

A general theory for nonlinear sufficient dimension reduction: formulation and estimation

KUANG-YAO LEE, BING LI, AND FRANCESCA CHIAROMONTE

Abstract

In this paper we introduce a general theory for nonlinear sufficient dimension reduction, and explore its ramifications and scope. This theory subsumes recent work employing reproducing kernel Hilbert spaces, and reveals many parallels between linear and nonlinear sufficient dimension reduction. Using these parallels we analyze the properties of existing methods and develop new ones. We begin by characterizing dimension reduction at the general level of σ -fields, and proceed to that of classes of functions, leading to the notions of sufficient, complete, and central dimension reduction classes. We show that, when it exists, the complete and sufficient class coincides with the central class, and can be unbiasedly and exhaustively estimated by a generalized slice inverse regression estimator (GSIR). When completeness does not hold, this estimator captures only part of the central class. However, in these cases we show that a generalized sliced average variance estimator (GSAVE) can capture a larger portion of the class. Both estimators require no numerical optimization, because they can be computed by spectral decomposition of linear operators. Finally, we compare our estimators with existing methods by simulation and on actual data sets.

Key words and phrases. Dimension reduction σ -field; Exhaustiveness; Generalized sliced average variance estimator; Generalized sliced inverse regression estimator; Heteroscedastic conditional covariance operator; Sufficient and complete dimension reduction classes; Unbiasedness.

1 Introduction

In this paper we propose a general theory for nonlinear sufficient dimension reduction (SDR), develop novel estimators, and investigate their properties under this theory. Along with these developments we also introduce a new conditional variance operator, which can potentially be used to generalize all second-order dimension reduction methods to the nonlinear case.

In its classical form, linear SDR seeks a low dimensional linear predictor that captures in full a regression relationship. Imagining a regression setting that comprises multiple predictor variables and multiple responses, let X and Y be random vectors of dimension p and q . If there is a matrix $\beta \in \mathbb{R}^{p \times d}$ with $d < p$ such that

$$Y \perp\!\!\!\perp X | \beta^T X, \tag{1}$$

then the subspace spanned by the columns of β is called a *sufficient dimension reduction (SDR) subspace*. Under mild conditions, the intersection of all such subspaces still satisfies (1), and is called the *central subspace*, denoted by $\mathcal{S}_{Y|X}$ – see Li (1991, 1992), Li and Duan (1989), Duan and Li (1991), Cook and Weisberg (1991), Cook (1994, 1998). A general condition for the existence of the central subspace is given by Yin, Li, and Cook (2008).

Several recent papers have combined sufficient dimension reduction and kernels – see Akaho (2001), Bach and Jordan (2002), Fukumizu, Bach, and Gretton (2007), Wu (2008), Wu, Liang and Mukherjee (2008), Hsing and Ren (2009), Yeh, Huang and Lee (2009), Zhu and Li (2011), and Li, Artemiou, and Li (2011). This proliferation of work, in addition to producing versatile methods for extracting nonlinear sufficient predictors, points towards a general synthesis between the notions of sufficiency at the core of SDR and the ability to encompass nonlinearity afforded by kernel mappings. To achieve this synthesis, explore its many ramifications and broad scope, and develop new estimators based on it, are the goals of this paper.

Specifically, we articulate a general formulation that comprises both linear and nonlinear SDR, and parallels the basic theoretical developments pioneered by Li (1991, 1992) and Cook (1994, 1996, 1998). This formulation allows us to study linear and nonlinear SDR comparatively and, somewhat surprisingly, to *relax* some stringent conditions required by linear SDR. For example, a linear conditional mean (Li, 1991; Cook, 1998) is no longer needed for unbiasedness, and the sufficient conditions for existence and uniqueness of the central subspace are far more general and transparent. Finally, our general formulation links linear and nonlinear SDR to the classical notions sufficiency, completeness, and minimal sufficiency, which brings insights and great clarity to the SDR theory.

Our developments, and the sections of this paper, can be summarized as follows. In Section 2, we build upon the ideas of Cook (2007) and Li, Artemiou, and Li (2011) to define an SDR σ -field as a sub σ -field \mathcal{G} of $\sigma(X)$ (the σ -field generated by X) such that $Y \perp\!\!\!\perp X|\mathcal{G}$, and the corresponding SDR class as the set of all square-integrable, \mathcal{G} -measurable functions. Under very mild conditions – much milder than the corresponding conditions for linear SDR (Yin, Li, and Cook, 2008) – we show that there exists a unique minimal σ -field $\mathcal{G}_{Y|X}$ that satisfies $Y \perp\!\!\!\perp X|\mathcal{G}_{Y|X}$, which we call the *central σ -field*. The set of all $\mathcal{G}_{Y|X}$ -measurable, square-integrable functions is named the *central class*.

In Section 3, we provide two additional definitions that generalize concepts in Cook (1998), Li, Zha, and Chiaromonte (2005), and Li, Artemiou and Li (2011): a class of functions is *unbiased* if its members are $\mathcal{G}_{Y|X}$ -measurable, and *exhaustive* if they generate $\mathcal{G}_{Y|X}$. Next, we show that the special class,

$$L_2(P_X) \ominus [L_2(P_X) \ominus L_2(P_Y)], \tag{2}$$

is unbiased, where $L_2(P_X)$ and $L_2(P_Y)$ are the spaces of square-integrable functions of X and Y . For reasons detailed in Section 3, we call this class the *regression class*.

In Section 4, we introduce the *complete dimension reduction class*: if $\mathcal{G} \subseteq \sigma(X)$ is a σ -field and for each \mathcal{G} -measurable $f \in L_2(P_X)$ we have

$$E(f(X)|Y) = 0 \text{ almost surely} \Rightarrow f(X) = 0 \text{ almost surely,}$$

then we say that the class of \mathcal{G} -measurable functions in $L_2(P_X)$ is complete. We prove that when a complete sufficient dimension reduction (CSDR) class exists it is unique and coincides with the central class. We further show that the CSDR class coincides with the regression class – which is therefore not just unbiased, but also exhaustive.

In Section 5 we establish a critical relationship between the regression class and a covariance operator linking X and Y and, based on this, we generalize sliced inverse regression (SIR; Li, 1991) to a method (GSIR) that can recover the regression class – and hence is unbiased and exhaustive under completeness. In Section 6, we consider the case where the central class is not complete, so that GSIR is unbiased but no longer exhaustive. By introducing a novel conditional variance operator, we generalize sliced average variance estimation (SAVE, Cook and Weisberg, 1991) to a method (GSAVE) that can recover a class larger than the regression class. Here, the situation is similar to that in the linear SDR wetting, where it is well known that

$$\text{SIR subspace} \subseteq \text{SAVE subspace} \subseteq \mathcal{S}_{Y|X}. \quad (3)$$

See Cook and Critchley (2000), Ye and Weiss (2003), Li, Zha, and Chiaromonte (2005), and Li and Wang (2007).

In Section 7 we develop algorithms for the sample versions of GSIR and GSAVE, and a cross-validation algorithm to determine regularizing parameters. In Section 8 we compare GSIR and GSAVE with some existing methods by simulation and on actual data sets. Section 9 contains some concluding remarks. Some highly technical developments are provided in a Supplement available online.

2 Sufficient dimension reduction σ -fields and classes

Let (Ω, \mathcal{F}, P) be a probability space, and $(\Omega_X, \mathcal{F}_X)$, $(\Omega_Y, \mathcal{F}_Y)$, and $(\Omega_{XY}, \mathcal{F}_{XY})$ be measurable spaces. For convenience, assume that $\Omega_{XY} = \Omega_X \times \Omega_Y$ and $\mathcal{F}_{XY} = \mathcal{F}_X \times \mathcal{F}_Y$. Let X , Y , and (X, Y) be random elements that take values in Ω_X , Ω_Y , and Ω_{XY} , with distributions P_X , P_Y , P_{XY} , which are dominated by σ -finite measures. Let

$$\sigma(X) = X^{-1}(\mathcal{F}_X), \quad \sigma(Y) = Y^{-1}(\mathcal{F}_Y), \quad \sigma(X, Y) = (X, Y)^{-1}(\mathcal{F}_{XY}),$$

and finally let $P_{Y|X}(\cdot|\cdot) : \mathcal{F}_Y \times \Omega_X \rightarrow \mathbb{R}$ be the conditional distribution of Y given X .

Definition 1 A sub σ -field \mathcal{G} of $\sigma(X)$ is an SDR σ -field for Y versus X if it satisfies

$$Y \perp\!\!\!\perp X | \mathcal{G}. \quad (4)$$

That is, if Y and X are independent given \mathcal{G} .

This definition is sufficiently general to accommodate the two cases of nonlinear sufficient dimension reduction that interest us the most. The first case is when $\Omega_X = \mathbb{R}^p$ and $\Omega_Y = \mathbb{R}^q$ for some positive integers p and q , and \mathcal{F}_X , \mathcal{F}_Y , and \mathcal{F}_{XY} are Borel σ -fields generated by the open sets in \mathbb{R}^p , \mathbb{R}^q , and \mathbb{R}^{p+q} . Clearly, in this case, the conditional independence (4) is a generalization of (1) for linear SDR: if we take $\mathcal{G} = \sigma(\beta^\top X)$ then (4) reduces to (1).

The second case is when X , or Y , or both of them, are random functions. In this case Definition 1 is a generalization of the linear SDR for functional data introduced by Ferré and Yao (2003), and Hsing and Ren (2009). Specifically, let $[a, b]$ be a closed interval, λ the Lebesgue measure, and $L_2(\lambda)$ the class of functions on $[a, b]$ that are square integrable with respect to λ . Let $\Omega_X = L_2(\lambda)$ and $\Omega_Y = \mathbb{R}$. In this case, each $X(\omega)$ is a function in $L_2(\lambda)$, which, depending on applications, could be, say, a growth curve or the fluctuation of a stock price. Let h_1, \dots, h_d be functions in $L_2(\lambda)$. Ferré and Yao (2003) considered the following functional dimension reduction problem

$$Y \perp\!\!\!\perp X | \langle X, h_1 \rangle_{L_2(\lambda)}, \dots, \langle X, h_d \rangle_{L_2(\lambda)}. \quad (5)$$

This generalizes linear SDR to the infinite dimensional case, but not to the nonlinear case, because $\langle X, h_1 \rangle_{L_2(\lambda)}, \dots, \langle X, h_d \rangle_{L_2(\lambda)}$ are linear in X . Hsing and Ren (2009) considered a more general setting where the sample paths $\{X_t(\omega) : t \in J\}$ need not lie within $L_2(\lambda)$. Still, their generalization is inherently linear in the same sense that problem (5) is linear. In contrast, our formulation in (4) allows an arbitrary sub σ -field of $\sigma(X)$, which need not be generated by linear functionals. Interestingly, as we will see Section 5, it is the relaxation of linearity that allows us to remove a restrictive linear conditional mean assumption used both in Ferré and Yao (2003) (Theorem 2.1), and in Hsing and Ren (2009) (assumption (IR2)).

The notion of sufficiency underlying SDR (Definition in (1) and (4)) is different from the classical notion of sufficiency, because \mathcal{G} is allowed to depend on any parameter in the joint distribution of P_{XY} . For example $\mathcal{G} = \sigma(\beta^\top X)$ depends on the parameter β (or rather, the meta-parameter $\text{span}(\beta)$) which characterizes the conditional distribution of $Y|X$. Nevertheless, both notions imply a reduction, or simplification, in the representation of a stochastic mechanism – the SDR one through a newly constructed predictor, and the classical one through a statistic. Indeed, it is partly by exploring and exploiting this similarity that we developed our theory of nonlinear SDR.

Obviously there are many sub σ -fields of X that satisfy (4), starting with $\sigma(X)$ itself – which induces no reduction. For maximal dimension reduction we seek the smallest such σ -field. As in the case of classical sufficiency, the minimal SDR σ -field does not universally exist, but exists under very mild assumptions. The next theorem gives the sufficient condition for the minimal SDR σ -field to uniquely exist. The proof echoes Bahadur (1954), which established the existence of the minimal sufficient σ -field in the classical setting.

