of the plots.
...	Extra arguments passed to boxplot .

Value

Returns `invisible{NULL}`.

See Also

[bsts AddRandomWalkHoliday](#)

Examples

```

trend <- cumsum(rnorm(730, 0, .1))
dates <- seq.Date(from = as.Date("2014-01-01"), length = length(trend),
  by = "day")
y <- zoo(trend + rnorm(length(trend), 0, .2), dates)

AddHolidayEffect <- function(y, dates, effect) {
  ## Adds a holiday effect to simulated data.
  ## Args:
  ##   y: A zoo time series, with Dates for indices.
  ##   dates: The dates of the holidays.
  ##   effect: A vector of holiday effects of odd length. The central effect is
  ##           the main holiday, with a symmetric influence window on either side.
  ## Returns:
  ##   y, with the holiday effects added.
  time <- dates - (length(effect) - 1) / 2
  for (i in 1:length(effect)) {
    y[time] <- y[time] + effect[i]
    time <- time + 1
  }
  return(y)
}

## Define some holidays.
memorial.day <- NamedHoliday("MemorialDay")
memorial.day.effect <- c(.3, 3, .5)
memorial.day.dates <- as.Date(c("2014-05-26", "2015-05-25"))
y <- AddHolidayEffect(y, memorial.day.dates, memorial.day.effect)

presidents.day <- NamedHoliday("PresidentsDay")
presidents.day.effect <- c(.5, 2, .25)
presidents.day.dates <- as.Date(c("2014-02-17", "2015-02-16"))
y <- AddHolidayEffect(y, presidents.day.dates, presidents.day.effect)

labor.day <- NamedHoliday("LaborDay")
labor.day.effect <- c(1, 2, 1)
labor.day.dates <- as.Date(c("2014-09-01", "2015-09-07"))
y <- AddHolidayEffect(y, labor.day.dates, labor.day.effect)

## The holidays can be in any order.
holiday.list <- list(memorial.day, labor.day, presidents.day)
number.of.holidays <- length(holiday.list)

## In a real example you'd want more than 100 MCMC iterations.
niter <- 100
ss <- AddLocalLevel(list(), y)
ss <- AddRegressionHoliday(ss, y, holiday.list = holiday.list)
model <- bst(y, state.specification = ss, niter = niter)

PlotHoliday(memorial.day, model)

```

Description

Functions to plot the results of a model fit using [mbsts](#).

Usage

```
## S3 method for class 'mbsts'
plot(x, y = c("means", "help"), ...)

PlotMbstsSeriesMeans(mbsts.object,
                     series.id = NULL,
                     same.scale = TRUE,
                     burn = SuggestBurn(.1, mbsts.object),
                     time,
                     show.actuals = TRUE,
                     ylim = NULL,
                     gap = 0,
                     cex.actuals = 0.2,
                     ...)
```

Arguments

x	An object of class mbsts .
y	A character string indicating the aspect of the model that should be plotted.
mbsts.object	An object of class mbsts .
series.id	Indicates which series should be plotted. An integer, logical, or character vector.
same.scale	Logical. If TRUE then all the series or state components will be plotted with the same scale on the vertical axis. If FALSE then each plot will get its own scale for the vertical axis.
burn	The number of MCMC iterations to discard as burn-in.
time	An optional vector of values to plot against. If missing, the default is to diagnose the time scale of the original time series.
show.actuals	Logical. If TRUE then actual values from the fitted series will be shown on the plot.
ylim	Limits for the vertical axis. If NULL these will be inferred from the state components and the same.scale argument. Otherwise all plots will be created with the same ylim values.
gap	The number of lines to leave between plots. This need not be an integer.
cex.actuals	Scale factor to use for plotting the raw data.
...	Additional arguments passed to PlotDynamicDistribution .

See Also[plot.bsts](#)

 plot.mbsts.prediction *Plot Multivariate Bsts Predictions*

Description

Plot the posterior predictive distribution from an [mbsts](#) prediction object.

Usage

```
## S3 method for class 'mbsts.prediction'
plot(x,
     y = NULL,
     burn = 0,
     plot.original = TRUE,
     median.color = "blue",
     median.type = 1,
     median.width = 3,
     interval.quantiles = c(.025, .975),
     interval.color = "green",
     interval.type = 2,
     interval.width = 2,
     style = c("dynamic", "boxplot"),
     ylim = NULL,
     series.id = NULL,
     same.scale = TRUE,
     gap = 0,
     ...)
```

Arguments

- | | |
|---------------|---|
| x | An object of class bsts.prediction created by calling <code>predict</code> on a bsts object. |
| y | A dummy argument necessary to match the signature of the <code>plot</code> generic function. This argument is unused. |
| plot.original | Logical or numeric. If TRUE then the prediction is plotted after a time series plot of the original series. If FALSE, the prediction fills the entire plot. If numeric, then it specifies the number of trailing observations of the original time series to plot in addition to the predictions. |
| burn | The number of observations you wish to discard as burn-in from the posterior predictive distribution. This is in addition to the burn-in discarded using predict.bsts . |
| median.color | The color to use for the posterior median of the prediction. |

median.type	The type of line (lty) to use for the posterior median of the prediction.
median.width	The width of line (lwd) to use for the posterior median of the prediction.
interval.quantiles	The lower and upper limits of the credible interval to be plotted.
interval.color	The color to use for the upper and lower limits of the 95% credible interval for the prediction.
interval.type	The type of line (lty) to use for the upper and lower limits of the 95% credible interval for of the prediction.
interval.width	The width of line (lwd) to use for the upper and lower limits of the 95% credible interval for of the prediction.
style	Either "dynamic", for dynamic distribution plots, or "boxplot", for box plots. Partial matching is allowed, so "dyn" or "box" would work, for example.
ylim	Limits on the vertical axis.
series.id	A factor, string, or integer used to indicate which of the multivariate series to plot. If NULL then predictions for all series will be plotted. If there are many series this can make the plot unreadable.
same.scale	Logical. If TRUE then all predictions are plotted with the same scale, and limits are drawn on the Y axis. If FALSE then each prediction is drawn to fill its plot region, and no tick marks are drawn on the y axis. If ylim is specified then it is used for all plots, and same.scale is ignored.
gap	The amount of space to leave between plots, measured in lines of text.
...	Extra arguments to be passed to PlotDynamicDistribution and lines .

