

PRINCIPAL SUPPORT VECTOR MACHINES FOR LINEAR AND NONLINEAR SUFFICIENT DIMENSION REDUCTION

BY BING LI¹, ANDREAS ARTEMIIOU AND LEXIN LI²

Pennsylvania State University, Michigan Technological University and North Carolina State University

We introduce a principal support vector machine (PSVM) approach that can be used for both linear and nonlinear sufficient dimension reduction. The basic idea is to divide the response variables into slices and use a modified form of support vector machine to find the optimal hyperplanes that separate them. These optimal hyperplanes are then aligned by the principal components of their normal vectors. It is proved that the aligned normal vectors provide an unbiased, \sqrt{n} -consistent, and asymptotically normal estimator of the sufficient dimension reduction space. The method is then generalized to nonlinear sufficient dimension reduction using the reproducing kernel Hilbert space. In that context, the aligned normal vectors become functions and it is proved that they are unbiased in the sense that they are functions of the true nonlinear sufficient predictors. We compare PSVM with other sufficient dimension reduction methods by simulation and in real data analysis, and through both comparisons firmly establish its practical advantages.

1. Introduction. With the increase of computer power in storing and processing data, high dimensional data have become increasingly prevalent across many disciplines. The demand for effective methods to extract useful information from such data has led inevitably to dimension reduction, an area that has undergone tremendous development during the past two decades.

Let \mathbf{X} be a p -dimensional predictor and Y be a response variable. In its classical form, sufficient dimension reduction (SDR) [Li (1991, 1992), Cook and Weisberg (1991), Cook (1998)] identifies a small number of linear combinations of predictors that can replace the original predictor vector \mathbf{X} without loss of information on the conditional distribution of Y given \mathbf{X} . In other words, the objective is to find a $p \times d$ ($d < p$) matrix $\boldsymbol{\eta}$ such that the following conditional independence holds:

$$(1) \quad Y \perp\!\!\!\perp \mathbf{X} | \boldsymbol{\eta}^\top \mathbf{X}.$$

In this relation, the identifiable parameter is the subspace spanned by the columns of $\boldsymbol{\eta}$, rather than $\boldsymbol{\eta}$ itself. The intersection of all subspaces satisfying (1), provided

Received October 2010; revised September 2011.

¹Supported in part by NSF Grants DMS-07-04621 and DMS-08-06058.

²Supported in part by NSF Grant DMS-11-06668.

MSC2010 subject classifications. 62-09, 62G08, 62H12.

Key words and phrases. Contour regression, invariant kernel, inverse regression, principal components, reproducing kernel Hilbert space, support vector machine.

itself satisfies (1), is called the central subspace, and is denoted by $\mathcal{S}_{Y|X}$ [Cook (1994)]. Cook (1996) and Yin, Li and Cook (2008) showed that $\mathcal{S}_{Y|X}$ uniquely exists under very mild conditions. Thus, we assume its existence throughout this article. Many methods have been proposed for this problem since the publication of the original works. See, for example, Cook and Li (2002), Xia et al. (2002), Yin and Cook (2002), Fung et al. (2002), Li, Zha and Chiaromonte (2005), Cook and Ni (2005), Li and Wang (2007), Li and Dong (2009).

A more general sufficient dimension reduction problem, as formulated in Cook (2007), is to seek an arbitrary function $\phi: \mathbb{R}^p \rightarrow \mathbb{R}^d$ such that

$$(2) \quad Y \perp\!\!\!\perp \mathbf{X} | \phi(\mathbf{X}).$$

We refer to this problem as *nonlinear sufficient dimension reduction*, and any one-to-one function of $\phi(\mathbf{X})$ as the *nonlinear sufficient predictor*. Several recent works pioneered estimation procedures for nonlinear dimension reduction of this type, including Wu (2008), Wu, Liang and Mukherjee (2008), Wang (2008) and Yeh, Huang and Lee (2009), by extending sliced inverse regression [SIR; Li (1991)] from different angles.

In this paper, we propose a sufficient dimension reduction method, to be called the principal support vector machine (PSVM), that can extract the sufficient predictors in both problems (1) and (2). Let $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$ be a sample of (\mathbf{X}, Y) . The basic idea of PSVM is to divide $\mathbf{X}_1, \dots, \mathbf{X}_n$ into several slices according to the values of the responses, and then use support vector machine [SVM; Vapnik (1998)] to find the optimal hyperplanes that separate these slices. The optimal hyperplanes are then aligned by applying principal component analysis to their normal vectors. We show that the principal components are, in fact, an unbiased estimator of the central subspace $\mathcal{S}_{Y|X}$. This idea is then extended to the nonlinear dimension reduction problem (2) via the reproducing kernel Hilbert space [RKHS; Aronszajn (1950), Hsing and Ren (2009)]. In this context, the normal vectors in the linear case become functions in the RKHS. It is shown that the normal functions thus derived are functions of ϕ in the general problem (2). This is, to our knowledge, the first result of this type.

Our proposal is noticeably different from the existing SDR methods in the following respects. First, PSVM is developed under, and for, a unified framework of linear and nonlinear sufficient dimension reduction. Such a standpoint allows us to formulate some theoretical properties, such as unbiasedness, more rigorously and generally than previous works. Second, PSVM improves the accuracy for sufficient dimension reduction, for the following reason. It is well known that a regression surface is more accurately estimated at the center of the data cloud than at the outskirts. However, an inverse regression based method, such as SIR, tends to down-weight the slice means near the center due to their shorter lengths. Since PSVM relies on separating hyperplanes rather than slice means, it makes better use of the central portion of the data than inverse regression. This improvement is clearly

demonstrated in our numerical studies. Finally, PSVM establishes a firm connection between sufficient dimension reduction and the acclaimed machine learning technique, support vector machine, both of which have been extensively used in high dimensional data analysis. This combination brings fresh insights and further advances to both subjects. Along with the theoretical development of PSVM, we develop a more complete asymptotic theory for SVM than previously given, and introduce the notion of invariant kernel for SVM. Meanwhile, we expect some inherent advantages of SVM to benefit sufficient dimension reduction estimation. For instance, SVM tends to be more robust against outliers than a typical moment method. This is because the separating hyperplanes are largely determined by the support vectors lying in the interior of the data cloud, as a result an observation far away from the data cloud has less influence than a typical moment-based estimator. In this sense, SVM behaves more like a median than a mean. It is also expected to help address several challenging issues facing the existing SDR methods, such as small- n -large- p and presence of categorical predictors. However, due to limited space these potential advantages cannot be fully discussed within this paper. Some of them, such as robustness and categorical predictors, are further explored in Artemiou (2010).

The rest of the paper is organized as follows. In Section 2, we illustrate the basic idea of PSVM by examples and figures, and give intuitions about why it works. In Section 3, we formally introduce the linear PSVM and study its population-level properties in terms of its unbiasedness as an estimator of the central subspace. In Section 4, we develop the estimation procedures for the linear PSVM, and describe how to implement it using standard SVM packages. In Section 5, we generalize the linear PSVM to the kernel PSVM to solve the nonlinear sufficient dimension reduction problem, and establish its unbiasedness in this general setting. In Section 6, we develop an algorithm to implement the kernel PSVM, and introduce the notion of invariant kernel. In Section 7, we study the asymptotic properties for the linear PSVM estimator. Though the identified subspaces are asymptotically consistent, they are almost surely incorrect for finite sample sizes. Thus, in Section 8, we compare the linear and kernel PSVM with other dimension reduction methods in finite sample by simulation. In Section 9, we apply it to analyze a data set concerning the recognition of vowels, and make further comparisons in the practical setting.

2. Principal support vector machine: The basic idea. The idea of the principal support vector machine arises from an interplay of several ideas: sliced inverse regression, support vector machine, and contour regression [Li, Zha and Chiaromonte (2005)]. In this section, we illustrate this idea by two simple examples that cover both linear and nonlinear dimension reduction. Throughout this paper, \mathbf{X} represents a random vector; X_r represents the r th component of \mathbf{X} ; \mathbf{X}_i represents the i th random vector from a sample $\mathbf{X}_1, \dots, \mathbf{X}_n$, and X_{ir} represents the r th component of \mathbf{X}_i .

First, consider the regression model

$$(3) \quad Y = f(X_1 + 2X_2) + \varepsilon,$$

where $\varepsilon \perp (X_1, X_2)$. This is a linear sufficient dimension reduction problem, in which the central subspace is spanned by $(1, 2)^\top \in \mathbb{R}^2$. Note that the contours for the regression function is the set $\{(x_1, x_2) : x_1 + 2x_2 = c\}$, which is uniquely associated with the vector $(1, 2)^\top$. Based on this intuition, Li, Zha and Chiaromonte (2005) introduced the contour regression, which estimates the contour directions by the directions in \mathbf{X} that are aligned with the smallest increments in Y .

Here, we propose to identify the contours by the separating hyperplanes derived from the support vector machine as applied to different slices of \mathbf{X} , formed according to the values of Y . Let $S_1 = \{\mathbf{X}_i : Y_i \leq c\}$ and $S_2 = \{\mathbf{X}_i : Y_i > c\}$ for some constant c . We use SVM to obtain the optimal separating hyperplane of S_1 and S_2 , and repeat the process to obtain several hyperplanes. Intuitively, the normals of these hyperplanes are roughly aligned with the directions in which the regression surface varies—directions that form the central subspace. We use the principal components of these normals to estimate the central subspace. A related idea is Loh (2002), who proposed to divide each individual predictor according to the mean of Y and assess the importance of that predictor by its degree of separation.

As an illustration, we generate 100 replications from model (3) where f is taken to be the identity mapping. We divide $\mathbf{X}_1, \dots, \mathbf{X}_{100}$ into 4 slices according to the 25th, 50th, 75th sample quantiles of Y_1, \dots, Y_n , as indicated in Figure 1 by differently colored dots. Application of SVM between these slices yields three hyperplanes, represented by the solid lines on the right panel, which closely resemble




FIG. 1. Linear contours for model $Y = 2X_1 + X_2 + \varepsilon$. Left panel: true contours; right panel: contours based on linear SVM. Contour levels are three evenly spaced sample quantiles of Y_1, \dots, Y_n . The sample size is $n = 100$.




FIG. 2. *Nonlinear contours for model $Y = X_1 + X_2^2 + \varepsilon$. Left panel: true contours; right panel: contours based on kernel SVM with gauss radial kernel. Contour levels are three evenly spaced sample quantiles of Y_1, \dots, Y_n . The sample size is $n = 100$.*

the contours derived from the true model, as shown on the left panel. Clearly, the normals of the three hyperplanes give close estimate of the central subspace.

We can apply the same idea to sufficient nonlinear dimension reduction. Let

$$(4) \quad Y = f(X_1 + X_2^2) + \varepsilon,$$

where f is an unknown function. The contours of this function are of the form $\{(x_1, x_2) : x_1 + x_2^2 = c\}$, which are no longer hyperplanes in \mathbb{R}^2 . However, if we map \mathbf{x} to a higher dimensional space of functions of \mathbf{x} that is rich enough to contain $x_1 + x_2^2$, then the contours become hyperplanes again. We can apply SVM at that level to find the optimal hyperplanes, and then map them back to the \mathbf{x} -space to extract the nonlinear predictor. Usually, in conjunction with mapping a low-dimensional regressor to a high-dimensional regressor, a Tikhonov-type regularization is applied, so that the overfitting tendency of increased dimension is counteracted by the regularization.

