PHILOSOPHICAL TRANSACTIONS OF — THE ROYAL SOCIETY A

MATHEMATICAL, PHYSICAL & ENGINEERING SCIENCES

Sufficient dimension reduction and prediction in regression

Kofi P. Adragni and R. Dennis Cook

References	**This article cites 26 articles, 3 of which can be accessed free** http://rsta.royalsocietypublishing.org/content/367/1906/4385.full.html#ref-list-1
Rapid response	Respond to this article http://rsta.royalsocietypublishing.org/letters/submit/roypta;367/1906/4385
Subject collections	Articles on similar topics can be found in the following collections statistics (30 articles)
Email alerting service	Receive free email alerts when new articles cite this article - sign up in the box at the top right-hand corner of the article or click **here**

To subscribe to *Phil. Trans. R. Soc. A* go to:
http://rsta.royalsocietypublishing.org/subscriptions

Sufficient dimension reduction and prediction in regression

By Kofi P. Adragni and R. Dennis Cook*

University of Minnesota, School of Statistics, 313 Ford Hall, 224 Church Street Southeast, Minneapolis, MN 55455, USA

Dimension reduction for regression is a prominent issue today because technological advances now allow scientists to routinely formulate regressions in which the number of predictors is considerably larger than in the past. While several methods have been proposed to deal with such regressions, principal components (PCs) still seem to be the most widely used across the applied sciences. We give a broad overview of ideas underlying a particular class of methods for dimension reduction that includes PCs, along with an introduction to the corresponding methodology. New methods are proposed for prediction in regressions with many predictors.

Keywords: lasso; partial least squares; principal components; principal component regression; principal fitted components

1. Introduction

Consider the frequently encountered goal of determining a rule $m(\mathbf{x})$ for predicting a future observation of a univariate response variable Y at the given value $\mathbf{x}$ of a $p \times 1$ vector $\mathbf{X}$ of continuous predictors. Assuming that Y is quantitative, continuous or discrete, the mean-squared error $E(Y - m(\mathbf{x}))^2$ is minimized by choosing $m(\mathbf{x})$ to be the mean $E(Y \mid \mathbf{X} = \mathbf{x})$ of the conditional distribution of $Y \mid (\mathbf{X} = \mathbf{x})$. Consequently, the prediction goal is often specialized immediately to the task of estimating the conditional mean function $E(Y \mid \mathbf{X})$ from the regression of Y on $\mathbf{X}$. When the response is categorical with sample space S_Y consisting of h categories $S_Y = \{C_1, \ldots, C_h\}$, the mean function is no longer a relevant quantity for prediction. Instead, given an observation $\mathbf{x}$ on $\mathbf{X}$, the predicted category C_* is usually taken to be the one with the largest conditional probability $C_* = \arg\max \Pr(C_k \mid \mathbf{X} = \mathbf{x})$, where the maximization is over S_Y. When pursuing the estimation of $E(Y \mid \mathbf{X})$ or $\Pr(C_k \mid \mathbf{X})$, it is nearly always worth while to consider predictions based on a function $R(\mathbf{X})$ of dimension less than p, provided that it captures all of the information that $\mathbf{X}$ contains about Y so that $E(Y \mid \mathbf{X}) = E(Y \mid R(\mathbf{X}))$. We can think of $R(\mathbf{X})$ as a function that concentrates the relevant information in $\mathbf{X}$. The action of replacing $\mathbf{X}$ with a lower-dimensional function $R(\mathbf{X})$ is called *dimension reduction*; it is called *sufficient dimension reduction* when $R(\mathbf{X})$ retains all the relevant information

about Y. A potential advantage of sufficient dimension reduction is that predictions based on an estimated R may be substantially less variable than those based on $\mathbf{X}$, without introducing worrisome bias. This advantage is not confined to predictions, but may accrue in other phases of a regression analysis as well.

One goal of this paper is to give a broad overview of ideas underlying sufficient dimension reduction for regression, along with an introduction to the corresponding methodology. Sections 1*a,b*, 2 and 3 are devoted largely to this review. Sufficient dimension-reduction methods are designed to estimate a population parameter called the central subspace, which is defined in §1*b*. Another goal of this paper is to describe a new method of predicting quantitative responses following sufficient dimension reduction; categorical responses will be discussed only for contrast. The focus of this paper shifts to prediction in §4, where we discuss four inverse regression models, describe the prediction methodology that stems from them and give simulation results to illustrate their behaviour. Practical implementation issues are discussed in §5, along with additional simulation results.

(*a*) *Dimension reduction*

There are many methods available for estimating $E(Y \mid \mathbf{X})$ based on a random sample $(Y_i, \mathbf{X}_i)$, $i = 1, \ldots, n$, from the joint distribution of Y and $\mathbf{X}$. If p is sufficiently small and n is sufficiently large, it may be possible to estimate $E(Y \mid \mathbf{X})$ adequately by using non-parametric smoothing (see Wand & Jones 1995). Otherwise, nearly all techniques for estimating $E(Y \mid \mathbf{X})$ employ some types of dimension reduction for $\mathbf{X}$, either estimated or imposed as an intrinsic part of the model or method.

Broadly viewed, dimension reduction has always been a central statistical concept. In the second half of the nineteenth century, 'reduction of observations' was widely recognized as a core goal of statistical methodology, and principal components (PCs) were emerging as a general method for the reduction of multivariate observations (Adcock 1878). PCs was established as a first reductive method for regression by the mid-1900s.

Dimension reduction for regression is a prominent issue today because technological advances now allow scientists to routinely formulate regressions in which p is considerably larger than in the past. This has complicated the development and fitting of regression models. Experience has shown that the standard iterative paradigm for model development guided by diagnostics (Cook & Weisberg 1982, p. 7) can be imponderable when applied with too many predictors. An added complication arises when p is larger than the number of observations n, leading to the so-called '$n < p$' problem. Standard methods of fitting and corresponding inference procedures may no longer be applicable in such regressions. These and related issues have caused a shift in the applied sciences towards a different regression genre with the goal of reducing the dimensionality of the predictor vector as a first step in the analysis. Although large-p regressions are perhaps mainly responsible for renewed interest, dimension-reduction methodology can be useful regardless of the size of p. For instance, it is often helpful to have an informative low-dimensional graphical summary of the regression to facilitate model building and gain insights. For

this goal, p may be regarded as large when it exceeds 2 or 3 because these bounds represent the limits of our ability to view a dataset in full using computer graphics. Subsequent references to 'large p' in this paper do not necessarily imply that $n < p$.

Reduction by PCs is ubiquitous in the applied sciences, particularly in bioinformatics applications, where PCs have been called 'eigen-genes' (Alter *et al.* 2000) in microarray data analyses and 'meta-$k$mers' in analyses involving DNA motifs. The report by Wegman *et al.* (2006) reiterates and makes clear that past influential analyses of data on global warming are flawed because of an inappropriate use of the PC methodology.

While PCs seem to be the dominant method of dimension reduction across the applied sciences, there are many other established and recent statistical methods that might be used to address large-p regressions, including factor analysis, inverse regression estimation (IRE) (Cook & Ni 2005), partial least squares (PLS), projection pursuit, seeded reductions (Cook *et al.* 2007), kernel methods (Fukumizu *et al.* 2009) and sparse methods like the lasso (Tibshirani 1996) that are based on penalization.

(b) Sufficient dimension reduction

Dimension reduction is a rather amorphous concept in statistics, changing its character and goals depending on context. Formulated specifically for regression, the following definition (Cook 2007) of a *sufficient reduction* will help in our pursuit of methods for reducing the dimension of **X** while en route to estimating $E(Y \mid \mathbf{X})$.

Definition 1.1. A reduction $R : \mathbb{R}^p \to \mathbb{R}^q$, $q \leq p$, is sufficient if it satisfies one of the following three statements:

   (i) inverse reduction, $\mathbf{X} \mid (Y, R(\mathbf{X})) \sim \mathbf{X} \mid R(\mathbf{X})$,
  (ii) forward reduction, $Y \mid \mathbf{X} \sim Y \mid R(\mathbf{X})$,
 (iii) joint reduction, $\mathbf{X} \perp\!\!\!\perp Y \mid R(\mathbf{X})$,

where $\perp\!\!\!\perp$ indicates independence, $\sim$ means identically distributed and $\mathbf{A} \mid \mathbf{B}$ refers to the random vector $\mathbf{A}$ given the vector $\mathbf{B}$.

Each of the three conditions in this definition conveys the idea that the reduction $R(\mathbf{X})$ carries all the information that $\mathbf{X}$ has about Y, and consequently all the information available to estimate $E(Y \mid \mathbf{X})$. They are equivalent when $(Y, \mathbf{X})$ has a joint distribution. In that case, we are free to determine a reduction inversely or jointly and then pass it to the conditional mean without additional structure: $E(Y \mid \mathbf{X}) = E(Y \mid R(\mathbf{X}))$. In some cases, there may be a direct connection between $R(\mathbf{X})$ and $E(Y \mid \mathbf{X})$. For instance, if $(Y, \mathbf{X})$ follows a non-singular multivariate normal distribution, then $R(\mathbf{X}) = E(Y \mid \mathbf{X})$ is a sufficient reduction, $E(Y \mid \mathbf{X}) = E\{Y \mid E(Y \mid \mathbf{X})\}$. This reduction is also *minimal sufficient*: if $T(\mathbf{X})$ is any sufficient reduction, then R is a function of T. Further, because of the nature of the multivariate normal distribution, it can be expressed as a linear combination of the elements of **X**: $R = \boldsymbol{\beta}^{\mathrm{T}} \mathbf{X}$ is minimal sufficient for some vector $\boldsymbol{\beta}$.

Inverse reduction by itself does not require the response Y to be random, and it is perhaps the only reasonable reductive route when Y is fixed by design. For instance, in discriminant analysis, $\mathbf{X} \mid Y$ is a random vector of features observed in one of a number of subpopulations indicated by the categorical response Y, and no discriminatory information will be lost if classifiers are restricted to R.

If we consider a generic statistical problem and reinterpret $\mathbf{X}$ as the total data D and Y as the parameter θ, then the condition for inverse reduction becomes $D \mid (\theta, R) \sim D \mid R$ so that R is a sufficient statistic. In this way, the definition of a sufficient reduction encompasses Fisher's (1922) classical definition of sufficiency. One difference is that sufficient statistics are observable, while a sufficient reduction may contain unknown parameters and thus needs to be estimated. For example, if $(\mathbf{X}, Y)$ follows a non-singular multivariate normal distribution, then $R(\mathbf{X}) = \boldsymbol{\beta}^{\mathrm{T}}\mathbf{X}$, and it is necessary to estimate $\boldsymbol{\beta}$.