Theorem 1 *Suppose that the family of probability measures $\{P_{X|Y}(\cdot|y) : y \in \Omega_Y\}$ is dominated by a σ -finite measure. Then there is a unique sub σ -field \mathcal{G}^* of $\sigma(X)$ such that*

1. $Y \perp\!\!\!\perp X|\mathcal{G}^*$,
2. if \mathcal{G} is a sub σ -field of $\sigma(X)$ such that $Y \perp\!\!\!\perp X|\mathcal{G}$, then $\mathcal{G}^* \subseteq \mathcal{G}$.

PROOF. Let $\Pi_y = P_{X|Y}(\cdot|y)$ and $\mathbb{P} = \{\Pi_y : y \in \Omega_Y\}$. Since \mathbb{P} is dominated by a σ -finite measure, it contains a countable subset $\mathbb{Q} = \{Q_k : k = 1, 2, \dots\}$ such that $\mathbb{Q} \equiv \mathbb{P}$, where \equiv means two families of measures dominating each other. Let $\{c_k : k = 1, 2, \dots\}$ be a sequence of positive numbers that sum to 1, and let $Q_0 = \sum_{k=1}^{\infty} c_k Q_k$. Then Q_0 is a probability measure on Ω_X such that $\{Q_0\} \equiv \mathbb{Q} \equiv \mathbb{P}$. Let $\pi_y = d\Pi_y/dQ_0$ and \mathcal{G} be a sub σ -field of $\sigma(X)$. We claim that the following statements are equivalent:

1. $Y \perp\!\!\!\perp X|\mathcal{G}$;
2. π_y is essentially measurable with respect to \mathcal{G} for all $y \in \Omega_Y$ modulo Q_0 .

Proof of 1 \Rightarrow 2. Let $B \in \mathcal{F}_X$. Then

$$E_{Q_0}(\pi_y(X)I_B(X)) = E_{\Pi_y}(I_B(X)) = E_{\Pi_y}[E_{\Pi_y}(I_B(X)|\mathcal{G})] = E_{Q_0}[E_{\Pi_y}(I_B(X)|\mathcal{G})\pi_y(X)].$$

By 1, $\Pi_y(B|\mathcal{G})$ is the same for all $y \in \Omega_Y$. Hence $\Pi_y(B|\mathcal{G}) = Q_k(B|\mathcal{G})$ for all k , which implies $\Pi_y(B|\mathcal{G}) = Q_0(B|\mathcal{G})$. Hence we can rewrite the right hand side of the above equalities as

$$E_{Q_0}[E_{Q_0}(I_B(X)|\mathcal{G})\pi_y(X)] = E_{Q_0}[I_B(X)E_{Q_0}(\pi_y(X)|\mathcal{G})].$$

Thus the following equality holds for all $B \in \mathcal{F}_X$:

$$E_{Q_0}(\pi_y(X)I_B(X)) = E_{Q_0}[I_B(X)E_{Q_0}(\pi_y(X)|\mathcal{G})],$$

which implies $\pi_y(X) = E_{Q_0}(\pi_y(X)|\mathcal{G})$ a.s. Q_0 .

Proof of 2 \Rightarrow 1. For any $A \in \mathcal{G}$,

$$\begin{aligned} E_{\Pi_y}[E_{Q_0}(I_B(X)|\mathcal{G})I_A(X)] &= E_{Q_0}[E_{Q_0}(I_B(X)|\mathcal{G})I_A(X)\pi_y(X)] \\ &= E_{Q_0}[I_B(X)I_A(X)E_{Q_0}(\pi_y(X)|\mathcal{G})]. \end{aligned}$$

By 2, $E_{Q_0}(\pi_y(X)|\mathcal{G}) = \pi_y(X)$. Hence the right hand side becomes

$$E_{Q_0}[I_B(X)I_A(X)\pi_y(X)] = E_{\Pi_y}[I_B(X)I_A(X)] = \Pi_y(X \in A \cap B).$$

Thus $E_{Q_0}(I_B(X)|\mathcal{G}) = Q_0(B|\mathcal{G})$ is the conditional probability $\Pi_y(B|\mathcal{G})$, which means $\Pi_y(B|\mathcal{G})$ does not depend on y . That is, 1 holds.

Now let \mathcal{G}^* be the intersection of all SDR σ -fields \mathcal{G} . Then \mathcal{G}^* is itself a σ -field. Moreover, since π_y is essentially measurable with respect to all SDR σ -fields for all $y \in \Omega_Y$, it is also essentially measurable with respect to \mathcal{G}^* for all $y \in \Omega_Y$. Consequently, \mathcal{G}^* is itself an SDR σ -field, which implies that it is also the smallest SDR σ -field. If \mathcal{G}^{**} is another smallest SDR σ -field, then we know $\mathcal{G}^* \subseteq \mathcal{G}^{**}$ and $\mathcal{G}^{**} \subseteq \mathcal{G}^*$. Thus \mathcal{G}^* is unique. \square

We can now naturally introduce the following definition:

Definition 2 *Suppose that the class of probability measures $\{P_{X|Y}(\cdot|y) : y \in \Omega_Y\}$ on Ω_X is dominated by a σ -finite measure. Then we call the σ -field \mathcal{G}^* in Theorem 1 the central σ -field for Y versus X , and denote it by $\mathcal{G}_{Y|X}$.*

Notably, this set up characterizes dimension reduction solely in terms of conditional independence. However, explicitly turning to functions and introducing an additional mild assumption of square integrability are very consequential for further development, because they allow us to work with structures such as orthogonality and projection.

Let $L_2(P_{XY})$, $L_2(P_X)$, and $L_2(P_Y)$ be the spaces of functions defined on Ω_{XY} , Ω_X , and Ω_Y that are square-integrable with respect to P_{XY} , P_X , and P_Y , respectively. Since constants are irrelevant for dimension reduction, we assume throughout that all functions in $L_2(P_X)$, $L_2(P_Y)$ and $L_2(P_{XY})$ have mean 0. Given a sub σ -field \mathcal{G} of $\sigma(X, Y)$, we use $\mathfrak{M}_{\mathcal{G}}$ to denote the class of all functions f in $L_2(P_{XY})$ such that $f(X)$ is \mathcal{G} -measurable. If \mathcal{G} is generated by a random vector, say X , then we use \mathfrak{M}_X to abbreviate $\mathfrak{M}_{\sigma(X)}$. It can be easily shown that, for any \mathcal{G} , $\mathfrak{M}_{\mathcal{G}}$ is a linear subspace of $L_2(P_{XY})$.

Definition 3 *Let \mathcal{G} be an SDR σ -field and $\mathcal{G}_{Y|X}$ be the central σ -field. Then $\mathfrak{M}_{\mathcal{G}}$ is called an SDR class, and $\mathfrak{M}_{\mathcal{G}_{Y|X}}$ is called the central class. The latter class is denoted by $\mathfrak{S}_{Y|X}$.*

The central class, comprising square-integrable functions that are measurable with respect to the central σ -field $\mathcal{G}_{Y|X}$, represents our generalization of the central space $\mathcal{S}_{Y|X}$ defined in linear SDR (see Introduction).

3 Unbiasedness and exhaustiveness

In linear SDR, the goal is to find a set of vectors that span $\mathcal{S}_{Y|X}$. If a matrix γ satisfies $\text{span}(\gamma) \subseteq \mathcal{S}_{Y|X}$, we say that γ is unbiased (Cook, 1998). If $\text{span}(\gamma) = \mathcal{S}_{Y|X}$, we say that γ

is exhaustive (Li, Zha, and Chiaromonte, 2005). Note that when $\text{span}(\gamma) \subseteq \mathcal{S}_{Y|X}$, $\gamma^\top X$ is a linear function of $\beta^\top X$, where β is any matrix such that $\text{span}(\beta) = \mathcal{S}_{Y|X}$; if $\text{span}(\gamma) = \mathcal{S}_{Y|X}$, then $\gamma^\top X$ is an injective linear transformation of $\beta^\top X$. In the nonlinear setting, we follow the same logic but remove the linear requirement. Part of the following definition was given in Li, Artemiou, and Li (2011).

Definition 4 *A class of functions in $L_2(P_X)$ is unbiased for $\mathfrak{S}_{Y|X}$ if its members are $\mathcal{G}_{Y|X}$ -measurable, and exhaustive for $\mathfrak{S}_{Y|X}$ if its members generate $\mathcal{G}_{Y|X}$.*

Next, we look into what type of functions are unbiased. The lemma below provides a characterization of the orthogonal complement of $\mathfrak{M}_{\mathcal{G}}$ that will be used many times in the subsequent development. Its proof is essentially the definition of the conditional expectation, and is omitted.

Lemma 1 *Suppose U is a random element defined on (Ω, \mathcal{F}) , \mathcal{G} is a sub σ -field of $\sigma(U)$, and $f \in L_2(P_U)$. Then f is orthogonal to $\mathfrak{M}_{\mathcal{G}}$ ($f \perp \mathfrak{M}_{\mathcal{G}}$) if and only if $E[f(U)|\mathcal{G}] = 0$.*

Note that \perp and \perp have different meanings: the former means independence; the latter means orthogonality. For two subspaces, say \mathcal{S}_1 and \mathcal{S}_2 , of a generic Hilbert space \mathcal{H} , we use $\mathcal{S}_1 \ominus \mathcal{S}_2$ to denote the subspace $\mathcal{S}_1 \cap \mathcal{S}_2^\perp$. The following theorem explicitly specifies a class of functions, which we call *regression class*, that is unbiased for $\mathfrak{S}_{Y|X}$.

Theorem 2 *If the family $\{\Pi_y : y \in \Omega_Y\}$ is dominated by a σ -finite measure, then*

$$L_2(P_X) \ominus [L_2(P_X) \ominus L_2(P_Y)] \subseteq \mathfrak{S}_{Y|X}. \quad (6)$$

PROOF. It is equivalent to show that $L_2(P_X) \ominus \mathfrak{S}_{Y|X} \subseteq L_2(P_X) \ominus L_2(P_Y)$. If $f \in L_2(P_X) \ominus \mathfrak{S}_{Y|X}$, then, by Lemma 1, $E[f(X)|\mathcal{G}_{Y|X}] = 0$. Since $\mathcal{G}_{Y|X}$ is a sufficient σ -field,

$$E[f(X)|Y] = E[E(f(X)|Y, \mathcal{G}_{Y|X})|Y] = E[E(f(X)|\mathcal{G}_{Y|X})|Y] = 0.$$

By Lemma 1 again, $f \perp \mathfrak{M}_Y$. Because $\mathfrak{M}_Y = L_2(P_Y)$, we have $f \in L_2(P_X) \ominus L_2(P_Y)$. \square

The intuition behind the term “regression class” is that $L_2(P_X) \ominus L_2(P_Y)$ resembles the residual in a regression problem; thus $L_2(P_X) \ominus [L_2(P_X) \ominus L_2(P_Y)]$ is simply the orthogonal complement of the “residual class.” Henceforth we write the regression class as $\mathfrak{C}_{Y|X}$.

4 Complete and sufficient dimension reduction classes

After showing that the regression class (2) is an unbiased, we investigate under what conditions it is also exhaustive for the central class $\mathfrak{S}_{Y|X}$. To this end we need to introduce the notion of complete classes of functions in $L_2(P_X)$.

Definition 5 Let $\mathcal{G} \subseteq \sigma(X)$ be a sub σ -field. The class $\mathfrak{M}_{\mathcal{G}}$ is said to be complete if, for any $g \in \mathfrak{M}_{\mathcal{G}}$,

$$E[g(X)|Y] = 0 \text{ a.s. } P \Rightarrow g(X) = 0 \text{ a.s. } P.$$

Again there are similarities and differences between completeness as defined here and in the classical setting. A complete and sufficient statistic in the classical setting is a rather restrictive concept, often associated with exponential families, the uniform distribution, or the order statistics. In contrast, completeness here is a rather general concept. To demonstrate this point, in the next two propositions we give two examples of complete and sufficient dimension reduction classes. In particular, the first shows that if Y is related to X through *any* regression model with additive error, then the subspace of $L_2(P_X)$ determined by the regression function is a complete and sufficient dimension reduction class. In the following, $[L_2(P_X)]^q$ denotes the q -fold Cartesian product of $L_2(P_X)$.

