

Problem Set 4

Advanced Forecasting, UNCC, 2019

Ercan Karadas

Due: February 20, 2019

Exercise 1

Daily electricity demand for Victoria, Australia, during 2014 is contained in `elecdaily`. The data for the first 20 days can be obtained as follows.

```
daily20 <- head(elecdaily,20)
```

- Plot the data and find the regression model for Demand with temperature as an explanatory variable. Why is there a positive relationship?
- Produce a residual plot. Is the model adequate? Are there any outliers or influential observations?
- Use the model to forecast the electricity demand that you would expect for the next day if the maximum temperature was 15° and compare it with the forecast if the with maximum temperature was 35° . Do you believe these forecasts?
- Give prediction intervals for your forecasts. The following R code will get you started:

```
autoplot(daily20, facets=TRUE)
daily20 %>%
  as.data.frame() %>%
  ggplot(aes(x=Temperature, y=Demand)) +
  geom_point() +
  geom_smooth(method="lm", se=FALSE)
fit <- tslm(Demand ~ Temperature, data=daily20)
checkresiduals(fit)
forecast(fit, newdata=data.frame(Temperature=c(15,35)))
```

- Plot Demand vs Temperature for all of the available data in `elecdaily`. What does this say about your model?

Exercise 2

Data set `mens400` contains the winning times (in seconds) for the men's 400 meters final in each Olympic Games from 1896 to 2016.

- Plot the winning time against the year. Describe the main features of the scatterplot.
- Fit a regression line to the data. Obviously the winning times have been decreasing, but at what *average* rate per year?
- Plot the residuals against the year. What does this indicate about the suitability of the fitted line?
- Predict the winning time for the men's 400 meters final in the 2020 Olympics. Give a prediction interval for your forecasts. What assumptions have you made in these calculations?

Exercise 3

Type `easter(ausbeer)` and interpret what you see.

Exercise 4

An elasticity coefficient is the ratio of the percentage change in the forecast variable (y) to the percentage change in the predictor variable (x). Mathematically, the elasticity is defined as $(dy/dx) \times (x/y)$. Consider the log-log model,

$$\log y = \beta_0 + \beta_1 \log x + \varepsilon.$$

Express y as a function of x and show that the coefficient β_1 is the elasticity coefficient.

Exercise 5

The data set `fancy` concerns the monthly sales figures of a shop which opened in January 1987 and sells gifts, souvenirs, and novelties. The shop is situated on the wharf at a beach resort town in Queensland, Australia. The sales volume varies with the seasonal population of tourists. There is a large influx of visitors to the town at Christmas and for the local surfing festival, held every March since 1988. Over time, the shop has expanded its premises, range of products, and staff.

- Produce a time plot of the data and describe the patterns in the graph. Identify any unusual or unexpected fluctuations in the time series.
- Explain why it is necessary to take logarithms of these data before fitting a model.
- Use R to fit a regression model to the logarithms of these sales data with a linear trend, seasonal dummies and a “surfing festival” dummy variable.
- Plot the residuals against time and against the fitted values. Do these plots reveal any problems with the model?
- Do boxplots of the residuals for each month. Does this reveal any problems with the model?
- What do the values of the coefficients tell you about each variable?
- What does the Breusch-Godfrey test tell you about your model?
- Regardless of your answers to the above questions, use your regression model to predict the monthly sales for 1994, 1995, and 1996. Produce prediction intervals for each of your forecasts.
- Transform your predictions and intervals to obtain predictions and intervals for the raw data.
- How could you improve these predictions by modifying the model?

Exercise 6

The `gasoline` series consists of weekly data for supplies of US finished motor gasoline product, from 2 February 1991 to 20 January 2017. The units are in “thousand barrels per day”. Consider only the data to the end of 2004.

- Fit a harmonic regression with trend to the data. Experiment with changing the number Fourier terms. Plot the observed gasoline and fitted values and comment on what you see.
- Select the appropriate number of Fourier terms to include by minimizing the AICc or CV value.

- c. Check the residuals of the final model using the `checkresiduals()` function. Even though the residuals fail the correlation tests, the results are probably not severe enough to make much difference to the forecasts and forecast intervals. (Note that the correlations are relatively small, even though they are significant.)
- d. To forecast using harmonic regression, you will need to generate the future values of the Fourier terms. This can be done as follows.

```
>> r > fc <- forecast(fit, newdata=fourier(x, K, h)) >>>
```

 where `fit` is the fitted model using `tslm`, `K` is the number of Fourier terms used in creating `fit`, and `h` is the forecast horizon required.

```
>> Forecast the next year of data.
```
- e. Plot the forecasts along with the actual data for 2005. What do you find?

Exercise 7

Data set `huron` gives the level of Lake Huron in feet from 1875 to 1972.