As in the linear case we generate 100 replications from model (4) and use the same set of quantiles to slice the response. The curves in the left panel in Figure 2 are the true contours computed from the function $y = x_1 + x_2^2$. Those in the right panel are obtained by first applying kernel SVM (with Gaussian radial basis) to find hyperplanes in \mathbb{R}^{100} and then mapping them back to \mathbb{R}^2 . Clearly, any function of (x_1, x_2) that generates the contours in the right panel would closely resemble the true predictor $x_1 + x_2^2$, modulo a monotone transformation.

3. PSVM for linear sufficient dimension reduction. We first develop PSVM for linear sufficient dimension reduction. We begin with a population-level formulation of SVM, since it is usually described at the sample level, which is not the

best way to set up our problem. For now, assume Y to be a binary random variable taking values -1 and 1 . The soft-margin SVM is defined through the following optimization:

$$(5) \quad \begin{aligned} &\text{minimize } \boldsymbol{\psi}^\top \boldsymbol{\psi} + \frac{\lambda}{n} \sum_{i=1}^n \xi_i \quad \text{among } (\boldsymbol{\psi}, t, \boldsymbol{\xi}) \in \mathbb{R}^p \times \mathbb{R} \times \mathbb{R}^n \\ &\text{subject to } \xi_i \geq 0, Y_i[\boldsymbol{\psi}^\top (\mathbf{X}_i - \bar{\mathbf{X}}) - t] \geq 1 - \xi_i, \quad i = 1, \dots, n, \end{aligned}$$

where λ is a positive constant often referred to as the “cost.” See Vapnik [(1998), page 411] for the intuitions behind this construction. If $(\boldsymbol{\psi}^*, t^*, \boldsymbol{\xi}^*)$ is the solution to (5), then the set $\{\mathbf{x} : \boldsymbol{\psi}^{*\top} \mathbf{x} = t^*\}$ is the optimal hyperplane that separates $\{\mathbf{X}_i : Y_i = -1\}$ and $\{\mathbf{X}_i : Y_i = 1\}$.

Although this representation defines the algorithm, it does not tell us what objective function is minimized at the population level. To see things more clearly, we first carry out the optimization for a fixed $(\boldsymbol{\psi}, t)$. This amounts to minimizing $\sum_{i=1}^n \xi_i$ subject to $\xi_i \geq \max\{0, 1 - Y_i[\boldsymbol{\psi}^\top (\mathbf{X}_i - \bar{\mathbf{X}}) - t]\}$. The optimal solution is $\xi_i^* = \{1 - Y_i[\boldsymbol{\psi}^\top (\mathbf{X}_i - \bar{\mathbf{X}}) - t]\}^+$ where $a^+ = \max(a, 0)$. Substituting ξ_i^* into (5), we have

$$(6) \quad \boldsymbol{\psi}^\top \boldsymbol{\psi} + \frac{\lambda}{n} \sum_{i=1}^n \{1 - Y_i[\boldsymbol{\psi}^\top (\mathbf{X}_i - \bar{\mathbf{X}}) - t]\}^+.$$

This corresponds to the following objective function at the population level:

$$(7) \quad \boldsymbol{\psi}^\top \boldsymbol{\psi} + \lambda E[1 - Y(\boldsymbol{\psi}^\top (\mathbf{X} - E\mathbf{X}) - t)]^+.$$

The hyperplane that minimizes this criterion can be viewed as that which best separates the *conditional distributions* of $\mathbf{X}|Y = -1$ and $\mathbf{X}|Y = 1$. Jiang, Zhang and Cai (2008) used a slight variation of representation (7) to derive the asymptotic distribution of SVM. We also use a representation similar to (7) but with two important modifications, as we describe below.

Return now to sufficient dimension reduction problem (1) where Y is an arbitrary random variable (in particular, it can be either continuous or categorical). Let Ω_Y be the support of Y and let A_1 and A_2 be disjoint subsets of Ω_Y . Let \tilde{Y} be the discrete random variable defined by

$$(8) \quad \tilde{Y} = I(Y \in A_2) - I(Y \in A_1).$$

We introduce the following objective function for linear SDR:

$$(9) \quad L(\boldsymbol{\psi}, t) = \boldsymbol{\psi}^\top \boldsymbol{\Sigma} \boldsymbol{\psi} + \lambda E\{1 - \tilde{Y}[\boldsymbol{\psi}^\top (\mathbf{X} - E\mathbf{X}) - t]\}^+,$$

where $\boldsymbol{\Sigma} = \text{var}(\mathbf{X})$. Compared with (7), we have made two modifications. First, we allow \tilde{Y} to take the value 0, so that we can use a pair of disjoint subsets that are not a partition of Ω_Y . Second, we have inserted $\boldsymbol{\Sigma}$ in the first term of (9). This

is so that the objective function transforms in a desired manner. We will return to this point in Section 6.

We now establish the unbiasedness of the normal vector for the optimal separating hyperplane in SVM as an estimator of the central subspace. Let F_n be the empirical distribution based on the sample $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$, F_0 be the true distribution of (\mathbf{X}, Y) , and \mathbf{T} be a statistic that can be expressed as a matrix-valued function of the distribution of (\mathbf{X}, Y) . In our context, we say $\mathbf{T}(F_n)$ is an unbiased estimator of $\mathcal{S}_{Y|\mathbf{X}}$, if it satisfies

$$(10) \quad \text{span}[\mathbf{T}(F_0)] \subseteq \mathcal{S}_{Y|\mathbf{X}}.$$

THEOREM 1. *Suppose $E(\mathbf{X}|\boldsymbol{\eta}^\top \mathbf{X})$ is a linear function of $\boldsymbol{\eta}^\top \mathbf{X}$, where $\boldsymbol{\eta}$ is as defined in (1). If $(\boldsymbol{\psi}^*, t^*)$ minimizes the objective function (9) among all $(\boldsymbol{\psi}, t) \in \mathbb{R}^p \times \mathbb{R}$, then $\boldsymbol{\psi}^* \in \mathcal{S}_{Y|\mathbf{X}}$.*

The linearity condition on $E(\mathbf{X}|\boldsymbol{\eta}^\top \mathbf{X})$ in the theorem is well known and generally assumed in the SDR literature. See, for example, Li and Duan (1989), Li (1991) and Li and Dong (2009). It implies

$$(11) \quad E(\boldsymbol{\psi}^\top \mathbf{X}|\boldsymbol{\eta}^\top \mathbf{X}) = \boldsymbol{\psi}^\top \mathbf{P}_\boldsymbol{\eta}^\top(\boldsymbol{\Sigma})\mathbf{X},$$

where $\mathbf{P}_\boldsymbol{\eta}(\boldsymbol{\Sigma})$ is the projection matrix $\boldsymbol{\eta}(\boldsymbol{\eta}^\top \boldsymbol{\Sigma} \boldsymbol{\eta})^{-1} \boldsymbol{\eta}^\top \boldsymbol{\Sigma}$ [Cook (1998)]. It is satisfied when \mathbf{X} is elliptically symmetric [Eaton (1986)], and is approximately satisfied when p is large [Hall and Li (1993)]. Interestingly, as we show in Section 5, this assumption is no longer needed for the unbiasedness in the more general setting of nonlinear sufficient dimension reduction.

Here we note that, though Theorem 1 is far from a trivial generalization, the type of argument used in the proof is somewhat standard in the SDR literature. See, for example, Li and Duan (1989), Cook (1998) and Cook and Li (2002). It is possible to extend the above theorem to more general objective functions. For example, the theorem still holds if $a \mapsto a^+$ in the objective function is replaced by any convex function $u(a)$.

4. Estimation procedure for linear PSVM.

4.1. *Estimation.* We propose two ways to generate the set of pairs of slices for PSVM. One, which we call “left versus right” (LVR), repeatedly divides the predictors into two groups according to a set of cutting points for the response. The other, which we call “one versus another” (OVA), partitions the predictors into several slices and pairs up all possible slices. We summarize the estimation procedure as follows.

1. Compute the sample mean $\bar{\mathbf{X}}$ and sample variance matrix $\hat{\boldsymbol{\Sigma}}$.

2. (LVR) Let $q_r, r = 1, \dots, h - 1$, be $h - 1$ dividing points. For example, they can be equally spaced sample percentiles of $\{Y_1, \dots, Y_n\}$. Let

$$(12) \quad \tilde{Y}_i^r = I(Y_i > q_r) - I(Y_i \leq q_r)$$

and let $(\hat{\psi}_r, \hat{t}_r), r = 1, \dots, h - 1$, be the minimizer of

$$(13) \quad \psi^\top \hat{\Sigma} \psi + \lambda E_n \{1 - \tilde{Y}^r [(\mathbf{X} - \bar{\mathbf{X}})^\top \psi - t]\}^+.$$

2'. (OVA) Apply SVM to each pair of slices from the h slices. More specifically, let $q_0 = \min\{Y_1, \dots, Y_n\}$ and $q_h = \max\{Y_1, \dots, Y_n\}$. For each (r, s) satisfying $1 \leq r < s \leq h$, let

$$\tilde{Y}_i^{rs} = I(q_{s-1} < Y_i \leq q_s) - I(q_{r-1} < Y_i \leq q_r).$$

Let $(\hat{\psi}_{rs}, \hat{t}_{rs})$ be the minimizer of the objective function

$$\psi^\top \hat{\Sigma} \psi + \lambda E_n \{1 - \tilde{Y}^{rs} [(\mathbf{X} - \bar{\mathbf{X}})^\top \psi - t]\}^+.$$

3. Let $\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_d$ be the d leading eigenvectors of either one of the matrices

$$(14) \quad \hat{\mathbf{M}}_n = \sum_{r=1}^{h-1} \hat{\psi}_r \hat{\psi}_r^\top \quad \text{or} \quad \hat{\mathbf{M}}_n = \sum_{r=1}^h \sum_{s=r+1}^h \hat{\psi}_{rs} \hat{\psi}_{rs}^\top.$$

We use subspace spanned by $\hat{\mathbf{v}} = (\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_d)$ to estimate $S_{Y|\mathbf{X}}$.

Based on our experiences, LVR works best when the response is a continuous variable, where Y being larger or smaller has a concrete physical meaning; OVA works best when the response is categorical, where the values of Y are simply labels of classes, such as different vowels in our example in Section 9. Our numerical studies also suggest that the estimation results are not overly sensitive to the choice of the number of slices h , though a larger h often works better.

Standard packages for SVM minimize the objective function (6) instead of (13). However, they can be modified to suit our procedure. Let $\boldsymbol{\zeta} = \hat{\Sigma}^{1/2} \boldsymbol{\psi}$ and $\mathbf{Z} = \hat{\Sigma}^{-1/2} (\mathbf{X} - \bar{\mathbf{X}})$. Then (13) becomes

$$(15) \quad \boldsymbol{\zeta}^\top \boldsymbol{\zeta} + \lambda E_n [1 - \tilde{Y}^r (\mathbf{Z}^\top \boldsymbol{\zeta} - t)]^+.$$

We can apply standard packages to minimize (15) to obtain $\hat{\boldsymbol{\zeta}}$, whose transformation $\hat{\Sigma}^{-1/2} \hat{\boldsymbol{\zeta}}$ is the desired minimizer of (13). We use the kernlab package in R to solve problem (15). See Karatzoglou and Meyer (2006) for an exposition of this package.

4.2. *Order determination.* Estimating the dimension d of the central subspace is a vital ingredient of sufficient dimension reduction estimation. Here, we propose a cross-validated BIC procedure [Schwarz (1978)] for this purpose. The BIC component of this procedure is an extension of a criterion introduced by Wang and

Yin (2008), and is also related to Zhu, Miao and Peng (2006). We refer to this combined procedure as CVBIC.

