Time series models 2008: Computer exercise I

Deadline: Monday Oct 27

October 9, 2008

A report on this computer exercise should be delivered to Marit Ulvmoen, no later than Monday Oct 27 (unless otherwise agreed upon). Email-submissions are accepted. The assumed program to use is R (see www.r-project.org), but the program following the book can be used at own risk. Groups no larger than two is accepted. Some hints are given at the end.

1 Problem 1

Create a new data set `ashort` that consists of the data in the file `http://www.math.ntnu.no/~ulvmoen/tma4285/airpass.txt` with the last twelve values deleted. Use R to find an ARIMA model for the logarithm of the data in `ashort`.

Your analysis should include

1. A logical explanation of the steps taken to find the chosen model
2. Approximate 95% bounds for the components of ϕ and θ
3. An examination of the residuals to check for whiteness as described in Section 1.6
4. A graph of the series `ashort` showing forecast of the next 12 values and 95% prediction bounds for the forecasts
5. Numerical values for the 12-step ahead forecast and the corresponding 95% prediction bounds
6. A table of the actual forecast errors, i.e., the true value (deleted from `airpass.txt`) minus the forecast value, for each of the twelve forecasts

Does the last value of `airpass.txt` lie within the corresponding 95% prediction bounds?

2 Problem 2

Repeat Problem 1, but instead of differencing, apply the classical decomposition method of the logarithms of the data in `ashort.txt` by deseasonalizing, subtracting a quadratic trend, and then finding an appropriate ARMA model for the residuals. Compare the twelve forecast errors found from this approach with those found in Problem 1.

3 Problem 3

A time series $\{X_t\}$ is differenced at lag 12, then at lag 1 to produce a zero-mean series $\{Y_t\}$ with the following sample ACF:

$$\hat{\rho}(12j) \approx (0.8)^j, j = 0, \pm 1, \pm 2, \ldots \quad (1)$$

$$\hat{\rho}(12j \pm 1) \approx (0.4)(0.8)^j, j = 0, \pm 1, \pm 2, \dots \quad (2)$$

$$\hat{\rho}(h) \approx 0, otherwise \quad (3)$$

and $\hat{\gamma}(0) = 25$.

1. Suggest a SARIMA model for $\{X_t\}$ specifying all parameters
2. For large n, express the one- and twelve-step linear predictors $P_n X_{n+1}$ and $P_n X_{n+12}$ in terms of X_t, $t = -12, -11, \dots, n$, and $Y_t - \hat{Y}_t$, $t = 1, \dots, n$
3. Find the mean squared errors of the predictors in 2.

Hints

- You can read the file into R using

```
airpass <- as.ts(scan("http://www.math.ntnu.no/~ulvmoen/tma4285/airpass.txt"))
```

- To get info about functions in R, read the help pages using the command ?function (?as.ts)

- Useful functions are

```
ts.plot(x)
diff(x,lag=y)
acf(x,type="cor")
acf(x,type="partial")
tsdiag(x)
arima(x,order=c(p,d,q),seasonal=list(order=c(a,b,c),period=s))
predict(x,n.ahead=y)
```

- You can create eps- and pdf-copies of your plots as shown on the screen, using commands like

```
library(R.utils) ## once only
dev.print(pdf,"file.pdf")
dev.print(eps,"file.eps")
```

- Although **acf(x)** plots the ACF, the actual *values* are returned in the functions call;

```
> a = acf(x)
> names(a)
[1] "acf"    "type"   "n.used" "lag"    "series" "snames"
> a$acf[1:10]
 [1] 1.0000000000 0.8912153009 0.7942161702 0.6954640762 0.6218270909
 [6] 0.5633687558 0.5074224632 0.4575082037 0.4140609409 0.3653829956
```

Similarly with the other functions.