



Published in final edited form as:

J Am Stat Assoc. 2012 ; 107(497): 168–179. doi:10.1080/01621459.2011.646925.

A Semiparametric Approach to Dimension Reduction

Yanyuan Ma [Professor] and

Department of Statistics, Texas A&M University, 3143 TAMU, College Station, TX 77843-3143
(ma@stat.tamu.edu)

Liping Zhu [Associate Professor]

School of Statistics and Management, Shanghai University of Finance and Economics, Shanghai 200433, China (zhu.liping@mail.shufe.edu.cn)

Abstract

We provide a novel and completely different approach to dimension-reduction problems from the existing literature. We cast the dimension-reduction problem in a semiparametric estimation framework and derive estimating equations. Viewing this problem from the new angle allows us to derive a rich class of estimators, and obtain the classical dimension reduction techniques as special cases in this class. The semiparametric approach also reveals that in the inverse regression context while keeping the estimation structure intact, the common assumption of linearity and/or constant variance on the covariates can be removed at the cost of performing additional nonparametric regression. The semiparametric estimators without these common assumptions are illustrated through simulation studies and a real data example. This article has online supplementary material.

Keywords

Estimating equations; Nonparametric regression; Robustness; Semiparametric methods; Sliced inverse regression

1. INTRODUCTION

Dimension reduction has been an active field of statistical research for the last 20 years and continues to be important due to the increasingly large amount of available covariates in various scientific areas. The goal of dimension reduction is to identify one or multiple directions represented by a matrix β , so that the response Y relates to the covariate vector \mathbf{x} only through a few linear combinations $\mathbf{x}^T\beta$. When the conditional distribution depends on $\mathbf{x}^T\beta$, it is a problem of estimating the central space (Cook 1998); when the conditional mean $E(Y|\mathbf{x})$ depends on $\mathbf{x}^T\beta$, it is a problem of estimating the central mean space (Cook and Li 2002).

Started with the ingenious idea of sliced inverse regression (SIR) in the seminal article by Li (1991), many highly effective methods in the area of dimension reduction have been developed. For identifying the central space, see, for example, sliced average variance estimation (SAVE) (Cook and Weisberg 1991) and directional regression (DR) (Li and Wang 2007) and their variations such as kernel inverse regression (Zhu and Fang 1996), CANCOR analysis (Fung et al. 2002), and so on. These methods extend the inverse

regression idea and are promising in recovering the central space. However, they all rely on certain conditions. These conditions mainly include the linearity condition, where $E(\mathbf{x} | \mathbf{x}^T\boldsymbol{\beta})$ is assumed to be a linear function of \mathbf{x} , and the constant variance condition, where $\text{cov}(\mathbf{x} | \mathbf{x}^T\boldsymbol{\beta})$ is assumed to be a constant matrix. These conditions are not always satisfied, and sometimes could imply stringent assumptions on the joint distribution of \mathbf{x} . To be precise, SIR requires the linearity condition; SAVE and DR require both the linearity condition and the constant variance condition. If the covariates do not satisfy these two conditions, current practice often relies on transformation (Box and Cox 1964) or reweighting (Cook and Nachtsheim 1994), which can restore these conditions sometimes. Li and Dong (2009) and Dong and Li (2010) successfully remove the linearity condition from the dimension-reduction problems while their estimators remain to be the inverse regression type. This is no doubt a great advancement. The residual issue is that they assumed $E(\mathbf{x} | \mathbf{x}^T\boldsymbol{\beta})$ to be a polynomial function of $\mathbf{x}^T\boldsymbol{\beta}$ and they still required the constant variance condition. These remaining requirements can still be stringent and difficult to check in practice. Zhu and Zeng (2006) introduced a dimension-reduction method to identify the central subspace through using Fourier transformations. Their method, however, requires one to estimate the joint probability density function (pdf) of \mathbf{x} , which is typically infeasible in a high-dimensional environment. To circumvent this difficulty, they assumed \mathbf{x} to be multivariate normal in implementations. Adapting the idea of minimum average variance estimation (MAVE) (Xia et al. 2002), Xia (2007) proposed a similar procedure (dMAVE) to recover the central space, and Wang and Xia (2008) proposed sliced regression (SR) for dimension reduction. However, their methods estimate the distribution function nonparametrically and heavily rely on the implicit assumption that all the covariates are continuous. Because all the existing dimension-reduction methods impose either the above two conditional moment conditions or distributional assumptions on the covariate vector in one form or another, new dimension-reduction methods which are free of any of these assumptions are highly in demand, particularly when some covariates are discrete or categorical.

Similarly, a large amount of literature exists for identifying the central mean space. For example, when the central mean space is one dimensional, Li and Duan (1989) suggested using the ordinary least squares (OLS), assuming \mathbf{x} to satisfy the linearity condition. Härdle and Stoker (1989) and Power, Stock, and Stoker (1989) proposed the average derivative estimation, which requires \mathbf{x} to be continuous. Horowitz and Härdle (1996) proposed a method that allows some covariates to be discrete; however, the number of levels in discrete variables cannot be large. Ichimura (1993) and Härdle, Hall, and Ichimura (1993) suggested using nonlinear least squares, which are essentially special cases of our proposal which will be described in the following. When the central mean space is possibly two or more dimensional, Xia et al. (2002) proposed the minimum average variance estimation method, which, similar to the constructive approach in Xia (2007) and the sliced regression in Wang and Xia (2008), requires \mathbf{x} to be continuous. Li (1992) and Cook and Li (2002) proposed the method of principal Hessian directions which requires \mathbf{x} to satisfy both the linearity condition and the constant variance condition, which are not always satisfied in practice. Assuming conditional normality of \mathbf{x} on Y , Cook and Forzani (2009) proposed a likelihood-based method, and in the context of classification, Hernández and Velilla (2005) estimated the dimension reduction space via minimizing a criterion function which involves kernel density estimation. Yin and Cook (2005); Yin, Li, and Cook (2008); and Park, Sriram, and Yin (2010) proposed a method to recover the dimension-reduction space via minimizing a Kullback–Leibler distance.

In this article, we provide a completely different viewpoint for looking at the dimension-reduction problems. Our approach is through semiparametrics, which has not been considered in the literature. By casting the dimensional-reduction problem in the semiparametric framework, the dimension-reduction problems become semiparametric

estimation problems. Therefore, powerful semiparametric estimation and inference tools become applicable. We use the geometric approach in Bickel et al. (1993) and Tsiatis (2006) to analyze these problems and derive the space of the influence functions. This enables us to construct a rich class of estimators. Many of the existing dimension-reduction methods turn out to be special cases in this class. In fact, the complete class of influence functions provide all the possible consistent estimators.

A direct consequence of the semiparametric analysis is the relaxation of the linearity condition and the constant variance condition. Using the semiparametric construction, we reveal that these conditions are not structurally necessary. Consequently, in all these existing dimension-reduction procedures, we can remove these two conditions, and instead replace the assumed quantity with nonparametric estimation of the corresponding conditional expectations. Thus, the semiparametric derivation allows us to obtain the dimension reduction spaces without any distributional assumption on the covariate vector. Another advantage of the semiparametric analysis is that we do not require all the covariates to be continuous.

In summary, the contributions of this article are:

1. We introduce a novel and drastically different approach to the dimension-reduction field. We anticipate to stimulate deeper and richer literature in this direction.
2. We derive the complete class of influence functions, which guarantees to yield all possible root- n consistent estimators for the column space of β . We demonstrate how to obtain several most popular dimension-reduction methods from this class. This further reveals the underlying connection between these different methods and provides a different and natural motivation for their construction.
3. We completely eliminate the linearity condition, the constant variance condition, the condition on the quadratic form of the covariates, or, in fact, any moment conditions on the covariates at all.
4. We eliminate the redundant continuity conditions on the covariates. The new approach can be readily used even when some covariates are categorical or discrete.

The outline of this article is the following. In Section 2, we describe the semiparametric approach to the central space estimation and derive a rich class of estimators. We establish their link and generalization to several existing dimension-reduction methods in Section 3. The analysis for central mean space estimation is given in Section 4. We explain the implementation details on estimation and on selecting the dimension of the central space/central mean space in Section 5. Extensive simulation studies are conducted in Section 6 to demonstrate the practical performance and the method is implemented in a real data example in Section 7. We finish the article with a brief discussion in Section 8. Technical derivations are collected in an appendix and the online supplementary document.

2. ESTIMATING THE CENTRAL SUBSPACE VIA SEMIPARAMETRICS

Let \mathbf{x} be a $p \times 1$ covariate vector and Y a univariate response variable. The goal of sufficient dimension reduction (Cook 1998) is to seek a matrix β such that

$$F(y|\mathbf{x})=F(y|\mathbf{x}^T\beta), \text{ for } y \in \mathbb{R}, \quad (1)$$

where $F(y|\mathbf{x}) \stackrel{\text{def}}{=} \Pr(Y \leq y|\mathbf{x})$ denotes the conditional distribution function of Y given \mathbf{x} . (1) implies that the response variable Y relates to \mathbf{x} only through linear combinations $\mathbf{x}^T\beta$. The

column space of β satisfying (1) is called a dimension-reduction subspace. Because the dimension-reduction subspace is not unique, our primary interest is the central subspace, which is defined as the intersection of all dimension-reduction subspaces, provided that the intersection itself is a dimension-reduction subspace (Cook 1998).

Following the convention in the area of dimension reduction, we denote by $S_{Y|\mathbf{x}}$ the central subspace and assume β to satisfy $\beta^T \text{cov}(\mathbf{x})\beta = \mathbf{I}_d$. Note that β is more restrictive than before, however to avoid introducing a new notation, we keep the same notation. We also assume $S_{Y|\mathbf{x}}$ exists and is unique. Here, the number of columns in β , denoted by d , is the dimension of $S_{Y|\mathbf{x}}$ and is often referred to as the structural dimension. Our goal is to find $S_{Y|\mathbf{x}}$ through finding $\beta \in \mathbb{R}^{p \times d}$ which satisfies (1). To focus on the main issues of the dimension-reduction problems, we assume throughout our article that the covariate vector \mathbf{x} satisfies $E(\mathbf{x}) = \mathbf{0}$ and $\text{cov}(\mathbf{x}) = \mathbf{I}_p$. This assumption causes no loss of generality, thanks to an invariance property of the central subspace (Cook 1998, p. 106).

