

Package ‘cts’

March 17, 2026

Title Continuous Time Autoregressive Models

Version 1.0-24

Date 2026-03-07

Description Provides tools for fitting continuous-time autoregressive (CAR) and complex CAR (CZAR) models for irregularly sampled time series using an exact Gaussian state-space formulation and Kalman filtering/smoothing. Implements maximum-likelihood estimation with stable parameterizations of characteristic roots, model selection via AIC, residual and spectral diagnostics, forecasting and simulation, and extraction of fitted state estimates. Methods are described in Wang (2013) <[doi:10.18637/jss.v053.i05](https://doi.org/10.18637/jss.v053.i05)>.

Suggests R.rsp

VignetteBuilder R.rsp

Maintainer Zhu Wang <zwang145@uthsc.edu>

License GPL (>= 2)

NeedsCompilation yes

Author Granville Tunnicliffe-Wilson [aut] (Fortran code),
Zhu Wang [cre, aut] (Fortran code, R port, updates, and maintainer),
Cleve Moler [ctb, cph] (LINPACK routines in src/d*),
Jack Dongarra [ctb, cph] (LINPACK routines via Netlib),
Jim Bunch [ctb, cph] (LINPACK routines),
G. W. Stewart [ctb, cph] (LINPACK routines),
John Nash [ctb]

Repository CRAN

Date/Publication 2026-03-17 19:00:10 UTC

Contents

asth	2
car	3
car_control	6
factab	8
kalsmoComp	9

Description

Fit a continuous AR model to an irregularly sampled univariate time series with the Kalman filter

Usage

```
car(x, y=NULL, scale = 1.5, order = 3, ctrl=car_control())
## S3 method for class 'car'
print(x, digits = 3, ...)
## S3 method for class 'car'
summary(object, ...)
## S3 method for class 'car'
plot(x, type=c("spec", "pred", "diag"),...)
## S3 method for class 'car'
predict(object, se.fit = TRUE, digits = 3, plot.it=TRUE,...)
## S3 method for class 'car'
spectrum(object, frmult=1, n.freq, plot.it = TRUE, na.action = na.fail, ...)
## S3 method for class 'car'
AIC(object, ..., verbose=TRUE, k=2)
## S3 method for class 'car'
tsdiag(object, gof.lag = 10, ...)
## S3 method for class 'car'
kalsmo(object)
```

Arguments

x	two column data frame or matrix with the first column being the sampled time and the second column being the observations at the first column; otherwise x is a numeric vector of sampled time. It can be a car object for S3 methods
y	not used if x has two columns; otherwise y is a numeric vector of observations at sampled time x.
scale	The kappa value referred to in the paper by Belcher et al. (1994). We now recommend selection of kappa along with the model order by using AIC. Also, it is suggested to choose kappa close to 2π times $1/\text{mean}.\text{delta}$ (reciprocal of the mean time between observations), though it is a good idea to explore somewhat lower and higher values to see whether the spectrum estimates were sensitive to this choice. Choosing kappa lower increases the risk of trying to estimate the spectrum beyond the effective Nyquist frequency of the data - though this does depend on the distribution of intersample times.
order	order of autoregression.
ctrl	control parameters used in predict and numerical optimization.
object	object of class car

type	in plot a character indicating the type of plot. type="spec", call spec; type="pred", call predict; type="diag", call diag
se.fit	Logical: should standard errors of prediction be returned?
digits	return value digits
plot.it	Logical: plot the forecast values?
gof.lag	the maximum number of lags for a Portmanteau goodness-of-fit test
frmult	numerical value, can be used to multiply the frequency range
n.freq	number of frequency
verbose	logical; if TRUE, prints the model selection table using message .
k	numeric penalty per parameter (as in AIC); default is 2.
na.action	NA action function.
...	further arguments to be passed to particular methods

Details

spectrum returns (and by default plots) the spectral density of the fitted model.

tsdiag is a generic diagnostic function for continuous AR model. It will generally plot the residuals, often standadized, the autocorrelation function of the residuals, and the p-values of a Portmanteau test for all lags up to gof.lag. The method for [car](#) object plots residuals scaled by the estimate of their (individual) variance, and use the Ljung–Box version of the portmanteau test.

AIC For continuous CAR model selection, t-statistic and AIC are calculated based on reparameterized coefficients phi and covariance matrix ecov. From the t-statistic, the final model is chosen such that if the true model order is less than the large value used for model estimation then for $i > \text{order}$ the deviations of the estimated parameters phi from their true value of 0 will be small. From the AIC, the final model is chosen based on the smallest AIC value. A table with t-statistic and AIC for the corresponding model order.

factab calculate characteristic roots and system frequency from the estimated reparameterized coefficients of CAR fits.

smooth computes components corresponding to the diagonal transition matrix with the Kalman smoother. This may not be stable for some data due to numerical inversion of matrix.

