

MSG500/MVE190

Linear Models - Lecture 7

Rebecka Jörnsten
Mathematical Statistics
University of Gothenburg/Chalmers University of Technology

November 14, 2012

1 RECAP

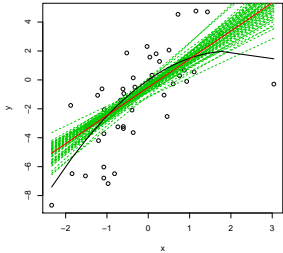
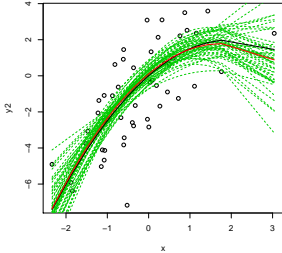
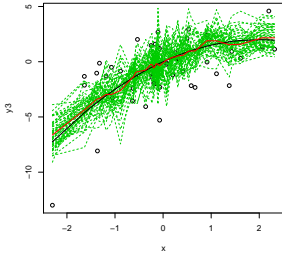
- Multivariate regression - more convenient to work with this in a matrix form
- $y = X\beta + \epsilon$ where X is a $n \times p$ matrix where each columns are the values for an x -variable, column of 1's in the first column for the intercept.
- Normal equations $X'X\beta = X'y$, solve for β if you can $\hat{\beta} = (X'X)^{-1}X'y$
- If x 's are correlated, then $\hat{\beta}_j$ depend not only the correlation between variable x_j and y but also all the other correlations $\text{corr}(x_k, y)$ - difficult to interpret multivariate regression parameters
- Correlated x 's also lead to numerical instabilities in the estimation due to the fact that the inverse of $X'X$ is unstable
- $E[\hat{\beta}] = \beta$ and $V(\hat{\beta}) = \sigma^2(X'X)^{-1}$.
- Three sources of estimation variance: noise level σ^2 , sample size and structure in X (variance of each x' and correlation between x 's).
- You can test the significance of $\hat{\beta}_j$ using a t-test: $t_j = \frac{(\hat{\beta})_j - 0}{\sqrt{\sigma^2(X'X)^{-1}_{jj}}}$ compare with t_{n-p} distribution
- Caution: if x 's are correlated the t-test can give misleading results. Also, look out for multiple testing issues!
- The goodness-of-fit F-test: $F_{obs} = \frac{(SS_T - RSS)/(p-1)}{RSS/(n-p)}$ compare with $F_{p-1, n-p}$. Use to test if at least one $\beta_j, j = 1, \dots, p-1$ is non-zero
- You can also use an F-test to compare any pair of nested models. $F_{obs} = \frac{(RSS_{simple} - RSS_{complex})/(\Delta p)}{RSS_{complex}/(n-p_{complex})}$, where Δp is the difference in number of parameters between the simple and complex models. Compare F_{obs} to $F_{\Delta p, n-p_{complex}}$.

2 Model Selection

The tools we have discussed so far are t-tests on individual slope estimates and F-test to compare two nested models. Last lecture we also looked at using the F-test to select between models in a backward search - starting with a model including all variables and dropping them one by one until the drop leads to a rejection in the F-test. Now we will look into other criteria that can be used to select between models.

The ultimate validation of a model is if it can be used to *predict* future observations, i.e. does the model generalize to the population under consideration? If you look back to lecture 1, I talked about how the safe option of picking as complex a model as possible was not so safe statistically. That was because while a complex model may have little or no bias it suffers from high estimation variance. This is what will hurt us when we do prediction. We don't know how much 'off' the true model our estimate

is, but with a high estimation variance this deviation can be quite substantial. This will result in poor prediction performance for the model. On the other hand, a simple model may make consistent errors in the prediction, but these predictions will not be too sensitive to individual data sets used for estimation of model parameters. This may in fact work better for prediction. Let's recap some things we talked about in lecture 1:

Model	Simple	added complexity	Flexible
	few parameters	...	many parameters
	rigid	...	adapts to data
	Linear	polynomial/nonlinear	Local average/smoother
			
Properties	Large bias, low variance	...	Small bias, high variance

Let us summarize this bias-variance tradeoff:

$$E[\hat{y}] - E[y], \text{ average fitted value across many data sets - the true model value} = \text{Bias}$$

$$V[\hat{y}] = \text{the estimation variance}$$

With more complex models we can adapt to the data and reduce the bias, but more complex models results in increased estimation variance. In terms of prediction both bias and variance work against us. Bias means we are mismatching our model to the true $y - X$ relationship, so even with infinite amounts of data our predictions would be off. Estimation variance has to do what happens with finite data and for a particular sample. The estimated model can be quite far from its ideal (what you would get for infinite data, or if you could average your estimated models over many finite data sets drawn from the same true model) so even if the model has no bias due to estimation variance our predictions based on an estimated model can be off. We want to control both bias and variance. We combine the two in the Mean Squared Error (MSE):

$$MSE = Bias^2 + Variance.$$

A good prediction model can thus be identified as one that minimizes MSE. To avoid confusion with the MSE defined as $RSS/(n - p)$ I will refer to the $MSE = Bias^2 + Variance$ as the *prediction MSE*.