Details

Plots the posterior predictive distribution described by `x` using a dynamic distribution plot generated by [PlotDynamicDistribution](#). Overlays the posterior median and 95% prediction limits for the predictive distribution.

Value

Returns NULL.

predict.bsts

Prediction for Bayesian Structural Time Series

Description

Generate draws from the posterior predictive distribution of a [bsts](#) object.

Usage

```
## S3 method for class 'bsts'
predict(object,
        horizon = 1,
        newdata = NULL,
        timestamps = NULL,
        burn = SuggestBurn(.1, object),
        na.action = na.exclude,
        olddata = NULL,
        olddata.timestamps = NULL,
        trials.or.exposure = 1,
        quantiles = c(.025, .975),
        seed = NULL,
        ...)
```

Arguments

object	An object of class <code>bsts</code> created by a call to the function <code>bsts</code> .
horizon	An integer specifying the number of periods into the future you wish to predict. If object contains a regression component then the forecast horizon is <code>nrow(X)</code> , and this argument is not used.
newdata	a vector, matrix, or data frame containing the predictor variables to use in making the prediction. This is only required if object contains a regression component. If a data frame, it must include variables with the same names as the data used to fit object. The first observation in newdata is assumed to be one time unit after the end of the last observation used in fitting object, and the subsequent observations are sequential time points. If the regression part of object contains only a single predictor then newdata can be a vector. If newdata is passed as a matrix it is the caller's responsibility to ensure that it contains the correct number of columns and that the columns correspond to those in <code>object\$coefficients</code> .
timestamps	A vector of time stamps (of the same type as the timestamps used to fit object), with one per row of newdata (or element of newdata, if newdata is a vector). The time stamps give the time points as which each prediction is desired. They must be interpretable as integer (0 or larger) time steps following the last time stamp in object. If NULL, then the requested predictions are interpreted as being at 1, 2, 3, ... steps following the training data.
burn	An integer describing the number of MCMC iterations in object to be discarded as burn-in. If <code>burn <= 0</code> then no burn-in period will be discarded.
na.action	A function determining what should be done with missing values in newdata.
olddata	This is an optional component allowing predictions to be made conditional on data other than the data used to fit the model. If omitted, then it is assumed that forecasts are to be made relative to the final observation in the training data. If olddata is supplied then it will be filtered to get the distribution of the next state before a prediction is made, and it is assumed that the first entry in newdata comes immediately after the last entry in olddata.

The value for `olddata` depends on whether or not `object` contains a regression component.

- If a regression component is present, then `olddata` is a `data.frame` including variables with the same names as the data used to fit `object`, including the response.
- If no regression component is present, then `olddata` is a vector containing historical values of a time series.

`olddata.timestamps`

A set of timestamps corresponding to the observations supplied in `olddata`. If `olddata` is `NULL` then this argument is not used. If `olddata` is supplied and this is `NULL` then trivial timestamps (1, 2, ...) will be assumed. Otherwise this argument behaves like the `timestamps` argument to the `bsts` function.

`trials.or.exposure`

For logit or Poisson models, the number of binomial trials (or the exposure time) to assume at each time point in the forecast period. This can either be a scalar (if the number of trials is to be the same for each time period), or it can be a vector with length equal to `horizon` (if the model contains no regression term) or `nrow(newdata)` if the model contains a regression term.

`quantiles`

A numeric vector of length 2 giving the lower and upper quantiles to use for the forecast interval estimate.

`seed`

An integer to use as the C++ random seed. If `NULL` then the C++ seed will be set using the clock.

...

This is a dummy argument included to match the signature of the generic `predict` function. It is not used.

Details

Samples from the posterior distribution of a Bayesian structural time series model. This function can be used either with or without contemporaneous predictor variables (in a time series regression).

If predictor variables are present, the regression coefficients are fixed (as opposed to time varying, though time varying coefficients might be added as state component). The predictors and response in the formula are contemporaneous, so if you want lags and differences you need to put them in the predictor matrix yourself.

If no predictor variables are used, then the model is an ordinary state space time series model.

Value

Returns an object of class `bsts.prediction`, which is a list with the following components.

`mean`

A vector giving the posterior mean of the prediction.

`interval`

A two (column/row?) matrix giving the upper and lower bounds of the 95 percent credible interval for the prediction.

`distribution`

A matrix of draws from the posterior predictive distribution. Each row in the matrix is one MCMC draw. Columns represent time.

Author(s)

Steven L. Scott

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts](#). [AddLocalLevel](#). [AddLocalLinearTrend](#). [AddSemilocalLinearTrend](#).

Examples

The number of MCMC draws in the following examples is artificially low.

```
## Making predictions when there is no regression component.
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 250)
pred <- predict(model, horizon = 12, burn = 100)
plot(pred)

## An example using the olddata argument.
full.pred <- pred
training <- window(y, end = c(1959, 12))
model <- bst(training, state.specification = ss, niter = 250)
## Predict the next 12 months.
pred <- predict(model, horizon = 12)
## Compare the predictions to the actual data.
plot(pred)
lines(as.numeric(y, col = "red", lty = 2, lwd = 2))

## Predict the 12 months of 1961 based on the posterior distribution
## of the model fit to data through 1959, but with state filtered
## through 1960.
updated.pred <- predict(model, horizon = 12, olddata = y)
par(mfrow = c(1, 2))
plot(full.pred, ylim = c(4, 7))
plot(updated.pred, ylim = c(4, 7))

## Examples including a regression component.
##
data(iclaims)
training <- initial.claims[1:402, ]
holdout1 <- initial.claims[403:450, ]
```

```

holdout2 <- initial.claims[451:456, ]

## Not run:

## This example puts the total run time over 5 seconds, which is a CRAN
## violation.

ss <- AddLocalLinearTrend(list(), training$iclaimsNSA)
ss <- AddSeasonal(ss, training$iclaimsNSA, nseasons = 52)
## In real life you'd want more iterations...
model <- bsts(iclaimsNSA ~ ., state.specification = ss, data =
              training, niter = 100)

## Predict the holdout set given the training set.
## This is really fast, because we can use saved state from the MCMC
## algorithm.
pred.full <- predict(model, newdata = rbind(holdout1, holdout2))

## Predict holdout 2, given training and holdout1.
## This is much slower because we need to re-filter the 'olddata' before
## simulating the predictions.
pred.update <- predict(model, newdata = holdout2,
                       olddata = rbind(training, holdout1))

## End(Not run)

```

predict.mbsts

Prediction for Multivariate Bayesian Structural Time Series

Description

Generate draws from the posterior predictive distribution of an `mbsts` object.