Value

A list of class "car" with the following elements:

n.used	The number of observations used in fitting
order	The order of the fitted model. This is chosen by the user.
np	The number of parameters estimated. This may include the mean and the observation noise ratio.
scale	The kappa value referred to in the paper of Belcher et al.
vri	If vri=1, estimate the observation noise ratio.
vr	The estimated observation noise ratio.
sigma2	The estimated innovation variance.

phi	The estimated reparameterized autoregressive parameters.
x.mean	The estimated mean of the series used in fitting and for use in prediction.
b	All estimated parameters, which include phi, and possibly x.mean and vr.
delb	The estimated standard error of b
essp	The estimated correlation matrix of b
ecov	The estimated covariance matrix of phi. See also AIC
rootr	The real part of roots of phi. See also AIC
rooti	The imaginary part of roots of phi. See also AIC
tim	The numeric vector of sampled time.
ser	The numeric vector of observations at sampled time tim.
filser	The filtered time series with the Kalman filter.
filvar	The estimated variance of Kalman filtered time series filser
sser	The smoothed time series with the Kalman smoother.
svar	The estimated variance of smoothed time series sser
stdred	The standardized residuals from the fitted model.
pretime	Time of predictions.
pred	Predictions for the pretime.
prv	Prediction variance of pred
pre2	Fitted values including pred for all the time series. See also fty.
prv2	Variance of fitted values including prv for all the time series. See also fty.
fty	Forecast type
tnit	Numeric vector: iteration
ss	Numeric vector: sum of squares for each tnit
bit	Matrix with rows for tnit and columns for parameter estimates
aic	AIC value for the fitted model
bic	BIC value for the fitted model

Author(s)

G. Tunnicliffe Wilson and Zhu Wang

References

- Belcher, J. and Hampton, J. S. and Tunnicliffe Wilson, G. (1994). Parameterization of continuous time autoregressive models for irregularly sampled time series data. *Journal of the Royal Statistical Society, Series B, Methodological*, **56**, 141–155
- Jones, Richard H. (1981). Fitting a continuous time autoregression to discrete data. *Applied Time Series Analysis II*, 651–682
- Wang, Zhu (2004). *The Application of the Kalman Filter to Nonstationary Time Series through Time Deformation*. PhD thesis, Southern Methodist University

Wang, Zhu and Woodward, W. A. and Gray, H. L. (2009). The Application of the Kalman Filter to Nonstationary Time Series through Time Deformation. *Journal of Time Series Analysis*, **30(5)**, 559-574.

Wang, Zhu (2013). cts: An R Package for Continuous Time Autoregressive Models via Kalman Filter. *Journal of Statistical Software*, **Vol. 53(5)**, 1–19. <https://www.jstatsoft.org/v53/i05>

See Also

[car_control](#) for predict and numerical optimization parameters, and [AIC](#) for model selection

Examples

```
data(V22174)
fit <- car(V22174,scale=0.2,order=7, ctrl=car_control(trace=TRUE))
summary(fit)
spectrum(fit)
tsdiag(fit)
AIC(fit)
factab(fit)
###fitted values vs observed values
ntim <- dim(V22174)[1]
plot(V22174[,1], V22174[,2], type="l")
points(V22174[,1], fit$pre2[1:ntim], col="red")
### alternatively
fit2 <- car(V22174,scale=0.2,order=7, ctrl=car_control(fty=3))
plot(V22174[,1], V22174[,2], type="l")
points(V22174[,1], fit2$pre2, col="red")

data(asth)
fit <- car(asth,scale=0.25,order=4, ctrl=car_control(n.ahead=10))
kalsmo(fit)
```

car_control

Parameters for Predict and Numerical Optimization in Kalman Filter

Description

Setup in the predict and numerical optimization in the Kalman filter algorithm.