How can we select models in practise? We can't actually compute the prediction MSE since we need to know the true model to compute the bias! We need to find a substitute for the prediction MSE that we can compute from observed data. Can we use the residual sum of squares (RSS) as a substitute? No! The $RSS = \sum_i (y_i - \hat{y}_i)^2$ always decreasing the more complex the model is (the more parameters to allow the model to be matched to the data). In fact, if we use n parameters we can get a perfect fit to the data, $RSS = 0$. That would make no sense for prediction though since we are then building a model to match the random error ϵ_i .

Let's revisit the South-African heart disease data. There are 11 explanatory variables and the outcome is the cholesterol level (ldl). Let's fit all possible subset models to this data. There are 11 models of size 1 (size = number of variables), $(11*10/2)$ models of size 2, etc. In Figure 1 I depict the RSS as a function of model size for all subset models.

It is more common to only look at the winning models of each size. In the right panel of Figure 1 I show you the so-called *RSS curve*. Note that this is decreasing with model size. The RSS (or MSE) is

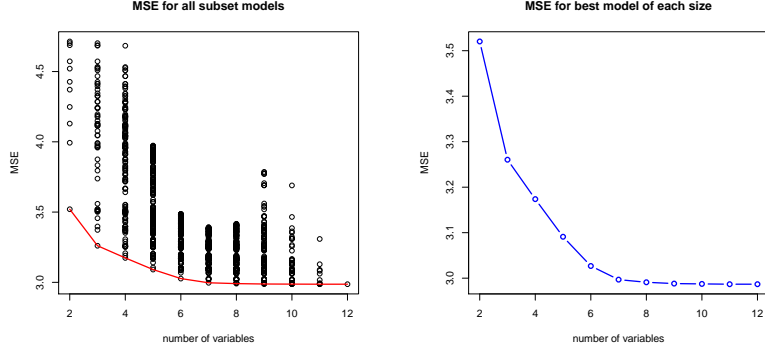


Figure 1: RSS as a function of model size for all subset models for the SA heart disease data. The blue line connects the best model of each size.

minimized for the largest model, in this case the one including all the variables. Thus, RSS is not a good substitute for prediction MSE.

2.1 Using training and test data

Let us assume we had access to another data set. Our *training* data is used to estimate model parameters. Our *test* data is used to check which models work well for prediction.

$$\text{Training data: } (X_i, y_i)_{i=1}^n, X_i = (x_{i1}, x_{i2}, \dots, x_{i,p-1})$$

$$\text{Test data: } (X_i, y_i^{new})_{i=1}^n, X_i = (x_{i1}, x_{i2}, \dots, x_{i,p-1})$$

The test data consist of new outcome data drawn from the same true model and at the same x -locations as the training data.

The true model $\beta^* = (\beta_0^*, \beta_1^*, \dots, \beta_{p-1}^*)$ is unknown to us. If the outcome is not related to some of the x -variables some of the $\beta_j^* = 0$. We can thus write

$$\text{Training data: } y_i = X_i \beta^* + \epsilon_i, \epsilon_i \sim N(0, \sigma^2)$$

$$\text{Test data: } y_i^{new} = X_i \beta^* + \epsilon_i^{new}, \epsilon_i^{new} \sim N(0, \sigma^2)$$

1. We enumerate all models $m = 1, \dots, M$. If we have $p - 1$ variables there are $M = 2^{p-1}$ possible subset models.
2. We fit each model m to the training data and obtain parameter estimates $\hat{\beta}(m)$ with corresponding fitted values $\hat{y}(m)_i, i = 1, \dots, n$
3. $RSS(m) = \sum_{i=1}^n (y_i - \hat{y}(m)_i)^2$
4. $MSE(m)_{train} = \frac{1}{n} RSS(m)$
5. $pMSE(m) = MSE(m)_{test} = \frac{1}{n} \sum_{i=1}^n (y_i^{new} - \hat{y}(m)_i)^2$

In Figure 2 I mimic the training/testing setup from above. I split the data in half - use one half as training and reserve the other part for testing. (Note, this is not the exact setup from above since I am not creating a test data at the precise same x locations in the training data.)

I compare the MSE on the training data to the pMSE (MSE on the test data). Note that while MSE on the training data is minimized for the largest model, the pMSE decreases up to a point ($p_{selected} = 9$) and then starts increasing. The pMSE decreases as long as the decrease in bias exceeds the increase in estimation variance. At some point, when we exceed the true model size, we are no longer gaining

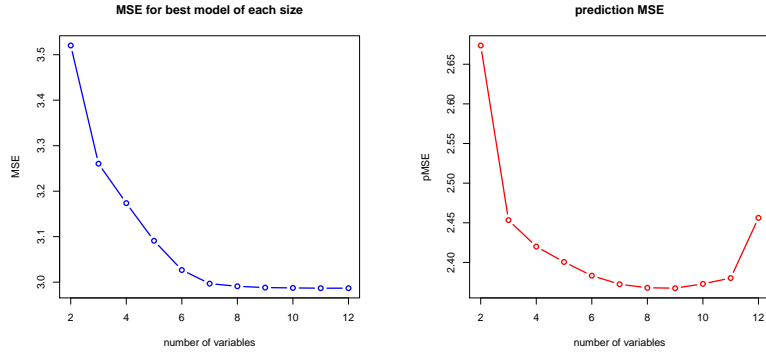


Figure 2: Left: RSS-curve. MSE on training data. Middle: pMSE-curve. MSE on test data.

in terms of bias but estimation variance increases since we are including unnecessary parameters in our model. Note, while the optimum size model is here identified as the minimizer of pMSE, this model may not be the minimum MSE model of the same size (see Table 1).

	Intercept	age	sbp	adiposity	obesity	typea
pMSE win	TRUE	FALSE	TRUE	TRUE	TRUE	TRUE
MSE of same size	TRUE	TRUE	FALSE	TRUE	TRUE	TRUE

	alcohol	alcind	tobacco	tobind	chd	famhist
pMSE win	TRUE	FALSE	TRUE	TRUE	TRUE	FALSE
MSE of same size	TRUE	FALSE	FALSE	TRUE	TRUE	TRUE

Table 1: Selected variables using pMSE

3 Demo 7

We will examine training and prediction MSE using a data set on car prices and gas mileage (how many miles/km per gallons/liters). We first read the data into R.

```
> cars <- data.frame(read.table("cars.dat", sep = "\t", header = T))
> print(dim(cars))

[1] 82 19

> print(names(cars))

[1] "mid.price"   "city.mpg"    "hw.mpg"      "airbagstd"   "cylnbr"      "engsize"
[7] "horsepwr"   "rpm.at.max"  "engrev.high" "mantrans.op" "fueltank"    "passengers"
[13] "length"     "width"       "uturn"       "rearroom"   "luggage"     "weight"
[19] "domestic"
```

There are, as you see, 82 different cars in this data set and 19 variables. We will focus on the price of the cars as a function of the other variables, which include mileage (miles per gallon in the city and on highway (compare liters per km in Sweden)), airbag standard (0 if not included, 1 if for the driver and 2 if passenger/side), the number of cylinders, the engine size, horsepower, maximum rpm of engine, manual transmission (yes=1, no=0), size of the fuel tank, passenger room, length and width of the car, a measure about the space needed to make u-turn with the car, size of the rear room of the car, luggage room, weight of the car, and finally an indicator if the car is domestic (US built=1) or not.

Let us examine the data set. I first split the data into a random training set and a random test set for illustration purposes - this means that you will see different results when you run the demo yourselves.

```
> ntrain <- 60
> ii <- sample(seq(1, dim(cars)[1]), ntrain)
> cars.train <- cars[ii, ]
> row.names(cars.train) <- seq(1, ntrain)
> cars.test <- cars[-ii, ]
> row.names(cars.test) <- seq(1, dim(cars.test)[1])
```

I first try some select pairwise scatter plots. Do this yourselves for other pairs of variables.

```
> par(mfrow = c(2, 2))
> plot(cars.train$mid, cars.train$ci, main = "citympg on price")
> plot(cars.train$hw, cars.train$ci, main = "citympg on hwmpg")
> plot(cars.train$le, cars.train$ci, main = "citympg on length")
> plot(cars.train$engsize, cars.train$ci, main = "citympg on enginesize")
> p <- locator()

> par(mfrow = c(2, 2))
> plot(cars.train$le, cars.train$wi, main = "width on length")
> plot(cars.train$wi, cars.train$we, main = "weight on width")
> plot(cars.train$le, cars.train$wh, main = "wheelbase on length")
> plot(cars.train$lu, cars.train$rea, main = "rearroom on luggage room")
> p <- locator()
```

In Figure 3 and 4 you can see that there are correlations between both the outcome and covariates as well as between the covariates themselves. We will focus on price as the outcome, but keep in mind that we have a collinearity problem.

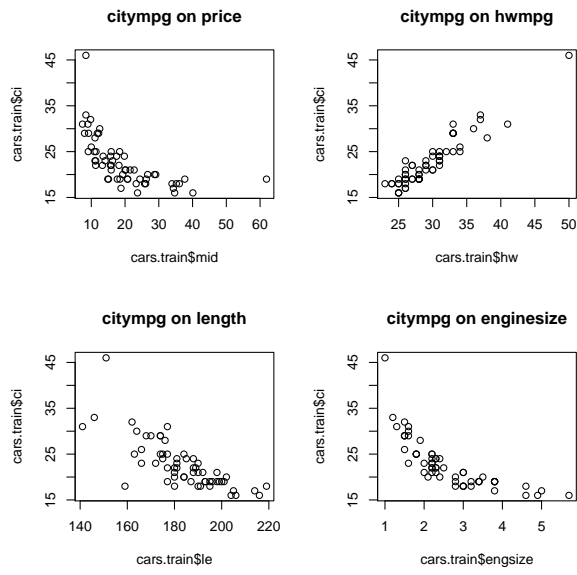


Figure 3: Scatter plots - 1

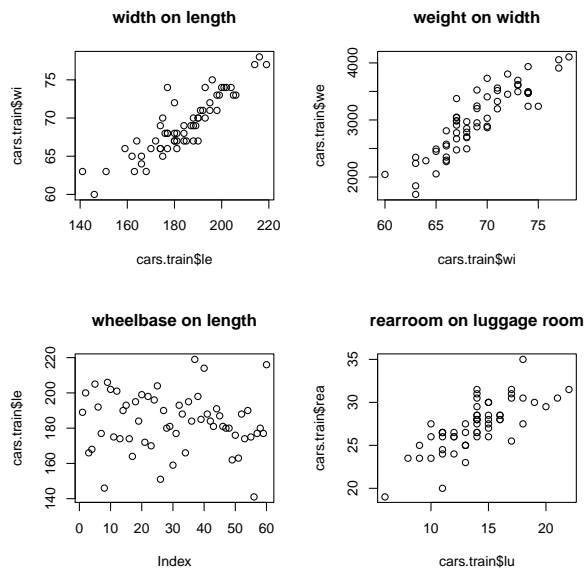


Figure 4: Scatter plots - 2

```

> par(mfrow = c(2, 2))
> plot(cars.train$ci, cars.train$mid, main = "price on citymg")
> plot(cars.train$we, cars.train$mid, main = "price on weight")
> plot(cars.train$ho, cars.train$mid, main = "price on horsepower")
> plot(cars.train$man, cars.train$mid, main = "price on transmission")
> p <- locator()