In model (1), the likelihood of one random observation (\mathbf{x}, Y) is

$$\eta_1(x)\eta_2(Y, x^T\beta),$$

where η_1 is a probability mass function (pmf) of \mathbf{x} or a pdf of \mathbf{x} , or a mixture, depending on whether \mathbf{x} contains discrete variables, and η_2 is the conditional pmf/pdf of Y on \mathbf{x} . Treating η_1, η_2 as infinite-dimensional nuisance parameters while β as the parameter of interest, this can be viewed as a semiparametric estimation problem. The essential idea in semiparametrics is to construct estimators through deriving influence functions. Influence functions can be viewed as normalized elements in a so called nuisance tangent space orthogonal complement Λ^\perp . Thus, if one can successfully derive Λ^\perp , one at least has the hope of characterizing the influence functions and constructing estimators. Because semiparametrics is not a familiar tool in the dimension-reduction community, we give a general and more precise explanation in Appendix 0. It is not a simple and straightforward tool to grasp and to master, and interested readers are encouraged to refer to Bickel et al. (1993) and Tsiatis (2006) for in-depth understanding. In model (1), using the semiparametric analysis, we characterize the space of the span of all possible score functions, had η_1, η_2 been replaced by all possible parametric submodels, to find the nuisance tangent space. We then derive its orthogonal complement to obtain

$$\begin{aligned} \Lambda^\perp &= \{f(Y, x) - E\{f(Y, x)|x^T\beta, Y\} : E\{f(Y, x)|x\} \\ &= E\{f(Y, x)|x^T\beta\}, \forall f(Y, x)\} \end{aligned}$$

See the detailed derivation in Appendix 1. The form of Λ^\perp permits many possibilities for constructing consistent estimating equations. For example, for any functions $\mathbf{g}(Y, \mathbf{x}^T\beta)$ and $\alpha(\mathbf{x})$, we can choose $\mathbf{f}(Y, \mathbf{x})$ to be

$$[g(Y, x^T\beta) - E\{g(Y, x^T\beta)|x^T\beta\}][\alpha(x) - E\{\alpha(x)|x^T\beta\}].$$

Here, $\mathbf{f}(Y, \mathbf{x})$ satisfies $E\{\mathbf{f}(Y, \mathbf{x}) | \mathbf{x}^T\beta, Y\} = \mathbf{0}$ and is thus a valid element in Λ^\perp . Therefore, a general class of estimating equations can be obtained using the sample version of

$$E([g(Y, x^T\beta) - E\{g(Y, x^T\beta)|x^T\beta\}] \times [\alpha(x) - E\{\alpha(x)|x^T\beta\}])=0. \quad (2)$$

The resulting estimate is obviously \sqrt{n} -consistent (Newey 1990).

Remark 1. Equation (2) is only one convenient way to construct elements in Λ^\perp . Other constructions are also possible. For example, an arbitrary linear combination

$$\sum_{i=1}^k [g_i(Y, x^T \beta) - E\{g_i(Y, x^T \beta) | x^T \beta\}] \times [\alpha_i(x) - E\{\alpha_i(x) | x^T \beta\}]$$

also provides a \sqrt{n} -consistent estimator because it is a valid element in Λ^\perp .

Remark 2. It is easy to see that solving (2) does not necessarily yield a unique solution. Theoretically, as long as we choose \mathbf{g} and α so that the matrix \mathbf{A} in Theorem 1 has rank $p(p - d)$ and \mathbf{B} is bounded, solving (2) can yield a basis of $S_{Y|x}$, although the basis may not be unique; and as long as the dimension-reduction problem is identifiable, such \mathbf{g} and α always exist. Regarding the issue of multiple solutions in practice, there are two aspects to it. First, in a typical estimating equation approach, multiple roots issue presents a challenge. This problem almost always exists in a finite sample, even under the condition that at the population level a unique solution exists. There is no established method to handle it as far as we know. Almost in all the situations, empirical methods are used to select the most sensible root among several roots. Second, uniquely in the context of dimension reduction, even if the targeted central space $S_{Y|x}$ is unique, its basis—which is what we solve for—is not. Fortunately, this level of multiple roots issue is not a concern. As any particular choice of the basis will yield the same space, and the space is what we really aim for.

In Appendix 2, we show that (2) has a double robustness property, in that the consistency is doubly assured by the term $\mathbf{g} - E\{\mathbf{g} | \mathbf{x}^T \beta\}$ and the term $\alpha - E\{\alpha | \mathbf{x}^T \beta\}$. We can misspecify either $E\{\mathbf{g}(Y, \mathbf{x}^T \beta) | \mathbf{x}^T \beta\}$ or $E\{\alpha(\mathbf{x}) | \mathbf{x}^T \beta\}$, the estimator obtained from (2) will still be consistent. Specifically, if we replace $E\{\alpha(\mathbf{x}) | \mathbf{x}^T \beta\}$ with an arbitrary function $\mathbf{h}(\mathbf{x}^T \beta)$, then (2) becomes

$$E\{[g(Y, x^T \beta) - E\{g(Y, x^T \beta) | x^T \beta\}]\{\alpha(x) - h(x^T \beta)\}\} = 0,$$

which still yields a consistent estimator. Similarly, if we replace $E\{\mathbf{g}(Y, \mathbf{x}^T \beta) | \mathbf{x}^T \beta\}$ with an arbitrary function $\mathbf{h}(\mathbf{x}^T \beta)$, then (2) becomes

$$E\{(\{g(Y, x^T \beta) - h(x^T \beta)\}[\alpha(x) - E\{\alpha(x) | x^T \beta\}])\} = 0, \quad (3)$$

which also yields a consistent estimator.

Remark 3. To ensure the consistency of the estimation of β , only one of the two expectations $E\{\mathbf{g}(Y, \mathbf{x}^T \beta) | \mathbf{x}^T \beta\}$ and $E\{\alpha(\mathbf{x}) | \mathbf{x}^T \beta\}$ can be misspecified. The other expectation needs to be calculated consistently, this typically requires specifying a correct parametric model or performing nonparametric regression.

Remark 4. In contrast to the existing literature on sufficient dimension reduction (Cook 1998), when using (2) or its misspecified versions to identify $S_{Y|x}$, no additional assumptions are made on the covariate vector. This means that (i) we do not need to assume a specific joint distribution for \mathbf{x} ; (ii) we do not need to assume the linearity condition or the constant variance condition; and (iii) we do not need to assume \mathbf{x} to be continuously distributed. Although Cook and Li (2005) also considered noncontinuous covariates, they had to assume a parametric model for the distribution of the covariate vector \mathbf{x} conditional on the response Y .

Remark 5. If we are willing to make additional parametric assumptions on $E\{\alpha(\mathbf{x}) \mid \mathbf{x}^T\beta\}$ in (3), for example, $E\{\alpha(\mathbf{x}) \mid \mathbf{x}^T\beta\} = \mathbf{C}\mathbf{x}$ or $E\{\alpha(\mathbf{x}) \mid \mathbf{x}^T\beta\} = \mathbf{C}$ for a quantity \mathbf{C} that may or may not depend on β , then we will no longer need to perform a nonparametric estimation of $E\{\alpha(\mathbf{x}) \mid \mathbf{x}^T\beta\}$ when using (3). Such assumptions greatly simplify the computation. We suspect this is the implicit motivation behind the linearity condition and the constant variance condition, which are widely used in the sufficient dimension reduction literature. We will explore this issue in detail in the next two sections.

The double robustness property further allows us to obtain a \sqrt{n} -consistent estimator without any undersmoothing requirement even when $d \geq 3$ through nonparametrically estimating both $E\{\mathbf{g}(Y, \mathbf{x}^T\beta) \mid \mathbf{x}^T\beta\}$ and $E\{\alpha(\mathbf{x}) \mid \mathbf{x}^T\beta\}$. We state this result in Theorem 1 and provide the proof in the online supplementary document.

Theorem 1. Under conditions (C1)–(C4) given in Appendix 4, the estimator $\widehat{\beta}$ obtained from the estimating equation

$$\sum_{i=1}^n [g(Y_i, x_i^T \widehat{\beta}) - \widehat{E}\{g(Y_i, x_i^T \widehat{\beta}) \mid x_i^T \widehat{\beta}\}] \times [\alpha(x_i) - \widehat{E}\{\alpha(x_i) \mid x_i^T \widehat{\beta}\}] = 0$$

satisfies

$$\sqrt{n} \text{Avec}(\widehat{\beta} - \beta) \rightarrow N(0, B)$$

in distribution, where

$$A = E[\partial \text{vec}(\{g(Y, x^T \beta) - E\{g(Y, x^T \beta) \mid x^T \beta\}\}) \times \{\alpha(x) - E\{\alpha(x) \mid x^T \beta\}\} \mid \partial \{\text{vec}(\beta)\}^T],$$

$$B = \text{cov}\{\text{vec}(\{g(Y, x^T \beta) - E\{g(Y, x^T \beta) \mid x^T \beta\}\}) \times \{\alpha(x) - E\{\alpha(x) \mid x^T \beta\}\}\}.$$

Here $\text{vec}(\mathbf{M})$ denotes the vector formed by concatenating the columns of \mathbf{M} .

3. CONNECTION WITH EXISTING METHODS

In this section, we will examine several popular existing sufficient dimension-reduction methods, and illustrate why they are special cases of the semiparametric estimation family. We will show that all these methods take advantage of the double robustness property. In addition, we will point out that the linearity condition and/or the constant variance condition are used in these methods to simplify the computation. To be specific, the linearity condition characterizes the mean of \mathbf{x} conditional on $\mathbf{x}^T\beta$ by assuming

$$E(x \mid x^T \beta) = \mathbf{P}x, \quad (4)$$

and the constant variance condition characterizes the variance-covariance matrix of \mathbf{x} conditional on $(\mathbf{x}^T\beta)$ by assuming

$$\text{cov}(x \mid x^T \beta) = \mathbf{Q}, \quad (5)$$

where $\mathbf{P} = \beta(\beta^T\beta)^{-1}\beta^T = \beta\beta^T$, $\mathbf{Q} = \mathbf{I}_p - \mathbf{P}$. Note that here the two conditions are given in the context where $E(\mathbf{x}) = \mathbf{0}$ and $\text{cov}(\mathbf{x}) = \mathbf{I}_p$. Here, both \mathbf{P} and \mathbf{Q} are symmetric matrices.

Before presenting the specific analysis on these methods, We first highlight two simple linear algebra results that will be used frequently in the remaining context. These are simple linear algebra results; hence, we only sketch the proofs in the online supplementary document.

Lemma 1. Assume Λ is a $p \times p$ symmetric matrix of rank d . If and only if β satisfies

$$\Lambda - P\Lambda P=0,$$

then the span of the columns in β is the eigenspace of Λ corresponding to the d nonzero eigenvalues.

Lemma 2. Assume Λ is a $p \times p$ symmetric nonnegative definite matrix of rank d . If and only if β satisfies

$$Q\Lambda Q=0,$$

then the span of the columns in β is the eigenspace of Λ corresponding to the d nonzero eigenvalues.

3.1 Sliced Inverse Regression

The classic SIR (Li 1991) requires \mathbf{x} to satisfy the linearity condition (4). It uses the eigenvectors associated with the d nonzero eigenvalues of the matrix $\Lambda_{\text{SIR}} \stackrel{\text{def}}{=} \text{cov}\{E(x|Y)\}$ to span $\mathcal{S}_{Y|\mathbf{x}}$. For ease of illustration, we assume that Λ_{SIR} has rank d hence excluding some degenerated cases.