Usage

```
car_control(fty=1, n.ahead=10, trace=FALSE, ari=TRUE, vri=FALSE, vr=0, pfi="MAPS",
ccv="CTES", lpv=TRUE, scc=TRUE, nit=40, opm=1, rgm=1, req=0.5, con=1.0e-5, rpe=1.0,
ivl=1.0e-2, fac=1.0e1, stl=1.0e-5, sml=1.0e2, gtl=1.0e5, kst=TRUE, fct=TRUE)
```

Arguments

fty	fty=1 forecast past the end. fty=2 forecast last L-steps. fty=3 forecast last L-steps updated (filtering)types. See also fct. If fct=TRUE, all time series is used to fit the model for fty=1 or 3. If fty=2, only the first (length of time - n.ahead) is used to fit the model. Thus, only the first (length of time - n.head) prediction values are the same for fty=1 or 3 vs fty=2. See also pre2, prv2.
n.ahead	number of steps ahead at which to predict.
trace	a logical value triggering printout of information during the fitting process, and major results for the fitted model.
ari	ari=TRUE: parameter starting values. ari=FALSE: they are taken as zero. This is obsolete.
vri	vri=FALSE, observation noise not included in the model. vri=TRUE, observation noise included
vr	0.5, initial value of observation noise ratio: only if vri=TRUE
pfi	always use the option pfi="MAPS".
ccv	ccv="CTES" for constant term estimation. ccv="MNCT" if mean correction, ccv=NULL if omitted.
lpv	lpv=TRUE always use this option.
scc	scc=TRUE always use this option.
nit	number of iterations.
opm	opm=1 always use this.
rgm	rgm=1 always use this.
req	root equality switch value.
con	convergence criterion.
rpe	relative size of parameter perturbations.
ivl	initial value of step size constraint parameter.
fac	step size constraint modification parameter. This value may be setup to fac=5 for better convergency.
stl	typical smallest step size parameter.
sml	typical small step size parameter.
gtl	typical greatest step size parameter.
kst	kst=TRUE to save estimated states.
fct	fct=TRUE to use all time series to fit the model.

Details

Objects returned by this function specify predict and numerical optimization parameters of the Kalman filter algorithms implemented in `car`, (via the `ctrl` argument).

Value

An object of class `car_control`, a list.

See Also

[car](#) for the usage

factab

Calculate Characteristic Roots and System Frequency

Description

Calculate characteristic roots and system frequency from the estimated reparameterized coefficients of CAR fits.

Usage

```
factab(object)
```

Arguments

object a fitted time-series CAR model

Value

A table with characteristic roots and frequencies for the corresponding model fit.

Author(s)

G. Tunnicliffe Wilson and Zhu Wang

References

- Belcher, J. and Hampton, J. S. and Tunnicliffe Wilson, G. (1994). Parameterization of continuous time autoregressive models for irregularly sampled time series data. *Journal of the Royal Statistical Society, Series B, Methodological*, **56**, 141–155
- Jones, Richard H. (1981). Fitting a continuous time autoregression to discrete data. *Applied Time Series Analysis II*, 651–682
- Wang, Zhu(2004). *The Application of the Kalman Filter to Nonstationary Time Series through Time Deformation*. PhD thesis, Southern Methodist University
- Wang, Zhu (2013). cts: An R Package for Continuous Time Autoregressive Models via Kalman Filter. *Journal of Statistical Software*, **Vol. 53(5)**, 1–19. <https://www.jstatsoft.org/v53/i05>

See Also

[car](#) and [kalsmo](#)

Examples

```
data(asth)
(fit <- car(asth, scale=0.25, order=4))
factab(fit)
```

Description

Estimate unobserved components with the Kalman smoother to a fitted CAR model.

Usage

```
kalsmoComp(x, comp = NULL, plot.it = TRUE, xlab= "time", ylab="", na.action = na.fail, ...)
```

Arguments

x	the result of estimated components by kalsmo .
comp	a numeric vector from which components are estimated
plot.it	plot the component?
xlab	name of xlab
ylab	name of ylab
na.action	how to handle NAs?
...	further graphical parameters.

Value

A component is computed from the estimated components for each root of the characteristic equation.

Author(s)

Zhu Wang

References

Belcher, J. and Hampton, J. S. and Tunnicliffe Wilson, G. (1994). Parameterization of continuous time autoregressive models for irregularly sampled time series data. *Journal of the Royal Statistical Society, Series B, Methodological*, **56**, 141–155

Wang, Zhu (2004). *The Application of the Kalman Filter to Nonstationary Time Series through Time Deformation*. PhD thesis, Southern Methodist University

Wang, Zhu (2013). cts: An R Package for Continuous Time Autoregressive Models via Kalman Filter. *Journal of Statistical Software*, **Vol. 53(5)**, 1–19. <https://www.jstatsoft.org/v53/i05>

See Also

[kalsmo](#)

Examples

```
data(asth)
kalsmoComp(kalsmo(car(asth,scale=0.25,order=4)),c(2,3))
```

plotSpecCar

Plotting Spectral Densities

Description

Internal function used in spectrum.car only. Not called by user. Plotting method for objects of class "spectrum.car".