Proposition 1 Suppose there exists a function $h \in [L_2(P_X)]^q$ such that

$$Y = h(X) + \varepsilon, \tag{7}$$

where $\varepsilon \perp\!\!\!\perp X$ and $E(\varepsilon) = 0$. Then $\mathfrak{M}_{h(X)}$ is a complete and sufficient dimension reduction class for Y versus X .

Note that, since $L_2(P_X)$ is centered, we have implicitly assumed that $E[h(X)] = 0$ (and hence $E(Y) = 0$). However, this does not entail any real loss of generality, because the proof below can be easily modified for the case where $L_2(P_X)$ is not centered.

PROOF. Suppose $m \in \mathfrak{M}_{h(X)}$ and $E[m(X)|Y] = 0$ a.s. P . Then there is a measurable function $g : \mathbb{R}^q \rightarrow \mathbb{R}$ such that $m = g \circ h$. Let $U = h(X)$. Then $E(g(U)|Y) = 0$ a.s. P . By Lemma 1, for any $f \in L_2(P_Y)$, we have $E[g(U)f(Y)] = 0$. In particular, $E[g(U)e^{it^\top Y}] = 0$. Because $U \perp\!\!\!\perp \varepsilon$, this implies

$$E[g(U)e^{it^\top U}]E(e^{it^\top \varepsilon}) = E[g(U)e^{it^\top U}e^{it^\top \varepsilon}] = E[g(U)e^{it^\top Y}] = 0.$$

Hence $E[g(U)e^{it^\top U}] = 0$. By the uniqueness of inverse Fourier transformation we see that $g(U) = 0$ a.s. P , which implies $m(X) = (g \circ h)(X) = 0$ a.s. P . \square

The expression in (7) covers many useful models in Statistics and Econometrics. For example, any homoscedastic parametric or nonparametric regression, such as the single index and the multiple index models (Ichimura and Lee, 1991; Härdle, Hall, and Ichimura 1993; Yin, Cook, and Li, 2008), are special cases of (7). Thus, complete and sufficient

dimension reduction classes exist for all those settings. The next proposition considers a type of inverse regression model, in which X is transformed into two components, one of which is related to Y by an inverse linear regression model, and the other independent of the rest of the data.

Proposition 2 *Suppose $q < p$, Ω_Y has a nonempty interior, and P_Y is dominated by the Lebesgue measure on \mathbb{R}^q . Suppose there exist functions $g \in [L_2(P_X)]^q$ and $h \in [L_2(P_X)]^{p-q}$ such that*

1. $g(X) = Y + \varepsilon$, where $Y \perp \varepsilon$, and $\varepsilon \sim N(0, \Sigma)$;
2. $\sigma(g(X), h(X)) = \sigma(X)$;
3. $h(X) \perp (Y, g(X))$;
4. the induced measure $P_{X \circ g^{-1}}$ is dominated by the Lebesgue measure on \mathbb{R}^q .

Then $\mathfrak{M}_{g(X)}$ is a complete sufficient dimension reduction class for Y versus X .

PROOF. Assumption 3 implies $Y \perp h(X)|g(X)$, which, by assumption 2, implies $Y \perp X|g(X)$. That is, $\mathfrak{M}_{g(X)}$ is an SDR class. Let $u \in \mathfrak{M}_{g(X)}$. Then $u = v \circ g$ for some measurable function $v : \mathbb{R}^q \rightarrow \mathbb{R}$. Let $U = g(X)$. Suppose that $E[v(U)|Y] = 0$ almost surely P . Because $Y \perp \varepsilon$, this implies $P_Y(\{y : Ev(y + \varepsilon) = 0\}) = 1$. In other words

$$\int_{\mathbb{R}^q} v(t) \frac{1}{(2\pi)^{q/2} |\Sigma|^{1/2}} e^{-(t-y)^\top \Sigma^{-1} (t-y)/2} dt = 0$$

a.s. P_Y . This implies

$$\int v(t) e^{-t^\top \Sigma^{-1} t/2} e^{y^\top \Sigma^{-1} t} dt = 0 \Rightarrow \int v(\Sigma s) e^{-s^\top \Sigma s/2} e^{y^\top s} ds = 0$$

a.s. P_Y , where $s = \Sigma^{-1} t$. Because Ω_Y contains an open set in \mathbb{R}^q and the above function of y is analytic, by the analytic continuation theorem, the above function is 0 everywhere on \mathbb{R}^q . Hence, by the uniqueness of inverse Laplace transformation, we have

$$v(\Sigma s) e^{-s^\top \Sigma s/2} = 0 \text{ almost surely } \lambda,$$

where λ is the Lebesgue measure on \mathbb{R}^q . But, because $e^{-s^\top \Sigma s/2} > 0$, we have $v(\Sigma s) = 0$ a.s. λ or equivalently $v(t) = 0$ a.s. λ . By the change of variable theorem,

$$\int_{v \circ g(x) \neq 0} dP_X = \int_{v(t) \neq 0} dP_{X \circ g^{-1}}.$$

By assumption 4, $P_{x \circ g^{-1}} \ll \lambda$. Hence the above integral is 0, implying $v \circ g(x) = 0$ a.s. P_X , or, equivalently, $v \circ g(X) = 0$ a.s. P . \square

Inverse regressions of this type are considered in Cook (2007), Cook and Forzani (2009), and Cook, Li, and Chiaromonte (2010) for linear SDR. The above two propositions show that a complete and sufficient dimension reduction class exists for a reasonably wide range of problems, including forward and inverse regressions of very general, non-parameterized form. The next theorem shows that when a complete and sufficient dimension reduction class exists, it is unique and coincides with the central class. Once again, the situation here echoes that in classical theory, where a complete and sufficient statistic, if it exists, coincides with the minimal sufficient statistic – see Lehmann (1981).

Theorem 3 *Suppose $\{\Pi_y : y \in \Omega_Y\}$ is dominated by a σ -finite measure, and \mathcal{G} is a sub σ -field of $\sigma(X)$. If $\mathfrak{M}_{\mathcal{G}}$ is a complete and sufficient dimension reduction class, then*

$$\mathfrak{M}_{\mathcal{G}} = \mathfrak{C}_{Y|X} = \mathfrak{S}_{Y|X}.$$

PROOF. If $f \perp \mathfrak{C}_{Y|X}$, then by Lemma 1, $E(f|Y) = 0$ which, because $\mathfrak{M}_{\mathcal{G}}$ is sufficient, implies

$$E[E(f|\mathcal{G})|Y] = 0.$$

Because $\mathfrak{M}_{\mathcal{G}}$ is complete and because $E(f|\mathcal{G}) \in \mathfrak{M}_{\mathcal{G}}$, we have $E(f|\mathcal{G}) = 0$. By Lemma 1, this implies $f \perp \mathfrak{M}_{\mathcal{G}}$. Thus we have proved $\mathfrak{M}_{\mathcal{G}} \subseteq \mathfrak{C}_{Y|X}$. However, by Theorem 2 we know that $\mathfrak{C}_{Y|X} \subseteq \mathfrak{S}_{Y|X} \subseteq \mathfrak{M}_{\mathcal{G}}$. This proves the desired equality. \square

5 Generalizations of SIR and their population-level properties

From the previous developments we see that the subspace $L_2(P_X) \ominus L_2(P_Y)$ of $L_2(P_X)$ plays a critical role in nonlinear SDR. Its orthogonal complement in $L_2(P_X)$ coincides with the central class $\mathfrak{S}_{Y|X}$ under completeness, and even without completeness it is guaranteed to be inside $\mathfrak{S}_{Y|X}$. It turns out that this subspace can be expressed as the range of a certain bounded linear operators. This representation ensures that estimation procedures can rely on simple spectral decompositions, rather than complicated numerical optimizations. We first introduce some covariance operators which are the building block of this approach.

5.1 Covariance operators

Since constants are irrelevant here (for example, f and $f + 3$ can be considered as the same function), we will speak of set relations modulo constants. If A and B are sets, then we

say $A \subseteq B$ modulo constants if for each $f \in A$ there is $c \in \mathbb{R}$ such that $f + c \in B$. We say that A is a dense subset of B modulo constants if (i) $A \subseteq B$ modulo constants and (ii) for each $f \in B$, there is a sequence $\{f_n\} \subseteq A$ and a sequence of constants $\{c_n\} \subseteq \mathbb{R}$ such that $\{f_n + c_n\} \subseteq A$ and $f_n + c_n \rightarrow f$ in the topology for B . Let \mathcal{H}_X and \mathcal{H}_Y be Hilbert spaces of functions of X and Y satisfying the conditions:

- (A) \mathcal{H}_X and \mathcal{H}_Y are dense subsets of $L_2(P_X)$ and $L_2(P_Y)$ modulo constants;
- (B) there are constants $C_1 > 0$ and $C_2 > 0$ such that $\text{var}[f(X)] \leq C_1 \|f\|_{\mathcal{H}_X}^2$ and $\text{var}[g(Y)] \leq C_2 \|g\|_{\mathcal{H}_Y}^2$.

Although we will later take \mathcal{H}_X and \mathcal{H}_Y to be reproducing kernel Hilbert spaces (RKHS), our theory is not restricted to such spaces. In particular, we do not require the evaluation functionals (such as $f \mapsto f(x)$ from \mathcal{H}_X to \mathbb{R}) to be continuous.

For two generic Hilbert spaces \mathcal{H}_1 and \mathcal{H}_2 , let $\mathcal{B}(\mathcal{H}_1, \mathcal{H}_2)$ denote the class of all bounded linear operators from \mathcal{H}_1 to \mathcal{H}_2 and let $\mathcal{B}(\mathcal{H}_1)$ abbreviate $\mathcal{B}(\mathcal{H}_1, \mathcal{H}_1)$. We denote the range of a linear operator A by $\text{ran}A$, the kernel of A by $\ker A$, and the closure of $\text{ran}A$ by $\overline{\text{ran}A}$. Under assumption (B), the symmetric bilinear form $u : \mathcal{H}_X \times \mathcal{H}_X \rightarrow \mathbb{R}$ defined by $u(f, g) = \text{cov}[f(X), g(X)]$ is bounded and thus induces an operator $M_{XX} \in \mathcal{B}(\mathcal{H}_X)$ that satisfies $\langle f, M_{XX} g \rangle_{\mathcal{H}_X} = u(f, g)$. Similarly, the bounded bilinear form $(f, g) \mapsto \text{cov}[f(Y), g(Y)]$ from $\mathcal{H}_Y \times \mathcal{H}_Y$ to \mathbb{R} defines an operator $M_{YY} \in \mathcal{B}(\mathcal{H}_Y)$. Let \mathcal{G}_X and \mathcal{G}_Y represent the subspaces $\overline{\text{ran}M_{XX}}$ and $\overline{\text{ran}M_{YY}}$.