```

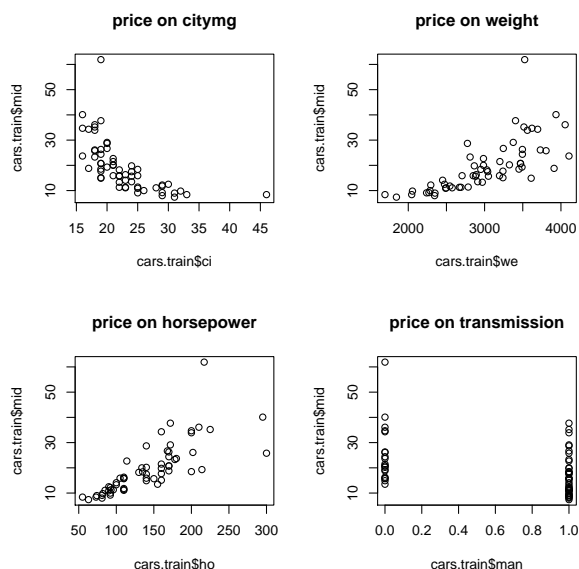


Figure 5: Price vs city milage, weight, horsepower and manual transmission.

In Figure 5 I depict some pairwise relationships between car prices and some other variables. Do you spot any problems with the modeling assumptions? I believe we have a skewed distribution for the outcome (salary, price and similar data often has this long-tailed appearance) and increasing variance with the expected value of the outcome. In addition, the relationship between price and mileage does not appear to be linear. I try some transformations (log of price, and inverse of mileage) to see if we can fix the problem.

```

> par(mfrow = c(2, 2))
> plot(1/cars.train$ci, log(cars.train$mid), main = "log(price) on 1/citympg")
> plot(cars.train$we, log(cars.train$mid), main = "log(price) on weight")
> plot(cars.train$ho, log(cars.train$mid), main = "log(price) on horsepower")
> plot(cars.train$man, log(cars.train$mid), main = "log(price) on transmission")
> p <- locator()

```

From Figure 6 we see that the transformations did a pretty good job. Try some other transformations at home.

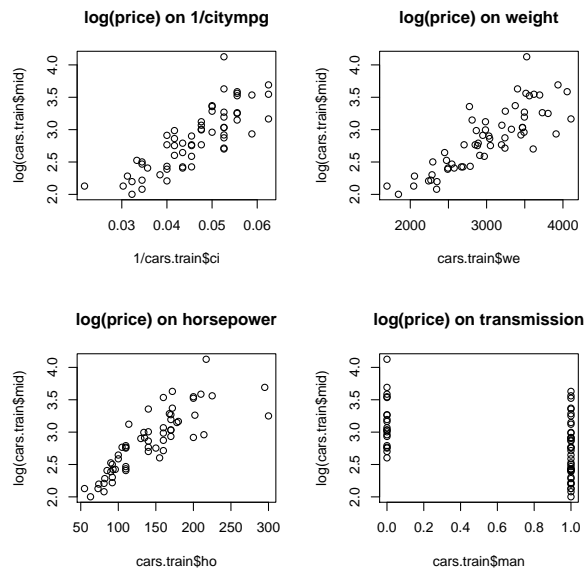


Figure 6: $\log(\text{price})$ vs $1/(\text{city milage})$, weight, horsepower and manual transmission.

I decide to use this transformations and rename and transform objects in the data frame as follows:

```
> cars.train[, 1] <- log(cars.train$mid)
> cars.train[, 2] <- 1/cars.train$ci
> cars.train[, 3] <- 1/cars.train$hw
> cars.test[, 1] <- log(cars.test$mid)
> cars.test[, 2] <- 1/cars.test$ci
> cars.test[, 3] <- 1/cars.test$hw
> names(cars.train)[c(1, 2, 3)] <- c("log.mid.price", "city.gpm", "hw.gpm")
> names(cars.test)[c(1, 2, 3)] <- c("log.mid.price", "city.gpm", "hw.gpm")
```

We are now ready to try a linear model fit:

```
> mm <- lm(log.mid.price ~ city.gpm + hw.gpm + airbagstd + cylnbr + engsize +
+   horsepwr + rpm.at.max + engrev.high + mantrans.op + fueltank + passengers +
+   length + width + uturn + rearroom + luggage + weight + domestic, data = cars.train)
> print(ms <- summary(mm))
```

Call:

```
lm(formula = log.mid.price ~ city.gpm + hw.gpm + airbagstd +
    cylnbr + engsize + horsepwr + rpm.at.max + engrev.high +
    mantrans.op + fueltank + passengers + length + width + uturn +
    rearroom + luggage + weight + domestic, data = cars.train)
```

