STATS 531
Winter, 2018
Midterm Exam

Name: ________________________ **UMID #:** ________________________

There are 3 sections (A, B and C) containing a total of 30 points. Points will be awarded for clearly explained and accurate answers.

Only pens and/or pencils should be out of your bag for the duration of the exam. You may not use access any electronic device, notes, or books during the exam.

Section	Points	Score
A	9	
B	14	
C	7	
Total	30	

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 $F_1, F_2, \ldots, F_{N+1}$, measured in milliseconds (1ms is 10^{-3}s). We take as our time series $x_n^* = F_{n+1} - F_n$ with $n = 1, \ldots, N$. 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.

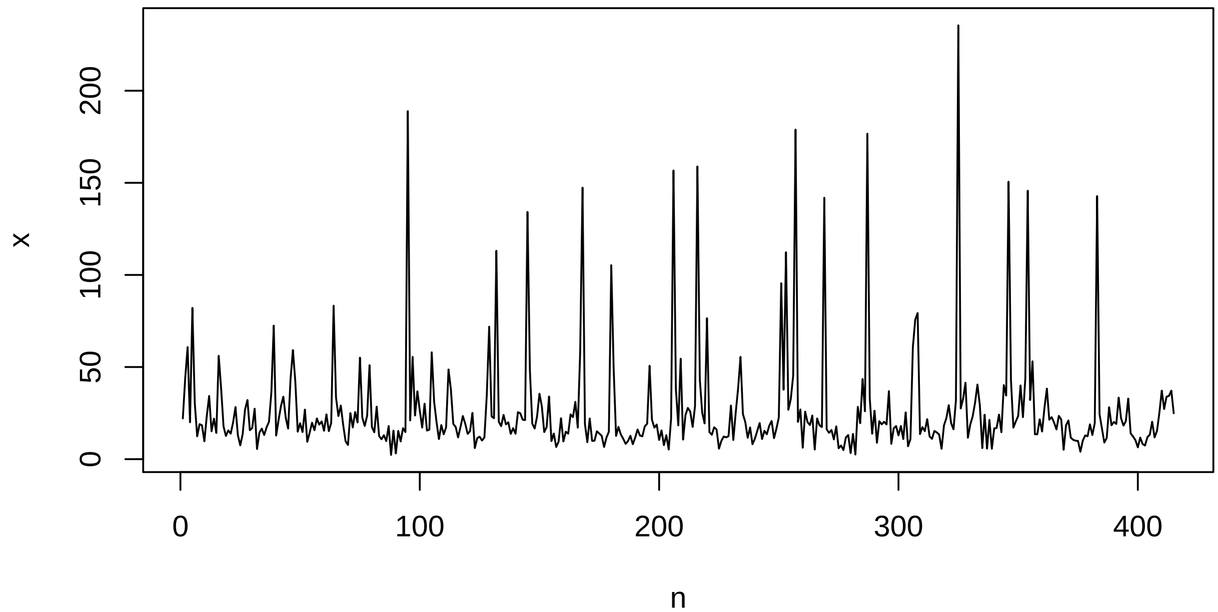

Figure 1: Time series $x_{1:N}^*$ of time (in milliseconds) between subsequent firings of a monkey neuron.

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.

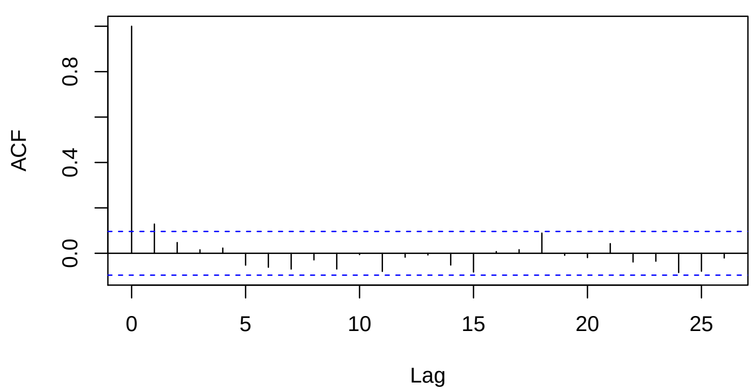

Figure 2: Sample autocorrelation function of $x_{1:N}^*$.