Definition 6 *Suppose conditions (A) and (B) are satisfied. We define the covariance operators $\Sigma_{XX} : \mathcal{G}_X \rightarrow \mathcal{G}_X$, $\Sigma_{YY} : \mathcal{G}_Y \rightarrow \mathcal{G}_Y$, and $\Sigma_{YX} : \mathcal{G}_X \rightarrow \mathcal{G}_Y$ through the relations:*

$$\langle f, \Sigma_{XX} g \rangle_{\mathcal{G}_X} = \langle f, g \rangle_{L_2(P_X)}, \quad \langle f, \Sigma_{YY} g \rangle_{\mathcal{G}_Y} = \langle f, g \rangle_{L_2(P_Y)}, \quad \langle f, \Sigma_{YX} g \rangle_{\mathcal{G}_Y} = \langle f, g \rangle_{L_2(P_Y)}.$$

These operators are essentially the same as those introduced by Fukumizu, Back, and Jordan (2004, 2009), except that here we do not assume \mathcal{H}_X and \mathcal{H}_Y to be RKHS. By Baker (1973), Theorem 1, there is a unique operator $R_{YX} \in \mathcal{B}(\mathcal{G}_X, \mathcal{G}_Y)$ such that $\Sigma_{YX} = \Sigma_{YY}^{1/2} R_{YX} \Sigma_{XX}^{1/2}$. We call R_{YX} the *correlation operator*. In order to connect these operators with the central class, which is an $L_2(P_X)$ -object, we need to extend the domains of $\Sigma_{XX}^{1/2}$ and $\Sigma_{YY}^{1/2}$ from \mathcal{G}_X and \mathcal{G}_Y to $L_2(P_X)$ and $L_2(P_Y)$. The following extension theorem is important and nontrivial, but since the material presented here can be understood without its proof we relegate it to the online Supplement to this article.

Theorem 4 *Under assumptions (A) and (B), there exist unique isomorphisms*

$$\tilde{\Sigma}_{XX}^{1/2} : L_2(P_X) \rightarrow \mathcal{G}_X, \quad \tilde{\Sigma}_{YY}^{1/2} : L_2(P_Y) \rightarrow \mathcal{G}_Y$$

that agree with $\Sigma_{XX}^{1/2}$ and $\Sigma_{YY}^{1/2}$ on \mathcal{G}_X and \mathcal{G}_Y in the sense that, for all $f \in \mathcal{G}_X$ and $g \in \mathcal{G}_Y$,

$$\tilde{\Sigma}_{XX}^{1/2}(f - Ef) = \Sigma_{XX}^{1/2}f, \quad \tilde{\Sigma}_{YY}^{1/2}(g - Eg) = \Sigma_{YY}^{1/2}g.$$

Furthermore, for any $f \in L_2(P_X)$, $g \in L_2(P_Y)$ we have

$$\langle \tilde{\Sigma}_{YY}^{1/2}g, R_{YX}\tilde{\Sigma}_{XX}^{1/2}f \rangle_{\mathcal{G}_Y} = \text{cov}[g(Y), f(X)]. \quad (8)$$

The easiest way to understand equality (8) is through the special case where $f = f' - E(f')$, $g = g' - E(g')$ where $f' \in \mathcal{G}_X$, $g' \in \mathcal{G}_Y$. In this case,

$$\langle \tilde{\Sigma}_{YY}^{1/2}g, R_{YX}\tilde{\Sigma}_{XX}^{1/2}f \rangle_{\mathcal{G}_Y} = \langle \Sigma_{YY}^{1/2}g', R_{YX}\Sigma_{XX}^{1/2}f' \rangle_{\mathcal{G}_Y} = \langle g', \Sigma_{YX}f' \rangle_{\mathcal{G}_Y} = \text{cov}[f(X), g(Y)].$$

The theorem also implies that, for all $f, g \in L_2(P_X)$ and $s, t \in L_2(P_Y)$,

$$\begin{aligned} \langle \tilde{\Sigma}_{XX}^{1/2}g, \tilde{\Sigma}_{XX}^{1/2}f \rangle_{\mathcal{G}_X} &= \langle g, f \rangle_{L_2(P_X)} = \text{cov}[g(X), f(X)], \\ \langle \tilde{\Sigma}_{YY}^{1/2}s, \tilde{\Sigma}_{YY}^{1/2}t \rangle_{\mathcal{G}_Y} &= \langle s, t \rangle_{L_2(P_Y)} = \text{cov}[s(Y), t(Y)]. \end{aligned}$$

5.2 Generalized SIR

The results of the last subsection allow us to characterize $L_2(P_X) \ominus L_2(P_Y)$ in terms of extended covariance operators, which is the key to developing its estimator. Recall that classical SIR (Li, 1991) for linear SDR is based on the matrix

$$[\text{var}(X)]^{-1}\text{var}[E(X|Y)]. \quad (9)$$

Under the linear conditional mean assumption requiring that $E(X|\beta^T X)$ be linear in X for any matrix β spanning $\mathcal{S}_{Y|X}$, the re-scaled “inverse” conditional mean $[\text{var}(X)]^{-1}E(X|Y)$ is contained in this space. To generalize this to the nonlinear setting, we first introduce a conditional mean operator.

Definition 7 We call the operator $\tilde{\Sigma}_{YY}^{-1/2}R_{YX}\tilde{\Sigma}_{XX}^{1/2} : L_2(P_X) \rightarrow L_2(P_Y)$ the conditional expectation operator, and denote it by $E_{X|Y}$.

The relation between the conditional expectation operator and conditional expectations is elucidated by the next proposition, which is followed by an important corollary.

Proposition 3 Under conditions (A) and (B), we have

1. for any $f \in L_2(P_X)$, $E_{X|Y}f = E(f(X)|Y)$;
2. for any $g \in L_2(P_Y)$, $E_{X|Y}^*g = E(g(Y)|X)$.

PROOF. For any $g \in L_2(P_Y)$,

$$\langle E_{X|Y}f, g \rangle_{L_2(P_Y)} = \langle \tilde{\Sigma}_{YY}^{-1/2} R_{YX} \tilde{\Sigma}_{XX}^{1/2} f, g \rangle_{L_2(P_Y)} = \langle R_{YX} \tilde{\Sigma}_{XX}^{1/2} f, \tilde{\Sigma}_{YY}^{1/2} g \rangle_{\mathcal{H}_Y} = \text{cov}(f(X), g(Y)),$$

where the last equality follows from (8). Hence $\text{cov}(f(X) - (E_{X|Y}f)(Y), g(Y)) = 0$. By the definition of conditional expectation, $E_{X|Y}f = E(f(X)|Y)$, which proves 1. Assertion 2 follows from the fact that $\tilde{\Sigma}_{YY}^{-1/2}$ and $\tilde{\Sigma}_{XX}^{1/2}$ are isomorphisms, and $R_{YX}^* = R_{XY}$. \square

Corollary 1 *Under conditions (A) and (B), for any $f, g \in L_2(P_X)$,*

$$\langle g, E_{X|Y}^* E_{X|Y} f \rangle_{L_2(P_X)} = \text{cov}[E(g(X)|Y), E(f(X)|Y)]. \quad (10)$$

Moreover, $E_{X|Y}^* E_{X|Y} \in \mathcal{B}(L_2(P_X))$, and its norm is no greater than 1.

PROOF. We have

$$\langle g, E_{X|Y}^* E_{X|Y} f \rangle_{L_2(P_X)} = \langle E_{X|Y} g, E_{X|Y} f \rangle_{L_2(P_Y)} = \langle E(g(X)|Y), E(f(X)|Y) \rangle_{L_2(P_Y)},$$

which is the right hand side of (10). Moreover, since $\tilde{\Sigma}_{XX}^{1/2}$ is isomorphic, we have

$$E_{X|Y}^* E_{X|Y} = (\tilde{\Sigma}_{YY}^{-1/2} R_{YX} \tilde{\Sigma}_{XX}^{1/2})^* (\tilde{\Sigma}_{YY}^{-1/2} R_{YX} \tilde{\Sigma}_{XX}^{1/2}) = \tilde{\Sigma}_{XX}^{-1/2} R_{XY} R_{YX} \tilde{\Sigma}_{XX}^{1/2}.$$

Hence $\|E_{X|Y}^* E_{X|Y}\| \leq \|\tilde{\Sigma}_{XX}^{-1/2}\| \|R_{XY}\| \|R_{YX}\| \|\tilde{\Sigma}_{XX}^{1/2}\|$. Because $\tilde{\Sigma}_{XX}^{1/2}$ and $\tilde{\Sigma}_{XX}^{-1/2}$ are isomorphisms, their norms are both 1. By Baker (1972, Theorem 1), $\|R_{YX}\| \leq 1$. Hence $\|E_{X|Y}^* E_{X|Y}\| \leq 1$. \square

From this corollary we see that the quadratic form

$$f \mapsto \langle f, E_{X|Y}^* E_{X|Y} f \rangle_{L_2(P_X)}, \quad L_2(P_X) \times L_2(P_X) \rightarrow \mathbb{R}$$

generalizes the matrix $\text{var}[E(X|Y)]$ of the linear case, which is the essential ingredient of SIR for linear SDR. It is then not surprising that the operator $E_{X|Y}^* E_{X|Y}$ is closely connected to the central class for nonlinear SDR, as shown in the following theorem.

Theorem 5 *If conditions (A) and (B) are satisfied and $\mathfrak{S}_{Y|X}$ is complete, then*

$$\overline{\text{ran}}(E_{X|Y}^* E_{X|Y}) = \mathfrak{S}_{Y|X}.$$

PROOF. By Lemma 1, $f \in \mathfrak{C}_{Y|X}$ if and only if $f \in L_2(P_X)$ and $E(f|Y) = 0$. By Proposition 3, this happens if and only if $f \in \ker E_{X|Y}$. This shows $\ker E_{X|Y} = \mathfrak{C}_{Y|X}^\perp$. However, because $\ker(E_{X|Y}) = \ker(E_{X|Y}^* E_{X|Y})$, we have

$$\overline{\text{ran}}(E_{X|Y}^* E_{X|Y}) = [\ker(E_{X|Y}^* E_{X|Y})]^\perp = (\ker E_{X|Y})^\perp = (\mathfrak{C}_{Y|X}^\perp)^\perp = \mathfrak{C}_{Y|X}.$$

Since $\mathfrak{S}_{Y|X}$ is complete, we have $\mathfrak{C}_{Y|X} = \mathfrak{S}_{Y|X}$, as desired. \square

Note that, unlike in classical SIR for linear SDR, here we do not have to consider an analogue to the rescaling $[\text{var}(X)]^{-1}$ in (9). This is because the $L_2(P_X)$ -inner product absorbs the marginal variance in the predictor vector. We refer to the sample estimator based on $\overline{\text{ran}}(E_{X|Y}^* E_{X|Y})$ (see Section 7.2) as *generalized SIR*, or GSIR. The GSIR estimator is related to kernel canonical component analysis (KCCA) introduced by Bach and Jordan (2002) (see also Fukumizu, Bach, and Gretton (2007)). In Section 7.2 we will explore similarities and differences between these two methods.

5.3 Kernel SIR

We now turn to another nonlinear SDR estimator, which was proposed by Wu (2008) and further studied by Yeh, Huang, and Lee (2009), called *kernel sliced inverse regression* (KSIR). In our setting, the population-level description of this estimator is as follows. Let \mathcal{H}_X be a Hilbert space satisfying (A) and (B) (in this case an RKHS, but this assumption is unnecessary). Let $T : \mathcal{H}_X \rightarrow L_2(P_X)$ be the centering transformation $T(f) = f - E(f)$. Let J_1, \dots, J_h be a partition of Ω_Y , and let $\mu_1, \dots, \mu_h \in \overline{\text{ran}} T$ be the Riesz representations of the linear functionals

$$T_j : \overline{\text{ran}} T \rightarrow \mathbb{R}, \quad g \mapsto E(g(X)|Y \in J_i), \quad i = 1, \dots, h.$$

In our language, KSIR uses (the sample version of) the subspace $\text{span}(\Sigma_{X|X}^{-1}\mu_1, \dots, \Sigma_{X|X}^{-1}\mu_h)$ to estimate $\mathfrak{S}_{Y|X}$. The next theorem shows that any such representation must be a member of $\mathfrak{C}_{Y|X}$, and thus of $\mathfrak{S}_{Y|X}$ (since $\mathfrak{C}_{Y|X} \subseteq \mathfrak{S}_{Y|X}$) – which implies that KSIR is unbiased.