In some regressions, $R(\mathbf{X})$ may be a nonlinear function of $\mathbf{X}$, and in extreme cases no reduction may be possible, so all sufficient reductions are one-to-one functions of $\mathbf{X}$ and thus equivalent to $R(\mathbf{X}) = \mathbf{X}$. Most often, we encounter multi-dimensional reductions consisting of several linear combinations $R(\mathbf{X}) = \boldsymbol{\eta}^{\mathrm{T}}\mathbf{X}$, where $\boldsymbol{\eta}$ is an unknown $p \times q$ matrix, $q \leq p$, that must be estimated from the data. Linear reductions may be imposed to facilitate progress, as in the moment-based approach reviewed in §3a. They can also arise as a natural consequence of modelling restrictions, as we will see in §3b. If $\boldsymbol{\eta}^{\mathrm{T}}\mathbf{X}$ is a sufficient linear reduction, then so is $(\boldsymbol{\eta}\mathbf{A})^{\mathrm{T}}\mathbf{X}$ for any $q \times q$ full-rank matrix $\mathbf{A}$. Consequently, only the subspace span$(\boldsymbol{\eta})$ spanned by the columns of $\boldsymbol{\eta}$ can be identified—span$(\boldsymbol{\eta})$ is called a *dimension-reduction subspace*. If span$(\boldsymbol{\eta})$ is a dimension-reduction subspace, then so is span$(\boldsymbol{\eta}, \boldsymbol{\eta}_1)$ for any $p \times q_1$ matrix $\boldsymbol{\eta}_1$. If span$(\boldsymbol{\eta}_1)$ and span$(\boldsymbol{\eta}_2)$ are both dimension-reduction subspaces, then under mild conditions so is their intersection span$(\boldsymbol{\eta}_1) \cap \text{span}(\boldsymbol{\eta}_2)$ (Cook 1996, 1998). Consequently, the inferential target in sufficient dimension reduction is often taken to be the central subspace $\mathcal{S}_{Y|X}$, defined as the intersection of all dimension-reduction subspaces (Cook 1994, 1996, 1998). A *minimal sufficient linear reduction* is then of the form $R(\mathbf{X}) = \boldsymbol{\eta}^{\mathrm{T}}\mathbf{X}$, where the columns of $\boldsymbol{\eta}$ now form a basis for $\mathcal{S}_{Y|X}$. We assume that the central subspace exists throughout this paper, and use $d = \dim(\mathcal{S}_{Y|X})$ to denote its dimension.

The ideas of a sufficient reduction and the central subspace can be used to further our understanding of the existing methodology and to guide the development of a new methodology. In §§2 and 3, we consider how sufficient reductions arise in three contexts: forward linear regression, inverse moment-based reduction and inverse model-based reduction.

2. Reduction in forward linear regression

The standard linear regression model $Y = \beta_0 + \boldsymbol{\beta}^{\mathrm{T}}\mathbf{X} + \epsilon$, with $\epsilon \perp\!\!\!\perp \mathbf{X}$ and $E(\epsilon) = 0$, implies that $\mathcal{S}_{Y|X} = \text{span}(\boldsymbol{\beta})$ and thus that $R(\mathbf{X}) = \boldsymbol{\beta}^{\mathrm{T}}\mathbf{X}$ is minimal sufficient. The assumption of a linear regression then automatically focuses our interest on $\boldsymbol{\beta}$, which can be estimated straightforwardly using ordinary least squares (OLS) when n is sufficiently large, and it may appear that there is little to be gained from dimension reduction. However, dimension reduction has been used in linear regression to improve on the OLS estimator of $\boldsymbol{\beta}$ and to deal with $n < p$

regressions. One approach consists of regressing Y on $\mathbf{X}$ in two steps. The first is the reduction step: reduce $\mathbf{X}$ linearly to $\mathbf{G}^{\mathrm{T}}\mathbf{X}$ using some methodology that produces $\mathbf{G} \in \mathbb{R}^{p \times q}$, $q \leq p$. The second step consists of using OLS to estimate the mean function $E(Y \mid \mathbf{G}^{\mathrm{T}}\mathbf{X})$ for the reduced predictors. To describe the resulting estimator $\hat{\boldsymbol{\beta}}_{\mathbf{G}}$ of $\boldsymbol{\beta}$ and establish notation for later sections, let $\mathbb{Y}$ be the $n \times 1$ vector of centred responses, let $\bar{\mathbf{X}} = \sum_{i=1}^{n} \mathbf{X}/n$ denote the sample mean vector, let $\mathbb{X}$ be the $n \times p$ matrix with rows $(\mathbf{X}_i - \bar{\mathbf{X}})^{\mathrm{T}}$, $i = 1, \ldots, n$, let $\hat{\boldsymbol{\Sigma}} = \mathbb{X}^{\mathrm{T}}\mathbb{X}/n$ denote the usual estimator of $\boldsymbol{\Sigma} = \mathrm{var}(\mathbf{X})$, let $\hat{\mathbf{C}} = \mathbb{X}^{\mathrm{T}}\mathbb{Y}/n$, which is the usual estimator of $\mathbf{C} = \mathrm{cov}(\mathbf{X}, Y)$, and let $\hat{\boldsymbol{\beta}}_{\mathrm{ols}} = \hat{\boldsymbol{\Sigma}}^{-1}\hat{\mathbf{C}}$ be the vector of coefficients from the OLS fit of Y on $\mathbf{X}$. Then (Cook & Forzani 2009*b*)

$$\hat{\boldsymbol{\beta}}_{\mathbf{G}} = \mathbf{P}_{\mathbf{G}(\hat{\boldsymbol{\Sigma}})}\hat{\boldsymbol{\beta}}_{\mathrm{ols}} = \mathbf{G}(\mathbf{G}^{\mathrm{T}}\hat{\boldsymbol{\Sigma}}\mathbf{G})^{-1}\mathbf{G}^{\mathrm{T}}\hat{\mathbf{C}}. \tag{2.1}$$

This estimator, which is the projection $\mathbf{P}_{\mathbf{G}(\hat{\boldsymbol{\Sigma}})}$ of $\hat{\boldsymbol{\beta}}_{\mathrm{ols}}$ onto $\mathrm{span}(\mathbf{G})$ in the $\hat{\boldsymbol{\Sigma}}$ inner product, does not require computation of $\hat{\boldsymbol{\Sigma}}^{-1}$ if $q < p$ and thus could be useful when $n < p$, depending on the size of q. In any case, the estimator of $E(Y \mid \mathbf{X})$ is

$$\hat{E}(Y \mid \mathbf{X}) = \bar{\mathbf{Y}} + \hat{\boldsymbol{\beta}}_{\mathbf{G}}^{\mathrm{T}}(\mathbf{X} - \bar{\mathbf{X}}). \tag{2.2}$$

If $\mathbf{G} = \mathbf{I}_p$, then $\hat{\boldsymbol{\beta}}_{\mathbf{G}} = \hat{\boldsymbol{\beta}}_{\mathrm{ols}}$, which achieves nothing beyond $\hat{\boldsymbol{\beta}}_{\mathrm{ols}}$. If we choose the columns of $\mathbf{G}$ to be the first q eigenvectors of $\hat{\boldsymbol{\Sigma}}$, then $\mathbf{G}^{\mathrm{T}}\mathbf{X}$ consists of the first q PCs and $\hat{\boldsymbol{\beta}}_{\mathbf{G}}$ is the standard principal component regression (PCR) estimator. Setting $\mathbf{G} = (\hat{\mathbf{C}}, \hat{\boldsymbol{\Sigma}}\hat{\mathbf{C}}, \ldots, \hat{\boldsymbol{\Sigma}}^{q-1}\hat{\mathbf{C}})$ yields the PLS estimator with q factors (Helland 1990). Eliminating predictors by using an information criterion like Akaike AIC or Bayesian BIC (§5*a*) can result in a $\mathbf{G}$ with rows selected from the identity matrix $\mathbf{I}_p$, and again we obtain a reduction in $\mathbf{X}$ prior to estimation of $\boldsymbol{\beta}$. If $\mathrm{span}(\mathbf{G})$ is a consistent estimator of a dimension-reduction subspace $\mathcal{S}$, then $\hat{\boldsymbol{\beta}}_{\mathbf{G}}$ may be a reasonable estimator of $\boldsymbol{\beta}$ because $\mathrm{span}(\boldsymbol{\beta}) \subseteq \mathcal{S} \subseteq \mathbb{R}^p$, recalling that the intersection of any two dimension-reduction subspaces is itself a dimension-reduction subspace. However, while these estimators are well known, $\mathrm{span}(\mathbf{G})$ may not be a consistent estimator of a dimension-reduction subspace without an additional structure, even if the linear model is accurate. The PCR estimator depends on $\mathbf{G}$ only through the marginal distribution of $\mathbf{X}$, and this alone cannot guarantee that $\mathrm{span}(\mathbf{G})$ is consistent. The performance of the PLS estimator depends on the relationship between $\mathbf{C}$ and the eigenstructure of $\boldsymbol{\Sigma}$ (Naik & Tsai 2000).

A somewhat different approach is based on estimating $\boldsymbol{\beta}$ by using a penalized objective function like that for the lasso (Tibshirani 1996). The lasso estimator is

$$\hat{\boldsymbol{\beta}}_{\mathrm{lasso}} = \arg\min_{\boldsymbol{\beta}} \left\{ \sum_{i=1}^{n} (Y_i - \bar{Y} - \boldsymbol{\beta}^{\mathrm{T}}(\mathbf{X}_i - \bar{\mathbf{X}}))^2 + \lambda \sum_{j=1}^{p} |\beta_j| \right\},$$

where β_j is the jth element of $\boldsymbol{\beta}$, $j = 1, \ldots, p$, and the tuning parameter λ is often chosen by cross validation. Several elements of $\hat{\boldsymbol{\beta}}_{\mathrm{lasso}}$ are typically zero, which corresponds to setting the rows of $\mathbf{G}$ to be the rows of the identity matrix $\mathbf{I}_p$ corresponding to the non-zero elements of $\hat{\boldsymbol{\beta}}_{\mathrm{lasso}}$. However, with this $\mathbf{G}$, we do

not necessarily have $\hat{\boldsymbol{\beta}}_{\text{lasso}} = \hat{\boldsymbol{\beta}}_{\mathbf{G}}$, although the two estimators are often similar. Consequently, the methodology based on penalization does not fit exactly the general form given in equation (2.1).