Usage

```

## S3 method for class 'mbsts'
predict(object,
        horizon = 1,
        newdata = NULL,
        timestamps = NULL,
        burn = SuggestBurn(.1, object),
        na.action = na.exclude,
        quantiles = c(.025, .975),
        seed = NULL,
        ...)

```

Arguments

`object` An object of class `mbsts`.

horizon	An integer specifying the number of periods into the future you wish to predict. If object contains a regression component then the forecast horizon is nrow(newdata) and this argument is not used.
newdata	A vector, matrix, or data frame containing the predictor variables to use in making the prediction. This is only required if object contains a regression component. If a data frame, it must include variables with the same names as the data used to fit object. The first observation in newdata is assumed to be one time unit after the end of the last data used in fitting object, and the subsequent observations are sequential time points. If the regression part of object contains only a single predictor then newdata can be a vector. If newdata is passed as a matrix it is the caller's responsibility to ensure that it contains the correct number of columns and that the columns correspond to those in object\$coefficients.
timestamps	A vector of time stamps (of the same type as the timestamps used to fit object), with one per row of newdata (or element of newdata, if newdata is a vector). The time stamps give the time points as which each prediction is desired. They must be interpretable as integer (0 or larger) time steps following the last time stamp in object. If NULL, then the requested predictions are interpreted as being at 1, 2, 3, ... steps following the training data.
burn	An integer describing the number of MCMC iterations in object to be discarded as burn-in. If burn <= 0 then no burn-in period will be discarded.
na.action	A function determining what should be done with missing values in newdata.
quantiles	A numeric vector of length 2 giving the lower and upper quantiles to use for the forecast interval estimate.
seed	An integer to use as the C++ random seed. If NULL then the C++ seed will be set using the clock.
...	Not used. Present to match the signature of the default predict method.

Details

The prediction is based off of samples taken from the posterior distribution of a multivariate Bayesian structural time series model.

As an added convenience, means and interval estimates are produced from the posterior predictive distribution.

Value

Returns an object of class mbsts.prediction, which is a list.

Author(s)

Steven L. Scott

See Also

[mbsts.predict](#), [bsts.plot](#), [mbsts.prediction](#)

quarter	<i>Find the quarter in which a date occurs</i>
---------	--

Description

Returns the quarter and year in which a date occurs.

Usage

```
Quarter(date)
```

Arguments

date A vector convertible to `POSIXlt`. A `Date` or character is fine.

Value

A numeric vector identifying the quarter that each element of date corresponds to, expressed as a number of years since 1900. Thus Q1-2000 is 100.00, and Q3-2007 is 107.50.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
Quarter(c("2008-02-29", "2008-04-29"))
# [1] 108.00 108.25
```

regression.holiday	<i>Regression Based Holiday Models</i>
--------------------	--

Description

Add a regression-based holiday model to the state specification.

Usage

```
AddRegressionHoliday(
  state.specification = NULL,
  y,
  holiday.list,
  time0 = NULL,
  prior = NULL,
  sdy = sd(as.numeric(y), na.rm = TRUE))
```

```
AddHierarchicalRegressionHoliday(
  state.specification = NULL,
  y,
  holiday.list,
  coefficient.mean.prior = NULL,
  coefficient.variance.prior = NULL,
  time0 = NULL,
  sdy = sd(as.numeric(y), na.rm = TRUE))
```

Arguments

<code>state.specification</code>	A list of state components that you wish to add to. If omitted, an empty list will be assumed.
<code>holiday.list</code>	A list of objects of type <code>Holiday</code> . The width of the influence window should be the same number of days for all the holidays in this list. If the data contains many instances of holidays with different window widths, then multiple instances <code>HierarchicalRegressionHolidayModel</code> can be used as long as all holidays in the same state component model have the same sized window width.
<code>y</code>	The time series to be modeled, as a numeric vector convertible to <code>xts</code> . This state model assumes <code>y</code> contains daily data.
<code>prior</code>	An object of class <code>NormalPrior</code> describing the expected variation among daily holiday effects.
<code>coefficient.mean.prior</code>	An object of type <code>MvnPrior</code> giving the hyperprior for the average effect of a holiday in each day of the influence window.
<code>coefficient.variance.prior</code>	An object of type <code>InverseWishartPrior</code> describing the prior belief about the variation in holiday effects from one holiday to the next.
<code>time0</code>	An object convertible to <code>Date</code> containing the date of the initial observation in the training data. If omitted and <code>y</code> is a <code>zoo</code> or <code>xts</code> object, then <code>time0</code> will be obtained from the index of <code>y[1]</code> .
<code>sdy</code>	The standard deviation of the series to be modeled. This will be ignored if <code>y</code> is provided, or if all the required prior distributions are supplied directly.

Details

The model assumes that

$$y_t = \beta_{d(t)} + \epsilon_t$$

The regression state model assumes vector of regression coefficients β contains elements $\beta_d \sim N(0, \sigma)$.

The `HierarchicalRegressionHolidayModel` assumes β is composed of holiday-specific sub-vectors $\beta_h \sim N(b_0, V)$, where each β_h contains coefficients describing the days in the influence window of holiday `h`. The hierarchical version of the model treats b_0 and V as parameters to be learned, with prior distributions

$$b_0 \sim N(\bar{b}, \Omega)$$

and

$$V \sim IW(\nu, S)$$

where IW represents the inverse Wishart distribution.

Value

Returns a list with the elements necessary to specify a local linear trend state model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.RandomWalkHolidayStateModel](#). [SdPrior](#) [NormalPrior](#)

Examples

