

Chapter 9

1D Regression

... estimates of the linear regression coefficients are relevant to the linear parameters of a broader class of models than might have been suspected.

Brillinger (1977, p. 509)

After computing $\hat{\beta}$, one may go on to prepare a scatter plot of the points $(\hat{\beta}x_j, y_j)$, $j = 1, \dots, n$ and look for a functional form for $g(\cdot)$.

Brillinger (1983, p. 98)

Regression is the study of the conditional distribution $Y|\mathbf{x}$ of the response Y given the $k \times 1$ vector of nontrivial predictors \mathbf{x} . The scalar Y is a random variable and \mathbf{x} is a random vector. In Chapter 5, a special case of regression was the multiple linear regression model $Y_i = w_{i,1}\eta_1 + w_{i,2}\eta_2 + \dots + w_{i,p}\eta_p + e_i = \mathbf{w}_i^T \boldsymbol{\eta} + e_i$ for $i = 1, \dots, n$. In this chapter, the subscript i is often suppressed and the multiple linear regression model is written as $Y = \alpha + x_1\beta_1 + \dots + x_k\beta_k + e = \alpha + \boldsymbol{\beta}^T \mathbf{x} + e$ where $k = p - 1$. The primary difference is the separation of the constant term α and the nontrivial predictors \mathbf{x} . In Chapter 5, $w_{i,1} \equiv 1$ for $i = 1, \dots, n$. Taking $Y = Y_i$, $\alpha = \eta_1$, $\beta_j = \eta_{j+1}$, and $x_j = w_{i,j+1}$ and $e = e_i$ for $j = 1, \dots, k = p - 1$ shows that the two models are equivalent. The change in notation was made because the distribution of the nontrivial predictors is very important for the theory of the more general regression models.

Definition 9.1: In a *1D regression model*, Y is conditionally independent of \mathbf{x} given the *sufficient predictor* $SP = h(\mathbf{x})$, written

$$Y \perp\!\!\!\perp \mathbf{x} | h(\mathbf{x}), \quad (9.1)$$

where the real valued function $h : \mathbb{R}^p \rightarrow \mathbb{R}$.

This chapter will primarily consider 1D regression models where $h(\mathbf{x}) = \alpha + \boldsymbol{\beta}^T \mathbf{x}$. An important 1D regression model, introduced by Li and Duan (1989), has the form

$$Y = g(\alpha + \boldsymbol{\beta}^T \mathbf{x}, e) \quad (9.2)$$

where g is a bivariate (inverse link) function and e is a zero mean error that is independent of \mathbf{x} . The constant term α may be absorbed by g if desired.

Special cases of the 1D regression model (9.1) include many important *generalized linear models* (GLMs) and the additive error *single index model*

$$Y = m(\alpha + \beta^T \mathbf{x}) + e. \quad (9.3)$$

Typically m is the conditional mean or median function. For example if all of the expectations exist, then

$$E[Y|\mathbf{x}] = E[m(\alpha + \beta^T \mathbf{x})|\mathbf{x}] + E[e|\mathbf{x}] = m(\alpha + \beta^T \mathbf{x}).$$

The *multiple linear regression model* is an important special case where m is the identity function: $m(\alpha + \beta^T \mathbf{x}) = \alpha + \beta^T \mathbf{x}$. Another important special case of 1D regression is the *response transformation model* where

$$g(\alpha + \beta^T \mathbf{x}, e) = t^{-1}(\alpha + \beta^T \mathbf{x} + e) \quad (9.4)$$

and t^{-1} is a one to one (typically monotone) function. Hence

$$t(Y) = \alpha + \beta^T \mathbf{x} + e.$$

If Y_i is an observed survival time, then many *survival regression models*, including the Cox (1972) *proportional hazards model*, are 1D regression models.

Definition 9.2. *Regression* is the study of the conditional distribution of $Y|\mathbf{x}$. Focus is often on the *mean function* $E(Y|\mathbf{x})$ and/or the *variance function* $\text{VAR}(Y|\mathbf{x})$. There is a distribution for each value of $\mathbf{x} = \mathbf{x}_o$ such that $Y|\mathbf{x} = \mathbf{x}_o$ is defined. For a 1D regression with $h(\mathbf{x}) = \beta^T \mathbf{x}$,

$$E(Y|\mathbf{x} = \mathbf{x}_o) = E(Y|\beta^T \mathbf{x} = \beta^T \mathbf{x}_o) \equiv M(\beta^T \mathbf{x}_o) \quad \text{and}$$

$$\text{VAR}(Y|\mathbf{x} = \mathbf{x}_o) = \text{VAR}(Y|\beta^T \mathbf{x} = \beta^T \mathbf{x}_o) \equiv V(\beta^T \mathbf{x}_o)$$

where M is the *kernel mean function* and V is the *kernel variance function*.

Notice that the mean and variance functions depend on the *same* linear combination if the 1D regression model is valid. This dependence is typical of GLMs where M and V are known kernel mean and variance functions that depend on the family of GLMs. See Cook and Weisberg (1999a, section 23.1). A *heteroscedastic regression model*

$$Y = M(\beta_1^T \mathbf{x}) + \sqrt{V(\beta_2^T \mathbf{x})} e \quad (9.5)$$

is a 1D regression model if $\beta_2 = c\beta_1$ for some scalar c .

In multiple linear regression, the difference between the response Y_i and the estimated conditional mean function $\hat{\alpha} + \hat{\beta}^T \mathbf{x}_i$ is the residual. For more general regression models this difference may not be the residual, and the “discrepancy” $Y_i - M(\hat{\beta}^T \mathbf{x}_i)$ may not be estimating the error e_i . To guarantee that the residuals are estimating the errors, the following definition is used when possible.

Definition 9.3: Cox and Snell (1968). Let the errors e_i be iid with pdf f and assume that the regression model $Y_i = g(\mathbf{x}_i, \boldsymbol{\eta}, e_i)$ has a unique solution for e_i :

$$e_i = h(\mathbf{x}_i, \boldsymbol{\eta}, Y_i).$$

Then the i th *residual*

$$\hat{e}_i = h(\mathbf{x}_i, \hat{\boldsymbol{\eta}}, Y_i)$$

where $\hat{\boldsymbol{\eta}}$ is a consistent estimator of $\boldsymbol{\eta}$.

Example 9.1. Let $\boldsymbol{\eta} = (\alpha, \boldsymbol{\beta}^T)^T$. If $Y = m(\alpha + \boldsymbol{\beta}^T \mathbf{x}) + e$ where m is known, then $e = Y - m(\alpha + \boldsymbol{\beta}^T \mathbf{x})$. Hence $\hat{e}_i = Y_i - m(\hat{\alpha} + \hat{\boldsymbol{\beta}}^T \mathbf{x}_i)$ which is the usual definition of the i th residual for such models.

Dimension reduction can greatly simplify our understanding of the conditional distribution $Y|\mathbf{x}$. If a 1D regression model is appropriate, then the k -dimensional vector \mathbf{x} can be replaced by the 1-dimensional scalar $\boldsymbol{\beta}^T \mathbf{x}$ with “no loss of information about the conditional distribution.” Cook and Weisberg (1999a, p. 411) define a *sufficient summary plot* (SSP) to be a plot that contains all the sample regression information about the conditional distribution $Y|\mathbf{x}$ of the response given the predictors.

Definition 9.4: For a 1D regression model, a *sufficient summary plot* is a plot of $h(\mathbf{x})$ versus Y . A *response plot* or *estimated sufficient summary plot* (ESSP) is a plot of the *estimated sufficient predictor* (ESP) versus Y . If $h(\mathbf{x}) = \boldsymbol{\beta}^T \mathbf{x}$, then $Y \perp \mathbf{x} | (a + c\boldsymbol{\beta}^T \mathbf{x})$ for any constants a and $c \neq 0$. Hence $a + c\boldsymbol{\beta}^T \mathbf{x}$ is a SP with $ESP = \tilde{\alpha} + \tilde{\boldsymbol{\beta}}^T \mathbf{x}$ where $\tilde{\boldsymbol{\beta}}$ is an estimator of $c\boldsymbol{\beta}$ for some nonzero constant c .

If there is only one predictor x , then the plot of x versus Y is both a sufficient summary plot and a response plot, but generally only a response plot can be made. Since a can be any constant, $a = 0$ is often used. The following section shows how to use the OLS regression of Y on \mathbf{x} to obtain an ESP. If we plot the fitted values and the ESP versus Y , the plots are called fit–response and ESP–response plots. For multiple linear regression, these two plots are the same.

9.1 Estimating the Sufficient Predictor

Some notation is needed before giving theoretical results. Let \mathbf{x} , \mathbf{a} , \mathbf{t} , and $\boldsymbol{\beta}$ be $k \times 1$ vectors where only \mathbf{x} is random.

Definition 9.5: Cook and Weisberg (1999a, p. 431). The predictors \mathbf{x} satisfy the condition of *linearly related predictors* with 1D structure if

$$E[\mathbf{x}|\boldsymbol{\beta}^T \mathbf{x}] = \mathbf{a} + \mathbf{t}\boldsymbol{\beta}^T \mathbf{x}. \quad (9.6)$$

If the predictors \mathbf{x} satisfy this condition, then for any given predictor x_j ,

$$E[x_j|\boldsymbol{\beta}^T \mathbf{x}] = a_j + t_j\boldsymbol{\beta}^T \mathbf{x}.$$

Notice that $\boldsymbol{\beta}$ is a fixed $k \times 1$ vector. If \mathbf{x} is elliptically contoured (EC) with 1st moments, then the assumption of linearly related predictors holds since

$$E[\mathbf{x}|\mathbf{b}^T \mathbf{x}] = \mathbf{a}_b + \mathbf{t}_b\mathbf{b}^T \mathbf{x}$$

for *any* nonzero $k \times 1$ vector \mathbf{b} . The condition of linearly related predictors is impossible to check since $\boldsymbol{\beta}$ is unknown, but the condition is far weaker than the assumption that \mathbf{x} is EC. The stronger EC condition is often used since there are checks for whether this condition is reasonable, e.g. use the DD plot. The following proposition gives an equivalent definition of linearly related predictors. Both definitions are frequently used in the dimension reduction literature.

Theorem 9.1. The predictors \mathbf{x} are linearly related iff

$$E[\mathbf{b}^T \mathbf{x}|\boldsymbol{\beta}^T \mathbf{x}] = a_b + t_b\boldsymbol{\beta}^T \mathbf{x} \quad (9.7)$$

for any $k \times 1$ constant vector \mathbf{b} where a_b and t_b are constants that depend on \mathbf{b} .

Proof. Suppose that the assumption of linearly related predictors holds. Then

$$E[\mathbf{b}^T \mathbf{x}|\boldsymbol{\beta}^T \mathbf{x}] = \mathbf{b}^T E[\mathbf{x}|\boldsymbol{\beta}^T \mathbf{x}] = \mathbf{b}^T \mathbf{a} + \mathbf{b}^T \mathbf{t}\boldsymbol{\beta}^T \mathbf{x}.$$

Thus the result holds with $a_b = \mathbf{b}^T \mathbf{a}$ and $t_b = \mathbf{b}^T \mathbf{t}$.

Now assume that Equation (9.7) holds. Take $\mathbf{b}_i = (0, \dots, 0, 1, 0, \dots, 0)^T$, the vector of zeroes except for a one in the i th position. Then Equation (9.6) holds since $E[\mathbf{x}|\boldsymbol{\beta}^T \mathbf{x}] = E[\mathbf{I}_k \mathbf{x}|\boldsymbol{\beta}^T \mathbf{x}] =$

$$E\left[\begin{pmatrix} \mathbf{b}_1^T \mathbf{x} \\ \vdots \\ \mathbf{b}_k^T \mathbf{x} \end{pmatrix} \mid \boldsymbol{\beta}^T \mathbf{x}\right] = \begin{pmatrix} a_1 + t_1\boldsymbol{\beta}^T \mathbf{x} \\ \vdots \\ a_k + t_k\boldsymbol{\beta}^T \mathbf{x} \end{pmatrix} \equiv \mathbf{a} + \mathbf{t}\boldsymbol{\beta}^T \mathbf{x}. \quad \square$$

Following Cook (1998a, p. 143-144), assume that there is an objective function

$$L_n(a, \mathbf{b}) = \frac{1}{n} \sum_{i=1}^n L(a + \mathbf{b}^T \mathbf{x}_i, Y_i) \quad (9.8)$$

where $L(u, v)$ is a bivariate function that is a convex function of the first argument u . Assume that the estimate $(\hat{a}, \hat{\mathbf{b}})$ of (a, \mathbf{b}) satisfies

$$(\hat{a}, \hat{\mathbf{b}}) = \arg \min_{a, \mathbf{b}} L_n(a, \mathbf{b}). \quad (9.9)$$

For example, the ordinary least squares (OLS) estimator uses

$$L(a + \mathbf{b}^T \mathbf{x}, Y) = (Y - a - \mathbf{b}^T \mathbf{x})^2.$$

Maximum likelihood type estimators such as those used to compute GLMs and Huber's M -estimator also work, as does the Wilcoxon rank estimator. Assume that the population analog (α^*, β^*) is the unique minimizer of $E[L(a + \mathbf{b}^T \mathbf{x}, Y)]$ where the expectation exists and is with respect to the joint distribution of $(Y, \mathbf{x}^T)^T$. For example, (α^*, β^*) is unique if $L(u, v)$ is strictly convex in its first argument. The following result is a useful extension of Brillinger (1977, 1983).

Theorem 9.2 (Li and Duan 1989, p. 1016): Assume that the \mathbf{x} are linearly related predictors, that $(Y_i, \mathbf{x}_i^T)^T$ are iid observations from some joint distribution with $\text{Cov}(\mathbf{x}_i)$ nonsingular. Assume $L(u, v)$ is convex in its first argument and that β^* is unique. Assume that $Y \perp\!\!\!\perp \mathbf{x} | \beta^T \mathbf{x}$. Then $\beta^* = c\beta$ for some scalar c .

Proof. See Li and Duan (1989) or Cook (1998a, p. 144).

Remark 9.1. This theorem basically means that if the 1D regression model is appropriate and if the condition of linearly related predictors holds, then the (e.g. OLS) estimator $\hat{\mathbf{b}} \equiv \hat{\beta}^* \approx c\beta$. Li and Duan (1989, p. 1031) show that under additional conditions, $(\hat{a}, \hat{\mathbf{b}})$ is asymptotically normal. In particular, the OLS estimator frequently has a \sqrt{n} convergence rate. If the OLS estimator $(\hat{\alpha}, \hat{\beta})$ satisfies $\hat{\beta} \approx c\beta$ when model (9.1) holds, then the response plot of

$$\hat{\alpha} + \hat{\beta}^T \mathbf{x} \text{ versus } Y$$

can be used to visualize the conditional distribution $Y | (\alpha + \beta^T \mathbf{x})$ provided that $c \neq 0$.

Remark 9.2. If $\hat{\mathbf{b}}$ is a consistent estimator of β^* , then certainly

$$\beta^* = c\mathbf{x}\beta + \mathbf{u}_g$$

where $\mathbf{u}_g = \beta^* - c\mathbf{x}\beta$ is the bias vector. Moreover, the bias vector $\mathbf{u}_g = \mathbf{0}$ if \mathbf{x} is elliptically contoured under the assumptions of Theorem 9.2. This result suggests that the bias vector might be negligible if the distribution of the predictors is close to being EC. **Often if no strong nonlinearities are present among the predictors**, the bias vector is small enough so that $\hat{\mathbf{b}}^T \mathbf{x}$ is a useful ESP.

Remark 9.3. Suppose that the 1D regression model is appropriate and $Y \perp\!\!\!\perp \mathbf{x} | \beta^T \mathbf{x}$. Then $Y \perp\!\!\!\perp \mathbf{x} | c\beta^T \mathbf{x}$ for any nonzero scalar c . If $Y = g(\beta^T \mathbf{x}, e)$ and both g and β are unknown, then $g(\beta^T \mathbf{x}, e) = h_{a,c}(a + c\beta^T \mathbf{x}, e)$ where

$$h_{a,c}(w, e) = g\left(\frac{w - a}{c}, e\right)$$

for $c \neq 0$. In other words, if g is unknown, we can estimate $c\beta$ but we can not determine c or β ; i.e., we can only estimate β up to a constant.

A very useful result is that if $Y = m(x)$ for some function m , then m can be visualized with both a plot of x versus Y and a plot of cx versus Y if $c \neq 0$. In fact, there are only three possibilities, if $c > 0$ then the two plots are nearly identical: except the labels of the horizontal axis change. (The two plots are usually not exactly identical since plotting controls to “fill space” depend on several factors and will change slightly.) If $c < 0$, then the plot appears to be flipped about the vertical axis. If $c = 0$, then $m(0)$ is a constant, and the plot is basically a dot plot. Similar results hold if $Y_i = g(\alpha + \beta^T \mathbf{x}_i, e_i)$ if the errors e_i are small. OLS often provides a useful estimator of $c\beta$ where $c \neq 0$, but OLS can result in $c = 0$ if g is symmetric about the population median of $\alpha + \beta^T \mathbf{x}$.

Definition 9.6. If the 1D regression model (9.1) holds with $h(\mathbf{x}) = \alpha + \beta^T \mathbf{x}$, and OLS is used, then the ESP may be called the *OLS ESP* and the response plot may be called the *OLS response plot*. Other estimators, such as SIR, may have similar labels.