Theorem 6 *If (A) and (B) hold, then $\mu_j \in \mathfrak{C}_{Y|X}$.*

PROOF. By condition (A), $\overline{\text{ran}} T = L_2(P_X)$. If $f \in L_2(P_X) \ominus L_2(P_Y) \subseteq \overline{\text{ran}} T$, then, by Lemma 1, $E(f|Y) = 0$. Hence $\langle f, \mu_i \rangle_{L_2(P_X)} = E[f(X)|Y \in J_i] = 0$. \square

Yeh, Huang, and Lee (2009) give another unbiasedness proof for KSIR, but they assume that the spanning functions of \mathcal{H}_X , say f_1, \dots, f_m , satisfy the linear conditional mean assumption. That is, for any $f \in \mathcal{H}_X$, $E(f|f_1, \dots, f_m)$ has the form $c_0 + c_1 f_1 + \dots + c_m f_m$ for some $c_0, \dots, c_m \in \mathbb{R}$. This condition is an analogue of the linear conditional mean assumption for linear SDR (see, for example, Li (1991) and Cook and Li (2002)). Interestingly, our result no longer relies on this assumption. The reason Yeh, Huang and Lee need the assumption in the first place is that they define the central class (Definition 1 of Yeh, Huang,

and Lee, 2009) as the linear subspace spanned by h_1, \dots, h_d in $\text{span}(f_1, \dots, f_m)$ such that

$$Y \perp\!\!\!\perp X | h_1(X), \dots, h_d(X); \quad (11)$$

whereas we define the central class as the class of all measurable functions of h_1, \dots, h_d . Indeed, in the nonlinear setting there is no reason to restrict to this linear span formulation, since the conditional independence (11) only relies on the σ -field generated by h_1, \dots, h_d .

6 Beyond completeness: generalized SAVE

We now turn to the more general problem of estimating the central class when it is not complete, in which case the regression class may be a proper subset of the central class. We will generalize SAVE (Cook and Weisberg, 1991) to the nonlinear case and show that it can recover functions beyond the regression class.

The setting here is different from that for GSIR in two respects. First, since we now deal with the location-invariant quantity $f(X) - E[f(X)|Y]$, we no longer need to define the conditional mean operator through the centered L_2 -spaces $L_2(P_Y)$ and $L_2(P_X)$. Second, we now define relevant operators through L_2 -spaces instead of RKHSs, which is more convenient in this context. Let $L'_2(P_X)$ and $L'_2(P_Y)$ denote the noncentered L_2 -spaces. Define the noncentered conditional mean operator $E'_{X|Y} : L'_2(P_X) \rightarrow L'_2(P_Y)$ through

$$\langle g, E'_{X|Y} f \rangle_{L'_2(P_Y)} = E(g(Y)f(X)), \quad f \in L'_2(P_X), \quad g \in L'_2(P_Y). \quad (12)$$

By the same argument of Proposition 3, $E'_{X|Y} f = E(f(X)|Y)$. To generalize SAVE, we introduce a new type of conditional variance operator.

Definition 8 *For each $y \in \Omega_Y$, the bilinear form*

$$L_2(P_X) \times L_2(P_X) \rightarrow \mathbb{R}, \quad (f, g) \mapsto (E'_{X|Y}(fg) - E'_{X|Y} f E'_{X|Y} g)(y)$$

uniquely defines an operator $V_{X|Y}(y) \in \mathcal{B}(L_2(P_X))$ via the Riesz representation. We call the random operator

$$V_{X|Y} : \Omega_Y \rightarrow \mathcal{B}(L_2(P_X)), \quad y \mapsto V_{X|Y}(y)$$

the heteroscedastic conditional variance operator given Y .

The operator $V_{X|Y}$ is different from the conditional variance operator $\Sigma_{X|Y}$ introduced by Fukumizu, Bach and Jordan (2004, 2009). In a sense, $\Sigma_{X|Y}$ is a generalization of $E[\text{var}(X|Y)]$ rather than $\text{var}(X|Y)$, because $\langle f, \Sigma_{X|Y} f \rangle_{\mathcal{H}_X} = E[\text{var}(f(X)|Y)]$. Note that $E[\text{var}(f(X)|Y)]$ becomes $\text{var}(f(X)|Y)$ only when the latter is nonrandom. So $\Sigma_{X|Y}$ might

be called a *homoscedastic* conditional variance operator. In contrast, $\langle f, V_{X|Y}f \rangle_{L_2(P_X)}$ gives directly the conditional variance $\text{var}[f(X)|Y]$, hence the term heteroscedastic conditional variance operator. Here, we should also stress that $E'_{X|Y}$ is defined between noncentered $L'_2(P_X)$ and $L'_2(P_Y)$, whereas $V_{X|Y}(y)$ is defined between centered $L_2(P_X)$ and $L_2(P_X)$.

We now define the expectation of a generic random operator $A : \Omega_Y \rightarrow \mathcal{B}(L_2(P_X))$. For each $f \in L_2(P_X)$ and $x \in \Omega_X$, the mapping $y \mapsto (A(y)f)(x)$ defines a random variable. Its expectation defines a function $x \mapsto \int_{\Omega_Y} (A(y)f)(x) P_Y(dy)$, which is a member of $L_2(P_X)$. Denoting this member as \tilde{f} , we define the nonrandom operator $L_2(P_X) \rightarrow L_2(P_X)$, $f \mapsto \tilde{f}$ as the expectation $E(A)$. We now consider the operator

$$S = E(V - V_{X|Y})^2 : L_2(P_X) \rightarrow L_2(P_X), \quad (13)$$

where $V : L_2(P_X) \rightarrow L_2(P_X)$ is the (unconditional) covariance operator defined by

$$\langle f, Vg \rangle_{L_2(P_X)} = \text{cov}(f(X), g(X)).$$

This operator is similar to $\tilde{\Sigma}_{XX}$ in Section 5 except that it is not defined through RKHS. The operator S is an extension of the SAVE matrix (Cook and Weisberg, 1991)

$$\Sigma^{-1} E[\text{var}(X) - \text{var}(X|Y)]^2 \Sigma^{-1}. \quad (14)$$

Let β be a basis matrix of the central subspace $\mathcal{S}_{Y|X}$ of linear SDR. Cook and Weisberg show that if $E(X|\beta^T X)$ is linear in $\beta^T X$ and $\text{var}(X|\beta^T X)$ is nonrandom, then the column space of (14) is contained in $\mathcal{S}_{Y|X}$. The next theorem generalizes this result, but without requiring an analogue of the linear conditional mean assumption.

Theorem 7 *Suppose that conditions (A) and (B) are satisfied, and $\text{var}[f(X)|\mathcal{G}_{Y|X}]$ is non-random for any $f \in \mathfrak{S}_{Y|X}^\perp$. Then $\overline{\text{ran}} S \subseteq \mathfrak{S}_{Y|X}$.*

PROOF. Let $f \perp \mathfrak{S}_{Y|X}$. We claim that for any $y \in \Omega_Y$,

$$\langle f, [V - V_{X|Y}(y)]f \rangle_{L_2(P_X)} = 0. \quad (15)$$

Because $Y \perp\!\!\!\perp X | \mathcal{G}_{Y|X}$, we have

$$\text{var}(f(X)|Y) = \text{var}(E(f(X)|\mathcal{G}_{Y|X})|Y) + E(\text{var}(f(X)|\mathcal{G}_{Y|X})|Y).$$

Because, by Lemma 1, $E(f(X)|\mathcal{G}_{Y|X})$ is constant, the first term is 0. Because $\text{var}(f(X)|\mathcal{G}_{Y|X})$ is non-random, the second term is $\text{var}(f(X)|\mathcal{G}_{Y|X})$. Hence

$$\text{var}(f(X)|Y) = \text{var}(f(X)|\mathcal{G}_{Y|X}).$$

Similarly,

$$\text{var}(f(X)) = \text{var}(E(f(X)|\mathcal{G}_{Y|X})) + E(\text{var}(f(X)|\mathcal{G}_{Y|X})) = \text{var}(f(X)|\mathcal{G}_{Y|X}).$$

Therefore $\text{var}(f(X)|Y) = \text{var}(f(X))$, which implies (15). Since $V - V_{X|Y}(y)$ is self-adjoint, (15) implies $f \in \ker[V - V_{X|Y}(y)]$. Hence

$$\langle f, [V - V_{X|Y}(y)]^2 f \rangle_{L_2(P_X)} = 0.$$

Now integrate both sides of this equation to obtain

$$\begin{aligned} \int_{\Omega_Y} \langle f, (V - V_{X|Y}(y))^2 f \rangle_{L_2(P_X)} P_Y(dy) &= \langle f, \int_{\Omega_Y} (V - V_{X|Y}(y))^2 f P_Y(dy) \rangle_{L_2(P_X)} \\ &= \langle f, (E(V - V_{X|Y})^2) f \rangle_{L_2(P_X)} = 0. \end{aligned}$$

Hence $f \in \ker E(V - V_{X|Y})^2$, as desired. \square

Similar to the case of GSIR, we do not need to employ the rescaling by Σ^{-1} in (14) when generalizing SAVE, because the $L_2(P_X)$ -inner product absorbs any marginal variance. We call the estimator derived from $\overline{\text{ran}} S$ (see Section 7.3) *generalized SAVE*, or GSAVE. The next theorem shows that GSAVE can recover functions outside $\mathfrak{C}_{Y|X}$.

Theorem 8 *If conditions (A) and (B) are satisfied, then $\mathfrak{C}_{Y|X} \subseteq \overline{\text{ran}} S$.*

PROOF. Since S is self adjoint, it suffices to show that $\ker S \subseteq \mathfrak{C}_{Y|X}^\perp$. For any $f \in \ker S$,

$$\int_{\Omega_Y} \langle f, (V - V_{X|Y}(y))^2 f \rangle_{L_2(P_X)} P_Y(dy) = 0.$$

Hence $\langle f, (V - V_{X|Y}(y))^2 f \rangle_{L_2(P_X)} = 0$ a.s. P_Y , which implies $(V - V_{X|Y}(y))f = 0$ a.s. P_Y . Then

$$\int_{\Omega_Y} \langle f, (V - V_{X|Y}(y))f \rangle_{L_2(P_X)} P_Y(dy) = 0.$$

By Definition 8, the left hand side is $\text{var}[f(X)] - E[\text{var}(f(X)|Y)] = \text{var}[E(f(X)|Y)]$. Hence $\text{var}[E(f(X)|Y)] = 0$, which implies $E[f(X)|Y] = E[f(X)] = 0$. By Lemma 1, we have $f \in L_2(P_X) \ominus L_2(P_Y) = \mathfrak{C}_{Y|X}^\perp$, as desired. \square

Combining Theorems 7 and 8 we see that

$$\mathfrak{C}_{Y|X} \subseteq \overline{\text{ran}} S \subseteq \mathfrak{S}_{Y|X}, \tag{16}$$

which is analogous to the relation (3) in the classical setting. Thus we can expect GSAVE to discover functions outside the class $\mathfrak{C}_{Y|X}$, just as we can expect SAVE to discover vectors outside the space spanned by SIR.

7 Algorithms

We now develop algorithms for the sample versions of GSIR and GSAVE, together with a cross-validation scheme to select parameters in the GSIR and GSAVE algorithms. These sample versions involve representing the operators in Theorems 5 and 7 as matrices. To formulate the algorithms we need to introduce coordinate representations of functions and operators, which we adopt with modifications from Horn and Johnson (1985, P. 31); see also Li, Chun, and Zhao (2011).

Throughout this section, A^\dagger represents the Moore-Penrose inverse of a matrix A , $A^{\dagger\alpha}$ represents $(A^\dagger)^\alpha$, I_n denotes the $n \times n$ identity matrix, $\mathbf{1}_n$ denotes the vector in \mathbb{R}^n whose entries are all 1, and $Q = I_n - \mathbf{1}_n \mathbf{1}_n^\top / n$. Let $\kappa_X : \Omega_X \times \Omega_X \rightarrow \mathbb{R}$ be a positive definite function. Also, let K_X be the $n \times n$ the Gram matrix $\{\kappa_X(X_i, X_j) : i, j = 1, \dots, n\}$, G_X its centered versions QK_XQ , and L_X the Gram matrices with intercept; that is, $L_X = (\mathbf{1}_n, K_X)^\top$. Finally, define κ_Y, K_Y, G_Y, L_Y in the same manner for Y .