Let $\hat{\mathbf{M}}_n$ be one of the matrices in (14), and let $\lambda_i(\hat{\mathbf{M}}_n)$ be its i th largest eigenvalue. Let $G_n(k) = \sum_{i=1}^k \lambda_i(\hat{\mathbf{M}}_n) - c_1(n)c_2(k)$, where $c_1(n)$ is a sequence of positive numbers or random variables that converge (in probability) to 0, and $c_2(k)$ is a nonrandom increasing function of k . Let \hat{d} be the maximizer of $G_n(k)$ over $\{0, \dots, p\}$. In Section 7, we show that $P(\hat{d} = d) \rightarrow 1$. The standard choices of $c_1(n)$ and $c_2(k)$ are $c_1(n) \propto n^{-1/2} \log(n)$ and $c_2(k) = k$, so that the penalty term is $c_0 n^{-1/2} \log(n)k$, where $c_0 > 0$ is a constant (or random variable) of order $O(1)$ [or $O_P(1)$]. Since the eigenvalues $\lambda_i(\hat{\mathbf{M}}_n)$ may differ for different problems, it is sensible to make c_0 comparable to their magnitude. One reasonable choice is to make c_0 proportional to $\lambda_1(\hat{\mathbf{M}}_n)$, leading to the following BIC-type criterion:

$$(16) \quad \sum_{i=1}^k \lambda_i(\hat{\mathbf{M}}_n) - a\lambda_1(\hat{\mathbf{M}}_n)n^{-1/2} \log(n)k.$$

We now turn to the choice of a . Though this choice does not affect the consistency of \hat{d} , it does affect its finite-sample performance. Moreover, from our experience this choice is also sensitive to p , d , and the regression model. For these reasons, it is important to have a systematic way of choosing a . The SVM used in our setting suggests naturally the cross-validation, because the former provides a set of labels to validate. We outline the CVBIC procedure as follows, using LVR as an illustration.

First, divide the data into a training set and a testing set, denoted by

$$\{(\hat{\mathbf{X}}_1, \hat{Y}_1), \dots, (\hat{\mathbf{X}}_{n_1}, \hat{Y}_{n_1})\}, \quad \{(\check{\mathbf{X}}_1, \check{Y}_1), \dots, (\check{\mathbf{X}}_{n_2}, \check{Y}_{n_2})\}.$$

Apply the PSVM to the training set with dividing points q_1, \dots, q_{h-1} to obtain a set normal vectors $\hat{\psi}_1, \dots, \hat{\psi}_{h-1}$. Let $\hat{\mathbf{M}}_{n_1} = \sum_{i=1}^{h-1} \hat{\psi}_i \hat{\psi}_i^\top$. Second, for a fixed a , maximize the criterion (16), with $\hat{\mathbf{M}}_n$ replaced by $\hat{\mathbf{M}}_{n_1}$, to obtain an integer k . Let $\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_k$ be the k leading eigenvectors of $\hat{\mathbf{M}}_{n_1}$ and transform the testing predictors $\check{\mathbf{X}}_i$ to $\check{\mathbf{X}}_i^{(k)} = (\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_k)^\top \check{\mathbf{X}}_i$, $i = 1, \dots, n_2$. Third, let $\check{L}_i = I(\check{Y}_i > q_r) - I(\check{Y}_i \leq q_r)$ be the true label of \check{Y}_i in the testing set. Apply SVM to $(\check{\mathbf{X}}_1^{(k)}, \check{L}_1), \dots, (\check{\mathbf{X}}_{n_2}^{(k)}, \check{L}_{n_2})$ to predict $\check{L}_1, \dots, \check{L}_{n_2}$. Repeat this process for all dividing points and record the total number of misclassifications. The optimal a is the one that minimizes the total number of misclassifications. Finally, substitute the optimal a into (16) and maximize it again using the full data to estimate d . In Section 8.3, we investigate the numerical performance of CVBIC under a variety of combinations of p , d , n and regression models.

4.3. *Special features of linear PSVM.* As we conclude the exposition of the linear PSVM, we mention some special features of this method. One is that it

shares the similar limitation with SIR when dealing with regression functions that are symmetric about the origin. If the regression function is $f(\|\mathbf{X}\|)$, then all slices of the form $\{\mathbf{X}_i : Y_i \in S\}$ are roughly concentric spheres in \mathbb{R}^p , which no hyperplane in \mathbb{R}^p can separate. However, as we shall see in Sections 5 and 8, this is remedied by the kernel PSVM, because when mapped into higher dimensional feature space the slices become linear again.

Another is that when dealing asymmetric regression functions, the linear PSVM tends to work better than SIR for the following reason. Recall that SIR is based, roughly, on the principal components of the slice mean vectors of the form $E(\mathbf{X}|Y \in S) - E(\mathbf{X})$, where S is an interval in Ω_Y . This determines that it downweights the slice means near the center of the data cloud, where the Euclidean norm of $E(\mathbf{X}|Y \in S) - E(\mathbf{X})$ is smaller. However, it is well known that the regression function $E(Y|\mathbf{X})$ tends to be more accurately estimated near the center of the data cloud [see, e.g., [Kutner, Nachtsheim and Neter \(2004\)](#), Section 2.4]. In comparison, the linear PSVM relies on the normals of the separating hyperplanes of the slices, which does not downweight the data near the center. As we will see from our simulation studies in Section 8, this brings substantial improvement to the estimate. We should point out, however, that there is an important exception. As shown in [Cook \(2007\)](#) and [Cook and Forzani \(2008\)](#), under the assumption that Y has a finite support and $\mathbf{X}|Y$ has a conditional multivariate normal distribution where $\text{var}(\mathbf{X}|Y)$ is independent of Y , SIR is the maximum likelihood estimate of the central subspace. In this case, no regular estimate can be more efficient than SIR. The mentioned advantage of linear PSVM applies mainly to the forward regression setting where the conditional distribution of $\mathbf{X}|Y$ is typically non-Gaussian.

5. Kernel PSVM for nonlinear dimension reduction. In this section, we extend the PSVM to nonlinear sufficient dimension reduction as defined by (2). We first develop the objective function by generalizing the linear PSVM objective function (9), and then establish the unbiasedness of the proposed nonlinear PSVM estimator.

Before proceeding further, we note that the function ϕ in relation (2) is not unique in the strict sense, but is unique modulo injective transformations. Again, the situation is parallel to linear sufficient dimension reduction problem (1), where $\eta^\top \mathbf{X}$ is only unique modulo injective linear transformations. Any injective linear transformation of $\eta^\top \mathbf{X}$ is an equivalent linear predictor, because it does not change the linear subspace. Likewise, for nonlinear SDR, any injective transformation of ϕ is an equivalent sufficient predictor, because it does not change conditional independence (2).

Let \mathcal{H} be a Hilbert space of functions of \mathbf{X} . In analogy to the linear objective function (9), consider $\Lambda : \mathcal{H} \times \mathbb{R} \rightarrow \mathbb{R}^+$ defined by

$$(17) \quad \Lambda(\psi, t) = \text{var}[\psi(\mathbf{X})] + \lambda E[1 - \tilde{Y}(\psi(\mathbf{X}) - E\psi(\mathbf{X}) - t)]^+,$$

where \tilde{Y} is as defined in (8). To see that this is indeed a generalization of (9), consider the bilinear form from $\mathcal{H} \times \mathcal{H}$ to \mathbb{R} defined by $b(f_1, f_2) = \text{cov}[f_1(\mathbf{X}), f_2(\mathbf{X})]$. Under the assumption that the mapping

$$(18) \quad \mathcal{H} \rightarrow L_2(P_{\mathbf{X}}), \quad f \mapsto f$$

is continuous, the bilinear form b induces a bounded and self-adjoint operator $\Sigma: \mathcal{H} \rightarrow \mathcal{H}$ such that $\langle f_1, \Sigma f_2 \rangle_{\mathcal{H}} = b(f_1, f_2)$, where $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ is the inner product in \mathcal{H} . See, for example, Conway (1990), Theorem 2.2, and Fukumizu, Bach and Jordan (2004). The objective function (17) can now be rewritten as

$$(19) \quad \Lambda(\psi, t) = \langle \psi, \Sigma \psi \rangle_{\mathcal{H}} + \lambda E[1 - \tilde{Y}(\psi(\mathbf{X}) - E\psi(\mathbf{X}) - t)]^+.$$

Thus, $\Lambda(\psi, t)$ is a generalization of $L(\boldsymbol{\psi}, t)$ with the matrix $\boldsymbol{\Sigma}$ replaced by the operator Σ , the linear function $\boldsymbol{\psi}^\top \mathbf{X}$ replaced by an arbitrary function ψ in \mathcal{H} , and the inner product in \mathbb{R}^p replaced by the inner product in \mathcal{H} . For the usual kernel SVM, the population-level objective function is

$$\langle \psi, \psi \rangle_{\mathcal{H}} + \lambda E[1 - \tilde{Y}(\psi(\mathbf{X}) - E\psi(\mathbf{X}) - t)]^+.$$

Comparing with (19), we see a parallel modification to the linear case. The significance of this modification is further discussed in Section 6.

We now establish that, if (ψ^*, t^*) is the minimizer of $\Lambda(\psi, t)$, then ψ^* is necessarily a function of the sufficient predictor $\boldsymbol{\phi}(\mathbf{X})$ in the nonlinear problem problem (2). This is a generalization of the notion *unbiasedness* in the linear setting. Our definition of unbiasedness (10) in the linear sufficient dimension reduction setting is equivalent to

$$(20) \quad [\mathbf{T}(F_0)]^\top \mathbf{X} \text{ is a linear function of } \boldsymbol{\eta}^\top \mathbf{X}.$$

It is the statement (20) that is more readily generalized to the nonlinear sufficient dimension reduction setting: we simply require ψ to be a function of the sufficient predictor $\boldsymbol{\phi}(\mathbf{X})$ in (2). The following definition makes this notion rigorous. For a generic random element \mathbf{U} , let $\sigma\{\mathbf{U}\}$ denote the σ -field generated by \mathbf{U} .

DEFINITION 1. A function $\psi \in \mathcal{H}$ is unbiased for nonlinear sufficient dimension reduction (2) if it has a version that is measurable $\sigma\{\boldsymbol{\phi}(\mathbf{X})\}$.

The reason that we only require a version of ψ to be measurable $\sigma\{\boldsymbol{\phi}(\mathbf{X})\}$ is that the L_2 -metric ignores measure zero sets.

THEOREM 2. *Suppose the mapping (18) is continuous and:*

1. \mathcal{H} is a dense subset of $L_2(P_{\mathbf{X}})$,
2. $Y \perp\!\!\!\perp \mathbf{X} | \boldsymbol{\phi}(\mathbf{X})$.

If (ψ^, t^*) minimizes (19) among all $(\psi, t) \in \mathcal{H} \times \mathbb{R}$, then $\psi^*(\mathbf{X})$ is unbiased.*

Condition 1 is satisfied by some commonly used reproducing kernel Hilbert spaces. For example, if \mathcal{G} is a reproducing kernel Hilbert space based on the Gaussian radial basis, then the collection of functions $\{c + g : c \in \mathbb{R}, g \in \mathcal{G}\}$ is dense in $L_2(P_{\mathbf{X}})$. See Fukumizu, Bach and Jordan (2009).

