

Homework #5 - Fitting DLMs

Answer Key

Contents

Background	1
Question 1	1
Question 2	5
Question 3	9
Question 4	12
Question 5	16

Background

Here are the answers for the homework problems on fitting Dynamic Linear Models (DLMs). Begin by getting the data

```
## get S-R data; cols are:
## 1: brood yr (brood.yr)
## 2: number of spawners (Sp)
## 3: number of recruits (Rec)
## 4: PDO during first summer at sea (PDO.t2)
## 5: PDO during first winter at sea (PDO.t3)
load("KvichakSockeye.RData")
## head of data file
head(SRdata)

##   brood.yr   Sp   Rec PDO.t2 PDO.t3
## 1    1952 5970 17310  -0.61  -0.61
## 2    1953  320   520  -1.48  -2.66
## 3    1954  240   750  -2.05  -1.26
## 4    1955  250  1280   0.01   0.11
## 5    1956 9443 39036   0.86   0.37
## 6    1957 2843  4091  -0.25   0.29
```

Question 1

Begin by fitting a reduced form of Equation 15 that includes only a time-varying level (α_t) and observation error (v_t). That is,

$$\begin{aligned}\log(R_t) &= \alpha_t + \log(S_t) + v_t \\ \log(R_t/S_t) &= \alpha_t + v_t\end{aligned}\tag{1}$$

This model assumes no density-dependent survival in that the number of recruits is an ascending function of spawners. Plot the ts of α_t and note the AICc for this model. Also plot appropriate model diagnostics.

Answer

The stock-recruit model here is a random walk observed with error, which we have seen a lot in class. To see the equivalency, instead write the observation model as

$$y_t = x_t + v_t\tag{2}$$

where $y_t = \log(R_t/S_t)$ and $x_t = \alpha_t$. The process model is then

$$x_t = x_{t-1} + w_t\tag{3}$$

The first thing we need is to compute the response variable $y_t = \log(R_t/S_t)$.

```
## Time series of ln(R/S)
lnRS <- log(SRdata$Rec/SRdata$Sp)
dat <- matrix(lnRS, nrow=1)
## number of years of data--we'll need this later
TT <- length(lnRS)
```


Now we can set up the DLM as a level-only model (i.e., a random walk with observation error) and fit it with MARSS.

```
library(MARSS)
## MARSS model defn
## for process eqn
BB <- matrix(1)
UU <- matrix(0)
QQ <- matrix("q")
## for observation eqn
ZZ <- matrix(1)
AA <- matrix(0)
RR <- matrix("r")
## only need starting values for regr parameters
inits_list <- list(x0=matrix(1))
## list of model matrices & vectors
mod_list <- list(B=BB, U=UU, Q=QQ, Z=ZZ, A=AA, R=RR, tinitx=0)
## fit DLM
Q1 <- MARSS(dat, inits=inits_list, model=mod_list)
```

```

## Success! abstol and log-log tests passed at 22 iterations.
## Alert: conv.test.slope.tol is 0.5.
## Test with smaller values (<0.1) to ensure convergence.
##
## MARSS fit is
## Estimation method: kem
## Convergence test: conv.test.slope.tol = 0.5, abstol = 0.001
## Estimation converged in 22 iterations.
## Log-likelihood: -48.67502
## AIC: 103.35   AICc: 104.0559
##
##      Estimate
## R.r      0.271
## Q.q      0.318
## x0.x0    0.953
## Initial states (x0) defined at t=0
##
## Standard errors have not been calculated.
## Use MARSSparamCIs to compute CIs and bias estimates.
## plot the time-varying level
plot.ts(t(Q1$states), ylab=expression(alpha[italic(t)]))

```



```


## get AICc
Q1$AICc
## [1] 104.0559

```

And finally examine some diagnostic plots:


```
## get list of Kalman filter output
```

```
kf_out <- MARSS::MARSSkfss(Q1)
## forecast errors
innov <- kf_out$Innov
## Q-Q plot of forecast errors
qqnorm(t(innov), main="", pch=16, col="blue")
## add y=x line for easier interpretation
qqline(t(innov))
```



```
## plot ACF of innovations
acf(t(innov), lag.max=10, main="ACF for Q1 residuals")
```

ACF for Q1 residuals



The residuals seem to be reasonably well behaved in that there appear normal with no significant autocorrelation.