Example 9.2. Suppose that $\mathbf{x}_i \sim N_3(\mathbf{0}, \mathbf{I}_3)$ and that

$$Y = m(\beta^T \mathbf{x}) + e = (x_1 + 2x_2 + 3x_3)^3 + e.$$

Then a 1D regression model $Y \perp\!\!\!\perp \mathbf{x} | \beta^T \mathbf{x}$ holds with $\beta = (1, 2, 3)^T$. Figure 9.1 shows the sufficient summary plot of $\beta^T \mathbf{x}$ versus Y , and Figure 9.2 shows the sufficient summary plot of $-\beta^T \mathbf{x}$ versus Y . Notice that the functional form m appears to be cubic in both plots and that both plots can be smoothed by eye or with a scatterplot smoother such as *lowess*. The two figures were generated with the following *R* commands.

```
X <- matrix(rnorm(300), nrow=100, ncol=3)
SP <- X%*%1:3
Y <- (SP)^3 + rnorm(100)
```

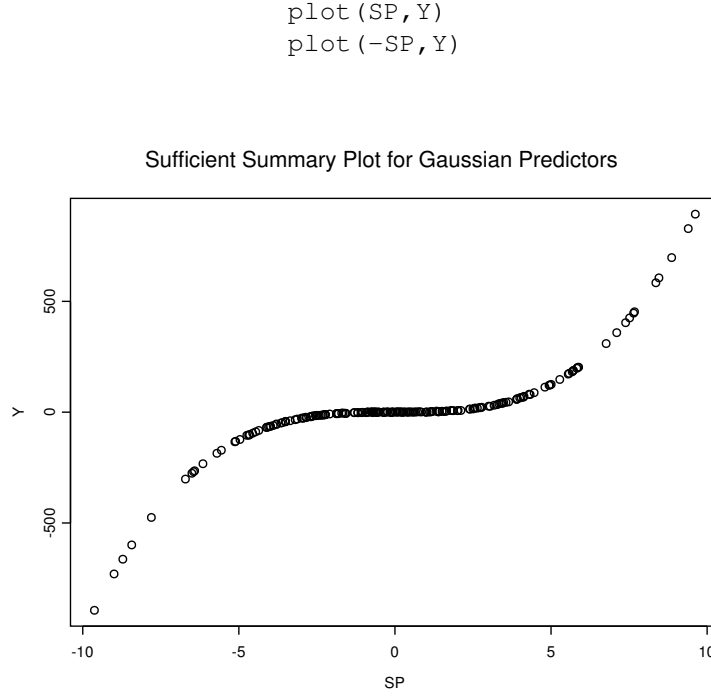


Fig. 9.1 SSP for $m(u) = u^3$

We particularly want to use the OLS estimator $(\hat{\alpha}, \hat{\beta})$ to produce an estimated sufficient summary plot. This estimator is obtained from the usual multiple linear regression of Y_i on \mathbf{x}_i , but *we are not assuming that the multiple linear regression model holds*; however, we are hoping that the 1D regression model $Y \perp\!\!\!\perp \mathbf{x} | \beta^T \mathbf{x}$ is a useful approximation to the data and that $\hat{\beta} \approx c\beta$ for some nonzero constant c . In addition to Theorem 9.2, nice results exist if the single index model is appropriate. Recall that

$$\text{Cov}(\mathbf{x}, \mathbf{Y}) = E[(\mathbf{x} - E(\mathbf{x}))(\mathbf{Y} - E(\mathbf{Y}))^T].$$

Definition 9.7. Suppose that $(Y_i, \mathbf{x}_i^T)^T$ are iid observations and that the positive definite $k \times k$ matrix $\text{Cov}(\mathbf{x}) = \Sigma_X$ and the $k \times 1$ vector $\text{Cov}(\mathbf{x}, Y) = \Sigma_{X,Y}$. Let the OLS estimator $(\hat{\alpha}, \hat{\beta})$ be computed from the multiple linear regression of Y on \mathbf{x} plus a constant. Then $(\hat{\alpha}, \hat{\beta})$ estimates the population quantity $(\alpha_{OLS}, \beta_{OLS})$ where

$$\beta_{OLS} = \Sigma_X^{-1} \Sigma_{X,Y}. \quad (9.10)$$

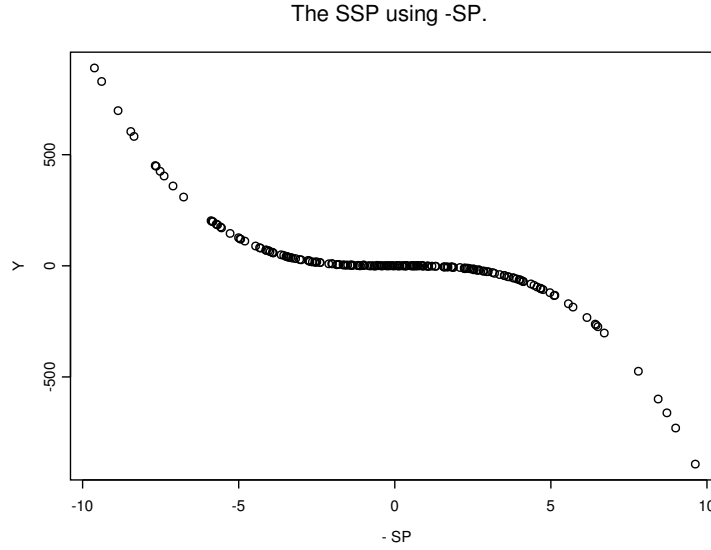


Fig. 9.2 Another SSP for $m(u) = u^3$

The following notation will be useful for studying the OLS estimator. Let the sufficient predictor $\mathbf{z} = \boldsymbol{\beta}^T \mathbf{x}$ and let $\mathbf{w} = \mathbf{x} - E(\mathbf{x})$. Let $\mathbf{r} = \mathbf{w} - (\boldsymbol{\Sigma}_X \boldsymbol{\beta}) \boldsymbol{\beta}^T \mathbf{w}$.

Theorem 9.3. In addition to the conditions of Definition 9.7, also assume that $Y_i = m(\boldsymbol{\beta}^T \mathbf{x}_i) + e_i$ where the zero mean constant variance iid errors e_i are independent of the predictors \mathbf{x}_i . Then

$$\boldsymbol{\beta}_{OLS} = \boldsymbol{\Sigma}_X^{-1} \boldsymbol{\Sigma}_{X,Y} = c_{m,X} \boldsymbol{\beta} + \mathbf{u}_{m,X} \quad (9.11)$$

where the scalar

$$c_{m,X} = E[\boldsymbol{\beta}^T (\mathbf{x} - E(\mathbf{x})) m(\boldsymbol{\beta}^T \mathbf{x})] \quad (9.12)$$

and the bias vector

$$\mathbf{u}_{m,X} = \boldsymbol{\Sigma}_X^{-1} E[m(\boldsymbol{\beta}^T \mathbf{x}) \mathbf{r}]. \quad (9.13)$$

Moreover, $\mathbf{u}_{m,X} = \mathbf{0}$ if \mathbf{x} is from an EC distribution with nonsingular $\boldsymbol{\Sigma}_X$, and $c_{m,X} \neq 0$ unless $\text{Cov}(\mathbf{x}, Y) = \mathbf{0}$. If the multiple linear regression model holds, then $c_{m,X} = 1$, and $\mathbf{u}_{m,X} = \mathbf{0}$.

The proof of the above result is outlined in Problem 9.2 using an argument due to Aldrin, Bølviken, and Schweder (1993). See related results in Stoker (1986) and Cook, Hawkins, and Weisberg (1992). If the 1D regression model is appropriate, then typically $\text{Cov}(\mathbf{x}, Y) \neq \mathbf{0}$ unless $\boldsymbol{\beta}^T \mathbf{x}$ follows a symmetric distribution and m is symmetric about the median of $\boldsymbol{\beta}^T \mathbf{x}$.

Definition 9.8. Let $(\hat{\alpha}, \hat{\beta})$ denote the OLS estimate obtained from the OLS multiple linear regression of Y on \mathbf{x} . The *OLS view* is a response plot of $a + \hat{\beta}^T \mathbf{x}$ versus Y . Typically $a = 0$ or $a = \hat{\alpha}$.

Remark 9.4. All of this awkward notation and theory leads to a remarkable result, perhaps first noted by Brillinger (1977, 1983) and called the *1D Estimation Result* by Cook and Weisberg (1999a, p. 432). The result is that if the 1D regression model $Y \perp \mathbf{x} | \beta^T \mathbf{x}$ is appropriate, then *the OLS view will frequently be a useful estimated sufficient summary plot* (ESSP). Hence the OLS predictor $\hat{\beta}^T \mathbf{x}$ is a useful *estimated sufficient predictor* (ESP).

Although the OLS view is frequently a good ESSP if no strong nonlinearities are present in the predictors and if $c_{m,X} \neq 0$ (e.g. the sufficient summary plot of $\beta^T \mathbf{x}$ versus Y is not approximately symmetric), even better estimated sufficient summary plots can be obtained by using ellipsoidal trimming. This topic is discussed in the following section and follows Olive (2002) closely.

9.2 Visualizing 1D Regression

Cook and Nachtsheim (1994) and Cook (1998a, p. 152) demonstrate that the bias $\mathbf{u}_{m,X}$ can often be made small by ellipsoidal trimming. To perform ellipsoidal trimming, an estimator (T, \mathbf{C}) is computed where T is a $k \times 1$ multivariate location estimator and \mathbf{C} is a $k \times k$ symmetric positive definite dispersion estimator. Then the i th squared Mahalanobis distance is the random variable

$$D_i^2 = (\mathbf{x}_i - T)^T \mathbf{C}^{-1} (\mathbf{x}_i - T) \quad (9.14)$$

for each vector of observed predictors \mathbf{x}_i . If the ordered distances $D_{(j)}$ are unique, then j of the \mathbf{x}_i are in the hyperellipsoid

$$\{\mathbf{x} : (\mathbf{x} - T)^T \mathbf{C}^{-1} (\mathbf{x} - T) \leq D_{(j)}^2\}. \quad (9.15)$$

The i th case $(Y_i, \mathbf{x}_i^T)^T$ is trimmed if $D_i > D_{(j)}$. Thus if $j \approx 0.9n$, then about 10% of the cases are trimmed.

We suggest that the estimator (T, \mathbf{C}) should be the classical sample mean and covariance matrix $(\bar{\mathbf{x}}, \mathbf{S})$ or a robust multivariate location and dispersion estimator such as RFCH. See Section 10.7. When $j \approx n/2$, the RFCH estimator attempts to make the volume of the hyperellipsoid given by Equation (9.15) small.

Ellipsoidal trimming seems to work for at least three reasons. The trimming divides the data into two groups: the *trimmed cases* and the *remaining cases* (\mathbf{x}_M, Y_M) where $M\%$ is the amount of trimming, e.g. $M = 10$ for 10% trimming. If the distribution of the predictors \mathbf{x} is EC then the distribution

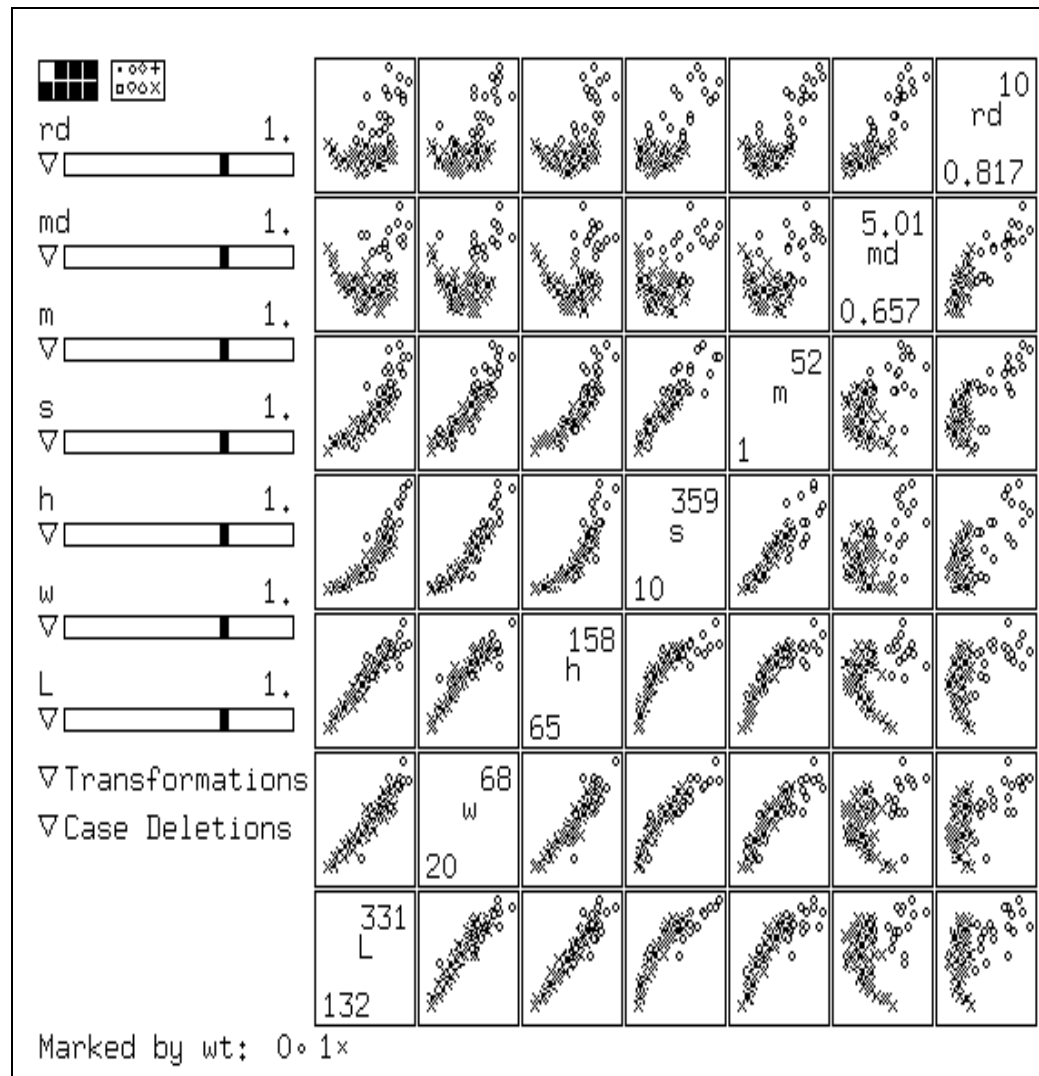


Fig. 9.3 Scatterplot for Mussel Data, o Corresponds to Trimmed Cases

of \mathbf{x}_M still retains enough symmetry so that the bias vector is approximately zero. If the distribution of \mathbf{x} is not EC, then the distribution of \mathbf{x}_M will often have enough symmetry so that the bias vector is small. In particular, trimming often removes strong nonlinearities from the predictors and the weighted predictor distribution is more nearly elliptically symmetric than the predictor distribution of the entire data set (Winsor's principle: "all data are roughly Gaussian in the middle"). Secondly, under heavy trimming, the mean function of the remaining cases may be more linear than the mean func-

tion of the entire data set. Thirdly, if $|c|$ is very large, then the bias vector may be small relative to $c\beta$. Trimming sometimes inflates $|c|$. From Theorem 9.3, any of these three reasons should produce a better estimated sufficient predictor.

For example, examine Figure 4.4. The data are not EC, but the data within the resistant covering ellipsoid are approximately EC.

Example 9.3. Cook and Weisberg (1999a, p. 351, 433, 447) gave a data set on 82 mussels sampled off the coast of New Zealand. The variables are the *muscle mass* M in grams, the *length* L and *height* H of the shell in mm, the *shell width* W and the *shell mass* S . The robust and classical Mahalanobis distances were calculated, and Figure 9.3 shows a scatterplot matrix of the mussel data, the RD_i 's, and the MD_i 's. Notice that many of the subplots are nonlinear. The cases marked by open circles were given weight zero by the FMCD algorithm, and the linearity of the retained cases has increased. Note that only one trimming proportion is shown and that a heavier trimming proportion would increase the linearity of the cases that were not trimmed.

The two ideas of using ellipsoidal trimming to reduce the bias and choosing a view with a smooth mean function and smallest variance function can be combined into a graphical method for finding the estimated sufficient summary plot and the estimated sufficient predictor. Trim the $M\%$ of the cases with the largest Mahalanobis distances, and then compute the OLS estimator $(\hat{\alpha}_M, \hat{\beta}_M)$ from the cases that remain. Use $M = 0, 10, 20, 30, 40, 50, 60, 70, 80$, and 90 to generate ten plots of $\hat{\beta}_M^T \mathbf{x}$ versus Y using all n cases. In analogy with the Cook and Weisberg procedure for visualizing 1D structure with two predictors, the plots will be called “trimmed views.” Notice that $M = 0$ corresponds to the OLS view.

Definition 9.9. The *best trimmed view* is the trimmed view with a smooth mean function and the smallest variance function and is the estimated sufficient summary plot. If $M^* = E$ is the percentage of cases trimmed that corresponds to the best trimmed view, then $\hat{\beta}_E^T \mathbf{x}$ is the estimated sufficient predictor.

The following examples illustrate the *R/Splus* function `trviews` that is used to produce the ESSP. If *R* is used instead of *Splus*, the command

```
library(MASS)
```

needs to be entered to access the function `cov.mcd` called by `trviews`. The function `trviews` is used in Problem 9.6. Also notice the `trviews` estimator is basically the same as the `tvreg` estimator described in Section 11.3. The `tvreg` estimator can be used to simultaneously detect whether the data is following a multiple linear regression model or some other single index model. Plot $\hat{\alpha}_E + \hat{\beta}_E^T \mathbf{x}$ versus Y and add the identity line. If the plotted points follow the identity line then the MLR model is reasonable, but if the

plotted points follow a nonlinear mean function, then a nonlinear single index model may be reasonable.

Example 9.2 continued. The command

```
trviews(X, Y)
```

produced the following output.

