ARMA processes

Simulation of ARMA

Simulate a time series of length 100 using an ARMA model (you may choose the number and values of parameters in both the AR and the MA part of the model)

- Fit various ARMA models with different number of parameters and compare the AIC to choose a model - do you get the same order of the model as was used in the simulation?
- Estimate the parameters in the chosen model - if you got the right order, are the estimates then close to the parameters used in the simulation?

Simulate data (here an ARMA(1,1) model used with parameters alpha1 = 0.9 and beta1 = 0.9):

```
x <- arima.sim(model = list(ar=0.9,ma=0.9), n = 100)
plot(x)
```

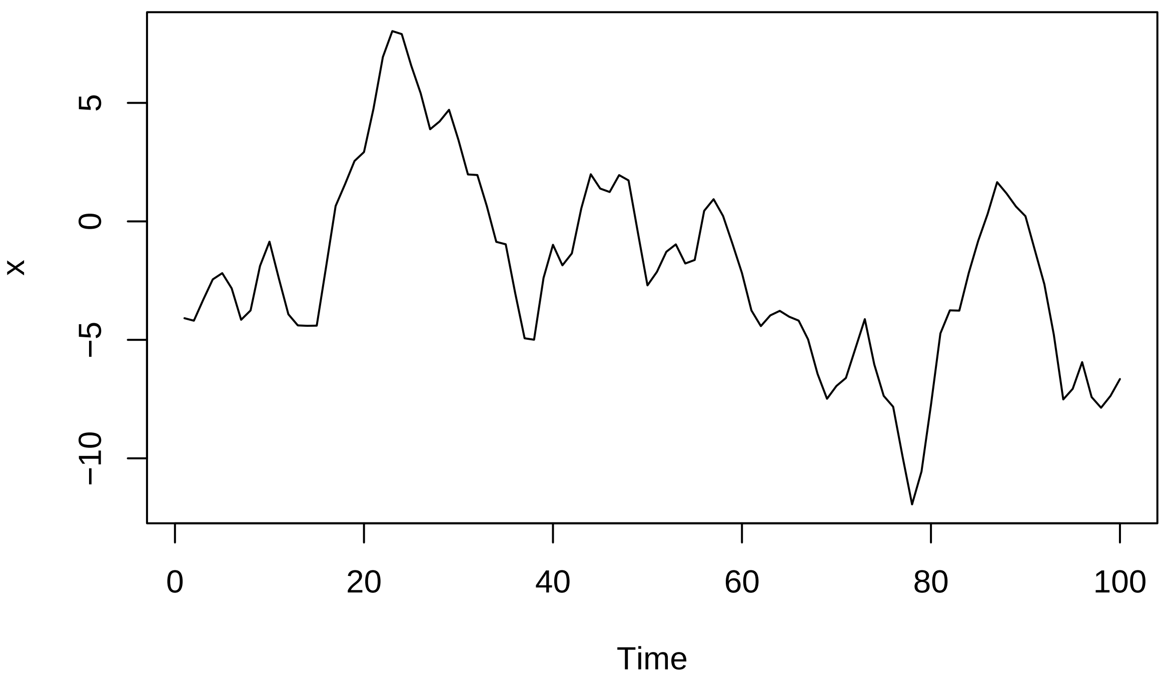

Try out various models, and look for the minimal AIC:

```
fit10 <- arima(x,order=c(1,0,0))
fit01 <- arima(x,order=c(0,0,1))
fit11 <- arima(x,order=c(1,0,1))
fit20 <- arima(x,order=c(2,0,0))
fit02 <- arima(x,order=c(0,0,2))
AIC(fit10); AIC(fit01); AIC(fit11); AIC(fit20); AIC(fit02)
```

```
## [1] 342.0329

## [1] 439.8934

## [1] 282.9398
```

```
## [1] 308.252
```

```
## [1] 363.8852
```