Usage

```
plotSpecCar(x, add = FALSE, ci = 0.95, log = "dB", xlab = "frequency",
ylab = NULL, type = "l", main = NULL, sub = NULL,...)
```

Arguments

x	an object of class "spectrum.car".
add	logical. If TRUE, add to already existing plot.
ci	Coverage probability for confidence interval. Plotting of the confidence bar is omitted unless ci is strictly positive.
log	If "dB", plot on log10 (decibel) scale (as S-PLUS), otherwise use conventional log scale or linear scale. Logical values are also accepted. The default is "yes" unless options(ts.S.compat = TRUE) has been set, when it is "dB".
xlab	the x label of the plot.
ylab	the y label of the plot.
type	the type of plot to be drawn, defaults to lines.
main	overall title for the plot.
sub	a sub title for the plot.
...	further graphical parameters.

Value

plot of spectral density from continuous time autoregressive model

Author(s)

G. Tunncliffe Wilson and Zhu Wang

References

- Belcher, J. and Hampton, J. S. and Tunnicliffe Wilson, G. (1994). Parameterization of continuous time autoregressive models for irregularly sampled time series data. *Journal of the Royal Statistical Society, Series B, Methodological*, **56**, 141–155
- Jones, Richard H. (1981). Fitting a continuous time autoregression to discrete data. *Applied Time Series Analysis II*, 651–682
- Wang, Zhu (2004). *The Application of the Kalman Filter to Nonstationary Time Series through Time Deformation*. PhD thesis, Southern Methodist University
- Wang, Zhu (2013). cts: An R Package for Continuous Time Autoregressive Models via Kalman Filter. *Journal of Statistical Software*, **Vol. 53(5)**, 1–19. <https://www.jstatsoft.org/v53/i05>

See Also

[spectrum.car](#)

plotSpecLs

Plotting Lomb-Scargle Periodogram

Description

Plotting method for objects of class "spec.ls".

Usage

```
plotSpecLs(x, add = FALSE, ci = 0.95, log = c("yes", "dB", "no"), xlab
= "frequency", ylab = NULL, type = "l", main = NULL, sub = NULL, ...)
```

Arguments

x	an object of class "spec.ls".
add	logical. If TRUE, add to already existing plot.
ci	Coverage probability for confidence interval. Plotting of the confidence bar is omitted unless ci is strictly positive.
log	If "dB", plot on log10 (decibel) scale (as S-PLUS), otherwise use conventional log scale or linear scale. Logical values are also accepted. The default is "yes" unless options(ts.S.compat = TRUE) has been set, when it is "dB".
xlab	the x label of the plot.
ylab	the y label of the plot.
type	the type of plot to be drawn, defaults to lines.
main	overall title for the plot.
sub	a sub title for the plot.
...	Further graphical parameters.

Value

Called for its side effect of producing a plot. Invisibly returns `x` (an object of class "spec.ls").

See Also

[spec.ls](#)

spec.ls	<i>Estimate Spectral Density of an Irregularly Sampled Time Series by a Smoothed Periodogram</i>
---------	--

Description

The most commonly used method of computing the spectrum on unevenly spaced time series is periodogram analysis, see Lomb (1975) and Scargle (1982). The Lomb-Scargle method for unevenly spaced data is known to be a powerful tool to find, and test significance of, weak periodic signals. The Lomb-Scargle periodogram possesses the same statistical properties of standard power spectra.

Usage

```
spec.ls(x, y=NULL, spans = NULL, kernel = NULL, taper = 0.1, pad = 0,
fast = TRUE, type = "lomb", demean = FALSE, detrend = TRUE, plot.it = TRUE,
na.action = na.fail, ...)
```

Arguments

<code>x</code>	two column data frame or matrix with the first column being the sampled time and the second column being the observations at the first column; otherwise <code>x</code> is a numeric vector of sampled time.
<code>y</code>	not used if <code>x</code> has two columns; otherwise <code>y</code> is a numeric vector of observations at sampled time <code>x</code> .the time at which <code>x</code> is observed
<code>spans</code>	vector of odd integers giving the widths of modified Daniell smoothers to be used to smooth the periodogram.
<code>kernel</code>	alternatively, a kernel smoother of class "tskernel".
<code>taper</code>	proportion of data to taper. A split cosine bell taper is applied to this proportion of the data at the beginning and end of the series.
<code>pad</code>	proportion of data to pad. Zeros are added to the end of the series to increase its length by the proportion <code>pad</code> .
<code>fast</code>	logical; if TRUE, pad the series to a highly composite length.
<code>type</code>	Lomb-Scargle spectrum of Fourier transformation spectrum
<code>demean</code>	logical. If TRUE, subtract the mean of the series.
<code>detrend</code>	logical. If TRUE, remove a linear trend from the series. This will also remove the mean.
<code>plot.it</code>	plot the periodogram?
<code>na.action</code>	NA action function.
<code>...</code>	graphical arguments passed to <code>plotSpecLs</code> .