- a. Plot the data and comment on its features.
- b. Fit a linear regression and compare this to a piecewise linear trend model with a knot at 1915.
- c. Generate forecasts from these two models for the period upto 1980 and comment on these.

Exercise 8 (Optional)

Using matrix notation it was shown that if $\mathbf{y} = \mathbf{X}\beta + \varepsilon$, where $\mathbf{e}$ has mean $\mathbf{0}$ and variance matrix $\sigma^2\mathbf{I}$, the estimated coefficients are given by $\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$ and a forecast is given by $\hat{y} = \mathbf{x}^*\hat{\beta} = \mathbf{x}^*(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$ where $\mathbf{x}^*$ is a row vector containing the values of the regressors for the forecast (in the same format as $\mathbf{X}$), and the forecast variance is given by $\text{var}(\hat{y}) = \sigma^2 [1 + \mathbf{x}^*(\mathbf{X}'\mathbf{X})^{-1}(\mathbf{x}^*)']$.

Consider the simple time trend model where $y_t = \beta_0 + \beta_1 t$. Using the following results,

$$\sum_{t=1}^T t = \frac{1}{2}T(T+1), \quad \sum_{t=1}^T t^2 = \frac{1}{6}T(T+1)(2T+1)$$

derive the following expressions:

- a. $\mathbf{X}'\mathbf{X} = \frac{1}{6} \begin{bmatrix} 6T & 3T(T+1) \\ 3T(T+1) & T(T+1)(2T+1) \end{bmatrix}$
- b. $(\mathbf{X}'\mathbf{X})^{-1} = \frac{2}{T(T^2-1)} \begin{bmatrix} (T+1)(2T+1) & -3(T+1) \\ -3(T+1) & 6 \end{bmatrix}$
- c. $\hat{\beta}_0 = \frac{2}{T(T-1)} \left[(2T+1) \sum_{t=1}^T y_t - 3 \sum_{t=1}^T ty_t \right]$
- d. $\hat{\beta}_1 = \frac{6}{T(T^2-1)} \left[2 \sum_{t=1}^T ty_t - (T+1) \sum_{t=1}^T y_t \right]$
- d. $\text{Var}(\hat{y}_t) = \hat{\sigma}^2 \left[1 + \frac{2}{T(T-1)} \left(1 - 4T - 6h + 6 \frac{(T+h)^2}{T+1} \right) \right]$

Exercise 9

Plot the following figures using *only* `ggplot2` functionalities. All of them are from chapter 5, so look at the data details from the book.


a.

```
autoplot(uschange[,c("Consumption", "Income")]) +  
  ylab("% change") + xlab("Year")
```




b.

```
fit.cons <- tslm(Consumption ~ Income, data=uschange)  
uschange %>%  
  as.data.frame %>%  
  ggplot(aes(x=Income, y=Consumption)) +  
    ylab("Consumption (quarterly % change)") +  
    xlab("Income (quarterly % change)") +  
    geom_point() +  
    geom_smooth(method="lm", se=FALSE)
```



c.

```
autoplot(uschange[,3:5], facets = TRUE, colour=TRUE) +  
  ylab("") + xlab("Year") +  
  guides(colour="none")
```



Exercise 10

The same question as before, but with `ausbeer` data in chapter 5.4.

a.

```
beer <- window(ausbeer, start=1992)
autoplot(beer) + xlab("Year") + ylab("megalitres") +
  ggtitle("Australian quarterly beer production")
```




b. You need this

```
fit.beer <- tslm(beer ~ trend + season)
```


then replicate this in `ggplot2`

```
autoplot(beer, series="Data") +
  autolayer(fitted(fit.beer), series="Fitted") +
  xlab("Year") + ylab("Megalitres") +
  ggtitle("Quarterly Beer Production")
```



C.

```
cbind(Data=beer, Fitted=fitted(fit.beer)) %>%
  as.data.frame %>%
  ggplot(aes(x=Data, y=Fitted, colour=as.factor(cycle(beer)))) +
  geom_point() +
  ylab("Fitted") + xlab("Actual values") +
  ggtitle("Quarterly beer production") +
  scale_colour_brewer(palette="Dark2", name="Quarter") +
  geom_abline(intercept=0, slope=1)
```



Exercise 11

Consider a regression model with a linear trend and quarterly dummy variables,

$$y_t = \beta_0 + \beta_1 t + \sum_{k=2}^4 \beta_k d_{k,t} + \varepsilon_t$$

We estimated this model for `ausbeer` dataset:

```
beer2 <- window(ausbeer, start=1992)
fit.beer <- tslm(beer2 ~ trend + season)
summary(fit.beer)
```

Estimate the coefficient vector $\beta = (\beta_0, \dots, \beta_4)$ by explicitly defining the data matrix $\mathbf{X}$ and then using the formula

$$\hat{\beta} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$$