Pursuing dimension reduction based on linear regression may not produce useful results if the model is not accurate, particularly if the distribution of $Y \mid \mathbf{X}$ depends on more than one linear combination of the predictors. There are many diagnostic and remedial methods available to improve linear regression models when p is not too large. Otherwise, application of these methods can be quite burdensome.

3. Inverse reduction

Inverse regression $\mathbf{X} \mid Y$ provides an alternative approach to estimating a sufficient reduction. It can deal straightforwardly with regressions that depend on more than one linear combination of the predictors, and does not necessarily suffer from the modelling problems that plague forward regression when p is large. There are two general paradigms for determining a sufficient reduction inversely. The first is by specifying a parametric model for the inverse regression of $\mathbf{X}$ on Y, as discussed in §§3*b* and 4. In this model-based approach, minimal sufficient reductions can in principle be determined from the model itself. For example, we saw previously that $E(Y \mid \mathbf{X})$ is a sufficient reduction when $(Y, \mathbf{X})$ is normally distributed. The second, which is discussed §3*a*, is the moment-based approach in which derived moment relations are used to estimate a sufficient reduction by way of the central subspace. Model accuracy is nearly always an issue in the model-based approach, while efficiency is worrisome in the moment-based approach. Some of the best moment-based methods have turned out to be quite inefficient in relatively simple settings (see Cook & Forzani 2009*a*, for an instance of this inefficiency).

(*a*) *Moment-based inverse reduction*

In contrast to model-based reduction, there is no law to guide the choice of $R(\mathbf{X})$ in moment-based reduction. However, progress is still possible by restricting consideration to multi-dimensional linear reduction and pursuing estimation of the central subspace $\mathcal{S}_{Y|X}$, as discussed in §1*b*. Recall that $d = \dim(\mathcal{S}_{Y|X})$ and that the columns of the $p \times d$ matrix $\boldsymbol{\eta}$ are a basis for $\mathcal{S}_{Y|X}$.

Sliced inverse regression (SIR; Li 1991) and sliced average variance estimation (SAVE; Cook & Weisberg 1991) were the first moment-based methods proposed for dimension reduction. Although the concept of the central subspace was not developed until a few years after SIR and SAVE were introduced, it is now known that these methods in fact estimate $\mathcal{S}_{Y|X}$ under two key conditions: (i) $E(\mathbf{X} \mid \boldsymbol{\eta}^{\mathrm{T}}\mathbf{X})$ is a linear function of $\mathbf{X}$ (*linearity condition*) and (ii) $\mathrm{var}(\mathbf{X} \mid \boldsymbol{\eta}^{\mathrm{T}}\mathbf{X})$ is a non-random matrix (*constant covariance condition*). We forgo discussion of these conditions, which involve only the marginal distribution of $\mathbf{X}$, as they are well known and widely regarded as mild. A good recent discussion of them was given by Li & Wang (2007). Under the linearity condition, $E(\mathbf{X} \mid Y) - E(\mathbf{X}) \in \boldsymbol{\Sigma}\mathcal{S}_{Y|X}$, which is the population foundation for SIR. Under the linearity and constant covariance conditions, $\mathrm{span}(\boldsymbol{\Sigma} - \mathrm{var}(\mathbf{X} \mid Y)) \in \boldsymbol{\Sigma}\mathcal{S}_{Y|X}$,

which is population basis for SAVE. When the response is categorical, $E(\mathbf{X} \mid C_k)$ can be estimated straightforwardly as the average predictor vector $\bar{\mathbf{X}}_k$ in category C_k. The SIR estimator of $\mathcal{S}_{Y|X}$, which requires $n > p$, is then the span of the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}^{-1}\mathbf{M}\mathbf{M}^{\mathrm{T}}$, where $\mathbf{M}$ is the $p \times h$ matrix with columns $\bar{\mathbf{X}}_k - \bar{\mathbf{X}}$. Continuous responses are treated by slicing the observed range of Y into h categories C_k and then applying the method for a categorical response. The SAVE estimator uses a similar construction. Routines for computing SIR, SAVE and other moment-based estimates of $\mathcal{S}_{Y|X}$ are available in the R package 'dr' (http://cran.us.r-project.org/web/packages/dr/index.html) and in the ARC software (www.stat.umn.edu/arc/software.html).

Both SIR and SAVE provide $\sqrt{n}$ consistent estimators of $\mathcal{S}_{Y|X}$ under standard conditions, but by itself consistency does not guarantee good performance in practice. It is known that SIR has difficulty finding directions that are associated with certain types of nonlinear trends in $E(Y \mid \mathbf{X})$. SAVE was developed in response to this limitation, but its ability to find linear trends is generally inferior to SIR's. Several moment-based methods have been developed in an effort to improve on the estimates of $\mathcal{S}_{Y|X}$ provided by SIR and SAVE. Using the same population foundations as SIR, Cook & Ni (2005) developed an asymptotically optimal method of estimating $\mathcal{S}_{Y|X}$ called IRE. Ye & Weiss (2003) and Zhu *et al.* (2005) attempted to combine the advantages of SIR and SAVE by using linear combinations of them. Cook & Forzani (2009*a*) used a likelihood-based objective function to develop a method called LAD (likelihood acquired directions) that apparently dominates all dimension-reduction methods based on the same population foundations as SIR and SAVE. These methods have been developed and studied mostly in regressions where $p \ll n$, although there are some results for other settings (Li 2007; Li & Yin 2008). SIR, SAVE, IRE and LAD come with a range of inference capabilities, including methods for estimating d and tests of conditional independence hypotheses such as Y is independent of $\mathbf{X}_1$ given $\mathbf{X}_2$, where we have partitioned $\mathbf{X} = (\mathbf{X}_1^{\mathrm{T}}, \mathbf{X}_2^{\mathrm{T}})^{\mathrm{T}}$.

Moment-based sufficient dimension-reduction methods provide estimates of the minimal sufficient linear reduction, but they are not designed specifically for prediction and do not produce predictive methods *per se*. Instead, once $\hat{R} = \hat{\boldsymbol{\eta}}^{\mathrm{T}}\mathbf{X}$ has been determined, it is considered fixed and standard model development methods are typically used to estimate $E(Y \mid \hat{\boldsymbol{\eta}}^{\mathrm{T}}\mathbf{X})$ and thereby obtain predictions. Model-based sufficient reduction methods allow a more direct route to estimation of $E(Y \mid \mathbf{X})$.

(b) Model-based inverse reduction

Model-based sufficient dimension reduction is relatively new (Cook 2007). There are several useful characteristics of this approach, which may become clear in the sections that follow. One is that a model for $\mathbf{X} \mid Y$ can itself be inverted to provide a method for estimating the forward mean function $E(Y \mid \mathbf{X})$ without specifying a model for the full joint distribution of $(\mathbf{X}, Y)$. For convenience, we denote the densities of $\mathbf{X}$ and $\mathbf{X} \mid Y$ by $g(\mathbf{X})$ and $g(\mathbf{X} \mid Y)$, and so on, keeping in mind that the symbol g indicates a different density in each case. Because densities will always appear together with their arguments, this should

cause no ambiguity. We assume that $R(\mathbf{X})$ has a density as well. With these understandings, we have

$$
\left.\begin{aligned}
E\{Y \mid \mathbf{X} = \mathbf{x}\} &= \frac{E\{Yg(\mathbf{x} \mid Y)\}}{E\{g(\mathbf{x} \mid Y)\}} \\
&= E\{Y \mid R(\mathbf{x})\} \\
&= \frac{E\{Yg(R(\mathbf{x}) \mid Y)\}}{E\{g(R(\mathbf{x}) \mid Y)\}},
\end{aligned}\right\}
\tag{3.1}
$$

where all right-hand side expectations are with respect to the marginal distribution of Y. Equation $(3.1)_1$ provides a relationship between the mean function $E\{Y \mid \mathbf{X}\}$ and the conditional density of $\mathbf{X} \mid Y$, while equation $(3.1)_2$ is a restatement of the equality of the mean functions for the regressions of Y on $\mathbf{X}$ and $R(\mathbf{X})$. The final equality $(3.1)_3$ establishes a relationship between these forward mean functions and the conditional density of $R \mid Y$, and provides a method to estimate $E(Y \mid \mathbf{X})$:

$$
\left.\begin{aligned}
\hat{E}\{Y \mid \mathbf{X} = \mathbf{x}\} &= \sum_{i=1}^{n} w_i(\mathbf{x}) \, Y_i \\
w_i(\mathbf{x}) &= \frac{\hat{g}(\hat{R}(\mathbf{x}) \mid Y_i)}{\sum_{i=1}^{n} \hat{g}(\hat{R}(\mathbf{x}) \mid Y_i)},
\end{aligned}\right\}
\tag{3.2}
$$

where $\hat{g}$ denotes an estimated density and $\hat{R}$ is the estimated reduction. This estimator is reminiscent of a non-parametric kernel estimator (Simonoff 1996, ch. 4), but there are important differences. The weights in a kernel estimator do not depend on the response, while the weights w_i here do. Kernel weights typically depend on the full vector of predictors $\mathbf{X}$, while the weights here depend on $\mathbf{X}$ only through the estimated reduction $\hat{R}(\mathbf{x})$. If d is small, the estimator (3.2) may avoid the curse of dimensionality. Multivariate kernels are usually taken to be the product of univariate kernels, corresponding here to constraining the components of R to be independent. Finally, there is no explicit bandwidth in our weights as they are determined entirely from $\hat{g}$, which eliminates the need for bandwidth estimation by, for example, cross validation.

The success of this approach depends on obtaining good estimators of the reduction and of its conditional density. In the next section, we address these issues by using normal models for the conditional distribution of $\mathbf{X} \mid Y$. The models in §§4a,b and d were introduced by Cook (2007). The model in §4c is from Cook & Forzani (2009b). Our discussion includes a review of the results that are needed to use equation (3.2).

