

SUFFICIENT DIMENSION REDUCTION BASED ON AN ENSEMBLE OF MINIMUM AVERAGE VARIANCE ESTIMATORS

BY XIANGRONG YIN¹ AND BING LI²

University of Georgia and Pennsylvania State University

We introduce a class of dimension reduction estimators based on an ensemble of the minimum average variance estimates of functions that characterize the central subspace, such as the characteristic functions, the Box–Cox transformations and wavelet basis. The ensemble estimators exhaustively estimate the central subspace without imposing restrictive conditions on the predictors, and have the same convergence rate as the minimum average variance estimates. They are flexible and easy to implement, and allow repeated use of the available sample, which enhances accuracy. They are applicable to both univariate and multivariate responses in a unified form. We establish the consistency and convergence rate of these estimators, and the consistency of a cross validation criterion for order determination. We compare the ensemble estimators with other estimators in a wide variety of models, and establish their competent performance.

1. Introduction. Sufficient dimension reduction [Li (1991, 1992), Cook and Weisberg (1991), Cook (1994, 1996)] is a methodology for reducing the dimension of predictors while preserving its regression relation with a response. The reduction is achieved by projecting the raw predictors on to a lower-dimensional subspace. Let (X, Y) be a pair of random vectors of dimensions p and s . In this section we tentatively assume $s = 1$. Let \mathcal{S} denote a subspace of \mathbb{R}^p , and let $P_{\mathcal{S}}$ denote the orthogonal projection on to \mathcal{S} . If Y and X are independent conditioning on $P_{\mathcal{S}}X$, then $P_{\mathcal{S}}X$ can be used as the predictor without loss of regression information. Such subspaces \mathcal{S} are called dimension reduction subspaces. The intersection of all such subspaces \mathcal{S} , if itself satisfies the conditional independence, is called the central subspace [Cook (1994)], and is denoted by $\mathcal{S}_{Y|X}$. Under mild conditions [Cook (1996), Yin, Li and Cook (2008)], the central subspace is well defined and is unique.

A closely related concept is the notion of central mean subspace [Cook and Li (2002)], which is the intersection of all subspaces such that $E(Y|X) = E(Y|P_{\mathcal{S}}X)$. This subspace is written as $\mathcal{S}_{E(Y|X)}$. Evidently, if conditional distribution of Y

Received April 2011; revised October 2011.

¹Supported in part by NSF Grant DMS-08-06120.

²Supported in part by NSF Grants DMS-07-04621 and DMS-08-06058.

MSC2010 subject classifications. Primary 62G08, 62B05; secondary 62H12.

Key words and phrases. Central mean subspace, central subspace, characterizing family, characteristic function, gradient estimation, projective resampling, wavelets.

given X depends on X only through $E(Y|X)$, then $\mathcal{S}_{Y|X} = \mathcal{S}_{E(Y|X)}$. However, if this conditional distribution also depends on other functions of X , such as $\text{var}(Y|X)$, then $\mathcal{S}_{E(Y|X)}$ is a proper subspace of $\mathcal{S}_{Y|X}$. Cook and Li (2002) noted that several previously introduced dimension reduction methods, such as the ordinary least squares [Li and Duan (1989), Duan and Li (1991)] and principal Hessian directions [Li (1992), Cook (1998)], actually estimates the central mean subspaces; whereas some other pre-existing estimates, such as the sliced inverse regression (SIR), the SIR-II [Li (1991)] and the sliced average variance estimator (SAVE) [Cook and Weisberg (1991)], can recover additional directions in the central subspace.

Yin and Cook (2002) extended central mean subspace to central moment subspace, based on the relation $E(Y^k|X) = E(Y^k|P_{\mathcal{S}}X)$, which is written as $\mathcal{S}_{E(Y^k|X)}$. This provides us with a graduation between the central mean subspace and the central subspace. That is, for sufficiently large k , the subspace spanned by $\{\mathcal{S}_{E(Y^\ell|X)}, \ell = 1, \dots, k\}$ approaches the central subspace. Zhu and Zeng (2006) showed that the central mean subspaces for $E(e^{tY}|X), t \in \mathbb{R}$, when put together, recovers the central subspace, and exploited this relation to develop a Fourier transformation method to estimate the central subspace. Here and throughout, we use i to denote the imaginary unit $\sqrt{-1}$. More recently, Zeng and Zhu (2010) developed a general integral transform method. Both papers hint at the following fact: if one can estimate the central mean subspace of $E[f(X)|Y]$ for sufficiently many functions f , then one can recover the central subspace.

In a seminal paper, Xia et al. (2002) introduced a dimension reduction method, called the minimum average variance estimator (MAVE), based on estimation of the gradient of the conditional expectation $E(Y|X)$. This method has three main advantages: (1) it estimates the central mean subspaces exhaustively; (2) it does not impose strong assumptions on the distribution of X ; (3) its computation can be broken down into iterations between two quadratic optimization steps, each of which having an explicit solution. However, a drawback of this method is that it cannot estimate directions outside the central mean subspace. For example, it cannot recover directions in the conditional variance function $\text{var}(Y|X)$. To remedy this deficiency, Xia (2007) and Wang and Xia (2008), respectively, proposed density MAVE (DMAVE) and sliced regression (SR) that can exhaustively estimate the central subspace. The former is based the estimation of the gradients of the functions $E\{H[(Y - a)/b]|X\}$, where H is a known probability density function, $a \in (-\infty, \infty)$ and $b \in (0, \infty)$. The latter is based on the estimation of the gradients of the functions $E[I(Y \leq c)|X]$, where c is an arbitrary constant. Here, again, we see the echo of the same basic fact that estimating the central mean subspaces for a rich enough family of functions is equivalent to estimating the central subspace itself.

The ensemble approach introduced in this paper is based on the same fact, but it is more general, more flexible and, in the numerical examples we considered,

more efficient. In broad outlines the procedure can be described as follows. Consider a general family \mathfrak{F} of functions of Y . For each $f \in \mathfrak{F}$, let $\mathcal{S}_{E[f(Y)|X]}$ denote the central mean subspace for the conditional mean $E[f(Y)|X]$. We say that \mathfrak{F} characterizes the central subspace if the subspace spanned by the collection of subspaces $\{\mathcal{S}_{E[f(Y)|X]} : f \in \mathfrak{F}\}$ is equal to the central subspace. We introduce a probability measure on \mathfrak{F} , and randomly sample functions f_1, \dots, f_m from \mathfrak{F} according to this probability. We then assemble the central mean subspaces $\mathcal{S}_{E[f_\ell(Y)|X]}$, $\ell = 1, \dots, m$, together to recover the central subspace.

In principle, the ensemble approach can be used in conjunction with any estimators of the central mean subspace to recover the central subspace, such as the ordinary least squares, the principal Hessian directions, MAVE and its two variants: the outer product of gradients (OPG) and the refined MAVE (RMAVE). In this paper we focus on its combination with MAVE and its variants, and refer to this combination the MAVE (OPG or RMAVE) ensemble. We show that these ensemble estimators exhaustively estimate the central subspace and that the RMAVE ensemble has the same convergence rate as RMAVE itself. We also introduce a cross validation criterion to determine the dimension of the central subspace, and establish its consistency. Through a number of simulation experiments, most of which are based on published models, we demonstrate the superb performance of the RMAVE ensemble based on the family $\mathfrak{F} = \{e^{tY} : t \in \mathbb{R}\}$. We also explore other characterizing families, such as the Box–Cox transformations and wavelet basis.

The rest of the paper is organized as follows. In Section 2 we investigate what types of family \mathfrak{F} can characterize the central subspace. We introduce the MAVE ensemble in Section 3 and outline the parallel developments for OPG ensemble and the RMAVE ensemble in Section 4. In Section 5 we introduce a cross validation criterion for order determination and discuss the choices of the characterizing family \mathfrak{F} , with emphasis on the characteristic function and the Box–Cox transformations. In Section 6 we establish the consistency and derive the convergence rate of the RMAVE ensemble, and establish the consistency of the cross validation estimator. In Section 7 we conduct simulation comparisons between the RMAVE ensemble and other estimators in a large variety of models. Some concluding remarks are made in Section 8.

2. Characterizing the central subspace. The basic fact that underlies our approach is that the dimension reduction subspaces for the conditional means $E[f(Y)|X]$, when combined in unison, can recover the dimension reduction subspace for Y versus X . For this idea to work, the family of f needs to be sufficiently rich, and in this section we rigorously pose and study this characterization problem.

Let X be a p -dimensional random vector defined on Ω_X and Y be an s -dimensional random vector defined on Ω_Y . Let \mathfrak{F} be a family of functions $f : \Omega_Y \rightarrow \mathbb{F}$, where \mathbb{F} can be the set of real numbers \mathbb{R} or complex numbers \mathbb{C} . Let $\mathcal{S}_{E[f(Y)|X]}$

denote the central mean subspace for the conditional mean $E[f(Y)|X]$, as defined in Cook and Li (2002) and Yin and Cook (2002). That is, $\mathcal{S}_{E[f(Y)|X]}$ is the intersection of all subspaces of \mathbb{R}^p such that

$$(1) \quad E[f(Y)|X] = E[f(Y)|P_{\mathcal{S}}X].$$

Let $\mathcal{S}_{Y|X}$ denote the central subspace of Y versus X as defined in Cook (1994). That is, $\mathcal{S}_{Y|X}$ is the intersection of all linear subspaces of \mathbb{R}^p such that

$$(2) \quad Y \perp\!\!\!\perp X|P_{\mathcal{S}}X.$$

Note that here we do allow Y to be a random vector; whereas the mentioned previous works assume Y to be a scalar. This relaxation is made possible by the transformation f , which takes value in the scalar field \mathbb{F} .

DEFINITION 2.1. Let \mathfrak{F} be a family of measurable \mathbb{F} -valued functions defined on Ω_Y . If

$$(3) \quad \text{span}\{\mathcal{S}_{E[f(Y)|X]} : f \in \mathfrak{F}\} = \mathcal{S}_{Y|X},$$

then we say the family \mathfrak{F} characterizes the central subspace.

