

Univariate ARIMA Forecasts

Al Nosedal
University of Toronto

April 2, 2019


Random Walk Model

Random series are sometimes building blocks for other time series models. The model we now discuss, the random walk model, is an example of this. This type of behavior is typical of stock price data. For example, the graph in the next slide shows monthly Dow Jones averages from January 1988 through March 1992. (See the file dow.txt)

Reading data

```
dow<-read.table(file="dow.txt",header=TRUE);  
  
names(dow);  
  
## [1] "Date" "Dow"
```

Time Series Plot



Loading R Library

```
library(forecast);  
  
## This is forecast 8.2  
## Want to stay up-to-date? Read the Hyndsight blog:  
## https://robjhyndman.com/hyndsight/
```


Fitting ARIMA Model

```
auto.arima(dow$Dow);  
  
## Series: dow$Dow  
## ARIMA(0,1,0) with drift  
##  
## Coefficients:  
##          drift  
##          26.0014  
## s.e.    11.8516  
##  
## sigma^2 estimated as 7166:  log likelihood=-292.37  
## AIC=588.74   AICc=589   BIC=592.56
```

Time Series Forecast

```
dow.arima<-auto.arima(dow$Dow);  
  
fcast<-forecast(dow.arima);  
  
plot(fcast);
```

Forecasts from ARIMA(0,1,0) with drift




The intervals shown are 80% and 95% prediction intervals.

The time series object *LakeHuron* (datasets) has annual depth measurements at a specific site on Lake Huron.

Reading data

```
LH<-LakeHuron;
```

Time Series Plot



Loading R Library

```
library(forecast);
```

(since we already did this for our previous example, we wouldn't need to reload it)


Fitting ARIMA Model

```
auto.arima(LH);  
  
## Series: LH  
## ARIMA(0,1,0)  
##  
## sigma^2 estimated as 0.5588: log likelihood=-109.11  
## AIC=220.22   AICc=220.26   BIC=222.79
```

Time Series Forecast

```
LH.arima<-auto.arima(LH);  
  
LH.fcast<-forecast(LH.arima);  
  
plot(LH.fcast);
```

Forecasts from ARIMA(0,1,0)



The intervals shown are 80% and 95% prediction intervals.