To obtain SIR as a semiparametric estimator, we set $\mathbf{g}(Y, \mathbf{x}^T\beta) = E(\mathbf{x} | Y)$ and $\alpha(\mathbf{x}) = \mathbf{x}^T$ in (2). The linearity condition promises a parametric form $E\{\alpha(\mathbf{x}) | \mathbf{x}^T\beta\} = \mathbf{x}^T\mathbf{P}$, hence we can use the misspecified version (3) while selecting $\mathbf{h}(\mathbf{x}^T\beta) = \mathbf{0}$. This choice of \mathbf{g} , α , and \mathbf{h} in (3) yields $E\{E(\mathbf{x} | Y)\mathbf{x}^T\}(\mathbf{I}_p - \mathbf{P}) = \mathbf{0}$, or equivalently, $\Lambda_{\text{SIR}}\mathbf{Q} = \mathbf{0}$. Because Λ_{SIR} is nonnegative definite and has rank d , hence $\Lambda_{\text{SIR}}\mathbf{Q} = \mathbf{0}$ is equivalent to $\mathbf{Q}\Lambda_{\text{SIR}}\mathbf{Q} = \mathbf{0}$. Lemma 2 indicates that this is equivalent to obtaining the eigenspace of Λ_{SIR} to span $\mathcal{S}_{Y|\mathbf{x}}$.

The above semiparametric derivation of SIR indicates clearly that the linearity condition is not structurally necessary for constructing SIR. When the linearity condition (4) does not hold, we simply lose the convenience of replacing $E\{\alpha(\mathbf{x}) | \mathbf{x}^T\beta\}$ by $\mathbf{x}^T\mathbf{P}$, everything else remains unchanged. Specifically, in this case, SIR becomes

$$\text{SIR: } E[\{E(x|Y) - E\{E(x|Y)|x^T\beta\}\}\{x - E(x|x^T\beta)\}^T]=0,$$

where $E(\cdot | \mathbf{x}^T\beta)$ and $E(\cdot | Y)$ need to be estimated nonparametrically. This is what we propose as the semiparametric generalization of SIR in the absence of the linearity condition.

3.2 Sliced Average Variance Estimation

The SAVE (Cook and Weisberg 1991) assumes both the linearity condition (4) and the constant variance condition (5). Similar to SIR, SAVE uses the eigenvectors associated with

the d nonzero eigenvalues of a matrix Λ_{SAVE} to span $S_{Y|\mathbf{x}}$, where

$$\Lambda_{\text{SIR}} \stackrel{\text{def}}{=} E[(I_p - \text{cov}(x|Y))^2].$$

To obtain SAVE from the semiparametric approach, we define $\mathbf{g}_1(Y, \mathbf{x}^T\beta) = \mathbf{I}_p - \text{cov}(\mathbf{x} | Y)$, $\mathbf{g}_2(Y, \mathbf{x}^T\beta) = \mathbf{g}_1(Y, \mathbf{x}^T\beta)E(\mathbf{x} | Y)$ and $\alpha_1(\mathbf{x}) = -\mathbf{x}\{\mathbf{x} - E(\mathbf{x} | \mathbf{x}^T\beta)\}^T$, $\alpha_2(\mathbf{x}) = \mathbf{x}^T$. As we have

pointed out in Remark 1, $\sum_{i=1}^2 \{g_i(Y, \mathbf{x}^T\beta) - E(g_i|x^T\beta)\}\{\alpha_i(x) - E(\alpha_i|x^T\beta)\}$ is an element in Λ^\perp , hence it yields a valid semiparametric estimating equation. In this construction, taking advantage of the double robustness, we are allowed to misspecify $E(\mathbf{g}_1 | \mathbf{x}^T\beta) = \mathbf{0}$ and $E(\mathbf{g}_2 | \mathbf{x}^T\beta) = \mathbf{0}$. Some algebra then yields

$$\text{SAVE: } E\{(I_p - \text{cov}(x|Y))\{x - E(x|Y)\}\{x - E(x|x^T\beta)\}^T - \text{cov}(x|x^T\beta)\} = 0. \quad (6)$$

The linearity condition (4) and the constant variance condition (5) further allow us to replace $E(\mathbf{x} | \mathbf{x}^T\beta)$ and $\text{cov}(\mathbf{x} | \mathbf{x}^T\beta)$ with $\mathbf{P}\mathbf{x}$ and \mathbf{Q} , which directly simplifies (6) to

$$\Lambda_{\text{SAVE}}\mathbf{Q} = 0.$$

Because Λ_{SAVE} is nonnegative definite, solving $\Lambda_{\text{SAVE}}\mathbf{Q} = \mathbf{0}$ is equivalent to solving $\mathbf{Q}\Lambda_{\text{SAVE}}\mathbf{Q} = \mathbf{0}$, which is equivalent to SAVE because of Lemma 2.

Relaxing the linearity condition and the constant variance condition is now obvious. Because (6) is obtained without using these two conditions, we can simply use (6) as the condition-free semiparametric generalization of SAVE, while in implementation, we need to estimate $E(\mathbf{x} | \mathbf{x}^T\beta)$ and $\text{cov}(\mathbf{x} | \mathbf{x}^T\beta)$ nonparametrically.

3.3 Directional Regression

Like SAVE, the DR (Li and Wang 2007) also assumes both the linearity condition (4) and the constant variance condition (5). It uses the eigenvectors associated with the d nonzero eigenvalues of the matrix Λ_{DR} to span $S_{Y|\mathbf{x}}$. Here,

$$\Lambda_{\text{DR}} \stackrel{\text{def}}{=} E[\{2I_p - A(Y, \tilde{Y})\}^2], A(Y, \tilde{Y}) = E\{(x - \tilde{x})(x - \tilde{x})^T | Y, \tilde{Y}\}, \text{ and } (\tilde{x}, \tilde{Y}) \text{ is an independent copy of } (\mathbf{x}, Y).$$

To obtain DR from the semiparametric approach, we choose $\mathbf{g}_1(Y, \mathbf{x}^T\beta) = \mathbf{I}_p - E(\mathbf{x}\mathbf{x}^T | Y)$, $\mathbf{g}_2(Y, \mathbf{x}^T\beta) = E\{E(\mathbf{x} | Y)E(\mathbf{x}^T | Y)\}E(\mathbf{x} | Y)$, $\mathbf{g}_3(Y, \mathbf{x}^T\beta) = E\{E(\mathbf{x}^T | Y)E(\mathbf{x} | Y)\}E(\mathbf{x} | Y)$, $\alpha_1(\mathbf{x}) = -\mathbf{x}\{\mathbf{x} - E(\mathbf{x} | \mathbf{x}^T\beta)\}^T$ and $\alpha_2(\mathbf{x}) = \alpha_3(\mathbf{x}) = \mathbf{x}^T$. Remark 1 indicates that

$\sum_{j=1}^3 \{g_j(Y, \mathbf{x}^T\beta) - E(g_j|x^T\beta)\}\{\alpha_j(x) - E(\alpha_j|x^T\beta)\}$ is an element in Λ^\perp . Taking advantage of the double robustness property, we misspecify $E(\mathbf{g}_j | \mathbf{x}^T\beta) = \mathbf{0}$, for $j = 1, 2, 3$. The subsequent estimating equation is therefore the sample version of

DR:

$$E[\{I_p - E(\mathbf{x}\mathbf{x}^T | Y)\}\{-\mathbf{x}\mathbf{x}^T + \mathbf{x}E(\mathbf{x}^T | \mathbf{x}^T\beta) + \text{cov}(\mathbf{x} | \mathbf{x}^T\beta)\}] + E[E\{E(\mathbf{x} | Y)E(\mathbf{x}^T | Y)\}E(\mathbf{x} | Y)\{x^T - E(x^T | \mathbf{x}^T\beta)\}] + E[E\{E(\mathbf{x}^T | Y)E(\mathbf{x} | Y)\}E(\mathbf{x} | Y)\{x^T - E(x^T | \mathbf{x}^T\beta)\}]$$

When both the linearity condition (4) and the constant variance condition (5) hold, we can insert $E(\mathbf{x} | \mathbf{x}^T\beta) = \mathbf{P}\mathbf{x}$ and $\text{cov}(\mathbf{x} | \mathbf{x}^T\beta) = \mathbf{Q}$ in (7). Some algebra then leads to the equivalence between (7) and

$$0 = E[\{I_p - E(\mathbf{x}\mathbf{x}^T | Y)\}^2]\mathbf{Q} + [E\{E(\mathbf{x} | Y)E(\mathbf{x}^T | Y)\}]^2\mathbf{Q} + E\{E(\mathbf{x}^T | Y)E(\mathbf{x} | Y)\}E\{E(\mathbf{x} | Y)E(\mathbf{x}^T | Y)\}\mathbf{Q} \\ = 2^{-1}\Lambda_{\text{DR}}\mathbf{Q},$$

where the last equality is due to Li and Wang (2007). Because Λ_{DR} is nonnegative definite, solving $\Lambda_{DR}\mathbf{Q} = \mathbf{0}$ is equivalent to solving $\mathbf{Q}\Lambda_{DR}\mathbf{Q} = \mathbf{0}$, which is equivalent to DR because of Lemma 2.

Similar to SAVE, as (7) is obtained without any of the linearity or constant variance condition, it can thus be used as a semiparametric generalization of DR.

Li and Dong (2009) and Dong and Li (2010) extended SIR, SAVE, and DR to the case when the linearity condition (4) is violated while the constant variance condition (5) is true. Similar analysis shows that these are also special cases of the semiparametric approach. These results are available in the supplementary document.

4. SEMIPARAMETRIC ESTIMATION OF THE CENTRAL MEAN SUBSPACE

In situations when one only concerns about the conditional mean of the response given the predictors, Cook and Li (2002) introduced the notion of the central mean subspace. They defined the column space of β as a mean dimension-reduction subspace if β satisfies

$$E(Y|x) = E(Y|x^T\beta).$$

The intersection of all mean dimension-reduction subspaces is defined as the central mean subspace, denoted by $S_{E(Y|x)}$, if the intersection itself is also a mean dimension-reduction subspace. The conditional mean model assumes the mean of Y conditional on \mathbf{x} relies on $\mathbf{x}^T\beta$ only. In other words, \mathbf{x} contributes to the conditional mean of Y only through $\mathbf{x}^T\beta$. Our main interest is to estimate $S_{E(Y|x)}$, or equivalently, a basis matrix β which spans $S_{E(Y|x)}$.