Exactly which model has the lowest AIC depends on the simulation, so here we just take the ARMA(1,1) model:

```
fit11
```

```
##
## Call:
## arima(x = x, order = c(1, 0, 1))
##
## Coefficients:
##          ar1     ma1  intercept
##       0.8990  0.8892    -2.5274
## s.e.  0.0415  0.0567     1.6160
##
## sigma^2 estimated as 0.876:  log likelihood = -137.47,  aic = 282.94
```

co2 data analysis

Last time we considered the co2 dataset. The code below removes trend and seasonality and extracts the noise term.

```
data<-co2
data<-ts(data,frequency = 12)
random<-na.omit(decompose(data)$random)
plot(random)
```

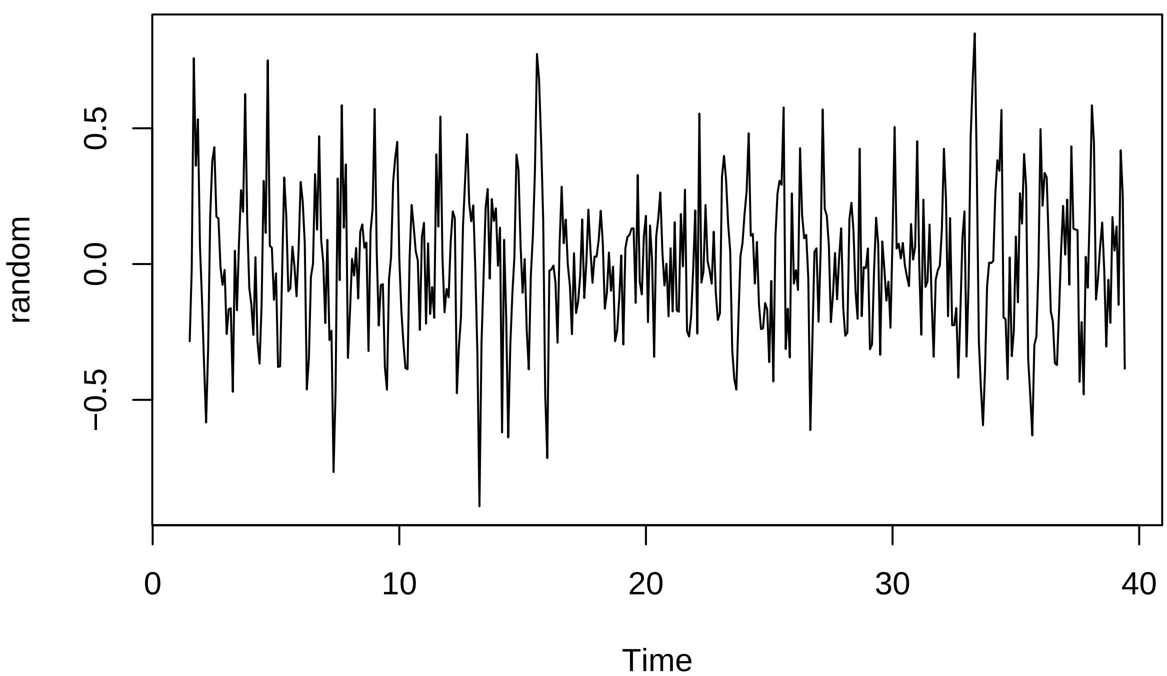

- Find the best ARMA(p,q) model with p and q at most 2 for the random component.
- Plot the acf of the residuals - does it look like white noise?
- Is the fitted model stationary? (for an ARMA model to be stationary, it is enough that the AR term is stationary)

```r
fit10 <- arima(random,order=c(1,0,0))
fit01 <- arima(random,order=c(0,0,1))
fit11 <- arima(random,order=c(1,0,1))
fit20 <- arima(random,order=c(2,0,0))
fit02 <- arima(random,order=c(0,0,2))
fit12 <- arima(random,order=c(1,0,2))
fit21 <- arima(random,order=c(2,0,1))
fit22 <- arima(random,order=c(2,0,2))
AIC(fit10); AIC(fit01); AIC(fit11); AIC(fit20); AIC(fit02); AIC(fit12); AIC(fit21); AIC(fit22)
```

```
## [1] 5.458649
## [1] 18.88495
## [1] 6.947824
## [1] 6.342985
## [1] -3.220705
## [1] -1.291071
## [1] -89.83909
## [1] -102.5696
```