Intercept	X1	X2	X3
0.6701255	3.133926	4.031048	7.593501
Intercept	X1	X2	X3
1.101398	8.873677	12.99655	18.29054
Intercept	X1	X2	X3
0.9702788	10.71646	15.40126	23.35055
Intercept	X1	X2	X3
0.5937255	13.44889	23.47785	32.74164
Intercept	X1	X2	X3
1.086138	12.60514	25.06613	37.25504
Intercept	X1	X2	X3
4.621724	19.54774	34.87627	48.79709
Intercept	X1	X2	X3
3.165427	22.85721	36.09381	53.15153
Intercept	X1	X2	X3
5.829141	31.63738	56.56191	82.94031
Intercept	X1	X2	X3
4.241797	36.24316	70.94507	105.3816
Intercept	X1	X2	X3
6.485165	41.67623	87.39663	120.8251

The function generates 10 trimmed views. The first plot trims 90% of the cases while the last plot does not trim any of the cases and is the OLS view. To advance a plot, press the right button on the mouse (in *R*, highlight stop rather than continue). After all of the trimmed views have been generated, the output is presented. For example, the 5th line of numbers in the output corresponds to $\hat{\alpha}_{50} = 1.086138$ and $\hat{\beta}_{50}^T$ where 50% trimming was used. The second line of numbers corresponds to 80% trimming while the last line corresponds to 0% trimming and gives the OLS estimate $(\hat{\alpha}_0, \hat{\beta}_0^T) = (\hat{\alpha}, \hat{\beta})$. The trimmed views with 50% and 90% trimming were very good. We decided that the view with 50% trimming was the best. Hence $\hat{\beta}_E = (12.60514, 25.06613, 37.25504)^T \approx 12.5\beta$. The best view is shown in Figure 9.4 and is nearly identical to the sufficient summary plot shown in Figure 9.1. Notice that the OLS estimate $= (41.68, 87.40, 120.83)^T \approx 42\beta$. The OLS view is Figure 1.5 in Chapter 1, and is again very similar to the sufficient summary plot, but it is not quite as smooth as the best trimmed view.

The plot of the estimated sufficient predictor versus the sufficient predictor is also informative. Of course this plot can usually only be generated for

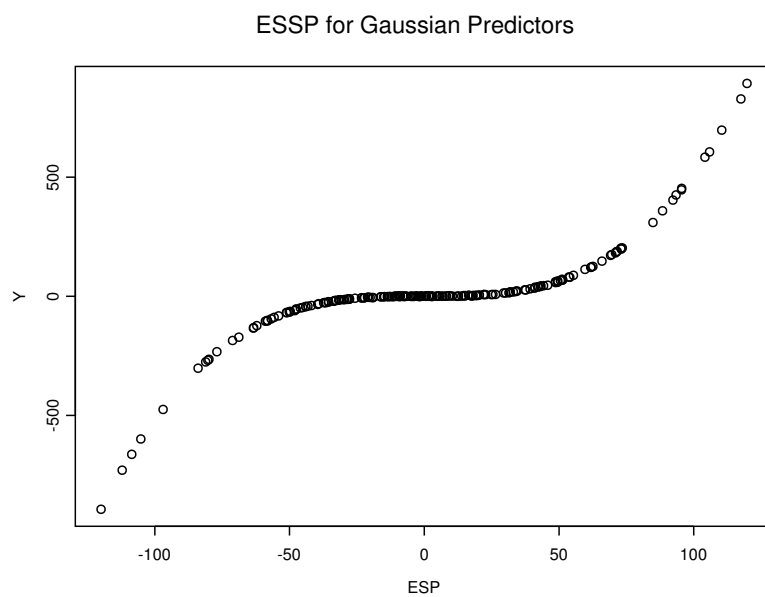


Fig. 9.4 Best View for Estimating $m(u) = u^3$

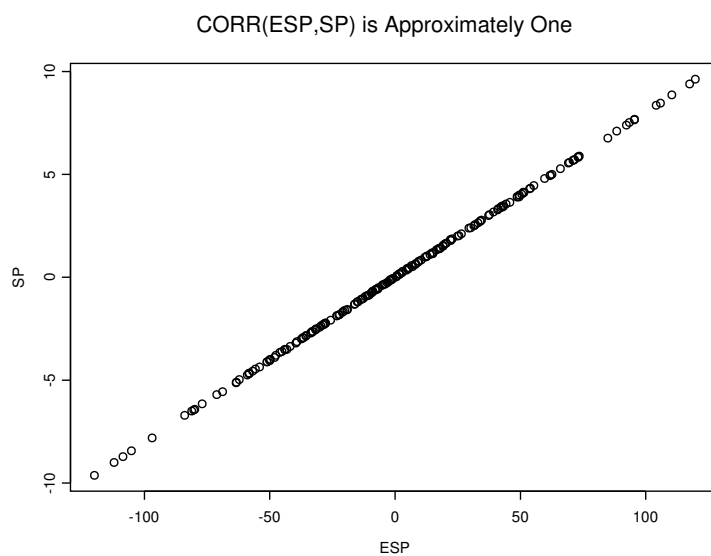


Fig. 9.5 The angle between the SP and the ESP is nearly zero.

simulated data since β is generally unknown. If the plotted points are highly correlated (with $|\text{corr}(\text{ESP}, \text{SP})| > 0.95$) and follow a line through the origin, then the estimated sufficient summary plot is nearly as good as the sufficient summary plot. The simulated data used $\beta = (1, 2, 3)^T$, and the commands

```
SP <- X %**% 1:3
ESP <- X %**% c(12.60514, 25.06613, 37.25504)
plot(ESP, SP)
```

generated the plot shown in Figure 9.5.

Example 9.4. An artificial data set with 200 trivariate vectors \mathbf{x}_i was generated. The marginal distributions of $x_{i,j}$ are iid lognormal for $j = 1, 2$, and 3. Since the response $Y_i = \sin(\beta^T \mathbf{x}_i) / \beta^T \mathbf{x}_i$ where $\beta = (1, 2, 3)^T$, the random vector \mathbf{x}_i is not elliptically contoured and the function m is strongly nonlinear. Figure 9.6d shows the OLS view and Figure 9.7d shows the best trimmed view. Notice that it is difficult to visualize the mean function with the OLS view, and notice that the correlation between Y and the ESP is very low. By focusing on a part of the data where the correlation is high, it may be possible to improve the estimated sufficient summary plot. For example, in Figure 9.7d, temporarily omit cases that have ESP less than 0.3 and greater than 0.75. From the untrimmed cases, obtained the ten trimmed estimates $\hat{\beta}_{90}, \dots, \hat{\beta}_0$. Then using *all of the data*, obtain the ten views. The best view could be used as the ESSP.

Application 9.1. Suppose that a 1D regression analysis is desired on a data set, use the trimmed views as an exploratory data analysis technique to visualize the conditional distribution $Y | \beta^T \mathbf{x}$. The best trimmed view is an estimated sufficient summary plot. If the single index model (9.3) holds, the function m can be estimated from this plot using parametric models or scatterplot smoothers such as `lowess`. Notice that Y can be predicted visually using *up and over lines*.

Table 9.1 Estimated Sufficient Predictors Coefficients Estimating $c(1, 2, 3)^T$

method	b_1	b_2	b_3
OLS View	0.0032	0.0011	0.0047
90% Trimmed OLS View	0.086	0.182	0.338
SIR View	-0.394	-0.361	-0.845
10% Trimmed SIR VIEW	-0.284	-0.473	-0.834
SAVE View	-1.09	0.870	-0.480
40% Trimmed SAVE VIEW	0.256	0.591	0.765
PHD View	-0.072	-0.029	-0.0097
90% Trimmed PHD VIEW	-0.558	-0.499	-0.664
LMSREG VIEW	-0.003	-0.005	-0.059
70% Trimmed LMSREG VIEW	0.143	0.287	0.428

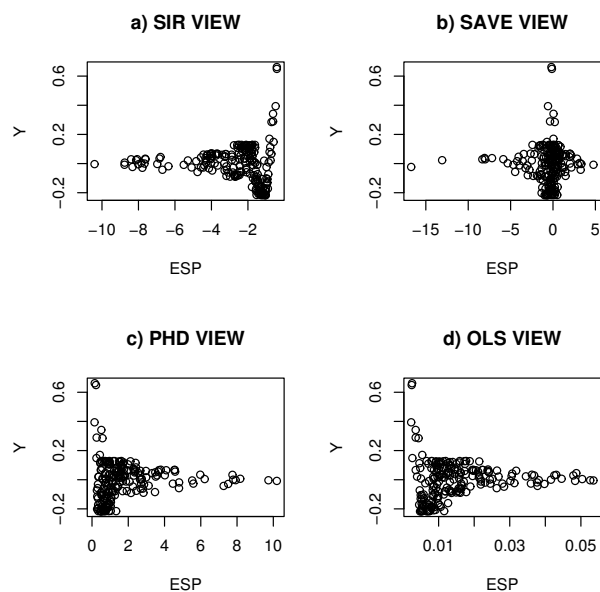


Fig. 9.6 Estimated Sufficient Summary Plots Without Trimming

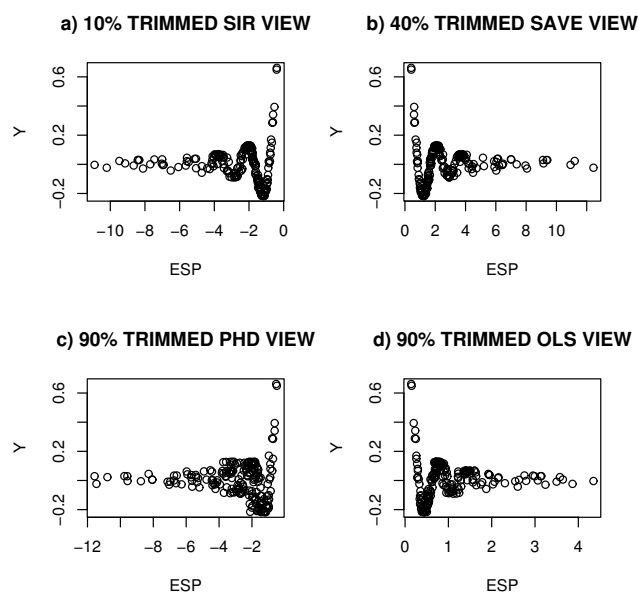


Fig. 9.7 1D Regression with Trimmed Views

Application 9.2. The best trimmed view can also be used as a diagnostic for linearity and monotonicity.

For example in Figure 9.4, if $\text{ESP} = 0$, then $\hat{Y} = 0$ and if $\text{ESP} = 100$, then $\hat{Y} = 500$. Figure 9.4 suggests that the mean function is monotone but not linear, and Figure 9.7 suggests that the mean function is neither linear nor monotone.

Application 9.3. Assume that a known 1D regression model is assumed for the data. Then the best trimmed view is a model checking plot and can be used as a diagnostic for whether the assumed model is appropriate.

The trimmed views are sometimes useful even when the assumption of linearly related predictors fails. Cook and Li (2002) summarize when competing methods such as the OLS view, sliced inverse regression (SIR), principal Hessian directions (PHD), and sliced average variance estimation (SAVE) can fail. All four methods frequently perform well if there are no strong nonlinearities present in the predictors.

Example 9.4 (continued). Figure 9.6 shows that the response plots for SIR, PHD, SAVE, and OLS are not very good while Figure 9.7 shows that trimming improved the SIR, SAVE and OLS methods.

One goal for future research is to develop better methods for visualizing 1D regression. Trimmed views seem to become less effective as the number of predictors $k = p - 1$ increases. Consider the sufficient predictor $\text{SP} = x_1 + \cdots + x_k$. With the $\sin(\text{SP})/\text{SP}$ data, several trimming proportions gave good views with $k = 3$, but only one of the ten trimming proportions gave a good view with $k = 10$. In addition to problems with dimension, it is not clear which covariance estimator and which regression estimator should be used. We suggest using the RFCH estimator with OLS, and preliminary investigations suggest that the classical covariance estimator gives better estimates than `cov.mcd`. But among the many *Splus* regression estimators, `lmsreg` often worked well. Theorem 9.2 suggests that strictly convex regression estimators such as OLS will often work well, but there is no theory for the robust regression estimators.

Example 9.4 continued. Replacing the OLS trimmed views by alternative MLR estimators often produced good response plots, and for single index models, the `lmsreg` estimator often worked the best. Figure 9.8 shows a scatterplot matrix of Y , ESP and SP where the sufficient predictor $\text{SP} = \beta^T \mathbf{x}$. The ESP used ellipsoidal trimming with `cov.mcd` and with `lmsreg` instead of OLS. The top row of Figure 9.8 shows that the estimated sufficient summary plot and the sufficient summary plot are nearly identical. Also the correlation of the ESP and the SP is nearly one. Table 9.1 shows the estimated sufficient predictor coefficients \mathbf{b} when the sufficient predictor coefficients are $c(1, 2, 3)^T$. Only the SIR, SAVE, OLS and `lmsreg` trimmed

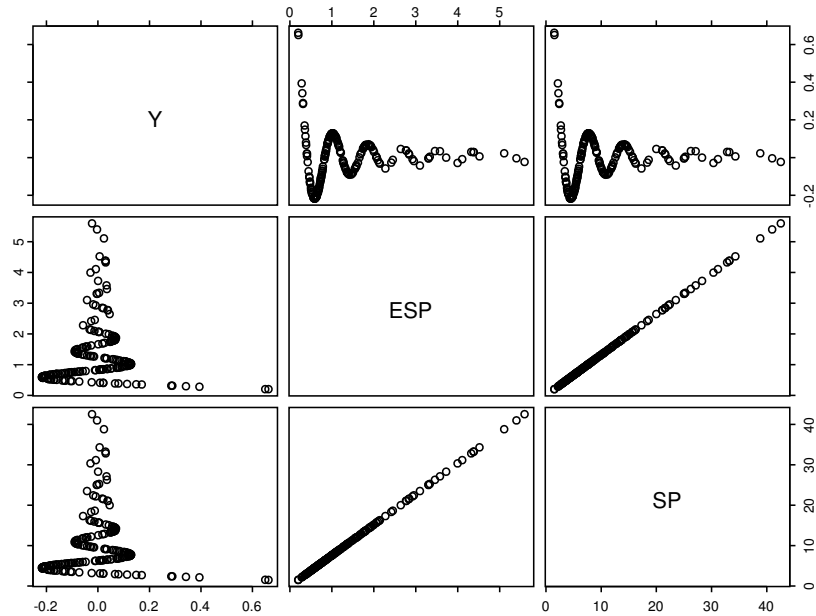


Fig. 9.8 1D Regression with `lmsreg`

views produce estimated sufficient predictors that are highly correlated with the sufficient predictor.

Figure 9.9 helps illustrate why ellipsoidal trimming works. This view used 70% trimming and the open circles denote cases that were trimmed. The highlighted squares correspond to the cases $(\mathbf{x}_{70}, Y_{70})$ that were not trimmed. Note that the highlighted cases are far more linear than the data set as a whole. Also `lmsreg` will give half of the highlighted cases zero weight, further linearizing the function. In Figure 9.9, the `lmsreg` constant $\hat{\alpha}_{70}$ is included, and the plot is simply the response plot of the weighted `lmsreg` fitted values versus Y . The vertical deviations from the line through the origin are the “residuals” $Y_i - \hat{\alpha}_{70} - \hat{\beta}_{70}^T \mathbf{x}$ and at least half of the highlighted cases have small residuals.

9.3 Predictor Transformations

Even if the multiple linear regression model is valid, a model based on a subset of the predictor variables depends on the predictor distribution. If the predictors are linearly related (e.g. EC), then the submodel mean function

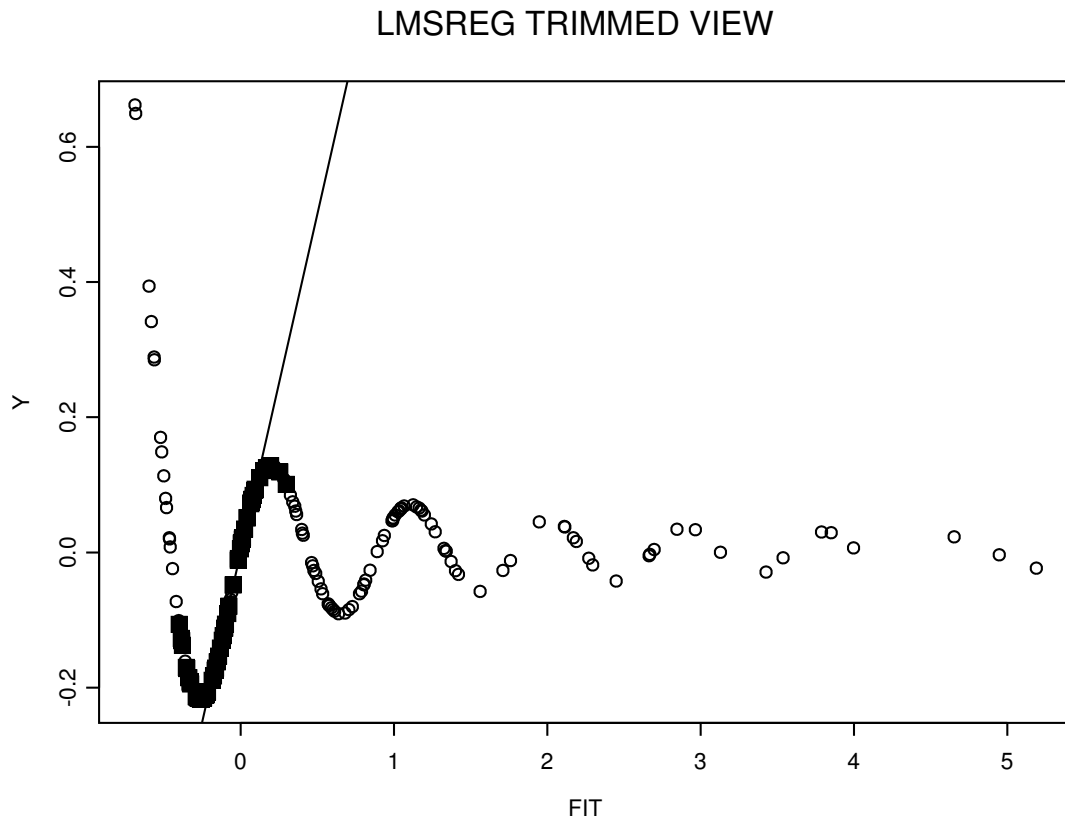


Fig. 9.9 The Weighted `lmsreg` Fitted Values Versus Y

is often well behaved, but otherwise the submodel mean function could be nonlinear and the submodel variance function could be nonconstant.