Let F_Y denote the distribution of Y , and let $L_2(F_Y)$ be the class of functions $f(Y)$ with finite variances, together with the inner product $\langle f_1, f_2 \rangle = E[f_1(Y)f_2(Y)]$. Let $L_1(F_Y)$ be the class of functions $f(Y)$ such that $E|f(Y)| < \infty$, together with the norm $E|f(Y)|$. We denote the subspace on the left-hand side of (3) by $\mathcal{S}(\mathfrak{F})$. Note that $E[f(Y)|X]$ is finite if $f \in L_1(F_Y)$.

LEMMA 2.1. Suppose that $\mathfrak{F} \subseteq L_1(F_Y)$. Then the following assertions hold:

- (1) $\mathcal{S}(\mathfrak{F}) \subseteq \mathcal{S}_{Y|X}$.
- (2) If (1) being satisfied for all $f \in \mathfrak{F}$ implies (2), then $\mathcal{S}_{Y|X} \subseteq \mathcal{S}(\mathfrak{F})$.

Before proving this lemma we first note the following fact. If \mathcal{S} , \mathcal{S}_1 and \mathcal{S}_2 are linear subspaces of \mathbb{R}^p , then

$$(4) \quad \mathcal{S}_1 \subseteq \mathcal{S}_2 \quad \text{if and only if} \quad \{\mathcal{S} : \mathcal{S} \text{ contains } \mathcal{S}_2\} \subseteq \{\mathcal{S} : \mathcal{S} \text{ contains } \mathcal{S}_1\}.$$

This can be easily seen by taking intersection on both sides of the equality.

PROOF OF LEMMA 2.1. (1) Let \mathcal{S} be a subspace of \mathbb{R}^p that contains $\mathcal{S}_{Y|X}$. Then (2) holds, and consequently (1) holds for all $f \in \mathfrak{F}$. This implies that \mathcal{S} contains $\mathcal{S}_{E[f(Y)|X]}$ for all $f \in \mathfrak{F}$. Since \mathcal{S} is a linear subspace, it must contain $\mathcal{S}(\mathfrak{F})$. Hence

$$\{\mathcal{S} : \mathcal{S} \text{ contains } \mathcal{S}_{Y|X}\} \subseteq \{\mathcal{S} : \mathcal{S} \text{ contains } \mathcal{S}(\mathfrak{F})\},$$

which, by (4), proves part 1.

(2) Let \mathcal{S} be a subspace of \mathbb{R}^p that contains $\mathcal{S}(\mathfrak{F})$. Then (1) holds for all $f \in \mathfrak{F}$. By assumption this implies (2), and consequently \mathcal{S} contains $\mathcal{S}_{Y|X}$. Hence

$$\{\mathcal{S} : \mathcal{S} \text{ contains } \mathcal{S}(\mathfrak{F})\} \subseteq \{\mathcal{S} : \mathcal{S} \text{ contains } \mathcal{S}_{Y|X}\},$$

which, by (4), implies $\mathcal{S}_{Y|X} \subseteq \mathcal{S}(\mathfrak{F})$. \square

Let \mathfrak{B} be the family of measurable indicator functions of Y . That is, $\mathfrak{B} = \{I_B : B \text{ is a Borel set in } \Omega_Y\}$. Note that $\mathfrak{B} \subseteq L_2(F_Y)$.

THEOREM 2.1. *If \mathfrak{F} is a subset of $L_2(F_Y)$ that is dense in \mathfrak{B} , then \mathfrak{F} characterizes the central subspace.*

PROOF. Because \mathfrak{F} is a subset of $L_2(F_Y)$, it is also a subset of $L_1(F_Y)$. Hence, by Lemma 2.1, it suffices to show that (1) being satisfied for all $f \in \mathfrak{F}$ implies (2).

Let \mathcal{S} be a subspace such that (1) holds for all $f \in \mathfrak{F}$, and let B be a Borel set in Ω_Y . Because \mathfrak{F} is dense in \mathfrak{B} there is a sequence $\{f_k\} \subseteq \mathfrak{F}$ such that $\lim_{k \rightarrow \infty} E[I_B(Y) - f_k(Y)]^2 = 0$. For any $g \in L_2(F_X)$ we have

$$\begin{aligned} & E\{g(X)[I_B(Y) - E(I_B(Y)|P_{\mathcal{S}}X)]\} \\ (5) \quad & = E\{g(X)[I_B(Y) - E(f_k(Y)|P_{\mathcal{S}}X)]\} \\ & \quad + E\{g(X)E[f_k(Y) - I_B(Y)|P_{\mathcal{S}}X]\}. \end{aligned}$$

The square of the second term on the right is no more than

$$E[g^2(X)]E\{E^2[f_k(Y) - I_B(Y)|P_{\mathcal{S}}X]\} \leq E[g^2(X)]E\{E[f_k(Y) - I_B(Y)]^2\} \rightarrow 0.$$

Since $f_k \in \mathfrak{F}$ we have $E[f_k(Y)|P_{\mathcal{S}}X] = E[f_k(Y)|X]$. Hence the first term on the right-hand side of (5) can be rewritten as

$$(6) \quad E\{g(X)[I_B(Y) - E(I_B(Y)|X)]\} + E\{g(X)E[I_B(Y) - f_k(Y)|X]\}.$$

The first term is 0 by the definition of conditional expectation. The square of the second term in (6) is no more than

$$E[g^2(X)]E\{E^2[I_B(Y) - f_k(Y)|X]\} \leq E[g^2(X)]E\{[I_B(Y) - f_k(Y)]^2\} \rightarrow 0.$$

Since the left-hand side of (5) does not depend on k , and the right-hand side converges to 0 as $k \rightarrow \infty$, we have $E\{g(X)[I_B(Y) - E(I_B(Y)|P_{\mathcal{S}}X)]\} = 0$. By the definition of conditional expectation the above being true for all $g \in L_2(F_X)$ implies $E[I_B(Y)|P_{\mathcal{S}}X] = E[I_B(Y)|X]$ almost surely. Since B is an arbitrary Borel set in Ω_X , this implies $Y \perp\!\!\!\perp X|P_{\mathcal{S}}X$. \square

This theorem synthesizes several recently developed methods in the literature, and also anticipates useful new ways to combine central mean subspaces into the central subspace. The following examples demonstrate its potential.

EXAMPLE 2.1 (Polynomials). Let $\mathfrak{F} = \{Y^t : t = 1, 2, \dots\}$. Then $\mathcal{S}_{E[f(Y)|X]} = \mathcal{S}_{E(Y^t|X)}$. This is the type of dimension reduction subspaces studied by Cook and Li (2002), Yin and Cook (2002) and more recently Zhu and Zhu (2009). If the conditional moment generating function $E(e^{tY}|X)$ is finite in an open interval that contains 0, then \mathfrak{F} is dense in $L_2(F_Y)$, and hence characterizes $\mathcal{S}_{Y|X}$.

EXAMPLE 2.2 (Kernel density). Let $b > 0$ and H be a symmetric probability density function defined on \mathbb{R} . Let $\mathfrak{F} = \text{span}\{b^{-1}H[(y - t)/b] : t \in \mathbb{R}, b \in \mathbb{R}^+\}$. Xia (2007) proposed a DMAVE method that, in effect, recovers $\mathcal{S}_{Y|X}$ by estimating $\mathcal{S}_{E[f(Y)|X]}$ for $f \in \mathfrak{F}$. This family is dense in $L_2(F_Y)$ when H is the normal density. See, for example, Fukumizu, Bach and Jordan (2009).

EXAMPLE 2.3 (Slices). Let $\mathfrak{F} = \{I_{(-\infty, t)}(y) : t \in \mathbb{R}\}$. Then \mathfrak{F} is clearly dense in \mathfrak{B} . The method proposed by Wang and Xia (2008) is based on the estimation of $\mathcal{S}_{E[f(Y)|X]}$ for f in this family.

EXAMPLE 2.4 (Box–Cox transformations). Let Y be a nonnegative random variable, and consider the family of transformations

$$(7) \quad f_t(y) = \begin{cases} \frac{y^t - 1}{t}, & t \neq 0, \\ \log(y), & t = 0. \end{cases}$$

This is the Box–Cox transformation [Box and Cox (1964)]. This family characterizes the central subspace because it contains the family in Example 2.1.

EXAMPLE 2.5 (Characteristic function). Let $\mathfrak{F} = \{e^{itY} : t \in \mathbb{R}\}$, where $\iota = \sqrt{-1}$. Note that $E(e^{itY}|X)$ is simply the conditional characteristic function of $Y|X$. It is well known that this family is dense in $L_2(F_Y)$. It is used by Zhu and Zeng (2006) to recover the central mean subspace and central subspace, respectively, based on the assumption that X is multivariate normal. This family is also our focus when we implement the ensemble estimators.

EXAMPLE 2.6 (Haar wavelets). Let

$$\psi(t) = \begin{cases} 1, & t \in [0, 1/2), \\ -1, & t \in [1/2, 1), \\ 0, & \text{otherwise.} \end{cases}$$

Consider the family $\mathfrak{F} = \{1\} \cup \{\psi(2^n y - k) : k = 0, \dots, 2^n - 1; n = 1, 2, \dots\}$, where the 1 in $\{1\}$ represents the function of y that always takes the value 1. This is the famous Haar basis often used in wavelet estimators. See, for example, Donoho and Johnstone (1994) and Antoniadis and Fan (2001). The Haar basis is obviously dense in \mathfrak{B} and hence characterizes the central subspace.