It is important to note that in this more general setting we no longer require any linearity assumption that resembles the one assumed in Theorem 1. In contrast, the kernel sliced inverse regression developed by Wu (2008) and Wu, Liang and Mukherjee (2008), and functional sliced inverse regression by Hsing and Ren (2009) all require a version of the linearity condition to hold in the reproducing kernel Hilbert space.

The notion of unbiasedness for sufficient dimension reduction is more akin to Fisher consistency than to unbiasedness in the classical setting. While unbiasedness in the classical setting can exclude many useful statistics, Fisher consistency often guarantees correct asymptotic behavior without putting undue restrictions on the expectation. Moreover, an estimator that is not Fisher consistent is clearly undesirable, because it is guaranteed *not* to converge to the true parameter. For these reasons unbiasedness for linear SDR is a useful criterion, even though some useless estimators (such as $\mathbf{0}$) are unbiased. Unbiasedness for nonlinear SDR plays the parallel role, except that it only requires the estimator to be an arbitrary, rather than a linear, function of the true predictor. This relaxation also allows us to establish the unbiasedness of PSVM without evoking the linearity condition.

Theorem 2 assumes that $\Lambda(\psi, t)$ attains its minimum in $\mathcal{H} \times \mathbb{R}$. We think this is a reasonable assumption for the following reasons. As shown below, $\Lambda(\psi, t)$ is lower semicontinuous with respect to the weak topology in $\mathcal{H} \times \mathbb{R}$. Since any closed, bounded, and convex set in a Hilbert space is compact with respect to the weak topology [Weidmann (1980), Theorem 4.25, Conway (1990), Corollary V.1.5], by the generalized Weierstrass theorem [Kurdila and Zabarankin (2005), Section 7.3], $\Lambda(\psi, t)$ attains its minimum within any such set in $\mathcal{H} \times \mathbb{R}$. The next proposition establishes this fact. Let \mathcal{H}' be the Hilbert space $\mathcal{H} \times \mathbb{R}$ endowed with the inner product $\langle \psi_1, \psi_2 \rangle_{\mathcal{H}} + t_1 t_2$.

PROPOSITION 1. *If \mathcal{H} is an RKHS with its kernel κ satisfying $E\kappa(\mathbf{X}, \mathbf{X}) < \infty$, then $\Lambda(\psi, t)$ is lower semicontinuous with respect to the weak topology in \mathcal{H}' , and attains its minimum in any closed, bounded, and convex set in \mathcal{H}' .*

6. Estimation of kernel PSVM and invariant kernel. The purpose of this section is twofold. First, because we have modified $\langle \psi, \psi \rangle_{\mathcal{H}}$ to $\langle \psi, \Sigma \psi \rangle_{\mathcal{H}}$ in the kernel SVM objective function, we can no longer use the standard SVM packages to solve for ψ^* . Therefore, we reformulate the minimization of $\Lambda(\psi, t)$ as quadratic programming that can be solved by available computer packages. Second, in deriving this quadratic programming problem, we gain more insights into the meaning and significance of this modification. As we shall see, by replacing $\langle \psi, \psi \rangle_{\mathcal{H}}$ by $\langle \psi, \Sigma \psi \rangle_{\mathcal{H}}$, we are in effect making SVM invariant with respect to the

marginal distribution of \mathbf{X} . Intuitively, since we are using SVM to make inference about the conditional distribution of $Y|\mathbf{X}$, it is plausible that the procedure does not depend on the marginal distribution of \mathbf{X} .

Let \mathcal{H} be a linear space of functions from $\Omega_{\mathbf{X}}$ to \mathbb{R} spanned by $\mathcal{F}_n = \{\psi_1, \dots, \psi_k\}$. The choice of these functions will be discussed later, but it will ensure $E_n[\psi_i(\mathbf{X})] = 0$, so that $\psi_i(\mathbf{x}) = \psi_i(\mathbf{x}) - E_n\psi_i(\mathbf{X})$. Let

$$(21) \quad \Psi = \begin{pmatrix} \psi_1(\mathbf{X}_1) & \cdots & \psi_k(\mathbf{X}_1) \\ \vdots & \ddots & \vdots \\ \psi_1(\mathbf{X}_n) & \cdots & \psi_k(\mathbf{X}_n) \end{pmatrix}.$$

Then the sample version of the objective function (19) is

$$(22) \quad \hat{\Lambda}(\mathbf{c}) = n^{-1} \mathbf{c}^\top \Psi^\top \Psi \mathbf{c} + \lambda n^{-1} \sum_{i=1}^n [1 - \tilde{Y}_i(\Psi_i^\top \mathbf{c} - t)]^+,$$

where $\Psi_i^\top = (\psi_1(\mathbf{X}_i), \dots, \psi_k(\mathbf{X}_i))$ and $\mathbf{c} \in \mathbb{R}^k$. We minimize $\hat{\Lambda}(\mathbf{c})$ among all \mathbf{c} .

In the following, $\tilde{\mathbf{y}} = (\tilde{y}_1, \dots, \tilde{y}_n)^\top$ and $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\xi} \in \mathbb{R}^n$. The symbol \leq represents componentwise inequality. The symbol \odot represents the Hadamard product between matrices. For a matrix \mathbf{A} of full column rank, $\mathbf{P}_\mathbf{A}$ is the projection $\mathbf{A}(\mathbf{A}^\top \mathbf{A})^{-1} \mathbf{A}^\top$. The symbols $\mathbf{0}$ and $\mathbf{1}$ represent, respectively, the n -dimensional vectors whose entries are 0 and 1.

THEOREM 3. *If \mathbf{c}^* minimizes $\hat{\Lambda}(\mathbf{c})$ over \mathbb{R}^k , then $\mathbf{c}^* = \frac{1}{2}(\Psi^\top \Psi)^{-1} \Psi^\top (\tilde{\mathbf{y}} \odot \boldsymbol{\alpha}^*)$, where $\boldsymbol{\alpha}^*$ is the solution to the quadratic programming problem:*

$$(23) \quad \begin{aligned} & \text{maximize } \mathbf{1}^\top \boldsymbol{\alpha} - \frac{1}{4}(\boldsymbol{\alpha} \odot \tilde{\mathbf{y}})^\top \mathbf{P}_\Psi (\boldsymbol{\alpha} \odot \tilde{\mathbf{y}}) \\ & \text{subject to } \mathbf{0} \leq \boldsymbol{\alpha} \leq \lambda \mathbf{1}, \boldsymbol{\alpha}^\top \tilde{\mathbf{y}} = 0. \end{aligned}$$

Note that the quadratic programming problem (23) differs from that of the standard kernel SVM, where the projection \mathbf{P}_Ψ is replaced by the kernel matrix $\mathbf{K}_n = \{\kappa(i, j) : i, j = 1, \dots, n\}$ for some positive definite bivariate mapping $\kappa : \Omega_{\mathbf{X}} \times \Omega_{\mathbf{X}} \rightarrow \mathbb{R}$. The kernel matrix \mathbf{K}_n uniquely determines the sample estimate of the covariance operator Σ , which bears the information about the shape of the marginal distribution of \mathbf{X} . By replacing \mathbf{K}_n with \mathbf{P}_Ψ , we are, in effect, removing the information about \mathbf{X} . For this reason we call the matrix \mathbf{P}_Ψ an *invariant kernel*.

For the function class \mathcal{H} , we use the reproducing kernel Hilbert space based on the mapping κ . Common choices of κ include the polynomial kernel $\kappa(\mathbf{x}_1, \mathbf{x}_2) = (\mathbf{x}_1^\top \mathbf{x}_2 + c)^r$, where r is a positive integer, and the Gaussian radial kernel $\kappa(\mathbf{x}_1, \mathbf{x}_2) = e^{-\gamma \|\mathbf{x}_1 - \mathbf{x}_2\|^2}$, where $\gamma > 0$. Let

$$(24) \quad \mathcal{H}_\kappa = \{c_0 + c_1 \kappa(\cdot, \mathbf{X}_1) + \dots + c_n \kappa(\cdot, \mathbf{X}_n) : c_0, \dots, c_n \in \mathbb{R}\}$$

with inner product specified by $\langle \kappa(\cdot, \mathbf{a}), \kappa(\cdot, \mathbf{b}) \rangle = \kappa(\mathbf{a}, \mathbf{b})$. In the standard kernel SVM, it is a common practice to use all functions in \mathcal{H}_κ as \mathcal{H} . However, the invariant nature of our kernel, \mathbf{P}_Ψ , determines that we cannot use all those functions,

because if so then \mathbf{P}_Ψ becomes nearly an identity matrix (note that if \mathbf{P}_Ψ were an identity matrix then the objective function in (23) would become independent of $\mathbf{X}_1, \dots, \mathbf{X}_n$). We instead use the principal functions of the linear operator Σ_n , as defined by $\langle \psi_1, \Sigma_n \psi_2 \rangle = \text{cov}_n[\psi_1(\mathbf{X}), \psi_2(\mathbf{X})]$, as our basis \mathcal{F}_n . Here $\text{cov}_n(\cdot, \cdot)$ denotes sample covariance. Let $\mathbf{Q}_n = \mathbf{I}_n - \mathbf{J}_n/n$, where \mathbf{I}_n is the $n \times n$ identity matrix and \mathbf{J}_n is the $n \times n$ matrix whose entries are 1. The next proposition tells us how to find the eigenfunctions of Σ_n . Its proof is easy and omitted.

PROPOSITION 2. *Let $\mathbf{w} = (w_1, \dots, w_n)$, $\psi_{\mathbf{w}} = \sum w_i [\kappa(\mathbf{x}, \mathbf{X}_i) - E_n \kappa(\mathbf{x}, \mathbf{X})]$. The following statements are equivalent:*

1. \mathbf{w} is an eigenvector of the matrix $\mathbf{Q}_n \mathbf{K}_n \mathbf{Q}_n$ with eigenvalue λ ;
2. $\psi_{\mathbf{w}}$ is an eigenfunction of the operator Σ_n with eigenvalue λ/n .

If $\lambda \neq 0$, then either statement implies $(\psi_{\mathbf{w}}(\mathbf{X}_1), \dots, \psi_{\mathbf{w}}(\mathbf{X}_n)) = \lambda \mathbf{w}^\top$.

Although the eigenvectors of $\mathbf{Q}_n \mathbf{K}_n \mathbf{Q}_n$ and the eigenfunctions of Σ_n are similar objects, it is the latter that can be evaluated at any \mathbf{x} , not just the observed $\mathbf{X}_1, \dots, \mathbf{X}_n$. This property is important for prediction. Essentially, we use the first k eigenfunctions ϕ_1, \dots, ϕ_k of Σ_n as the functions in \mathcal{F}_n . This is equivalent to using $\{a_1 \phi_1, \dots, a_k \phi_k\} \equiv \{\psi_1, \dots, \psi_k\}$ for any nonzero a_1, \dots, a_k . We choose a_i to satisfy $a_i (\psi_i(\mathbf{X}_1), \dots, \psi_i(\mathbf{X}_n))^\top = \mathbf{w}_i$, where \mathbf{w}_i is the eigenvector of $\mathbf{Q}_n \mathbf{K}_n \mathbf{Q}_n$, corresponding to its i th eigenvalue λ_i . Thus $a_i = 1/\lambda_i$. With this choice, Ψ is simply $(\mathbf{w}_1, \dots, \mathbf{w}_k)$. The choice of number of basis functions, k , should allow sufficient flexibility but not as large as n ; our experiences indicate that the choice of k in the range $n/3 \sim 2n/3$ works well. We summarize the kernel PSVM estimation procedure as follows.