```
trend <- cumsum(rnorm(730, 0, .1))
dates <- seq.Date(from = as.Date("2014-01-01"), length = length(trend), by = "day")
y <- zoo(trend + rnorm(length(trend), 0, .2), dates)

AddHolidayEffect <- function(y, dates, effect) {
  ## Adds a holiday effect to simulated data.
  ## Args:
  ##   y: A zoo time series, with Dates for indices.
  ##   dates: The dates of the holidays.
  ##   effect: A vector of holiday effects of odd length. The central effect is
  ##           the main holiday, with a symmetric influence window on either side.
  ## Returns:
  ##   y, with the holiday effects added.
  time <- dates - (length(effect) - 1) / 2
  for (i in 1:length(effect)) {
    y[time] <- y[time] + effect[i]
    time <- time + 1
  }
  return(y)
}

## Define some holidays.
memorial.day <- NamedHoliday("MemorialDay")
memorial.day.effect <- c(.3, 3, .5)
memorial.day.dates <- as.Date(c("2014-05-26", "2015-05-25"))
y <- AddHolidayEffect(y, memorial.day.dates, memorial.day.effect)
```

```

presidents.day <- NamedHoliday("PresidentsDay")
presidents.day.effect <- c(.5, 2, .25)
presidents.day.dates <- as.Date(c("2014-02-17", "2015-02-16"))
y <- AddHolidayEffect(y, presidents.day.dates, presidents.day.effect)

labor.day <- NamedHoliday("LaborDay")
labor.day.effect <- c(1, 2, 1)
labor.day.dates <- as.Date(c("2014-09-01", "2015-09-07"))
y <- AddHolidayEffect(y, labor.day.dates, labor.day.effect)

## The holidays can be in any order.
holiday.list <- list(memorial.day, labor.day, presidents.day)

## In a real example you'd want more than 100 MCMC iterations.
niter <- 100

## Fit the model
ss <- AddLocalLevel(list(), y)
ss <- AddRegressionHoliday(ss, y, holiday.list = holiday.list)
model <- bst(y, state.specification = ss, niter = niter)

## Plot all model state components.
plot(model, "comp")

## Plot the specific holiday state component.
plot(ss[[2]], model)

## Try again with some shrinkage. With only 3 holidays there won't be much
## shrinkage.
ss2 <- AddLocalLevel(list(), y)

## Plot the specific holiday state component.
ss2 <- AddHierarchicalRegressionHoliday(ss2, y, holiday.list = holiday.list)
model2 <- bst(y, state.specification = ss2, niter = niter)

plot(model2, "comp")
plot(ss2[[2]], model2)

```

regularize.timestamps *Produce a Regular Series of Time Stamps*

Description

Given an set of timestamps that might contain duplicates and gaps, produce a set of timestamps that has no duplicates and no gaps.

Usage

```
RegularizeTimestamps(timestamps)
```

```

## Default S3 method:
RegularizeTimestamps(timestamps)

## S3 method for class 'numeric'
RegularizeTimestamps(timestamps)

## S3 method for class 'Date'
RegularizeTimestamps(timestamps)

## S3 method for class 'POSIXt'
RegularizeTimestamps(timestamps)

```

Arguments

`timestamps` A set of (possibly irregular or non-unique) timestamps. This could be a set of integers (like 1, 2, , 3...), a set of numeric like (1945, 1945.083, 1945.167, ...) indicating years and fractions of years, a [Date](#) object, or a [POSIXt](#) object.

If the argument is NULL a NULL will be returned.

Value

A set of regularly spaced timestamps of the same class as the argument (which might be NULL).

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```

first <- as.POSIXct("2015-04-19 08:00:04")
monthly <- seq(from = first, length.out = 24, by = "month")
skip.one <- monthly[-8]
has.duplicates <- monthly
has.duplicates[2] <- has.duplicates[3]

reg1 <- RegularizeTimestamps(skip.one)
all.equal(reg1, monthly) ## TRUE

reg2 <- RegularizeTimestamps(has.duplicates)
all.equal(reg2, monthly) ## TRUE

```

residuals.bsts

Residuals from a bsts Object

Description

Residuals (or posterior distribution of residuals) from a bsts object.

Usage

```
## S3 method for class 'bsts'
residuals(object,
  burn = SuggestBurn(.1, object),
  mean.only = FALSE,
  ...)
```

Arguments

<code>object</code>	An object of class <code>bsts</code> created by the function of the same name.
<code>burn</code>	The number of MCMC iterations to discard as burn-in.
<code>mean.only</code>	Logical. If TRUE then the mean residual for each time period is returned. If FALSE then the full posterior distribution is returned.
<code>...</code>	Not used. This argument is here to comply with the signature of the generic <code>residuals</code> function.

Value

If `mean.only` is TRUE then this function returns a vector of residuals with the same "time stamp" as the original series. If `mean.only` is FALSE then the posterior distribution of the residuals is returned instead, as a matrix of draws. Each row of the matrix is an MCMC draw, and each column is a time point. The colnames of the returned matrix will be the timestamps of the original series, as text.

See Also

[bsts](#), [plot.bsts](#).

rsxfs

Retail sales, excluding food services

Description

A monthly time series of retail sales in the US, excluding food services. In millions of dollars. Seasonally adjusted.

Usage

```
data(rsxfs)
```

Format

zoo time series

Source

FRED. See <http://research.stlouisfed.org/fred2/series/RXFS>

Examples

```
data(rsxfs)
plot(rsxfs)
```

shark	<i>Shark Attacks in Florida.</i>
-------	----------------------------------

Description

An annual time series of shark attacks and fatalities in Florida.

Usage

```
data(shark)
```

Format

zoo time series

Source

From Jeffrey Simonoff "Analysis of Categorical Data". http://people.stern.nyu.edu/jsimonof/AnalCatData/Data/Comma_sep

Examples

```
data(shark)
head(shark)
```

shorten	<i>Shorten long names</i>
---------	---------------------------

Description

Removes common prefixes and suffixes from character vectors.

Usage

```
Shorten(words)
```

Arguments

words A character vector to be shortened.