For 1D regression models, the presence of strong nonlinearities among the predictors can invalidate inferences. A necessary condition for \mathbf{x} to have an EC distribution (or for no strong nonlinearities to be present among the predictors) is for each marginal plot of the scatterplot matrix of the predictors to have a linear or ellipsoidal shape if n is large.

One of the most useful techniques in regression is to remove gross nonlinearities in the predictors by using predictor transformations. Power transformations are particularly effective. A multivariate version of the Box–Cox transformation due to Velilla (1993) can cause the distribution of the transformed predictors to be closer to multivariate normal, and the Cook and Nachtsheim (1994) procedure can cause the distribution to be closer to ellip-

tical symmetry. Marginal Box-Cox transformations also seem to be effective. Power transformations can also be selected with slider bars in *Arc*.

Sections 5.1 gives several rules for predictor transformations, including the unit rule, log rule, and ladder rule. As an illustration of the ladder rule and log rule, in Figure 9.14c, small values of Y and large values of FESP need spreading, and using $\log(Y)$ would make the plot more linear.

9.4 Variable Selection

A standard problem in 1D regression is variable selection, also called subset or model selection. Assume that $Y \sim \mathbf{x} | (\alpha + \boldsymbol{\beta}^T \mathbf{x})$, that a constant is always included, and that $\mathbf{x} = (x_1, \dots, x_{p-1})^T$ are the $p - 1$ nontrivial predictors, which we assume to be of full rank. Then *variable selection* is a search for a subset of predictor variables that can be deleted without important loss of information. This section follows Olive and Hawkins (2005) closely.

Variable selection for the 1D regression model is very similar to variable selection for the multiple linear regression model (see Section 5.3). To clarify ideas, assume that there exists a subset S of predictor variables such that if \mathbf{x}_S is in the 1D model, then none of the other predictors are needed in the model. Write E for these ('extraneous') variables not in S , partitioning $\mathbf{x} = (\mathbf{x}_S^T, \mathbf{x}_E^T)^T$. Then

$$SP = \alpha + \boldsymbol{\beta}^T \mathbf{x} = \alpha + \boldsymbol{\beta}_S^T \mathbf{x}_S + \boldsymbol{\beta}_E^T \mathbf{x}_E = \alpha + \boldsymbol{\beta}_S^T \mathbf{x}_S. \quad (9.16)$$

The extraneous terms that can be eliminated given that the subset S is in the model have zero coefficients.

Now suppose that I is a candidate subset of predictors, that $S \subseteq I$ and that O is the set of predictors not in I . Then

$$SP = \alpha + \boldsymbol{\beta}^T \mathbf{x} = \alpha + \boldsymbol{\beta}_S^T \mathbf{x}_S = \alpha + \boldsymbol{\beta}_S^T \mathbf{x}_S + \boldsymbol{\beta}_{(I/S)}^T \mathbf{x}_{I/S} + \mathbf{0}^T \mathbf{x}_O = \alpha + \boldsymbol{\beta}_I^T \mathbf{x}_I,$$

(if I includes predictors from E , these will have zero coefficients). For any subset I that contains the subset S of relevant predictors, the correlation

$$\text{corr}(\alpha + \boldsymbol{\beta}^T \mathbf{x}_i, \alpha + \boldsymbol{\beta}_I^T \mathbf{x}_{I,i}) = 1. \quad (9.17)$$

This observation, which is true regardless of the explanatory power of the model, suggests that variable selection for 1D regression models is simple in principle. For each value of $j = 1, 2, \dots, p - 1$ nontrivial predictors, keep track of subsets I that provide the largest values of $\text{corr}(\text{ESP}, \text{ESP}(I))$. Any such subset for which the correlation is high is worth closer investigation and consideration. To make this advice more specific, use the *rule of thumb* that a candidate subset of predictors I is worth considering if the sample correlation

of ESP and $ESP(I)$ satisfies

$$\text{corr}(\tilde{\alpha} + \tilde{\beta}^T \mathbf{x}_i, \tilde{\alpha}_I + \tilde{\beta}_I^T \mathbf{x}_{I,i}) = \text{corr}(\tilde{\beta}^T \mathbf{x}_i, \tilde{\beta}_I^T \mathbf{x}_{I,i}) \geq 0.95. \quad (9.18)$$

The difficulty with this approach is that fitting all of the possible submodels involves substantial computation. An exception to this difficulty is multiple linear regression where there are efficient “leaps and bounds” algorithms for searching all subsets when OLS is used (see Furnival and Wilson 1974). Since OLS often gives a useful ESP, the following all subsets procedure can be used for 1D models when $p \leq 20$. Forward selection and backward elimination can be for much larger values of p .

- Fit a full model using the methods appropriate to that 1D problem to find the ESP $\hat{\alpha} + \hat{\beta}^T \mathbf{x}$.
- Find the OLS ESP $\hat{\alpha}_{OLS} + \hat{\beta}_{OLS}^T \mathbf{x}$.
- If the 1D ESP and the OLS ESP have “a strong linear relationship” (for example $|\text{corr}(\text{ESP}, \text{OLS ESP})| \geq 0.95$), then infer that the 1D problem is one in which OLS may serve as an adequate surrogate for the correct 1D model fitting procedure.
- Use computationally fast OLS variable selection procedures such as forward selection, backward elimination and the leaps and bounds all subsets algorithm along with the Mallows (1973) C_p criterion to identify predictor subsets I containing k variables (including the constant) with $C_p(I) \leq \min(2k, p)$.
- Perform a final check on the subsets that satisfy the C_p screen by using them to fit the 1D model.

For a 1D model, the response, ESP and vertical discrepancies $V = Y - ESP$ are important. When the multiple linear regression (MLR) model holds, the fitted values are the ESP: $\hat{Y} = ESP$, and the vertical discrepancies are the residuals.

Definition 9.10. a) The plot of $\tilde{\alpha}_I + \tilde{\beta}_I^T \mathbf{x}_{I,i}$ versus $\tilde{\alpha} + \tilde{\beta}^T \mathbf{x}_i$ is called an *EE plot* (also called an FF plot for MLR).

b) The plot of discrepancies $Y_i - \tilde{\alpha}_I - \tilde{\beta}_I^T \mathbf{x}_{I,i}$ versus $Y_i - \tilde{\alpha} - \tilde{\beta}^T \mathbf{x}_i$ is called a *VV plot* (also called an RR plot for MLR).

c) The plots of $\tilde{\alpha}_I + \tilde{\beta}_I^T \mathbf{x}_{I,i}$ versus Y_i and of $\tilde{\alpha} + \tilde{\beta}^T \mathbf{x}_i$ versus Y_i are called *estimated sufficient summary plots* or *response plots*.

Many numerical methods such as forward selection, backward elimination, stepwise and all subset methods using the C_p criterion (Jones 1946, Mallows 1973), have been suggested for variable selection. The four plots in Definition 9.10 contain valuable information to supplement the raw numerical results of these selection methods. Particular uses include:

- The key to understanding which plots are the most useful is the observation that a plot of w versus z is used to *visualize the conditional distribution of z given w* . Since the 1D regression is the study of the conditional distribution of Y given $\alpha + \beta^T \mathbf{x}$, the response plot is used to visualize this conditional distribution and should always be made. A major problem with variable selection is that deleting important predictors can change the functional form of the model. In particular, if a multiple linear regression model is appropriate for the full model, linearity may be destroyed if important predictors are deleted. When the single index model (9.3) holds, m can be visualized with a response plot. Adding visual aids such as the estimated parametric mean function $m(\hat{\alpha} + \hat{\beta}^T \mathbf{x})$ can be useful. If an estimated nonparametric mean function $\hat{m}(\hat{\alpha} + \hat{\beta}^T \mathbf{x})$ such as lowess follows the parametric curve closely, then often numerical goodness of fit tests will suggest that the model is good. See Chambers, Cleveland, Kleiner, and Tukey (1983, p. 280) and Cook and Weisberg (1999a, p. 425, 432). For variable selection, *the response plots from the full model and submodel should be very similar if the submodel is good*.
- Sometimes outliers will influence numerical methods for variable selection. Outliers tend to stand out in at least one of the plots. An EE plot is useful for variable selection because the correlation of $\text{ESP}(I)$ and ESP is important. The EE plot can be used to quickly check that the correlation is high, that the plotted points fall about some line, that the line is the identity line, and that the correlation is high because the relationship is linear, rather than because of outliers.
- Numerical methods may include too many predictors. Investigators can examine the p-values for individual predictors, but the assumptions needed to obtain valid p-values are often violated; however, the OLS t tests for individual predictors are meaningful since deleting a predictor changes the C_p value by $t^2 - 2$ where t is the test statistic for the predictor. See Section 9.5, Daniel and Wood (1980, p. 100-101) and the following two remarks.

Remark 9.5. Variable selection with the C_p criterion is closely related to the partial F test that uses test statistic F_I . Suppose that the full model contains p predictors including a constant and the submodel I includes k predictors including a constant. If $n \geq 10p$, then the submodel I is “interesting” if $C_p(I) \leq \min(2k, p)$.

To see this claim notice that *the following results are properties of OLS and hold even if the data does not follow a 1D model*. If the candidate model of \mathbf{x}_I has k terms (including the constant), then

$$F_I = \frac{SSE(I) - SSE}{(n - k) - (n - p)} \bigg/ \frac{SSE}{n - p} = \frac{n - p}{p - k} \left[\frac{SSE(I)}{SSE} - 1 \right]$$

where SSE is the “residual” sum of squares from the full model and $SSE(I)$ is the “residual” sum of squares from the candidate submodel. Then

$$C_p(I) = \frac{SSE(I)}{MSE} + 2k - n = (p - k)(F_I - 1) + k \quad (9.19)$$

where MSE is the “residual” mean square for the full model. Let $ESP(I) = \hat{\alpha}_I + \hat{\beta}_I^T \mathbf{x}_I$ be the ESP for the submodel and let $V_I = Y - ESP(I)$ so that $V_{I,i} = Y_i - \hat{\alpha}_I - \hat{\beta}_I^T \mathbf{x}_{I,i}$. Let ESP and V denote the corresponding quantities for the full model. Using Proposition 5.1 and Remarks 5.1 and 5.2 with $\text{corr}(r, r_I)$ replaced by $\text{corr}(V, V_I)$, it can be shown that

$$\text{corr}(V, V_I) = \sqrt{\frac{SSE}{SSE(I)}} = \sqrt{\frac{n - p}{C_p(I) + n - 2k}} = \sqrt{\frac{n - p}{(p - k)F_I + n - p}}.$$

It can also be shown that $C_p(I) \leq 2k$ corresponds to $\text{corr}(V, V_I) \geq d_n$ where

$$d_n = \sqrt{1 - \frac{p}{n}}.$$

Notice that for a fixed value of k , the submodel I_k that minimizes $C_p(I)$ also maximizes $\text{corr}(V, V_I)$. If $C_p(I) \leq 2k$ and $n \geq 10p$, then $0.948 \leq \text{corr}(V, V_I)$, and both $\text{corr}(V, V_I) \rightarrow 1.0$ and $\text{corr}(\text{OLS ESP}, \text{OLS ESP}(I)) \rightarrow 1.0$ as $n \rightarrow \infty$. Hence the plotted points in both the VV plot and the EE plot will cluster about the identity line (see Proposition 5.1 vi).

Remark 9.6. Suppose that the OLS ESP and the standard ESP are highly correlated: $|\text{corr}(\text{ESP}, \text{OLS ESP})| \geq 0.95$. Then often OLS variable selection can be used for the 1D data, and using the p-values from OLS output seems to be a useful benchmark. To see this, suppose that $n > 5p$ and first consider the model I_i that deletes the predictor X_i . Then the model has $k = p - 1$ predictors including the constant, and the test statistic is t_i where

$$t_i^2 = F_{I_i}.$$

Using (9.19) and $C_p(I_{full}) = p$, it can be shown that

$$C_p(I_i) = C_p(I_{full}) + (t_i^2 - 2).$$

Using the screen $C_p(I) \leq \min(2k, p)$ suggests that the predictor X_i should not be deleted if

$$|t_i| > \sqrt{2} \approx 1.414.$$

If $|t_i| < \sqrt{2}$ then the predictor can probably be deleted since C_p decreases.

More generally, it can be shown that $C_p(I) \leq 2k$ iff

$$F_I \leq \frac{p}{p - k}.$$

Now k is the number of terms in the model including a constant while $p - k$ is the number of terms set to 0. As $k \rightarrow 0$, the partial F test will reject $H_0: \beta_O = \mathbf{0}$ (i.e., say that the full model should be used instead of the submodel I) unless F_I is not much larger than 1. If p is very large and $p - k$ is very small, then the partial F test will tend to suggest that there is a model I that is about as good as the full model even though model I deletes $p - k$ predictors.

The $C_p(I) \leq k$ screen tends to overfit. We simulated multiple linear regression and single index model data sets with $p = 8$ and $n = 50, 100, 1000$ and 10000 . The true model S satisfied $C_p(S) \leq k$ for about 60% of the simulated data sets, but S satisfied $C_p(S) \leq 2k$ for about 97% of the data sets.

In many settings, not all of which meet the Li–Duan sufficient conditions, the full model OLS ESP is a good estimator of the sufficient predictor. If the fitted full 1D model $Y \mid \mathbf{x} | (\alpha + \beta^T \mathbf{x})$ is a useful approximation to the data and if $\hat{\beta}_{OLS}$ is a good estimator of $c\beta$ where $c \neq 0$, then a subset I will produce a response plot similar to the response plot of the full model if $\text{corr}(\text{OLS ESP}, \text{OLS ESP}(I)) \geq 0.95$. Hence the response plots based on the full and submodel ESP can both be used to visualize the conditional distribution of Y .

Assuming that a 1D model holds, a common assumption made for variable selection is that the fitted full model ESP is a good estimator of the sufficient predictor, and the usual numerical and graphical checks on this assumption should be made. To see that this assumption is weaker than the assumption that the OLS ESP is good, notice that if a 1D model holds but $\hat{\beta}_{OLS}$ estimates $c\beta$ where $c = 0$, then the $C_p(I)$ criterion could wrongly suggest that all subsets I have $C_p(I) \leq 2k$. Hence we also need to check that $c \neq 0$.

There are several methods for checking the OLS ESP, including: a) if an ESP from an alternative fitting method is believed to be useful, check that the ESP and the OLS ESP have a strong linear relationship: for example that $|\text{corr}(\text{ESP}, \text{OLS ESP})| \geq 0.95$. b) Often examining the OLS response plot shows that a 1D model is reasonable. For example, if the data are tightly clustered about a smooth curve, then a single index model may be appropriate. c) Verify that a 1D model is appropriate using graphical techniques given by Cook and Weisberg (1999a, p. 434-441). d) Verify that \mathbf{x} has an EC distribution with nonsingular covariance matrix and that the mean function $m(\alpha + \beta^T \mathbf{x})$ is not symmetric about the median of the distribution of $\alpha + \beta^T \mathbf{x}$. Then results from Li and Duan (1989) suggest that $c \neq 0$.

Condition a) is both the most useful (being a direct performance check) and the easiest to check. A standard fitting method should be used when available (e.g., for parametric 1D models such as GLMs). Conditions c) and d) need \mathbf{x} to have a continuous multivariate distribution while the predictors can be factors for a) and b). Using trimmed views results in an ESP that can sometimes cause condition b) to hold when d) is violated.

To summarize, variable selection procedures, originally meant for MLR, can often be used for 1D data. If the fitted full 1D model $Y \approx x|(\alpha + \beta^T x)$ is a useful approximation to the data and if $\hat{\beta}_{OLS}$ is a good estimator of $c\beta$ where $c \neq 0$, then a subset I is good if $\text{corr}(\text{OLS ESP}, \text{OLS ESP}(I)) \geq 0.95$. If n is large enough, Remark 9.5 implies that this condition will hold if $C_p(I) \leq 2k$ or if $F_I \leq 1$. This result suggests that within the (large) subclass of 1D models where the OLS ESP is useful, the OLS partial F test is robust (asymptotically) to model misspecifications in that $F_I \leq 1$ correctly suggests that submodel I is good. The OLS t tests for individual predictors are also meaningful since if $|t| < \sqrt{2}$ then the predictor can probably be deleted since C_p decreases while if $|t| \geq 2$ then the predictor is probably useful even when the other predictors are in the model. Section 9.5 provides related theory, and the following examples help illustrate the above discussion.

