

STATS 531
Winter, 2018
Midterm Exam

We investigate some data from neurophysiology. Neurons communicate by generating pulses of electrical charge known as *firing events*. An electrode implanted (painlessly) into a monkey's brain records a sequence of firing events for an individual neuron cell. Suppose the firing times are F_1, F_2, \dots, F_{N+1} , measured in milliseconds (1ms is 10^{-3} s). We take as our time series $x_n^* = F_{n+1} - F_n$ with $n = 1, \dots, N$. This is the series of times intervals between firing events. The data, with $N = 415$, are plotted in Fig. 1. We wish to model $x_{1:N}^*$ in order to quantify the behavior of the neuron, to later compare it with other neurons and investigate the effects of experimental treatments.




Figure 1: Time series $x_{1:N}^*$ of time (in milliseconds) between subsequent firings of a monkey neuron.

SECTION A. We start with a linear time series analysis of $x_{1:N}^*$. The sample autocorrelation function of $x_{1:N}^*$ is shown in Fig. 2.

A1. [2 pts] What does Fig. 2 suggest to you about suitable ARMA models to model $x_{1:N}^*$, and why?

There is no single clear-cut answer. The sample ACF is within the dashed lines after lag 1, but shows some indication of decreasing like a damped oscillation (an AR(2) property). The fairly rapid decay of the sample ACF to values close to zero is consistent with mean stationarity but doesn't give much evidence for or against covariance stationarity. Anything between AR(1) and ARMA(2,2) can be defended by inspecting Fig. 2.

A2. [2 pts] Another way to select a model is by comparing AIC values. A table of AIC values is shown in Table 1. What ARMA model(s) would you consider based on this table, and why?

ARMA(2,2) has the lowest AIC, so is favored by this criterion. ARMA(2,1) and




Figure 2: Sample autocorrelation function of $x_{1:N}^*$.

	MA0	MA1	MA2	MA3
AR0	3966.0	3961.5	3962.7	3964.7
AR1	3961.1	3962.6	3964.6	3966.6
AR2	3962.7	3960.5	3959.8	3961.7
AR3	3964.6	3965.5	3962.6	3968.4

Table 1: AIC values from fitting ARMA(p,q) models to $x_{1:N}^*$.

ARMA(1,0) are simpler models which also have promising AIC values. Other considerations, such as proximity to reducibility, non-invertibility or non-causality, may also play a role in model selection. We will also see, later in the exam, that in fact none of these models are particularly appropriate.

A3. [2 pts] Find the log likelihood of an ARMA(2,1) model, and explain your calculation.

AIC = $-2\lambda + 2k$, where λ is the maximized log likelihood and k is the number of parameters. Here, $k = 5$ since the parameter vector is $(\phi_1, \phi_2, \theta_1, \mu, \sigma^2)$, so

$$\lambda = \frac{3960.5 - 10}{-2} = -1975.25$$

.

A4. [3 pts] Does the table of AIC values contain any evidence for or against the claim that the likelihood is correctly calculated and maximized? Explain.

The table is inconsistent — adding a parameter can only increase the maximized log-likelihood, i.e. the AIC can only increase by ≤ 2 . Compare ARMA(3,3) to ARMA(3,2). This can only come about by imperfect likelihood calculation and/or maximization. We may suspect a problem with likelihood maximization.

SECTION B. Fitting an ARMA(2,2) model gives the following *R* output.

```
arma22 <- arima(x,order=c(2,0,2)) ; arma22

##
## Call:
## arima(x = x, order = c(2, 0, 2))
```

```
##
## Coefficients:
##      ar1      ar2      ma1      ma2  intercept
##      1.6009 -0.6445 -1.4982  0.5219   26.4163
## s.e.  0.1886  0.1839  0.2104  0.2094    0.7954
##
## sigma^2 estimated as 791.7:  log likelihood = -1973.88,  aic = 3959.76
```

B1. [4 pts]. Write out the fitted model, carefully stating all the assumptions behind the model used by R to generate this output.

The fitted model is

$$X_n = 26.4 + 1.60(X_{n-1} - 26.4) - 0.64(X_{n-2} - 26.4) + \epsilon_n - 1.50\epsilon_{n-1} + 0.5\epsilon_{n-2}$$

where $\{\epsilon_n\}$ is white noise with standard deviation **28.1**. The likelihood calculation also assumes that $\{\epsilon_n\}$ is Gaussian, i.e. an independent sequence with $\epsilon_n \sim N(0, 28.1^2)$.




Figure 3: Sample autocorrelation function of the residuals from fitting an ARMA(2,2) model to $x_{1:N}^*$.

B2. [3 pts] Fig. 3 shows the ACF of the residuals from fitting an ARMA(2,2) model. Comment on which modeling assumptions are investigated by this figure, and whether they are consistent with the data.

Fig. 3 supports the assumptions that the driving noise process is uncorrelated and has no trend. It does not reveal anything about the assumptions that the driving noise is (i) constant variance or (ii) Gaussian.