Question 2

Fit the full model specified by Equation 15. For this model, obtain the time series of α_t , which is an estimate of the stock productivity in the absence of density-dependent effects. How do these estimates of productivity compare to those from the previous question? Plot the ts of α_t and note the AICc for this model. Also plot appropriate model diagnostics. (*Hint*: If you don't want a parameter to vary with time, what does that say about its process variance?)

Answer

Now we need to fit a DLM with a time-varying level (intercept), but time-invariant slope. Begin by obtaining the time series of spawners to use as the covariate.

```
Sp <- matrix(SRdata$Sp, nrow=1) # /10000
```

Set up the MARSS model structure so α varies with time, but not β , which means $q = 0$ for β . This means that \mathbf{Q} should be

$$\mathbf{Q} = \begin{bmatrix} q_\alpha & 0 \\ 0 & 0 \end{bmatrix} \quad (4)$$


```

## number of regr coefs
m <- 2
## MARSS model defn
## for process eqn
B <- diag(m)          ## 2x2; Identity
U <- matrix(0,nrow=m,ncol=1)  ## 2x1; both elements = 0
Q <- matrix(list(0),m,m)    ## 2x2; all 0 for now
Q[1,1] <- "q_alpha"       ## 2x2; diag = (q1,q2)
## for observation eqn
Z <- array(NA, c(1,m,TT))  ## NxMxT; empty for now
Z[1,1,] <- rep(1,TT)      ## Nx1; 1's for intercept
Z[1,2,] <- Sp             ## Nx1; regr variable
A <- matrix(0)           ## 1x1; scalar = 0
R <- matrix("r")        ## 1x1; scalar = r
## only need starting values for regr parameters
inits_list <- list(x0=matrix(c(0, 0), nrow=m))
## list of model matrices & vectors
mod_list <- list(B=B, U=U, Q=Q, Z=Z, A=A, R=R)
## list of control params
con_list <- list(maxit=2000)
## fit DLM
Q2 <- MARSS(dat, inits=inits_list, model=mod_list, control=con_list)

## Success! abstol and log-log tests passed at 21 iterations.
## Alert: conv.test.slope.tol is 0.5.
## Test with smaller values (<0.1) to ensure convergence.
##
## MARSS fit is
## Estimation method: kem
## Convergence test: conv.test.slope.tol = 0.5, abstol = 0.001
## Estimation converged in 21 iterations.
## Log-likelihood: -48.5829
## AIC: 105.1658   AICc: 106.3779
##
##           Estimate
## R.r       2.68e-01
## Q.q_alpha 3.19e-01
## x0.X1     9.86e-01
## x0.X2     -8.07e-06
## Initial states (x0) defined at t=0
##
## Standard errors have not been calculated.
## Use MARSSparamCIs to compute CIs and bias estimates.

## plot the time-varying level
plot.ts(Q2$states[1,], ylab=expression(alpha[italic(t)]))

```



There does not appear to be any model support for a time-varying α .

Now let's check the AICc value.


```
## get AIC
Q2$AICc
## [1] 106.3779
```

The AICc value for this model is much lower for that in Q1.

Let's check out some diagnostic plots for the model in Q2. First we get the model innovations from our fitted MARSS object.


```
## get list of Kalman filter output
kf_out <- MARSS::MARSSkfss(Q2)
## forecast errors
innov <- kf_out$Innov

## Q-Q plot of forecast errors
qqnorm(t(innov), main="", pch=16, col="blue")
## add y=x line for easier interpretation
qqline(t(innov))
```



```
## plot ACF of innovations
acf(t(innov), lag.max=10, main="ACF for Q2 residuals")
```

ACF for Q2 residuals



The diagnostics indicate that there is some significant autocorrelation in the residuals at lags 1, 5 and 6. The autocorrelation at lag 1 is likely due to environmental factors whereas the autocorrelation at lags 5 & 6 is perhaps a reflection of the dominant age classes of these fish.