In this paper we only consider parametric characterizing families \mathfrak{F} . That is, \mathfrak{F} is of the form $\{f_t : t \in \Omega_T\}$ where Ω_T is a subset of a Euclidean space \mathbb{R}^q . All the characterizing families in the above examples are parametric. In the following, for a sequence of subspaces $\{\mathcal{S}_k\}$ and a subspace \mathcal{S} , we say $\lim_{k \rightarrow \infty} \mathcal{S}_k = \mathcal{S}$ if $\lim_{k \rightarrow \infty} \|P_{\mathcal{S}_k} - P_{\mathcal{S}}\| = 0$, where $\|\cdot\|$ is a matrix norm, such as the operator norm or the Frobenius matrix norm. The two norms are topologically equivalent, and makes no difference in asymptotic analysis. See, for example, [Li, Zha and Chiaromonte \(2005\)](#). Note that we are interested in $\mathcal{S}_{Y|X}$ via $\text{span}\{\mathcal{S}_{E[f_t(Y)|X]} : t \in \Omega_T\}$, then a question arises is whether we can recover $\mathcal{S}_{Y|X}$ from a finite $t \in \Omega_T$. Indeed, we can. [Theorem 2.2](#) below demonstrates that, with probability 1, the central subspace can be characterized by a finite number of functions in a characterizing family. In essence, it relies on the following fact: if a sequence of subspaces \mathcal{S}_m converges to another subspace \mathcal{S} from within, then the norm $\|\mathcal{S}_m - \mathcal{S}\|$ is discrete in nature; that is, if this norm converges to 0 then it must be identically 0 for large m . This phenomenon is also noticed in [Yin, Li and Cook \(2008\)](#). The next lemma, albeit simple, reveals this discrete nature of dimension reduction.

LEMMA 2.2. *Let $\mathcal{S}_1 \subseteq \mathcal{S}_2$ be two subspaces of \mathbb{R}^p . Then $\|P_{\mathcal{S}_2} - P_{\mathcal{S}_1}\|$ is either 0 or no less than 1.*

PROOF. If $\mathcal{S}_1 = \mathcal{S}_2$, then $\|P_{\mathcal{S}_2} - P_{\mathcal{S}_1}\| = 0$. If $\mathcal{S}_1 \neq \mathcal{S}_2$, then direct difference $\mathcal{S}_2 \ominus \mathcal{S}_1$ is nonempty. We know that in this case $P_{\mathcal{S}_2} - P_{\mathcal{S}_1} = P_{\mathcal{S}_2 \ominus \mathcal{S}_1}$. Let v be a unit vector in $\mathcal{S}_2 \ominus \mathcal{S}_1$. Then

$$\|P_{\mathcal{S}_2 \ominus \mathcal{S}_1}\|^2 \geq \|vv^\top\|^2 = \sum_{i=1}^p \sum_{j=1}^p (v_i v_j)^2 = \sum_{i=1}^p v_i^2 \sum_{j=1}^p v_j^2 = 1.$$

This completes the proof. \square

Let $B_0 = (\beta_1, \dots, \beta_{d_0})$ be an orthogonal basis for the central subspace, $\mathcal{S}_{Y|X}$, whose dimension is d_0 . In the following, we will randomly sample T_1, \dots, T_m from Ω_T . In this setting, we assume that these random elements are defined on a measurable space (Ω, \mathcal{A}) . Then Ω_T is interpreted as the range of the mapping $T_i : \Omega \rightarrow \Omega_T$. We denote a generic member of Ω by ω .

THEOREM 2.2. *Suppose that \mathfrak{F} characterizes the central subspace, T_1, T_2, \dots is an i.i.d. sequence of random variables supported on Ω_T and, for each integer m , $B(T_1, \dots, T_m)$ is an orthogonal basis matrix of $\text{span}\{\mathcal{S}_{E[f_{T_i}(Y)|X]} : i = 1, \dots, m\}$. Then the following event has probability 1:*

$$\{\omega \in \Omega : \text{there is an integer } m_0(\omega) \text{ such that,} \\ \text{for all } m \geq m_0(\omega), \text{span}(B(T_1(\omega), \dots, T_m(\omega))) = \text{span}(B_0)\}.$$

PROOF. For $u = 1, \dots, d_0$, let Λ_u be a subset of $\{t \in \Omega_T\}$ such that $\beta_u \notin \text{span}\{\mathcal{S}_{E[f_t(Y)|X]}, t \in \Lambda_u\}$. If $P(T \in \Lambda_u) = 1$ for some u , then \mathfrak{F} does not characterize $\mathcal{S}_{Y|X}$, which is a contradiction. Hence $P(T \in \Lambda_u) < 1$ for $u = 1, \dots, d_0$. Let

$$(8) \quad \delta(T_1, \dots, T_m) = \|B(T_1, \dots, T_m)B^\top(T_1, \dots, T_m) - B_0B_0^\top\|.$$

Note that $\delta(T_1, \dots, T_m) \neq 0$ if and only if, for some $u \in \{1, \dots, d_0\}$, T_1, \dots, T_m all belongs to Λ_u . This is the event $\bigcup_{u=1}^{d_0} \bigcap_{k=1}^m \{T_k \in \Lambda_u\}$, and has probability

$$P\left(\bigcup_{u=1}^{d_0} \bigcap_{k=1}^m \{T_k \in \Lambda_u\}\right) \leq \sum_{u=1}^{d_0} [P(T \in \Lambda_u)]^m.$$

Since

$$\sum_{m=1}^{\infty} \sum_{u=1}^{d_0} [P(T \in \Lambda_u)]^m = \sum_{u=1}^{d_0} \frac{P(T \in \Lambda_u)}{1 - P(T \in \Lambda_u)} < \infty,$$

we have, by the first Borel–Cantelli lemma, with probability 1,

$$(9) \quad \lim_{m \rightarrow \infty} \delta(T_1, \dots, T_m) = 0.$$

Since $\text{span}[B(T_1, \dots, T_m)] \subseteq \text{span}(B_0)$, by Lemma 2.2, event (9) occurs if and only if $\delta(T_1, \dots, T_m)$ becomes 0 for sufficiently large m . Thus, with probability 1, there exists an $m_0(\omega)$ such that for $m > m_0(\omega)$, $\text{span}[B(T_1(\omega), \dots, T_m(\omega))] = \text{span}(B_0)$. \square

3. MAVE ensemble. We first describe our method at the population level, and then develop the estimation procedure at the sample level. The idea underlying MAVE can be outlined as follows. Assume that the central mean subspace $\mathcal{S}_{E(Y|X)}$ has dimension $d_0 < p$. Let β be a $p \times d_0$ matrix such that $\text{span}(\beta) = \mathcal{S}_{E(Y|X)}$. Then

$$\partial E(Y|X = x)/\partial x = \beta[\partial E(Y|\beta^\top X = u)/\partial u].$$

Since the vector on the right always belongs to $\text{span}(\beta)$, we can recover $\text{span}(\beta)$ by estimating the gradient of $E(Y|X = x)$. This is achieved by local linear regression. Let K_h be a probability density function defined on \mathbb{R}^p where h is proportional to the square root of the largest eigenvalue of the variance matrix under K_h . Let f_X be the density of X . Consider the objective function

$$\int_{\Omega_X} E\{[Y - a(x) - b^\top(x)B^\top(X - x)]^2 K_h(X - x)\} f_X(x) dx,$$

where $a : \Omega_X \rightarrow \mathbb{R}$, $b : \Omega_X \rightarrow \mathbb{R}^{d_0}$, $B \in \mathbb{R}^{p \times d_0}$. Let $(a_h^*(\cdot), b_h^*(\cdot), B_h^*)$ be the minimizer of the above function over all possible functions a, b and constant matrices B , then it can be shown that $\lim_{h \rightarrow 0} \text{span}(B_h^*) = \mathcal{S}_{E(Y|X)}$. See Xia et al. (2002).

We now describe at the population level how to assemble a collection of MAVEs to recover the central subspace. Let $\mathfrak{F} = \{f_t : t \in \Omega_T\}$ be a parametric characterizing family. Throughout we assume $\mathbb{F} = \mathbb{C}$, though the subsequent statements are true also for $\mathbb{F} = \mathbb{R}$ by simply discarding the imaginary part. Let $f_t(y, 1)$ and $f_t(y, 2)$ denote the real and imaginary parts of $f_t(y)$. That is, $f_t(y) = f_t(y, 1) + \iota f_t(y, 2)$. Let T be a random vector defined on Ω_T , with distribution F_T . Applying the MAVE procedure to the transformed response $f_t(Y)$ and integrating with respect to the distribution F_T leads to the following population-level objective function:

$$(10) \quad \sum_{\ell=1}^2 \int_{\Omega_T \times \Omega_X} E\{[f_t(Y, \ell) - a_\ell(x) - b_\ell^\top(x)B^\top(X - x)]^2 K_h(X - x)\} dF_X(x) dF_T(t).$$

We minimize this function over all \mathbb{R} -valued functions $a_\ell(\cdot)$, $\ell = 1, 2$, all \mathbb{R}^{d_0} -valued functions $b_\ell(\cdot)$, $\ell = 1, 2$ and all $p \times d_0$ constant matrices B .

At the sample level, suppose that $(X_1, Y_1), \dots, (X_n, Y_n)$ are independent copies of (X, Y) . Let $K_0(\cdot)$ be a symmetric probability density function define on \mathbb{R} . For any $v \in \mathbb{R}^p$ and $h \in \mathbb{R}^+$, let $K_h(v) = h^{-p} K_0(\|v\|/h)$. Let

$$w_{ij}(h) = K_h(X_i - X_j) / \sum_{u=1}^n K_h(X_u - X_j).$$

Let T_1, \dots, T_m be an independent sample from F_T . Mimicking (10) we minimize the sample-level objective function

$$(11) \quad \sum_{\ell=1}^2 \sum_{k=1}^m \sum_{j=1}^n \sum_{i=1}^n \rho_j w_{ij}(h) [f_{T_k}(Y_i, \ell) - a_{jk}(\ell) - b_{jk}^\top(\ell)B^\top(X_i - X_j)]^2$$

over scalars $\{a_{jk}(\ell) : j = 1, \dots, n, k = 1, \dots, m, \ell = 1, 2\}$, d_0 -dimensional vectors $\{b_{jk}(\ell) : j = 1, \dots, n; k = 1, \dots, m, \ell = 1, 2\}$ and $p \times d_0$ matrices B . The coefficients $\{\rho_j : j = 1, \dots, n\}$ are trimming constants. Their purpose is to exclude those X 's with too few observations around, which are unreliable. Let $\rho : \mathbb{R} \rightarrow \mathbb{R}$ be a function with a bounded second derivative such that $\rho(v) > 0$ if $v > v_0$ and $\rho(v) = 0$ if $v \leq v_0$, for some small $v_0 > 0$. We take $\rho_j = \rho(n^{-1} \sum_{i=1}^n K_h(X_i - X_j))$. The bandwidth h is taken to be proportional to $n^{-1/(p+4)}$, which is the optimal bandwidth in the sense of mean integrated squared errors. For more details about the trimming constants and the bandwidth, see Xia et al. (2002), Fan, Yao and Cai (2003), Wang and Xia (2008).

