STAT 720 sp 2019 hw 3

due on Wednesday, March 20th, 2019

Please make use of my R package `tscourse` to complete the homework. You can install the package with the following commands:

```
library(devtools)
devtools::install_github("gregorkb/tscourse")
library(tscourse)
```

This pulls the package from where it resides in my github repository. You will first need to install the `devtools` package using `install.packages("devtools")`.

1. Do problems 5.1, 5.2, and 5.4 of B&D Intro.

2. Let $\{X_t, t \in \mathbb{Z}\}$ be the $\text{ARMA}(1,1)$ times series defined by

$$X_t - \phi X_{t-1} = Z_t + \theta Z_{t-1},$$

where $\{Z_t, t \in \mathbb{Z}\}$ is $\text{WN}(0, \sigma^2)$ and $\theta + \phi \neq 0$ and $|\phi| \neq 1$.

(a) Give an expression for the spectral density of the time series in terms of ϕ, θ, and σ^2.

(b) Choose some values of ϕ, θ, and σ^2 and generate a realization from the model of length $n = 200$. Then compute from the data the periodogram ordinates at the Fourier frequencies

$$\frac{k}{n} \cdot 2\pi, \quad k = -\left\lfloor \frac{n-1}{2} \right\rfloor, \ldots, \left\lfloor \frac{n}{2} \right\rfloor$$

and plot the ordinates, scaled by $1/(2\pi)$, against the frequencies with the true spectral density function overlaid. Note: Since the periodogram and the spectral density are even functions, it is only necessary to consider their values over the range $[0, \pi]$. My plot looks like the following (the spectral density will change depending on what values you choose for ϕ, θ, and σ^2):

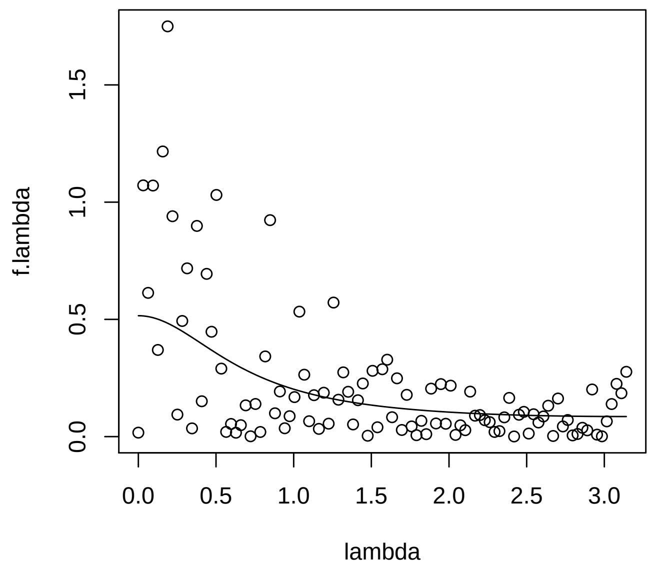

(c) Try smoothing the periodogram ordinates using the function `ksmooth()` under default settings. This takes local averages, much like the way we previously estimated trends. Add to the plot from part (b) the curve resulting from locally averaging the periodogram. My plot looks like the following:

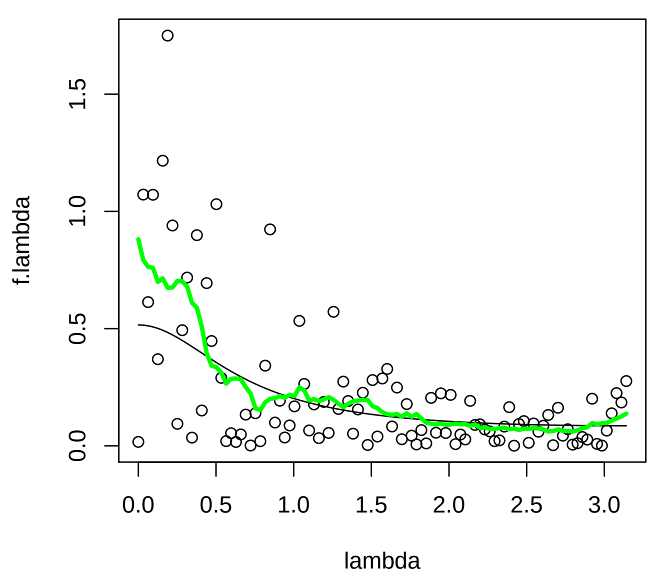

(d) Compute the maximum likelihood estimators of ϕ and θ based on the data. Then plug these estimates into the expression from part (a) to obtain a maximum-likelihood-based estimator of the spectral density. Add to the plot from part (b) the curve corresponding to this estimator of the spectral density. My plot looks like the following:

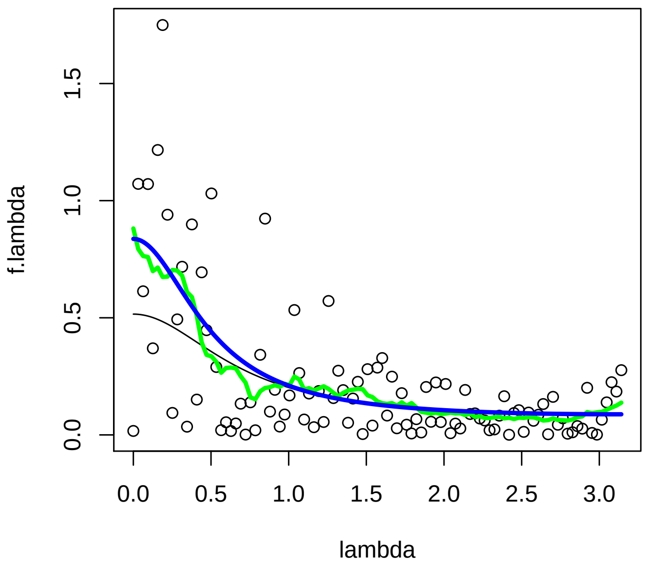

(e) Now compute a lag-window estimator of the spectral density. In particular, compute the estimator

$$\hat{f}(\lambda) = \frac{1}{2\pi} \sum_{h=-L}^{L} w(h/L)\hat{\gamma}_n(h)e^{-\iota h\lambda}, \quad -\pi < \lambda \leq \pi,$$

using the lag-window function given by

$$w(x) = \begin{cases} 1 - 6x^2 + 6|x|^3, & |x| < 1/2 \\ 2(1-|x|)^3, & 1/2 \leq |x| \leq 1 \\ 0, & |x| > 1, \end{cases}$$

which is the Parzen window. Compute this estimator using $L = \lfloor \sqrt{n} \rfloor$, and add to the plot from part (b) the resulting curve. *I you wish, you may submit one single plot for parts* (b)–(e). My plot looks like the following:

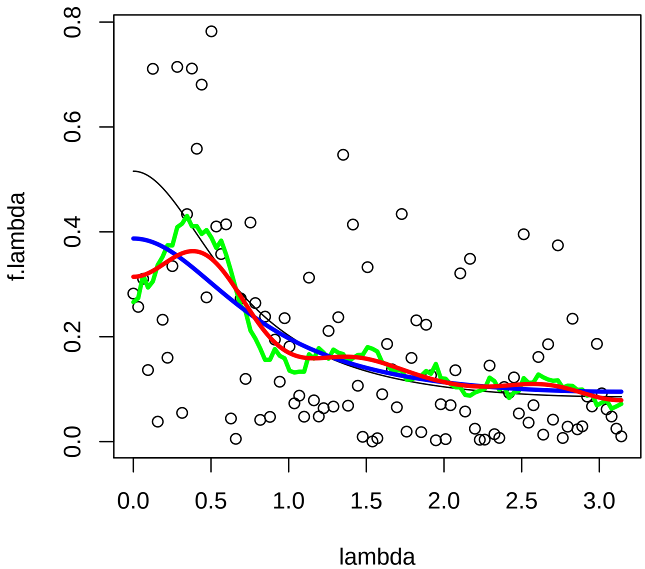

(f) According to Theorem 10.3.2 of B&D Theory, under some conditions, the periodogram ordinates at any finite number of frequencies $\lambda_1, \ldots, \lambda_m$ should asymptotically behave as independent exponentially distributed random variables with means $2\pi f(\lambda_1), \ldots, 2\pi f(\lambda_m)$, respectively, as $n \to \infty$. For the dataset generated in part (b), standardize the periodogram ordinates at the Fourier frequencies according to

$$I_n(\lambda_j)/(2\pi f(\lambda_j)), \quad j = 1, \ldots, n,$$

where $\lambda_1, \ldots, \lambda_n$ are the Fourier frequencies. Then make a quantile-quantile plot of these values to see if they follow the exponential distribution with mean 1. Note that you only need to consider the frequencies in the interval $[0, \pi]$. My plot looks like this:

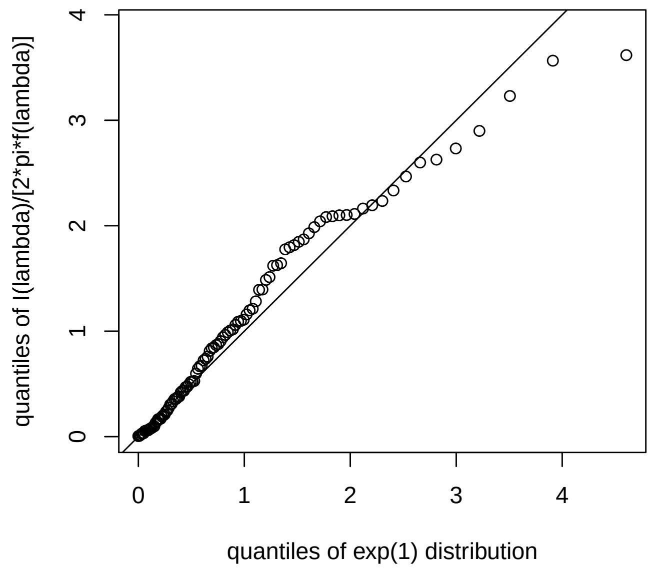

(g) Since

$$f(0) = \frac{1}{2\pi} \sum_{h=-\infty}^{\infty} \gamma(h) e^{-\iota h(0)} = \frac{1}{2\pi} \sum_{h=-\infty}^{\infty} \gamma(h),$$

we have

$$\sqrt{n}(\bar{X} - \mu) \to \text{Normal}(0, 2\pi f(0)) \text{ in distribution}$$

as $n \to \infty$, where $\bar{X} = n^{-1}(X_1, \dots, X_n)$ and $\mu = \mathbb{E}\bar{X}$. Run the following simulation: For $n = 200$, generate 500 realizations of length n from a causal invertible $\text{ARMA}(1,1)$ process with some mean $\mu \neq 0$. For each of the 500 datasets, compute the lag-window estimator

$$\hat{f}(0) = \frac{1}{2\pi} \sum_{h=-L}^{L} w(h/L)\hat{\gamma}_n(h)$$

of $f(0)$, with $w(\cdot)$ and L chosen as in part (e), and use it to construct a 95% confidence interval for μ. Report the coverage of your interval over the 500 simulated datasets and provide your code.

3. Let $\{X_t, t \in \mathbb{Z}\}$ be the causal $\text{AR}(p)$ time series given by

$$X_t - \phi_1 X_{t-1} - \cdots - \phi_p X_{t-p} = Z_t,$$

where $\{Z_t, t \in \mathbb{Z}\}$ is $\text{WN}(0, \sigma^2)$, and consider the sequence of Yule-Walker estimators

$$\hat{\boldsymbol{\phi}}_k = (\hat{\phi}_{k,1}, \dots, \hat{\phi}_{k,k})^T = \hat{\boldsymbol{\Gamma}}_k^{-1} \hat{\boldsymbol{\gamma}}_k,$$

where $\hat{\boldsymbol{\Gamma}}_k = (\hat{\gamma}_n(i-j))_{1 \le i,j \le k}$ and $\hat{\boldsymbol{\gamma}}_k = (\hat{\gamma}_n(1), \dots, \hat{\gamma}_n(k))^T$ for $k = 1, \dots, n$, where $\hat{\gamma}_n(\cdot)$ is the sample autocovariance function from a realization $X_1, \dots, X_n$ of the time series.