1. (Optional) Marginally standardize $\mathbf{X}_1, \dots, \mathbf{X}_n$. Let $\hat{\mu}_r$ and $\hat{\sigma}_r^2$ be the sample mean and sample variance X_{1r}, \dots, X_{nr} . Reset X_{ir} to be $(X_{ir} - \hat{\mu}_r)/\hat{\sigma}_r$. The purpose of this step is so that the kernel κ treats different components of \mathbf{X}_i more or less equally. This step can be omitted if the components of \mathbf{X}_i have similar variances.

2. Choose a kernel κ and the number of basis functions k (say $k = n/2$). Compute $\Psi = (\mathbf{w}_1, \dots, \mathbf{w}_k)$ and \mathbf{P}_Ψ from $\mathbf{Q}_n \mathbf{K}_n \mathbf{Q}_n$.

3. Divide the sample according to LVR or OVA, each yielding a set of slices. For each pair of slices, solve the quadratic programming problem in Theorem 3 using the \mathbf{P}_Ψ computed from step 2. This gives coefficient vectors $\mathbf{c}_1^*, \dots, \mathbf{c}_{\tilde{h}}^* \in \mathbb{R}^k$, where $\tilde{h} = h - 1$ for LVR and $\tilde{h} = \binom{h}{2}$ for OVA.

4. Compute the first d eigenvectors, $\mathbf{v}_1, \dots, \mathbf{v}_d$, of the matrix $\sum_{s=1}^{\tilde{h}} \mathbf{c}_s^* \mathbf{c}_s^{*\top}$. Denote the r th component of \mathbf{v}_s as v_{sr} .

5. The s th sufficient predictor evaluated at \mathbf{x} is $v_{s1} \psi_1(\mathbf{x}) + \dots + v_{sk} \psi_k(\mathbf{x})$, where $\psi_r(\mathbf{x}) = \lambda_r^{-1} \sum_{i=1}^n w_{ri} [\kappa(\mathbf{x}, \mathbf{X}_i) - E_n \kappa(\mathbf{x}, \mathbf{X})]$. If step 1 is used, then \mathbf{x} should be marginally standardized by the $\hat{\mu}_r$ and $\hat{\sigma}_r$ computed from that step.

Many computing packages are available to solve the quadratic programming problem in step 3. We use the `ipop` program in the `kernlab` package in R. See Karatzoglou et al. (2004). If the Gaussian radial kernel is used in step 2, then we recommend choosing γ as

$$(25) \quad \gamma = 1/\tau^2, \quad \tau = \frac{1}{\binom{n}{2}} \sum_{i < j, j=2}^n \|\mathbf{X}_i - \mathbf{X}_j\|.$$

Alternatively, we can use the population version of the above quantity,

$$(26) \quad \gamma = 1/(E\|\mathbf{X} - \mathbf{X}'\|^2),$$

where \mathbf{X} and \mathbf{X}' are independent $N(\mathbf{0}, \mathbf{I}_p)$ random vectors. This quantity can be easily evaluated by Monte Carlo. In Section 8, we use (26) for large-scale simulations to avoid repeated evaluations of (25), whereas in Section 9 we use (25) for the real data analysis, where it needs to be calculated only once. Some authors recommend sample median in (25). See Gretton et al. (2005) and Fukumizu, Bach and Jordan (2009). This does not make a significant difference in our examples.

7. Asymptotic analysis of linear PSVM. In this section, we give a comprehensive asymptotic analysis of linear PSVM estimator introduced in Sections 3 and 4. This is developed in three parts. First, we derive the influence function for the normal vector $\hat{\boldsymbol{\psi}}$ based on two slices. In this part, we employ some asymptotic properties of SVM developed recently by Jiang, Zhang and Cai (2008). In the second part, we derive the asymptotic distribution of the linear PSVM estimator, $(\hat{\boldsymbol{v}}_1, \dots, \hat{\boldsymbol{v}}_d)$, defined in Section 4.1. In the third part, we establish the consistency of the order determination criterion introduced in Section 4.2.

7.1. Influence function for support vector machine. The asymptotic results of Jiang, Zhang and Cai (2008) are largely applicable here except for three places: our SVM involves an additional $\boldsymbol{\Sigma}$; our λ is fixed but the λ in their paper depends on n ; they did not derive the explicit form of the hessian matrix—and hence neither the asymptotic variance—but we are interested in the explicit asymptotic distribution. The first two points are minor but the third needs nontrivial additional work. We only consider the case where \tilde{Y} is defined through a partition $\{A_1, A_2\}$ of Ω_Y . Thus, our results only apply to the LVR scheme. The asymptotic analysis the OVA scheme can be carried out similarly, and is omitted.

We first develop some notation. Let $\boldsymbol{\theta} = (\boldsymbol{\psi}^\top, t)^\top$, $\mathbf{Z} = (\mathbf{X}^\top, \tilde{Y})^\top$, $\mathbf{X}^* = (\mathbf{X}^\top, -1)^\top$, $\boldsymbol{\Sigma}^* = \text{diag}(\boldsymbol{\Sigma}, 0)$. Then

$$(27) \quad \boldsymbol{\psi}^\top \boldsymbol{\Sigma} \boldsymbol{\psi} + \lambda[1 - \tilde{Y}(\mathbf{X}^\top \boldsymbol{\psi} - t)]^+ = \boldsymbol{\theta}^\top \boldsymbol{\Sigma}^* \boldsymbol{\theta} - \lambda(1 - \boldsymbol{\theta}^\top \mathbf{X}^* \tilde{Y})^+.$$

We denote this function by $m(\boldsymbol{\theta}, \mathbf{Z})$. Let $\Omega_{\mathbf{Z}}$ be the support of \mathbf{Z} and let $\mathbf{h}: \Theta \times \Omega_{\mathbf{Z}} \rightarrow \mathbb{R}^r$ be a function of $(\boldsymbol{\theta}, \mathbf{Z})$. Let $D_{\boldsymbol{\theta}}$ denote the $(p + 1)$ -dimensional column vector of differential operators $(\partial/\partial\theta_1, \dots, \partial/\partial\theta_{p+1})^\top$. The next theorem gives the gradient of the support vector machine objective function $E[m(\boldsymbol{\theta}, \mathbf{Z})]$.

THEOREM 4. *Suppose, for each $\tilde{y} = -1, 1$, the distribution of $\mathbf{X}|\tilde{Y} = \tilde{y}$ is dominated by the Lebesgue measure and $E(\|\mathbf{X}\|^2) < \infty$. Then*

$$(28) \quad D_{\theta} E[m(\theta, \mathbf{Z})] = (2\boldsymbol{\psi}^{\top} \boldsymbol{\Sigma}, 0)^{\top} - \lambda E[\mathbf{X}^* \tilde{Y} I(1 - \theta^{\top} \mathbf{X}^* \tilde{Y} > 0)].$$

We now present the hessian matrix of support vector machine, which leads to the asymptotic variance of $\hat{\theta}$. To our knowledge, this is the first time that the asymptotic variance is explicitly given. This result is then used to derive the asymptotic distribution of the linear PSVM estimator.

THEOREM 5. *Suppose \mathbf{X} has a convex and open support and its conditional distributions given $\tilde{Y} = 1$ and $\tilde{Y} = -1$ are dominated by the Lebesgue measure. Suppose, moreover:*

1. *for any linearly independent $\boldsymbol{\psi}, \boldsymbol{\delta} \in \mathbb{R}^p$, $\tilde{y} = -1, 1$, and $v \in \mathbb{R}$, the following function is continuous:*

$$u \mapsto E(\mathbf{X}^* | \boldsymbol{\psi}^{\top} \mathbf{X} = u, \boldsymbol{\delta}^{\top} \mathbf{X} = v, \tilde{Y} = \tilde{y}) f_{\boldsymbol{\psi}^{\top} \mathbf{X} | \boldsymbol{\delta}^{\top} \mathbf{X}, \tilde{Y}}(u | v, \tilde{y});$$

2. *for any $i = 1, \dots, p$, and $\tilde{y} = -1, 1$, there is a nonnegative function $c_i(v, \tilde{y})$ with $E[c_i(V, \tilde{Y}) | \tilde{Y}] < \infty$ such that*

$$v E(X_i | \boldsymbol{\psi}^{\top} \mathbf{X} = u, \boldsymbol{\delta}^{\top} \mathbf{X} = v, \tilde{Y} = \tilde{y}) f_{\boldsymbol{\psi}^{\top} \mathbf{X} | \boldsymbol{\delta}^{\top} \mathbf{X}, \tilde{Y}}(u | v, \tilde{y}) \leq c_i(v, \tilde{y});$$

3. *there is a nonnegative function $c_0(v, \tilde{y})$ with $E[c_0(V, \tilde{Y}) | \tilde{Y}] < \infty$ such that $f_{\boldsymbol{\psi}^{\top} \mathbf{X} | \boldsymbol{\delta}^{\top} \mathbf{X}, \tilde{Y}}(u | v, \tilde{y}) \leq c_0(v, \tilde{y})$.*

Then the function $\theta \mapsto D_{\theta} E[m(\theta, \mathbf{Z})]$ is differentiable in all directions with derivative matrix

$$2 \text{diag}(\boldsymbol{\Sigma}, 0) + \lambda \sum_{\tilde{y}=-1,1} P(\tilde{Y} = \tilde{y}) f_{\boldsymbol{\psi}^{\top} \mathbf{X} | \tilde{Y}}(t + \tilde{y} | \tilde{y}) E(\mathbf{X}^* \mathbf{X}^{*\top} | \boldsymbol{\psi}^{\top} \mathbf{X} = t + \tilde{y}).$$

Furthermore, if the function $(\boldsymbol{\psi}, t) \mapsto f_{\boldsymbol{\psi}^{\top} \mathbf{X} | \tilde{Y}}(t + \tilde{y} | \tilde{y}) E(\mathbf{X}^ \mathbf{X}^{*\top} | \boldsymbol{\psi}^{\top} \mathbf{X} = t + \tilde{y})$ is continuous, then $D_{\theta}[m(\theta, \mathbf{Z})]$ is jointly differentiable with respect to θ .*

Joint differentiability and directional differentiability are sometimes referred to as Frechet differentiability and Gateaux differentiability. The latter is generally weaker than the former. In a finite-dimensional space, having continuous directional derivative in all directions implies joint differentiability [Bickel et al. (1993), page 453]. The next theorem gives the influence function for support vector machine.

THEOREM 6. *If the conditions in Theorems 4 and 5 are satisfied, then*

$$\hat{\theta} = \theta_0 - \mathbf{H}^{-1} \{ (2\boldsymbol{\psi}_0^{\top} \boldsymbol{\Sigma}, 0)^{\top} - \lambda E_n[\mathbf{X}^* \tilde{Y} I(1 - \tilde{Y} \theta_0^{\top} \mathbf{X}^* > 0)] \} + o_P(n^{-1/2}),$$

where \mathbf{H} is hessian matrix given by Theorem 5.

The proof is similar to that of Jiang, Zhang and Cai (2008) and is omitted. Alternatively, one can prove it by applying Theorem 5.23 of van der Vaart (1998).