Residuals:

```
      Min       1Q   Median       3Q      Max
-0.37621 -0.09170 -0.00774  0.09307  0.38837
```

Coefficients:

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept)  1.632e+00  1.632e+00   1.000  0.323282
city.gpm     2.632e+01  1.329e+01   1.980  0.054392 .
hw.gpm       9.028e+00  1.875e+01   0.482  0.632716
airbagstd    1.974e-01  4.770e-02   4.140  0.000169 ***
cylnbr      -8.654e-03  5.666e-02  -0.153  0.879364
engsize     -1.321e-01  1.186e-01  -1.113  0.272151
horsepwr     3.476e-03  2.030e-03   1.712  0.094417 .
rpm.at.max  -2.010e-05  1.016e-04  -0.198  0.844110
engrev.high  4.070e-05  1.043e-04   0.390  0.698418
mantrans.op -1.706e-01  8.473e-02  -2.014  0.050635 .
fueltank    -1.592e-03  2.418e-02  -0.066  0.947839
passengers  -2.655e-02  6.868e-02  -0.387  0.701081
length      2.811e-03  4.601e-03   0.611  0.544664
width      -8.273e-03  2.378e-02  -0.348  0.729687
uturn      -1.975e-02  1.709e-02  -1.155  0.254652
rearroom     8.706e-03  1.688e-02   0.516  0.608866
luggage     2.383e-02  1.892e-02   1.259  0.215039
weight      1.429e-05  2.629e-04   0.054  0.956912
domestic    -1.685e-01  7.987e-02  -2.109  0.041075 *
```

Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.1857 on 41 degrees of freedom

Multiple R-squared: 0.8928, Adjusted R-squared: 0.8458

F-statistic: 18.98 on 18 and 41 DF, p-value: 2.02e-14

```
> par(mfrow = c(2, 2))
> plot(mm)
> p <- locator()
```

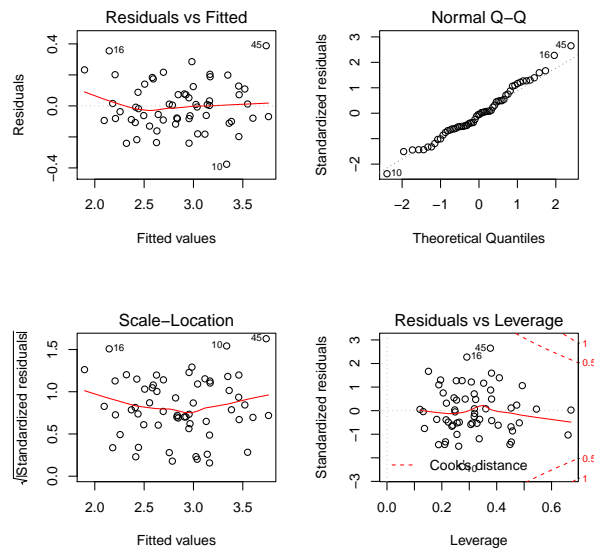


Figure 7: Diagnostic plots

From the model summary we can see that the R-squared is 0.893. Check the model summary for the Goodness-of-fit test and the p-values of individual coefficients. Are there any surprises? Do you see the collinearity problem having an impact in this model summary? In Figure 7 I display the standard diagnostic plots. We will revisit them one by one below, in addition to some other plots, to check for outliers.

Identifying outliers

In the plots below, I identify potential outliers in the data set.

```
> par(mfrow = c(1, 1))
> plot(mm$fit, mm$res, xlab = "fitted values", ylab = "residuals")
> abline(h = 0)
> id1 <- identify(mm$fit, mm$res, pos = T)
```




Figure 8: Residuals vs fitted values

```
> qq <- seq(0.5/ntrain, (ntrain - 0.5)/ntrain, length = ntrain)
> normq <- qnorm(p = qq)
> rsort <- sort(rstandard(mm))
> rlist <- sort.list(rstandard(mm))
> plot(normq, rsort, xlab = "Theoretical quantiles", ylab = "Standardized residuals")
> qr <- quantile(rstandard(mm))
> qn <- quantile(qnorm(p = qq))
> b <- (qr[4] - qr[2]) / (qn[4] - qn[2])
> a <- qr[4] - b * qn[4]
> abline(a, b)
> id2 <- identify(normq, sort(rstandard(mm)), label = rlist, pos = T)

> plot(mm$fit, abs(rstandard(mm)), xlab = "fitted values", ylab = "|standardized residuals|")
> id3 <- identify(mm$fit, abs(rstandard(mm)), pos = T)

> lm1 <- lm.influence(mm)
> cooks <- cooks.distance(mm)
> plot(cooks, main = "Cook's Distance", type = "h")
> abline(h = qf(0.95, 1, mm$df), lty = 2)
> idc <- identify(cooks, pos = T)

> plot(lm1$hat, main = "Leverage")
> idlev <- identify(lm1$hat, pos = T)
```




Figure 9: QQplot

We also examine the leverage and what the effect is of dropping an observation, on both slope estimates and the standard error estimate $\hat{\sigma}$.