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

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?

	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 $x^*_{1:N}$.

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

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.

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.

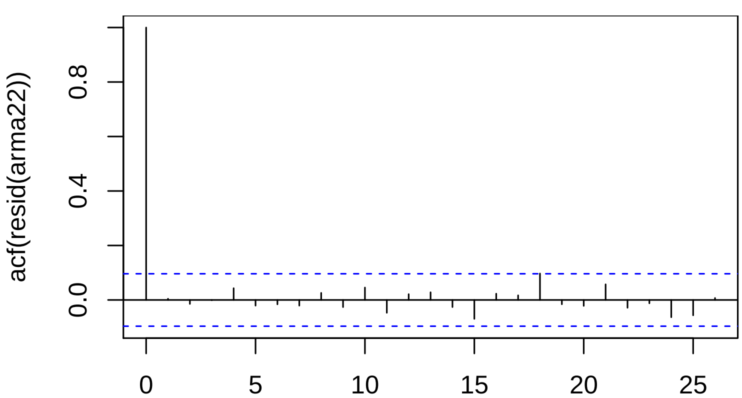

Figure 3: Sample autocorrelation function of the residuals from fitting an ARMA(2,2) model to $x^*_{1:N}$.

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.

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?

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 $x^*_{1:N}$?

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

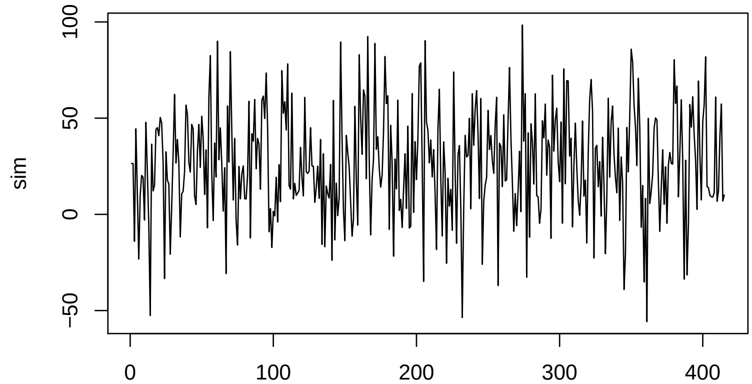

Figure 4: A simulation from the fitted ARMA(2,2) model