(a) The partial autocorrelation function (pacf) at lag k is the value $\hat{\phi}_{k,k}$ for $k = 1, 2, \dots$ We often examine a plot of the pacf to determine the appropriate order of an autoregressive model. Make a plot of the pacf at lags $k = 1, \dots, 20$ based on the Lake Huron data from the R data set `LakeHuron` and add to the plot two horizontal lines at the heights $-n^{-1/2}1.96$ and $n^{-1/2}1.96$. Note that you can extract the values $\hat{\phi}_{1,1}, \dots, \hat{\phi}_{20,20}$ from the output of the Durbin-Levinson algorithm when it is run using the sample autocovariance function. My plot looks like this:

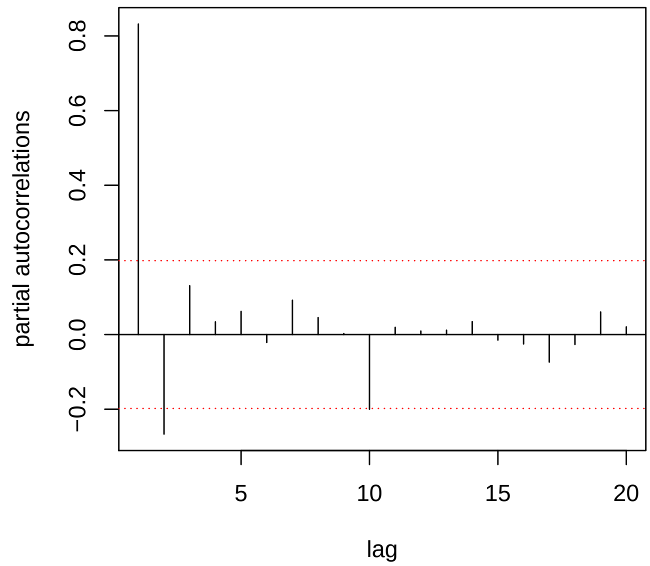

(b) A convention is to find the smallest lag at which the partial autocorrelations fall outside of the interval $(-n^{-1/2}1.96, n^{-1/2}1.96)$ and to fit an autoregressive model of this order to the data. The purpose of this question is for us to understand why this makes sense.

From Theorem 8.1.2 of B&D Theory, we have, for any $k > p$,

$$\sqrt{n}(\hat{\boldsymbol{\phi}}_k - \boldsymbol{\phi}_k) \to \text{Normal}(0, \boldsymbol{\Gamma}_k^{-1}\sigma^2) \quad \text{in distribution}$$

as $n \to \infty$, where $\boldsymbol{\Gamma}_k = (\gamma(i-j))_{1 \le i,j \le k}$ and $\boldsymbol{\phi}_k = (\phi_1, \dots, \phi_p, \mathbf{0}_{k-p}^T)^T$, where $\mathbf{0}_{k-p}$ is the $(k-p) \times 1$ vector with all entries equal to 0, and, in particular,

$$\sqrt{n}\hat{\phi}_{k,k} \to N(0,1) \quad \text{in distribution}$$

as $n \to \infty$.

i. Show that for any $k > p$, entry (k,k) of the matrix $\boldsymbol{\Gamma}_k^{-1}$ is equal to $(\gamma(0) - \boldsymbol{\phi}_p^T \boldsymbol{\gamma}_p)^{-1}$. *Hint: Consider inverting the matrix*

$$\begin{bmatrix} \boldsymbol{\Gamma}_{k-1} & \tilde{\boldsymbol{\gamma}}_{k-1} \\ \tilde{\boldsymbol{\gamma}}_{k-1}^T & \gamma(0) \end{bmatrix},$$

where $\tilde{\boldsymbol{\gamma}}_{k-1} = (\gamma(k-1), \dots, \gamma(1))^T$, *with a block-matrix inversion formula.*

ii. Show that $\gamma(0) - \boldsymbol{\phi}_p^T \boldsymbol{\gamma}_p = \sigma^2$, so that entry (k, k) of the matrix $\boldsymbol{\Gamma}_k^{-1}\sigma^2$ is equal to 1 for any $k > p$.

(c) Now explain why it is reasonable to follow the convention described in part (b) for choosing the order of autoregressive model to fit to the data.

(d) Make a choice for the order of autoregressive model appropriate for the Lake Huron data and compute the Yule-Walker estimators of the autoregressive parameters.

(e) Compute the maximum likelihood estimators of ϕ_1 and ϕ_2 and compare these values to the Yule-Walker estimates.

(f) Use your fitted model to obtain one-step-ahead predictions for $X_1, \ldots, X_n, X_{n+1}$ for the the Lake Huron time series. Make a plot of the original time series with the one-step-ahead predictions overlaid. My plot looks like this: