MONASH University

MONASH BUSINESS SCHOOL

Forecasting: principles and practice

Rob J Hyndman

3.1 Time series cross-validation

Outline

Cross-validation

Traditional evaluation

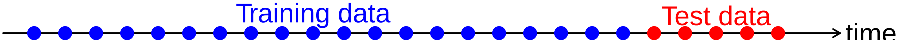

Leave-one-out cross-validation

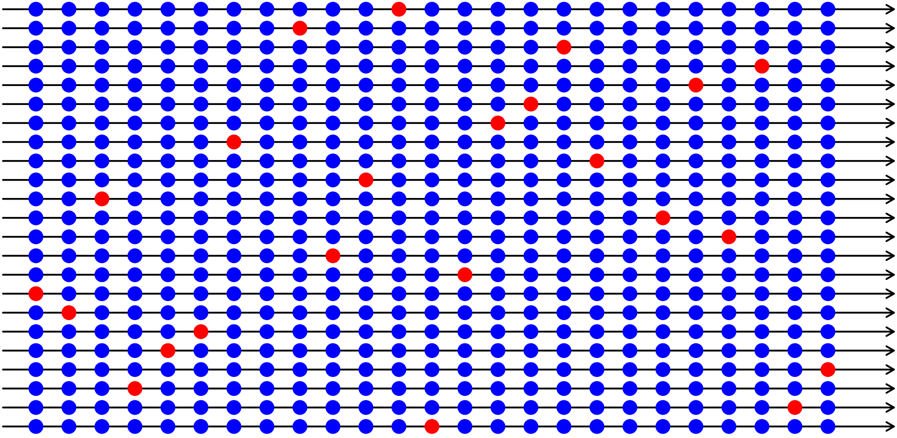

Cross-validation

Time series cross-validation

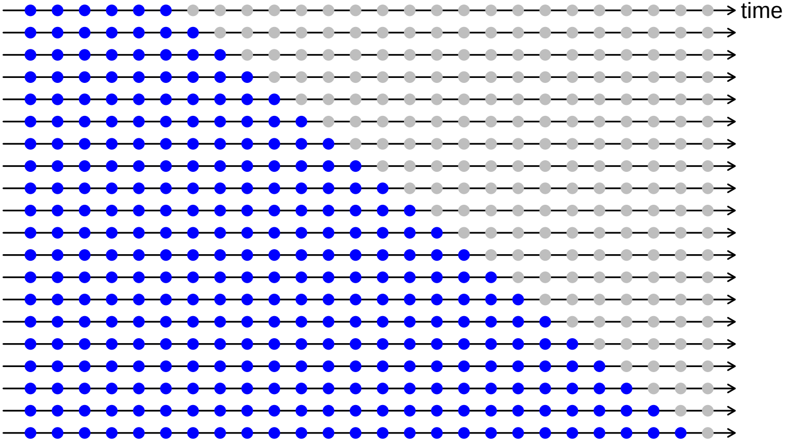

Cross-validation

Time series cross-validation

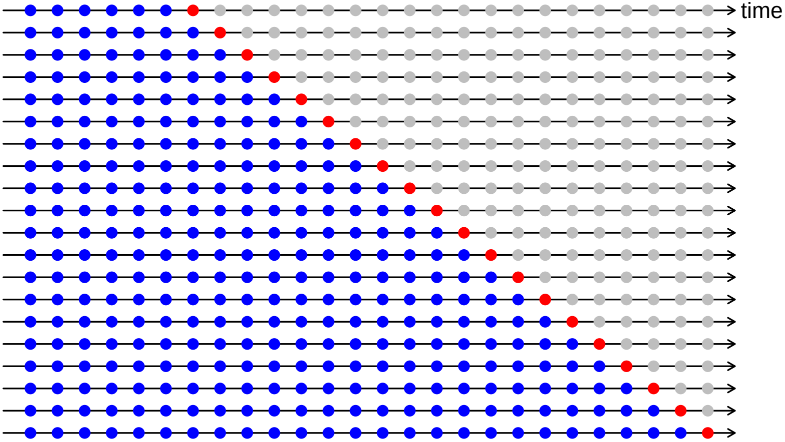

$h = 1$

Cross-validation

Time series cross-validation

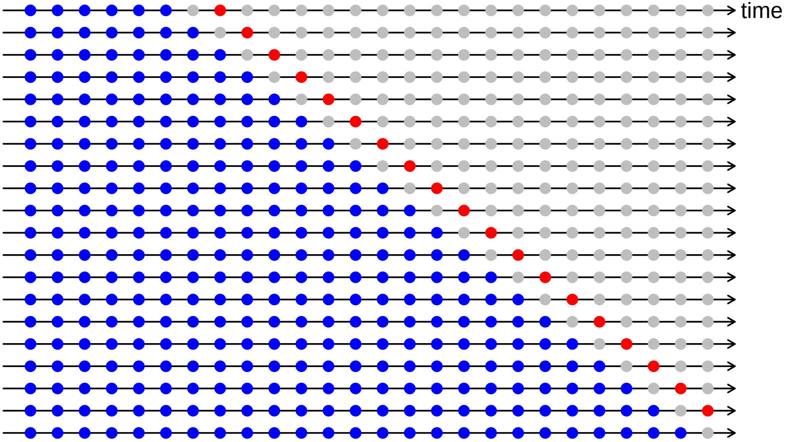

$h = 2$

Cross-validation

Time series cross-validation

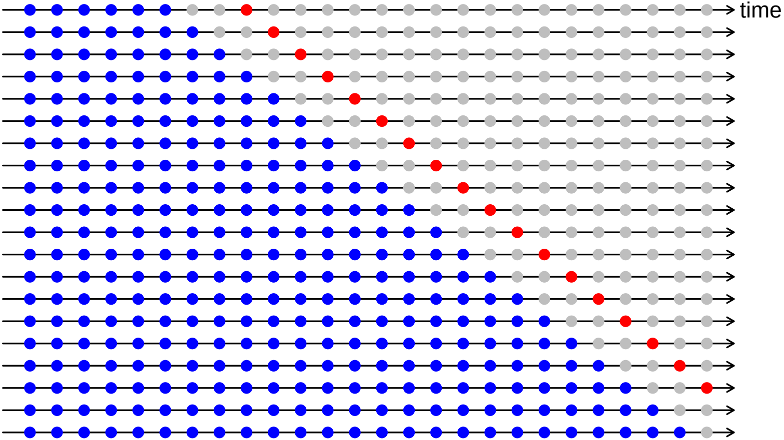

$h = 3$

Cross-validation

Time series cross-validation

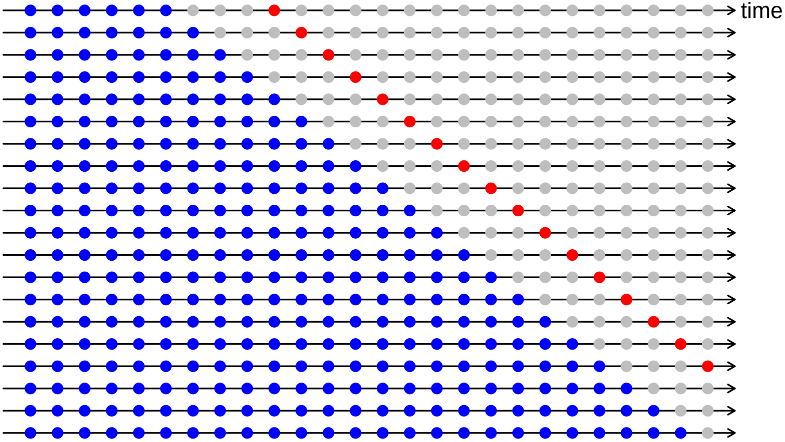