To facilitate the semiparametric analysis, we write the conditional mean model as

$$Y = \ell(x^T\beta) + \varepsilon, \quad (8)$$

where $\ell(x^T\beta) \stackrel{\text{def}}{=} E(Y|x^T\beta)$ is an unspecified smooth function and $E(\varepsilon | \mathbf{x}) = 0$. We emphasize that because we make no assumptions on ε other than conditional mean zero, (8) is equivalent to the central mean subspace model. For the conditional mean model (8), the likelihood of one random observation (\mathbf{x}, Y) is

$$\eta_1(x)\eta_2\{Y - \ell(x^T\beta), x\},$$

where η_1 has the same meaning as in Section 2, $\ell(x^T\beta)$ is the mean function of Y conditional on \mathbf{x} (or equivalently, on $\mathbf{x}^T\beta$), and η_2 is the conditional pmf/pdf of the residual $\varepsilon = Y - E(Y | \mathbf{x}^T\beta)$ on \mathbf{x} . Here, η_2 satisfies $E(\varepsilon | \mathbf{x}) = 0$, and is otherwise unconstrained. Similarly, treating η_1 , η_2 , and ℓ as nuisance parameters while β as the parameter of interest, following the semiparametric analysis in Appendix 3, we obtain the nuisance tangent space orthogonal complement to be the class of the form

$$\Lambda^\perp = \{ \{Y - E(Y|x^T\beta)\} \{ \alpha(x) - E(\alpha|x^T\beta) \} : \forall \alpha \}.$$

Similar to (1), the form of Λ^\perp allows us to take any $\alpha(\mathbf{x})$ to obtain

$$\{Y - E(Y|x^T\beta)\} \{ \alpha(x) - E(\alpha(x)|x^T\beta) \}$$

as a valid element in Λ^\perp . Hence, a general class of estimating equations can be obtained using the sample version of

$$E(\{Y - E(Y|x^T\beta)\}[\alpha(x) - E\{\alpha(x)|x^T\beta\}])=0. \quad (9)$$

Similar to Appendix 2, we can easily show that (9) has a double robustness property, in that we can misspecify either $E(Y | \mathbf{x}^T\beta)$ or $E\{\alpha(\mathbf{x}) | \mathbf{x}^T\beta\}$, the resulting estimator from (9) will still yield a consistent estimating equation.

4.1 Ordinary Least Squares

We first inspect the OLS (Li and Duan 1989) method, where the linearity condition (4) is assumed to hold. The OLS method uses $\text{cov}(\mathbf{x}, Y)$ to infer a subspace of the column space of β in (8).

From the semiparametric approach, we let $\alpha(\mathbf{x}) = \mathbf{x}$ in (9). Taking advantage of the double robustness, we misspecify $E(Y | \mathbf{x}^T\beta) = 0$. Then (9) reduces to

$$0=E(xY) - PE(xY)=E(xY) - \beta(\beta^T\beta)^{-1}\beta^TE(xY).$$

Note that $(\beta^T\beta)^{-1}\beta^TE(\mathbf{x}Y)$ is a $d \times 1$ vector. This directly yields $\text{cov}(\mathbf{x}, Y)$ as a one-dimensional subspace of $S_{E(Y|\mathbf{x})}$, which is exactly the OLS estimation.

When the linearity condition (4) does not hold, (9) has the form

$$E[\{Y - E(Y|x^T\beta)\}\{x - E(x|x^T\beta)\}]=0$$

and can still be used to estimate β . A simple treatment is to set $E(\mathbf{x} | \mathbf{x}^T\beta) = 0$ and solve the sample version of the above display to obtain β .

4.2 Principal Hessian Directions

The principal Hessian directions (PHD) method (Li 1992) assumes both the linearity condition (4) and the constant variance condition (5). It uses the eigenvectors associated with d nonzero eigenvalues of Λ_{PHD} to form a basis of $S_{E(Y|\mathbf{x})}$. Here,

$\Lambda_{\text{PHD}} \stackrel{\text{def}}{=} E[\{Y - E(Y)\}xx^T]$. To obtain PHD from the semiparametric approach, we let $\alpha(\mathbf{x}) = \mathbf{xx}^T$ in (9). Taking advantage of the double robustness, we misspecify $E(Y | \mathbf{x}^T\beta) = E(Y)$. The linearity condition (4) and the constant variance condition (5) yield a simplification $\alpha(\mathbf{x}) - E\{\alpha(\mathbf{x}) | \mathbf{x}^T\beta\} = \mathbf{xx}^T - \mathbf{Q} - \mathbf{Pxx}^T\mathbf{P}$, hence (9) reduces to

$$E[\{Y - E(Y)\}(xx^T - \mathbf{Pxx}^T\mathbf{P})]=\Lambda_{\text{PHD}} - \mathbf{P}\Lambda_{\text{PHD}}\mathbf{P}=0.$$

Lemma 1 indicates that this is equivalent to the PHD method. When either (4) or (5) is not true, we can use

$$E[\{Y - E(Y|x^T\beta)\}\{xx^T - E(xx^T|x^T\beta)\}]=0$$

to estimate β , where we calculate $E(Y | \mathbf{x}^T\beta)$ and $E(\mathbf{xx}^T | \mathbf{x}^T\beta)$ nonparametrically. We can opt to misspecify $E(Y | \mathbf{x}^T\beta) = 0$ or preferably $E(\mathbf{xx}^T | \mathbf{x}^T\beta) = \mathbf{0}$ to simplify the computation. The second simplification

$$\text{PHD}: E[\{Y - E(Y|x^T\beta)\}xx^T] = 0$$

will be considered as the semiparametric generalization of the PHD method without extra conditions.

5. IMPLEMENTATION

To focus on delivering the main message, we have avoided detailing the implementation details in practice, which we explain now.

The proposed semiparametric counterparts of SIR, SAVE, DR, PHD, and so on, which were respectively denoted by semi-SIR, semi-SAVE, semi-DR, and semi-PHD for ease of subsequent illustration, all have the similar components of estimating conditional expectation and solving estimating equations. Because the number of the estimating equations sometimes is larger than the number of parameters, the implementation is through minimizing their Frobenius norm. We use the familiar Newton-Raphson procedure to numerically obtain the minimizer. This essentially means that at the j th iteration $\beta^{(j)}$, we perform nonparametric conditional estimation using, say, a kernel regression method to evaluate an estimating equation at $\beta^{(j)}$, and use numerical difference to obtain the derivative of the estimating equation with respect to β evaluated at $\beta^{(j)}$. We then update $\beta^{(j)}$. This process iterates until the difference between two consecutive candidates are sufficiently small.

Because semi-DR has the most complicated form, we now outline an algorithm for semi-DR as a concrete illustration. We use $\beta^{(j)}$ to denote the value of $\widehat{\beta}$ at the j th iteration, and use $K(\cdot)$ to denote a kernel function. Let $K_h(\cdot) = K(\cdot/h)/h$ for any bandwidth h . We use the Epanechnikov kernel function in the implementation. Assume the observations are (\mathbf{x}_i, Y_i) for $i = 1, \dots, n$.

1. Nonparametrically estimate $E(\mathbf{x}\mathbf{x}^T | Y)$ and $E(\mathbf{x} | Y)$ using

$$\widehat{E}(xx^T | Y) = \frac{\sum_{i=1}^n x_i x_i^T K_h(Y - Y_i)}{\sum_{i=1}^n K_h(Y - Y_i)} \text{ and}$$

$$\widehat{E}(x | Y) = \frac{\sum_{i=1}^n x_i K_h(Y - Y_i)}{\sum_{i=1}^n K_h(Y - Y_i)}.$$

The bandwidth can be selected using the classic cross-validation procedure. We evaluate the above estimations at Y_i for $i = 1, \dots, n$.

2. Form $A_i = I_p - \widehat{E}(xx^T | Y_i)$, $b_i = n^{-1} \sum_{j=1}^n \{\widehat{E}(x | Y_j) \widehat{E}(x^T | Y_j)\} \widehat{E}(x | Y_i)$ and $c_i = n^{-1} \sum_{j=1}^n \{\widehat{E}(x^T | Y_j) \widehat{E}(x | Y_j)\} \widehat{E}(x | Y_i)$.
3. Pick an arbitrary starting value $\beta^{(1)}$, for example, the classical DR estimate.
4. At the j th iteration, form $t_i = x_i^T \beta^{(j)}$ for $i = 1, \dots, n$.
5. Nonparametrically estimate $E(\mathbf{x} | \mathbf{x}^T \beta^{(j)})$ and $\text{cov}(\mathbf{x} | \mathbf{x}^T \beta^{(j)})$ using

$$\widehat{E}(x|x^T\beta^{(j)}) = \frac{\sum_{i=1}^n x_i K_h(x^T\beta^{(j)} - t_i)}{\sum_{i=1}^n K_h(x^T\beta^{(j)} - t_i)}$$

and

$$\widehat{\text{cov}}(x|x^T\beta^{(j)}) = \frac{\sum_{i=1}^n x_i x_i^T K_h(x^T\beta^{(j)} - t_i)}{\sum_{i=1}^n K_h(x^T\beta^{(j)} - t_i) - \widehat{E}(x|x^T\beta^{(j)})\widehat{E}(x^T|x^T\beta^{(j)})}$$

We use a same bandwidth to estimate the above two quantities to ensure that $\widehat{\text{cov}}(x|x^T\beta^{(j)})$ is positive definite. The bandwidth can be selected using the cross-validation procedure. Evaluate the above estimations at $\mathbf{x}^T\beta^{(j)} = \mathbf{t}_1, \dots, \mathbf{t}_n$.

6. Form the sample version of the left-hand side of (7):

$$r(\beta^{(j)}) = \frac{1}{n} \sum_{i=1}^n [A_i \{-x_i x_i^T + x_i \widehat{E}(x^T|t_i) + \widehat{\text{cov}}(x|t_i)\} + (b_i + c_i) \{x_i - \widehat{E}(x_i|t_i)\}^T].$$

Update $\beta^{(j)}$ using the following Newton-Raphson steps.

7. Form the derivative of $\{\|\mathcal{r}(\beta)\|^2\} / \{\text{vec}(\beta)\}$ evaluated at $\beta^{(j)}$ through numerical difference. Specifically, let δ be a small number and decide an order for the elements in $\beta^{(j)}$. Let $\beta_{k+}^{(j)} = \beta^{(j)} + \delta e_k$ and $\beta_{k-}^{(j)} = \beta^{(j)} - \delta e_k$. Here e_k has the same size as β but has 1 in the k th entry and zero elsewhere. Repeat the procedure in Step 5 to obtain $r(\beta_{k+}^{(j)})$ and $r(\beta_{k-}^{(j)})$. Set the k th row of $\{\|\mathcal{r}(\beta)\|^2\} / \{\text{vec}(\beta)\}$ to be $\{\|r(\beta_{k+}^{(j)})\|^2 - \|r(\beta_{k-}^{(j)})\|^2\} / (2\delta)$. Repeat this for all the entries in $\beta^{(j)}$. For simplicity, we denote the resulting first derivative $\{\|\mathcal{r}(\beta^{(j)})\|^2\} / \{\text{vec}(\beta)\}$. Following similar procedures, we obtain the second derivative ${}^2\{\|\mathcal{r}(\beta)\|^2\} / \{\text{vec}(\beta)\} \{\text{vec}(\beta)\}^T$.
Practically, $\delta = 0.001$ is sufficiently small to obtain a close approximation of the derivatives. If more precision is desired, one can simply opt for smaller δ .
8. Update to obtain

$$\text{vec}(\beta^{(j+1)}) = \text{vec}(\beta^{(j)}) - \left[\frac{\partial^2 \{\|\mathcal{r}(\beta^{(j)})\|^2\}}{\partial \{\text{vec}(\beta^{(j)})\} \partial \{\text{vec}(\beta^{(j)})\}^T} \right]^{-1} \times \frac{\partial \{\|\mathcal{r}(\beta^{(j)})\|^2\}}{\partial \{\text{vec}(\beta^{(j)})\}}$$

9. Repeat Steps 4 to 8 until convergence.

For identification, \mathbf{x} must contain at least one continuous variable (Ichimura 1993; Horowitz and Hardle 1996). The above algorithm is applicable if the support of $\mathbf{x}^T\beta$ is connected. When the support consists of several disjoint regions, we need to carry out the kernel estimation in Step 5 in each region.