```
> plot(lm1$coeff[, 4], main = "change in slope 4")
> id4 <- identify(lm1$coeff[, 4], pos = T)

> plot(lm1$coeff[, 7], main = "change in slope 7")
> id7 <- identify(lm1$coeff[, 7], pos = T)

> plot(lm1$sig, main = "change in sigma")
> ids <- identify(lm1$sig, pos = T)
```

We collect on the outlier information and drop the observation most frequently identified (or more than one if they equally frequently identified):

```
> indvec <- sort(c(id1$ind, rlist[id2$ind], id3$ind, idlev$ind, id4$ind, id7$ind,
+   ids$ind))
> print(table(indvec))

indvec
10 16 45
 3  1  6

> maxid <- max(table(indvec))
> indout <- unique(indvec)[table(indvec) == max(table(indvec))]
```

Here, this is observation 45. Note, when you do this yourselves you may not need to drop an observation, or perhaps drop several. If the latter, just run the code again after dropping the first observation or pick more than one to drop in one go by manually adjusting `indout`




Figure 10: Absolute standardized residuals vs fitted values




Figure 11: Cook's distance: $D_i = e^2 \frac{h_{ii}}{p * MSE * (1 - h_{ii})^2}$

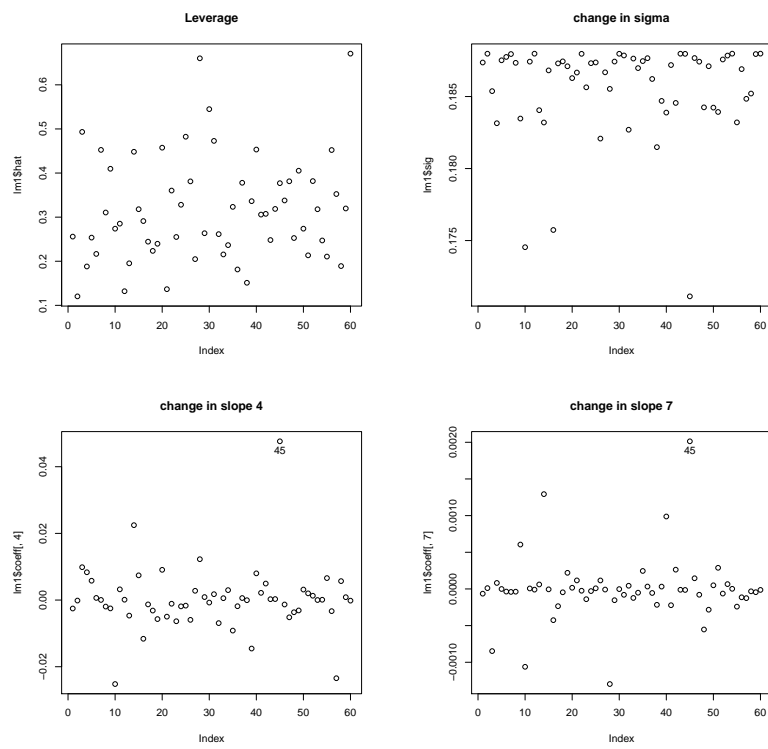


Figure 12: Top Left: Leverage. Top Right: Effect on standard error. Lower panels: impact on slope estimates.

Updating the model

We can now update the model without the outlier:

```
> mmb <- lm(log.mid.price ~ city.gpm + hw.gpm + airbagstd + cylnbr + engsize +
+   horsepwr + rpm.at.max + engrev.high + mantrans.op + fueltank + passengers +
+   length + width + uturn + rearroom + luggage + weight + domestic, data = cars.train,
+   subset = -indout)
> print(summary(mmb))
```

Call:

```
lm(formula = log.mid.price ~ city.gpm + hw.gpm + airbagstd +
    cylnbr + engsize + horsepwr + rpm.at.max + engrev.high +
    mantrans.op + fueltank + passengers + length + width + uturn +
    rearroom + luggage + weight + domestic, data = cars.train,
    subset = -indout)
```

Residuals:

```
      Min       1Q   Median       3Q      Max
-0.29365 -0.07898 -0.01359  0.10554  0.31400
```

Coefficients:

```
              Estimate Std. Error t value Pr(>|t|)
(Intercept) -3.259e-02  1.612e+00  -0.020  0.98397
city.gpm     2.894e+01  1.228e+01   2.356  0.02345 *
hw.gpm       4.233e+00  1.736e+01   0.244  0.80859
airbagstd    1.498e-01  4.698e-02   3.189  0.00277 **
cylnbr       2.492e-02  5.351e-02   0.466  0.64400
engsize     -1.227e-01  1.094e-01  -1.121  0.26883
horsepwr     1.463e-03  1.998e-03   0.732  0.46838
rpm.at.max   6.518e-05  9.821e-05   0.664  0.51071
engrev.high  5.523e-05  9.626e-05   0.574  0.56934
mantrans.op -7.442e-02  8.495e-02  -0.876  0.38628
fueltank     1.191e-03  2.231e-02   0.053  0.95769
passengers  -1.918e-02  6.335e-02  -0.303  0.76366
length       6.727e-04  4.305e-03   0.156  0.87661
width        4.283e-05  2.211e-02   0.002  0.99846
uturn       -1.985e-03  1.692e-02  -0.117  0.90719
rearroom     1.422e-02  1.568e-02   0.907  0.36996
luggage      1.401e-02  1.777e-02   0.788  0.43512
weight       1.168e-04  2.449e-04   0.477  0.63598
domestic    -1.423e-01  7.417e-02  -1.918  0.06224 .
---
```