7.2. *Asymptotic distribution of $(\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_d)$.* Consider a fixed division point q_r , where $r \in \{1, \dots, h - 1\}$. Let \tilde{Y}^r be as defined in (12), and $\mathbf{Z}^r = (\mathbf{X}^\top, \tilde{Y}^r)^\top$. Let $\boldsymbol{\theta}_{0r} = (\boldsymbol{\psi}_{0r}^\top, t_{0r})^\top$ be the minimizer of $E[m(\boldsymbol{\theta}, \mathbf{Z}^r)]$, and $\hat{\boldsymbol{\theta}}_r = (\hat{\boldsymbol{\psi}}_r^\top, \hat{t}_r)^\top$ be the minimizer of $E_n[m(\boldsymbol{\theta}, \mathbf{Z}^r)]$. Let \mathbf{H}_r be the hessian matrix of $E[m(\boldsymbol{\theta}, \mathbf{Z}^r)]$, and let \mathbf{F}_r be the first p rows of \mathbf{H}_r^{-1} . By Theorem 6,

$$(29) \quad \hat{\boldsymbol{\psi}}_r = \boldsymbol{\psi}_{0r} - \mathbf{s}_r(\boldsymbol{\theta}_{0r}, \mathbf{Z}^r) + o_P(n^{-1/2}),$$

where $\mathbf{s}_r(\boldsymbol{\theta}, \mathbf{Z}^r) = \mathbf{F}_r[(2\boldsymbol{\psi}^\top \boldsymbol{\Sigma}, 0)^\top - \lambda \mathbf{X}^* \tilde{Y}^r I(1 - \tilde{Y}^r \boldsymbol{\theta}^\top \mathbf{X}^*) > 0]$. Let

$$\hat{\mathbf{M}}_n = \sum_{r=1}^{h-1} \hat{\boldsymbol{\psi}}_r \hat{\boldsymbol{\psi}}_r^\top, \quad \mathbf{M}_0 = \sum_{r=1}^{h-1} \boldsymbol{\psi}_{0r} \boldsymbol{\psi}_{0r}^\top.$$

For a matrix $\mathbf{A} \in \mathbb{R}^{r_1 \times r_2}$, let $\mathbf{K}_{r_1, r_2} \in \mathbb{R}^{r_1 r_2 \times r_1 r_2}$ be the commutation matrix defined by the relation $\mathbf{K}_{r_1, r_2} \text{vec}(\mathbf{A}) = \text{vec}(\mathbf{A}^\top)$. See Magnus and Neudecker (1979). Two properties of \mathbf{K}_{r_1, r_2} that will prove useful for our purpose are that $\mathbf{K}_{r_1, r_2} = \mathbf{K}_{r_2, r_1}^\top$ and that for any $\mathbf{B} \in \mathbb{R}^{r_3 \times r_4}$,

$$(30) \quad \mathbf{A} \otimes \mathbf{B} = \mathbf{K}_{r_1, r_3}(\mathbf{B} \otimes \mathbf{A})\mathbf{K}_{r_4, r_2}.$$

We now present the asymptotic distribution of $\hat{\mathbf{M}}_n$.

THEOREM 7. *Under the assumptions in Theorems 4 and 5, $\sqrt{n} \text{vec}(\hat{\mathbf{M}}_n - \mathbf{M}_0)$ converges to multivariate normal with mean $\mathbf{0}$ and variance*

$$(\mathbf{I}_{p^2} + \mathbf{K}_{p,p}) \sum_{r=1}^{h-1} \sum_{t=1}^{h-1} (\boldsymbol{\psi}_{0r} \boldsymbol{\psi}_{0t}^\top \otimes \boldsymbol{\Lambda}_{rt})(\mathbf{I}_{p^2} + \mathbf{K}_{p,p}),$$

where $\boldsymbol{\Lambda}_{rt} = E[\mathbf{s}_r(\boldsymbol{\theta}_{0r}, \mathbf{Z}^r) \mathbf{s}_t^\top(\boldsymbol{\theta}_{0t}, \mathbf{Z}^t)]$.

This result leads directly to the asymptotic distribution of $\hat{\mathbf{V}} = (\hat{\mathbf{v}}_1, \dots, \hat{\mathbf{v}}_d)$. Since, by Theorem 1, $\text{span}(\mathbf{M}_0) \subseteq \mathcal{S}_{Y|X}$, we have $\text{rank}(\mathbf{M}_0) \leq d$. We make the working assumption that $\text{rank}(\mathbf{M}_0) = d$. This means we exclude the situations where the regression surface is symmetric about the origin. Since \mathbf{M}_0 is positive semi-definite, it has the spectral decomposition $\mathbf{U}\mathbf{D}\mathbf{U}^\top$, where \mathbf{U} is a $p \times d$ matrix whose columns are the eigenvectors of \mathbf{M}_0 corresponding to nonzero eigenvalues, and \mathbf{D} is a $d \times d$ diagonal matrix with diagonal elements being the nonzero eigenvalues. The following corollary is a direct consequence of Theorem 7 and Bura and Pfeiffer (2008). Its proof is omitted.

COROLLARY 1. *Under the assumptions in Theorems 4 and 5 and $\text{rank}(\mathbf{M}_0) = d$, $\sqrt{n} \text{vec}(\hat{\mathbf{V}} - \mathbf{V}_0) \xrightarrow{D} N(\mathbf{0}, \boldsymbol{\Upsilon})$, where $\boldsymbol{\Upsilon}$ is the $pd \times pd$ matrix*

$$(\mathbf{D}^{-1} \mathbf{U}^\top \otimes \mathbf{I}_p)(\mathbf{I}_{p^2} + \mathbf{K}_{p,p}) \sum_{r=1}^{h-1} \sum_{t=1}^{h-1} (\boldsymbol{\psi}_{0r} \boldsymbol{\psi}_{0t}^\top \otimes \boldsymbol{\Lambda}_{rt})(\mathbf{I}_{p^2} + \mathbf{K}_{p,p})(\mathbf{U}\mathbf{D}^{-1} \otimes \mathbf{I}_p).$$

It is possible to refine the PSVM estimator by introducing weights to $\hat{\mathbf{M}}_n$. Take the LVR scheme for example. Let $\Psi = (\hat{\psi}_1, \dots, \hat{\psi}_{h-1})$. Let \mathbf{A} be an $h - 1$ by $h - 1$ matrix. Rather than working with $\hat{\mathbf{M}}_n$, we could base the spectral decomposition on a weighted matrix $\hat{\mathbf{M}}_n(\mathbf{A}) = \Psi^\top \mathbf{A} \Psi$. Let $\hat{\mathbf{v}}(\mathbf{A}) = (\hat{v}_1(\mathbf{A}), \dots, \hat{v}_d(\mathbf{A}))$ be the first d eigenvectors of $\hat{\mathbf{M}}_n(\mathbf{A})$. One way to determine the optimal \mathbf{A} is by minimizing a real-valued monotone function (say trace) of the asymptotic variance matrix of $\text{vec}[\hat{\mathbf{v}}(\mathbf{A})]$, which can be extracted from the asymptotic distribution. This type of argument was used in Li (2000, 2001) to construct optimal estimating equations. Alternatively, one can develop an optimal procedure using the minimum distance approach introduced by Cook and Ni (2005). We leave these to future research.

7.3. *Consistency of the BIC-type criterion.* In the following, we say a sequence of random variables W_n converges in probability to infinity ($W_n \xrightarrow{P} \infty$) if, for any $K > 0$, $\lim_{n \rightarrow \infty} P(|W_n| > K) = 1$. Let \hat{d} is the maximizer of $G_n(k)$ over $\{0, \dots, p\}$ as defined in Section 4.2.

THEOREM 8. *Suppose $P(c_1(n) > 0) = 1$, $c_1(n) \xrightarrow{P} 0$, $n^{1/2}c_1(n) \xrightarrow{P} \infty$, and $c_2(k)$ is an increasing function of k . Under the conditions in Theorems 1, 4 and 5 and $\text{rank}(\mathbf{M}_0) = d$, we have $\lim_{n \rightarrow \infty} P(\hat{d} = d) = 1$.*

Note that we have again made the working assumption $\text{rank}(\mathbf{M}_0) = d$. However, even when this assumption is violated the theorem still holds with d replaced by the rank of \mathbf{M}_0 .

8. Simulation studies. In this section, we compare the linear and kernel PSVM with four other methods based on the idea of inverse regression: SIR, the sliced average variance estimator [SAVE; Cook and Weisberg (1991)], directional regression [DR; Li and Wang (2007)], and kernel sliced inverse regression [Wu (2008)]. We also investigate the performance of the CVBIC for order determination.

8.1. *Linear dimension reduction.* We use the following models:

$$\text{Model I: } Y = X_1/[0.5 + (X_2 + 1)^2] + \sigma\varepsilon,$$

$$\text{Model II: } Y = X_1(X_1 + X_2 + 1) + \sigma\varepsilon,$$

$$\text{Model III: } Y = (X_1^2 + X_2^2)^{1/2} \log(X_1^2 + X_2^2)^{1/2} + \sigma\varepsilon,$$

where $\mathbf{X} \sim N(0, \mathbf{I}_p)$, $p = 10, 20, 30$, $\varepsilon \sim N(0, 1)$ and $\sigma = 0.2$. The sample size n is taken to be 100. The first two models, which are taken from Li (1991), are asymmetric about 0, but the last one is symmetric about 0. As we have discussed in Section 4.3, linear PSVM, like SIR, does not work when the regression surface

is symmetric about 0. The first two examples show how the linear PSVM compares with other methods in the situations where it works. The purpose of the last model is to provide a benchmark of error when it fails, so that we can gauge how the kernel PSVM improves the situation in the next comparison.

To evaluate the performance of each method, we use the distance measure suggested by Li, Zha and Chiaromonte (2005). Specifically, let \mathcal{S}_1 and \mathcal{S}_2 be two subspaces of \mathbb{R}^p . Then

$$(31) \quad \text{dist}(\mathcal{S}_1, \mathcal{S}_2) = \|\mathbf{P}_{\mathcal{S}_1} - \mathbf{P}_{\mathcal{S}_2}\|,$$

where $\mathbf{P}_{\mathcal{S}_1}$ and $\mathbf{P}_{\mathcal{S}_2}$ are orthogonal projections on to \mathcal{S}_1 and \mathcal{S}_2 , and $\|\cdot\|$ is a matrix norm. In the following the Frobenius norm is used.

For SAVE and DR, we use $h = 4$ slices, and for SIR, we use $h = 8$ slices, having roughly the same number of points. Our choices of h are in line with the usual practice in the SDR literature for such a sample size. For methods such as SAVE and DR that involve the second-order inverse moment, h is suggested to be chosen smaller than that for methods such as SIR which only involve the first-order inverse moment [Li and Zhu (2007)]. For the linear PSVM, the cost λ is taken to be 1. The number of division points (q_r) is 20. We have tried some other numbers of division points and obtained very similar results. In general, our experiences suggest that a relatively large number of dividing points is preferable. The results are presented in Table 1. The entries are of the form $a(b)$ where a is the mean, and b is the standard deviation of the distance criterion (31) calculated from 200 simulated samples. The last row in Table 1 records the CPU time (in seconds) each method uses for Model I with $p = 10$ (on a Dell OptiPlex 745 desktop computer with speed 2.66 GHz).