4. Normal inverse models

Let $\mathbf{X}_y$ denote a random vector distributed as $\mathbf{X} \mid (Y = y)$, and assume that $\mathbf{X}_y$ is normally distributed with mean $\boldsymbol{\mu}_y$ and constant variance matrix $\boldsymbol{\Delta} > 0$, where the inequality means that $\boldsymbol{\Delta}$ is positive definite. Let $\bar{\boldsymbol{\mu}} = E(\mathbf{X})$ and let $\boldsymbol{\Gamma} \in \mathbb{R}^{p \times d}$

denote a basis matrix whose columns form a basis for the d-dimensional subspace $\mathcal{S}_{\boldsymbol{\Gamma}} = \operatorname{span}\{\boldsymbol{\mu}_y - \bar{\boldsymbol{\mu}} \mid y \in S_Y\}$, where S_Y denotes the sample space of Y. Then we can write (Cook 2007)

$$\mathbf{X}_y = \bar{\boldsymbol{\mu}} + \boldsymbol{\Gamma}\boldsymbol{v}_y + \boldsymbol{\varepsilon}, \tag{4.1}$$

where $\boldsymbol{\varepsilon}$ is independent of Y and normally distributed with mean 0 and covariance matrix $\boldsymbol{\Delta}$, and $\boldsymbol{v}_y = (\boldsymbol{\Gamma}^{\mathrm{T}}\boldsymbol{\Gamma})^{-1}\boldsymbol{\Gamma}^{\mathrm{T}}(\boldsymbol{\mu}_y - \bar{\boldsymbol{\mu}}) \in \mathbb{R}^d$; we assume that $\operatorname{var}(\boldsymbol{v}_Y) > 0$. The basis matrix $\boldsymbol{\Gamma}$ is not identifiable in this model, as for any full-rank $d \times d$ matrix $\mathbf{A}$ we can always obtain an equivalent parametrization as $\boldsymbol{\Gamma}\boldsymbol{v}_y = (\boldsymbol{\Gamma}\mathbf{A}^{-1})(\mathbf{A}\boldsymbol{v}_y)$. However, $\operatorname{span}(\boldsymbol{\Gamma})$ is identifiable and estimable, and for this reason we assume without loss of generality that $\boldsymbol{\Gamma}$ is a semiorthogonal matrix, $\boldsymbol{\Gamma}^{\mathrm{T}}\boldsymbol{\Gamma} = \mathbf{I}_d$.

Model (4.1) represents the fact that the translated conditional means $\boldsymbol{\mu}_y - \bar{\boldsymbol{\mu}}$ fall in the d-dimensional subspace $\mathcal{S}_{\boldsymbol{\Gamma}}$. Under model (4.1), $R(\mathbf{X}) = \boldsymbol{\Gamma}^{\mathrm{T}}\boldsymbol{\Delta}^{-1}\mathbf{X}$ is minimal sufficient (Cook & Forzani 2009b), and the goal is thus to estimate $\boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}} = \{\boldsymbol{\Delta}^{-1}\mathbf{z} : \mathbf{z} \in \mathcal{S}_{\boldsymbol{\Gamma}}\}$. As the minimal sufficient reduction is linear, it follows that $\mathcal{S}_{Y|X} = \boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}}$. In other words, in the class of models represented by equation (4.1), moment-based and model-based dimension reduction coincide in the population.

As a convenient notation for describing estimators of $\boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}}$, let $\mathcal{S}_d(\mathbf{A}, \mathbf{B})$ denote the span of $\mathbf{A}^{-1/2}$ times the first d eigenvectors of $\mathbf{A}^{-1/2}\mathbf{B}\mathbf{A}^{-1/2}$, where $\mathbf{A}$ and $\mathbf{B}$ are symmetric matrices and $\mathbf{A}$ is non-singular. The subspace $\mathcal{S}_d(\mathbf{A}, \mathbf{B})$ can also be described as the span of $\mathbf{A}^{-1}$ times the first d eigenvectors of $\mathbf{B}$. We refer to errors having covariance matrix $\boldsymbol{\Delta} = \sigma^2\mathbf{I}_p$ as *isotonic*. Isotonic models are models with isotonic errors. For notational simplicity, we will use $\mathbf{X}_i$ when referring to observations rather than the more awkward notation $\mathbf{X}_{y_i}$.

(a) The principal component model

The isotonic version of model (4.1) is called the *PC model* because the maximum likelihood estimator (MLE) of $\boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}} = \mathcal{S}_{\boldsymbol{\Gamma}}$ is $\mathcal{S}_d(\mathbf{I}_p, \hat{\boldsymbol{\Sigma}})$ and thus the d components of $\hat{R}(\mathbf{X})$ are simply the first d PCs. This relatively simple result is due to the nature of $\boldsymbol{\Delta}$. As the errors are isotonic, the contours of $\boldsymbol{\Delta}$ are circular. When the signal $\boldsymbol{\Gamma}\boldsymbol{v}_y$ is added, the contours of $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}\operatorname{var}(\boldsymbol{v}_Y)\boldsymbol{\Gamma}^{\mathrm{T}} + \sigma^2\mathbf{I}_p$ become p-dimensional ellipses with their longest d axes spanning $\mathcal{S}_{\boldsymbol{\Gamma}}$. The MLE $\hat{\sigma}^2$ of σ^2 is $\hat{\sigma}^2 = \sum_{j=d+1}^{p} \hat{\lambda}_j / p$, where $\hat{\lambda}_1 > \cdots > \hat{\lambda}_{d+1} \geq \cdots \geq \hat{\lambda}_p$ are the eigenvalues of $\hat{\boldsymbol{\Sigma}}$, and the MLE of $\bar{\boldsymbol{\mu}}$ is simply $\bar{\mathbf{X}}$.

We performed a small simulation to provide some intuition. Observations on $\mathbf{X}$ were generated as $\mathbf{X}_y = \boldsymbol{\Gamma}^*\boldsymbol{v}_y^* + \boldsymbol{\varepsilon}$, where $\boldsymbol{v}_y^*$ was sampled uniformly from the boundary of the square $[-1, 1]^2$, the elements of the $p \times 2$ matrix $\boldsymbol{\Gamma}^*$ were sampled independently from a standard normal distribution, and the error vector $\boldsymbol{\varepsilon}$ was sampled from a normal distribution with mean 0 and variance matrix $\mathbf{I}_p$. In terms of model (4.1), $\bar{\boldsymbol{\mu}} = 0$, $\boldsymbol{\Gamma} = \boldsymbol{\Gamma}^*(\boldsymbol{\Gamma}^{*\mathrm{T}}\boldsymbol{\Gamma}^*)^{-1/2}$ and $\boldsymbol{v}_y = (\boldsymbol{\Gamma}^{*\mathrm{T}}\boldsymbol{\Gamma}^*)^{1/2}\boldsymbol{v}_y^*$. This sampling process, which does not require an explicit choice of Y, was repeated $n = 80$ times for various values of p. Figure 1 shows plots of the first two PCs for four values of p. We see that for small p the square is not recognizable, but for larger values of p the square is quite clear. In figure 1d, there are $p = 500$ predictors,

Figure 1. Plots of the estimated sufficient reduction from the PC model with $n = 80$ observations and varying number of predictors p: (*a*) $p = 3$; (*b*) $p = 5$; (*c*) $p = 25$; (*d*) $p = 500$.

while the number of observations is still $n = 80$. The sides of the estimated square in figure 1*d* do not align with the coordinate axes because the method is designed to estimate only the subspace $\boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}}$, which is equal to $\mathcal{S}_{\boldsymbol{\Gamma}}$ with isotonic errors.

Turning to prediction, let $\hat{\boldsymbol{\Gamma}}$ denote the $p \times d$ matrix with columns consisting of the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}$. Then, the weights (3.2) can be written as

$$\left.\begin{aligned} w_i(\mathbf{x}) &\propto \exp\{-(2\hat{\sigma}^2)^{-1}\|\hat{\boldsymbol{\Gamma}}^{\mathrm{T}}(\mathbf{x} - \mathbf{X}_i)\|^2\} \\ &= \exp\{-(2\hat{\sigma}^2)^{-1}\|\hat{R}(\mathbf{x}) - \hat{R}(\mathbf{X}_i)\|^2\}, \end{aligned}\right\} \tag{4.2}$$

where $\hat{R}(\mathbf{x}) = \hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\mathbf{x}$ is the estimated reduction. In this case, $\hat{E}(Y \mid \mathbf{X} = \mathbf{x})$ is in the form of a normal product kernel density estimator with bandwidth $\hat{\sigma}$. Here and in other weights, we assumed that d is known. Methods for choosing d in the context of prediction are discussed in §5.

The method of estimating $E(Y \mid \mathbf{X})$ characterized by equation (4.2) is distinct from the standard PCR estimator given in equation (2.2) with the columns of $\mathbf{G}$ being the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}$. Predictions from equations (4.2) and (2.2) both use PCs, but the way in which they are used is quite different. Model (4.1), which leads to PCs as the MLE of the sufficient reduction, has a very general mean function, but its error structure is restrictive. Nevertheless, this error structure has recently been used in studies of gene expression ($\mathbf{X}$) that are complicated by

Figure 2. Comparison of prediction errors: (*a*) $\mathbf{v}_Y^* = 2.5\,Y$; (*b*) $\mathbf{v}_Y^* = Y + Y^2 + Y^3$. The line passing through the plotted points corresponds to predictions from the inverse model using equation (3.2) with weights (4.2). The other line is for PCR predictions (2.2).

stratification and heterogeneity (Leek & Storey 2007). On the other hand, the usual linear model has a restrictive mean function, and under that model alone there seems to be no clear rationale for reduction by PCs (Cox 1968).

We performed a small simulation study to illustrate the relative behaviour of predictions using equations (2.2) and (4.2). The basic simulation scenario $\mathbf{X}_y = \boldsymbol{\Gamma}^* \mathbf{v}_Y^* + \boldsymbol{\varepsilon}$ was the same as that leading to figure 1, but in this case we set $n = 200$, $d = 1$, $\boldsymbol{\Gamma}^* = (1, \ldots, 1)^{\mathrm{T}}$, $\mathrm{var}(\boldsymbol{\varepsilon}) = 5^2 \mathbf{I}_p$ and $\mathbf{v}_Y^* = 2.5\,Y$, where Y is a standard normal random variable. In this set-up, $(\mathbf{X}, Y)$ has a multivariate normal distribution, so both equations (2.2) and (4.2) are appropriate, but the predictions from PCR (2.2) use the fact that the mean function $E(Y \mid \mathbf{X})$ is linear, while predictions using equation (4.2) do not. For each of 100 datasets generated in this way, the mean-squared prediction error PE_k for the estimated mean function $\hat{E}(Y \mid \mathbf{X})$ was determined by sampling 200 new observations $(\mathbf{X}^*, Y^*)$:

$$\mathrm{PE}_k = \sum_{i=1}^{200} \frac{(Y_i^* - \hat{E}(Y \mid \mathbf{X} = \mathbf{X}_i^*))^2}{200}, \quad k = 1, \ldots, 100. \tag{4.3}$$

The final prediction error was then determined by averaging, $\sum_{k=1}^{100} \mathrm{PE}_k / 100$. The prediction errors are shown in figure 2*a* for $3 \le p \le 150$. While the PCR predictions (2.2) perform a bit better at all values of p, the loss when using predictions from the inverse model is negligible. Figure 2*b* was constructed in the same way, except that we set $\mathbf{v}_Y^* = Y + Y^2 + Y^3$. In this case, the predictions (4.2) from the inverse model are much better than the PCR predictions because they automatically adapt to $E(Y \mid \mathbf{X})$ regardless of the nature of $\mathbf{v}_Y$.