```
Signif. codes:  0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1
```

Residual standard error: 0.1711 on 40 degrees of freedom

Multiple R-squared: 0.8986, Adjusted R-squared: 0.853

F-statistic: 19.7 on 18 and 40 DF, p-value: 1.81e-14

```
> par(mfrow = c(2, 2))
> plot(mmb)
> p <- locator()
```

In Figure 13 we can see the diagnostic plot after we drop the outlier(s). Compare with Figure 7 above. You should also compare the model summary: has the fit improved? in what sense?

Stepwise model selection

We try a step wise model selection with the updated model as the starting point. Note, the option `trace=F` tells R to only show us the final model.

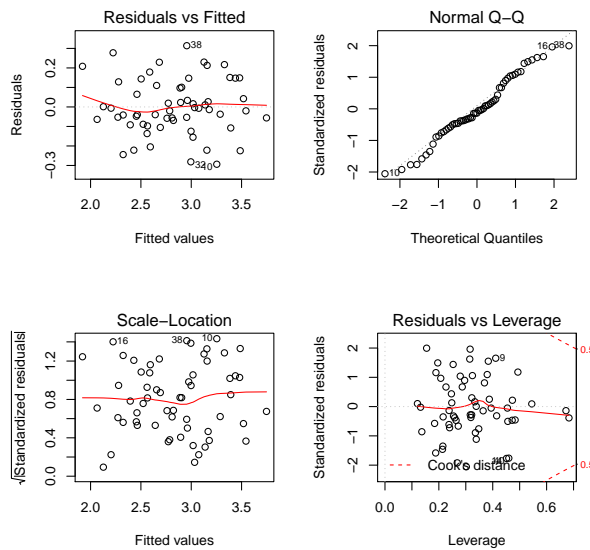


Figure 13: Updated model fit: diagnostic plots

```
> selectstep <- step(mmb, trace = F)
```

We can use the selected model for prediction. First, we look at the updated model summary after selection:

```
> print(summary(selectstep))
```

Call:

```
lm(formula = log.mid.price ~ city.gpm + airbagstd + engsize +
    horsepwr + mantrans.op + luggage + domestic, data = cars.train,
    subset = -indout)
```

Residuals:

Min	1Q	Median	3Q	Max
-0.31048	-0.09376	-0.01446	0.10227	0.33623

Coefficients:

	Estimate	Std. Error	t value	Pr(> t)	
(Intercept)	1.0081506	0.1964829	5.131	4.53e-06	***
city.gpm	32.7339789	5.7974054	5.646	7.31e-07	***
airbagstd	0.1538940	0.0344034	4.473	4.33e-05	***
engsize	-0.1526909	0.0563171	-2.711	0.009111	**
horsepwr	0.0029042	0.0007772	3.737	0.000472	***
mantrans.op	-0.0980543	0.0650805	-1.507	0.138067	
luggage	0.0253647	0.0103841	2.443	0.018084	*
domestic	-0.1908148	0.0490933	-3.887	0.000294	***

 Signif. codes: 0 '***' 0.001 '**' 0.01 '*' 0.05 '.' 0.1 ' ' 1

Residual standard error: 0.1567 on 51 degrees of freedom
 Multiple R-squared: 0.8916, Adjusted R-squared: 0.8767
 F-statistic: 59.91 on 7 and 51 DF, p-value: < 2.2e-16

To compute predictions for the test data we do the following:

```
> predval <- predict(selectstep, newdata = cars.test)
> prederror <- sum((cars.test[, 1] - predval)^2)
```


For comparison, what happens when we try to fit a model to the test data? We use the full model fitted to the test data as a starting point for backward model selection. We compute the training error (fiterror) below for the selected model on the test data (this is just the residual sum of square for the backward selected model on the test data).

```
> mmtest <- lm(log.mid.price ~ city.gpm + hw.gpm + airbagstd + cylnbr + engsize +
+   horsepwr + rpm.at.max + engrev.high + mantrans.op + fueltank + passengers +
+   length + width + uturn + rearroom + luggage + weight + domestic, data = cars.test)
> print(selectttest <- step(mmtest, trace = F))
```

Call:

```
lm(formula = log.mid.price ~ city.gpm + hw.gpm + airbagstd +
    cylnbr + engsize + engrev.high + fueltank + passengers +
    length + width + rearroom + luggage + domestic, data = cars.test)
```

Coefficients:

(Intercept)	city.gpm	hw.gpm	airbagstd	cylnbr	engsize	engrev.high
5.7634630	20.9197752	32.9105730	0.1969852	-0.1725995	0.4795992	-0.0005791
fueltank	passengers	length	width	rearroom	luggage	domestic
-0.1574263	-0.4340470	0.0556435	-0.1047648	-0.0567480	-0.0501851	-0.4342678

```
> fiterror <- sum(summary(selectttest)$res^2)
> p <- locator()
```

Compare the selected models on the training and test data sets. Are they the same?

We can also compare the prediction error to the error sum of squares obtained from the selected model on the test data. Here, the prediction error, using the selected model from the training data, is 1.502, whereas the error sum of squares, using the test data both for model fitting and prediction is 0.156. What does this tell you?

Using prediction error as a model selection tool

Let us find out which model is actually best for prediction (on *this* test data). We first enumerate and fit all models to the training data and collect the MSEs for this model fits.