Question 3

Fit the model specified by Equation 16 with the summer PDO index as the covariate (PDO.t2). What is the mean level of productivity? Plot the ts of δ_t and note the AICc for this model. Also plot appropriate model diagnostics.

Answer

Now we need to fit a DLM so α and β are time-invariant, but δ varies by year. This means that \mathbf{Q} should be

$$\mathbf{Q} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & q_\delta \end{bmatrix} \quad (5)$$

```
## number of regr coefs
m <- 3
## MARSS model defn
## for process eqn
B <- diag(m) ## 2x2; Identity
U <- matrix(0,nrow=m,ncol=1)
Q <- matrix(list(0),m,m)
## place delta last--it's the only one to time-vary
Q[3,3]=("q_delta")
## for observation eqn
Z <- array(NA, c(1,m,TT)) ## NxMxT; empty for now
Z[1,1,] <- rep(1,TT) ## 1's for intercept
Z[1,2,] <- SRdata[,2] ## Sp regr variable
Z[1,3,] <- SRdata[,4] ## summer PDO regr variable
A <- matrix(0) ## 1x1; scalar = 0
R <- matrix("r") ## 1x1; scalar = r
## only need starting values for regr parameters
inits_list <- list(x0=matrix(c(0,0,0), nrow=m))
## list of model matrices & vectors
mod_list <- list(B=B, U=U, Q=Q, Z=Z, A=A, R=R)
## list of control params
con_list <- list(maxit=2000, allow.degen=TRUE)
## fit DLM
Q3 <- MARSS(dat, inits=inits_list, model=mod_list, control=con_list)

## Success! abstol and log-log tests passed at 616 iterations.
## Alert: conv.test.slope.tol is 0.5.
## Test with smaller values (<0.1) to ensure convergence.
##
## MARSS fit is
## Estimation method: kem
## Convergence test: conv.test.slope.tol = 0.5, abstol = 0.001
```

```


## Estimation converged in 616 iterations.
## Log-likelihood: -47.68255
## AIC: 105.3651   AICc: 107.2401
##
##           Estimate
## R.r       7.20e-01
## Q.q_delta 0.00e+00
## x0.X1     7.20e-01
## x0.X2     -1.72e-05
## x0.X3     4.52e-02
## Initial states (x0) defined at t=0
##
## Standard errors have not been calculated.
## Use MARSSparamCIs to compute CIs and bias estimates.

## mean productivity
mean(Q3$states[1,])

## [1] 0.7197999

## plot the time-varying effect of PDO
plot.ts(Q3$states[3,], ylab=expression(delta[italic(t)]))

```



Here, as in Q2, there does not appear to be any model support for a time-varying δ .

```

## get AIC
Q3$AICc


## [1] 107.2401

```

Again, this model is better than that in Q1, but not better than that for Q2.


Here are some diagnostic plots:

```
## get list of Kalman filter output
kf_out <- MARSS::MARSSkfss(Q3)
## forecast errors
innov <- kf_out$Innov
## Q-Q plot of forecast errors
qqnorm(t(innov), main="", pch=16, col="blue")
## add y=x line for easier interpretation
qqline(t(innov))
```



```
## plot ACF of innovations
acf(t(innov), lag.max=10, main="ACF for Q3 residuals")
```

ACF for Q3 residuals



There is some indication that our model is not adequately accounting for autocorrelation in the residuals (i.e., significant correlation at lag=1).

Question 4

Fit the model specified by Equation 16 with the winter PDO index as the covariate (PDO.t3). What is the mean level of productivity? Plot the ts of δ_t and note the AICc for this model. Also plot appropriate model diagnostics.