7.1 Coordinate representation

Let \mathcal{H} be a finite-dimensional Hilbert space with spanning system $\mathcal{B} = \{b_1, \dots, b_n\}$. For an $f \in \mathcal{H}$, let $[f]_{\mathcal{B}} \in \mathbb{R}^n$ denote the coordinates of f relative to \mathcal{B} ; that is, $f = \sum_{i=1}^n ([f]_{\mathcal{B}})_i b_i$. Let $b : \Omega_X \rightarrow \mathbb{R}^n$ denote the \mathbb{R}^n -valued function $(b_1, \dots, b_n)^\top$. Then we can write $f = [f]_{\mathcal{B}}^\top b$. Let $A : \mathcal{H} \rightarrow \mathcal{H}'$, where \mathcal{H}' is another finite-dimensional Hilbert spaces with spanning system $\mathcal{C} = \{c_1, \dots, c_m\}$, and let $c = (c_1, \dots, c_m)^\top$. Then, for $f \in \mathcal{H}$,

$$Af = A(b^\top [f]_{\mathcal{B}}) = (Ab_1, \dots, Ab_n)[f]_{\mathcal{B}} = (c^\top [Ab_1]_{\mathcal{C}}, \dots, c^\top [Ab_n]_{\mathcal{C}})[f]_{\mathcal{B}}.$$

Thus, if we let $c[A]_{\mathcal{B}} = ([Ab_1]_{\mathcal{C}}, \dots, [Ab_n]_{\mathcal{C}})$, then $Af = c^\top (c[A]_{\mathcal{B}})[f]_{\mathcal{B}}$. In other words,

$$[Af]_{\mathcal{C}} = (c[A]_{\mathcal{B}})[f]_{\mathcal{B}}.$$

Furthermore, if $A_1 : \mathcal{H}' \rightarrow \mathcal{H}''$ is another linear operator, where \mathcal{H}'' is a third finite-dimensional Hilbert space with spanning system \mathcal{D} , then, by a similar argument,

$$\mathcal{D}[A_1 A]_{\mathcal{B}} = (\mathcal{D}[A_1]_{\mathcal{C}})(c[A]_{\mathcal{B}}).$$

Since the spanning systems in the domain and range of an operator are self evident in the following discussion, we will write $c[A]_{\mathcal{B}}$ and $[f]_{\mathcal{B}}$ simply as $[A]$ and $[f]$.

Suppose $A \in \mathcal{B}(\mathcal{H})$ is self-adjoint. It can be shown that, for any $\alpha > 0$, $[A^\alpha] = [A]^\alpha$. Depending on the choice of the spanning system of \mathcal{H} , it is possible that A is invertible and yet $[A]$ is singular, but it is generally true that $A^{-\alpha} = [A]^{\dagger\alpha}$. Throughout this section the square brackets $[\cdot]$ will be used exclusively for denoting coordinate representations.

7.2 Algorithm for GSIR

At the sample level, P_X is replaced by the empirical measure $P_{n,X}$; \mathcal{H}_X is the RKHS spanned by $\mathcal{B}_X = \{\kappa_X(\cdot, X_1), \dots, \kappa_X(\cdot, X_n)\}$ with inner product $\langle f, g \rangle_{\mathcal{H}_X} = [f]^\top K_X [g]$, where $[\cdot]$ is coordinate with respect to \mathcal{B}_X . The space $L_2(P_{n,X})$ is spanned by $\kappa_X(\cdot, X_i) - E_n \kappa_X(X, X_i)$, $i = 1, \dots, n$ with inner product $\langle f, g \rangle_{L_2(P_{n,X})} = \text{cov}_n[f(X), g(X)] = n^{-1}[f]^\top K_X Q K_X [g]$. The operator M_{XX} is defined through the relation $\langle f, M_{XX} g \rangle_{\mathcal{H}_X} = \text{cov}_n(f(X), g(X))$; that is,

$$[f]^\top K_X [M_{XX}][g] = n^{-1}[f]^\top K_X Q K_X [g].$$

Since $[f]$ and $[g]$ are arbitrary members of \mathbb{R}^n , the above implies $[M_{XX}] = n^{-1}QK_X$. Then any $f \in \text{ran } M_{XX} \equiv \mathcal{G}_X$ can be written as $M_{XX} g$ for some $g \in \mathcal{H}_X$, which implies $[f] = QK_X [g] = Q[f]$. Consequently, for any $f, g \in \mathcal{G}_X$, $\langle f, g \rangle_{\mathcal{H}_X} = [f]^\top G_X [g]$.

Let us now find the matrix representations of Σ_{XX} , Σ_{YY} , and Σ_{YX} . In the following, h_X represents the function $x \mapsto (\kappa_X(x, X_1), \dots, \kappa_X(x, X_n))^\top$. For any $f \in \mathcal{G}_X$, we have

$$\Sigma_{XX} f = M_{XX} f = h_X^\top [M_{XX}][f] = n^{-1}h_X^\top Q K_X Q [f] = n^{-1}h_X^\top G_X [f].$$

Hence $[\Sigma_{XX} f] = [\Sigma_{XX}][f] = n^{-1}G_X [f]$. Since this is true for all $[f] \in \text{span}(Q)$, we have $[\Sigma_{XX}] = n^{-1}G_X$. By the same argument we can show that

$$\begin{aligned} [\tilde{\Sigma}_{XX}] &= n^{-1}G_X, \quad [\Sigma_{YY}] = [\tilde{\Sigma}_{YY}] = n^{-1}G_Y, \quad [\Sigma_{YX}] = [\tilde{\Sigma}_{YX}] = n^{-1}G_X, \\ [E_{X|Y}] &= G_Y^\dagger G_X G_X^{\dagger 1/2} G_X^{1/2}. \end{aligned} \quad (17)$$

Theorem 5 suggests that we use $\overline{\text{ran}}(E_{X|Y}^* E_{X|Y})$ to estimate $\mathfrak{S}_{Y|X}$. Since $E_{X|Y}^* E_{X|Y}$ is an operator on $L_2(P_{n,X})$ to $L_2(P_{n,X})$, the vectors in $\overline{\text{ran}}(E_{X|Y}^* E_{X|Y})$ can be found by

$$\text{maximizing } \langle f, E_{X|Y}^* E_{X|Y} f \rangle_{L_2(P_{n,X})} = \|E_{X|Y} f\|_{L_2(P_{n,Y})}^2 \quad \text{subject to } \langle f, f \rangle_{L_2(P_{n,X})} = 1.$$

The coordinate representation of this problem is

$$\text{maximizing } [f]^\top [E_{X|Y}]^\top G_Y^2 [E_{X|Y}][f] \quad \text{subject to } [f]^\top G_X^2 [f] = 1.$$

The optimal solution is $[f] = G_X^\dagger \phi$, where ϕ are the leading eigenvectors of the matrix

$$G_X^\dagger [E_{X|Y}]^\top G_Y^2 [E_{X|Y}] G_X^\dagger = G_X^\dagger G_X^{1/2} G_X^{\dagger 1/2} G_X G_Y^\dagger G_Y^2 G_Y^\dagger G_X G_X^{\dagger 1/2} G_X^{1/2} G_X^\dagger. \quad (18)$$

To enhance accuracy we replace the Moore-Penrose inverses G_X^\dagger and G_Y^\dagger by the ridge-regression-type regularized inverses $(G_X + \epsilon_X I_n)^{-1}$ and $(G_Y + \epsilon_Y I_n)^{-1}$. We summarize the algorithm as follows.

1. Select the parameters $\gamma_X, \gamma_Y, \epsilon_X, \epsilon_Y$ using the algorithm in Section 7.4.

2. Compute the matrix

$$(G_X + \epsilon_X I_n)^{-3/2} G_X^{3/2} (G_Y + \epsilon_Y I_n)^{-1} G_Y^2 (G_Y + \epsilon_Y I_n)^{-1} G_X^{3/2} (G_X + \epsilon_X I_n)^{-3/2}$$

and its first d eigenvectors ϕ_1, \dots, ϕ_d of this matrix.

3. Form the sufficient predictors at x $\phi_i^\top (G_X + \epsilon_X I_n)^{-1} h_X(x)$, $i = 1, \dots, d$.

GSIR estimation is similar to the kernel canonical correlation analysis (KCCA) developed by Akaho (2001), Bach and Jordan (2002), and Fukumizu, Bach, and Gretton (2007). In our notation, KCCA maximizes

$$\langle g, \Sigma_{YX} f \rangle_{L_2(P_Y)} = [g]^\top G_Y G_X [f]$$

subject to $\langle g, \Sigma_{YY} g \rangle_{L_2(P_Y)} = [g]^\top G_Y^2 [g] = 1$ and $\langle f, \Sigma_{XX} f \rangle_{L_2(P_X)} = [f]^\top G_X^2 [f] = 1$. The optimal solution for $[f]$ is $[f] = (G_X + \epsilon_X I_n)^{-1} \phi$, where ϕ is one of the first d eigenvectors of

$$(G_X + \epsilon_X I_n)^{-1} G_X G_Y (G_Y + \epsilon_Y I_n)^{-2} G_Y G_X (G_X + \epsilon_X I_n)^{-1}.$$

We will compare GSIR and KCCA in Section 8.

7.3 Algorithm for GSAVE

We first derive the sample-level representation of the operator $V_{X|Y}(y)$. The sample version of the noncentered L_2 -classes $L'_2(P_{n,X})$ and $L'_2(P_{n,Y})$ are spanned by

$$\mathcal{C}_X = \{1, \kappa_X(\cdot, X_1), \dots, \kappa_X(\cdot, X_n)\}, \quad \mathcal{C}_Y = \{1, \kappa_Y(\cdot, Y_1), \dots, \kappa_Y(\cdot, Y_n)\}, \quad (19)$$

respectively. Let $[\cdot]$ represent the coordinates relative to these spanning systems. Then, for any $f \in L'_2(P_{n,X})$, $(f(X_1), \dots, f(X_n))^\top = L_X^\top [f]$. The operator $E'_{X|Y}$ is defined through the relation $\langle g, E'_{X|Y} f \rangle_{L'_2(P_{n,Y})} = E_n(g(Y)f(X))$, which yields the representation

$$[E'_{X|Y}] = (L_Y L_Y^\top)^\dagger (L_Y L_X^\top). \quad (20)$$

Let ℓ_Y denote the function $y \mapsto (1, \kappa_Y(y, Y_1), \dots, \kappa_Y(y, Y_n))^\top$, and let ℓ_X denote the same function of x . For any $f, g \in L'_2(P_{n,X})$,

$$\begin{aligned} & \{E'_{X|Y}(fg) - (E'_{X|Y} f)(E'_{X|Y} g)\}(y) \\ &= \ell_Y^\top(y) [E'_{X|Y}] [fg] - [f]^\top [E'_{X|Y}]^\top \ell_Y(y) \ell_Y^\top(y) [E'_{X|Y}] [g]. \end{aligned} \quad (21)$$

For any X_i , $f(X_i)g(X_i)$ can be expressed as the i th entry of the vector $L_X^\top [f] \odot L_X^\top [g]$, which is the same as $L_X^\top (L_X L_X^\top)^\dagger L_X (L_X^\top [f] \odot L_X^\top [g])$, where \odot is the Hadamard product. Thus we have the coordinate representation

$$[fg] = (L_X L_X^\top)^\dagger L_X (L_X^\top [f] \odot L_X^\top [g]). \quad (22)$$

Substituting (20) and (22) into (21) we see that, for any $f, g \in L'_2(P_{n,x})$,

$$\langle f, V_{X|Y}(y)g \rangle_{L'_2(P_{n,x})} = [f]^\top L_X (\text{diag} C_Y(y) - C_Y(y) C_Y^\top(y)) L_X^\top [g] \equiv [f]^\top L_X \Lambda(y) L_X^\top [g], \quad (23)$$

where $C_Y(y) = L_Y^\top (L_Y L_Y^\top)^\dagger \ell_Y(y)$.