Example 9.5. This example illustrates that the plots are useful for general 1D regression models such as the response transformation model. Consider the data set on 82 mussels in Example 9.3. The response Y is the *muscle mass* in grams, and the four predictors are the *logarithms of the shell length, width, height and mass*. The logarithm transformation was used to remove strong nonlinearities that were evident in a scatterplot matrix of the untransformed predictors. The C_p criterion suggests using $\log(\text{width})$ and $\log(\text{shell mass})$ as predictors. The EE and VV plots are shown in Figure 9.10ab. The response

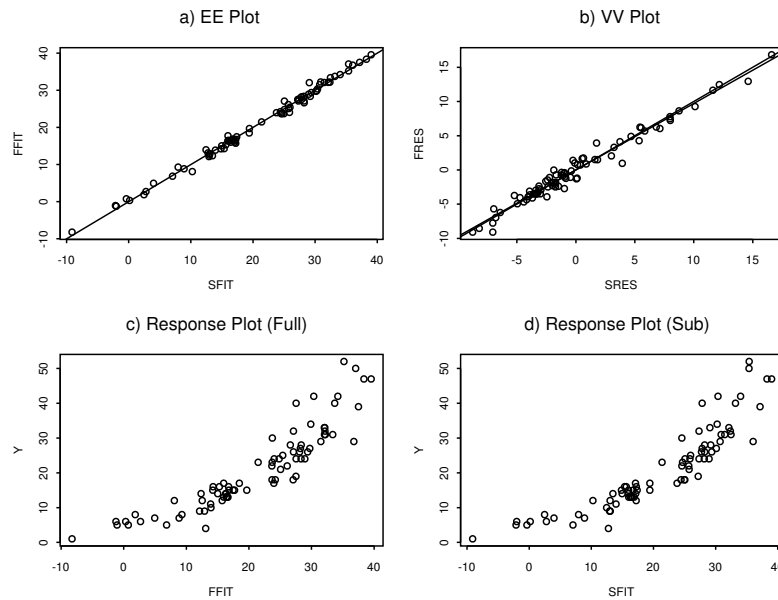


Fig. 9.10 Mussel Data with Muscle Mass as the Response

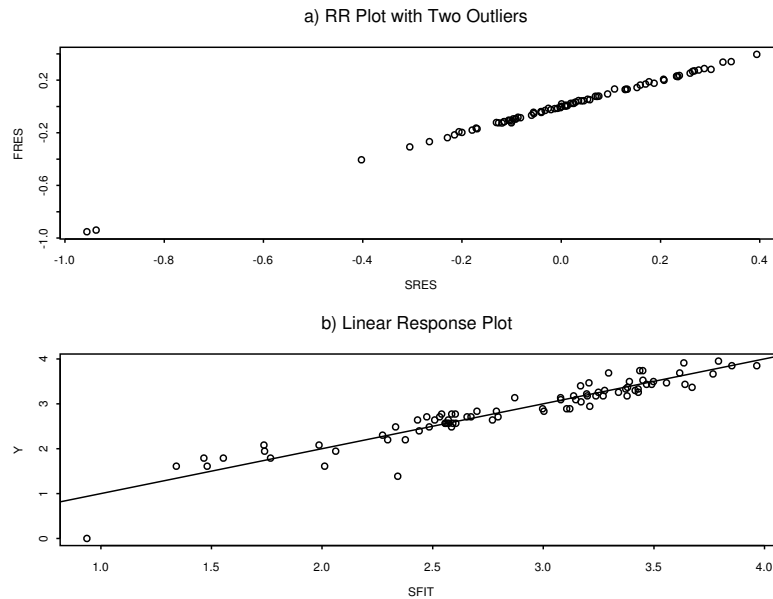


Fig. 9.11 Mussel Data with $\log(\text{Muscle Mass})$ as the Response

plots based on the full and submodel are shown in Figure 9.10cd and are nearly identical, but not linear.

When $\log(\text{muscle mass})$ is used as the response, the C_p criterion suggests using $\log(\text{height})$ and $\log(\text{shell mass})$ as predictors (the correlation between $\log(\text{height})$ and $\log(\text{width})$ is very high). Figure 9.11a shows the RR plot and 2 outliers are evident. These outliers correspond to the two cases with large negative residuals in the response plot shown in Figure 9.11b. After deleting the outliers, the C_p criterion still suggested using $\log(\text{height})$ and $\log(\text{shell mass})$ as predictors. The p-value for including $\log(\text{height})$ in the model was 0.03, and making the FF and RR plots after deleting $\log(\text{height})$ suggests that $\log(\text{height})$ may not be needed in the model.

Example 9.6. According to Li (1997), the predictors in the Boston housing data of Harrison and Rubinfeld (1978) have a nonlinear quasi-helix relationship which can cause regression graphics methods to fail. Nevertheless, the graphical diagnostics can be used to gain interesting information from the data. The response $Y = \log(\text{CRIM})$ where CRIM is the per capita crime rate by town. The predictors used were $x_1 =$ proportion of residential land zoned for lots over 25,000 sq.ft., $\log(x_2)$ where x_2 is the proportion of non-retail business acres per town, $x_3 =$ Charles River dummy variable (= 1 if tract bounds river; 0 otherwise), $x_4 = \text{NOX} =$ nitric oxides concentration (parts per 10 million), $x_5 =$ average number of rooms per dwelling, $x_6 =$

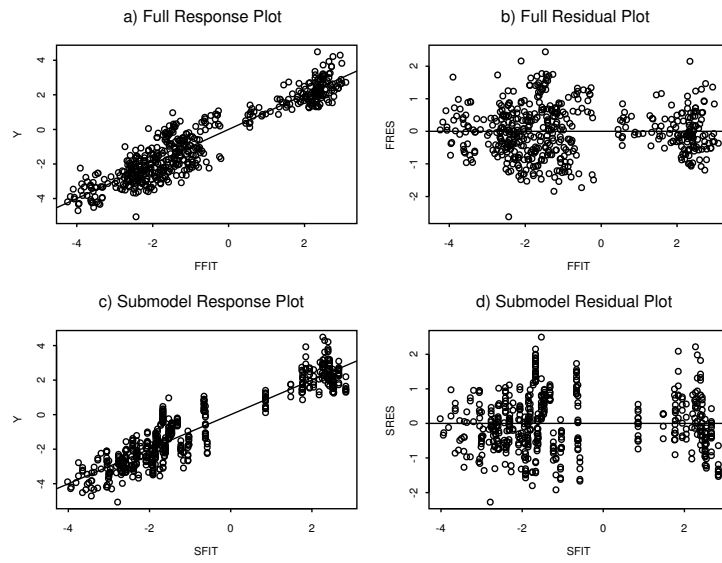


Fig. 9.12 Response and Residual Plots for Boston Housing Data

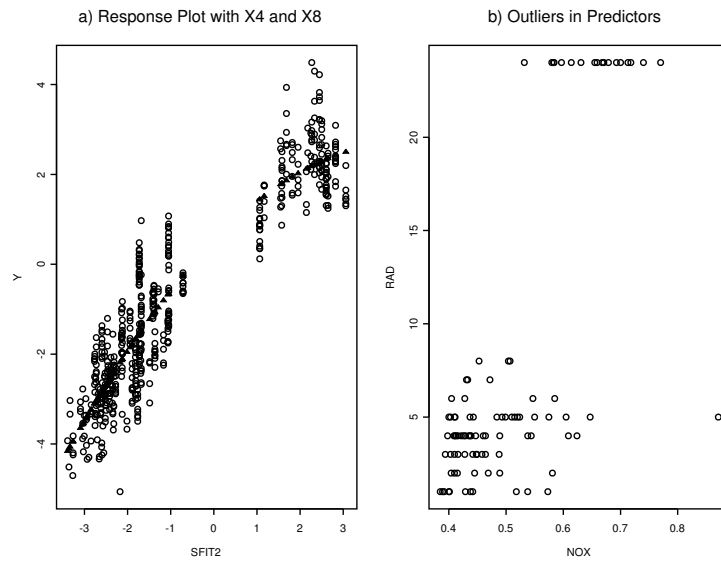


Fig. 9.13 Relationships between NOX, RAD and $Y = \log(\text{CRIM})$

proportion of owner-occupied units built prior to 1940, $\log(x_7)$ where x_7 = weighted distances to five Boston employment centers, $x_8 = RAD$ = index of accessibility to radial highways, $\log(x_9)$ where x_9 = full-value property-tax rate per \$10,000, x_{10} = pupil-teacher ratio by town, $x_{11} = 1000(Bk - 0.63)^2$ where Bk is the proportion of blacks by town, $\log(x_{12})$ where x_{12} = % lower status of the population, and $\log(x_{13})$ where x_{13} = median value of owner-occupied homes in \$1000's. The full model has 506 cases and 13 nontrivial predictor variables.

Figure 9.12ab shows the response plot and residual plot for the full model. The residual plot suggests that there may be three or four groups of data, but a linear model does seem reasonable. Backward elimination with C_p suggested the “min C_p submodel” with the variables $x_1, \log(x_2), NOX, x_6, \log(x_7), RAD, x_{10}, x_{11}$ and $\log(x_{13})$. The full model had $R^2 = 0.878$ and $\hat{\sigma} = 0.7642$. The C_p submodel had $C_p(I) = 6.576, R_I^2 = 0.878$, and $\hat{\sigma}_I = 0.762$. Deleting $\log(x_7)$ resulted in a model with $C_p = 8.483$ and the smallest coefficient p-value was 0.0095. The FF and RR plots for this model (not shown) looked like the identity line. Examining further submodels showed that NOX and RAD were the most important predictors. In particular, the OLS coefficients of x_1, x_6 and x_{11} were orders of magnitude smaller than those of NOX and RAD. The submodel including a constant, NOX, RAD and $\log(x_2)$ had $R^2 = 0.860, \hat{\sigma} = 0.811$ and $C_p = 67.368$. Figure 9.12cd shows the response plot and residual plot for this submodel.

Although this submodel has nearly the same R^2 as the full model, the residuals show more variability than those of the full model. Nevertheless, we can examine the effect of NOX and RAD on the response by deleting $\log(x_2)$. This submodel had $R^2 = 0.842, \hat{\sigma} = 0.861$ and $C_p = 138.727$. Figure 9.13a shows that the response plot for this model is no longer linear. The residual plot (not shown) also displays curvature. Figure 9.13a shows that there are two groups, one with high Y and one with low Y . There are three clusters of points in the plot of NOX versus RAD shown in Figure 9.13b (the single isolated point in the southeast corner of the plot actually corresponds to several cases). The two clusters of high NOX and high RAD points correspond to the cases with high per capita crime rate.

The tiny filled in triangles in Figure 9.13a represent the fitted values for a quadratic. We added NOX^2, RAD^2 and $NOX * RAD$ to the full model and again tried variable selection. Although the full quadratic in NOX and RAD had a linear response plot, the submodel with NOX, RAD and $\log(x_2)$ was very similar. For this data set, NOX and RAD seem to be the most important predictors, but other predictors are needed to make the model linear and to reduce residual variation.

In the Boston housing data, now let $Y = CRIM$. Since $\log(Y)$ has a linear relationship with the predictors, Y should follow a nonlinear 1D regression model. Consider the full model with predictors $\log(x_2), x_3, x_4, x_5, \log(x_7), x_8, \log(x_9)$ and $\log(x_{12})$. Regardless of whether Y or $\log(Y)$ is used as the

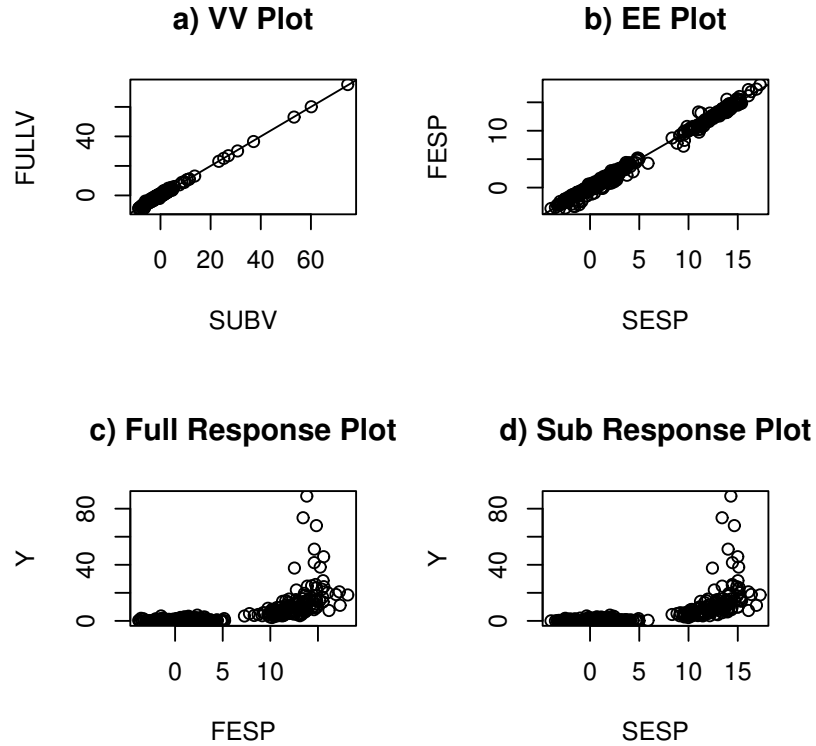


Fig. 9.14 Boston Housing Data: Nonlinear 1D Regression Model

response, the minimum C_p model from backward elimination used a constant, $\log(x_2)$, x_4 , $\log(x_7)$, x_8 and $\log(x_{12})$ as predictors. If Y is the response, then the model is nonlinear and $C_p = 5.699$. Remark 9.5 suggests that if $C_p \leq 2k$, then the points in the VV plot should tightly cluster about the identity line even if a multiple linear regression model fails to hold. Figure 9.14 shows the VV and EE plots for the minimum C_p submodel. The response plots for the full model and submodel are also shown. Note that the clustering in the VV plot is indeed higher than the clustering in the EE plot. Note that the response plots are highly nonlinear but are nearly identical.

Example 9.7. This insulation data was contributed by Ms. Spector. A box with insulation was heated for 20 minutes then allowed to cool down. The response variable $Y = \text{temperature}$ in middle of box was taken at time 0, 5, ..., 40. The *type* of insulation was a factor with type 1 = no insulation, 2 = corn pith, 3 = fiberglass, 4 = styrofoam and 5 = bubbles. There were 45 temperature measurements, one for each time type combination. The mea-

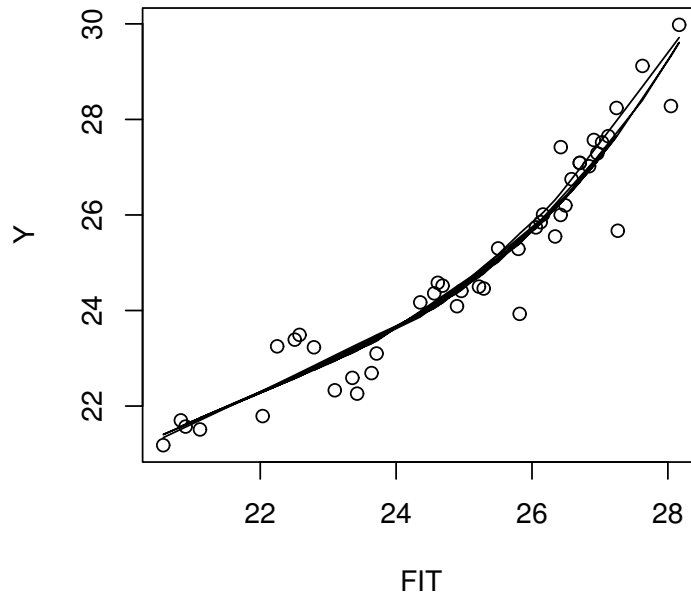


Fig. 9.15 Response Plot for Insulation Data

surements were averages of ten trials and starting temperatures were close but not exactly equal.

The model using time, $(\text{time})^2$, type, and the interactions type:time and type:(time)² had $E(Y|\mathbf{x}) \approx (\mathbf{x}^T \boldsymbol{\beta})^2$. A second model used time, $(\text{time})^2$ and type, and rather awkward *R* code for producing the response plot in Figure 9.15 is shown below. The solid curve corresponds to $(\mathbf{x}^T \hat{\boldsymbol{\beta}}, (\mathbf{x}^T \hat{\boldsymbol{\beta}})^2) = (\text{FIT}, (\text{FIT})^3)$ where $\hat{\boldsymbol{\beta}}$ is the OLS estimator from regressing Y on $\mathbf{x}^T = (1, \text{time}, (\text{time})^2, \text{type})$. The thin curve corresponds to lowess. Since the two curves correspond, $E(Y|\mathbf{x}) \approx (\mathbf{x}^T \boldsymbol{\beta})^3$ or $Y = m(\mathbf{x}^T \boldsymbol{\beta}) + e$ where $m(w) = w^3$. See Problem 9.10 for producing the response plot in *Arc*.

```
#assume the insulation data is loaded
ftype <- as.factor(insulation[,2])
zi <- as.data.frame(insulation)
iout <- lm(ytemp~time+I(time^2)+ftype,data=zi)
FIT <- iout$fit
Y <- insulation[,1]
plot(FIT,Y)
```

```

lines(lowess(FIT,Y)) #get (FIT, (FIT)^3) curve
zx <- FIT
z <- lsfit(cbind(zx, zx^2, zx^3), Y)
zfit <- Y-z$resid
lines(FIT, zfit)

```

9.5 Inference

This section follows Chang and Olive (2007, 2010) closely. Inference can be performed for trimmed views if M is chosen without using the response, e.g. if the trimming is done with a DD plot, and the dimension reduction (DR) method such as OLS or sliced inverse regression (SIR) is performed on the data $(Y_{Mi}, \mathbf{x}_{Mi})$ that remains after trimming $M\%$ of the cases with ellipsoidal trimming based on the FCH or RFCH estimator.