Another important component in dimension reduction is to decide the dimension d of $S_{Y|\mathbf{x}}$ or $S_{E|Y|\mathbf{x}}$. To this end, the bootstrap procedure described by Dong and Li (2010) adapts the idea of Ye and Weiss (2003) by taking into account the variation of the covariates, and is

very flexible. This is what we recommend for use in our context. Specifically, we can decide d through the following procedure. Let $\lambda_1, \dots, \lambda_k$ be the nonzero eigenvalues of

$$\{\text{var}(u)\}^{-1/2} \text{cov}(u, v^T) \{\text{var}(v)\}^{-1} \text{cov}(u, v^T) \{\text{var}(u)\}^{-1/2},$$

where \mathbf{u} and \mathbf{v} are two generic random vectors. In addition, we let

$$r^2(u, v) = k^{-1} \sum_{i=1}^k \lambda_i.$$

For any working dimension $k = 1, \dots, p - 1$, we let $\widehat{\beta}_k$ be the estimate based on the original sample, and $\widehat{\beta}_{k,b}$ be the estimate based on the b th bootstrap sample, for $b = 1, \dots, B$. We then estimate the structural dimension d by maximizing over $k = 1, \dots, p - 1$

$$\bar{r}_k^2 = \frac{1}{B} \sum_{b=1}^B r^2(\widehat{\beta}_k^T x, \widehat{\beta}_{k,b}^T x). \quad (10)$$

Interested readers are referred to Dong and Li (2010) for more details.

6. SIMULATION STUDY

In this section, we conduct simulation studies to evaluate the performance of different estimation procedures. Unless otherwise stated, we repeat the experiments 500 times each with sample size $n = 200$ and reduced space dimension $d = 2$. Throughout the simulations, we used the Epanechnikov kernel and fixed $\widehat{\sigma}(3n/4)^{-1/(d+4)}$ as bandwidth, where $\widehat{\sigma}$ is the robust estimation of the standard deviation of $\mathbf{x}^T \beta$, which is the default bandwidth selector implemented in Matlab routine **ksdensity**. We choose the predictor dimension p to be 6 and 12, and consider the following two cases for the covariate vector $\mathbf{x} = (X_1, \dots, X_p)^T$.

Case 1: We generate $(X_1, X_2)^T$ (corresponding to the case $p = 6$) and $X_1, X_2, X_7, \dots, X_p)^T$ (corresponding to the case $p = 12$) from normal population with mean zero and variance-covariance matrix $(\sigma_{ij})_{(p-4) \times (p-4)}$ where $\sigma_{ij} = 0.5^{|i-j|}$. We generate X_3 and X_4 from nonlinear models: $X_3 = |X_1 + X_2| + |X_1| \epsilon_1$, and $X_4 = |X_1 + X_2|^2 + |X_2| \epsilon_2$, where ϵ_j 's are independently generated from the standard normal population, X_5 from a Bernoulli distribution with success probability $\exp(X_2) / \{1 + \exp(X_2)\}$, and X_6 from another Bernoulli distribution with success probability $\Phi(X_2)$, where $\Phi(\cdot)$ denotes the cumulative distribution function of standard normal population. Note that in this case, both the linearity condition (4) and the constant variance condition (5) are violated.

Case 2: We generate \mathbf{x} from normal population with mean zero and variance-covariance matrix $(\sigma_{ij})_{p \times p}$ where $\sigma_{ij} = 0.5^{|i-j|}$. Note that in this case, the covariates satisfy both the linearity condition (4) and the constant variance condition (5).

Let β be a basis matrix of $S_{Y|\mathbf{x}}$ and $\widehat{\beta}$ be its estimate. To assess the estimation accuracy of $\widehat{\beta}$, we use the Euclidean distance between $\widehat{\beta}$ and β , defined as the Frobenius norm of the matrix $\widehat{\beta}(\widehat{\beta}^T \widehat{\beta})^{-1} \widehat{\beta}^T - \beta(\beta^T \beta)^{-1} \beta^T$. In both cases, the distance ranges from zero to two, and a smaller distance indicates a better estimate.

6.1 Example 1

We generate the response variable using the following four different models:

$$\begin{aligned} \text{model (I): } Y &= (x^T \beta_1) / \{0.5 + (x^T \beta_2 + 1.5)^2\} + 0.5 \varepsilon; \\ \text{model (II): } Y &= (x^T \beta_1)^2 + 2|x^T \beta_2| + 0.1|x^T \beta_2| \varepsilon; \\ \text{model (III): } Y &= \exp(x^T \beta_1) + 2(x^T \beta_2 + 1)^2 + |x^T \beta_1| \varepsilon; \\ \text{model (IV): } Y &= (x^T \beta_1)^2 + (x^T \beta_2)^2 + 0.5 \varepsilon, \end{aligned}$$

where β_1 and β_2 are $p \times 1$ vectors with their first six components being $(1, 1, 1, 1, 1, 1)^T / \sqrt{6}$ and $(1, -1, 1, -1, 1, -1)^T / \sqrt{6}$, respectively. When $p = 12$, the rest components of β_1 and β_2 are identically zero. The error term ε has a standard normal distribution. Models (I)–(IV) are chosen to compare, respectively, SIR, SAVE, DR, and PHD with their semiparametric counterparts. We also include dMAVE (Xia 2007) and MAVE (Xia et al. 2002) into our comparison as they are often used as a benchmark. To make a fair comparison on the core methodologies of these proposals, we estimate the kernel matrices of the classical SIR, SAVE, DR by using kernel smoothing rather than the usual slicing estimation. This allows us to avoid selecting the number of slices which usually adversely affects the performance. Thus, SIR, SAVE, and DR are implemented in their improved form.

The boxplots of the Euclidean distances are reported in Figure 1. The results under Case 1 are presented in panels (A) and (C), where we show the boxplots of Euclidean distances for different estimation procedures when both the linearity condition (4) and the constant variance condition (5) are violated. In this case, we can see that the semiparametric estimates are substantially more accurate than their classical dimension-reduction counterparts across all four models, indicating the significant improvement when the violation of these conditions is taken into account. Note that dMAVE has very large variability in these settings and sometimes even perform worse than the classical dimension-reduction methods. This is because some covariates are discrete which violates the continuity requirement of dMAVE. Because the semiparametric estimation procedures do not require continuous covariates, and they do not rely on any conditional moment or distributional assumptions, their performance dominates the competitors in all scenarios. The results for Case 2 are in panels (B) and (D), which contain the boxplots of Euclidean distances when both the linearity condition (4) and the constant variance condition (5) are satisfied. We can see that, surprisingly, our semiparametric proposals are still superior to their classical counterparts in these particular examples. This reminds us of a quite interesting phenomenon where sometimes, estimating a quantity even if it is known brings gain (Henmi and Eguchi 2004). However, whether this gain is real, in general, or just in these specific examples deserves further theoretical investigation. On the other hand, our estimators seem to perform comparably with dMAVE, which usually produces very accurate estimates when all the covariates are continuous, such as the case here. Following the request of a referee, we also experimented the simulations with $p = 50$, which corresponds to 96 free parameters with the sample size $n = 200$. Unfortunately, none of the methods examined above can provide any reasonable results in this case. We believe that to handle p very large in comparison with n , additional assumptions such as sparsity is needed. The sparsity assumption assumes many rows of β are zeros, hence the corresponding variables in \mathbf{x} simply have no effect on the response variable Y . Many existing methods are available to take advantage of the sparsity assumption (see, e.g., Li 2007; Wang and Wang 2010). Combining the sparsity techniques and the dimension reduction methods are promising future work.

6.2 Example 2

Next, we perform simulations to demonstrate the performance of the bootstrap procedure conjunct with the semiparametric methods in estimating the structural dimension d . We continue to use models (I)–(IV) in Example 1, and generate the predictors \mathbf{x} from case 1 when $p = 6$. We generated 100 datasets of sample size $n = 200$, with bootstrap size $B = 100$.

We report the relative frequency of the bootstrap selected dimensions for models (I)–(IV) in Table 1. It can be easily seen that with at least 95% accuracy the bootstrap method correctly chose the dimension, which we consider quite satisfactory.

7. REAL DATA APPLICATION

We illustrate further our semiparametric proposals through a dataset concerning the employees' salary in the Fifth National Bank of Springfield (Albright, Winston, and Zappe 1999). The aim of this study is to understand how an employee's salary associates with his/her personal characteristics. To this end, an employee's annual salary is the response variable Y , and six covariates are possibly associated with the salary: the employee's current job level, where a larger number indicates a higher rank (X_1); the employee's working experience at current bank, which is measured by the number of years of the employment (X_2); the employee's age (X_3); the experience of an employee at another bank prior to working at the Fifth National, which is measured by the number of years at other banks (X_4); the employee's gender (X_5); and a binary variable indicating whether the employee's job is computer related (X_6). We removed an obvious outlier in this dataset, leaving 207 observations in the subsequent analysis.

Under various working dimensions $k = 1, \dots, 5$, we implement the proposed semi-DR method given in (7) in association with the bootstrap method described in Section 5, with bootstrap size $B = 1000$. The results are presented as boxplots in Figure 2. It is clear that bootstrap method favors $d = 1$, which indicates that the six covariates affect the salary through one single linear combination direction. We present the estimation of this direction in Table 2. For comparison, we also implement the dMAVE and the classical DR with kernel smoothing. We emphasize here that dMAVE requires the covariates to be continuous, which is not the case for gender (X_5), and DR requires linearity condition and constant variance condition, which are not satisfied in this example (see Figure 3). Hence, both results should not be fully trusted.