$h = 4$

Cross-validation

Time series cross-validation

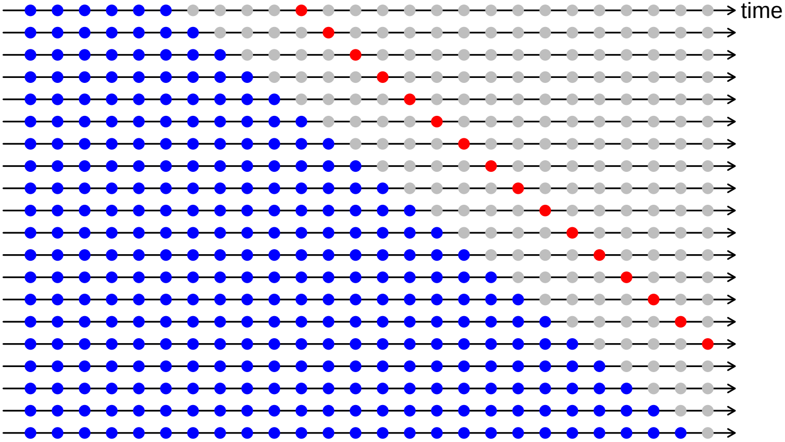

$h = 5$

Cross-validation

Time series cross-validation

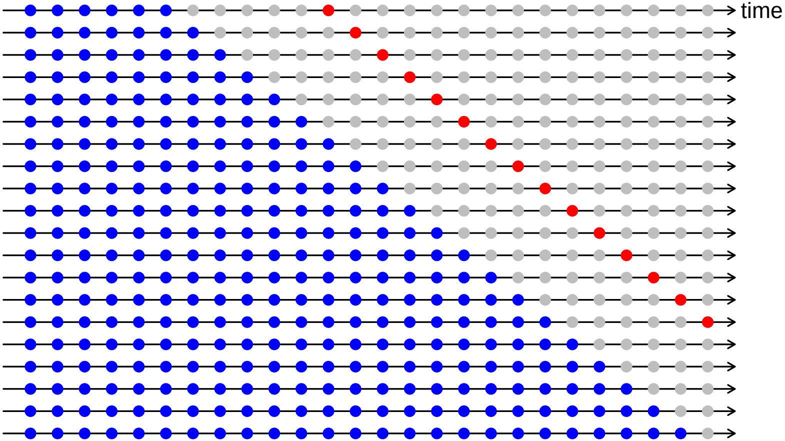

$h = 6$

Cross-validation

Time series cross-validation

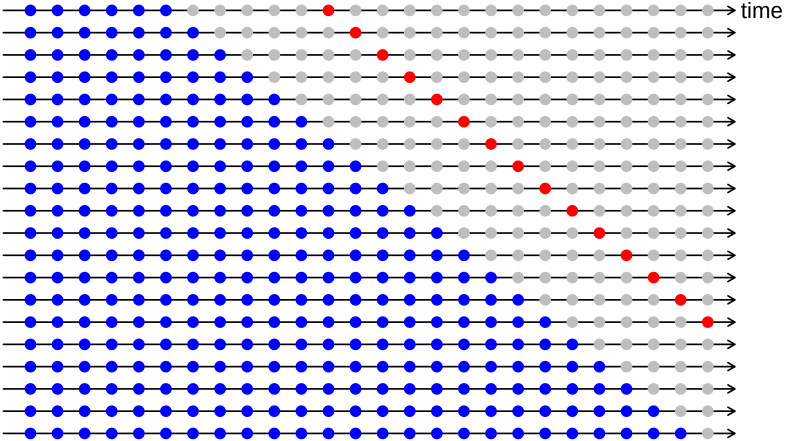

Also known as “Evaluation on a rolling forecast origin”

Some connections

Cross-sectional data

- Minimizing the AIC is asymptotically equivalent to minimizing MSE via leave-one-out cross-validation. (Stone, 1977).

Time series cross-validation

- Minimizing the AIC is asymptotically equivalent to minimizing MSE via one-step cross-validation. (Akaike, 1969, 1973).

Some connections

Cross-sectional data