A rather appealing aspect of this procedure is that the minimization of the objective function (11) can be broken down into iterations between two steps, each of which is a quadratic optimization problem having an explicit solution. More specifically, for a fixed $B \in \mathbb{R}^{p \times d_0}$, minimize (11) over $a_{jk}(\ell), b_{jk}(\ell)$ for

$j = 1, \dots, n, k = 1, \dots, m, \ell = 1, 2$. Note that, for each triplet (j, k, ℓ) , the summand of (11),

$$(12) \quad \sum_{i=1}^n \rho_j w_{ij}(h) [f_{T_k}(Y_i, \ell) - a_{jk}(\ell) - b_{jk}^\top(\ell) B^\top (X_i - X_j)]^2$$

depends on and only on $a_{jk}(\ell), b_{jk}(\ell)$. As a result, minimizing (11) jointly is equivalent to minimizing (12) individually. This is a least-squares problem whose solution is

$$\begin{pmatrix} \hat{a}_{jk}(\ell) \\ \hat{b}_{jk}(\ell) \end{pmatrix} = \left[\sum_{i=1}^n w_{ij}(h) \rho_j \Delta_{ij}(B) \Delta_{ij}^\top(B) \right]^{-1} \left[\sum_{i=1}^n w_{ij} \rho_j \Delta_{ij}(B) f_{T_k}(Y_i) \right],$$

where $\Delta_{ij}(B) = (1, (X_i - X_j)^\top B)^\top$.

For fixed $a_{jk}(\ell), b_{jk}(\ell), j = 1, \dots, n, k = 1, \dots, m, \ell = 1, 2$, the minimization of (11) is again a least-squares problem. The solution is

$$\begin{aligned} \text{vec}(\hat{B}) &= \left[\sum \rho_j w_{ij}(h) (b_{jk}(\ell) \otimes (X_i - X_j))(b_{jk}(\ell) \otimes (X_i - X_j))^\top \right]^{-1} \\ &\quad \times \left[\sum \rho_j w_{ij}(h) (b_{jk}(\ell) \otimes (X_i - X_j))(f_{T_k}(Y_i, \ell) - a_{jk}(\ell)) \right], \end{aligned}$$

where the summation is over

$$(13) \quad (i, j, k, \ell) \in \{1, \dots, n\} \times \{1, \dots, n\} \times \{1, \dots, m\} \times \{1, 2\}.$$

Thus, starting with an initial estimate of \hat{B}_0 of $S_{Y|X}$, which, for example, can be the OPG ensemble described in the next section, we iterate between the above two steps until convergence. More specifically, let $\hat{B}^{(r)}$ be the estimate at the r th iteration. We stop when $\|P_{\hat{B}^{(r)}} - P_{\hat{B}^{(r+1)}}\|$ is smaller than some preassigned constant, such as 10^{-6} . The subspace $\text{span}(\hat{B}^{(r+1)})$ is the estimate of $S_{Y|X}$. We call this procedure the MAVE ensemble and the integer m the ensemble size.

4. Variations of MAVE ensemble. Besides MAVE, Xia et al. (2002) also introduced two companion estimators: the outer product of gradients (OPG) and a refinement of MAVE (RMAVE). The former only involves eigen decompositions and is very easy to compute. It is in general less accurate than MAVE, but can be used as an initial estimate for MAVE. The latter involves iterations of steps, each similar to MAVE. It is more accurate than MAVE, and can take MAVE as its initial estimate. In this section we develop parallel generalizations of these methods, which we call the OPG ensemble and the RMAVE ensemble.

4.1. *OPG ensemble.* Let $\mathfrak{F}, F_T, T_1, \dots, T_m$ and $w_{ij}(h)$ be as defined in previous sections. For each j, k, ℓ , we minimize the objective function

$$\sum_{i=1}^n w_{ij}(h) [f_{T_k}(Y_i, \ell) - a - b^\top (X_i - X_j)]^2$$

over $(a, b) \in \mathbb{R} \times \mathbb{R}^p$ for each $j = 1, \dots, n, k = 1, \dots, m$ and $\ell = 1, 2$. This is a least-squares problem and its solution can be written down explicitly, as

$$\begin{pmatrix} \hat{a}_{jk}(\ell) \\ \hat{b}_{jk}(\ell) \end{pmatrix} = \left[\sum_{i=1}^n w_{ij} \Delta_{ij}(I_p) \Delta_{ij}^\top(I_p) \right]^{-1} \left[\sum_{i=1}^n w_{ij} f_{T_k}(Y_i, \ell) \Delta_{ij}(I_p) \right].$$

We then construct the following OPG matrix [Xia et al. (2002)]:

$$\sum_{\ell=1}^2 \sum_{k=1}^m \sum_{j=1}^n \rho_j \hat{b}_{jk}(\ell) \hat{b}_{jk}^\top(\ell).$$

We use the d_0 eigenvectors of this matrix corresponding to its largest eigenvalues as an estimate of $\mathcal{S}_{Y|X}$. This estimate shares the desirable property of OPG. Numerically, all it needs is the calculation of least squares estimate and principal components, none of which involves numerical optimization. As such it is very easy to compute and does not run into local minimum problem, making it an ideal initial estimate for MAVE ensemble.

4.2. *RMAVE ensemble.* The idea of RMAVE is to use an existing consistent estimate of β to reduce the dimension of the kernel function, so that smoothing is carried out over a d_0 -dimensional, rather than a p -dimensional subspace. When d_0 is small, this can mitigate the effect of the curse of dimensionality. In particular, when $d_0 \leq 3$, it achieves the \sqrt{n} -convergence rate.

Let $H_h : \mathbb{R}^{d_0} \rightarrow \mathbb{R}^+$ be the d_0 -dimensional kernel function $h^{-d_0} K_0(\|v\|/h)$, where v is a d_0 -dimensional vector. We minimize the objective function

$$(14) \quad \sum \rho_j [f_{T_k}(Y_i, \ell) - a_{jk}(\ell) - b_{jk}(\ell)^\top B^\top (X_i - X_j)]^2 H_h[B^\top (X_i - X_j)],$$

where the summation is over the indices in (13). Notice that, if we fix the B in the kernel H_h , then the objective function is similar to MAVE, and can be computed by iterations between two least squares problems, as described in Section 3. We can then substitute the updated B and repeat the process until convergence. The algorithm for RMAVE ensemble is summarized as follows.

Let \hat{B}_0 be an initial estimate of $\mathcal{S}_{Y|X}$. For example, we can use the MAVE ensemble to calculate the initial estimate. Set $h_0 \propto n^{-1/(p+4)}$ and $r = 1$.

(1) At step r , let $h_r = \max\{\zeta h_{r-1}, \bar{h}\}$, where $\zeta \in (1/2, 1)$ and $\bar{h} = \bar{h}_0 \times n^{-1/(d_0+4)}$. Note that h_r is a decreasing sequence that converges to \bar{h} . So \bar{h} is the final bandwidth. The purpose of starting with a wider bandwidth and narrowing it gradually is to avoid being trapped in a local minimum at an early stage, as well as to achieve a faster rate of consistency. The proportionality constant \bar{h}_0 can be selected by the rules as suggested by Scott (1992). Let

$$v_{ij}(h_r) = H_{h_r}[\hat{B}^{(r)}(X_i - X_j)] / \sum_{i=1}^n H_{h_r}[\hat{B}^{(r)}(X_i - X_j)].$$

Let

$$\rho_{jr} = \rho \left(n^{-1} \sum_{i=1}^n H_{h_r} [(B^{(r)})^\top (X_i - X_j)] \right),$$

where $\rho : \mathbb{R} \rightarrow \mathbb{R}$ is as defined in Section 3.

(2) Use the two-stage iteration procedure described in Section 3, with $w_{ij}(h)$ and ρ_j therein replaced by $v_{ij}(h_r)$ and ρ_{jr} , respectively, to compute $\hat{B}^{(r)}$. Note that w_{ij} in Section 3 is computed from a p -dimensional kernel, whereas $w_{ij}(h_r)$ here is computed from a d_0 -dimensional kernel.

(3) Standardize $\hat{B}^{(r)}$ so that it is a semiorthogonal matrix. That is, let

$$\hat{B}^{(r)} \leftarrow \hat{B}^{(r)} [\hat{B}^{(r)} (\hat{B}^{(r)})^\top]^{-1/2}.$$

(4) If $\|\hat{B}^{(r)} (\hat{B}^{(r)})^\top - \hat{B}^{(r-1)} (\hat{B}^{(r-1)})^\top\|$ is less than a preassigned small number, say 10^{-6} , then stop and set $\hat{B} = \hat{B}^{(r)}$. Otherwise set $r \leftarrow r + 1$ and return to 1.