To see the dimension-reduction effect from a different perspective, in Figure 4, we present the scatterplots of Y versus $x^T \widehat{\beta}$, where $\widehat{\beta}$ denotes the estimate obtained from Semi-DR, dMAVE, and DR. We can see that the estimates obtained from Semi-DR and dMAVE exhibit a more obvious pattern than DR in that the data cloud appears more compact. We also include in Figure 4 three curves (the dashed lines) fitted from a quadratic function on $(x_i^T \widehat{\beta}, Y_i), i=1, \dots, n$. These fit the data cloud rather well and illustrate a homoscedastic error pattern. Thus, we further used the fitted curves to perform a cross-validation procedure to calculate the prediction error. The resulting prediction errors from Semi-DR, dMAVE and DR are, respectively, 21.3191, 23.4481, and 47.0288, which further illustrate the advantage of the Semi-DR method. All these evidences support the conclusion that Semi-DR has successfully reduced the dimension from $p = 6$ to $d = 1$ and has found the right linear combination of the covariates in terms of establishing the association between salary and an employee's characteristic.

8. DISCUSSION

It is worth mentioning that (2) and (9) only differ in the first component. For (2), this component can be constructed from any function of Y and $\mathbf{x}^T\beta$, namely, any $\mathbf{g}(Y, \mathbf{x}^T\beta)$, while for (9), the only valid choice is $\mathbf{g}(Y, \mathbf{x}^T\beta) = Y$. Thus, the family of estimators provided by (2) is much richer than the one provided by (9). In fact, it includes the family of (9) as a subfamily of estimators. This is easy to understand intuitively as model (1) assumes more structures than model (8), consequently one can have more ways to construct estimators for (1) and, therefore, model (1) has a larger subspace Λ^\perp . Because of this relation, any estimator for model (8) is necessarily also an estimator for (1). This also agrees with the dimension-reduction result that the conditional mean subspace $\mathcal{S}_{E(Y|\mathbf{x})}$ is a subspace of the conditional subspace $\mathcal{S}_{Y|\mathbf{x}}$, because if (1) holds for β , then (8) also holds for β or a submatrix of β .

For illustration, we have chosen to derive SIR, SAVE, and DR as examples for the central subspace problem and OLS and PHD as examples for the central mean subspace from the semiparametric approach, mainly due to their popularity in the dimension-reduction literature. Because Λ^\perp contains all the influence functions, hence every root- n consistent method must correspond to a special choice of the functions in Λ^\perp . In this sense, the families given in Sections 2 and 4 are complete.

Finally, we acknowledge here that the linearity condition and the constant variance condition, when assumed to hold only at the true value of β , can be mild and often hold approximately, especially when p is large (Hall and Li 1993). However, our simulations in Case 1 have shown that this approximation may not be sufficient to justify making and using these assumptions. In other words, if we assume the linearity condition and/or constant variance condition to hold exactly at the true β while, in fact, they are only approximately true, the subsequent estimation of the central subspace or central mean subspace can be quite different. In addition, even if the linearity and constant variance conditions do hold exactly, our limited numerical results seem to indicate that it might be still beneficial not to use them if one is concerned more about the estimation quality than the computational cost. Investigation on whether this is indeed a general phenomenon or some special isolated instances is a worthy endeavor.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

Yanyuan Ma's work was supported by the National Science Foundation (DMS-0906341) and the National Institute of Neurological Disorders and Stroke (R01-NS073671). Liping Zhu's work was supported by the Natural Science Foundation of China (11071077).

APPENDIX: SOME TECHNICAL DERIVATIONS

A.1. General Introduction to Semiparametrics

Consider the Hilbert space H consisting of all the mean zero, finite variances, length m vector functions of \mathbf{x} , Y , where the inner product between two functions h, g is defined as $E(h^T g)$. Here and in the following definitions, all the expectations are calculated under the true distribution. The nuisance tangent space Λ is a subspace of H defined as the mean squared closure of all the elements of the form $B S$, where S is an arbitrary nuisance score vector function, and B is any conformable matrix with m rows. Here the nuisance score

vector functions are calculated conventionally in every possible valid parameterization of the infinite-dimensional nuisance parameter, where a “valid parameterization” means that there exists one parameter value which yields the truth. Furthermore, Λ^\perp is defined to be the orthogonal complement of Λ in H .

Semiparametric theory ensures that every regular, asymptotic linear, root n consistent (RAL) estimator corresponds to an influence function, and every influence function is a normalized element in Λ^\perp , where the normalization is to ensure $E(\phi S_\beta^T) = I_m$. Here, S_β is the usual score vector with respect to the parameter of interest β , and we use ϕ to denote an influence function and I_m to denote the size m identity matrix. This link between the complete family of RAL estimators and the space Λ^\perp allows one to derive estimators through characterizing Λ^\perp and identifying members in Λ^\perp . Every explicit identification of one function in Λ^\perp ensures the discovery of one estimator. This is the semiparametric approach that drives all our derivations.

A.2. The Derivation of Λ^\perp in Model (1)

Denote the nuisance tangent space corresponding to η_1 and η_2 , respectively, Λ_1 and Λ_2 . We have

$$\begin{aligned} \Lambda_1 &= \{f(x) : \forall f \text{ such that } E(f) = 0\}, \\ \Lambda_2 &= \{f(Y, \beta^T x) : \forall f \text{ such that } E(f|x) = E(f|\beta^T x) = 0\}. \end{aligned}$$

Obviously, $\Lambda_1 \perp \Lambda_2$, hence $\Lambda = \Lambda_1 \oplus \Lambda_2$. It is easy to see that $\Lambda_1^\perp = \{f(Y, x) : E(f|x) = 0\} \supseteq \Lambda_2$. We now show that $\Lambda_2^\perp = \{f(Y, x) : E(f|\beta^T x, Y) \text{ is a function of } \beta^T x \text{ only}\}$. Obviously, functions having the required conditional expectation property are certainly elements in Λ_2^\perp . To show that elements in Λ_2^\perp have to satisfy the conditional expectation requirement, consider any $f(Y, x) \in \Lambda_2^\perp$. We let $\mathbf{g} = E(\mathbf{f} | \beta^T \mathbf{x}, Y) - E(\mathbf{f} | \beta^T \mathbf{x})$. Obviously, $\mathbf{g} \in \Lambda_2$ hence $E(\mathbf{g}^T \mathbf{f}) = 0$. On the other hand,

$$\begin{aligned} 0 &= E(\mathbf{g}^T \mathbf{f}) = E\{ \mathbf{g}^T E(\mathbf{f} | \beta^T x, Y) \} \\ &= E(\mathbf{g}^T \mathbf{g}) + E\{ \mathbf{g}^T E(\mathbf{f} | \beta^T x) \} \\ &= E(\mathbf{g}^T \mathbf{g}) + E\{ E(\mathbf{g} | \beta^T x)^T E(\mathbf{f} | \beta^T x) \} \\ &= E(\mathbf{g}^T \mathbf{g}) + 0. \end{aligned}$$

Hence \mathbf{g} itself should be zero. This means \mathbf{f} indeed satisfies the conditional expectation requirement.

We now show that

$$\begin{aligned} \Lambda^\perp &= \Lambda_1^\perp \cap \Lambda_2^\perp = \{f(Y, x) - E(f|\beta^T x, Y) : E(f|x) \\ &= E(f|\beta^T x) \forall f\}. \end{aligned}$$

To see this, note that its form simply says for any element in Λ^\perp , it should satisfy the condition $E(\mathbf{g} | \beta^T \mathbf{x}, Y) = 0$. Let us denote the set $A = \{\mathbf{f}(Y, \mathbf{x}) - E(\mathbf{f} | \beta^T \mathbf{x}, Y) : E(\mathbf{f} | \mathbf{x}) = E(\mathbf{f} | \beta^T \mathbf{x}) \forall \mathbf{f}\}$. Obviously, $A \subset \Lambda_2^\perp$. Because

$$\begin{aligned}
 & E\{f(Y, x) - E(f|\beta^T x, Y)|x\} \\
 &= E\{f(Y, x)|x\} - E\{f|\beta^T x, Y\}|x\} \\
 &= E(f|x) - E\{E(f|\beta^T x, Y)|\beta^T x\} \\
 &= E(f|x) - E(f|\beta^T x) = 0,
 \end{aligned}$$

hence $A \subset \Lambda_1^\perp$ as well. Hence, $A \subset \Lambda^\perp$. On the other hand, for any $\mathbf{f} \in \Lambda^\perp$, because $f \in \Lambda_2^\perp$, we have $E(\mathbf{f} | \beta^T \mathbf{x}, Y) = \mathbf{a}(\beta^T \mathbf{x})$ for some \mathbf{a} . Writing this out, we obtain

$$\begin{aligned}
 a(\beta^T x) &= \frac{\int_{\beta^T x = \beta^T x} f(Y, x) \eta_1(x) \eta_2(Y|\beta^T x) d\mu(x)}{\int_{\beta^T x = \beta^T x} \eta_1(x) \eta_2(Y|\beta^T x) d\mu(x)} \\
 &= \frac{\int_{\beta^T x = \beta^T x} f(Y, x) \eta_1(x) d\mu(x)}{\int_{\beta^T x = \beta^T x} \eta_1(x) d\mu(x)}.
 \end{aligned}$$

Now, we have

$$\begin{aligned}
 a(\beta^T x) &= \int a(\beta^T x) \eta_2(Y|\beta^T x) d\mu(Y) \\
 &= \frac{\int_{\beta^T x = \beta^T x} f(Y, x) \eta_1(x) d\mu(x)}{\int_{\beta^T x = \beta^T x} \eta_1(x) d\mu(x)} \eta_2(Y|\beta^T x) d\mu(Y) \\
 &= \frac{\int \int_{\beta^T x = \beta^T x} f(Y, x) \eta_1(x) \eta_2(Y|\beta^T x) d\mu(x) d\mu(Y)}{\int_{\beta^T x = \beta^T x} \eta_1(x) d\mu(x)} \\
 &= \frac{\int_{\beta^T x = \beta^T x} \int f(Y, x) \eta_2(Y|\beta^T x) d\mu(Y) \eta_1(x) d\mu(x)}{\int_{\beta^T x = \beta^T x} \eta_1(x) d\mu(x)} = 0
 \end{aligned}$$

because $E(\mathbf{f} | \mathbf{x}) = 0$ due to $f \in \Lambda_1^\perp$. Thus, elements in Λ^\perp indeed have the form $\mathbf{f}(Y, \mathbf{x}) - E(\mathbf{f} | Y, \beta^T \mathbf{x})$. The requirement of these elements belonging to Λ_1^\perp generates the second requirement of A , and we obtain $\Lambda^\perp \subset A$. This completes the derivation of Λ^\perp .