First we review some theoretical results for the DR methods OLS and SIR and give the main theoretical result for OLS. Let

$$\text{Cov}(\mathbf{x}) = E[(\mathbf{x} - E(\mathbf{x}))(\mathbf{x} - E(\mathbf{x}))^T] = \boldsymbol{\Sigma}_{\mathbf{x}}$$

and $\text{Cov}(\mathbf{x}, Y) = E[(\mathbf{x} - E(\mathbf{x}))(Y - E(Y))] = \boldsymbol{\Sigma}_{\mathbf{x}Y}$. Let the OLS estimator be $(\hat{\alpha}_{OLS}, \hat{\boldsymbol{\beta}}_{OLS})$. Then the population coefficients from an OLS regression of Y on \mathbf{x} are

$$\alpha_{OLS} = E(Y) - \boldsymbol{\beta}_{OLS}^T E(\mathbf{x}) \quad \text{and} \quad \boldsymbol{\beta}_{OLS} = \boldsymbol{\Sigma}_{\mathbf{x}}^{-1} \boldsymbol{\Sigma}_{\mathbf{x}Y}. \quad (9.20)$$

Let the data be (Y_i, \mathbf{x}_i) for $i = 1, \dots, n$. Let the $p \times 1$ vector $\boldsymbol{\eta} = (\alpha, \boldsymbol{\beta}^T)^T$, let \mathbf{X} be the $n \times p$ OLS design matrix with i th row $(1, \mathbf{x}_i^T)$, and let $\mathbf{Y} = (Y_1, \dots, Y_n)^T$. Then the OLS estimator $\hat{\boldsymbol{\eta}} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{Y}$. The sample covariance of \mathbf{x} is

$$\hat{\boldsymbol{\Sigma}}_{\mathbf{x}} = \frac{1}{n-1} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T \quad \text{where the sample mean } \bar{\mathbf{x}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i.$$

Similarly, define the sample covariance of \mathbf{x} and Y to be

$$\hat{\boldsymbol{\Sigma}}_{\mathbf{x}Y} = \frac{1}{n} \sum_{i=1}^n (\mathbf{x}_i - \bar{\mathbf{x}})(Y_i - \bar{Y}) = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i Y_i - \bar{\mathbf{x}} \bar{Y}.$$

The first result shows that $\hat{\boldsymbol{\eta}}$ is a consistent estimator of $\boldsymbol{\eta}$.

i) Suppose that $(Y_i, \mathbf{x}_i^T)^T$ are iid random vectors such that $\boldsymbol{\Sigma}_{\mathbf{x}}^{-1}$ and $\boldsymbol{\Sigma}_{\mathbf{x}Y}$ exist. Then

$$\hat{\alpha}_{OLS} = \bar{Y} - \hat{\boldsymbol{\beta}}_{OLS}^T \bar{\mathbf{x}} \xrightarrow{D} \alpha_{OLS}$$

and

$$\hat{\beta}_{OLS} = \frac{n}{n-1} \hat{\Sigma}_{\mathbf{x}}^{-1} \hat{\Sigma}_{\mathbf{x}Y} \xrightarrow{D} \beta_{OLS} \text{ as } n \rightarrow \infty.$$

Some notation is needed for the following results. Many 1D regression models have an error e with

$$\sigma^2 = \text{Var}(e) = E(e^2). \quad (9.21)$$

Let \hat{e} be the error residual for e . Let the population OLS residual

$$v = Y - \alpha_{OLS} - \beta_{OLS}^T \mathbf{x} \quad (9.22)$$

with

$$\tau^2 = E[(Y - \alpha_{OLS} - \beta_{OLS}^T \mathbf{x})^2] = E(v^2), \quad (9.23)$$

and let the OLS residual be

$$r = Y - \hat{\alpha}_{OLS} - \hat{\beta}_{OLS}^T \mathbf{x}. \quad (9.24)$$

Typically the OLS residual r is not estimating the error e and $\tau^2 \neq \sigma^2$, but the following results show that the OLS residual is of great interest for 1D regression models.

Assume that a 1D model holds, $Y \perp\!\!\!\perp \mathbf{x} | (\alpha + \beta^T \mathbf{x})$, which is equivalent to $Y \perp\!\!\!\perp \mathbf{x} | \beta^T \mathbf{x}$. Then under regularity conditions, results ii) – iv) below hold.

ii) Li and Duan (1989): $\beta_{OLS} = c\beta$ for some constant c .

iii) Li and Duan (1989) and Chen and Li (1998):

$$\sqrt{n}(\hat{\beta}_{OLS} - c\beta) \xrightarrow{D} N_{p-1}(\mathbf{0}, \mathbf{C}_{OLS}) \quad (9.25)$$

where

$$\mathbf{C}_{OLS} = \Sigma_{\mathbf{x}}^{-1} E[(Y - \alpha_{OLS} - \beta_{OLS}^T \mathbf{x})^2 (\mathbf{x} - E(\mathbf{x}))(\mathbf{x} - E(\mathbf{x}))^T] \Sigma_{\mathbf{x}}^{-1}. \quad (9.26)$$

iv) Chen and Li (1998): Let \mathbf{A} be a known full rank constant $k \times (p-1)$ matrix. If the null hypothesis $H_0: \mathbf{A}\beta = \mathbf{0}$ is true, then

$$\sqrt{n}(\mathbf{A}\hat{\beta}_{OLS} - c\mathbf{A}\beta) = \sqrt{n}\mathbf{A}\hat{\beta}_{OLS} \xrightarrow{D} N_k(\mathbf{0}, \mathbf{A}\mathbf{C}_{OLS}\mathbf{A}^T)$$

and

$$\mathbf{A}\mathbf{C}_{OLS}\mathbf{A}^T = \tau^2 \mathbf{A}\Sigma_{\mathbf{x}}^{-1}\mathbf{A}^T. \quad (9.27)$$

Notice that $\mathbf{C}_{OLS} = \tau^2 \Sigma_{\mathbf{x}}^{-1}$ if $v = Y - \alpha_{OLS} - \beta_{OLS}^T \mathbf{x} \perp\!\!\!\perp \mathbf{x}$ or if the MLR model holds. If the MLR model holds, $\tau^2 = \sigma^2$.

To create test statistics, the estimator

$$\hat{\tau}^2 = \text{MSE} = \frac{1}{n-p} \sum_{i=1}^n r_i^2 = \frac{1}{n-p} \sum_{i=1}^n (Y_i - \hat{\alpha}_{OLS} - \hat{\beta}_{OLS}^T \mathbf{x}_i)^2$$

will be useful. The estimator $\hat{\mathbf{C}}_{OLS} =$

$$\hat{\Sigma}_{\mathbf{x}}^{-1} \left[\frac{1}{n} \sum_{i=1}^n [(Y_i - \hat{\alpha}_{OLS} - \hat{\beta}_{OLS}^T \mathbf{x}_i)^2 (\mathbf{x}_i - \bar{\mathbf{x}})(\mathbf{x}_i - \bar{\mathbf{x}})^T] \right] \hat{\Sigma}_{\mathbf{x}}^{-1} \quad (9.28)$$

can also be useful. Notice that for general 1D regression models, the OLS MSE estimates τ^2 rather than the error variance σ^2 .

v) Result iv) suggests that a test statistic for $H_0 : \mathbf{A}\beta = \mathbf{0}$ is

$$W_{OLS} = n \hat{\beta}_{OLS}^T \mathbf{A}^T [\mathbf{A} \hat{\Sigma}_{\mathbf{x}}^{-1} \mathbf{A}^T]^{-1} \mathbf{A} \hat{\beta}_{OLS} / \hat{\tau}^2 \xrightarrow{D} \chi_k^2, \quad (9.29)$$

the chi-square distribution with k degrees of freedom.

Before presenting the main theoretical result, some results from OLS MLR theory are needed. Let the $p \times 1$ vector $\boldsymbol{\eta} = (\alpha, \beta^T)^T$, the known $k \times p$ constant matrix $\tilde{\mathbf{A}} = [\mathbf{a} \ \mathbf{A}]$ where \mathbf{a} is a $k \times 1$ vector, and let \mathbf{c} be a known $k \times 1$ constant vector. Following Seber and Lee (2003, p. 99–106), the usual F statistic for testing $H_0 : \tilde{\mathbf{A}}\boldsymbol{\eta} = \mathbf{c}$ is

$$F_0 = \frac{(SSE(H_0) - SSE)/k}{SSE/(n-p)} = \quad (9.30)$$

$$(\tilde{\mathbf{A}}\hat{\boldsymbol{\eta}} - \mathbf{c})^T [\tilde{\mathbf{A}}(\mathbf{X}^T \mathbf{X})^{-1} \tilde{\mathbf{A}}^T]^{-1} (\tilde{\mathbf{A}}\hat{\boldsymbol{\eta}} - \mathbf{c}) / (k \hat{\tau}^2)$$

where $MSE = \hat{\tau}^2 = SSE/(n-p)$, $SSE = \sum_{i=1}^n r_i^2$ and

$$SSE(H_0) = \sum_{i=1}^n r_i^2(H_0)$$

is the minimum sum of squared residuals subject to the constraint $\tilde{\mathbf{A}}\boldsymbol{\eta} = \mathbf{c}$. If H_0 is true, the MLR model holds and the errors e_i are iid $N(0, \sigma^2)$, then $F_0 \sim F_{k, n-p}$, the F distribution with k and $n-p$ degrees of freedom. Also, if $Z_n \sim F_{k, n-p}$, then

$$Z_n \xrightarrow{D} \chi_k^2/k \quad (9.31)$$

as $n \rightarrow \infty$.

The main theoretical result of this section is Theorem 9.4 below. This theorem and (9.31) suggest that OLS output, originally meant for testing with the MLR model, can also be used for testing with many 1D regression data sets. Without loss of generality, let the 1D model $Y \mid \mathbf{x} | (\alpha + \beta^T \mathbf{x})$ be written as

$$Y \mid \mathbf{x} | (\alpha + \beta_R^T \mathbf{x}_R + \beta_O^T \mathbf{x}_O)$$

where the reduced model is $Y \mid \mathbf{x} | (\alpha + \beta_R^T \mathbf{x}_R)$ and \mathbf{x}_O denotes the terms outside of the reduced model. Notice that OLS ANOVA F test corresponds to $H_0: \beta = \mathbf{0}$ and uses $\mathbf{A} = \mathbf{I}_{p-1}$. The tests for $H_0: \beta_i = 0$ use $\mathbf{A} =$

$(0, \dots, 0, 1, 0, \dots, 0)$ where the 1 is in the i th position and are equivalent to the OLS t tests. The test $H_0: \beta_O = \mathbf{0}$ uses $\mathbf{A} = [\mathbf{0} \ \mathbf{I}_j]$ if β_O is a $j \times 1$ vector, and the test statistic (9.30) can be computed with the OLS partial F test: run OLS on the full model to obtain SSE and on the reduced model to obtain $SSE(R) \equiv SSE(H_0)$.

In the theorem below, it is crucial that $H_0: \mathbf{A}\beta = \mathbf{0}$. Tests for $H_0: \mathbf{A}\beta = \mathbf{1}$, say, may not be valid even if the sample size n is large. Also, confidence intervals corresponding to the t tests are for $c\beta_i$, and are usually not very useful when c is unknown.

Theorem 9.4. Assume that a 1D regression model (9.1) holds with $h(\mathbf{x}) = \alpha + \beta^T \mathbf{x}$ and that Equation (9.29) holds when $H_0: \mathbf{A}\beta = \mathbf{0}$ is true. Then the test statistic (9.30) satisfies

$$F_0 = \frac{n-1}{kn} W_{OLS} \xrightarrow{D} \chi_k^2/k$$

as $n \rightarrow \infty$.

Proof. Notice that by (9.29), the result follows if $F_0 = (n-1)W_{OLS}/(kn)$. Let $\tilde{\mathbf{A}} = [\mathbf{0} \ \mathbf{A}]$ so that $H_0: \tilde{\mathbf{A}}\boldsymbol{\eta} = \mathbf{0}$ is equivalent to $H_0: \mathbf{A}\beta = \mathbf{0}$. Following Seber and Lee (2003, p. 106),

$$(\mathbf{X}^T \mathbf{X})^{-1} = \begin{pmatrix} \frac{1}{n} + \bar{\mathbf{x}}^T \mathbf{D}^{-1} \bar{\mathbf{x}} & -\bar{\mathbf{x}}^T \mathbf{D}^{-1} \\ -\mathbf{D}^{-1} \bar{\mathbf{x}} & \mathbf{D}^{-1} \end{pmatrix} \quad (9.32)$$

where the $(p-1) \times (p-1)$ matrix

$$\mathbf{D}^{-1} = [(n-1)\hat{\boldsymbol{\Sigma}}_{\mathbf{x}}]^{-1} = \hat{\boldsymbol{\Sigma}}_{\mathbf{x}}^{-1}/(n-1). \quad (9.33)$$

Using $\tilde{\mathbf{A}}$ and (9.32) in (9.30) shows that $F_0 =$

$$(\mathbf{A}\hat{\boldsymbol{\beta}}_{OLS})^T \begin{bmatrix} \mathbf{0} & \mathbf{A} \end{bmatrix} \begin{pmatrix} \frac{1}{n} + \bar{\mathbf{x}}^T \mathbf{D}^{-1} \bar{\mathbf{x}} & -\bar{\mathbf{x}}^T \mathbf{D}^{-1} \\ -\mathbf{D}^{-1} \bar{\mathbf{x}} & \mathbf{D}^{-1} \end{pmatrix} \begin{pmatrix} \mathbf{0}^T \\ \mathbf{A}^T \end{pmatrix}^{-1} \mathbf{A}\hat{\boldsymbol{\beta}}_{OLS}/(k\hat{\tau}^2),$$

and the result follows from (9.33) after algebra. \square

For SIR, the theory is more complicated. Following Chen and Li (1998), SIR produces eigenvalues $\hat{\lambda}_i$ and associated SIR directions $\hat{\boldsymbol{\beta}}_{i,SIR}$ for $i = 1, \dots, p-1$. The SIR directions $\hat{\boldsymbol{\beta}}_{i,SIR}$ for $i = 1, \dots, d$ are used for dD regression.

vi) Chen and Li (1998): For a 1D regression and vector \mathbf{A} , a test statistic for $H_0: \mathbf{A}\beta_1 = \mathbf{0}$ is

$$W_S = n\hat{\boldsymbol{\beta}}_{1,SIR}^T \mathbf{A}^T [\mathbf{A}\hat{\boldsymbol{\Sigma}}_{\mathbf{x}}^{-1} \mathbf{A}^T]^{-1} \mathbf{A}\hat{\boldsymbol{\beta}}_{1,SIR}/[(1-\hat{\lambda}_1)/\hat{\lambda}_1] \xrightarrow{D} \chi_1^2. \quad (9.34)$$

Ellipsoidal trimming can be used to create outlier resistant dimension reduction (DR) methods that can give useful results when the assumption of linearly related predictors (9.6) is violated. To perform ellipsoidal trimming,

a robust estimator of multivariate location and dispersion (T, \mathbf{C}) is computed and used to create the Mahalanobis distances $D_i(T, \mathbf{C})$. The i th case (Y_i, \mathbf{x}_i) is trimmed if $D_i > D_{(j)}$. For example, if $j \approx 0.9n$, then about $M\% = 10\%$ of the cases are trimmed, and a DR method can be computed from the cases that remain.

For theory and outlier resistance, the choice of (T, \mathbf{C}) and M are important. Chang and Olive (2007) used the MBA estimator $(T_{MBA}, \mathbf{C}_{MBA})$ for (T, \mathbf{C}) , but we would now use the RFCH estimator because of its combination of speed, robustness and theory. The classical Mahalanobis distance uses $(T, \mathbf{C}) = (\bar{\mathbf{x}}, \hat{\Sigma}\mathbf{x})$. Denote the robust distances by RD_i and the classical distances by MD_i . Then the DD plot of the MD_i versus the RD_i can be used to choose M . Chapter 11 showed that the plotted points in the DD plot will follow the identity line with zero intercept and unit slope if the predictor distribution is multivariate normal (MVN), and will follow a line with zero intercept but non-unit slope for a large class of non-MVN elliptically contoured distributions that have a nonsingular covariance matrix. Delete $M\%$ of the cases with the largest RFCH distances so that the remaining cases follow the identity line (or some line through the origin) closely. Let $(Y_{Mi}, \mathbf{x}_{Mi})$ denote the data that was not trimmed where $i = 1, \dots, n_M$. Then apply the DR method on these n_M cases.

As long as M is chosen only using the predictors, DR theory will apply if the data (Y_M, \mathbf{x}_M) satisfies the regularity conditions. For example, if the MLR model is valid and the errors are iid $N(0, \sigma^2)$, then the OLS estimator

$$\hat{\boldsymbol{\eta}}_M = (\mathbf{X}_M^T \mathbf{X}_M)^{-1} \mathbf{X}_M^T \mathbf{Y}_M \sim N_p(\boldsymbol{\eta}, \sigma^2 (\mathbf{X}_M^T \mathbf{X}_M)^{-1}).$$

More generally, let $\hat{\boldsymbol{\beta}}_{DM}$ denote a DR estimator applied to $(Y_{Mi}, \mathbf{x}_{Mi})$ and assume that

$$\sqrt{n_M}(\hat{\boldsymbol{\beta}}_{DM} - c_M \boldsymbol{\beta}) \xrightarrow{D} N_{p-1}(\mathbf{0}, \mathbf{C}_{DM})$$

where \mathbf{C}_{DM} is nonsingular. Let $\phi_M = \lim_{n \rightarrow \infty} n/n_M$. Then

$$\sqrt{n}(\hat{\boldsymbol{\beta}}_{DM} - c_M \boldsymbol{\beta}) = \frac{\sqrt{n}}{\sqrt{n_M}} \sqrt{n_M}(\hat{\boldsymbol{\beta}}_{DM} - c_M \boldsymbol{\beta}) \xrightarrow{D} N_{p-1}(\mathbf{0}, \phi_M \mathbf{C}_{DM}). \quad (9.35)$$

If $H_0: \mathbf{A}\boldsymbol{\beta} = \mathbf{0}$ is true and $\hat{\mathbf{C}}_{DM}$ is a consistent estimator of \mathbf{C}_{DM} , then

$$W_{DM} = n_M \hat{\boldsymbol{\beta}}_{DM}^T \mathbf{A}^T [\mathbf{A} \hat{\mathbf{C}}_{DM} \mathbf{A}^T]^{-1} \mathbf{A} \hat{\boldsymbol{\beta}}_{DM} / \hat{\tau}_M^2 \xrightarrow{D} \chi_k^2.$$

Notice that $M = 0$ corresponds to the full data set and $n_0 = n$.

The tradeoff is that if low amounts of trimming do not work, then larger amounts of trimming sometimes greatly improve DR methods, but large amounts of trimming will have large efficiency losses if low amounts of trim-

ming work since $n/n_M \geq 1$ and the diagonal elements of \mathbf{C}_{DM} typically become larger with M .

Trimmed views can also be used to select $M \equiv M_{TV}$. If the MLR model holds and OLS is used, then the resulting trimmed views estimator $\hat{\beta}_{M,TV}$ is \sqrt{n} consistent, but need not be asymptotically normal.