Value

The argument words is returned, after common prefixes and suffixes have been removed. If all arguments are identical then no shortening is done.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[bsts.mixed.](#)

Examples

```
Shorten(c("/usr/common/foo.tex", "/usr/common/barbarian.tex"))  
# returns c("foo", "barbarian")
```

```
Shorten(c("hello", "hellobye"))  
# returns c("", "bye")
```

```
Shorten(c("hello", "hello"))  
# returns c("hello", "hello")
```

```
Shorten(c("", "x", "xx"))  
# returns c("", "x", "xx")
```

```
Shorten("abcde")  
# returns "abcde"
```

simulate.fake.mixed.frequency.data

Simulate fake mixed frequency data

Description

Simulate a fake data set that can be used to test mixed frequency code.

Usage

```
SimulateFakeMixedFrequencyData(nweeks,  
                                xdim,  
                                number.nonzero = xdim,  
                                start.date = as.Date("2009-01-03"),  
                                sigma.obs = 1.0,  
                                sigma.slope = .5,  
                                sigma.level = .5,  
                                beta.sd = 10)
```

Arguments

nweeks	The number of weeks of data to simulate.
xdim	The dimension of the predictor variables to be simulated.
number.nonzero	The number nonzero coefficients. Must be less than or equal to xdim.
start.date	The date of the first simulated week.
sigma.obs	The residual standard deviation for the fine time scale model.
sigma.slope	The standard deviation of the slope component of the local linear trend model for the fine time scale data.
sigma.level	The standard deviation of the level component fo the local linear trend model for the fine time scale data.
beta.sd	The standard deviation of the regression coefficients to be simulated.

Details

The simulation begins by simulating a local linear trend model for nweeks to get the trend component.

Next a nweeks by xdim matrix of predictor variables is simulated as IID normal(0, 1) deviates, and a xdim-vector of regression coefficients is simulated as IID normal(0, beta.sd). The product of the predictor matrix and regression coefficients is added to the output of the local linear trend model to get fine.target.

Finally, fine.target is aggregated to the month level to get coarse.target.

Value

Returns a list with the following components

coarse.target	A zoo time series containing the monthly values to be modeled.
fine.target	A zoo time series containing the weekly observations that aggregate to coarse.target.
predictors	A zoo matrix corresponding to fine.target containing the set of predictors variables to use in bstsmixed prediction.
true.beta	The vector of "true" regression coefficients used to simulate fine.target.
true.sigma.obs	The residual standard deviation that was used to simulate fine.target.
true.sigma.slope	The value of sigma.slope used to simulate fine.target.
true.sigma.level	The value of sigma.level use to simulate fine.target.
true.trend	The combined contribution of the simulated latent state on fine.target, including regression effects.
true.state	A matrix containin the fine-scale state of the model being simulated. Columns represent time (weeks). Rows correspond to regression (a constant 1), the local linear trend level, the local linear trend slope, the values of fine.target, and the weekly partial aggregates of coarse.target.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts.mixed](#), [AddLocalLinearTrend](#),

Examples

```
fake.data <- SimulateFakeMixedFrequencyData(nweeks = 100, xdim = 10)
plot(fake.data$coarse.target)
```

spike.slab.ar.prior *Spike and Slab Priors for AR Processes*

Description

Returns a spike and slab prior for the parameters of an AR(p) process.

Usage

```
SpikeSlabArPrior(
  lags,
  prior.inclusion.proBABILITIES =
    GeometricSequence( lags, initial.value = .8, discount.factor = .8),
  prior.mean = rep(0, lags),
  prior.sd =
    GeometricSequence(lags, initial.value = .5, discount.factor = .8),
  sdy,
  prior.df = 1,
  expected.r2 = .5,
  sigma.upper.limit = Inf,
  truncate = TRUE)
```

Arguments

lags A positive integer giving the maximum number of lags to consider.

prior.inclusion.proBABILITIES
A vector of length lags giving the prior probability that the corresponding AR coefficient is nonzero.

prior.mean	A vector of length lags giving the prior mean of the AR coefficients. This should almost surely stay set at zero.
prior.sd	A vector of length lags giving the prior standard deviations of the AR coefficients, which are modeled as a-priori independent of one another.
sd	The sample standard deviation of the series being modeled.
expected.r2	The expected fraction of variation in the response explained by this AR proces.
prior.df	A positive number indicating the number of observations (time points) worth of weight to assign to the guess at expected.r2.
sigma.upper.limit	A positive number less than infinity truncates the support of the prior distribution to regions where the residual standard deviation is less than the specified limit. Any other value indicates support over the entire positive real line.
truncate	If TRUE then the support of the distribution is truncated to the region where the AR coefficients imply a stationary process. If FALSE the coefficients are unconstrained.

Value

A list of class SpikeSlabArPrior containing the information needed for the underlying C++ code to instantiate this prior.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

state.sizes

Compute state dimensions

Description

Returns a vector containing the size of each state component (i.e. the state dimension) in the state vector.

Usage

```
StateSizes(state.specification)
```

Arguments

state.specification

A list containing state specification components, such as would be passed to [bsts](#).

Value

A numeric vector giving the dimension of each state component.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
y <- rnorm(1000)
state.specification <- AddLocalLinearTrend(list(), y)
state.specification <- AddSeasonal(state.specification, y, 7)
StateSizes(state.specification)
```

StateSpecification *Add a state component to a Bayesian structural time series model*

Description

Add a state component to the `state.specification` argument in a [bsts](#) model.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

References

Harvey (1990), "Forecasting, structural time series, and the Kalman filter", Cambridge University Press.

Durbin and Koopman (2001), "Time series analysis by state space methods", Oxford University Press.

See Also

[bsts](#), [SdPrior](#), [NormalPrior](#), [Ar1CoefficientPrior](#)

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 500)
pred <- predict(model, horizon = 12, burn = 100)
plot(pred)
```

SuggestBurn	<i>Suggested burn-in size</i>
-------------	-------------------------------

Description

Suggest the size of an MCMC burn in sample as a proportion of the total run.

Usage

```
SuggestBurn(proportion, bsts.object)
```

Arguments

proportion	The proportion of the MCMC run to discard as burn in.
bsts.object	An object of class bsts .

Value

An integer number of iterations to discard.

See Also

[bsts](#)

summary.bsts	<i>Summarize a Bayesian structural time series object</i>
--------------	---

Description

Print a summary of a [bsts](#) object.