5. Order determination and choices of \mathfrak{F} . In describing the foregoing algorithms we have assumed d_0 , the dimension of the central subspace, to be known. In practice this dimension must also be estimated. We now propose a cross validation method to estimate d_0 . Let \hat{B} be the estimate of $\mathcal{S}_{Y|X}$ for a fixed working dimension d . Then the leave-one-out fitted value of $f_{T_k}(Y_j, \ell)$, for $j = 1, \dots, n$, $k = 1, \dots, m$ and $\ell = 1, 2$, is

$$\hat{\mu}_{kj}(d, \ell) = \sum_{i \neq j} K_h [\hat{B}^\top (X_i - X_j)] f_{T_k}(Y_i, \ell) / \sum_{i \neq j} K_h [\hat{B}^\top (X_i - X_j)].$$

The corresponding cross validation value is

$$CV(d) = \frac{1}{2mn} \sum_{\ell=1}^2 \sum_{k=1}^m \sum_{j=1}^n [f_{T_k}(Y_j, \ell) - \hat{\mu}_{kj}(d, \ell)]^2.$$

To include the trivial case of $d = 0$, we define $\hat{\mu}_{kj}(0, \ell)$ to be

$$(n - 1)^{-1} \sum_{i \neq j} f_{T_k}(Y_i, \ell),$$

so that $CV(d)$ is defined for all $d = 0, \dots, p$. The structural dimension d_0 is estimated by

$$\hat{d}_0 = \arg \min \{ CV(d) : d = 0, \dots, p \}.$$

As we have mentioned in Section 2, there are many possible choices for \mathfrak{F} . In this paper we pay special attention to two families: the family determined by the characteristic function, as discussed in Example 2.5, and the family that corresponds to the Box–Cox transformations, as discussed in Example 2.4. That is,

$$\mathfrak{F}_C = \{ e^{t^\top y} : t \in \mathbb{R} \}, \quad \mathfrak{F}_B = \{ f_t : t \in \mathbb{R} \},$$

where f_i is as defined in (7). An advantage of the family \mathfrak{F}_C is that its members are bounded functions, and as such are relatively robust against the outliers in Y . Moreover, it requires virtually no condition on the distribution of Y . Also note that when t ranges over \mathbb{R}^s , the function $e^{t^\top y}$ fully recovers the joint information of the random vector Y . In this respect the ensemble estimators are akin to Projective Resampling [Li, Wen and Zhu (2008)]. However, here the univariate and multivariate responses are treated in a unified manner: we simply replace $e^{t^\top y}$ by $e^{t^\top Y}$, whereas in projective resampling the multivariate response is treated differently from the univariate response.

The family \mathfrak{F}_B requires Y to be nonnegative. When Y is not nonnegative, we make the transformation $Y_i - \min\{Y_1, \dots, Y_n\} + 0.5$ before applying the Box–Cox transformation. An advantage of this family is that often a few fixed functions in \mathfrak{F}_B would do a reasonably good job. In our simulation studies we have used $t \in \{-2, -1.5, -1, -0.5, 0, 0.5, 1, 1.5, 2\}$, as one typically uses for Box–Cox transformation. Note, however, if one uses such a finite, fixed set, then the corresponding \mathfrak{F}_B is not guaranteed to characterize the central subspace, unless the distribution of Y satisfies some special conditions. Alternative transformations such as those proposed by Manly (1976), John and Draper (1980) and Bickel and Doksum (1981), for instance, that do not require Y to be positive may be used to form different family \mathfrak{F} .

Henceforth we indicate an ensemble estimator based on a family \mathfrak{F} by attaching \mathfrak{F} to the name of the original estimator, such as MAVE- \mathfrak{F}_C or RMAVE- \mathfrak{F}_B . To implement RMAVE- \mathfrak{F}_C , we modify the code for the sliced regression in Wang and Xia (2008) based on a gradual descending algorithm; the random vectors T_1, \dots, T_m are an independent sample from $N(0, I_s)$. To implement RMAVE- \mathfrak{F}_B , we adopt the code for RMAVE by Xia et al. (2002) and use the fixed set of t mentioned earlier.

6. Consistency and convergence rate. In this section we investigate the asymptotic behavior of RMAVE ensemble based on \mathfrak{F}_C . We will study the convergence rate, assuming the structural dimension d_0 is known, and then the consistency of the estimator of d_0 . The asymptotic analysis proceeds in two steps. In the first step (Theorem 6.1), we establish the convergence rate for a fixed set of functions $\{f_{t_1}, \dots, f_{t_m}\}$ in \mathfrak{F} . In the second step (Theorem 6.2), we investigate the asymptotic behavior when $m \rightarrow \infty$. The first step is not fundamentally different from the asymptotic results for DMAVE and SR as developed in Xia (2007) and Wang and Xia (2008). We have therefore relegated the proof to an external Appendix. The second step is a novel development and is presented in detail. Although here we only consider RMAVE- \mathfrak{F}_C , we have no doubt that the development can be extended to other characterizing families. For any finite set $\{t_1, \dots, t_m\} \subseteq \Omega_T$, let $\hat{B}(t_1, \dots, t_m)$ denote the RMAVE- \mathfrak{F}_C estimator described in Section 4.2, and let $B(t_1, \dots, t_m)$ be a basis matrix of $\text{span}\{\mathcal{S}_{E[f_i(Y)|X]} : i =$

$1, \dots, m\}$ and B be a generic matrix with p rows. Without loss of generality, assume these matrices to be semiorthogonal.

We need to make the following regularity assumptions, which are similar to those made in Xia (2007) and Wang and Xia (2008).

(C1) *Marginal distribution of X* : The random vector X has a bounded support; its density function $g(x)$ has a bounded second derivative; the functions

$$(u, B) \mapsto E(X|B^\top X = u), \quad (u, B) \mapsto E(XX^\top|B^\top X = u)$$

have bounded derivatives for $u \in \mathbb{R}^{d_0}$ and $B \in \{B : \|BB^\top - B_0B_0^\top\| \leq c\}$, where $c > 0$.

(C2) *Conditional distribution function of Y given $B^\top X$* : The conditional density function $g(y|u)$ of Y given $B^\top X$ has a bounded fourth-order derivative with respect to x and u as B is in a small neighbor of B_0 .

(C3) *Identifiability of minimum*: For any semiorthogonal matrix $B \in \mathbb{R}^{p \times d}$, any constant $c > 0$ and a set $\{t_1, \dots, t_m\} \subseteq \Omega_T$,

$$\inf_{\{B : \|BB^\top - B_0B_0^\top\| \geq c\}} \sum_{\ell=1}^2 \sum_{k=1}^m E[E(f_{t_k}(Y, \ell)|B^\top X) - E(f_{t_k}(Y, \ell)|B_0^\top X)]^2 > 0.$$

(C4) *Kernel function*: The function K_0 is a symmetric univariate density with bounded second derivative and a compact support.

(C5) *Bandwidth*: For a working dimension d , the bandwidths $\{h_r : r = 0, 1, \dots\}$ satisfy $h_0 \propto n^{-1/(p+4)}$, $h_r = \max\{\varsigma h_{r-1}, \bar{h}\}$ with $1/2 < \varsigma < 1$ and $\bar{h} \propto n^{-1/(d+4)}$.

The following theorem gives the convergence rate of RMAVE- \mathfrak{F}_C for a fixed set of functions in \mathfrak{F} and a fixed d_0 . Let $d(t_1, \dots, t_m)$ be the dimension of the space spanned by $\{\mathcal{S}_{E[f_{t_i}(Y)|X]} : i = 1, \dots, m\}$.

THEOREM 6.1. *Suppose conditions (C1), (C2), (C4) and (C5) are satisfied, (C3) holds for $\{t_1, \dots, t_m\} \subseteq \Omega_T$ and set $d = d(t_1, \dots, t_m)$. Then, as $n \rightarrow \infty$,*

$$\begin{aligned} (15) \quad & \|\hat{B}(t_1, \dots, t_m)\hat{B}^\top(t_1, \dots, t_m) - B(t_1, \dots, t_m)B^\top(t_1, \dots, t_m)\| \\ & = O_P[\bar{h}^4 + \log n / (n\bar{h}^{d_0}) + n^{-1/2}]. \end{aligned}$$

PROOF. Let

$$(16) \quad \bar{h}(t_1, \dots, t_m) \propto n^{-1/[d(t_1, \dots, t_m)+4]}.$$

Then, by arguments similar to those used in Xia et al. (2002), Xia (2007) and Wang and Xia (2008), it can be shown that

$$\begin{aligned} & \|\hat{B}(t_1, \dots, t_m)\hat{B}^\top(t_1, \dots, t_m) - B(t_1, \dots, t_m)B^\top(t_1, \dots, t_m)\| \\ & = O_P(\bar{h}^4(t_1, \dots, t_m) + \log n / [n\bar{h}^{d(t_1, \dots, t_m)}(t_1, \dots, t_m)] + n^{-1/2}). \end{aligned}$$

See the external Appendix. By (16), the right-hand side is of the order

$$O_P(n^{-4/[d(t_1, \dots, t_m)+4]} + (\log n)n^{-4/[d(t_1, \dots, t_m)+4]} + n^{-1/2}).$$

Since the function in $O_P(\cdot)$ is increasing in $d(t_1, \dots, d_m)$, which is no more than d_0 , relation (15) holds. \square

Note that we are interested in $S_{Y|X}$ instead of $\text{span}\{\mathcal{S}_{E[f_{i_j}(Y)|X]} : i = 1, \dots, m\}$. The next theorem shows that, under the conditions no stronger than Theorem 6.1, RMAVE- \mathfrak{F}_C recovers the central subspace at the same rate as does RMAVE itself.

THEOREM 6.2. *Suppose that conditions (C1)–(C5) hold, that T_1, \dots, T_m are an independent sample from Ω_T and that they are independent of $(X_1, Y_1), \dots, (X_n, Y_n)$. Let $\hat{B}(T_1, \dots, T_m)$ be the RMAVE- \mathfrak{F}_C estimator of $B(T_1, \dots, T_m)$. Then, for any $\varepsilon > 0$,*

$$(17) \quad \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P\left(\frac{\|\hat{B}(T_1, \dots, T_m)\hat{B}^\top(T_1, \dots, T_m) - B_0B_0^\top\|}{\hbar^4 + \log n/(n\hbar^{d_0}) + n^{-1/2}} > \varepsilon\right) = 0.$$

PROOF. By Theorem 2.2, we have that $\delta(T_1, \dots, T_m)$ becomes 0 for sufficiently large m . Consequently, $P(\liminf_{m \rightarrow \infty} \{\delta(T_1, \dots, T_m) = 0\}) = 1$. By Fatou’s lemma,

$$(18) \quad \liminf_{m \rightarrow \infty} P(\delta(T_1, \dots, T_m) = 0) \geq P\left(\liminf_{m \rightarrow \infty} \{\delta(T_1, \dots, T_m) = 0\}\right) = 1.$$

Thus we see that the bias term converges to 0 infinitely fast.