(*b*) *The isotonic principal fitted component model*

It will often be possible and useful to model the coordinate vectors as $\mathbf{v}_y = \boldsymbol{\beta}(\mathbf{f}_y - \bar{\mathbf{f}})$, where $\mathbf{f}_y \in \mathbb{R}^r$ is a known vector-valued function of y with linearly independent elements and $\boldsymbol{\beta} \in \mathbb{R}^{d \times r}$, $d \le \min(r, p)$, is an unrestricted rank-d

matrix. Under this model for $\mathbf{v}_y$, each coordinate X_{yj}, $j = 1, \ldots, p$, of $\mathbf{X}_y$ follows a linear model with predictor vector $\mathbf{f}_y$. Consequently, when Y is quantitative, we are able to use inverse response plots (Cook 1998, ch. 10) of X_{yj} versus y, $j = 1, \ldots, p$, to gain information about suitable choices for $\mathbf{f}_y$, which is an ability that is not generally available in the forward regression of Y on $\mathbf{X}$. When p is too large for such graphical inspection, the PC model can be used as a tool to aid in selecting $\mathbf{f}_y$. Under the PC model, the MLE of $\mathbf{v}_y$ for the ith observed case is $\hat{\mathbf{v}}_i = (\hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\hat{\boldsymbol{\Gamma}})^{-1}\hat{\boldsymbol{\Gamma}}^{\mathrm{T}}(\mathbf{X}_i - \bar{\mathbf{X}}) = \hat{\boldsymbol{\Gamma}}^{\mathrm{T}}(\mathbf{X}_i - \bar{\mathbf{X}})$; the second equality follows because $\hat{\boldsymbol{\Gamma}}$ is a semiorthogonal matrix. These vectors $\hat{\mathbf{v}}_i$ can then be plotted against Y_i, $i = 1, \ldots, n$, and used in the same way as the inverse response plots. In cases like figure 1d, brushing a histogram of the response while observing movement of the corresponding points around the square may be useful. We can also consider vectors $\mathbf{f}_y$ that contain a reasonably flexible set of basis functions, like polynomial terms in Y, which may also be useful when it is impracticable to apply graphical methods to all of the predictors. Piecewise polynomials could also be used.

In some regressions, there may be a natural choice for $\mathbf{f}_y$. Suppose for instance that Y is categorical, taking values in one of h categories C_k, $k = 1, \ldots, h$. We can then set $r = h - 1$ and specify the kth element of $\mathbf{f}_y$ to be $J(y \in C_k)$, where J is the indicator function. Another option consists of 'slicing' the observed values of a continuous Y into h bins (categories) C_k, $k = 1, \ldots, h$, and then specifying the kth coordinate of $\mathbf{f}_y$ as for the case of a categorical Y. This has the effect of approximating each conditional mean $E(X_{yj})$ as a step function of y with h steps,

$$E(X_{yj}) \approx \bar{\mu}_j + \sum_{k=1}^{h-1} \boldsymbol{\gamma}_j^{\mathrm{T}} \mathbf{b}_k \{ J(y \in C_k) - \Pr(Y \in C_k) \},$$

where $\boldsymbol{\gamma}_j^{\mathrm{T}}$ is the jth row of $\boldsymbol{\Gamma}$ and $\mathbf{b}_k$ is the kth column of $\boldsymbol{\beta}$.

Models with $\mathbf{v}_y = \boldsymbol{\beta}(\mathbf{f}_y - \bar{\mathbf{f}})$ are called *principal fitted component* (PFC) *models*. To describe the MLE of $\boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}}$ when the errors are isotonic, let $\mathbb{F}$ denote the $n \times r$ matrix with rows $(\mathbf{f}_i - \bar{\mathbf{f}})^{\mathrm{T}}$. Then, the $n \times p$ matrix of centred fitted vectors from the linear regression of $\mathbf{X}$ on $\mathbf{f}$ is $\hat{\mathbb{X}} = \mathbf{P}_{\mathbb{F}}\mathbb{X}$ and the sample covariance matrix of these fitted vectors is $\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}} = \mathbb{X}^{\mathrm{T}}\mathbf{P}_{\mathbb{F}}\mathbb{X}/n$, where $\mathbf{P}_{\mathbb{F}}$ denotes the linear operator that projects onto the subspace spanned by the columns of $\mathbb{F}$. Under this *isotonic PFC model*, the MLE of $\boldsymbol{\Delta}^{-1}\mathcal{S}_{\boldsymbol{\Gamma}}$ is $\mathcal{S}_d(\mathbf{I}_p, \hat{\boldsymbol{\Sigma}}_{\mathrm{fit}})$. Johnson (2008) gave sufficient conditions for this estimator to converge at the usual $\sqrt{n}$ rate.

Under the isotonic PFC model, the sufficient reduction is estimated as $\hat{R}(\mathbf{x}) = \hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\mathbf{x}$, where the columns of $\hat{\boldsymbol{\Gamma}}$ are the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$. The elements of this $\hat{R}$ are called PFCs. Additionally, the MLE of $\boldsymbol{\beta}$ is $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\mathbb{X}^{\mathrm{T}}\mathbb{F}(\mathbb{F}^{\mathrm{T}}\mathbb{F})^{-1}$ and the MLE of σ^2 is $\hat{\sigma}^2 = (\sum_{j=1}^{p} \hat{\lambda}_j - \sum_{j=1}^{d} \hat{\lambda}_j^{\mathrm{fit}})/p$, where the $\hat{\lambda}_j^{\mathrm{fit}}$ are the ordered eigenvalues of $\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$.

To describe the weights under this model, let $\mathbf{B}_{\mathrm{ols}} = \mathbb{X}^{\mathrm{T}}\mathbb{F}(\mathbb{F}^{\mathrm{T}}\mathbb{F})^{-1}$ be the coefficient matrix from the multivariate OLS fit of $\mathbf{X}$ on $\mathbf{f}$, and let the fitted vectors be denoted by $\hat{\mathbf{X}}_i = \bar{\mathbf{X}} + \mathbf{B}_{\mathrm{ols}}(\mathbf{f}_i - \bar{\mathbf{f}})$. The weights from equation (3.2) are then

$$w_i(\mathbf{x}) \propto \exp\{ -(2\hat{\sigma}^2)^{-1} \| \hat{R}(\mathbf{x}) - \hat{R}(\hat{\mathbf{X}}_i) \|^2 \}. \tag{4.4}$$

Figure 3. Comparison of prediction errors from PCR (equation (2.2)), inverse PCs (equation (4.2)), isotonic principal fitted components (iso. PFC) (equation (4.4)) and diagonal principal fitted components (diag. PFC) (equation (4.5)): (*a*) isotonic PFC, $\mathbf{f}_y = (y, y^2, y^3)^{\mathrm{T}}$; (*b*) diagonal PFC, $f_y = y$.

These weights are of the same general form as those in equation (4.2) from the PC model, but there are two important differences. First, the isotonic PFC reduction uses the eigenvectors from $\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$, while the PC reduction uses the eigenvectors from $\hat{\boldsymbol{\Sigma}}$. Second, the reduction is applied to the observed predictor vectors $\mathbf{X}_i$ in the PC weights (4.2), while the reduction is applied to the fitted predictor vectors $\hat{\mathbf{X}}_i$ in the isotonic PFC weights (4.4).

Shown in figure 3*a* are results from a simulation study to illustrate the potential advantages of using fitted components. The data were generated as for figure 2*b*. We used $\mathbf{f}_y = (y, y^2, y^3)^{\mathrm{T}}$ for the fitted model and, to broaden the scope, we increased the overall variability by reducing the sample size to $n = 50$ and increasing the conditional variance matrix to $\mathrm{var}(\boldsymbol{\varepsilon}) = 15^2 \mathbf{I}_p$. This is a highly variable regression, but all three methods respond quickly as the number of predictors increases. The PCR predictions do better than the inverse predictions from the PC model for relatively small p, but the situation is reversed for larger values of p. Most importantly, the prediction errors from isotonic PFC are the smallest at all values of p. Figure 3*b* is discussed in the next section.

(*c*) *The diagonal principal fitted component model*

The isotonic PFC model requires that, given the response, the predictors must be independent and have the same variance. While this model will be useful in some applications, the requirement of equal variances is restrictive relative to the range of applications in which reduction may be desirable. In this section, we expand the scope of application by permitting a diagonal error covariance matrix $\boldsymbol{\Delta} = \mathrm{diag}(\sigma_1^2, \ldots, \sigma_p^2)$. This allows for different measurement scales of the predictors, but still requires that they be conditionally independent. The estimated sufficient reduction for this model is $\hat{R}(\mathbf{x}) = \hat{\boldsymbol{\Gamma}}^{\mathrm{T}} \hat{\boldsymbol{\Delta}}^{-1} \mathbf{x}$, where $\hat{\boldsymbol{\Delta}}$ is

the MLE of $\boldsymbol{\Delta}$ and $\hat{\boldsymbol{\Gamma}}$ is any basis for the MLE of span($\boldsymbol{\Gamma}$). With this $\hat{R}$, the weights have the same form as the weights for the isotonic PFC model,

$$w_i(\mathbf{x}) \propto \exp\left\{-\tfrac{1}{2}\|\hat{R}(\mathbf{x}) - \hat{R}(\hat{\mathbf{X}}_i)\|^2\right\}, \tag{4.5}$$

where $\hat{\mathbf{X}}$ is as defined for equation (4.4).

It remains to determine $\hat{\boldsymbol{\Delta}}$ and $\hat{\boldsymbol{\Gamma}}$. For this model, we were unable to find a closed-form expression for these estimators. However, a straightforward alternating algorithm can be developed based on the following reasoning. If the inverse mean function is specified, then the variances σ_j^2 can be estimated by using the sample variances of the centred variables $\mathbf{X}_i - \bar{\boldsymbol{\mu}} - \boldsymbol{\Gamma}\boldsymbol{\beta}(\mathbf{f}_i - \bar{\mathbf{f}})$. If $\boldsymbol{\Delta}$ is specified, then we can standardize the predictor vector to obtain an isotonic PFC model in $\mathbf{Z} = \boldsymbol{\Delta}^{-1/2}\mathbf{X}$,

$$\mathbf{Z} = \boldsymbol{\Delta}^{-1/2}\bar{\boldsymbol{\mu}} + \boldsymbol{\Delta}^{-1/2}\boldsymbol{\Gamma}\boldsymbol{\beta}(\mathbf{f} - \bar{\mathbf{f}}) + \boldsymbol{\varepsilon}, \tag{4.6}$$

where $\boldsymbol{\varepsilon}$ is normal with mean $\mathbf{0}$ and variance matrix $\mathbf{I}_p$. Consequently, we can estimate span($\boldsymbol{\Gamma}$) as $\boldsymbol{\Delta}^{1/2}$ times the estimate $\widetilde{\boldsymbol{\Gamma}}$ of $\boldsymbol{\Delta}^{-1/2}\boldsymbol{\Gamma}$ from the isotonic model (4.6). Alternating between these two steps leads to the following algorithm.