Let $S_n : L_2(P_{n,x}) \rightarrow L_2(P_{n,x})$ be the operator $E_n(V - V_{X|Y}(Y))^2$. By Theorem 7, GSAVE is the class of functions $\overline{\text{ran}}(S)$. At the sample level, this corresponds to

$$\text{maximizing } \langle f, S_n f \rangle_{L_2(P_{n,x})} \quad \text{subject to } \langle f, f \rangle_{L_2(P_{n,x})} = 1. \quad (24)$$

By (23), for each $y \in \Omega_Y$, and $f, g \in L_2(P_{n,x})$, we have

$$\langle g, V_{X|Y}(y)f \rangle_{L_2(P_{n,x})} = [f]^\top L_X Q \Lambda(y) Q L_X^\top [g].$$

From this we deduce that $[V_{X|Y}(y)] = (L_X Q L_X^\top / n)^\dagger L_X Q \Lambda(y) Q L_X^\top$. By a similar derivation we find $[V] = (L_X Q L_X^\top / n)^\dagger (L_X Q L_X^\top / n)$. Hence

$$[V - V_{X|Y}(y)] = (L_X Q L_X^\top / n)^\dagger L_X Q (Q/n - \Lambda(y)) Q L_X^\top.$$

It follows that

$$\langle f, S_n f \rangle_{L_2(P_{n,x})} = E_n \{ [f]^\top L_X Q (Q/n - \Lambda(Y)) Q (Q/n - \Lambda(Y)) Q L_X^\top [f] \}.$$

To find $\text{ran}(S_n)$ we maximize the above subject to $[f]^\top (L_X Q L_X^\top / n) [f] = 1$. Again we use the regularized inverses instead of the Moore Penrose inverses to enhance performance. The algorithm is summarized as follows.

1. Determine $\gamma_X, \gamma_Y, \epsilon_X, \epsilon_Y$ using the algorithm in section 7.4.
2. Compute $C = L_Y^\top (L_Y L_Y^\top + \epsilon_Y I_{n+1})^{-1/2} L_Y$. Let C_i be the i th column of C . Compute $\Lambda_i = \text{diag}(C_i) - C_i C_i^\top$ and then compute $\Gamma_i = Q/n - \Lambda_i$ for $i = 1, \dots, n$.
3. Compute

$$n^{-1} \sum_{i=1}^n (L_X Q L_X^\top + \epsilon_X I_{n+1})^{-1/2} L_X Q \Gamma_i Q \Gamma_i Q L_X^\top (L_X Q L_X^\top + \epsilon_X I_{n+1})^{-1/2}$$

and the first d eigenvectors of this matrix, say ϕ_1, \dots, ϕ_d .

4. The sufficient predictors' values at any given $x \in \Omega_X$ are the set of d numbers

$$\ell_X^\top(x) (L_X Q L_X^\top + \epsilon_X I_{n+1})^{-1/2} Q \phi_i, \quad i = 1, \dots, d.$$

Here we should mention that, similar to SAVE for linear SDR, GSAVE works best for extracting predictors affecting the conditional variance of the response, but often not so well for extracting predictors affecting the conditional mean. However, we expect that other second-order methods for linear SDR, such as directional regression (Li and Wang, 2007) and the minimum discrepancy approach (Cook and Ni, 2005), will be amenable to similar generalizations to nonlinear SDR. These will be left for future research.

7.4 Cross-validation algorithm

We now develop a cross-validation scheme to determine the parameters γ_X , γ_Y , ϵ_X , ϵ_Y , which are used in the algorithms for both the GSIR and the GSAVE. We will only describe the algorithm for determining (γ_X, ϵ_X) ; that of (γ_Y, ϵ_Y) is completely analogous.

In the following, for a matrix A , $A_{-i,-j}$ represents the submatrix of A with its i th row and j th column removed, and $A_{-i,j}$ represents the j th column of A with the i th entry removed. Let $\mathcal{C}_Y^{-i} = \mathcal{C}_Y \setminus \{\kappa_Y(\cdot, Y_i)\}$, and define \mathcal{C}_X^{-i} similarly. Our cross validation strategy is to predict $f(Y_i)$ for each $f \in \mathcal{C}_Y^{-i}$, using the conditional mean operator developed from $(\mathcal{C}_X^{-i}, \mathcal{C}_Y^{-i})$. The regularized matrix representation of $E'_{Y|X}$ based on $(\mathcal{C}_X^{-i}, \mathcal{C}_Y^{-i})$ is

$$[E'_{Y|X}]_{-(i+1),-(i+1)} = \{(L_X)_{-(i+1),-i}(L_X)_{-(i+1),-i}^\top + \epsilon_X I_n\}^{-1}(L_X)_{-(i+1),-i}(L_Y)_{-(i+1),-i}^\top$$

The k th member f_k of \mathcal{C}_Y^{-i} is the function $e_k^\top(\ell_Y)_{-(i+1)}(\cdot)$ where e_k is the vector in \mathbb{R}^n whose k th entry is 1 and the remaining entries are 0. Therefore, the estimate of $E(f_k(Y)|X = x)$ based on \mathcal{C}_X^{-i} is

$$(\ell_X)_{-(i+1)}^\top(x)[E'_{Y|X}f_k]_{-(i+1)} = e_k^\top[E'_{Y|X}]_{-(i+1),-(i+1)}^\top(\ell_X)_{-(i+1)}(x),$$

and the prediction of $(f_1(Y_i), \dots, f_n(Y_i))^\top$ is $[E'_{Y|X}]_{-(i+1),-(i+1)}^\top(\ell_X)_{-(i+1)}(X_i)$. However, because $(\ell_X)_{-(i+1)}(X_i)$ is the vector $(L_X)_{-(i+1),i}$, and $(f_1(Y_i), \dots, f_n(Y_i))^\top$ is the vector $(L_Y)_{-(i+1),i}$, the difference between $(f_1(Y_i), \dots, f_n(Y_i))^\top$ and its prediction is

$$(L_Y)_{-(i+1),i} - [E'_{Y|X}]_{-(i+1),-(i+1)}^\top(L_X)_{-(i+1),i}.$$

To stress that this difference depends on $\gamma_X, \epsilon_X, \gamma_Y$, we denote it by $\Delta_i(\epsilon_X, \gamma_X, \gamma_Y)$. Our cross-validation criterion is defined as $\text{CV}(\gamma_X, \epsilon_X, \gamma_Y) = \sum_{i=1}^n \|\Delta_i(\gamma_X, \epsilon_X, \gamma_Y)\|^2$. Since the role of γ_Y is only to determine the set of functions to be predicted, we exclude it from the optimization process (for the determination of ϵ_X, γ_X). Moreover, as argued in Fukumizu, Bach, and Jordan (2009), the parameters ϵ_X and γ_X have similar smoothing effects and only one of them needs to be optimized. For these reasons we fix γ_Y and ϵ_X at

$$1/\sqrt{\gamma_{Y0}} = \binom{n}{2}^{-1} \sum_{i < j} |Y_i - Y_j|, \quad \epsilon_{X0} = 0.01, \quad (25)$$

and minimize $\text{CV}(\gamma_X, \epsilon_{X0}, \gamma_{Y0})$ over a grid for γ_X . The grid consists of 20 subintervals in $[\gamma_{X0}/3, 3\gamma_{X0}]$, equally spaced in the log scale, where γ_{X0} is calculated using the first formula in(25) with $|Y_i - Y_j|$ replaced by $\|X_i - X_j\|$. The rationale for this formula can be found in Li, Artemiou, and Li (2011).

The pair (γ_Y, ϵ_Y) is selected in the same way, except that ϵ_{Y0} is set to 0.001. This is because Y has dimension 1, so a weaker penalty is needed.

8 Simulations and data analysis

In this section we present simulation comparisons among GSIR, GSAVE, KSIR, and KCCA. For the reasons explained in the previous section, we compare GSIR with KSIR and KCCA in settings where the sufficient predictor appears in the conditional mean, and we compare GSAVE with GSIR, KSIR, and KCCA in settings where the sufficient predictor appears in the conditional variance. We also apply GSIR, KSIR, and KCCA to two real data sets.

8.1 Simulation comparisons

To make a comprehensive comparison of GSIR, KSIR and KCCA we consider three regression models; namely:

$$\begin{cases} \text{I:} & Y = (X_1^2 + X_2^2)^{1/2} \log(X_1^2 + X_2^2)^{1/2} + \varepsilon \\ \text{II:} & Y = X_1/(1 + e^{X_2}) + \varepsilon \\ \text{III:} & Y = \sin(\pi(X_1 + X_2)/10) + \varepsilon \end{cases}, \quad \varepsilon \perp\!\!\!\perp X, \varepsilon \sim N(0, 0.25), p = 10,$$

as well as three distributional scenarios for the predictor vector X ; namely: (A) independent gaussian predictors, (B) independent non-gaussian predictors, and (C) correlated gaussian predictors. In symbols:

$$\begin{cases} \text{A:} & X \sim N(0, I_p) \\ \text{B:} & X \sim (1/2)N(-1_p, I_p) + (1/2)N(1_p, I_p) \\ \text{C:} & X \sim N(0, 0.6I_p + 0.41_p1_p^T) \end{cases}$$

Note that the central σ -fields for the three models I, II, and III are generated by $X_1^2 + X_2^2$, $X_1/(1 + e^{X_2})$, and $\sin(\pi(X_1 + X_2)/10)$, respectively.

We assess the quality of an estimated sufficient predictor by its closeness to the true sufficient predictor and its closeness to the response. Since we are only interested in monotone functions of the predictor, we use Spearman's correlation as the measure of closeness. For each combination of the models and scenarios, we generate $n = 200$ observations on (X, Y) as the training data, and compute the first predicting function using the each of three methods. We then independently generate $m = 200$ observations on (X, Y) as the testing data, and evaluate the predicting functions at these points. Finally, we compute the mentioned Spearman's correlations from the testing data. This process is repeated $N = 200$ times. In Table 1 we list means and standard deviations of the Spearman's correlations computed using the $N = 200$ simulated samples. From the table we see that the performances of KCCA and GSIR are similar, and both are slightly better than KSIR.

Table 1: Comparison of KSIR, KCCA, and GSIR when sufficient predictors appear in the conditional means

Models		Spearman cor. with true predictor			Spearman cor. with response		
X	$Y X$	KSIR	KCCA	GSIR	KSIR	KCCA	GSIR
A	I	0.78 (0.05)	0.81 (0.04)	0.80 (0.05)	0.63 (0.06)	0.66 (0.05)	0.64 (0.05)
	II	0.81 (0.05)	0.90 (0.03)	0.91 (0.03)	0.56 (0.06)	0.61 (0.05)	0.62 (0.05)
	III	0.76 (0.06)	0.89 (0.04)	0.91 (0.03)	0.47 (0.07)	0.56 (0.05)	0.56 (0.05)
B	I	0.88 (0.02)	0.88 (0.02)	0.87 (0.02)	0.82 (0.03)	0.81 (0.03)	0.80 (0.03)
	II	0.89 (0.03)	0.93 (0.02)	0.93 (0.02)	0.71 (0.04)	0.74 (0.04)	0.74 (0.04)
	III	0.90 (0.02)	0.97 (0.01)	0.97 (0.01)	0.72 (0.04)	0.77 (0.03)	0.77 (0.03)
C	I	0.79 (0.04)	0.82 (0.04)	0.81 (0.04)	0.64 (0.05)	0.66 (0.05)	0.65 (0.05)
	II	0.83 (0.05)	0.86 (0.06)	0.88 (0.04)	0.56 (0.06)	0.59 (0.06)	0.60 (0.06)
	III	0.83 (0.06)	0.96 (0.02)	0.96 (0.02)	0.56 (0.06)	0.65 (0.04)	0.65 (0.04)

Next, we compare GSAVE, KSIR, KCCA, and GSIR when the predictors only affect the variance. We use the following models:

$$\begin{cases} \text{IV:} & Y = X_1 \varepsilon \\ \text{V:} & Y = (1/50)(X_1^3 + X_2^3) \varepsilon \\ \text{VI:} & Y = (X_1/(1 + e^{X_2})) \varepsilon \end{cases}$$

and again the scenarios (A), (B), and (C) for the distribution of X . The specifications of n, m, N, p are the same as in the previous comparison.