TABLE 1
Estimated means and simulation standard errors (in parentheses) of the distance measure (31) and mean computation times (in second) of linear sufficient dimension reduction methods

Models	p	SIR	SAVE	DR	Linear PSVM
I	10	0.84 (0.22)	1.55 (0.19)	1.02 (0.23)	0.65 (0.17)
	20	1.14 (0.18)	1.93 (0.05)	1.32 (0.17)	0.93 (0.16)
	30	1.31 (0.14)	1.96 (0.03)	1.48 (0.11)	1.17 (0.14)
II	10	1.20 (0.27)	1.43 (0.16)	1.17 (0.23)	0.85 (0.25)
	20	1.51 (0.19)	1.72 (0.15)	1.46 (0.14)	1.26 (0.23)
	30	1.67 (0.16)	1.84 (0.12)	1.63 (0.12)	1.58 (0.17)
III	10	1.80 (0.13)	0.87 (0.21)	0.85 (0.20)	1.65 (0.16)
	20	1.89 (0.08)	1.46 (0.20)	1.45 (0.20)	1.85 (0.10)
	30	1.93 (0.05)	1.72 (0.12)	1.71 (0.12)	1.93 (0.05)
Time		0.03	0.01	0.03	0.16

Table 1 shows that the linear PSVM consistently performs better than the other methods in all cases for models I and II. The intuition behind this improvement is explained in Section 4.3. Also, as expected, the linear PSVM and SIR do not perform well for Model III because of the symmetry of the regression function. However, as we will see in the next comparison, this defect is no longer present in the kernel PSVM. The linear PSVM requires more computing time than the classical methods, mainly because it needs to process more dividing points, and, for each dividing point, the full data (rather than a slice of data) are processed.

8.2. *Nonlinear dimension reduction.* As we have mentioned, Model III is symmetric about 0, and the linear PSVM fails. To a certain degree, the shape of regression surface of Model II is also symmetric about 0. We now use these two models to investigate the performance of the kernel PSVM for nonlinear sufficient dimension reduction. In terms of linear dimension reduction, Model III has two sufficient predictors, X_1, X_2 , but in terms of nonlinear dimension reduction, it has only one sufficient predictor, $(X_1^2 + X_2^2)^{1/2}$, or any monotone function of it. The kernel PSVM is designed to recover a monotone transformation of $(X_1^2 + X_2^2)^{1/2}$ without having to assume any regression model. In doing so, it solves two problems at one stroke—further reducing the dimension from 2 to 1, and avoiding the difficulty of SIR in dealing with symmetric responses.

To illustrate the idea, in Figure 3 we present the 2-D and 3-D scatter plots for Y versus the nonlinear and linear predictors obtained by different methods. The upper left panel is the 2-D scatter plot for Y versus the true nonlinear predictor $(X_1^2 + X_2^2)^{1/2}$; the upper right panel is the 2-D scatter plot of Y versus the first kernel PSVM predictor; the lower panels are 3-D scatter plots for Y versus the first two predictors from SAVE and DR. We can see that all three methods capture the right shape of the regression function, but kernel PSVM only requires one predictor and its sufficient plot appears sharper (bearing in mind that the upper right panel only has to resemble a *monotone transformation* of the upper left panel).

To make a more precise comparison, we need to design a new criterion that can compare one nonlinear predictor with two linear predictors; the criterion (31) is no longer suitable for this purpose. Since the nonlinear sufficient predictor estimates a monotone function of $(X_1^2 + X_2^2)^{1/2}$, we use the absolute value of Spearman’s correlation to measure their closeness, which is invariant under monotone transformation [Kutner, Nachtsheim and Neter (2004), page 87]. To measure the closeness between two linear predictors and the true nonlinear predictor $(X_1^2 + X_2^2)^{1/2}$, let $(U_{11}, \dots, U_{1n}), (U_{21}, \dots, U_{2n})$ represent the two linear predictors obtained by SAVE or DR. These predictors estimate linear combinations of X_{1i}, X_{2i} but do not specify X_{1i}, X_{2i} themselves. We therefore regress $T_i = X_{1i}^2 + X_{2i}^2$ on

$$\{(1, U_{1i}, U_{2i}, U_{1i}^2, U_{2i}^2) : i = 1, \dots, n\}.$$

If U_{1i} and U_{2i} are (linearly independent) linear combinations of X_{1i} and X_{2i} , then this regression is guaranteed to recover the true predictor T_i regardless of the




FIG. 3. Comparison between linear and nonlinear sufficient dimension reduction methods. Upper left panel: true nonlinear predictor $\sqrt{X_1^2 + X_2^2}$ versus Y ; upper right panel: first (nonlinear) PSVM predictor versus Y ; lower left: first two SAVE predictors versus Y ; lower right panel: first two DR predictors versus Y .

specific form of the linear combinations. Let $\hat{T}_1, \dots, \hat{T}_n$ be the fitted responses of this regression. We use the absolute values of Spearman’s correlation between T_i and \hat{T}_i to measure the performance of SAVE and DR.

We compute these numbers for 200 simulation samples, and tabulate their means and standard deviations in Table 2. Note that large numbers represent better performance, and all numbers are between 0 and 1. The SAVE and DR estimators are computed in exactly the same way as in the linear dimension reduction comparison. For the kernel PSVM, the cost is 1, the number of division points is still 20, the kernel is the Gaussian radial basis, and the number of principal eigenfunctions of Σ_n is taken to be 60. The parameter γ is calculated by (26), which are approximately 0.0526, 0.0257 and 0.0169 for $p = 10, 20, 30$, respectively. We see that the kernel PSVM actually performs better than SAVE and DR, even though it uses only one predictor. It also performs better than KSIR. Moreover, the accuracy

TABLE 2
Estimated means and simulation standard errors (in parentheses) of Spearman correlations of linear and nonlinear sufficient dimension reduction

<i>p</i>	Model II				Model III			
	SAVE	DR	KSIR	KPSVM	SAVE	DR	KSIR	KPSVM
10	0.53 (0.13)	0.67 (0.11)	0.88 (0.07)	0.92 (0.02)	0.79 (0.09)	0.79 (0.08)	0.89 (0.05)	0.90 (0.02)
20	0.37 (0.14)	0.53 (0.09)	0.68 (0.17)	0.86 (0.03)	0.56 (0.11)	0.57 (0.11)	0.59 (0.18)	0.81 (0.03)
30	0.30 (0.13)	0.43 (0.10)	0.55 (0.23)	0.83 (0.04)	0.47 (0.11)	0.48 (0.11)	0.42 (0.21)	0.77 (0.04)

of the kernel PSVM remains reasonably high for larger *p*, where the accuracies of SAVE, DR, and KSIR drop considerably.

8.3. *Estimation of structural dimension.* We now investigate the performance of the CVBIC order-determination procedure for a variety of combinations of (*p, d, n*). We still use models I and II, but, to include different *d*, we add the following models which both have *d* = 1:

$$\text{Model IV: } Y = X_1/[0.5 + (X_1 + 1)^2] + \sigma\varepsilon,$$

$$\text{Model V: } Y = X_1(2X_1 + 1) + \sigma\varepsilon.$$

These are derived from models I and II by replacing *X*₂ by *X*₁.

We apply CVBIC in conjunction with PSVM to Models I, II, IV and V, with (*d, n, p*) ranging over the set {1, 2} × {200, 300, 400, 500} × {10, 20, 30}. The training and testing sample sizes are *n*₁ = *n*₂ = *n*/2. We take 20 dividing points *q_r* as equally-spaced sample quantiles of $\hat{Y}_1, \dots, \hat{Y}_{n_1}$. As a comparison we also apply the order-determination procedure for SIR based on Theorem 5.1 of Li (1991) with significant level $\alpha = 0.05$. The results are presented in Table 3, where the entries are the percentage of correct estimation of *d* out of 200 simulated samples for each of the 48 combinations of (model, *p, n*). Table 3 shows that CVBIC works very well, with percentage of correct estimation reaching as high as 100% for sample size of 200 (training sample size 100). In almost all cases, PSVM compares favorably with SIR for order determination. Also clear from the table is the trend of increasing accuracy for both methods as *n* increases.

9. Application and further discussions. We now compare the kernel PSVM with SIR, SAVE, and DR in a real data analysis concerning recognition of vowels. The data can be found in the UCI Machine Learning Repository (<http://archive.ics.uci.edu/ml/datasets>). The response variable *Y* is a categorical variable of 11 levels,

TABLE 3
Rate of correct order determination by SIR and PSVM in %

Model	d	p	$n = 200$		$n = 300$		$n = 400$		$n = 500$	
			SIR	PSVM	SIR	PSVM	SIR	PSVM	SIR	PSVM
I	1	10	92	96	92	100	97	100	98	100
		20	80	82	95	96	96	100	94	100
		30	65	54	92	94	94	98	96	100
II	2	10	67	80	86	85	97	90	98	94
		20	36	64	66	84	85	86	96	84
		30	22	32	55	77	73	84	88	80
IV	1	10	39	82	60	84	71	93	75	95
		20	28	78	42	74	54	76	68	80
		30	15	78	33	84	44	78	60	80
V	2	10	93	100	96	100	96	100	97	100
		20	94	98	96	100	96	100	96	100
		30	96	98	96	99	96	100	92	100

representing different vowel sounds. The predictor \mathbf{X} is a 10-dimensional vector describing the features of a sound. For clear presentation, we select only three vowels: the sounds in *heed*, *head* and *hud*, with training and testing sample sizes being 144 and 126, respectively.

For each dimension reduction method, we find a set of sufficient predictors from the training data, and evaluate them at the testing set, resulting in a sufficient plot for the testing set. Given that the testing data are independent of the training data from which the sufficient predictors is derived, the degree of separation of the vowels in the sufficient plot objectively reflects the discriminating power of a dimension reduction method. The four scatter plots in Figure 4 present the first two predictors found by SIR (upper left panel), SAVE (upper right panel), DR (lower left panel), and the kernel PSVM (lower right panel). For the kernel PSVM, the OVA scheme is used. The basis functions are the first 40 eigenfunctions of the operator Σ_n derived from the Gaussian radial kernel, whose parameter γ is calculated by (25). The cost λ is 1. We have varied the number of eigenfunctions (from 10 to 60) and the cost (from 0.5 to 20), but they do not seem to result in significant difference in the degree of separation in the test data.

From Figure 4, we see that the kernel PSVM achieves much better separation of the three vowels in the test data than the other three methods. The second best performer is DR, followed by SIR and SAVE. It is also interesting to note that the various degrees of separation are also reflected in the sufficient plots; that is, the distance between *heed* and *hud* is larger than those between *heed* and *head*, and *head* and *hud*.




FIG. 4. First two predictors based on SIR, SAVE, DR, and the kernel PSVM plotted for the vowel recognition testing data set. Green, red and blue colors indicate the vowel sounds in heed, head, hud.

We would like to comment that classification, though important, is not the sole purpose for sufficient dimension reduction, and that linear and nonlinear sufficient dimension reductions have their own strengths in reducing, discriminating, visualizing, and interpreting high-dimensional data. To illuminate the point, consider an example where variation, rather than location, is the differentiating characteristic. Let Y be a bernoulli variable with $P(Y = 1) = P(Y = 0) = 1/2$ and

$$(\mathbf{X}|Y = y) \sim N\left(\mathbf{0}, \begin{pmatrix} \sigma^2(y)\mathbf{I}_2 & 0 \\ 0 & \mathbf{I}_{p-2} \end{pmatrix}\right),$$

where $\sigma^2(0) = 1$ and $\sigma^2(1) = 10$. Let $(\mathbf{X}_1, Y_1), \dots, (\mathbf{X}_n, Y_n)$ be a sample from this model, where $n = 200$ and $p = 10$. For simplicity, we fix the number of cases of $Y = 1$ at $n/2$, because this has no bearing on our problem. In this case, the central subspace is $\text{span}(\mathbf{e}_1, \mathbf{e}_2)$, where $\mathbf{e}_i = (0, \dots, 1, \dots, 0)^\top$ with the 1 occupying the i th position.

