

Autocorrelations and white noise tests

Georgi N. Boshnakov
University of Manchester

Abstract

The R package `sarima` provides functions, classes and methods for time series modelling with ARIMA and related models. The aim of the package is to provide consistent interface for the user. This document gives examples for white noise tests.

Keywords: arima, sarima, time series, S4, R.

1. Autocorrelations and related properties

The generic function `autocorrelations()` computes autocorrelations. What exactly is computed depends on the first argument. There are analogous functions for other second order characteristics, e.g. `partialAutocorrelations`, see the package documentation for the available functions.

The examples below use models used in the examples in Francq & Zakoian's book on GARCH models. It can be consulted for concepts and technical details.

```
R> n <- 100
R> ma2.model <- list(ma = c(0.56, -0.44))
R> xma2 <- arima.sim(ma2.model, n)
```

With time series or similar argument `autocorrelations` computes the sample quantity. There are similar functions for other quantities, the example below computes also partial autocorrelations:

```
R> xma2.acf <- autocorrelations(xma2, maxlag = 8)
R> class(xma2.acf)

[1] "SampleAutocorrelations"
attr(,"package")
[1] "sarima"

R> xma2.pacf <- partialAutocorrelations(xma2, maxlag = 8)
R> class(xma2.pacf)

[1] "SamplePartialAutocorrelations"
attr(,"package")
[1] "sarima"
```

If the argument is a model, the suitable theoretical properties are computed. In interactive use an ARMA model can be specified also as a list with components `ar`, `ma` and `sigma2`

```
R> xma2.tacf <- autocorrelations(ma2.model, maxlag = 8)
R> class(xma2.tacf)
```


```
[1] "Autocorrelations"
attr(,"package")
[1] "sarima"
```

```
R> xma2.tpacf <- partialAutocorrelations(ma2.model, maxlag = 8)
R> class(xma2.tpacf)
```

```
[1] "PartialAutocorrelations"
attr(,"package")
[1] "sarima"
```

A plot of the autocorrelation object `plot` can be used to check if the time series is white noise. Without further arguments, the confidence limits correspond to a null hypothesis of iid:

```
R> plot(xma2.acf)
```




If the time series is supplied, limits for null hypothesis of GARCH are given, as well. Here is an example. First we simulate a GARCH(1,1) time series and compute sample autocorrelations and partial autocorrelations.

```
R> n <- 5000
R> set.seed(124) # for reproducibility
R> x <- rgarch1p1(n, alpha = 0.3, beta = 0.55, omega = 1, n.skip = 1000)
R> x.acf <- autocorrelations(x)
R> x.pacf <- partialAutocorrelations(x)
```

This produces a plot for assessing if the time series is white noise. There are two sets of intervals: one computed under the stronger hypothesis that the time series is iid, the other under the hypothesis that the time series is GARCH.


```
R> plot(x.acf, data = x)
```



Notice that several autocorrelations seem significant under the iid hypothesis and in practice some ARMA or MA model would be tried. On the other hand, there is no evidence against the GARCH hypothesis, suggesting a GARCH model without any ARMA terms. Similar inference follows from the portmanteau tests, see next section.


The partial autocorrelation function can be used instead of the autocorrelations:

```
R> plot(x.pacf, data = x)
```



Arguments can be used to change some aspects of the graphics. Here we change the title:

```
R> plot(x.acf, data = x, main = "Autocorrelation test")
R> ## plot(x.pacf, data = x, main = "Partial autocorrelation test")
```