Again we need to fit a DLM so that α and β are time-invariant, but δ varies by year. As for Q3, \mathbf{Q} should be

$$\mathbf{Q} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & q_\delta \end{bmatrix} \quad (6)$$

```
## number of regr coefs
m <- 3
## MARSS model defn
## for process eqn
B <- diag(m) ## 2x2; Identity
U <- matrix(0,nrow=m,ncol=1)
Q <- matrix(list(0),m,m)
## place delta last--it's the only one to time-vary
Q[3,3]=("q_delta")
```

```


## for observation eqn
Z <- array(NA, c(1,m,TT)) ## NxMxT; empty for now
Z[1,1,] <- rep(1,TT) ## 1's for intercept
Z[1,2,] <- SRdata[,2] ## Sp regr variable
Z[1,3,] <- SRdata[,5] ## winter PDO regr variable
A <- matrix(0) ## 1x1; scalar = 0
R <- matrix("r") ## 1x1; scalar = r
## only need starting values for regr parameters
inits_list <- list(x0=matrix(c(0,0,0), nrow=m))
## list of model matrices & vectors
mod_list <- list(B=B, U=U, Q=Q, Z=Z, A=A, R=R)
## list of control params
con_list <- list(maxit=2000, allow.degen=TRUE)
## fit DLM
Q4 <- MARSS(dat, inits=inits_list, model=mod_list, control=con_list)

## Success! abstol and log-log tests passed at 608 iterations.
## Alert: conv.test.slope.tol is 0.5.
## Test with smaller values (<0.1) to ensure convergence.
##
## MARSS fit is
## Estimation method: kem
## Convergence test: conv.test.slope.tol = 0.5, abstol = 0.001
## Estimation converged in 608 iterations.
## Log-likelihood: -47.62552
## AIC: 105.251   AICc: 107.126
##
##           Estimate
## R.r       7.18e-01
## Q.q_delta 0.00e+00
## x0.X1     7.50e-01
## x0.X2     -2.03e-05
## x0.X3     6.45e-02
## Initial states (x0) defined at t=0
##
## Standard errors have not been calculated.
## Use MARSSparamCIs to compute CIs and bias estimates.
## mean productivity
mean(Q4$states[1,])

## [1] 0.74952

## plot the time-varying effect of PDO
plot.ts(Q4$states[3,], ylab=expression(delta[italic(t)]))

```




Again, it appears as though there is no data support for a time-varying effect of PDO.

```
## get AIC
Q4$AICc
## [1] 107.126
```

This model is not any better than that in Q2.

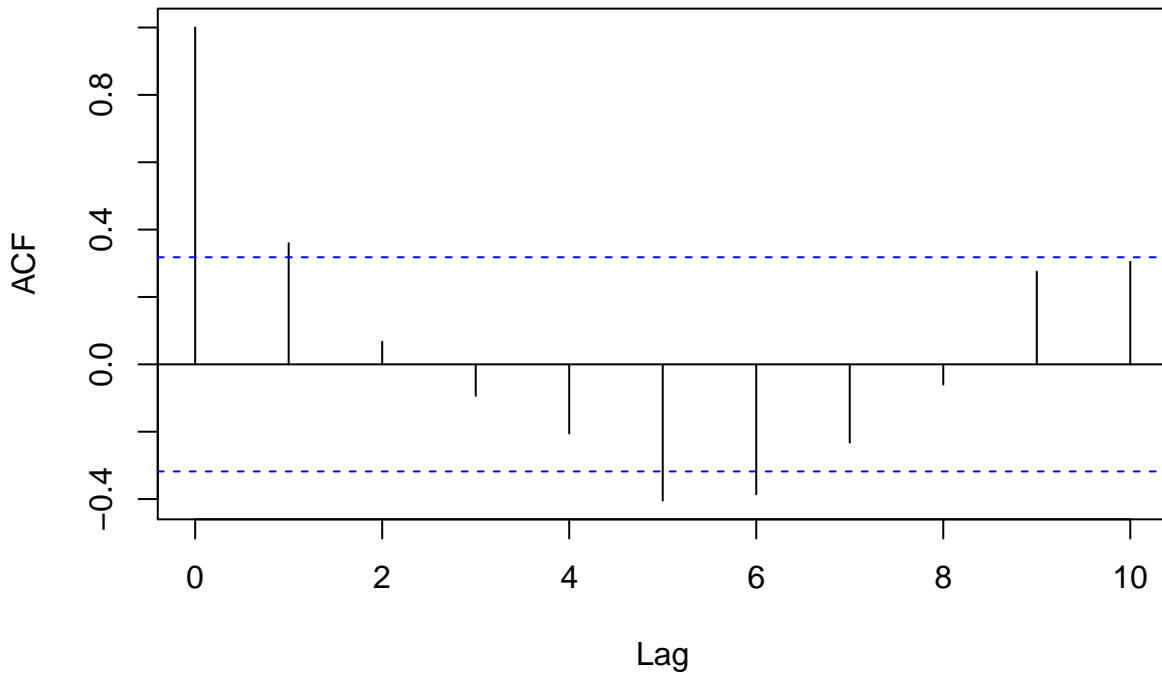
Here are some diagnostic plots:

```
## get list of Kalman filter output
kf_out <- MARSS::MARSSkfss(Q4)
## forecast errors
innov <- kf_out$Innov
## Q-Q plot of forecast errors
qqnorm(t(innov), main="", pch=16, col="blue")
## add y=x line for easier interpretation
qqline(t(innov))
```



```
## plot ACF of innovations
acf(t(innov), lag.max=10, main="ACF for Q4 residuals")
```