1. Fit the isotonic PFC model to the original data, getting initial estimates $\hat{\boldsymbol{\Gamma}}_{(1)}$ and $\hat{\boldsymbol{\beta}}_{(1)}$. The MLE of the intercept $\bar{\boldsymbol{\mu}}$ is always $\bar{\mathbf{X}}$, so there is no need to include this parameter explicitly in the algorithm.
2. For some small $\epsilon > 0$, repeat for $j = 1, 2 \ldots$ until $\mathrm{tr}\{(\hat{\boldsymbol{\Delta}}_{(j)} - \hat{\boldsymbol{\Delta}}_{(j+1)})^2\} < \epsilon$.
   (a) Calculate $\hat{\boldsymbol{\Delta}}_{(j)} = \mathrm{diag}\{(\mathbb{X} - \mathbb{F}\hat{\boldsymbol{\beta}}_{(j)}^{\mathrm{T}}\hat{\boldsymbol{\Gamma}}_{(j)}^{\mathrm{T}})^{\mathrm{T}}(\mathbb{X} - \mathbb{F}\hat{\boldsymbol{\beta}}_{(j)}^{\mathrm{T}}\hat{\boldsymbol{\Gamma}}_{(j)}^{\mathrm{T}})\}$.
   (b) Transform $\mathbf{Z} = \hat{\boldsymbol{\Delta}}_{(j)}^{-1/2}\mathbf{X}$.
   (c) Fit the isotonic PFC model to $\mathbf{Z}$, yielding estimates $\widetilde{\boldsymbol{\Gamma}}$, and $\widetilde{\boldsymbol{\beta}}$.
   (d) Backtransform the estimates to the original scale $\hat{\boldsymbol{\Gamma}}_{(j+1)} = \hat{\boldsymbol{\Delta}}_{(j)}^{1/2}\widetilde{\boldsymbol{\Gamma}}$, $\hat{\boldsymbol{\beta}}_{(j+1)} = \widetilde{\boldsymbol{\beta}}$.

Cook & Forzani (2009b) proposed a different algorithm for fitting the diagonal PFC model. The algorithm here is more reliable when $n < p$.

We again report partial results of a simulation study to illustrate the potential benefits of the diagonal PFC model. The scenario for generating the data was the same as that for figure 2a with $3 \leq p \leq 200$ and $n = 50$, modified to induce different conditional variances. The conditional variances $(\sigma_1^2, \ldots, \sigma_{200}^2)$ were generated once as the order statistics for a sample of size 200 from 60 times a chi-squared random variable with two degrees of freedom. The smallest order statistic was $\sigma_1^2 = 0.07$ and the largest was $\sigma_{200}^2 = 499$. We then used $(\sigma_1^2, \ldots, \sigma_p^2)$ for a regression with p predictors. In this way, the conditional variances increased as we added predictors. While it is unlikely that predictors would be so ordered in practice, this arrangement seems useful to highlight the relative behaviour of the methods. In each case, we fitted the isotonic and diagonal PFC model using $\mathbf{f}_y = y$. Predictions were again assessed using 200 new simulated observations, and the entire set-up was again replicated 100 times to obtain the average prediction errors shown in figure 3b. With $p = 3$ predictors, all methods perform well because

the first three conditional variances are similarly small. The prediction error for the diagonal PFC model changed very little as we added predictors, while the prediction errors for all other methods increased substantially. This happened because the MLEs from the diagonal PFC model can weight each predictor according to its conditional variance, while the predictors are treated equally by the other methods. Predictions from the PC model were consistently the worst. This is perhaps to be expected because the other three methods use mean functions that are appropriate for this simulation.

Predictions from the diagonal PFC model should perform well also when $\boldsymbol{\Delta} > 0$ is not a diagonal matrix, provided that the conditional predictor correlations are small to moderate. This is in agreement with Prentice & Zhao (1991, p. 830) who recommended in a related context that independence models generally should be adequate for a broad range of applications in which the dependences are not strong. The methodology discussed in the next section may be useful in the presence of strong conditional correlations.

(d) The principal fitted component model

In this section, we describe the estimator (3.2) for general $\boldsymbol{\Delta} > 0$, allowing the predictors to be conditionally dependent with different variances. The methods deriving from the isotonic and diagonal PFC models do not require $p < n$. In contrast, the methods of this section work best in data-rich regressions where $p \ll n$. The analysis relies on the MLEs given by Cook (2007, §7.2) and studied in more detail by Cook & Forzani (2009b).

The following notation will help describe the MLEs under the PFC model. Let $\hat{\boldsymbol{\Sigma}}_{\text{res}} = \hat{\boldsymbol{\Sigma}} - \hat{\boldsymbol{\Sigma}}_{\text{fit}} > 0$ be the sample covariance matrix of the residuals vectors $\mathbf{X}_i - \hat{\mathbf{X}}_i$ from the OLS fit of $\mathbf{X}$ on $\mathbf{f}$. Let $\hat{\omega}_1 > \cdots \geq \hat{\omega}_p$ and $\hat{\mathbf{V}} = (\hat{\mathbf{v}}_1, \ldots, \hat{\mathbf{v}}_p)$ be the eigenvalues and the corresponding matrix of eigenvectors $\hat{\mathbf{v}}_j$ of $\hat{\boldsymbol{\Sigma}}_{\text{res}}^{-1/2} \hat{\boldsymbol{\Sigma}}_{\text{fit}} \hat{\boldsymbol{\Sigma}}_{\text{res}}^{-1/2}$. Finally, let $\hat{\mathbf{K}}$ be a $p \times p$ diagonal matrix, with the first d diagonal elements equal to zero and the last $p - d$ diagonal elements equal to $\hat{\omega}_{d+1}, \ldots, \hat{\omega}_p$. Then, the MLE of $\boldsymbol{\Delta}$ is

$$\hat{\boldsymbol{\Delta}} = \hat{\boldsymbol{\Sigma}}_{\text{res}}^{1/2} \hat{\mathbf{V}} (\mathbf{I}_p + \hat{\mathbf{K}}) \hat{\mathbf{V}}^{\text{T}} \hat{\boldsymbol{\Sigma}}_{\text{res}}^{1/2}.$$

If $d = r$, then $\hat{\mathbf{K}} = 0$, and this estimator reduces to $\hat{\boldsymbol{\Delta}} = \hat{\boldsymbol{\Sigma}}_{\text{res}}$. Otherwise, $\hat{\boldsymbol{\Delta}}$ recovers information on variation due to overfitting ($r > d$). The MLE of $\boldsymbol{\beta}$ is $\hat{\boldsymbol{\beta}} = \hat{\boldsymbol{\Gamma}}^{\text{T}} \mathbf{P}_{\hat{\boldsymbol{\Gamma}}(\hat{\boldsymbol{\Delta}}^{-1})} \mathbf{B}_{\text{ols}}$, where $\mathbf{P}_{\hat{\boldsymbol{\Gamma}}(\hat{\boldsymbol{\Delta}}^{-1})} \mathbf{B}_{\text{ols}}$ is the projection of $\mathbf{B}_{\text{ols}}$ onto the span of $\hat{\boldsymbol{\Gamma}}$ in the $\hat{\boldsymbol{\Delta}}^{-1}$ inner product, and the MLE of $\boldsymbol{\Delta}^{-1}\boldsymbol{\Gamma}$ is $\mathcal{S}_d(\hat{\boldsymbol{\Delta}}, \hat{\boldsymbol{\Sigma}}_{\text{fit}})$, which is equal to the SIR estimator of $\mathcal{S}_{Y|X}$ when the response is categorical (Cook & Forzani 2009b).

Turning to prediction, we substituted these MLEs into the multivariate normal density for $\mathbf{X} \mid Y$ and simplified the resulting expression by in part ignoring proportionality constants not depending on the observation i to obtain the weights (3.2):

$$w_i(\mathbf{x}) \propto \exp\left\{-\tfrac{1}{2}(\mathbf{x} - \hat{\mathbf{X}}_i)^{\text{T}} [\hat{\boldsymbol{\Delta}}^{-1} \hat{\boldsymbol{\Gamma}} (\hat{\boldsymbol{\Gamma}}^{\text{T}} \hat{\boldsymbol{\Delta}}^{-1} \hat{\boldsymbol{\Gamma}})^{-1} \hat{\boldsymbol{\Gamma}}^{\text{T}} \hat{\boldsymbol{\Delta}}^{-1}](\mathbf{x} - \hat{\mathbf{X}}_i)\right\}, \qquad (4.7)$$

where $\hat{\mathbf{X}}_i$ is as defined for equation (4.4). We next discuss how these weights can be simplified to a more intuitive and computationally efficient form.