Usage

```
## S3 method for class 'bsts'
summary(object, burn = SuggestBurn(.1, object), ...)
```

Arguments

object	An object of class bsts created by the function of the same name.
burn	The number of MCMC iterations to discard as burn-in.
...	Additional arguments passed to summary.lm.spike if object has a regression component.

Value

Returns a list with the following elements.

residual.sd	The posterior mean of the residual standard deviation parameter.
prediction.sd	The standard deviation of the one-step-ahead prediction errors for the training data.
rsquare	Proportion by which the residual variance is less than the variance of the original observations.
relative.gof	Harvey's goodness of fit statistic. Let ν denote the one step ahead prediction errors, n denote the length of the series, and y denote the original series. The goodness of fit statistic is

$$1 - \frac{\sum_{i=1}^n \nu_i^2}{\sum_{i=2}^n n(\Delta y_i - \Delta \bar{y})^2}.$$

This statistic is analogous to R^2 in a regression model, but the reduction in sum of squared errors is relative to a random walk with a constant drift,

$$y_{t+1} = y_t + \beta + \epsilon_t,$$

which Harvey (1989, equation 5.5.14) argues is a more relevant baseline than a simple mean. Unlike a traditional R-square statistic, this can be negative.

size	Distribution of the number of nonzero coefficients appearing in the model
coefficients	If object contains a regression component then the output contains matrix with rows corresponding to coefficients, and columns corresponding to: <ul style="list-style-type: none"> • The posterior probability the variable is included. • The posterior probability that the variable is positive. • The conditional expectation of the coefficient, given inclusion. • The conditional standard deviation of the coefficient, given inclusion.

References

Harvey's goodness of fit statistic is from Harvey (1989) *Forecasting, structural time series models, and the Kalman filter*. Page 268.

See Also

[bsts](#), [plot.bsts](#), [summary.lm.spike](#)

Examples

```
data(AirPassengers)
y <- log(AirPassengers)
ss <- AddLocalLinearTrend(list(), y)
ss <- AddSeasonal(ss, y, nseasons = 12)
model <- bst(y, state.specification = ss, niter = 100)
summary(model, burn = 20)
```

`to.posixt`*Convert to POSIXt*

Description

Convert an object of class `Date` to class `POSIXct` without getting bogged down in timezone calculation.

Usage

```
DateToPOSIX(timestamps)
YearMonToPOSIX(timestamps)
```

Arguments

`timestamps` An object of class `yearmon` or `Date` to be converted to `POSIXct`.

Details

Calling `as.POSIXct` on another date/time object (e.g. `Date`) applies a timezone correction to the object. This can shift the time marker by a few hours, which can have the effect of shifting the day by one unit. If the day was the first or last in a month or year, then the month or year will be off by one as well.

Coercing the object to the character representation of a `Date` prevents this adjustment from being applied, and leaves the `POSIXt` return value with the intended day, month, and year.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

`turkish`*Turkish Electricity Usage*

Description

A daily time series of electricity usage in Turkey.

Usage

```
data(turkish)
```

Format

zoo time series

Source

https://robjhyndman.com/data/turkey_elec.csv

See Also

[bsts](#)

Examples

```
data(turkish)
plot(turkish)
```

week.ends

Check to see if a week contains the end of a month or quarter

Description

Returns a logical vector indicating whether the given week contains the end of a month or quarter.

Usage

```
WeekEndsMonth(week.ending)
WeekEndsQuarter(week.ending)
```

Arguments

week.ending A vector of class [Date](#). Each entry contains the date of the last day in a week.

Value

A logical vector indicating whether the given week contains the end of a month or a quarter.

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

See Also

[bsts.mixed](#).

Examples

```
week.ending <- as.Date(c("2011-10-01",
                        "2011-10-08",
                        "2011-12-03",
                        "2011-12-31"))
WeekEndsMonth(week.ending) == c(TRUE, FALSE, TRUE, TRUE)
WeekEndsQuarter(week.ending) == c(TRUE, FALSE, FALSE, TRUE)
```

weekday.names	<i>Days of the Week</i>
---------------	-------------------------

Description

A character vector listing the names the days of the week.

Usage

```
weekday.names
```

See Also

[month.name](#)

wide.to.long	<i>Convert Between Wide and Long Format</i>
--------------	---

Description

Convert a multivariate time series between wide and long formats. In "wide" format there is one row per time point, with series organized by columns. In "long" format there is one row per observation, with variables indicating the series and time point to which an observation belongs.

Usage

```
WideToLong(response, na.rm = TRUE)
LongToWide(response, series.id, timestamps)
```

Arguments

response	For WideToLong this is a matrix, with rows representing time and columns representing variables. This can be a zoo matrix with timestamps as an index. For LongToWide, response is a vector.
na.rm	If TRUE then missing values will be omitted from the returned data frame (their absence denoting missingness). Otherwise, missing values will be included as NA's.
series.id	A factor (or variable coercible to factor) of the same length as response, indicating the series to which each observation belongs.
timestamps	A variable of the same length as response, indicating the time period to which each observation belongs.

Value

LongToWide returns a zoo matrix with the time series in wide format. WideToLong returns a 3-column data frame with columns "time", "series", and "values".

Author(s)

Steven L. Scott <steve.the.bayesian@gmail.com>

Examples

```
data(gdp)
gdp.wide <- LongToWide(gdp$GDP, gdp$Country, gdp$Time)
gdp.long <- WideToLong(gdp.wide)
```