Details

The raw Lomb-Scargle periodogram for irregularly sampled time series is not a consistent estimator of the spectral density, but adjacent values are asymptotically independent. Hence a consistent estimator can be derived by smoothing the raw periodogram, assuming that the spectral density is smooth.

The series will be automatically padded with zeros until the series length is a highly composite number in order to help the Fast Fourier Transform. This is controlled by the `fast` and not the `pad` argument.

The periodogram at zero is in theory zero as the mean of the series is removed (but this may be affected by tapering): it is replaced by an interpolation of adjacent values during smoothing, and no value is returned for that frequency.

Value

A list object of class "spec.ls" with the following additional components:

kernel	The kernel argument, or the kernel constructed from spans.
df	The distribution of the spectral density estimate can be approximated by a chi square distribution with df degrees of freedom.
bandwidth	The equivalent bandwidth of the kernel smoother as defined by Bloomfield (1976, page 201).
taper	The value of the taper argument.
pad	The value of the pad argument.
detrend	The value of the detrend argument.
demean	The value of the demean argument.

The result is returned invisibly if `plot.it` is true.

Note

This is 'slow' program and a fast program may use FFT, see (Press et al, 1992)

Author(s)

Lomb-Scargle periodogram by Zhu Wang

References

- Bloomfield, P. (1976) *Fourier Analysis of Time Series: An Introduction*. Wiley.
- Lomb, N. R. (1976) Least-squares frequency-analysis of unequally spaced data. *Astrophysics and Space Science*, **39**,447-462
- W. H. Press and S. A. Teukolsky and W. T. Vetterling and B.P. Flannery (1992) *Numerical Recipes in C: The Art of Scientific Computing.*, Cambridge University Press, second edition.
- Scargle, J.D. (1982) Studies in astronomical time series analysis II. Statistical aspects of spectral analysis of unevenly spaced data, *The Astrophysical Journal*, **263**, volume 2, 835-853.

See Also

[spec.taper](#), [plotSpecLs](#), [fft](#)

Examples

```
data(V22174)
spec.ls(V22174)
```

```
data(asth)
spec.ls(asth)
```

V22174

Measurments of Relative Aboundance

Description

Measurments of relative abundance of an xxygen isotope in an ocean core.

Usage

```
data(V22174)
```

Format

The format is: num [1:164, 1:2] 6.13 8.39 10.64 12.90 15.16 ...

Source

Belcher, J. and Hampton, J. S. and Tunnicliffe Wilson, G. (1994). Parameterization of continuous time autoregressive models for irregularly sampled time series data. *Journal of the Royal Statistical Society, Series B, Methodological*, **56**, 141–155

References

Wang, Zhu (2013). cts: An R Package for Continuous Time Autoregressive Models via Kalman Filter. *Journal of Statistical Software*, **Vol. 53(5)**, 1–19. <https://www.jstatsoft.org/v53/i05>

Examples

```
data(V22174)
str(V22174) ; plot(V22174)
```

Index

- * **datasets**
 - asth, [2](#)
 - V22174, [14](#)
- * **hplot**
 - plotSpecCar, [10](#)
 - plotSpecLs, [11](#)
- * **misc**
 - car_control, [6](#)
- * **ts**
 - car, [3](#)
 - factab, [8](#)
 - kalsmoComp, [9](#)
 - plotSpecCar, [10](#)
 - plotSpecLs, [11](#)
 - spec.ls, [12](#)

AIC, [4–6](#)
AIC.car (car), [3](#)
asth, [2](#)

car, [3, 4, 7, 8](#)
car_control, [6, 6](#)

factab, [8](#)
fft, [14](#)

kalsmo, [8, 9](#)
kalsmo (car), [3](#)
kalsmoComp, [9](#)

message, [4](#)

plot.car (car), [3](#)
plotSpecCar, [10](#)
plotSpecLs, [11, 14](#)
predict.car (car), [3](#)
print.car (car), [3](#)

spec.ls, [12, 12](#)
spec.taper, [14](#)
spectrum (car), [3](#)

spectrum.car, [11](#)
summary.car (car), [3](#)

tsdiag (car), [3](#)

V22174, [14](#)