Let $\tilde{\mathbf{V}} = \hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\hat{\mathbf{V}}(\mathbf{I}_p + \hat{\mathbf{K}})^{-1/2}$. Then it is easy to show that $\hat{\boldsymbol{\Delta}}^{-1} = \tilde{\mathbf{V}}\tilde{\mathbf{V}}^{\mathrm{T}}$, and that the columns of $\hat{\boldsymbol{\Delta}}^{1/2}\tilde{\mathbf{V}}$ are the normalized eigenvectors of $\hat{\boldsymbol{\Delta}}^{-1/2}\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}\hat{\boldsymbol{\Delta}}^{-1/2}$. Let $\tilde{\mathbf{V}}_d$ and $\hat{\mathbf{V}}_d$ denote the $p \times d$ matrices consisting of the first d columns of $\tilde{\mathbf{V}}$ and $\hat{\mathbf{V}}$. Then, because the MLE of $\boldsymbol{\Delta}^{-1}\boldsymbol{\Gamma}$ is $\mathcal{S}_d(\hat{\boldsymbol{\Delta}}, \hat{\boldsymbol{\Sigma}}_{\mathrm{fit}})$, it follows that $\mathcal{S}_d(\hat{\boldsymbol{\Delta}}, \hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}) = \mathrm{span}(\boldsymbol{\Delta}^{-1/2}\boldsymbol{\Delta}^{1/2}\tilde{\mathbf{V}}) = \mathrm{span}(\tilde{\mathbf{V}}_d) = \mathrm{span}(\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\hat{\mathbf{V}}_d)$. Consequently, we may take $\hat{\boldsymbol{\Delta}}^{-1}\hat{\boldsymbol{\Gamma}} = \tilde{\mathbf{V}}_d = \hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\hat{\mathbf{V}}_d$, which implies the reduction

$$\hat{R}(\mathbf{x}) = \hat{\mathbf{V}}_d^{\mathrm{T}}\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\mathbf{x} = (\hat{\mathbf{v}}_1^{\mathrm{T}}\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\mathbf{x}, \ldots, \hat{\mathbf{v}}_d^{\mathrm{T}}\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\mathbf{x})^{\mathrm{T}}.$$

Using these results in equation (4.7), we have $\hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\hat{\boldsymbol{\Delta}}^{-1}\hat{\boldsymbol{\Gamma}} = (\hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\hat{\boldsymbol{\Delta}}^{-1})\hat{\boldsymbol{\Delta}}(\hat{\boldsymbol{\Delta}}^{-1}\hat{\boldsymbol{\Gamma}}) = \tilde{\mathbf{V}}_d^{\mathrm{T}}\hat{\boldsymbol{\Delta}}\tilde{\mathbf{V}}_d = \mathbf{I}_d$, and thus

$$\left.\begin{aligned} w_i(\mathbf{x}) &\propto \exp\left\{-\tfrac{1}{2}(\mathbf{x} - \hat{\mathbf{X}}_i)^{\mathrm{T}}[\hat{\boldsymbol{\Delta}}^{-1}\hat{\boldsymbol{\Gamma}}\hat{\boldsymbol{\Gamma}}^{\mathrm{T}}\hat{\boldsymbol{\Delta}}^{-1}](\mathbf{x} - \hat{\mathbf{X}}_i)\right\} \\ &= \exp\left\{-\tfrac{1}{2}(\mathbf{x} - \hat{\mathbf{X}}_i)^{\mathrm{T}}\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\hat{\mathbf{V}}_d\hat{\mathbf{V}}_d^{\mathrm{T}}\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}(\mathbf{x} - \hat{\mathbf{X}}_i)\right\} \\ &= \exp\left\{-\tfrac{1}{2}\|\hat{R}(\mathbf{x}) - \hat{R}(\hat{\mathbf{X}}_i)\|^2\right\}. \end{aligned}\right\} \tag{4.8}$$

The reduction in this case is applied to $\mathbf{x}$ and $\hat{\mathbf{X}}$ in the same way that it was for the weights from the isotonic and diagonal PFC models, but the reduction itself differs. A perhaps clearer connection between the PFC reduction and the isotonic PFC reduction can be seen by noting that $\mathrm{span}(\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1/2}\hat{\mathbf{V}}_d)$ is equal to the span of the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1}\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$. Under the isotonic PFC model, the reduction is computed by using the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$, while under the PFC model the reduction is computed using the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}^{-1}\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$. Additionally, as $\mathcal{S}_d(\hat{\boldsymbol{\Sigma}}_{\mathrm{res}}, \hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}) = \mathcal{S}_d(\hat{\boldsymbol{\Sigma}}, \hat{\boldsymbol{\Sigma}}_{\mathrm{res}})$, the reduction could be computed also as the first d eigenvectors of $\hat{\boldsymbol{\Sigma}}^{-1}\hat{\boldsymbol{\Sigma}}_{\mathrm{fit}}$.

5. Implementation

In this section, we present additional simulation results and provide some suggestions for issues that must be addressed prior to implementation. We consider only the isotonic, diagonal and general PFC models, and address the implementation issues first.

(a) Choice of d

Two general methods for choosing $d \in \{0, 1, \ldots, \min(r, p)\}$ have been studied in the literature. One is by using an information criterion such as the Akaike (AIC) and Bayesian (BIC) information criteria. Let $L(d_0)$ denote the value of the maximized log likelihood for any of the three inverse models under consideration fitted with a value d_0 for d. Then the dimension is selected that minimizes the information criterion $-2L(d_0) + h(n)g(d_0)$, where $h(n)$ is equal to $\log(n)$ for BIC and 2 for AIC, and $g(d_0) = p + d_0(p - d_0) + dr + D(\boldsymbol{\Delta})$ is the number of parameters to be estimated as a function of d_0. The first addend in this count corresponds to $\boldsymbol{\mu}$, the second to $\mathcal{S}_{\boldsymbol{\Gamma}}$ and the third to $\boldsymbol{\beta}$. The final term $D(\boldsymbol{\Delta})$ is the number of parameters required for $\boldsymbol{\Delta}$, which is equal to 1, p and $p(p + 1)/2$ for the isotonic, diagonal and PFC models. The second method is based on using the

likelihood ratio statistic $2(L(\min(r, p)) - L(d_0))$ in a sequential scheme to choose d. Using a common test level and starting with $d_0 = 0$, choose the estimate of d as the hypothesized value that is not rejected; see Cook (2007) and Cook & Forzani (2009a) for further discussion of these methods.

These methods for choosing d have been shown to work well in data-rich regressions where $p \ll n$, but can be unreliable otherwise. Consequently, in the context of this paper, we will determine d using the cross-validation method proposed by Hastie *et al.* (2001, ch. 7).

(b) Screening

When dealing with large-p regressions, there may be a possibility that a substantial subset $\mathbf{X}_2$ of the predictors is inactive; that is, $\mathbf{X}_2$ is independent of the response given the remaining predictors $\mathbf{X}_1$. In terms of PFC models, $\mathbf{X}_2$ is inactive if and only if the corresponding rows of $\boldsymbol{\Delta}^{-1}\boldsymbol{\Gamma}$ are all equal to zero. Addressing this possibility in terms of the general PFC models requires a data-rich regression, and is outside the scope of this paper (see Cook & Forzani 2009b for further discussion). For the isotonic and diagonal PFC models, the predictors $\mathbf{X}_2$ are inactive if and only if the corresponding rows of $\boldsymbol{\Gamma}$ are all equal to 0. In these models, we can test if the kth row of $\boldsymbol{\Gamma}$ is equal to 0 straightforwardly by using an F statistic to test if $\boldsymbol{\alpha} = 0$ in the univariate regression model of the kth predictor X_k on $\mathbf{f}$, $X_k = \mu + \boldsymbol{\alpha}^{\mathrm{T}}\mathbf{f} + \varepsilon$. This method with a common test level of 0.1 was used for screening in the simulations that follow. When $d = 1$ and $\mathbf{f} = y$, this method is equivalent to sure independence screening proposed by Fan & Lv (2008).

(c) Simulation results

To allow for some flexibility, we used the following general set-up as the basis for our simulations. Let $\mathbf{J}_k$ and $\mathbf{O}_k$ denote vectors of length k having all elements equal to 1 and 0. A response was generated from a uniform $(0, 3)$ distribution. Then with $\mathbf{f}_y = (y, \mathrm{e}^{2y})^{\mathrm{T}}$ a predictor vector of length p was generated using an integer $p_0 \leq p/2$ as $\mathbf{X}_y = \boldsymbol{\Gamma}^*\boldsymbol{\beta}^*\mathbf{f}_y + \boldsymbol{\varepsilon}$, where $\boldsymbol{\Gamma}^* = (\boldsymbol{\Gamma}_1^*, \boldsymbol{\Gamma}_2^*)$ with $\boldsymbol{\Gamma}_1^* = (\mathbf{J}_{p_0}^{\mathrm{T}}, \mathbf{O}_{p-p_0}^{\mathrm{T}})^{\mathrm{T}}$ and $\boldsymbol{\Gamma}_2^* = (\mathbf{O}_{p_0}^{\mathrm{T}}, \mathbf{J}_{p_0}^{\mathrm{T}}, \mathbf{O}_{p-2p_0}^{\mathrm{T}})^{\mathrm{T}}$, $\boldsymbol{\beta}^* = \mathrm{diag}(2/\sqrt{20}, 1/\sqrt{20})$ and $\boldsymbol{\varepsilon} \sim N(0, \boldsymbol{\Delta})$ with $\boldsymbol{\Delta} = \mathrm{diag}(\sigma_0^2\mathbf{J}_{p_0}^{\mathrm{T}}, \sigma_1^2\mathbf{J}_{p_0}^{\mathrm{T}}, \sigma_0^2\mathbf{J}_{p-2p_0}^{\mathrm{T}})$, $\sigma_0^2 = 2$ and $\sigma_1^2 = 20$. Datasets of n independent observations generated in this way have the following properties: (i) one set of p_0 predictors is linearly related to the response, another set of p_0 is nonlinearly related and the remaining $p - 2p_0$ predictors are independent of the response; (ii) the conditional covariance of $\mathbf{X} \mid Y$ has a diagonal structure; (iii) $\boldsymbol{\Gamma}$ is a $p \times 2$ matrix and thus the sufficient reduction is composed of $d = 2$ linear combinations of the predictors.

We present results from three cases of this simulation model. In each case, we fitted the diagonal PFC model with $\mathbf{f}_y = (y^1, \ldots, y^4)^{\mathrm{T}}$, so the fitted $\mathbf{f}_y$ was not the same as that used to generate the data. The prediction errors were determined using the method of equation (4.3).

Case 1. $p \ll n$. In this regression, we set $n = 400$, $p = 40$ and $p_0 = 20$, so all predictors are relevant and no screening was used. The results are shown in the third column of table 1. The first row shows the prediction errors when d is determined by the cross-validation (c.v.) method of Hastie *et al.* (2001, ch. 7). In this case, the estimated d can vary across the 100 replicate datasets. The next

Table 1. Prediction errors (4.3) for three simulation scenarios, with standard errors given in parentheses.

method	method class	case 1	case 2	case 3
PFC	$\hat{d}$ by c.v.	0.063 (0.0031)	0.07 (0.006)	0.090 (0.0047)
	$d = 1$	0.073 (0.0036)	0.08 (0.006)	0.092 (0.0053)
	$d = 2$ (true)	0.063 (0.0031)	0.07 (0.005)	0.087 (0.0049)
	$d = 3$	0.065 (0.0031)	0.07 (0.005)	0.089 (0.0046)
	$d = 4$	0.066 (0.0030)	0.08 (0.005)	0.093 (0.0045)
forward methods	lasso	0.18 (0.002)	0.19 (0.003)	0.26 (0.005)
	PLS	0.18 (0.002)	0.27 (0.004)	0.27 (0.004)

four rows show the prediction errors when d is fixed at the indicated value across all datasets. The final two rows show the prediction errors for PLS and the lasso applied straightforwardly to the linear regression of Y on $\mathbf{X}$ using the R packages 'pls' and 'lars'. We see that the PFC prediction errors are fairly stable across the choices for d and that they are roughly one-third of the prediction error for the forward methods.