The roots of the AR and the MA polynomials for the fitted ARMA(2,2) model are computed as follows:

```
AR_roots <- polyroot(c(1,-coef(arma22)[c("ar1","ar2")])) ; AR_roots
## [1] 1.24194+0.095596i 1.24194-0.095596i
MA_roots <- polyroot(c(1,coef(arma22)[c("ma1","ma2")])) ; MA_roots
## [1] 1.055662-0i 1.814977+0i
```

B3. [3 pts] Is there evidence for parameter redundancy? Do these roots raise any other potential concerns?

The AR roots are not very close to the MA roots, so there is no strong suggestion of parameter redundancy in the fitted model. One of the MA roots (1.055) is close to the unit circle, so the fitted model is close to non-invertibility. This could potentially cause numerical instability. It also suggests that the standard errors based on Fisher information may be unreliable, as we discovered in the class notes.

Simulations from the fitted ARMA(2,2) model were computed as follows:

```
arma22<-arima(x,order=c(2,0,2))
Nt<-length(x)
sim<-rep(0,Nt)
set.seed(1)
w<-rnorm(Nt,m=0,sd=sqrt(arma22$sigma2))
for(nt in 3:Nt){
  sim[nt]<-arma22$coef["ar1"]*sim[nt-1]+arma22$coef["ar2"]*sim[nt-2]+
    arma22$coef["ma1"]*w[nt-1]+arma22$coef["ma2"]*w[nt-2]+w[nt]
}
sim<-sim+arma22$coef["intercept"]
```

B4. [2 pts] Sample simulation output is shown in Fig. 4. What does a comparison of Fig. 4 with Fig. 1 say about ARMA modeling of $x_{1:N}^*$?

Almost all well-considered and comprehensible answers are acceptable here. The plots look quite different. Fig. 1 is always positive, and appears to have some regularity to the peaks. Fig. 4 has less pronounced peaks, is more symmetric about its mean, and occasionally becomes negative. The data do not resemble a Gaussian ARMA(2,2) process.




Figure 4: A simulation from the fitted ARMA(2,2) model

B5. [2 pts] Is the random process generated in B4 and plotted in Fig. 4 formally stationary? Answer yes or no, and explain.

Not quite. The initial values contradict having a constant variance (the initial variance is zero). However, the process is asymptotically stationary. One could make the simulations effectively stationary by throwing away some number (say 100) values at the start of the simulation.

SECTION C. We now investigate a logarithmic transformation of the data. Let $z_{1:N}^*$ be the \log_{10} transformation of $x_{1:N}^*$, so $z_n^* = \log_{10}(x_n^*)$ for $n \in 1:N$. Below is the R output from fitting an ARMA(2,2) model to $z_{1:N}^*$.

```
##
## Call:
## arima(x = z, order = c(2, 0, 2))
##
## Coefficients:
##      ar1      ar2      ma1      ma2  intercept
##      1.6974 -0.7249 -1.4647  0.4647      1.2925
## s.e.  0.0740  0.0718  0.0941  0.0939      0.0021
##
## sigma^2 estimated as 0.07637:  log likelihood = -56.79,  aic = 125.59
```




Figure 5: (a) Time plot of $z_{1:N}^*$. (b) Sample autocorrelation function of $z_{1:N}^*$.

C1. [2 pts] Is there any indication from Fig. 5 and the R fitted model output in Sections B and C that ARMA modeling is more successful after a log transformation? or less? Explain.

The model printouts do not directly tell us whether the transformation is appropriate: in particular, the AIC values are not directly comparable, though in principle one can figure out how to transform the likelihood appropriately to make them comparable. The sample ACF is also tricky to interpret as evidence for the success of the transformation. The additional dependence appearing in the sample ACF of $z_{1:N}^*$ is not a bad thing—in fact, it may be a clue that the log transformation is the right scale to see linear dependence in the data as measured by covariance. The time plot looks more symmetric on the transformed scale, and symmetry is a property of Gaussian ARMA models so this is encouraging.

C2. [2 pts] What do Figs. 3, 6 and 7 indicate about the success of the log transform?

Figs. 3 and 6 show that the residuals appear uncorrelated with or without the log transform. Fig. 7 shows that the residuals look much more like Gaussian white noise




Figure 6: Sample autocorrelation function of the residuals from fitting an ARMA(2,2) model to $z_{1:N}^*$.




Figure 7: (a) Residuals from fitting an ARMA(2,2) model to $x_{1:N}^*$. (b) Residuals from fitting an ARMA(2,2) model to the log transformed data, $z_{1:N}^*$.

after the transform. In particular, the residuals are asymmetric—with a long right tail—for Fig. 7(a), which is inconsistent with a Gaussian model. Since the models are fitted using Gaussian maximum likelihood, this is a good thing. We could confirm this finding by a normal quantile plot.

C3. [3 pts] Fig. 8 shows the smoothed periodogram of $z_{1:N}^*$. Find the frequency and period corresponding to the peak in the periodogram. Your answer should include the units of these quantities. Describe briefly what this peak leads you to conclude about how this monkey neuron behaves.

The peak at frequency approx. 0.06 cycles/firing event (or period approx. 16.7 firing events) is due to the neuron firing in a burst (several short inter-event times) and then being less active (long inter-event times) with a characteristic period of approx. 17 firing events.




Figure 8: Smoothed periodogram of $z_{1:N}^*$.