```
> yy <- cars.train[, 1]
> xx <- as.matrix(cars.train[, -1])
> yyt <- cars.test[, 1]
> xxt <- as.matrix(cars.test[, -1])
> rleaps <- regsubsets(xx, yy, int = T, nbest = 250, nvmax = dim(cars)[2],
+   really.big = T, method = c("ex"))
> cleaps <- summary(rleaps, matrix = T)
> tt <- apply(cleaps$which, 1, sum)
> mses <- cleaps$rss/length(yy)
```

We then plot the MSE for all the models as well as the MSE-curve (best models of each model size).

```
> plot(tt, mses, xlab = "number of variables", ylab = "MSE", main = "MSE for all subset models")
> tmin <- min(tt)
> tmax <- max(tt)
> tsec <- seq(tmin, tmax)
> msevec <- rep(0, length(tsec))
> for (tk in 1:length(tsec)) {
+   msevec[tk] <- min(mses[tt == tsec[tk]])
+ }
> lines(tsec, msevec, lwd = 2, col = 2)
> p <- locator()

> plot(tsec, msevec, xlab = "number of variables", ylab = "MSE", main = "MSE for best model of each
+   type = \"b\", col = 4, lwd = 2)
> p <- locator()
```

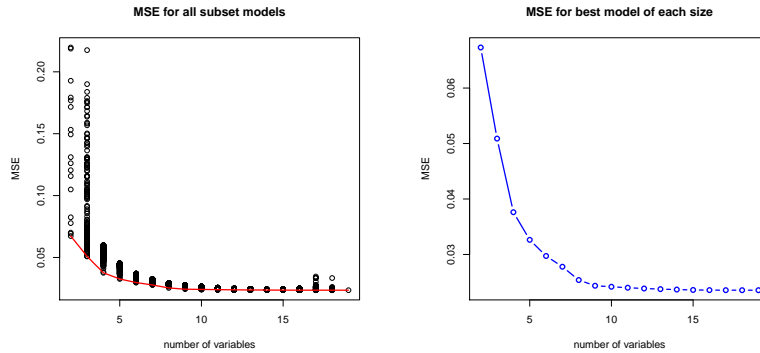


Figure 14: Left: MSE for all models fitted. Right: Best MSE models for each model size.

In Figure 14 you see that the MSE always decreases with model size.

We then apply all fitted models to the test data (without refitting, no updating of model coefficients) and collect pMSE values for each.

```
> pmses <- mses
> for (ta in (2:dim(cleaps$which)[1])) {
+   mmr <- lm(yy ~ xx[, cleaps$which[ta, -1] == T])
+   PEcp <- sum((yyt - cbind(rep(1, dim(xxt)[1]), xxt[, cleaps$which[ta,
+     -1] == T])) %*% mmr$coef)^2)/length(yyt)
+   pmses[ta] <- PEcp
+ }
> pmses[1] <- max(pmses)
> pmsevec <- rep(0, length(tsec))
> for (tk in 1:length(tsec)) {
+   pmsevec[tk] <- min(pmses[tt == tsec[tk]])
+ }

> plot(tsec, pmsevec, xlab = "number of variables", ylab = "pMSE", main = "prediction MSE",
+   type = "b", lwd = 2, col = 2)
```

We plot the pMSE values as a function of model size. Compare the pMSE curve in Figure 15 to the MSE curve in Figure 14. What do you see?

Finally, we want to compare the selected model (from training using backward selection) and the selected model that works best in terms of prediction.

```
> ptmin <- which.min(pmses)
> pmod <- cleaps$which[ptmin, ]
> winsize <- sum(pmod)
> mtmin <- which.min(mses[tt == sum(pmod)])
> mod <- (cleaps$which[tt == sum(pmod), ])[mtmin, ]
> print(names(mod[mod == T])[-1])

[1] "city.gpm" "airbagstd" "domestic"

> print(names(selectstep$model)[-1])

[1] "city.gpm" "airbagstd" "engsize" "horsepwr" "mantrans.op" "luggage"
[7] "domestic"
```

In the above summary you can compare the model selected on training data and the model that was actually best for prediction on *this* test data. What do you see? Repeat this exercise a couple of times and try to summarize your findings. Is the same model always best for prediction for all test data sets? What varies? The model size or the included variables or both?

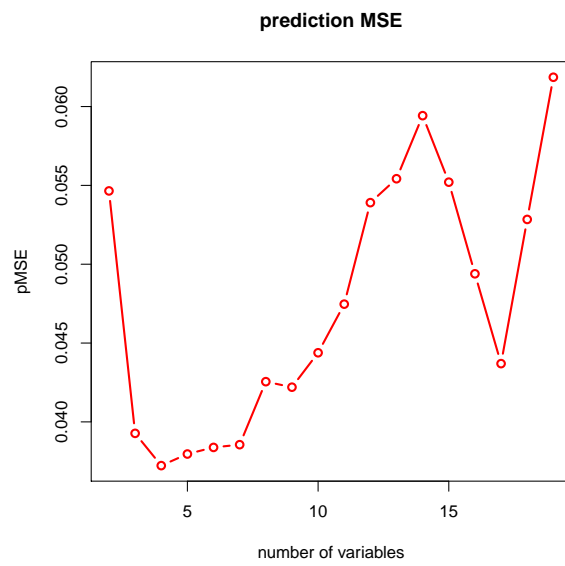


Figure 15: pMSE curve for the cars data