Adaptive trimming can be used to obtain an asymptotically normal estimator that may avoid large efficiency losses. First, choose an initial amount of trimming M_I by using, e.g., the DD plot or trimmed views. Let $\hat{\beta}$ denote the first direction of the DR method. Next compute $|corr(\hat{\beta}_M^T \mathbf{x}, \hat{\beta}_{M_I}^T \mathbf{x})|$ for $M = 0, 10, \dots, 90$ and find the smallest value $M_A \leq M_I$ such that the absolute correlation is greater than 0.95. If no such value exists, then use $M_A = M_I$. The resulting adaptive trimming estimator is asymptotically equivalent to the estimator that uses 0% trimming if $\hat{\beta}_0$ is a consistent estimator of $c_0\beta$ and if $\hat{\beta}_{M_I}$ is a consistent estimator of $c_{M_I}\beta$.

Detecting outlying \mathbf{x} is useful for any regression method, and now that effective methods such as RFCH are available, the DD plot should be used routinely. In a small simulation, the clean data $Y = (\alpha + \beta^T \mathbf{x})^3 + e$ where $\alpha = 1$, $\beta = (1, 0, 0, 0)^T$, $e \sim N(0, 1)$ and $\mathbf{x} \sim N_4(\mathbf{0}, \mathbf{I}_4)$. The outlier percentage γ was either 0% or 49%. The 2 clusters of outliers were about the same size and had $Y \sim N(0, 1)$, $\mathbf{x} \sim N_4(\pm 10(1, 1, 1, 1)^T, \mathbf{I}_4)$. Table 9.2 records the averages of $\hat{\beta}_i$ over 100 runs where the DR method used $M = 0$ or $M = 50\%$ trimming. SIR, SAVE and PHD were very similar except when $\gamma = 49$ and $M = 0$. When outliers were present, the average of $\hat{\beta}_{F,50} \approx c_F(1, 0, 0, 0)^T$ where c_F depended on the DR method and F was OLS, SIR, SAVE or PHD. The sample size $n = 1000$ was used although OLS gave reasonable estimates for much smaller sample sizes. The `rpac` function `drsim7` can be used to duplicate the simulation in *R*.

The following simulations show that ellipsoidal trimming based on the MBA estimator is useful for DR even when no outliers are present.

In the simulations, we used eight types of predictor distributions: d1) $\mathbf{x} \sim N_{p-1}(\mathbf{0}, \mathbf{I}_{p-1})$, d2) $\mathbf{x} \sim 0.6N_{p-1}(\mathbf{0}, \mathbf{I}_{p-1}) + 0.4N_{p-1}(\mathbf{0}, 25\mathbf{I}_{p-1})$, d3) $\mathbf{x} \sim 0.4N_{p-1}(\mathbf{0}, \mathbf{I}_{p-1}) + 0.6N_{p-1}(\mathbf{0}, 25\mathbf{I}_{p-1})$, d4) $\mathbf{x} \sim 0.9N_{p-1}(\mathbf{0}, \mathbf{I}_{p-1}) + 0.1N_{p-1}(\mathbf{0}, 25\mathbf{I}_{p-1})$, d5) $\mathbf{x} \sim LN(\mathbf{0}, \mathbf{I})$ where the marginals are iid log-normal(0,1), d6) $\mathbf{x} \sim MVT_{p-1}(3)$, d7) $\mathbf{x} \sim MVT_{p-1}(5)$ and d8) $\mathbf{x} \sim MVT_{p-1}(19)$. Here \mathbf{x} has a multivariate t distribution $\mathbf{x}_i \sim MVT_{p-1}(\nu)$ if $\mathbf{x}_i = \mathbf{z}_i / \sqrt{W_i/\nu}$ where $\mathbf{z}_i \sim N_{p-1}(\mathbf{0}, \mathbf{I}_{p-1})$ is independent of the chi-square random variable $W_i \sim \chi_\nu^2$. Of the eight distributions, only d5) is not elliptically contoured. The MVT distribution gets closer to the multivariate normal (MVN) distribution d1) as $\nu \rightarrow \infty$. The MVT distribution has first moments for $\nu \geq 2$ and second moments for $\nu \geq 3$. See Johnson and Kotz (1972, p. 134-135) and Press (2005, p. 136). All simulations used 1000 runs.

Table 9.2 DR Coefficient Estimation with Trimming

type	γ	M	$\hat{\beta}_1$	$\hat{\beta}_2$	$\hat{\beta}_3$	$\hat{\beta}_4$
SIR	0	0	.0400	.0021	-.0006	.0012
SIR	0	50	-.0201	-.0015	.0014	.0027
SIR	49	0	.0004	-.0029	-.0013	.0039
SIR	49	50	-.0798	-.0014	.0004	-.0015
SAVE	0	0	.0400	.0012	.0010	.0018
SAVE	0	50	-.0201	-.0018	.0024	.0030
SAVE	49	0	-.4292	-.2861	-.3264	-.3442
SAVE	49	50	-.0797	-.0016	-.0006	-.0024
PHD	0	0	.0396	-.0009	-.0071	-.0063
PHD	0	50	-.0200	-.0013	.0024	.0025
PHD	49	0	-.1068	-.1733	-.1856	-.1403
PHD	49	50	-.0795	.0023	.0000	-.0037
OLS	0	0	5.974	.0083	-.0221	.0008
OLS	0	50	4.098	.0166	.0017	-.0016
OLS	49	0	2.269	-.7509	-.7390	-.7625
OLS	49	50	5.647	.0305	.0011	.0053

The simulations were for single index models with $\alpha = 1$. Let the sufficient predictor $SP = \alpha + \beta^T \mathbf{x}$. Then the seven models considered were m1) $Y = SP + e$, m2) $Y = (SP)^2 + e$, m3) $Y = \exp(SP) + e$, m4) $Y = (SP)^3 + e$, m5) $Y = \sin(SP)/SP + 0.01e$, m6) $Y = SP + \sin(SP) + 0.1e$ and m7) $Y = \sqrt{|SP|} + 0.1e$ where $e \sim N(0, 1)$.

First, coefficient estimation was examined with $\beta = (1, 1, 1, 1)^T$, and for OLS the sample standard deviation (SD) of each entry $\hat{\beta}_{Mi,j}$ of $\hat{\beta}_{M,j}$ was computed for $i = 1, 2, 3, 4$ with $j = 1, \dots, 1000$. For each of the 1000 runs, the Chen and Li formula

$$SE_{cl}(\hat{\beta}_{Mi}) = \sqrt{n_M^{-1}(\hat{C}_M)_{ii}}$$

was computed where

$$\hat{C}_M = \hat{\Sigma}_{\mathbf{x}_M}^{-1} \left[\frac{1}{n_M} \sum_{i=1}^{n_M} [(Y_{Mi} - \hat{\alpha}_M - \hat{\beta}_M^T \mathbf{x}_{Mi})^2 (\mathbf{x}_{Mi} - \bar{\mathbf{x}}_M)(\mathbf{x}_{Mi} - \bar{\mathbf{x}}_M)^T] \right] \hat{\Sigma}_{\mathbf{x}_M}^{-1}$$

is the estimate (9.28) applied to (Y_M, \mathbf{x}_M) . The average of $\hat{\beta}_M$ and of $\sqrt{n}SE_{cl}$ were recorded as well as $\sqrt{n}SD$ of $\hat{\beta}_{Mi,j}$ under the labels $\bar{\beta}_M$, $\sqrt{n} \bar{SE}_{cl}$ and $\sqrt{n}SD$. Under regularity,

$$\sqrt{n} \bar{SE}_{cl} \approx \sqrt{n}SD \approx \sqrt{\frac{1}{1 - \frac{M}{100}} \text{diag}(\mathbf{C}_M)}$$

where \mathbf{C}_M is (9.26) applied to (Y_M, \mathbf{x}_M) .

Table 9.3 OLS Coefficient Estimation with Trimming

m	\mathbf{x}	M	$\bar{\beta}_M$	$\sqrt{n} \overline{SE}_{cl}$	$\sqrt{n} SD$
m2	d1	0	2.00, 2.01, 2.00, 2.00	7.81, 7.79, 7.76, 7.80	7.87, 8.00, 8.02, 7.88
m5	d4	0	-.03, -.03, -.03, -.03	.30, .30, .30, .30	.31, .32, .33, .31
m6	d5	0	1.04, 1.04, 1.04, 1.04	.36, .36, .37, .37	.41, .42, .42, .40
m7	d6	10	.11, .11, .11, .11	.58, .57, .57, .57	.60, .58, .62, .61

For MVN \mathbf{x} , MLR and 0% trimming, all three recorded quantities were near (1, 1, 1, 1) for $n = 60, 500$, and 1000. For 90% trimming and $n = 1000$, the results were $\bar{\beta}_{90} = (1.00, 1.00, 1.01, 0.99)$, $\sqrt{n} \overline{SE}_{cl} = (7.56, 7.61, 7.60, 7.54)$ and $\sqrt{n} SD = (7.81, 8.02, 7.76, 7.59)$, suggesting that $\hat{\beta}_{90}$ is asymptotically normal but inefficient.

For other distributions, results for 0 and 10% trimming were recorded as well as a “good” trimming value M_B . Results are “good” if all of the entries of both $\bar{\beta}_{M_B}$ and $\sqrt{n} \overline{SE}_{cl}$ were approximately equal, and if the theoretical $\sqrt{n} \overline{SE}_{cl}$ was close to the simulated $\sqrt{n} SD$. The results were good for MVN \mathbf{x} and all seven models, and the results were similar for $n = 500$ and $n = 1000$. The results were good for models m1 and m5 for all eight distributions. Model m6 was good for 0% trimming except for distribution d5 and model m7 was good for 0% trimming except for distributions d5, d6 and d7. Trimming usually helped for models m2, m3 and m4 for distributions d5 – d8. Some results are shown in Table 9.3 for $n = 500$.

For SIR with $h = 4$ slices $\bar{\beta}_M$ was recorded. The SIR results were similar to those for OLS, but often more trimming and larger sample sizes were needed than those for OLS. The results depended on h in that the largest sample sizes were needed for 2 slices and then for 3 slices.

Next testing was considered. Let F_M and W_M denote the OLS and SIR statistics (9.30) and (9.34) applied to the n_M cases (Y_M, \mathbf{x}_M) that remained after trimming. Ho was rejected for OLS if $F_M > F_{k, n_M - p}(0.95)$ and for SIR if $W_M > \chi_k^2(0.95)$. For SIR, 2 slices were used since using more than $h = 2$ slices rejected Ho too often. As h increased from 2 to 3 to 4, $\hat{\lambda}_1$ and the SIR chi-square test statistic W_0 rapidly increased. For $h > 4$ the increase was much slower.

For testing the nominal level was 0.05, and we recorded the proportion \hat{p} of runs where Ho was rejected. Since 1000 runs were used, the count $1000\hat{p} \sim \text{binomial}(1000, 1 - \delta_n)$ where $1 - \delta_n$ converges to the true large sample level $1 - \delta$. The standard error for the proportion is $\sqrt{\hat{p}(1 - \hat{p})}/1000 \approx 0.0069$ for $p = 0.05$. An observed coverage $\hat{p} \in (0.03, 0.07)$ suggests that there is no reason to doubt that the true level is 0.05.

Suppose a 1D model holds but $Y \not\perp \mathbf{x}$. Then the Y_i are iid and the model reduces to $Y = E(Y) + e = c_\alpha + e$ where $e = Y - E(Y)$. As a special case, if

Table 9.4 Rejection Proportions for $H_0: \beta = 0$

\mathbf{x}	n	F	SIR	n	F	SIR
d1	100	0.041	0.057	500	0.050	0.048
d2	100	0.050	0.908	500	0.045	0.930
d3	100	0.047	0.955	500	0.050	0.930
d4	100	0.045	0.526	500	0.048	0.599
d5	100	0.055	0.621	500	0.061	0.709
d6	100	0.042	0.439	500	0.036	0.472
d7	100	0.054	0.214	500	0.047	0.197
d8	100	0.044	0.074	500	0.060	0.077

Table 9.5 Rejection Proportions for $H_0: \beta_2 = 0$

m	\mathbf{x}	Test	70	60	50	40	30	20	10	0	ADAP
1	1	F	.061	.056	.062	.051	.046	.050	.044	.043	.043
1	1	W	.007	.013	.015	.020	.027	.032	.045	.056	.056
5	1	F	.019	.023	.019	.019	.020	.022	.027	.037	.029
5	1	W	.002	.003	.006	.005	.010	.014	.025	.055	.026
2	2	F	.023	.024	.026	.070	.183	.182	.142	.166	.040
2	2	W	.007	.010	.021	.067	.177	.328	.452	.576	.050
4	3	F	.027	.058	.096	.081	.071	.057	.062	.123	.120
4	3	W	.028	.069	.152	.263	.337	.378	.465	.541	.539
6	4	F	.026	.024	.030	.032	.028	.044	.051	.088	.088
6	4	W	.012	.009	.013	.016	.030	.040	.076	.386	.319
7	5	F	.058	.058	.053	.054	.046	.044	.051	.037	.037
7	5	W	.001	.000	.005	.005	.034	.080	.118	.319	.250
3	6	F	.021	.024	.019	.025	.025	.034	.080	.374	.036
3	6	W	.003	.008	.007	.021	.019	.041	.084	.329	.264
6	7	F	.027	.032	.023	.041	.047	.053	.052	.055	.055
6	7	W	.007	.006	.013	.022	.019	.025	.054	.176	.169

$Y = m(\alpha + \beta^T \mathbf{x}) + e$ and if $Y \perp \mathbf{x}$, then $Y = m(\alpha) + e$. For the corresponding test $H_0: \beta = \mathbf{0}$ versus $H_1: \beta \neq \mathbf{0}$, and the OLS F statistic (9.30) and SIR W statistic (9.34) are invariant with respect to a constant. This test is interesting since if H_0 holds, then the results do not depend on the 1D model (9.1), but only on the distribution of \mathbf{x} and the distribution of e . Since $\beta_{OLS} = c\beta$, power can be good if $c \neq 0$. The OLS test is equivalent to the ANOVA F test from MLR of Y on \mathbf{x} . Under H_0 , the test should perform well provided that the design matrix is nonsingular and the error distribution and sample size are such that the central limit theorem holds. Table 9.4 shows the results for OLS and SIR for $n = 100, 500$ and for the eight different distributions. Since the true model was linear and normal, the exact OLS level is 0.05 even for $n = 10$. Table 9.4 shows that OLS performed as expected while SIR only gave good results for MVN \mathbf{x} .

Next the test $H_0: \beta_2 = 0$ was considered. The OLS test is equivalent to the t test from MLR of Y on \mathbf{x} . The true model used $\alpha = 1$ and $\beta = (1, 0, 1, 1)^T$. To simulate adaptive trimming, $|\text{corr}(\hat{\beta}_M^T \mathbf{x}, \beta^T \mathbf{x})|$ was computed

for $M = 0, 10, \dots, 90$ and the initial trimming proportion M_I maximized this correlation. This process should be similar to choosing the best trimmed view by examining 10 plots. The rejection proportions were recorded for $M = 0, \dots, 90$ and for adaptive trimming. The seven models, eight distributions and sample sizes $n = 60, 150$, and 500 were used. Table 9.5 shows some results for $n = 150$.

For OLS, the test that used adaptive trimming had proportions ≤ 0.072 except for model m4 with distributions d2, d3, d4, d6, d7 and d8; m2 with d4, d6 and d7 for $n = 500$ and d6 with $n = 150$; m6 with d4 and $n = 60, 150$; m5 with d7 and $n = 500$ and m7 with d7 and $n = 500$. With the exception of m4, when the adaptive $\hat{p} > 0.072$, then 0% trimming had a rejection proportion near 0.1. Occasionally adaptive trimming was conservative with $\hat{p} < 0.03$. The 0% trimming worked well for m1 and m6 for all eight distributions and for d1 and d5 for all seven models. Models m2 and m3 usually benefited from adaptive trimming. For distribution d1, the adaptive and 0% trimming methods had identical \hat{p} for $n = 500$ except for m3 where the values were 0.038 and 0.042. Chang (2006) has much more extensive tables.

For SIR results were not as good. Adaptive trimming worked about as often as it failed, and failed for model m1. Also, 0% trimming performed well for all seven models for the MVN distribution d1, and there was always an M such the W_M did not reject H_0 too often.

9.6 Complements

Introductions to 1D regression and regression graphics are Cook and Weisberg (1999a, ch. 18, 19, and 20) and Cook and Weisberg (1999b), while Olive (2010) considers 1D regression. More advanced treatments are Cook (1998a) and Li (2000). Important papers include Brillinger (1977, 1983) and Li and Duan (1989). Formal testing procedures for the single index model are given by Simonoff and Tsai (2002) and Gao and Liang (1997). Li (1997) shows that OLS F tests can be asymptotically valid for model (9.2) if \mathbf{x} is multivariate normal and $\Sigma_{\mathbf{x}}^{-1} \Sigma_{\mathbf{x}Y} \neq \mathbf{0}$.

Let $\boldsymbol{\eta} = (\alpha, \boldsymbol{\beta}^T)^T$. Then the i th *Cook's distance*

$$CD_i = \frac{(\hat{\mathbf{Y}}_{(i)} - \hat{\mathbf{Y}})^T (\hat{\mathbf{Y}}_{(i)} - \hat{\mathbf{Y}})}{p \hat{\sigma}^2} = \frac{\|ESP(i) - ESP\|^2}{(p+1)MSE} \quad (9.36)$$

where $ESP(i) = \mathbf{X}^T \hat{\boldsymbol{\eta}}_{(i)}$ and the estimated sufficient predictor $ESP = \mathbf{X}^T \hat{\boldsymbol{\eta}}$ estimates $d\mathbf{x}_j^T \boldsymbol{\eta}$ for some constant d and $j = 1, \dots, n$. This fact suggests that Cook's distances and MD_i^2 still give useful information on cases that influence the estimated sufficient summary plot although MSE is estimating $E(r^2) = E[(Y - \alpha_{OLS} - \mathbf{x}^T \boldsymbol{\beta}_{OLS})^2] = \tau^2$.