We apply SAVE and the kernel PSVM and the results are shown in Figure 5, where the top panel shows the scatter plot for the true sufficient predictors X_1




FIG. 5. Variation as the differentiating characteristic. Blue \circ represents the $Y = 0$ cases and red $+$ represents the $Y = 1$ cases.

and X_2 , the lower left panel shows the first two SAVE predictors, and the lower right panel shows the boxplot of a single kernel PSVM predictor. Since for a single variable we cannot produce a scatter plot, for clarity we use a boxplot to represent the predictor. The value of the kernel PSVM predictor is represented by the height in the boxplot; the two boxes represents the two groups. All three plots are based on the testing data. What is interesting is that kernel PSVM in some sense “translates” the difference in variation into the difference in location. The intuitive reason is that there is a quadratic—and hence variance—component in the kernel mapping, but in the mapped high-dimensional space the variance component is treated as an augmented part of feature vector [as in (x, x^2)]. Of course this is only a simplification of the situation: there is still significant difference in variation in the kernel PSVM predictor for the two groups.

In this case, linear dimension reduction methods such as SAVE have a definite advantage, both for their clear separation of variation and for their good inter-




FIG. 6. Degrees of separation by SAVE (upper panels) and kernel PSVM (lower panels) for higher dimensions: $p = 60$ (left panels), $p = 80$ (middle panels) and $p = 100$ (right panels).

pretability. In the meantime, this example also shows that kernel PSVM is capable of differentiating variation, to the degree comparable to SAVE, but its interpretability is not as direct as SAVE.

Another desirable feature of the kernel PSVM is that its accuracy is more stable than the classical methods as the dimension p increases. Figure 6 shows the sufficient predictors derived from SAVE and kernel PSVM for $p = 60, 80, 100$ (from left to right). The upper panels are the scatter plots for the first two SAVE predictors, and the lower panels are the boxplots representing the single kernel PSVM predictor. Again, all plots are based on testing data. We see that SAVE gradually loses its discriminating power as p is increased to 100, whereas the discriminating power of kernel PSVM remains reasonably strong.

Acknowledgments. We are very grateful to three referees and an Associate Editor, whose many useful comments and suggestions greatly broadened and deepened an earlier version of this work.

REFERENCES

- ARONSAJN, N. (1950). Theory of reproducing kernels. *Trans. Amer. Math. Soc.* **68** 337–404. [MR0051437](#)
- ARTEMIOU, A. A. (2010). Topics on supervised and unsupervised dimension reduction. Ph.D. thesis, Pennsylvania State Univ., University Park, PA.
- BICKEL, P., KLAASSEN, C. A. J., RITOV, Y. and WELLNER, J. (1993). *Efficient and Adaptive Inference in Semi-Parametric Models*. Johns Hopkins Univ. Press, Baltimore.
- BURA, E. and PFEIFFER, R. (2008). On the distribution of the left singular vectors of a random matrix and its applications. *Statist. Probab. Lett.* **78** 2275–2280. [MR2462662](#)
- CONWAY, J. B. (1990). *A Course in Functional Analysis*, 2nd ed. *Graduate Texts in Mathematics* **96**. Springer, New York. [MR1070713](#)
- 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* 18–25. Amer. Statist. Assoc., Alexandria, VA.
- COOK, R. D. (1996). Graphics for regressions with a binary response. *J. Amer. Statist. Assoc.* **91** 983–992. [MR1424601](#)
- COOK, R. D. (1998). *Regression Graphics: Ideas for Studying Regressions Through Graphics*. Wiley, New York. [MR1645673](#)
- COOK, R. D. (2007). Fisher lecture: Dimension reduction in regression. *Statist. Sci.* **22** 1–26. [MR2408655](#)
- COOK, R. D. and FORZANI, L. (2008). Principal fitted components for dimension reduction in regression. *Statist. Sci.* **23** 485–501. [MR2530547](#)
- COOK, R. D. and LI, B. (2002). Dimension reduction for conditional mean in regression. *Ann. Statist.* **30** 455–474. [MR1902895](#)
- COOK, R. D. and NI, L. (2005). Sufficient dimension reduction via inverse regression: A minimum discrepancy approach. *J. Amer. Statist. Assoc.* **100** 410–428. [MR2160547](#)
- COOK, R. D. and WEISBERG, S. (1991). Discussion of “Sliced inverse regression for dimension reduction,” by K.-C. Li. *J. Amer. Statist. Assoc.* **86** 316–342.
- EATON, M. L. (1986). A characterization of spherical distributions. *J. Multivariate Anal.* **20** 272–276. [MR0866075](#)
- FUKUMIZU, K., BACH, F. R. and JORDAN, M. I. (2004). Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *J. Mach. Learn. Res.* **5** 73–99. [MR2247974](#)
- FUKUMIZU, K., BACH, F. R. and JORDAN, M. I. (2009). Kernel dimension reduction in regression. *Ann. Statist.* **37** 1871–1905. [MR2533474](#)
- FUNG, W. K., HE, X., LIU, L. and SHI, P. (2002). Dimension reduction based on canonical correlation. *Statist. Sinica* **12** 1093–1113. [MR1947065](#)
- GRETTON, A., BOUSQUET, O., SMOLA, A. and SCHÖLKOPF, B. (2005). Measuring statistical dependence with Hilbert–Schmidt norms. In *16th International Conference on Algorithmic Learning Theory* (S. Jain, H. U. Simon and E. Tomita, eds.). *Lecture Notes in Computer Science* **3734** 63–77. Springer, Berlin. [MR2255909](#)
- HALL, P. and LI, K.-C. (1993). On almost linearity of low-dimensional projections from high-dimensional data. *Ann. Statist.* **21** 867–889. [MR1232523](#)
- HSING, T. and REN, H. (2009). An RKHS formulation of the inverse regression dimension-reduction problem. *Ann. Statist.* **37** 726–755. [MR2502649](#)
- JIANG, B., ZHANG, X. and CAI, T. (2008). Estimating the confidence interval for prediction errors of support vector machine classifiers. *J. Mach. Learn. Res.* **9** 521–540. [MR2417245](#)
- KARATZOGLOU, A. and MEYER, D. (2006). Support vector machines in R. *J. Stat. Softw.* **15** 9.
- KARATZOGLOU, A., SMOLA, A., HORNİK, K. and ZEILEIS, A. (2004). Kernlab—an S4 package for kernel methods in R. *J. Stat. Software* **11** 9.

- KURDILA, A. J. and ZABARANKIN, M. (2005). *Convex Functional Analysis*. Birkhäuser, Basel. [MR2145612](#)
- KUTNER, M. H., NACHTSHEIM, C. J. and NETER, J. (2004). *Applied Linear Regression Models*, 4th ed. McGraw-Hill/Irwin, Boston.
- LI, K.-C. (1991). Sliced inverse regression for dimension reduction (with discussion). *J. Amer. Statist. Assoc.* **86** 316–342. [MR1137117](#)
- LI, K.-C. (1992). On principal Hessian directions for data visualization and dimension reduction: Another application of Stein's lemma. *J. Amer. Statist. Assoc.* **87** 1025–1039. [MR1209564](#)
- LI, B. (2000). Nonparametric estimating equations based on a penalized information criterion. *Canad. J. Statist.* **28** 621–639. [MR1793116](#)
- LI, B. (2001). On quasi likelihood equations with non-parametric weights. *Scand. J. Stat.* **28** 577–602. [MR1876501](#)
- LI, B. and DONG, Y. (2009). Dimension reduction for nonelliptically distributed predictors. *Ann. Statist.* **37** 1272–1298. [MR2509074](#)
- LI, K.-C. and DUAN, N. (1989). Regression analysis under link violation. *Ann. Statist.* **17** 1009–1052. [MR1015136](#)
- LI, B. and WANG, S. (2007). On directional regression for dimension reduction. *J. Amer. Statist. Assoc.* **102** 997–1008. [MR2354409](#)
- LI, B., ZHA, H. and CHIAROMONTE, F. (2005). Contour regression: A general approach to dimension reduction. *Ann. Statist.* **33** 1580–1616. [MR2166556](#)
- LI, Y. and ZHU, L.-X. (2007). Asymptotics for sliced average variance estimation. *Ann. Statist.* **35** 41–69. [MR2332268](#)
- LOH, W.-Y. (2002). Regression trees with unbiased variable selection and interaction detection. *Statist. Sinica* **12** 361–386. [MR1902715](#)
- MAGNUS, J. R. and NEUDECKER, H. (1979). The commutation matrix: Some properties and applications. *Ann. Statist.* **7** 381–394. [MR0520247](#)
- SCHWARZ, G. (1978). Estimating the dimension of a model. *Ann. Statist.* **6** 461–464. [MR0468014](#)
- VAN DER VAART, A. W. (1998). *Asymptotic Statistics*. *Cambridge Series in Statistical and Probabilistic Mathematics* **3**. Cambridge Univ. Press, Cambridge. [MR1652247](#)
- VAPNIK, V. N. (1998). *Statistical Learning Theory*. Wiley, New York. [MR1641250](#)
- WANG, Y. (2008). Nonlinear dimension reduction in feature space. Ph.D. thesis, Pennsylvania State Univ., University Park, PA.
- WANG, Q. and YIN, X. (2008). A nonlinear multi-dimensional variable selection method for high dimensional data: Sparse MAVE. *Comput. Statist. Data Anal.* **52** 4512–4520. [MR2432477](#)
- WEIDMANN, J. (1980). *Linear Operators in Hilbert Spaces*. *Graduate Texts in Mathematics* **68**. Springer, New York. [MR0566954](#)
- WU, H.-M. (2008). Kernel sliced inverse regression with applications to classification. *J. Comput. Graph. Statist.* **17** 590–610. [MR2528238](#)
- WU, Q., LIANG, F. and MUKHERJEE, S. (2008). Regularized sliced inverse regression for kernel models. Technical report, Duke Univ., Durham, NC.
- XIA, Y., TONG, H., LI, W. K. and ZHU, L.-X. (2002). An adaptive estimation of dimension reduction space. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 363–410. [MR1924297](#)
- YEH, Y.-R., HUANG, S.-Y. and LEE, Y.-Y. (2009). Nonlinear dimension reduction with kernel sliced inverse regression. *IEEE Transactions on Knowledge and Data Engineering* **21** 1590–1603.
- YIN, X. and COOK, R. D. (2002). Dimension reduction for the conditional k th moment in regression. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 159–175. [MR1904698](#)
- YIN, X., LI, B. and COOK, R. D. (2008). Successive direction extraction for estimating the central subspace in a multiple-index regression. *J. Multivariate Anal.* **99** 1733–1757. [MR2444817](#)

ZHU, L., MIAO, B. and PENG, H. (2006). On sliced inverse regression with high-dimensional covariates. *J. Amer. Statist. Assoc.* **101** 630–643. [MR2281245](#)

B. LI
DEPARTMENT OF STATISTICS
PENNSYLVANIA STATE UNIVERSITY
326 THOMAS BUILDING
UNIVERSITY PARK, PENNSYLVANIA 16802
USA
E-MAIL: bing@stat.psu.edu

A. ARTEMIOU
DEPARTMENT OF MATHEMATICAL SCIENCES
MICHIGAN TECHNOLOGICAL UNIVERSITY
FISHER HALL, ROOM 306
1400 TOWNSEND DRIVE
HOUGHTON, MICHIGAN 49931
USA
E-MAIL: aaartemio@mtu.edu

L. LI
DEPARTMENT OF STATISTICS
NORTH CAROLINA STATE UNIVERSITY
2311 STINSON DRIVE
CAMPUS BOX 8203
RALEIGH, NORTH CAROLINA 27695-8203
USA
E-MAIL: li@stat.ncsu.edu