- Minimizing the AIC is asymptotically equivalent to minimizing MSE via leave-one-out cross-validation. (Stone, 1977).

Time series cross-validation

- Minimizing the AIC is asymptotically equivalent to minimizing MSE via one-step cross-validation. (Akaike, 1969, 1973).

Time series cross-validation

Assume k is the minimum number of observations for a training set.

- Select observation $k + i$ for test set, and use observations at times $1, 2, \dots, k + i - 1$ to estimate model.
- Compute error on forecast for time $k + i$.
- Repeat for $i = 0, 1, \dots, T - k$ where T is total number of observations.
- Compute accuracy measure over all errors.

Example: Pharmaceutical sales

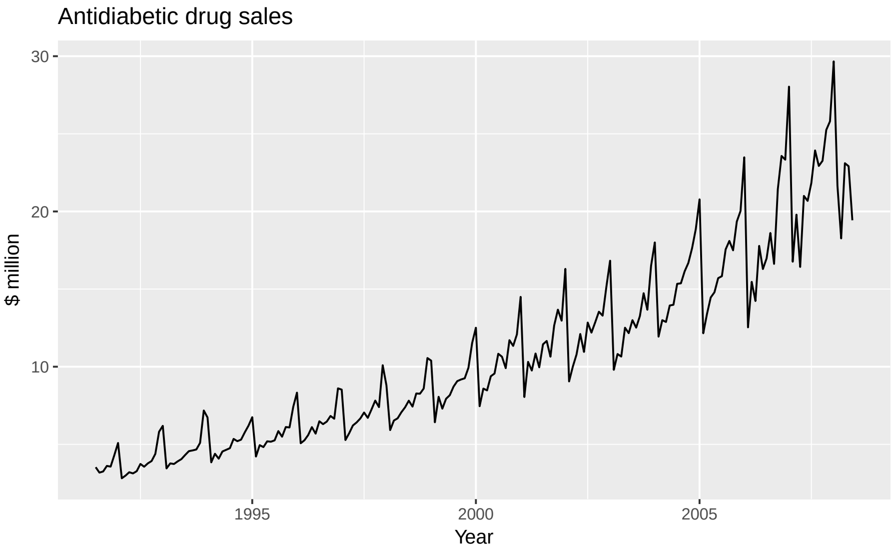

Example: Pharmaceutical sales

Which of these models is best?

- Linear model with trend and seasonal dummies applied to log data.
- ARIMA model applied to log data
- ETS model applied to original data

- Set $k = 48$ as minimum training set.
- Forecast 12 steps ahead based on data to time $k + i - 1$ for $i = 1, 2, \ldots, 156$.
- Compare MAE values for each forecast horizon.

Example: Pharmaceutical sales

Which of these models is best?

- Linear model with trend and seasonal dummies applied to log data.
- ARIMA model applied to log data
- ETS model applied to original data

- Set $k = 48$ as minimum training set.
- Forecast 12 steps ahead based on data to time $k + i - 1$ for $i = 1, 2, \ldots, 156$.
- Compare MAE values for each forecast horizon.

Example: Pharmaceutical sales

```
fc1 <- function(y, h) {
  fit <- tslm(y ~ trend + season)
  return(forecast(fit, h=h))
}
fc2 <- function(y, h) {
  fit <- auto.arima(y)
  return(forecast(fit, h=h))
}
fc3 <- function(y, h) {
  fit <- ets(y)
  return(forecast(fit, h=h))
}
e1 <- tsCV(a10, fc1, h=1)
e2 <- tsCV(a10, fc2, h=1)
e3 <- tsCV(a10, fc3, h=1)
mae1 <- mean(abs(e1))
mae2 <- mean(abs(e2))
mae3 <- mean(abs(e3))
```

- Repeat for each forecast horizon *h*.
- Inefficient because of re-fitting models

Example: Pharmaceutical sales

```
k <- 48
n <- length(a10)
mae1 <- mae2 <- mae3 <- matrix(NA,n-k-12,12)
for(i in 1:(n-k-12))
{
  xshort <- window(a10,end=1995+(5+i)/12)
  xnext <- window(a10,start=1995+(6+i)/12,end=1996+(5+i)/12)
  fit1 <- tslm(xshort ~ trend + season, lambda=0)
  fcast1 <- forecast(fit1,h=12)
  fit2 <- auto.arima(xshort,D=1, lambda=0)
  fcast2 <- forecast(fit2,h=12)
  fit3 <- ets(xshort)
  fcast3 <- forecast(fit3,h=12)
  mae1[i,] <- abs(fcast1[['mean']]-xnext)
  mae2[i,] <- abs(fcast2[['mean']]-xnext)
  mae3[i,] <- abs(fcast3[['mean']]-xnext)
}
```

Example: Pharmaceutical sales

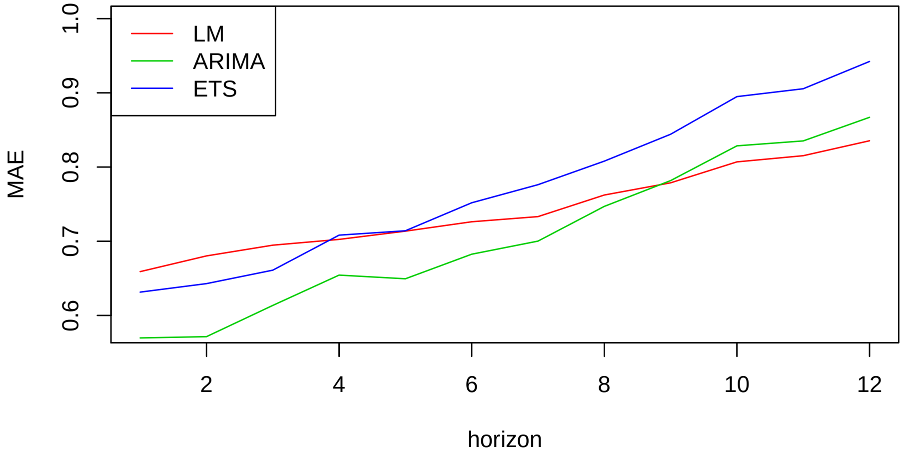

Variations on time series cross validation

- Keep training window of fixed length.

```
xshort <- window(a10,start=i+1/12,end=1995+(5+i)/12)
```

- Compute one-step forecasts in out-of-sample period.

```
for(i in 1:(n-k))
{
  xshort <- window(a10,end=1995+(5+i)/12)
  xlong <- window(a10,start=1995+(6+i)/12)
  fit2 <- auto.arima(xshort,D=1, lambda=0)
  fit2a <- Arima(xlong,model=fit2)
  fit3 <- ets(xshort)
  fit3a <- ets(xlong,model=fit3)
  mae2a[i,] <- abs(residuals(fit3a))
  mae3a[i,] <- abs(residuals(fit2a))
}
```

Outline

Lab Session 13