Next, let $a_n = \hbar^4 + \log n/(n\hbar^{d_0}) + n^{-1/2}$. We have

$$\begin{aligned} & a_n^{-1} \|\hat{B}(T_1, \dots, T_m)\hat{B}^\top(T_1, \dots, T_m) - B_0B_0^\top\| \\ & \leq a_n^{-1}\hat{\delta}_n(T_1, \dots, T_m) + a_n^{-1}\delta(T_1, \dots, T_m), \end{aligned}$$

where $\delta(T_1, \dots, T_m)$ is as defined before and

$$\begin{aligned} & \hat{\delta}_n(T_1, \dots, T_m) \\ & = \|\hat{B}(T_1, \dots, T_m)\hat{B}^\top(T_1, \dots, T_m) - B(T_1, \dots, T_m)B^\top(T_1, \dots, T_m)\|. \end{aligned}$$

Since $a_n \neq 0$, we have, by (18),

$$P(a_n^{-1}\delta(T_1, \dots, T_m) \neq 0) = P(\delta(T_1, \dots, T_m) \neq 0) \rightarrow 0 \quad \text{as } m \rightarrow \infty.$$

Since, despite its appearance, the term on the left does not depend on n , the above limit can be rewritten as

$$(19) \quad \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P(a_n^{-1}\delta(T_1, \dots, T_m) \neq 0) = 0.$$

By Theorem 6.1, for a fixed set t_1, \dots, t_m , $\lim_{n \rightarrow \infty} P(a_n^{-1}\hat{\delta}_n(t_1, \dots, t_m) > \varepsilon) = 0$. But because T_1, \dots, T_m and $(X_1, Y_1), \dots, (X_n, Y_n)$ are independent, this implies

$$\lim_{n \rightarrow \infty} P(a_n^{-1}\hat{\delta}_n(T_1, \dots, T_m) > \varepsilon | T_1 = t_1, \dots, T_m = t_m) = 0.$$

By the dominated convergence theorem, $\lim_{n \rightarrow \infty} P(a_n^{-1} \hat{\delta}_n(T_1, \dots, T_m) > \varepsilon) = 0$, and hence

$$(20) \quad \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P(a_n^{-1} \hat{\delta}_n(T_1, \dots, T_m) > \varepsilon) = 0.$$

Now combine (19) and (20) to prove (17). \square

Theorem 6.2 implies that if $d_0 \leq 3$, then \sqrt{n} -consistency can be achieved by taking $\bar{h} \propto n^{-1/(d_0+4)}$.

Next, we establish the consistency of the estimator of d_0 described in Section 5. Let $\hat{d}(t_1, \dots, t_m)$ be the cross validation estimator of $d(t_1, \dots, t_m)$. The proof of the following lemma can be found in the external Appendix.

LEMMA 6.1. *Suppose that conditions (C1), (C2), (C4) and (C5) hold, (C3) is satisfied for $\{t_1, \dots, t_m\} \subseteq \Omega_T$ and the bandwidth \bar{h}_d used for different dimension d satisfies $\bar{h}_d \propto n^{-1/(d+4)}$. Then we have*

$$\lim_{n \rightarrow \infty} P(\hat{d}(t_1, \dots, t_m) = d(t_1, \dots, t_m)) = 1.$$

We now consider the convergence to the structural dimension d_0 as $m \rightarrow \infty$. Let $\hat{d}(T_1, \dots, T_m)$ be the cross validation estimator of $d(T_1, \dots, T_m)$, which is the dimension for $B(T_1, \dots, T_m)$.

THEOREM 6.3. *Under the assumptions in Theorem 6.2 we have*

$$\lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P(\hat{d}(T_1, \dots, T_m) = d_0) = 1.$$

PROOF. Following the same argument that leads to (18) in the proof of Theorem 6.2, we can show that

$$(21) \quad \lim_{m \rightarrow \infty} P(d(T_1, \dots, T_m) = d_0) = \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P(d(T_1, \dots, T_m) = d_0) = 1.$$

As in the proof of Theorem 2.2, since $\text{span}[B(T_1, \dots, T_m)] \subseteq \text{span}(B_0)$, by Lemma 2.2, event (9) occurs if and only if $\delta(T_1, \dots, T_m)$ becomes 0 for sufficiently large m . By the definition of $\delta(T_1, \dots, T_m)$ in (8), for sufficiently large m , $d(T_1, \dots, T_m) = d_0$. Consequently, $P(\liminf_{m \rightarrow \infty} \{d(T_1, \dots, T_m) = d_0\}) = 1$. By Fatou’s lemma,

$$\liminf_{m \rightarrow \infty} P(d(T_1, \dots, T_m) = d_0) \geq P(\liminf_{m \rightarrow \infty} \{d(T_1, \dots, T_m) = d_0\}) = 1.$$

Thus $\lim_{m \rightarrow \infty} P(d(T_1, \dots, T_m) = d_0) = 1$. Since $d(T_1, \dots, T_m)$ does not depend on n , (21) holds.

Since T_1, \dots, T_m are independent of $(X_1, Y_1), \dots, (X_n, Y_n)$, Lemma 6.1 implies that $\lim_{n \rightarrow \infty} P(\hat{d}(T_1, \dots, T_m) = d(T_1, \dots, T_m) | T_1 = t_1, \dots, T_m = t_m) = 1$. By the

dominated convergence theorem, $\lim_{n \rightarrow \infty} P(\hat{d}(T_1, \dots, T_m) = d(T_1, \dots, T_m)) = 1$, which implies

$$(22) \quad \lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty} P(\hat{d}(T_1, \dots, T_m) = d(T_1, \dots, T_m)) = 1.$$

The desired assertion follows from (21) and (22). \square

Theorem 6.3 confirms that the proposed CV criterion is indeed consistent in selecting the dimension of the central subspace.

7. Simulation studies. In this section we compare the ensemble estimators, RMAVE- \mathfrak{F}_C and RMAVE- \mathfrak{F}_B , with existing methods such as SIR, SAVE, DMAVE, RMAVE and SR. For an estimate \hat{B} of B_0 , both assumed to be semiorthogonal without loss of generality, the estimation error is measured by $\Delta(\hat{B}, B_0) = \|\hat{B}\hat{B}^\top - B_0B_0^\top\|$, where $\|\cdot\|$ is the operator norm [Li, Zha and Chiaromonte (2005)]. For each setting, 100 replicates of the data are generated, unless stated otherwise.

EXAMPLE 7.1. The purpose of this example to demonstrate that the performance of RMAVE- \mathfrak{F}_C is very stable as the ensemble size m varies. Let

$$Y_i = \cos(2X_{i1}) - \cos(X_{i2}) + 0.2\varepsilon_i, \quad i = 1, \dots, n,$$

where ε_i is a standard normal random variable, and X_i is a random vector in \mathbb{R}^{10} . The random vector X_i is generated by $N(0, \Sigma_X)$, where the (i, j) th entry of Σ_X is $0.5^{|i-j|}$. For this model $d_0 = 2$ and $B_0 = (e_1, e_2) \in \mathbb{R}^{10 \times 2}$, where e_i is a vector whose i th entry is 1 and other entries are 0. The model was used in Li (1992) and Wang and Xia (2008).

In Figure 1 we plot the averages of $\Delta(\hat{B}, B_0)$ over the 100 simulated samples versus different ensemble sizes m , ranging from 5 to 50. The left panel corresponds to $n = 200$, and the right panel corresponds to $n = 400$. We see that the average error is quite stable as m varies: for $n = 200$ it is between 0.1806 and 0.1922, and for $n = 400$ it is between 0.1066 and 0.1086.

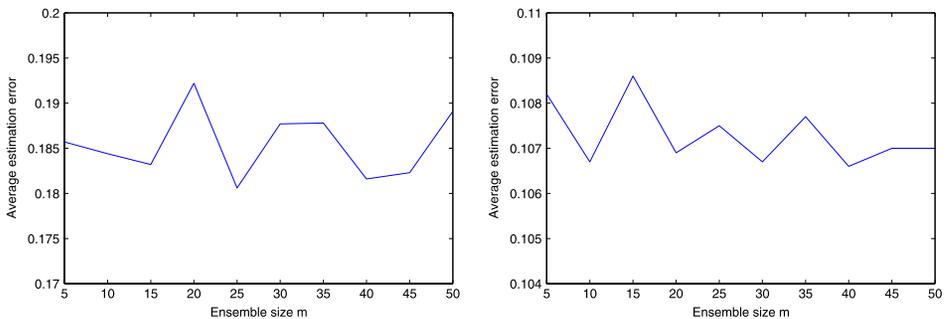


FIG. 1. Averaged $\Delta(\hat{B}, B_0)$ versus ensemble size m for Example 7.1. Left panel: $n = 200$. Right panel: $n = 400$.

TABLE 1

Comparisons for (Example 7.2). Entries are mean \pm standard error of $\Delta(\hat{B}_0, B_0)$ calculated from 100 simulation samples

p	RMAVE- \mathfrak{F}_C	SR	RMAVE- \mathfrak{F}_B	RMAVE	DMAVE	SIR	SAVE
10	0.388 ± 0.134	0.513 ± 0.192	0.878 ± 0.136	0.750 ± 0.170	0.409 ± 0.153	0.851 ± 0.115	0.903 ± 0.117
20	0.638 ± 0.150	0.844 ± 0.146	0.954 ± 0.060	0.880 ± 0.111	0.719 ± 0.159	0.947 ± 0.055	0.977 ± 0.034

EXAMPLE 7.2. The following regression model is a modification of Example 3 of Wang and Xia (2008):

$$Y_i = \frac{X_{i1}}{0.5 + (X_{i2} + 1.5)^2} + X_{i3}\varepsilon_i,$$

where ε_i and X_i are generated as in Example 7.1 with $n = 400$. In this case $d_0 = 3$ and $B_0 = (e_1, e_2, e_3) \in \mathbb{R}^{p \times 3}$. We use $m = 15$ for RMAVE- \mathfrak{F}_C and the number of slices $H = 5$ for SR, SIR and SAVE. Table 1 below indicates that RMAVE- \mathfrak{F}_C is the best performer, followed by DMAVE and SR.