A.3. The Double Robustness of (2)

Denote $E^*(\mathbf{g} | \mathbf{x}^T \beta)$ the misspecified function of $E(\mathbf{g} | \mathbf{x}^T \beta)$. We have

$$\begin{aligned}
 & E[\{g(Y, x^T \beta) - E^*(g|x^T \beta)\} \{\alpha(x) - E(\alpha|x^T \beta)\}] \\
 &= E(\{E\{g(Y, x^T \beta)|x\} - E^*(g|x^T \beta)\} \{\alpha(x) - E(\alpha|x^T \beta)\}) \\
 &= E[\{E(g|x^T \beta) - E^*(g|x^T \beta)\} \{\alpha(x) - E(\alpha|x^T \beta)\}] \\
 &= E[\{E(g|x^T \beta) - E^*(g|x^T \beta)\} E\{\alpha|x^T \beta\} - E(\alpha|x^T \beta) \{E(g|x^T \beta) - E^*(g|x^T \beta)\}] \\
 &= 0.
 \end{aligned}$$

Denote $E^*(\mathbf{a} | \mathbf{x}^T \beta)$ the misspecified function of $E(\mathbf{a} | \mathbf{x}^T \beta)$. We have

$$\begin{aligned} & E[\{g(Y, x^T\beta) - E(g|x^T\beta)\}\{\alpha(x) - E(\alpha|x^T\beta)\}] \\ &= E([E\{g(Y, x^T\beta)|x\} - E(g|x^T\beta)]\{\alpha(x) - E(\alpha|x^T\beta)\}) \\ &= E[\{E(g|x^T\beta) - E(g|x^T\beta)\}\{\alpha(x) - E(\alpha|x^T\beta)\}] = 0. \end{aligned}$$

A.4. The Derivation Λ^\perp in Model (8)

Denote the nuisance tangent space corresponding to η_1 , η_2 , and m , respectively, Λ_1 , Λ_2 , and Λ_m . Straightforward derivation yields

$$\begin{aligned} \Lambda_1 &= \{f(x): \forall f \text{ such that } F(f)=0\}, \\ \Lambda_2 &= \{f(\varepsilon, x): \forall f \text{ such that } E(f|x)=0, E(\varepsilon f|x)=0\}, \\ \Lambda_m &= \left\{ \frac{\eta'_{20,\varepsilon(\varepsilon,x)}}{\eta_{20(\varepsilon,x)}} h(x^T\beta): \forall h \right\}. \end{aligned}$$

Consequently, we have

$$\begin{aligned} \Lambda_1^\perp &= \{\alpha(\varepsilon, x): \forall \alpha \text{ such that } E(\alpha|x)=0\}, \\ (\Lambda_1 + \Lambda_2)^\perp &= \{\alpha(x)\varepsilon: \forall \alpha\}. \end{aligned}$$

To further derive $\Lambda^\perp \subset (\Lambda_1 + \Lambda_2)^\perp$, we first inspect an $\alpha(\mathbf{x})\varepsilon \in \Lambda^\perp$. For any $\mathbf{h}(\mathbf{x}^T\beta)$, $\alpha(\mathbf{x})$ satisfies

$$\begin{aligned} 0 &= \int \alpha(x)\varepsilon \frac{\eta'_{20,\varepsilon(\varepsilon,x)}}{\eta_{20(\varepsilon,x)}} h(x^T\beta)\eta_{10}(x)\eta_{20}(\varepsilon, x)d\mu(\varepsilon)d\mu(x) \\ &= \int \alpha(x)h(x^T\beta)\eta_{10}(x) \left\{ \int \varepsilon \eta'_{20\varepsilon}(\varepsilon, x)d\mu(\varepsilon) \right\} d\mu(x) \\ &= - \int \alpha(x)h(x^T\beta)\eta_{10}(x) \left\{ \int \eta_{20}(\varepsilon, x)d\mu(\varepsilon) \right\} d\mu(x) \\ &= - E\{\alpha(x)h(x^T\beta)\}. \end{aligned}$$

This implies $E\{\alpha(\mathbf{x}) | \mathbf{x}^T\beta\} = 0$. Therefore,

$$\begin{aligned} \Lambda^\perp &= [\{\alpha(x) - E(\alpha|x^T\beta)\}\varepsilon: \forall \alpha] \\ &= [Y - E(Y|x^T\beta)]\{\alpha(x) - E(\alpha|x^T\beta)\}: \forall \alpha]. \end{aligned}$$

A.5. Regularity Conditions for Theorem 1

(C1) The univariate kernel function $K(\cdot)$ is Lipschitz, has compact support. It satisfies

$$\begin{aligned} \int K(u)du &= 1, \int u^i K(u)du = 0, 1 \leq i \leq m-1, \\ 0 &\neq \int u^m K(u)du < \infty. \end{aligned}$$

The d -dimensional kernel function is a product of d univariate kernel functions, that is, $K_h(\mathbf{u}) = K(u/h)/h^d = \prod_{j=1}^d K_h(u_j) = \prod_{j=1}^d K(u_j/h)/h^d$ for $\mathbf{u} = (u_1, \dots, u_d)^T$. Here, we abuse the notation and use the same K regardless of the dimension of its argument.

(C2) Let $\mathbf{r}_1(\mathbf{x}^T\beta) = E\{\alpha(\mathbf{x}) | \mathbf{x}^T\beta\} f(\mathbf{x}^T\beta)$ and $\mathbf{r}_2(\mathbf{x}^T\beta) = E\{\mathbf{g}(Y, \mathbf{x}^T\beta) | \mathbf{x}^T\beta\} f(\mathbf{x}^T\beta)$. The m th derivatives of $\mathbf{r}_1(\mathbf{x}^T\beta)$, $\mathbf{r}_2(\mathbf{x}^T\beta)$ and $f(\mathbf{x}^T\beta)$ are locally Lipschitz-continuous.

(C3) The density functions of \mathbf{x} and $\mathbf{x}^T\beta$, denoted, respectively, by $f_{\mathbf{x}}(\mathbf{x})$ and $f(\mathbf{x}^T\beta)$, are bounded from below and above. Each entry in the matrices $E\{\alpha(\mathbf{x})\alpha^T(\mathbf{x}) \mid \mathbf{x}^T\beta\}$ and $E\{\mathbf{g}(Y, \mathbf{x}^T\beta)\mathbf{g}^T(Y, \mathbf{x}^T\beta) \mid \mathbf{x}^T\beta\}$ is locally Lipschitz-continuous and bounded from above as a function of $\mathbf{x}^T\beta$.

(C4) The bandwidth $h = O(n^{-\kappa})$ for $1/(4m) < \kappa < 1/(2d)$.

References

- Albright, SC.; Winston, WL.; Zappe, CJ. Data Analysis and Decision Making with Microsoft Excel. Pacific Grove, CA: Duxbury Press; 1999. [176]
- Bickel, PJ.; Klaassen, CAJ.; Ritov, Y.; Wellner, JA. Efficient and Adaptive Estimation for Semiparametric Models. Baltimore, MD: The Johns Hopkins University Press; 1993. [169]
- Box GEP, Cox DR. An Analysis of Transformations (with discussion). Journal of the Royal Statistical Society. 1964; 26:211–252. [168]. Series B
- Cook, RD. Regression Graphics: Ideas for Studying Regressions through Graphics. New York: Wiley; 1998. [168, 169, 170]
- Cook RD, Forzani L. Likelihood-Based Sufficient Dimension Reduction. Journal of the American Statistical Association. 2009; 104:197–208. [168].
- Cook RD, Li B. Dimension Reduction for Conditional mean in regression. The Annals of Statistics. 2002; 30:455–474. [168, 172].
- Cook RD, Li L. Dimension Reduction in Regressions with Exponential Family Predictors. Journal of Computational and Graphical Statistics. 2005; 18:774–791. [170].
- Cook RD, Nachtsheim CJ. Reweighting to Achieve Elliptically Contoured Covariates in Regression. Journal of the American Statistical Association. 1994; 89:592–599. [168].
- Cook RD, Weisberg S. Discussion of Sliced Inverse Regression for Dimension Reduction. Journal of the American Statistical Association. 1991; 86:28–33. [168, 171].
- Dong Y, Li B. Dimension Reduction for Non-Elliptically Distributed Predictors: Second-Order Moments. Biometrika. 2010; 97:279–294. [168, 172, 174].
- Fung WK, He X, Liu L, Shi P. Dimension Reduction Based on Canonical Correlation. Statistica Sinica. 2002; 12:1093–1113. [168].
- Hall P, Li KC. On Almost Linearity of Low Dimensional Projections from High Dimensional Data. The Annals of Statistics. 1993; 21:867–889. [176].
- Härdle W, Hall P, Ichimura H. Optimal Smoothing in Single-Index Models. The Annals of Statistics. 1993; 21:157–178. [168].
- Härdle W, Stoker TM. Investigating Smooth Multiple Regression by the Method of Average Derivatives. Journal of the American Statistical Association. 1989; 84:986–995. [168].
- Hernández A, Velilla S. Dimension Reduction in Nonparametric Kernel Discriminant Analysis. Journal of Computational and Graphical Statistics. 2005; 14:847–866. [168].
- Henmi M, Eguchi S. A Paradox Concerning Nuisance Parameters and Projected Estimating Functions. Biometrika. 2004; 91:929–941. [174].
- Horowitz JL, Härdle W. Direct Semiparametric Estimation of Single-Index Models with Discrete Covariates. Journal of the American Statistical Association. 1996; 91:1632–1639. [168, 173].
- Ichimura H. Semiparametric Least Squares (SLS) and Weighted SLS Estimation of Single-Index Models. Journal of Econometrics. 1993; 58:71–120. [168, 173].
- Li KC. Sliced Inverse Regression for Dimension Reduction (with discussion). Journal of the American Statistical Association. 1991; 86:316–342. [168, 171].
- Li KC. On Principal Hessian Directions for Data Visualization and Dimension Reduction: Another Application of Stein's Lemma. Journal of the American Statistical Association. 1992; 87:1025–1039. [168, 172].
- Li L. Sparse Sufficient Dimension Reduction. Biometrika. 2007; 94:603–613. [174].
- Li B, Dong Y. Dimension Reduction for Non-Elliptically Distributed Predictors. The Annals of Statistics. 2009; 37:1272–1298. [168, 172].