Index

- * **character**
 - shorten, 96
- * **chron**
 - aggregate.time.series, 27
 - aggregate.weeks.to.months, 28
 - bsts.options.Rd, 37
 - compare.bsts.models, 38
 - date.range, 40
 - estimate.time.scale, 47
 - extend.time, 48
 - format.timestamps, 49
 - get.fraction, 51
 - last.day.in.month, 57
 - match.week.to.month, 58
 - month.distance, 67
 - quarter, 90
 - regularize.timestamps, 93
 - state.sizes, 100
 - week.ends, 105
- * **datasets**
 - gdp, 50
 - goog, 52
 - iclaims, 56
 - new.home.sales, 69
 - rsxf, 95
 - shark, 96
 - turkish, 104
- * **models**
 - add.ar, 4
 - add.dynamic.regression, 5
 - add.local.level, 9
 - add.local.linear.trend, 10
 - add.random.walk.holiday, 14
 - add.seasonal, 16
 - add.semilocal.linear.trend, 18
 - add.shared.local.level, 20
 - add.static.intercept, 22
 - add.student.local.linear.trend, 23
 - add.trig, 25
 - auto.ar, 30
 - bsts, 31
 - dirm, 43
 - dirm-model-optoins, 46
 - HarveyCumulator, 53
 - max.window.width, 59
 - mbsts, 60
 - mixed.frequency, 65
 - one.step.prediction.errors, 69
 - predict.bsts, 84
 - regression.holiday, 90
 - simulate.fake.mixed.frequency.data, 97
 - StateSpecification, 101
 - to.posixt, 104
- * **package**
 - bsts-package, 3
- * **regression**
 - bsts, 31
 - dirm, 43
 - dirm-model-optoins, 46
 - HarveyCumulator, 53
 - max.window.width, 59
 - mbsts, 60
 - mixed.frequency, 65
 - one.step.prediction.errors, 69
 - predict.bsts, 84
 - simulate.fake.mixed.frequency.data, 97
 - StateSpecification, 101
 - to.posixt, 104
- acf, 42
- AcfDist (diagnostic-plots), 42
- add.ar, 4
- add.dynamic.regression, 5
- add.local.level, 9
- add.local.linear.trend, 10
- add.monthly.annual.cycle, 12
- add.random.walk.holiday, 14

- add.seasonal, 16
- add.semilocal.linear.trend, 18
- add.shared.local.level, 20
- add.static.intercept, 22
- add.student.local.linear.trend, 23
- add.trig, 25
- AddAr, 30
- AddAr (add.ar), 4
- AddAutoAr (auto.ar), 30
- AddDynamicRegression, 34, 45, 63
- AddDynamicRegression
 - (add.dynamic.regression), 5
- AddGeneralizedLocalLinearTrend
 - (add.semilocal.linear.trend), 18
- AddHierarchicalRegressionHoliday, 55, 56, 60
- AddHierarchicalRegressionHoliday
 - (regression.holiday), 90
- AddLocalLevel, 34, 45, 61, 63, 67, 70, 87
- AddLocalLevel (add.local.level), 9
- AddLocalLinearTrend, 33, 34, 43, 45, 63, 67, 70, 87, 99
- AddLocalLinearTrend
 - (add.local.linear.trend), 10
- AddMonthlyAnnualCycle
 - (add.monthly.annual.cycle), 12
- AddRandomWalkHoliday, 55, 56, 60, 80
- AddRandomWalkHoliday
 - (add.random.walk.holiday), 14
- AddRegressionHoliday, 55, 56, 60
- AddRegressionHoliday
 - (regression.holiday), 90
- AddSeasonal, 12, 33, 34, 43, 45, 61, 63
- AddSeasonal (add.seasonal), 16
- AddSemilocalLinearTrend, 34, 45, 63, 67, 70, 87
- AddSemilocalLinearTrend
 - (add.semilocal.linear.trend), 18
- AddSharedLocalLevel, 61
- AddSharedLocalLevel
 - (add.shared.local.level), 20
- AddStaticIntercept
 - (add.static.intercept), 22
- AddStudentLocalLinearTrend
 - (add.student.local.linear.trend), 23
- AddTrig (add.trig), 25
- aggregate.time.series, 27
- aggregate.weeks.to.months, 28
- AggregateTimeSeries, 29
- AggregateTimeSeries
 - (aggregate.time.series), 27
- AggregateWeeksToMonths
 - (aggregate.weeks.to.months), 28
- Ar1CoefficientPrior, 19, 101
- as.data.frame, 6, 33, 44, 61
- as.POSIXct, 104
- auto.ar, 30
- boxplot, 42, 80
- bsts, 3, 5, 8, 10, 12, 15, 17, 20–22, 24, 27, 31, 31, 33, 34, 39, 43–45, 57, 61, 63, 67, 69–71, 73, 75, 77–80, 83–87, 92, 95, 100–103, 105
- bsts-package, 3
- bsts.mixed, 48, 52, 53, 59, 76, 77, 97–99, 105
- bsts.mixed (mixed.frequency), 65
- bsts.options.Rd, 37
- bsts.prediction, 78, 83
- bsts.prediction (predict.bsts), 84
- bsts.prediction.errors, 39
- bsts.prediction.errors
 - (one.step.prediction.errors), 69
- BstsOptions, 33, 70
- BstsOptions (bsts.options.Rd), 37
- compare.bsts.models, 38
- CompareBstsModels
 - (compare.bsts.models), 38
- ConditionalZellnerPrior, 21
- contr.treatment, 33, 44, 62
- Date, 13, 14, 29, 40, 41, 44, 47–49, 51, 57, 59, 68, 73, 90, 91, 94, 104, 105
- date.range, 40
- DateRange (date.range), 40
- DateRangeHoliday (holiday), 54
- DateToPOSIX (to.posix), 104
- DayPlot (descriptive-plots), 41
- descriptive-plots, 41
- diagnostic-plots, 42
- dirm, 43, 44
- dirm-model-options, 46
- DirmModelOptions, 44

- DirmModelOptions (dirm-model-optoins),
46
- DoubleModel, 23, 24
- DynamicRegressionArOptions
(add.dynamic.regression), 5
- DynamicRegressionHierarchicalRandomWalkOptions
(add.dynamic.regression), 5
- DynamicRegressionOptions, 6
- DynamicRegressionOptions
(add.dynamic.regression), 5
- DynamicRegressionRandomWalkOptions
(add.dynamic.regression), 5

- estimate.time.scale, 47
- EstimateTimeScale
(estimate.time.scale), 47
- extend.time, 48
- ExtendTime (extend.time), 48

- FixedDateHoliday (holiday), 54
- format.timestamps, 49

- GammaPrior, 6, 7
- gdp, 50
- geometric.sequence, 50
- GeometricSequence (geometric.sequence),
50
- get.fraction, 51
- GetFractionOfDaysInInitialMonth
(get.fraction), 51
- GetFractionOfDaysInInitialQuarter
(get.fraction), 51
- GOOG (goog), 52
- goog, 52

