

A Simple Sweave Document using Caching

Roger D. Peng

May 1, 2007

First we load the packages which contain the data and functions that we will need later

```
> library(datasets)
> library(utils)
```

We fit a linear model to the `airquality` dataset. In this code chunk we will turn caching on so that we cache the `airquality` data frame and the `fit` object.

```
> data(airquality)
> fit <- lm(Ozone ~ Temp + Wind + Solar.R, data = airquality)
```

We can summarize the regression coefficients and their uncertainties in a table

```
> library(xtable)
> print(xtable(fit))
```

	Estimate	Std. Error	t value	Pr(> t)
(Intercept)	-64.3421	23.0547	-2.79	0.0062
Temp	1.6521	0.2535	6.52	0.0000
Wind	-3.3336	0.6544	-5.09	0.0000
Solar.R	0.0598	0.0232	2.58	0.0112

We can also plot some model diagnostics.

Do just a bivariate regression with ozone and temperature.

```
> fit <- lm(Ozone ~ Temp, data = airquality)
```

Plot the data with the regression line Hmm, things seem a little nonlinear here.

```
> par(mfcol = c(2, 2))
> plot(fit)
```

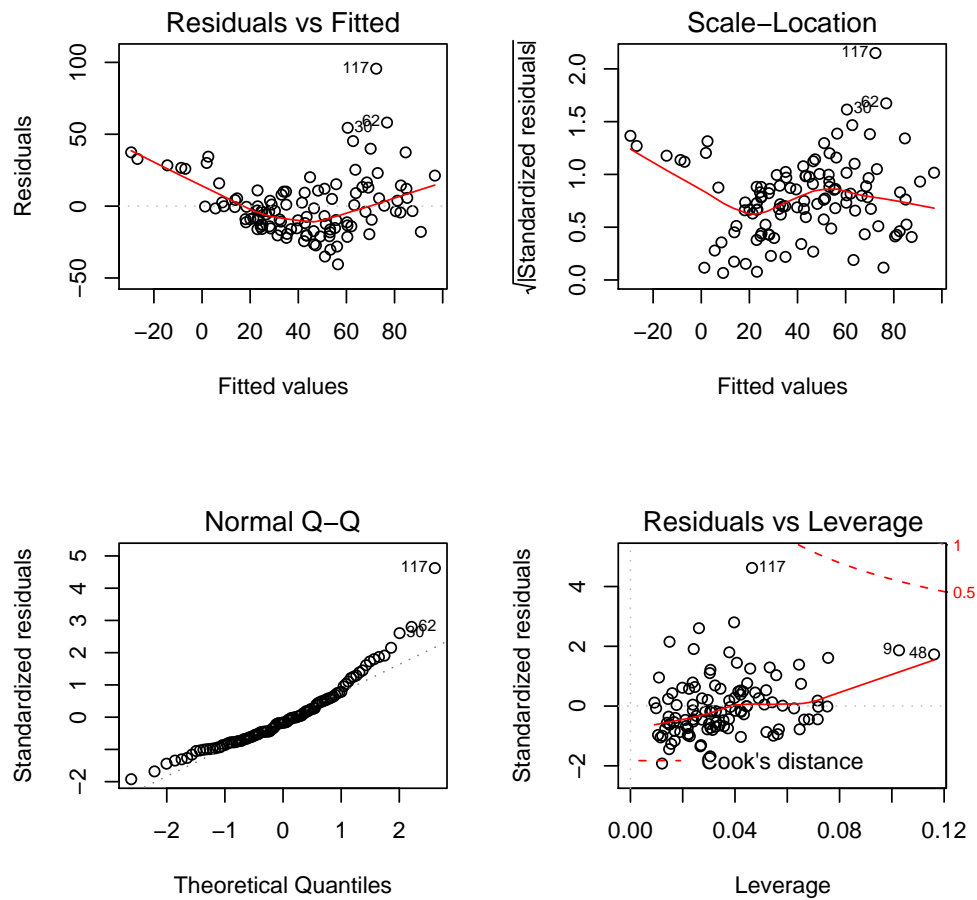


Figure 1: Linear model diagnostic plots

```
> par(mfcol = c(1, 1))  
> with(airquality, plot(Temp, Ozone))  
> abline(fit)
```




Figure 2: Temperature and ozone data with fitted regression line