ACF for Q4 residuals



As in Q3 there is some indication that our model is not adequately accounting for autocorrelation in the residuals (i.e., significant correlation at lags 1 & 4-5).

Question 5

Based on AICc, which of the models above is the most parsimonious? Is it well behaved (*i.e.*, are the model assumptions met)? Plot the model forecasts for the best model. Is this a good forecast model?

Here is a table of AICc values for all 4 models.

```
tbl_aicc <- data.frame(model=paste0("Q",seq(4)),
                      AICc=round(c(Q1$AICc,Q2$AICc,Q3$AICc,Q4$AICc),1))
tbl_aicc
##   model  AICc
## 1    Q1 104.1
## 2    Q2 106.4
## 3    Q3 107.2
## 4    Q4 107.1
```

The model we fit in Q1 appears to have the lowest AIC, so let's use that for forecasting.

Here's how to obtain the time series of forecasts (and their SE) for the best model.

```
## get list of Kalman filter output
kf_out <- MARSS::MARSSkfss(Q1)
## forecasts of regr parameters; 2xT matrix
eta <- kf_out$xtt1
## predictor variable (1's only for the intercept)
Z <- array(1, c(1,1,TT)) ## NxMxT; empty for now
## ts of E(forecasts)
fore_mean <- vector()
for(t in 1:TT) {
  fore_mean[t] <- Z[,,t] %*% eta[,t,drop=F]
}
## variance of regr parameters; 1x2xT array
Phi <- kf_out$Vtt1
## obs variance; 1x1 matrix
R_est <- coef(Q1, type="matrix")$R
## ts of Var(forecasts)
fore_var <- vector()
for(t in 1:TT) {
  tZ <- matrix(Z[,,t],1,1) ## transpose of Z
  fore_var[t] <- Z[,,t] %*% Phi[,,t] %*% tZ + R_est
}
```

And now we can plot them.


```
fup <- fore_mean+2*sqrt(fore_var)
flo <- fore_mean-2*sqrt(fore_var)
par(mar=c(4,4,0.1,0), oma=c(0,0,2,0.5))
ylims=c(min(flo),max(fup))
plot(SRdata$brood.yr, t(dat), type="p", pch=16, ylim=ylims,
```



```

col="blue", xlab="Year", ylab="ln(R/S)" ##, xaxt="n")
lines(SRdata$brood.yr, fore_mean, type="l", xaxt="n", ylab="", lwd=3)
lines(SRdata$brood.yr, fup)
lines(SRdata$brood.yr, flo)

```



Overall, the accuracy of the forecasts is a bit suspect as many observations are at 0.5+ log-units from the forecast. Although most of the observed $\ln(R/S)$ fell within the 95% forecast intervals, the intervals themselves are relatively large and span a large range of R/S .