Case 2. $n < p$. In this simulation, we set $n = 100$, $p_0 = 60$ and $p = 120$, so again all predictors are relevant and no screening was used. The results, which are shown in the fourth column of table 1, are qualitatively similar to those for case 1, except that the prediction error for the lasso is notably less than that for PLS.

Case 3. $n < p$. For this case, we set $n = 100$, $p_0 = 20$ and $p = 120$. Now there are only 40 relevant predictors out of 120 total predictors and, as discussed in §5*b*, screening was used prior to fitting the diagonal PFC model. The results shown in the final column of table 1 are qualitatively similar to the other two cases. Here, we were a bit surprised that the lasso did not perform better because it is designed specifically for sparse regressions. Evidently, the adaptability of the diagonal PFC model through the choices of $\mathbf{f}$ and d can be important for good predictions.

(*d*) Illustration

We conclude this section with an illustrative analysis of an economic regression using data from Enz (1991). The data and additional discussion were given by Cook & Weisberg (1999, p. 140). The response, which was measured in 1991, is the average minutes of labour required to buy a BigMac hamburger and french fries in $n = 45$ world cities. There are nine continuous predictors. We fitted the diagonal PFC model as discussed in §4*c*, and estimated the prediction error using leave-one-out cross validation. We set $d = 1$ to allow an even-handed comparison with forward methods.

The results are shown in table 2. Two entries are given for PFC. The first corresponds to fitting PFC with the cubic polynomial basis $\mathbf{f} = (y, y^2, y^3)^{\mathrm{T}}$ and the second uses a more elaborate piecewise constant polynomial with 10 slices for $\mathbf{f}$. PFC with the piecewise constant polynomial basis gives a better prediction error than PFC with the cubic polynomial basis, but both methods outperform the predictions by five forward methods, the elastic net (Zou &

Table 2. BigMac dataset.

methods	prediction error
PFC—cubic polynomial	933
PFC—piecewise constant	771
elastic net	1198
ridge regression	1211
lasso	1412
PLS	1426
OLS	2268

Hastie 2005), the lasso, ridge regression, PLS and OLS. The relative behaviour of PFC illustrated here is typical of our experiences, regardless of the number of predictors.

6. Discussion

The methodology introduced in this paper applies across a useful range of applications with continuous predictors that, given the response, are mildly correlated. We view the ability to choose $\mathbf{f}_y$ as an important feature because it allows a level of adaptability that is not really available when directly modelling the mean function $E(Y \mid \mathbf{X})$ in large-p regressions. The assumption of normal errors is not crucial, but the predictions may not perform well under extreme deviations from normality. For instance, we expect that it is possible to develop substantially better methodology for regressions in which all the predictors are binary, perhaps employing a quadratic exponential family for $\mathbf{X} \mid Y$ in place of the multivariate normal. Work along these lines is in progress.

It also seems possible to develop models 'between' the diagonal and general PFC models that allow for some of the conditional predictor correlations to be substantial but stop short of permitting a general $\boldsymbol{\Delta} > 0$. One route is to model the conditional covariance matrix as $\boldsymbol{\Delta} = \boldsymbol{\Gamma}\mathbf{M}\boldsymbol{\Gamma}^{\mathrm{T}} + \boldsymbol{\Gamma}_0\mathbf{M}_0\boldsymbol{\Gamma}_0^{\mathrm{T}}$, where $\boldsymbol{\Gamma}$ is the same as that in model (4.1), $(\boldsymbol{\Gamma}, \boldsymbol{\Gamma}_0)$ is an orthogonal $p \times p$ matrix, $\mathbf{M} > 0$ and $\mathbf{M}_0 > 0$ (Cook 2007). The minimal sufficient reduction in this case is $R(\mathbf{X}) = \boldsymbol{\Gamma}^{\mathrm{T}}\mathbf{X}$, which does not involve the $\mathbf{M}$s. This model for $\boldsymbol{\Delta}$ allows arbitrary correlations among the elements of R and among the elements of $\boldsymbol{\Gamma}_0^{\mathrm{T}}\mathbf{X}$, but requires that R and $\boldsymbol{\Gamma}_0^{\mathrm{T}}\mathbf{X}$ be conditionally independent. Predictions under this model in a data-rich regression can be obtained as an extension of results by Cook (2007). The behaviour of this model in $n < p$ regressions is under study.

References

Adcock, R. J. 1878 A problem in least squares. *Analyst* **5**, 53–54. (doi:10.2307/2635758)

Alter, O., Brown, P. & Botstein, D. 2000 Singular value decomposition for gene-wide expression data processing and modelling. *Proc. Natl. Acad. Sci. USA* **97**, 10 101–10 106. (doi:10.1073/pnas.97.18.10101)

Cook, R. D. 1994 Using dimension-reduction subspaces to identify important inputs in models of physical systems. In *Proc. Section on Physical and Engineering Sciences*, pp. 18–25. Alexandria, VA: American Statistical Association.

Cook, R. D. 1996 Graphics for regressions with a binary response. *J. Am. Stat. Assoc.* **91**, 983–992. (doi:10.2307/2291717)

Cook, R. D. 1998 *Regression graphics.* New York, NY: Wiley.

Cook, R. D. 2007 Fisher lecture: dimension reduction in regression (with discussion). *Stat. Sci.* **22**, 1–26. (doi:10.1214/088342306000000682)

Cook, R. D. & Forzani, L. 2009*a* Likelihood-based sufficient dimension reduction. *J. Am. Stat. Assoc.* **104**, 197–208. (doi:10.1198/jasa.2009.0106)

Cook, R. D. & Forzani, L. 2009*b* Principal fitted components for dimension reduction in regression. *Stat. Sci.* **23**, 485–501. (doi:10.1214/08-STS275)

Cook, R. D. & Ni, L. 2005 Sufficient dimension reduction via inverse regression: a minimum discrepancy approach. *J. Am. Stat. Assoc.* **100**, 410–428. (doi:10.1198/016214504000001501)

Cook, R. D. & Weisberg, S. 1982 *Residuals and influence in regression.* London, UK: Chapman and Hall. See http://www.stat.umn.edu/rir.

Cook, R. D. & Weisberg, S. 1991 Discussion of 'Sliced inverse regression for dimension reduction' by K. C. Li. *J. Am. Stat. Assoc.* **86**, 328–332. (doi:10.2307/2290564)

Cook, R. D. & Weisberg, S. 1999 *Applied regression including computing and graphics.* New York, NY: Wiley.

Cook. R. D., Li. B. & Chiaromonte, F. 2007 Dimension reduction without matrix inversion. *Biometrika* **94**, 569–584. (doi:10.1093/biomet/asm038)

Cox, D. R. 1968 Notes on some aspects of regression analysis. *J. R. Stat. Soc. A* **131**, 265–279.

Enz, R. 1991 *Prices and earnings around the globe.* Zurich, Switzerland: Union Bank of Switzerland.

Fan, J. & Lv, J. 2008 Sure independence screening for ultrahigh dimensional feature space (with discussion). *J. R. Stat. Soc. B* **70**, 849–911. (doi:10.1111/j.1467-9868.2008.00674.x)

Fisher, R. J. 1922 On the mathematical foundations of theoretical statistics. *Phil. Trans. R. Soc. Lond. A* **222**, 309–368. (doi:10.1098/rsta.1922.0009)

Fukumizu, K., Bach, F. R. & Jordan, M. I. 2009 Kernel dimension reduction in regression. *Ann. Stat.* **37**, 1871–1905. (doi:10.1214/08-AOS637)

Hastie, T., Tibshirani, R. & Friedman, J. 2001 *The elements of statistical learning—data mining, inference, and prediction.* New York, NY: Springer.

Helland, I. S. 1990 Partial least squares regression and statistical models. *Scand. J. Stat.* **17**, 97–114.

Johnson, O. 2008 Theoretical properties of Cook's PFC dimension reduction algorithm for linear regression. *Electron. J. Stat.* **2**, 807–827. (doi:10.1214/08-EJS255)

Leek, J. T. & Storey, J. D. 2007 Capturing heterogeneity in gene expression studies by surrogate variable analysis. *PLoS Genet.* **3**, 1724–1735. (doi:10.1371/journal.pgen.0030161)

Li, K.-C. 1991 Sliced inverse regression for dimension reduction. *J. Am. Stat. Assoc.* **86**, 316–342. (doi:10.2307/2290563)

Li, L. 2007 Sparse sufficient dimension reduction. *Biometrika* **94**, 603–613. (doi:10.1093/biomet/asm044)

Li, B. & Wang, S. 2007 On directional regression for dimension reduction. *J. Am. Stat. Assoc.* **102**, 997–1008. (doi:10.1198/016214507000000536)

Li, L. & Yin, X. 2008 Sliced inverse regression with regularizations. *Biometrics* **64**, 124–131. (doi:10.1111/j.1541-0420.2007.00836.x)

Naik, P. & Tsai, C.-L. 2000 Partial least squares estimator for single-index models. *J. R. Stat. Soc. B* **62**, 763–771.

Prentice, R. L. & Zhao, L. P. 1991 Estimating equations for parameters in means and covariances of multivariate discrete and continuous responses. *Biometrics* **47**, 825–838. (doi:10.2307/2532642)

Simonoff, J. S. 1996 *Smoothing methods in statistics.* New York, NY: Springer.

Tibshirani, R. 1996 Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. B* **58**, 267–288.

Wand, M. P. & Jones, M. C. 1995 *Kernel smoothing.* London, UK: Chapman and Hall.

Wegman, E. J., Scott, D. W. & Said, Y. H. 2006 Ad hoc committee report on the 'hockey stick' global climate reconstruction, Report to House Committee on Energy and Commerce, United States House of Representatives. See http://republicans.energycommerce.house.gov/108/home/ 07142006_Wegman_Report.pdf.

Ye, Z. & Weiss, R. 2003 Using the bootstrap to select one of a new class of dimension reduction methods. *J. Am. Stat. Assoc.* **98**, 968–978. (doi:10.1198/016214503000000927)

Zhu, L., Ohtaki, M. & Li, Y. 2005 On hybrid methods of inverse regression-based algorithms. *Comput. Stat. Data Anal.* **51**, 2621–2635. (doi:10.1016/j.csda.2006.01.005)

Zou, H. & Hastie, T. 2005 Regularization and variable selection via the elastic net. *J. R. Stat. Soc. B* **68**, 301–320.