- Li KC, Duan N. Regression Analysis Under Link Violation. *The Annals of Statistics*. 1989; 17:1009–1052. [168, 172].
- Li B, Wang S. On Directional Regression for Dimension Reduction. *Journal of the American Statistical Association*. 2007; 102:997–1008. [168, 171, 172].
- Newey WK. Semiparametric Efficiency Bounds. *Journal of Applied Econometrics*. 1990; 5:99–135. [170].
- Park JH, Sriram TN, Yin X. Dimension Reduction in Time Series. *Statistica Sinica*. 2010; 20:747–770. [168].
- Power JL, Stock JH, Stoker TM. Semiparametric Estimation of Index Coefficients. *Econometrica*. 1989; 51:1403–1430. [168].
- Tsiatis, AA. *Semiparametric Theory and Missing Data*. New York: Springer; 2006. [169]
- Wang J, Wang L. Sparse Supervised Dimension Reduction in High Dimensional Classification. *Electronic Journal of Statistics*. 2010; 4:914–931. [174].
- Wang H, Xia Y. Sliced Regression for Dimension Reduction. *Journal of the American Statistical Association*. 2008; 103:811–821. [168].
- Xia Y. A Constructive Approach to the Estimation of Dimension Reduction Directions. *The Annals of Statistics*. 2007; 35:2654–2690. [168, 174].
- Xia Y, Tong H, Li WK, Zhu LX. An Adaptive Estimation of Dimension Reduction Space (with discussion). *Journal of the Royal Statistical Society*. 2002; 64:363–410. [168, 174]. Series B
- Ye Z, Weiss RE. Using the Bootstrap to Select one of a New Class of Dimension Reduction Methods. *Journal of the American Statistical Association*. 2003; 98:968–979. [174].
- Yin X, Cook RD. Direction Estimation in Single-Index Regressions. *Biometrika*. 2005; 92:371–384. [168].
- Yin X, Li B, Cook RD. Successive Direction Extraction for Estimating the Central Subspace in a Multiple-Index Regression. *Journal of Multivariate Analysis*. 2008; 99:1733–1757. [168].
- Zhu LX, Fang KT. Asymptotics for Kernel Estimation of Sliced Inverse Regression. *The Annals of Statistics*. 1996; 3:1053–1068. [168].
- Zhu Y, Zeng P. Fourier Methods for Estimating the Central Subspace and the Central Mean Subspace in Regression. *Journal of the American Statistical Association*. 2006; 101:1638–1651. [168].

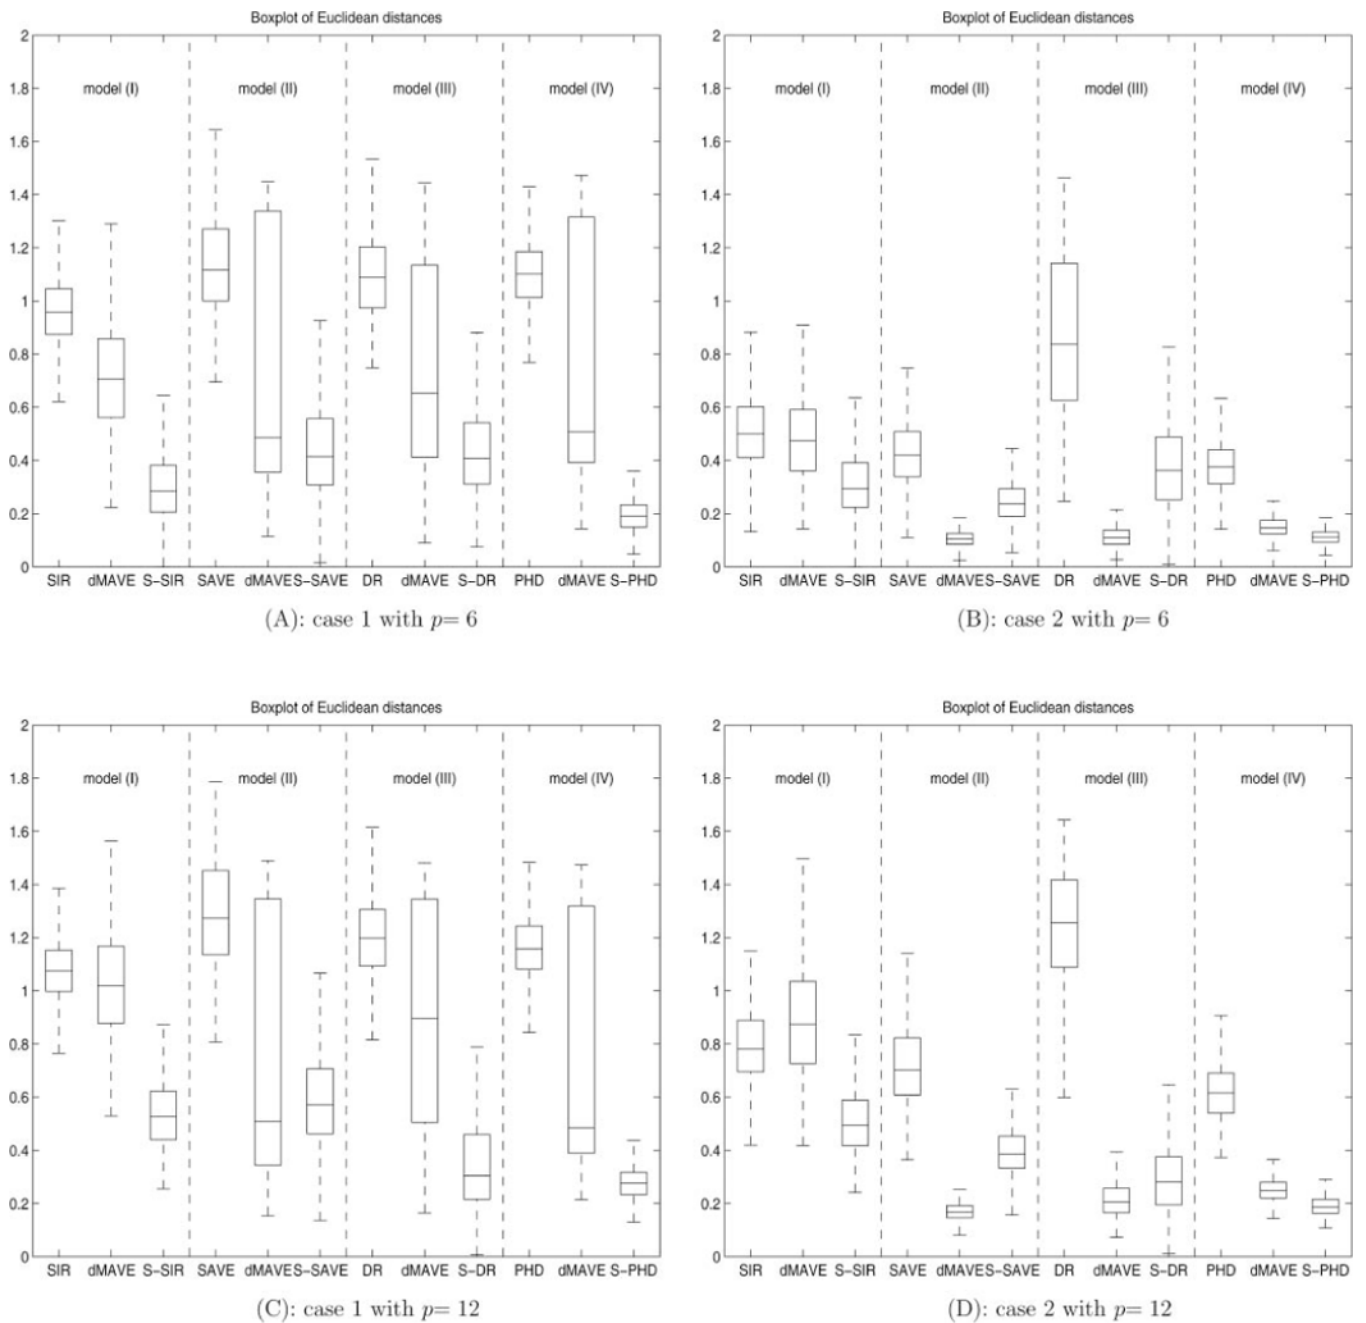


Figure 1. Boxplots of Euclidean distances for models (I)–(IV) with $p = 6$ and $p = 12$ in Example 1. Semi-SIR, Semi-SAVE, Semi-DR, and Semi-PHD are shortened to S-SIR, S-SAVE, S-DR, and S-PHD, respectively, in the labels.




Figure 2. Boxplot of the r_k^2 values defined in (10) from 1000 bootstrapped Semi-DR.

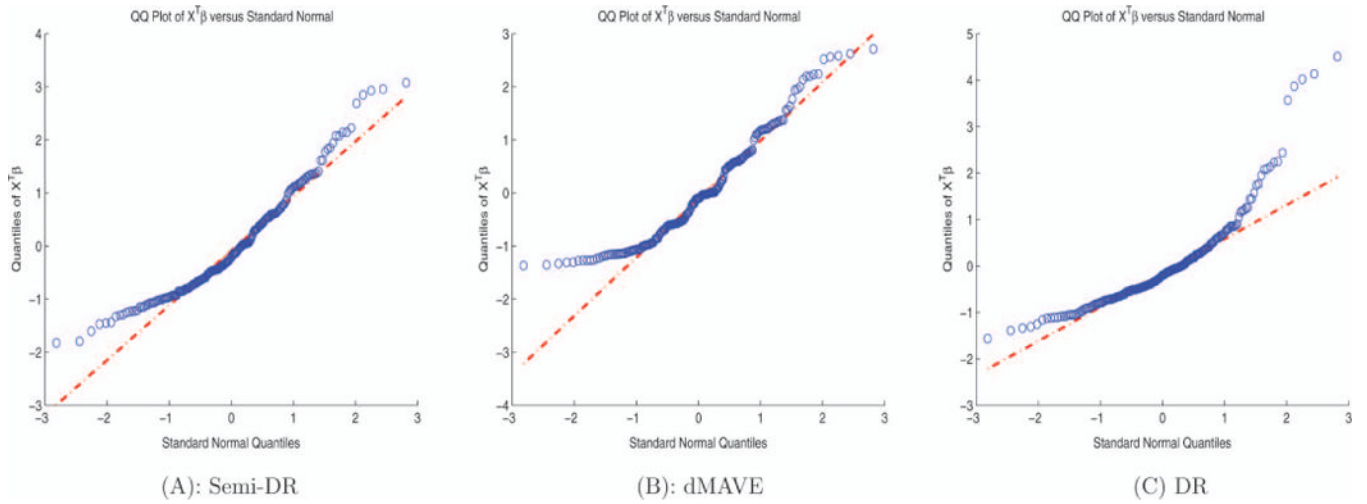


Figure 3. The Q–Q plot of $x^T \hat{\beta}_1$, with $\hat{\beta}_1$ estimated from Semi-DR, dMAVE, and DR. (The online version of this figure is in color.)




Figure 4.

Scatterplot of Y versus $x^T \hat{\beta}_1$, with $\hat{\beta}_1$ estimated from Semi-DR, dMAVE, and DR. The dash lines are fitted curves based on quadratic regression modeling.

Table 1

Relative frequency of the estimated dimension \hat{d}

Model	Method	$\hat{d} = 1$	$\hat{d} = 2$	$\hat{d} = 3$	$\hat{d} = 4$
(I)	Semi-SIR	1%	99%	0%	0%
(II)	Semi-SAVE	1%	97%	0%	2%
(III)	Semi-DR	3%	97%	0%	0%
(IV)	Semi-PHD	5%	95%	0%	0%

Table 2

Estimated coefficients of \mathbf{x}

Method	X_1	X_2	X_3	X_4	X_5	X_6
Semi-DR	0.7330	0.5794	0.2523	0.0147	0.1830	0.1723
dMAVE	0.8388	0.4404	0.2295	0.0417	0.1373	0.1711
DR	0.2673	0.8596	0.2829	-0.1023	0.3050	0.0782