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
```

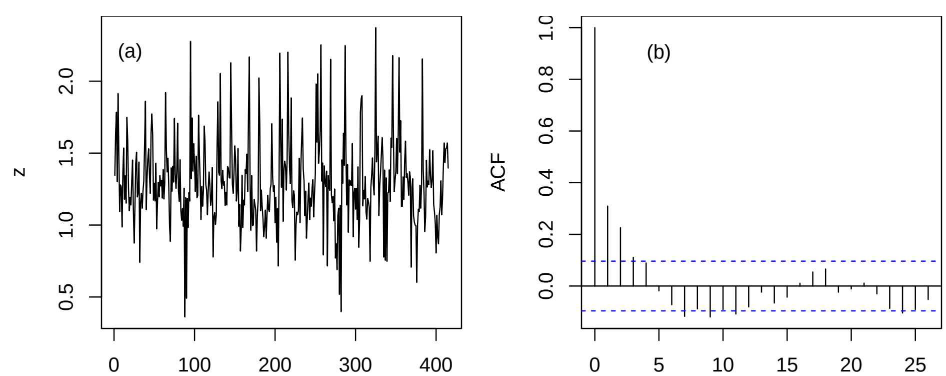

Figure 5: (a) Time plot of $z^*_{1:N}$. (b) Sample autocorrelation function of $z^*_{1:N}$.

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.

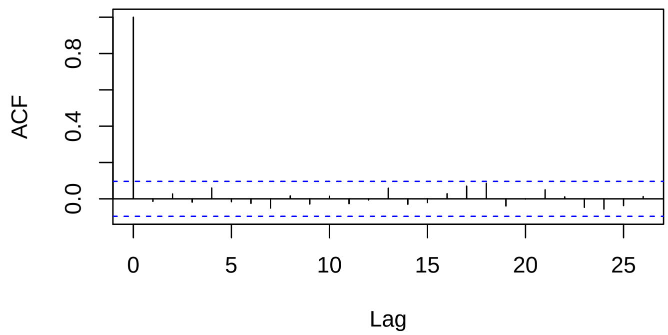

Figure 6: Sample autocorrelation function of the residuals from fitting an ARMA(2,2) model to $z^*_{1:N}$.

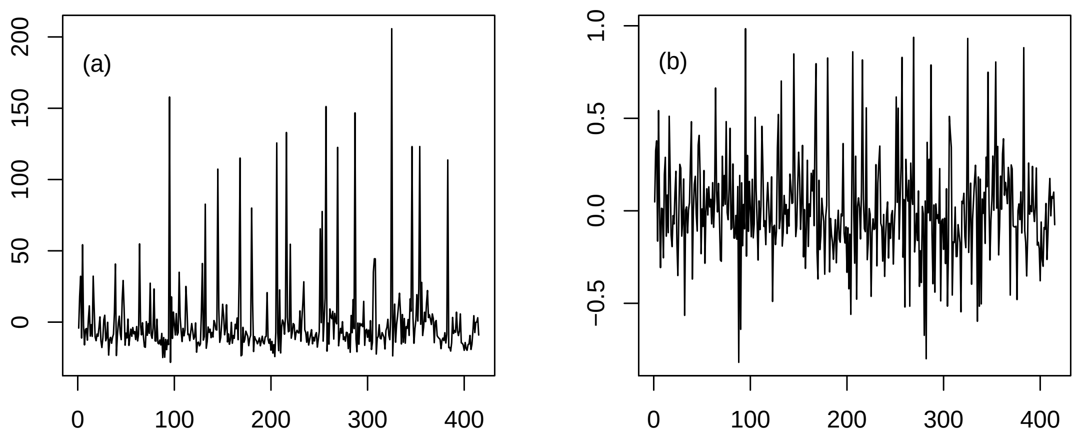

Figure 7: (a) Residuals from fitting an ARMA(2,2) model to $x^*_{1:N}$. (b) Residuals from fitting an ARMA(2,2) model to the log transformed data, $z^*_{1:N}$.

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

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.

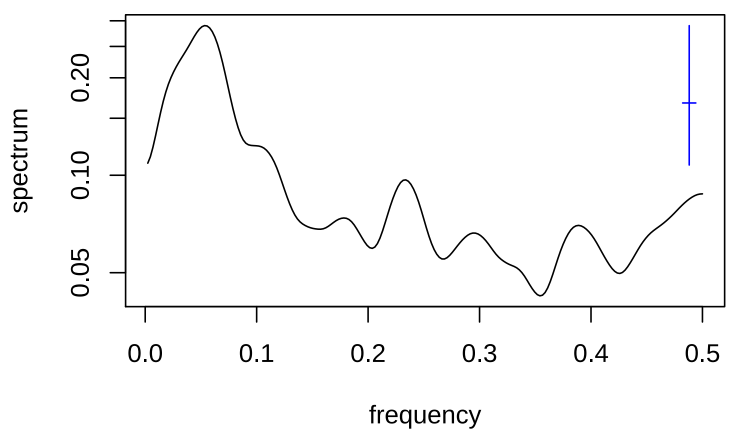

Figure 8: Smoothed periodogram of $z^*_{1:N}$.