- HarveyCumulator, 53
- HasDuplicateTimestamps
(format.timestamps), 49
- HierarchicalRegressionHolidayStateModel,
15, 80
- HierarchicalRegressionHolidayStateModel
(regression.holiday), 90
- Holiday, 14, 40, 59, 60, 80, 91
- Holiday (holiday), 54
- holiday, 54

- iclaims, 56
- initial.claims (iclaims), 56
- install.packages, 3

- InverseWishartPrior, 91
- IsRegular (format.timestamps), 49

- last.day.in.month, 57
- LastDayInMonth (last.day.in.month), 57
- LastWeekdayInMonthHoliday (holiday), 54
- lines, 41, 78, 84
- lm, 33, 43, 44, 62
- LongToWide (wide.to.long), 106

- MATCH, 58
- match, 58
- MATCH.NumericTimestamps, 58
- match.week.to.month, 58
- MatchWeekToMonth (match.week.to.month),
58
- max.window.width, 59
- MaxWindowWidth (max.window.width), 59
- mbsts, 60, 62, 82, 83, 88, 89
- mbsts.prediction (predict.mbsts), 88
- mixed.frequency, 65
- model.matrix.default, 33, 44, 62
- month.distance, 67
- month.name, 106
- MonthDistance (month.distance), 67
- MonthlyAnnualCycle, 74
- MonthlyAnnualCycle
(add.monthly.annual.cycle), 12
- MonthPlot (descriptive-plots), 41
- monthplot, 41
- months, 41, 73
- MvnPrior, 21, 27, 91

- named.holidays, 68
- NamedHoliday (holiday), 54
- new.home.sales, 69
- NoDuplicates (format.timestamps), 49
- NoGaps (format.timestamps), 49
- NormalPrior, 8, 10–13, 15, 17, 19–22, 24, 25,
91, 92, 101
- NthWeekdayInMonthHoliday (holiday), 54

- one.step.prediction.errors, 69

- plot, 41, 78, 79, 83
- plot.bsts, 32, 71, 83, 95, 103
- plot.bsts.mixed, 75
- plot.bsts.prediction, 77
- plot.bsts.predictors, 79

- plot.holiday, 80
- plot.lm.spike, 73, 75, 77, 79, 80
- plot.mbsts, 82
- plot.mbsts.prediction, 83, 89
- PlotBstsCoefficients, 74, 76
- PlotBstsCoefficients (plot.bsts), 71
- PlotBstsComponents, 74, 76
- PlotBstsComponents (plot.bsts), 71
- PlotBstsForecastDistribution, 74
- PlotBstsForecastDistribution (plot.bsts), 71
- PlotBstsMixedComponents (plot.bsts.mixed), 75
- PlotBstsMixedState, 76
- PlotBstsMixedState (plot.bsts.mixed), 75
- PlotBstsPredictionErrors, 74
- PlotBstsPredictionErrors (plot.bsts), 71
- PlotBstsPredictors (plot.bsts.predictors), 79
- PlotBstsResiduals, 74
- PlotBstsResiduals (plot.bsts), 71
- PlotBstsSize, 74, 77
- PlotBstsSize (plot.bsts), 71
- PlotBstsState, 74
- PlotBstsState (plot.bsts), 71
- PlotDynamicDistribution, 42, 74–80, 82, 84
- PlotDynamicRegression (plot.bsts), 71
- PlotHoliday (plot.holiday), 80
- PlotMbstsSeriesMeans (plot.mbsts), 82
- PlotMonthlyAnnualCycle, 74
- PlotMonthlyAnnualCycle (plot.bsts), 71
- PlotSeasonalEffect, 74
- PlotSeasonalEffect (plot.bsts), 71
- POSIXlt, 90
- POSIXt, 13, 40, 41, 44, 49, 73, 94
- predict, 86
- predict.bsts, 78, 83, 84, 89
- predict.mbsts, 88
- qqdist (diagnostic-plots), 42
- qqline, 42
- qqnorm, 42
- Quarter (quarter), 90
- quarter, 90
- quarters, 41, 73
- rainbow, 39
- RandomWalkHolidayStateModel (add.random.walk.holiday), 14
- regression.holiday, 90
- RegressionHolidayStateModel, 15, 80
- RegressionHolidayStateModel (regression.holiday), 90
- regularize.timestamps, 93
- RegularizeTimestamps (regularize.timestamps), 93
- residuals.bsts, 94
- retail.sales (rsxf), 95
- RSXFS (rsxf), 95
- rsxf, 95
- SdPrior, 5, 7, 8, 10–14, 17, 19–25, 27, 31, 33, 34, 45, 63, 67, 70, 92, 101
- shark, 96
- Shorten (shorten), 96
- shorten, 96
- simulate.fake.mixed.frequency.data, 97
- SimulateFakeMixedFrequencyData (simulate.fake.mixed.frequency.data), 97
- spike.slab.ar.prior, 99
- SpikeSlabArPrior, 30
- SpikeSlabArPrior (spike.slab.ar.prior), 99
- SpikeSlabPrior, 33, 34, 44, 45, 62, 63, 66, 67, 70
- state.sizes, 100
- state.specification, 33, 43
- state.specification (StateSpecification), 101
- StateSizes (state.sizes), 100
- StateSpecification, 101
- SuggestBurn, 102
- summary.bsts, 102
- summary.lm.spike, 102, 103
- TimeSeriesBoxplot, 74, 76
- to.posixt, 104
- ts, 32, 41
- turkish, 104
- week.ends, 105
- weekday.names, 106
- weekdays, 41, 73
- WeekEndsMonth (week.ends), 105
- WeekEndsQuarter (week.ends), 105

`wide.to.long`, [106](#)
`WideToLong` (`wide.to.long`), [106](#)
`xts`, [14](#), [32](#), [91](#)
`yearmon`, [41](#), [104](#)
`YearMonToPOSIX` (`to.posixt`), [104](#)
`YearPlot` (`descriptive-plots`), [41](#)
`zoo`, [13](#), [14](#), [29](#), [32](#), [33](#), [41](#), [65](#), [91](#), [98](#), [106](#)