2. White noise tests

`whiteNoiseTest` computes portmanteau statistics. When argument `ci` is `TRUE`, it also computes the intervals for the individual correlations visualised by the above plots. Here we carry out iid tests using the method of Li-McLeod and Ljung-Box:

```
R> x.iid <- whiteNoiseTest(x.acf, h0 = "iid", nlags = c(5,10,20), x = x, method = "LiMcLeod")
R> x.iid
```

```
$test
      ChiSq DF      pvalue
[1,]  8.416634  5 0.13471982
[2,] 21.396085 10 0.01849482
[3,] 29.570841 20 0.07711244
attr(,"method")
[1] "LiMcLeod"
```

```
$ci
      int
```

```

[1,] -0.02771808 0.02771808
[2,] -0.02771808 0.02771808
[3,] -0.02771808 0.02771808
[4,] -0.02771808 0.02771808
[5,] -0.02771808 0.02771808
[6,] -0.02771808 0.02771808
[7,] -0.02771808 0.02771808
[8,] -0.02771808 0.02771808
[9,] -0.02771808 0.02771808
[10,] -0.02771808 0.02771808
[11,] -0.02771808 0.02771808
[12,] -0.02771808 0.02771808
[13,] -0.02771808 0.02771808
[14,] -0.02771808 0.02771808
[15,] -0.02771808 0.02771808
[16,] -0.02771808 0.02771808
[17,] -0.02771808 0.02771808
[18,] -0.02771808 0.02771808
[19,] -0.02771808 0.02771808
[20,] -0.02771808 0.02771808
attr(,"level")
[1] 0.95

```

```

R> x.iid2 <- whiteNoiseTest(x.acf, h0 = "iid", nlags = c(5,10,20), x = x, method = "LjungB
R> x.iid2

```

```

$test
      ChiSq DF      pvalue
[1,]  8.423474  5 0.13438990
[2,] 21.419418 10 0.01835152
[3,] 29.589845 20 0.07677784
attr(,"method")
[1] "LjungBox"

```

```

$ci
      int
[1,] -0.02771808 0.02771808
[2,] -0.02771808 0.02771808
[3,] -0.02771808 0.02771808
[4,] -0.02771808 0.02771808
[5,] -0.02771808 0.02771808
[6,] -0.02771808 0.02771808
[7,] -0.02771808 0.02771808
[8,] -0.02771808 0.02771808
[9,] -0.02771808 0.02771808
[10,] -0.02771808 0.02771808
[11,] -0.02771808 0.02771808

```

```

[12,] -0.02771808 0.02771808
[13,] -0.02771808 0.02771808
[14,] -0.02771808 0.02771808
[15,] -0.02771808 0.02771808
[16,] -0.02771808 0.02771808
[17,] -0.02771808 0.02771808
[18,] -0.02771808 0.02771808
[19,] -0.02771808 0.02771808
[20,] -0.02771808 0.02771808
attr("level")
[1] 0.95

```

If the iid hypothesis is retained this implies uncorrelatedness, as well. But it is also a rather common practice when the iid hypothesis is rejected, to take this as evidence that the time series is correlated and go on to fit ARMA models. This is fine for many applications but not always.

In particular, if the intent is to fit GARCH-type models the following test is more suitable. Here we test the weaker assumption that the underlying time series is GARCH:

```

R> x.garch <- whiteNoiseTest(x.acf, h0 = "garch", nlags = c(5,10,20), x = x)
R> x.garch

```

\$test

	h	Q	pval
[1,]	5	3.941196	0.5579130
[2,]	10	10.951332	0.3613222
[3,]	20	17.332780	0.6312711

\$ci

	int
[1,]	-0.04666052 0.04666052
[2,]	-0.04434802 0.04434802
[3,]	-0.04677694 0.04677694
[4,]	-0.03992968 0.03992968
[5,]	-0.03891428 0.03891428
[6,]	-0.03907986 0.03907986
[7,]	-0.03515187 0.03515187
[8,]	-0.03573208 0.03573208
[9,]	-0.03439233 0.03439233
[10,]	-0.03183056 0.03183056
[11,]	-0.03219418 0.03219418
[12,]	-0.03246280 0.03246280
[13,]	-0.03327740 0.03327740
[14,]	-0.03005510 0.03005510
[15,]	-0.03036421 0.03036421
[16,]	-0.03227652 0.03227652

[17,] -0.03039006 0.03039006
[18,] -0.03120361 0.03120361
[19,] -0.03219531 0.03219531
[20,] -0.03092271 0.03092271

Affiliation:

Georgi N. Boshnakov
School of Mathematics
The University of Manchester
Oxford Road, Manchester M13 9PL, UK
URL: <http://www.maths.manchester.ac.uk/~gb/>