EXAMPLE 7.3. The following model is taken from Zhu and Zeng (2006), Example 3:

$$Y_i = I_{[\beta_1^\top X_i + \sigma\varepsilon_i > 1]} + 2I_{[\beta_2^\top X_i + \sigma\varepsilon_i > 0]},$$

where ε_i is a standard normal random variable, $\sigma = 0.2$ and $X_i \sim N(0, I_p)$. The regression coefficients are $\beta_1 = e_1 + \dots + e_4$ and $\beta_2 = e_{p-3} + \dots + e_p$. Thus we have $d_0 = 2$ and $B_0 = (\beta_1, \beta_2)$. The specifications for n, m, H are the same as Example 7.2. Table 2 below reports the results. In this case DMAVE is the top performer, with RMAVE- \mathfrak{F}_C as a close second.

TABLE 2

Comparisons for (Example 7.3). Entries are mean \pm standard error of $\Delta(\hat{B}_0, B_0)$ calculated from 100 simulation samples

p	RMAVE- \mathfrak{F}_C	SR	RMAVE- \mathfrak{F}_B	RMAVE	DMAVE	SIR	SAVE
10	0.101 ± 0.023	0.149 ± 0.044	0.156 ± 0.047	0.149 ± 0.041	0.100 ± 0.019	0.256 ± 0.048	0.294 ± 0.063
20	0.155 ± 0.028	0.246 ± 0.059	0.243 ± 0.055	0.243 ± 0.098	0.148 ± 0.020	0.339 ± 0.046	0.510 ± 0.096

TABLE 3
*Percentage of correctly estimated d_0 using the CV criterion
 combined with RMAVE- \mathfrak{F}_C (Example 7.4)*

Model	$n = 100$	$n = 200$	$n = 400$
A	0.95	1.00	1.00
B	0.83	1.00	1.00
C	0.39	0.60	0.79

EXAMPLE 7.4. This example is to investigate the effectiveness of the CV criterion for order determination introduced in Section 5, as used in conjunction with RMAVE- \mathfrak{F}_C , in the spirit similar to Example 4 of Wang and Xia (2008). We consider the following three models:

Model A: $Y_i = (X_i^\top \beta)^{-1} + 0.2\varepsilon_i, \quad \beta_0 = e_1 + \dots + e_4,$

Model B: $Y_i = \cos(2X_{i1}) - \cos(X_{i2}) + 0.2\varepsilon_i,$

Model C: $Y_i = X_{i1}/[0.5 + (X_{i2} + 1.5)^2] + X_{i3}^2\varepsilon_i,$

where X is generated as in Example 7.1. We take $p = 10, m = 15$. Table 3 shows that, as the sample size n increases, the percentages of correctly identified dimensions quickly approach to 100% for all three models, which is comparable with the results in Wang and Xia (2008). Our results for the first two models show substantial improvement over the corresponding results in Wang and Xia (2008) for $n = 100$ and $n = 200$. A possible explanation of this improvement is that RMAVE- \mathfrak{F}_C allows us to make repeated use of the sample of responses, with each repetition exploring a different aspect of the central subspace. In other words the ensemble approach makes fuller use of the data than dividing them into slices.

EXAMPLE 7.5. The four models in this example are the same as those used in Example 5 of Wang and Xia (2008):

Model D: $Y_i = (X_i^\top \beta)^{-1} + 0.2\varepsilon_i, \quad \beta = e_1 + \dots + e_4;$

Model E: $Y_i = 0.1(X_i^\top \beta + \varepsilon_i)^3, \quad \beta = e_1 + \dots + e_4;$

Model F: $Y_i = \exp(X_i^\top \beta) \times \varepsilon_i, \quad \beta = e_1 + 0.5e_2 + e_3;$

Model G: $Y_i = \text{sign}(2X_{i1} + \varepsilon_{i1}) \times \log|2X_{i2} + 4 + \varepsilon_{i2}|.$

Here, X, n, m, H are the same as specified in Example 7.1. Table 4 below reports the result for $n = 400$. We see that RMAVE- \mathfrak{F}_C again consistently outperforms other estimators in all four models, though SR is quite close to it in some cases.

TABLE 4

Comparisons for (Example 7.5). Entries are mean \pm standard error of $\Delta(\hat{B}_0, B_0)$ calculated from 100 simulation samples

p	Model	RMAVE- \mathfrak{F}_C	SR	RMAVE- \mathfrak{F}_B	RMAVE	DMAVE	SIR	SAVE
10	D	0.052	0.082	0.994	0.983	0.313	0.306	0.243
		± 0.014	± 0.024	± 0.008	± 0.067	± 0.412	± 0.072	± 0.075
	E	0.053	0.059	0.609	0.058	0.054	0.141	0.154
		± 0.016	± 0.015	± 0.173	± 0.015	± 0.014	± 0.037	± 0.040
	F	0.163	0.178	0.869	0.798	0.208	0.225	0.230
		± 0.043	± 0.055	± 0.979	± 0.141	± 0.073	± 0.056	± 0.067
	G	0.198	0.217	0.387	0.349	0.931	0.242	0.601
		± 0.053	± 0.059	± 0.178	± 0.109	± 0.084	± 0.058	± 0.230
20	D	0.080	0.122	0.955	0.996	0.501	0.399	0.363
		± 0.016	± 0.023	± 0.006	± 0.006	± 0.459	± 0.070	± 0.068
	E	0.068	0.092	0.694	0.085	0.075	0.179	0.236
		± 0.015	± 0.018	± 0.200	± 0.017	± 0.014	± 0.030	± 0.061
	F	0.250	0.281	0.927	0.879	0.323	0.303	0.414
		± 0.048	± 0.052	± 0.058	± 0.103	± 0.080	± 0.053	± 0.083
	G	0.299	0.321	0.507	0.535	0.973	0.312	0.957
		± 0.048	± 0.054	± 0.146	± 0.166	± 0.034	± 0.050	± 0.060

EXAMPLE 7.6. As we noted before, the family \mathfrak{F}_C is particularly useful for recovering directions in the $\mathcal{S}_{Y|X}$ that do not belong to $\mathcal{S}_{E(Y|X)}$, and when Y contains outliers. This example indicates that in the case where $\mathcal{S}_{Y|X} = \mathcal{S}_{E(Y|X)}$ and Y contains no outliers—conditions favorable to RMAVE. Consider the model

$$Y_i = \arcsin(1/(1 + |0.5 + X_{i1}|)) + 0.2\varepsilon_i,$$

where ε_i and $X_i \in \mathbb{R}^{10}$ are generated as in Example 7.1. Note that in this case both $\mathcal{S}_{Y|X}$ and $\mathcal{S}_{E(Y|X)}$ are spanned by e_1 . We take $m = 15$, $n = 400$ and the number of slices for SR, SIR and SAVE equal to 5. However, Table 5 indicates that RMAVE- \mathfrak{F}_C is slightly better than RMAVE.

EXAMPLE 7.7. The main point of this example is to demonstrate numerically the \sqrt{n} -consistency of RMAVE- \mathfrak{F}_C , which we have shown analytically in Section 6 for $d_0 \leq 3$. A secondary point is to reconfirm the stability of this estimator against the change of ensemble size m , for a wide range of sample sizes n , which we have demonstrated in Example 7.1 for two sample sizes ($n = 200, 400$). This second point also provides us intuition about the double limits, $\lim_{m \rightarrow \infty} \lim_{n \rightarrow \infty}$, we took in Section 6.

TABLE 5
 Comparisons for (Example 7.6). Entries are mean \pm standard error of $\Delta(\hat{B}_0, B_0)$ calculated from 100 simulation samples

p	RMAVE- \mathfrak{F}_C	SR	RMAVE- \mathfrak{F}_B	RMAVE	DMAVE	SIR	SAVE
10	0.098 ± 0.023	0.114 ± 0.031	0.128 ± 0.035	0.103 ± 0.026	0.132 ± 0.031	0.434 ± 0.134	0.310 ± 0.086
20	0.131 ± 0.026	0.155 ± 0.033	0.180 ± 0.038	0.143 ± 0.028	0.207 ± 0.062	0.590 ± 0.106	0.547 ± 0.132

Here we adopt the approach of Wang and Xia (2008), Example 8. We use model D in Example 7.5. In Figure 2 we plot the averaged $\Delta(\hat{B}, B_0)$ against $1/\sqrt{n}$ for $m = 15, 30, 50, 100$. The value of $1/\sqrt{n}$ ranges from 0.045 to 0.1, corresponding to sample sizes $n = 100, 200, 300, 400, 500$ in reverse order. We can see that the curves are roughly straight lines passing through the origin, which confirms the \sqrt{n} -consistency. We also see that the performance of RMAVE- \mathfrak{F}_C is very stable as m changes, across different sample sizes.

EXAMPLE 7.8. Finally, we investigate the performance of \mathfrak{F}_C for multivariate responses. The model is taken from Li, Wen and Zhu (2008), Model 4.4. Here $p = 6, s = 5$. The predictor X_i is generated from $N(0, I_6)$. The error ε_i is generated from $N(0, \Sigma)$, where $\Sigma = \text{diag}(\Sigma_1, \Sigma_2)$, in which

$$\Sigma_1 = \begin{pmatrix} 1 & -1/2 \\ -1/2 & 1/2 \end{pmatrix}, \quad \Sigma_2 = \begin{pmatrix} 1/2 & 0 & 0 \\ 0 & 1/3 & 0 \\ 0 & 0 & 1/4 \end{pmatrix}.$$

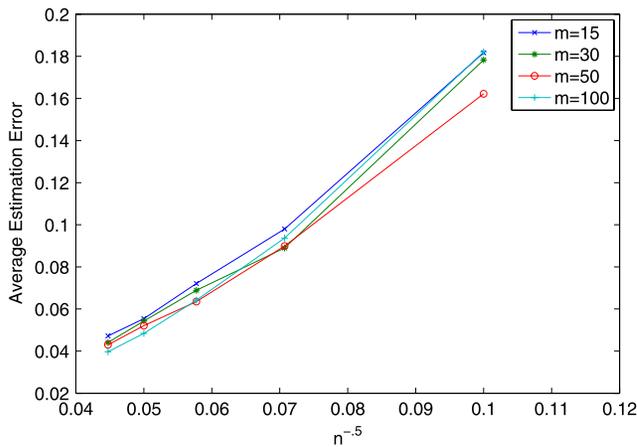


FIG. 2. Numerical demonstration of \sqrt{n} -consistency for RMAVE- \mathfrak{F}_C , as described in Example 7.7.