There are many ways to estimate 1D models, including maximum likelihood for parametric models. The literature for estimating $c\beta$ when model (9.2) holds is growing, and Cook and Li (2002) summarize when competing methods such as ordinary least squares (OLS), sliced inverse regression (SIR), principal Hessian directions (PHD), and sliced average variance estimation (SAVE) can fail. All four methods frequently perform well if there are no strong nonlinearities present in the predictors. Cook and Ni (2005) provides theory for inverse regression methods such as SAVE. Further information about these and related methods can be found, for example, in Brillinger (1977, 1983), Chen and Li (1998), Cook (1998ab, 2004), Cook and Critchley (2000), Cook and Weisberg (1991, 1999ab), Li (1991, 1992, 2000) and Li and Zhu (2007).

Several papers have suggested that outliers and strong nonlinearities need to be removed from the predictors. See Brillinger (1991), Cook (1998a, p. 152), Cook and Nachtsheim (1994), Heng-Hui (2001), Li and Duan (1989, p. 1011, 1041, 1042) and Li (1991, p. 319). Outlier resistant methods for general methods such as SIR are less common, but see Gather, Hilker and Becker (2001, 2002) (where FMCD should be replaced by RFCH) and Cížek and Härdle (2006). Trimmed views were introduced by Olive (2002, 2004b).

Section 9.4 follows Olive and Hawkins (2005) closely. The literature on numerical methods for variable selection in the OLS multiple linear regression model is enormous, and the literature for other given 1D regression models is also growing. See Naik and Tsai (2001) and Kong and Xia (2007).

Section 9.5 followed Chang and Olive (2007, 2010) closely. More examples and much more simulations are in Chang (2006). Severini (1998) discusses when OLS output is relevant for the Gaussian additive error single index model. Also see Yoo, Patterson and Datta (2009).

The mussel data set is included as the file *mussel.lsp* in the *Arc* software and can be obtained from the web site (<http://www.stat.umn.edu/arc/>). The Boston housing data can be obtained from the STATLIB website (<http://lib.stat.cmu.edu/datasets/boston>). Both data sets can be obtained from the text website.

9.7 Problems

9.1. Refer to Definition 9.3 for the Cox and Snell (1968) definition for residuals, but replace η by β .

- a) Find \hat{e}_i if $Y_i = \mu + e_i$ and $T(Y)$ is used to estimate μ .
- b) Find \hat{e}_i if $Y_i = \mathbf{x}_i^T \beta + e_i$.
- c) Find \hat{e}_i if $Y_i = \beta_1 \exp[\beta_2(x_i - \bar{x})]e_i$ where the e_i are iid exponential(1) random variables and \bar{x} is the sample mean of the x_i 's.

d) Find \hat{e}_i if $Y_i = \mathbf{x}_i^T \boldsymbol{\beta} + e_i / \sqrt{w_i}$.

9.2*. (Aldrin, Bølviken, and Schweder 1993). Suppose

$$Y = m(\boldsymbol{\beta}^T \mathbf{x}) + e \quad (9.37)$$

where m is a possibly unknown function and the zero mean errors e are independent of the predictors. Let $z = \boldsymbol{\beta}^T \mathbf{x}$ and let $\mathbf{w} = \mathbf{x} - E(\mathbf{x})$. Let $\boldsymbol{\Sigma}_{\mathbf{x}, Y} = \text{Cov}(\mathbf{x}, Y)$, and let $\boldsymbol{\Sigma}_{\mathbf{x}} = \text{Cov}(\mathbf{x}) = \text{Cov}(\mathbf{w})$. Let $\mathbf{r} = \mathbf{w} - (\boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta}) \boldsymbol{\beta}^T \mathbf{w}$.

a) Recall that $\text{Cov}(\mathbf{x}, Y) = E[(\mathbf{x} - E(\mathbf{x}))(\mathbf{Y} - E(Y))^T]$ and show that $\boldsymbol{\Sigma}_{\mathbf{x}, Y} = E(\mathbf{w}Y)$.

b) Show that $E(\mathbf{w}Y) = \boldsymbol{\Sigma}_{\mathbf{x}, Y} = E[(\mathbf{r} + (\boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta}) \boldsymbol{\beta}^T \mathbf{w}) m(z)] =$

$$E[m(z)\mathbf{r}] + E[\boldsymbol{\beta}^T \mathbf{w} m(z)] \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta}.$$

c) Using $\boldsymbol{\beta}_{OLS} = \boldsymbol{\Sigma}_{\mathbf{x}}^{-1} \boldsymbol{\Sigma}_{\mathbf{x}, Y}$, show that $\boldsymbol{\beta}_{OLS} = c(\mathbf{x}) \boldsymbol{\beta} + \mathbf{u}(\mathbf{x})$ where the constant

$$c(\mathbf{x}) = E[\boldsymbol{\beta}^T (\mathbf{x} - E(\mathbf{x})) m(\boldsymbol{\beta}^T \mathbf{x})]$$

and the bias vector $\mathbf{u}(\mathbf{x}) = \boldsymbol{\Sigma}_{\mathbf{x}}^{-1} E[m(\boldsymbol{\beta}^T \mathbf{x}) \mathbf{r}]$.

d) Show that $E(\mathbf{w}z) = \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta}$. (Hint: Use $E(\mathbf{w}z) = E[(\mathbf{x} - E(\mathbf{x})) \boldsymbol{\beta}^T \boldsymbol{\beta}] = E[(\mathbf{x} - E(\mathbf{x}))(\mathbf{x}^T - E(\mathbf{x}^T) + E(\mathbf{x}^T)) \boldsymbol{\beta}]$.)

e) Assume $m(z) = z$. Using d), show that $c(\mathbf{x}) = 1$ if $\boldsymbol{\beta}^T \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta} = 1$.

f) Assume that $\boldsymbol{\beta}^T \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta} = 1$. Show that $E(z\mathbf{r}) = E(\mathbf{r}z) = \mathbf{0}$. (Hint: Find $E(\mathbf{r}z)$ and use d).)

g) Suppose that $\boldsymbol{\beta}^T \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta} = 1$ and that the distribution of \mathbf{x} is multivariate normal. Then the joint distribution of z and \mathbf{r} is multivariate normal. Using the fact that $E(z\mathbf{r}) = \mathbf{0}$, show $\text{Cov}(\mathbf{r}, z) = \mathbf{0}$ so that z and \mathbf{r} are independent. Then show that $\mathbf{u}(\mathbf{x}) = \mathbf{0}$.

(Note: the assumption $\boldsymbol{\beta}^T \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta} = 1$ can be made without loss of generality since if $\boldsymbol{\beta}^T \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\beta} = d^2 > 0$ (assuming $\boldsymbol{\Sigma}_{\mathbf{x}}$ is positive definite), then $y = m(d(\boldsymbol{\beta}/d)^T \mathbf{x}) + e \equiv m_d(\boldsymbol{\eta}^T \mathbf{x}) + e$ where $m_d(u) = m(du)$, $\boldsymbol{\eta} = \boldsymbol{\beta}/d$ and $\boldsymbol{\eta}^T \boldsymbol{\Sigma}_{\mathbf{x}} \boldsymbol{\eta} = 1$.)

9.3. Suppose that you have a statistical model where both fitted values and residuals can be obtained. For example this is true for time series and for nonparametric regression models such as $Y = f(x_1, \dots, x_p) + e$ where $\hat{Y} = \hat{f}(x_1, \dots, x_p)$ and the residual $\hat{e} = Y - \hat{f}(x_1, \dots, x_p)$. Suggest graphs for variable selection for such models.

Output for Problem 9.4.

BEST SUBSET REGRESSION MODELS FOR CRIM

(A) LogX2 (B) X3 (C) X4 (D) X5 (E) LogX7 (F) X8 (G) LogX9 (H) LogX12
 3 "BEST" MODELS FROM EACH SUBSET SIZE LISTED.

k	CP	ADJUSTED R SQUARE	R SQUARE	RESID SS	MODEL VARIABLES
1	379.8	0.0000	0.0000	37363.2	INTERCEPT ONLY
2	36.0	0.3900	0.3913	22744.6	F
2	113.2	0.3025	0.3039	26007.8	G
2	191.3	0.2140	0.2155	29310.8	E
3	21.3	0.4078	0.4101	22039.9	E F
3	25.0	0.4036	0.4059	22196.7	F H
3	30.8	0.3970	0.3994	22442.0	D F
4	17.5	0.4132	0.4167	21794.9	C E F
4	18.1	0.4125	0.4160	21821.0	E F H
4	18.8	0.4117	0.4152	21850.4	A E F
5	10.2	0.4226	0.4272	21402.3	A E F H
5	10.8	0.4219	0.4265	21427.7	C E F H
5	12.0	0.4206	0.4252	21476.6	A D E F
6	5.7	0.4289	0.4346	21125.8	A C E F H
6	9.3	0.4248	0.4305	21279.1	A C D E F
6	10.3	0.4237	0.4294	21319.9	A B E F H
7	6.3	0.4294	0.4362	21065.0	A B C E F H
7	6.3	0.4294	0.4362	21066.3	A C D E F H
7	7.7	0.4278	0.4346	21124.3	A C E F G H
8	7.0	0.4297	0.4376	21011.8	A B C D E F H
8	8.3	0.4283	0.4362	21064.9	A B C E F G H
8	8.3	0.4283	0.4362	21065.8	A C D E F G H
9	9.0	0.4286	0.4376	21011.8	A B C D E F G H

9.4. The output above is for the Boston housing data from software that does all subsets variable selection. The full model is a 1D transformation model with response variable $Y = \text{CRIM}$ and uses a constant and variables A, B, C, D, E, F, G and H. (Using $\log(\text{CRIM})$ as the response would give an MLR model.) From this output, what is the best submodel? Explain briefly.

9.5*. a) Show that $C_p(I) \leq 2k$ if and only if $F_I \leq p/(p - k)$.

b) Using (9.19), find $E(C_p)$ and $\text{Var}(C_p)$ assuming that an MLR model is appropriate and that H_0 (the reduced model I can be used) is true.

c) Using (9.19), $C_p(I_{full}) = p$ and the notation in Section 9.4, show that

$$C_p(I_i) = C_p(I_{full}) + (t_i^2 - 2).$$

R Problems Some *R* code for homework problems is at (<http://parker.ad.siu.edu/Olive/robRhw.txt>).

Warning: Use a command like `source("G:/rpack.txt")` to download the programs. See Preface or Section 11.2. Typing the name of the `rpack` function, e.g. `trviews`, will display the code for the function. Use the `args` command, e.g. `args(trviews)`, to display the needed arguments for the function.

9.6. Use the following *R* commands to make 100 $N_3(\mathbf{0}, I_3)$ cases and 100 trivariate non-EC cases.

```
n3x <- matrix(rnorm(300), nrow=100, ncol=3)
ln3x <- exp(n3x)
```

In *R*, type the command `library(MASS)`.

a) Using the commands `pairs(n3x)` and `pairs(ln3x)` and include both scatterplot matrices in *Word*. (Click on the plot and hit *Ctrl* and *c* at the same time. Then go to *file* in the *Word* menu and select *paste*.) Are strong nonlinearities present among the MVN predictors? How about the non-EC predictors? (Hint: a box or ball shaped plot is linear.)

b) Make a single index model and the sufficient summary plot with the following commands

```
ncy <- (n3x%%1:3)^3 + 0.1*rnorm(100)
plot(n3x%%(1:3), ncy)
```

and include the plot in *Word*.

c) The command `trviews(n3x, ncy)` will produce ten plots. To advance the plots, click on the *rightmost mouse button* (and in *R* select *stop*) to advance to the next plot. The last plot is the OLS view. Include this plot in *Word*.

d) After all 10 plots have been looked at the output will show 10 estimated predictors. The last estimate is the OLS (least squares) view and might look like

```
Intercept      X1      X2      X3
4.417988 22.468779 61.242178 75.284664
```

If the OLS view is a good estimated sufficient summary plot, then the plot created from the command (leave out the intercept)

```
plot(n3x%%c(22.469, 61.242, 75.285), n3x%%1:3)
```

should cluster tightly about some line. Your linear combination will be different than the one used above. Using your OLS view, include the plot using the command above (but with your linear combination) in *Word*. Was this plot linear? Did some of the other trimmed views seem to be better than the OLS view, that is, did one of the trimmed views seem to have a smooth mean function with a smaller variance function than the OLS view?

e) Now type the *R* command

```
lncy <- (ln3x%%1:3)^3 + 0.1*rnorm(100).
```

Use the command *trviews(ln3x,lncy)* to find the best view with a smooth mean function and the smallest variance function. This view should not be the OLS view. Include your best view in *Word*.

f) Get the linear combination from your view, say $(94.848, 216.719, 328.444)^T$, and obtain a plot with the command

```
plot(ln3x%%c(94.848, 216.719, 328.444), ln3x%%1:3).
```

Include the plot in *Word*. If the plot is linear with high correlation, then your response plot in e) should be good.

9.7. (At the beginning of your *R* session, use *source("G:/rpack.txt")* command and the *library(MASS)* command.)

a) Perform the commands

```
> nx <- matrix(rnorm(300), nrow=100, ncol=3)
> lnx <- exp(nx)
> SP <- lnx%%1:3
> lnsincy <- sin(SP)/SP + 0.01*rnorm(100)
```

For parts b), c) and d) below, to make the best trimmed view with *trviews*, *ctrviews* or *lmsviews*, you may need to use the function twice. The first view trims 90% of the data, the next view trims 80%, etc. The last view trims 0% and is the OLS view (or *lmsreg* view). Remember to advance the view with the rightmost mouse button (and in *R*, highlight “stop”). Then click on the plot and next simultaneously hit *Ctrl* and *c*. This makes a copy of the plot. Then in *Word*, use the menu commands “Copy>paste.”

b) Find the best trimmed view with OLS and *covfch* with the following commands and include the view in *Word*.

```
> trviews(lnx, lnsincy)
```

(With *trviews*, suppose that 40% trimming gave the best view. Then instead of using the procedure above b), you can use the command

```
> essp(lnx, lnsincy, M=40)
```

to make the best trimmed view. Then click on the plot and next simultaneously hit *Ctrl* and *c*. This makes a copy of the plot. Then in *Word*, use the menu commands “Copy>paste”. Click the rightmost mouse button (and in *R*, highlight “stop”) to return the command prompt.)

c) Find the best trimmed view with OLS and (\bar{x}, S) using the following commands and include the view in *Word*. See the paragraph above b).

```
> ctrviews(lnx, lnsincy)
```

d) Find the best trimmed view with `lmsreg` and `cov.mcd` using the following commands and include the view in *Word*. See the paragraph above b).

```
> lmsviews(lnx, lnsincy)
```

e) Which method or methods gave the best response plot? Explain briefly.

9.8. Warning: this problem may take too much time. This problem is like Problem 9.7 but uses many more single index models.

a) Make some prototype functions with the following commands.

```
> nx <- matrix(rnorm(300), nrow=100, ncol=3)
> SP <- nx%%1:3
> ncuby <- SP^3 + rnorm(100)
> nexpy <- exp(SP) + rnorm(100)
> nlinsy <- SP + 4*sin(SP) + 0.1*rnorm(100)
> nsincy <- sin(SP)/SP + 0.01*rnorm(100)
> nsiny <- sin(SP) + 0.1*rnorm(100)
> nsqrty <- sqrt(abs(SP)) + 0.1*rnorm(100)
> nsqy <- SP^2 + rnorm(100)
```

b) Make sufficient summary plots similar to Figures 9.1 and 9.2 with the following commands and include both plots in *Word*.

```
> plot(SP, ncuby)
> plot(-SP, ncuby)
```

c) Find the best trimmed view with the following commands (first type `library(MASS)` if you are using *R*). Include the view in *Word*.

```
> trviews(nx, ncuby)
```

You may need to use the function twice. The first view trims 90% of the data, the next view trims 80%, etc. The last view trims 0% and is the OLS view. Remember to advance the view with the rightmost mouse button (and in *R*, highlight “stop”). Suppose that 40% trimming gave the best view. Then use the command

```
> essp(nx, ncuby, M=40)
```

to make the best trimmed view. Then click on the plot and next simultaneously hit *Ctrl* and *c*. This makes a copy of the plot. Then in *Word*, use the menu commands “Copy>paste”.

d) To make a plot like Figure 9.5, use the following commands. Let $tem = \hat{\beta}$ obtained from the *trviews* output. In Example 9.2 (continued), tem can be obtained with the following command.

```
> tem <- c(12.60514, 25.06613, 37.25504)
```

Include the plot in *Word*.

```
> ESP <- nx%%tem
> plot(ESP, SP)
```

e) Repeat b), c) and d) with the following data sets.

- i) nx and nexpy
- ii) nx and nlinsy
- iii) nx and nsincy
- iv) nx and nsiny
- v) nx and nsqrty
- vi) nx and nsqy

Enter the following commands to do parts vii) to x).

```
> lnx <- exp(nx)
> SP <- lnx%%1:3
> lncuby <- (SP/3)^3 + rnorm(100)
> lnlsinsy <- SP + 10*sin(SP) + 0.1*rnorm(100)
> lnsincy <- sin(SP)/SP + 0.01*rnorm(100)
> lnsiny <- sin(SP/3) + 0.1*rnorm(100)
> ESP <- lnx%%tem
```

- vii) lnx and lncuby
- viii) lnx and lnlsinsy
- ix) lnx and lnsincy
- x) lnx and lnsiny

9.9. Warning: this problem may take too much time. Repeat Problem 9.8 but replace `trviews` with a) `lmsviews`, b) `symviews` (that creates views that sometimes work even when symmetry is present), c) `ctrviews` and d) `sirviews`.

Except for part a), the *essp* command will not work. Instead, for the best trimmed view, click on the plot and next simultaneously hit *Ctrl* and *c*. This makes a copy of the plot. Then in *Word*, use the menu commands “Copy>paste”.