ARMA(2,2) looks best

```r
res<-fit22$resid
acf(res)
```

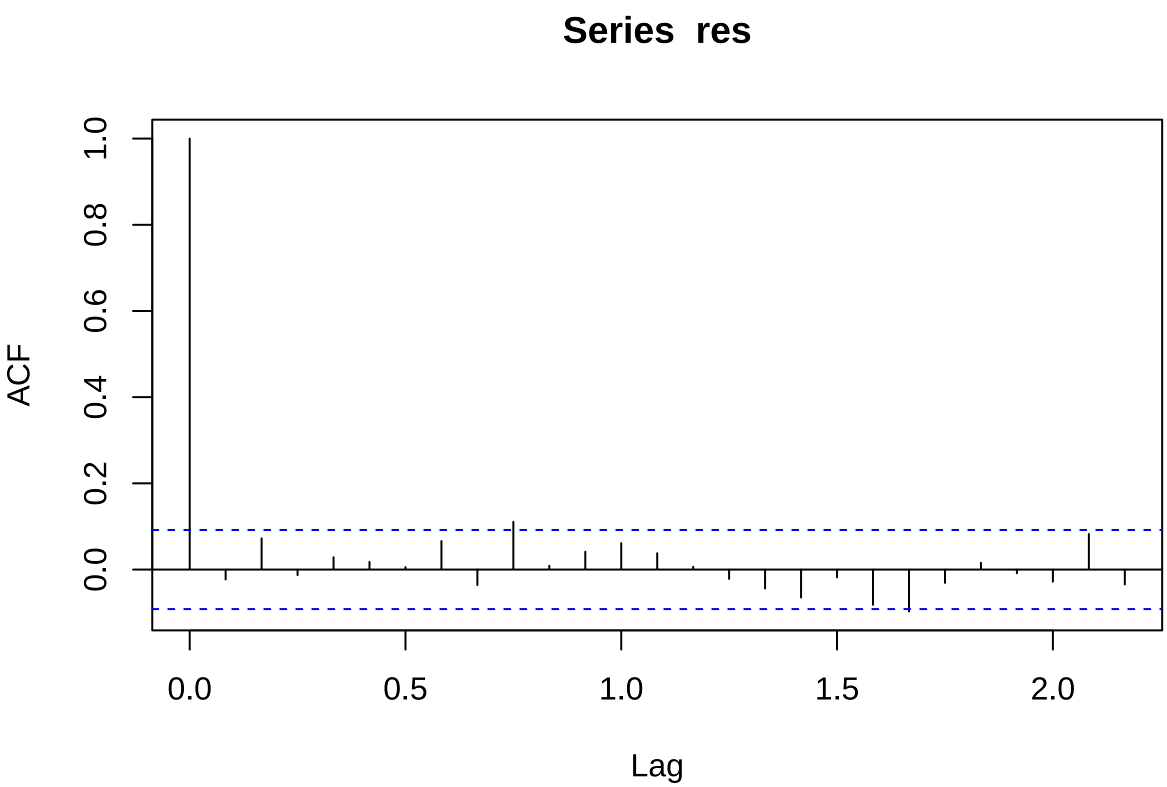

Looks like white noise

- Is the fitted model stationary? (for an ARMA model to be stationary, it is enough that the AR part is stationary, see 6.5.1 in [CM])

```r
abs(polyroot(c(1,- 1.4926390741, +0.7048490163)))
```

```
## [1] 1.19111 1.19111
```

Stationary, since all roots greater than 1