TABLE 6
Comparison of different estimators for multivariate Y
(Example 7.8, n = 100)

PR-SIR	RMAVE- \mathfrak{F}_C	PR-RMAVE
0.276 ± 0.088	0.248 ± 0.088	0.206 ± 0.071

The 5-dimensional response random vector Y is generated as:

$$\begin{aligned}
 Y_{i1} &= X_{i2} + 3X_{i2}/[0.5 + (X_{i1} + 1.5)^2] + \varepsilon_i, \\
 Y_{i2} &= X_{i1} + e^{0.5X_{i2}} + \varepsilon_{i2}, \\
 Y_{i3} &= X_{i1} + X_{i2} + \varepsilon_{i3}, \\
 Y_{i4} &= \varepsilon_{i4}, \\
 Y_{i5} &= \varepsilon_{i5}.
 \end{aligned}$$

For a fair comparison, we use the Frobenius norm instead of the operator norm for $\Delta(\hat{B}, B_0)$, the former of which was used by Li, Wen and Zhu (2008). Table 6 shows the results for $n = 100$, averaged over 1,000 simulated samples. The columns PR-SIR and PR-RMAVE refer to projective resampling used in conjunction with the SIR and RMAVE, respectively. See Li, Wen and Zhu (2008). The numbers in the PR-SIR column is taken from that paper. For RMAVE- \mathfrak{F}_C we use $m = 15,000$ random directions; for PR-RMAVE, we use 1,000 random directions.

In this case PR-RMAVE performs the best among the three estimators. Note that in this example the central subspace and the central mean subspace coincide, which is the most favorable scenario for methods derived from RMAVE.

8. Discussion. In this paper we introduce a general method for combining estimators of a family of central mean subspaces into a single estimator of the central subspace using the MAVE-type procedures as basic estimators for the central mean subspaces. Different combinations of the characterizing families and MAVE-type procedures result in a class of new estimators of the central subspace, which we call the ensemble estimators. Ensemble estimators exhaustively estimate the central subspace and are relatively easy to compute. The algorithm for estimation can be broken down into iterations of quadratic optimization steps, whose solutions have the least-square form. The ensemble estimators do not require special treatment for multivariate responses, because the characterizing nature of \mathfrak{F} automatically takes into account the multivariate information in the response. Ensemble estimators allow repeated use of the available sample of responses, and by doing so enhance the estimation accuracy. They do not require dividing the sample into slices, which not only simplifies the operation but also avoid sensitivity

to the number of slices. Ensemble estimators have the same convergence rates as their corresponding MAVE-type estimators. In particular, the RMAVE ensemble has the \sqrt{n} -rate when the structural dimension d_0 is no more than 4.

An important problem is the choice of \mathfrak{F} . At this stage we do not yet have a good theory to generate a universal criterion that can work across families. One theoretical difficulty in devising a general criterion to choose among different families \mathfrak{F} is that different transformations of the response result in different scales that cannot be meaningfully compared. For example, if we use cross validation of prediction errors to choose among families, then we face the problem that the prediction errors in different families have different meanings. At this stage, we suspect that any general criterion capable of choosing among different families must be intrinsic to the probabilistic relation between X and Y , as reflected in the conditional distribution of Y given X , rather than specific to any form of transformation of the response.

Our empirical knowledge seems to indicate that bounded transformations, such as \mathfrak{F}_C and SR, are preferable to unbounded transformations, such as power transformation (Example 2.1) and Box–Cox transformation (Example 2.4), especially when the model permits extreme values in the response. A bounded characterizing family of transformations serve the dual purposes of comprehensively describing the central subspace and decreasing the leverage of the extreme response values. In addition, the transformations in \mathfrak{F}_C make full reuse of the data at each resampling. In this respect it is rather similar to the bootstrap, except that the resampling is done by random projection. Indeed, this is the very spirit of ensemble estimator we would like to advocate in this paper, and it is this aspect that distinguishes the ensemble estimators from other sufficient dimension reduction estimators. Finally, in the majority of examples we considered in Section 7, the \mathfrak{F}_C -ensemble estimator consistently outperforms other methods. In light of these empirical evidences, we regard the \mathfrak{F}_C -family as the overall best performer among the families we considered.

The choice of number m is also important. Theorem 2.2 indicates that a large enough m will guarantee the exhaustive recovering of the central subspace, regardless of the characterizing family used. In practice, however, different families require different choices of m . For the family \mathfrak{F}_C , we recommend to choose m as large as computationally feasible, because adding a new function in \mathfrak{F}_C amounts to reusing the data one more time. Since the functions in \mathfrak{F}_C are bounded and smooth, the ensemble estimator is stable as more functions are included.

The general formulation of the ensemble estimators also provides a synthesis and fresh insights for many recently developed methods. In particular, it unifies the central mean subspace [Cook and Li (2002)], the central moment subspace [Yin and Cook (2002)], Fourier transform estimators [Zhu and Zeng (2006)], dMAVE [Xia (2007)] and sliced regression [Wang and Xia (2008)] in a coherent system. Although in this paper we have focussed on MAVE ensemble and its variations, the ensemble approach can potentially be combined with any estimator of the central

mean subspaces to recover the central subspace, such as the OLS [Li and Duan (1989), Duan and Li (1991)], pHd [Li (1992), Cook (1998)] and Iterative Hessian Transformations [Cook and Li (2002, 2004)].

Finally, the ensemble approach can also be used with other characterizing families which we cannot fully explore within this paper, but which may be especially useful for some applications. One example is the wavelet basis, such as the Haar basis briefly described in Example 2.6. Such families are highly effective for handling response variables that have sharp discontinuities, which frequently arise in image analysis and pattern recognition [Donoho and Johnstone (1994)]. We leave further exploration of these possibilities to future research.

Acknowledgments. We are grateful to an Associate Editor and two referees for their thoughtful and insightful reviews.

REFERENCES

- ANTONIADIS, A. and FAN, J. (2001). Regularization of wavelet approximations. *J. Amer. Statist. Assoc.* **96** 939–955.
- BICKEL, P. J. and DOKSUM, K. A. (1981). An analysis of transformations revisited. *J. Amer. Statist. Assoc.* **76** 296–311. [MR0624332](#)
- BOX, G. E. P. and COX, D. R. (1964). An analysis of transformations. *J. Roy. Statist. Soc. Ser. B* **26** 211–252. [MR0192611](#)
- COOK, R. D. (1994). Using dimension reduction subspaces to identify important inputs in models of physical systems. In *Proceedings of the 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). Principal Hessian directions revisited (with discussion). *J. Amer. Statist. Assoc.* **93** 84–100. [MR1614584](#)
- COOK, R. D. and LI, B. (2002). Dimension reduction for conditional mean in regression. *Ann. Statist.* **30** 455–474. [MR1902895](#)
- COOK, R. D. and LI, B. (2004). Determining the dimension of iterative Hessian transformation. *Ann. Statist.* **32** 2501–2531. [MR2153993](#)
- COOK, R. D. and WEISBERG, S. (1991). Discussion of “Sliced inverse regression for dimension reduction,” by K. C. Li. *J. Amer. Statist. Assoc.* **86** 328–332.
- DONOHO, D. L. and JOHNSTONE, I. M. (1994). Ideal spatial adaptation by wavelet shrinkage. *Biometrika* **81** 425–455. [MR1311089](#)
- DUAN, N. and LI, K.-C. (1991). Slicing regression: A link-free regression method. *Ann. Statist.* **19** 505–530. [MR1105834](#)
- FAN, J., YAO, Q. and CAI, Z. (2003). Adaptive varying-coefficient linear models. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **65** 57–80. [MR1959093](#)
- FUKUMIZU, K., BACH, F. R. and JORDAN, M. I. (2009). Kernel dimension reduction in regression. *Ann. Statist.* **37** 1871–1905. [MR2533474](#)
- JOHN, J. A. and DRAPER, N. R. (1980). An alternative family of transformations. *J. Appl. Stat.* **29** 190–197.
- LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *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, K.-C. and DUAN, N. (1989). Regression analysis under link violation. *Ann. Statist.* **17** 1009–1052. [MR1015136](#)
- LI, B., WEN, S. and ZHU, L. (2008). On a projective resampling method for dimension reduction with multivariate responses. *J. Amer. Statist. Assoc.* **103** 1177–1186. [MR2462891](#)
- LI, B., ZHA, H. and CHIAROMONTE, F. (2005). Contour regression: A general approach to dimension reduction. *Ann. Statist.* **33** 1580–1616. [MR2166556](#)
- MANLY, B. F. (1976). Exponential data transformation. *The Statistician* **25** 37–42.
- SCOTT, D. W. (1992). *Multivariate Density Estimation: Theory, Practice, and Visualization*. Wiley, New York. [MR1191168](#)
- WANG, H. and XIA, Y. (2008). Sliced regression for dimension reduction. *J. Amer. Statist. Assoc.* **103** 811–821. [MR2524332](#)
- XIA, Y. (2007). A constructive approach to the estimation of dimension reduction directions. *Ann. Statist.* **35** 2654–2690. [MR2382662](#)
- XIA, Y., TONG, H., LI, W. K. and ZHU, L.-X. (2002). An adaptive estimation of dimension reduction space (with discussion). *J. R. Stat. Soc. Ser. B Stat. Methodol.* **64** 363–410. [MR1924297](#)
- 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](#)
- ZENG, P. and ZHU, Y. (2010). An integral transform method for estimating the central mean and central subspaces. *J. Multivariate Anal.* **101** 271–290. [MR2557633](#)
- ZHU, Y. and ZENG, P. (2006). Fourier methods for estimating the central subspace and the central mean subspace in regression. *J. Amer. Statist. Assoc.* **101** 1638–1651. [MR2279485](#)
- ZHU, L.-P. and ZHU, L.-X. (2009). Dimension reduction for conditional variance in regressions. *Statist. Sinica* **19** 869–883. [MR2514192](#)

DEPARTMENT OF STATISTICS
UNIVERSITY OF GEORGIA
204 STATISTICS BUILDING
ATHENS, GEORGIA 30602-1952
USA
E-MAIL: xryin@stat.uga.edu

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