Table 2: Comparison of KSIR, KCCA, GSIR, and GSAVE when sufficient predictors appear in conditional variances

Models		Spearman's correlation with true predictors			
X	$Y X$	GSAVE	KSIR	KCCA	GSIR
A	IV	0.89 (0.08)	0.10 (0.07)	0.36 (0.22)	0.41 (0.23)
	V	0.73 (0.19)	0.09 (0.07)	0.17 (0.13)	0.20 (0.14)
	VI	0.84 (0.09)	0.10 (0.08)	0.25 (0.17)	0.27 (0.17)
B	IV	0.87 (0.08)	0.10 (0.07)	0.43 (0.25)	0.53 (0.25)
	V	0.88 (0.06)	0.09 (0.07)	0.11 (0.08)	0.11 (0.08)
	VI	0.76 (0.15)	0.27 (0.11)	0.61 (0.13)	0.64 (0.13)
C	IV	0.76 (0.20)	0.11 (0.07)	0.23 (0.16)	0.26 (0.18)
	V	0.82 (0.14)	0.10 (0.07)	0.11 (0.09)	0.12 (0.09)
	VI	0.73 (0.15)	0.15 (0.10)	0.41 (0.17)	0.44 (0.17)

Because the sufficient predictors appear in the conditional variance $\text{var}(Y|X)$ only, it is less meaningful to measure the closeness between the estimated sufficient predictor and the response. So in Table 2 we only report the means and standard deviations of Spearman’s correlations between the estimated and true sufficient predictors. We see that GSAVE performs substantially better than the other methods. The discrepancy can be explained by the fact that KSIR, KCCA, and GSIR depend completely on $E[\text{var}(f(X)|Y)]$, whereas GSAVE extracts more information from $\text{var}(f(X)|Y)$.

8.2 Data analysis

We first consider the *faces data*, available at <http://waldron.stanford.edu/isomap/datasets.html>. This data set contains 698 images of the same sculpture of a face photographed at different angles and with different lighting directions. The predictor comprises 64×64 image pixels (thus $p = 4096$) and the response comprises horizontal rotation, vertical rotation, and lighting direction measurements (thus $q = 3$). We use this data to demonstrate that the first three sufficient predictors estimated by KCCA and GSIR can effectively capture the 3-variate response. We use $n = 558$ of the images selected at random (roughly 80%) as training data, and the remaining $m = 140$ images as testing data. For each method, we estimate the first three predictor functions from the training data, and evaluate them on the testing data. The left panel of Figure 1 is the perspective plot of the first three KCCA predictors evaluated on the 140 testing images, and the right panel is the counterpart for GSIR. We did not include KSIR in this comparison because in its proposed form it cannot handle multivariate responses. The perspective plots indicate that nearby regions in the 3-D cubes have similar patterns of left-right rotation, up-down rotation and lighting direction, while distant regions have discernibly different patterns. This reflects the ability of the three sufficient predictors to capture the 3-variate responses.

Next, we apply KSIR, KCCA, and GSIR to the *handwritten digits data*, available at <http://www.cs.nyu.edu/~roweis/data.html>. This data set contains 2000 images of $p = 16 \times 16$ pixels showing handwritten digits from 0 to 9 – the response is thus categorical with 10 levels. We use 1000 images as training data and 1000 as testing data. Again, for each method we estimate the first three sufficient predictors on the training data, and evaluate them on the testing data. Results are presented in the three perspective plots in Figure 2 – for visual clarity, these plots include only 100 randomly selected points from the 1000 in the testing data. The plots show that all three methods provide low-dimensional representations in which the digits are well separated.




Figure 1: First 3 sufficient predictors by KCCA (left panel) and GSIR (right panel), computed from 558 training images, and evaluated on 140 testing images – faces data.

9 Concluding Remarks

In this article we described a novel and very general theory of sufficient dimension reduction. This theory allowed us to combine linear and nonlinear SDR into a coherent system, to link them with classical statistical sufficiency, and to subsume several existing nonlinear SDR methods into a unique framework.

Our developments thus revealed important and previously unexplored properties of SDR methods. For example, unbiasedness of various nonlinear extensions of SIR proposed in recent literature was proved under the stringent linear conditional mean assumption. We were able to show that these methods are all unbiased under virtually no assumption, and that GSIR is exhaustive under the completeness assumption. We were also able to show that nonlinear extensions of SIR are in general *not* exhaustive when completeness is not satisfied, and that in these cases GSAVE can recover a larger portion of the central class. These insights could not have been obtained without paralleling linear and nonlinear SDR as allowed by our new theory.

In addition to achieving theoretical synthesis and important insights on SDR methods, we introduced a new *heteroscedastic* conditional variance operator – which is more general than the (homoscedastic) conditional variance operator in Fukumizu, Bach, and Jordan (2004, 2009). This operator was crucial to generalizing SAVE to the nonlinear GSAVE, and thus to exploit dependence information in the conditional variance to improve upon the performance of the nonlinear extensions of SIR. We have no doubt that the heteroscedastic conditional variance operator can be used to generate nonlinear extensions of other second order SRD methods such as contour regression (Li, Zha, and Chiaromonte, 2005), directional




Figure 2: First 3 sufficient predictors by KSIR (upper-left panel), KCCA (upper-right panel) and GSIR (lower panel), computed on 1000 training images, and evaluated on 1000 testing images – handwritten digits data.

regression (Li and Wang, 2007), SIR-II (Li, 1991), and other F2M methods (Cook and Forzani, 2009). These extensions will be the topic of future work.

More generally, it is our hope that the clarity and simplicity that classical notions lend to the formulation of dimension reduction, as well as the transparent parallels we were able to draw between linear and nonlinear SDR, will provide fertile grounds for much research to come.

As we put forward a general theory that encompasses both linear and nonlinear SDR, it is also important to point out that linear SDR has its special values that cannot be replaced by nonlinear SDR via kernel mapping, one of which is its preservation of the original coordinates and as a result its strong interpretability. For example, when mapped

to higher dimension spaces, kernel methods can sometimes interpret difference in variances in the original coordinates as location separation in the transformed coordinates, which can be undesirable depending on the goal and emphasis of particular applications. For further discussion and an example of this point, see Li, Artemiou, and Li (2011).

Acknowledgments

We would like to thank two referees and an Associate Editor for their insightful comments and useful suggestions, which led to significant improvements to this paper. In particular, the consideration of nonlinear sufficient dimension reduction for functional data is inspired by the comments of two referees. Our research is supported in part by National Science Foundation grants DMS-0806058 and DMS-1106815 awarded to Bing Li.

References

1. Akaho, S. (2001). A kernel method for canonical correlation analysis. In *Proceedings of the International Meeting of the Psychometric Society (IMPS2001)*. Springer, Tokyo.
2. Bach, F. R. and Jordan. M. I. (2002). Kernel independent component analysis. *Journal of Machine Learning Research*, **3**, 1-48.
3. Bahadur, R. R. (1954). Sufficiency and Statistical Decision Functions. *The Annals of Mathematical Statistics*, **25**, 3, 423-462.
4. Baker, C. R. (1973). Joint measures and cross-covariance operators. *Transactions of the American Mathematical Society*, **186**, 273–289.
5. Cook, R. D. (1994). Using dimension-reduction subspaces to identify important inputs in models of physical systems. In *1994 Proceedings of the Section on Physical and Engineering Sciences*, Alexandria, VA: American Statistical Association, 18–25.
6. Cook, R. D. (1998). Principle Hessian Directions Revisited. *Journal of American Statistical Association*, **93**, 441, 84-94.
7. Cook, R. D. (1998). *Regression Graphics: Ideas for Studying Regressions through Graphics*. New York: *Wiley*.
8. Cook, R. D. (2007). Fisher Lecture: Dimension Reduction in Regression. *Statistical Science*, **22**, 1–40.
9. Cook, R. D. and Forzani, L. (2009). Likelihood-Based Sufficient Dimension Reduction. *Journal of the American Statistical Association*, **104**, 485, 197-208.

10. Cook, R. D. and Li, B. (2002). Dimension reduction for the conditional mean in regression. *Annals of Statistics*, 30, 455–474.
11. Cook, R.D., Li, B. and Chiaromonte, F. (2010). Envelope models for parsimonious and efficient multivariate linear regression (with discussion). *Statistica Sinica*, **20**, 927–1010.
12. Cook, R.D. and Ni, L. (2005). Sufficient dimension reduction via inverse regression: A minimum discrepancy approach. *Journal of the American Statistical Association*, **100**, 410–428.
13. Cook, R. D. and Weisberg, S. (1991). Discussion of “Sliced inverse regression for dimension reduction”. *Journal of the American Statistical Association*. **86**, 316–342.
14. Duan, N. and Li, K.-C. (1991). A Bias Bound for Applying Linear Regression to a General Linear Model. *Statistica Sinica*. **1**, 127–136.
15. Fukumizu, K., Bach, F. R. and Jordan M. I. (2004). Dimensionality reduction for supervised learning with reproducing kernel Hilbert spaces. *Journal of Machine Learning Research*, **5**, 73–99.
16. Fukumizu, K., Bach, F. R. and Gretton, A. (2007). Statistical Consistency of Kernel Canonical Correlation Analysis. *Journal of Machine Learning Research*, **8**, 361–383.
17. Fukumizu, K., Bach, F. R. and Jordan M. I. (2009). Kernel dimension reduction in regression. *Annals of Statistics*, **4** 1871–1905.
18. Härdle, W., Hall, P. and Ichimura, H. (1993). Optimal Smoothing in Single-Index Models. *Annals of Statistics*, **21(1)**,157-178.
19. Hsing, T. and Ren, H. (2009). An RKHS formulation of the inverse regression dimension-reduction problem. *The Annals of Statistics*, **37** 726–755.
20. Horn, R. A. and Johnson, C. R. (1985). *Matrix Analysis*. Cambridge University Press.
21. Ichimura, H. and L. Lee, (1991). Semiparametric least squares estimation of multiple index models: single equation estimation. W.A. Barnett, J.L. Powell, and G. Tauchen, eds., *Nonparametric and semiparametric methods in econometrics and statistics*, Cambridge: Cambridge University Press.
22. Lehmann, E. L. (1981). An interpretation of completeness and Basus theorem. *Journal of the American Statistical Association*, **76**, 335-340.

23. Li, B., Artemiou, A. and Li, L. (2011). Principal support vector machines for linear and nonlinear sufficient dimension reduction. *Annals of Statistics*. To appear.
24. Li, B., Chun, H. and Zhao, H. (2011). Sparse estimation of conditional graphical models with application to gene networks. *Journal of the American Statistical Association*. To appear.
25. Li, B. and Wang, S. (2007). On directional regression for dimension reduction. *Journal of the American Statistical Association*, **102**, 997–1008.
26. Li, B., Zha, H., and Chiaromonte, F. (2005). Contour regression: a general approach to dimension reduction. *Annals of Statistics*, **33**, 1580-1616.
27. Li, K. -C. (1991). Sliced inverse regression for dimension reduction (with discussion). *Journal of the American Statistical Association*, **86**, 316–342.
28. Li, K. -C. (1992). On principal Hessian directions for data visualization and dimension reduction: another application of Stein’s lemma. *Journal of the American Statistical Association*, **86**, 316–342.
29. Li, K. -C. and Duan, N. (1989). Regression analysis under link violation. *Annals of Statistics*, **17**, 1009–1052.
30. Wu, H. M. (2008). Kernel sliced inverse regression with applications on classification. *Journal of Computational and Graphical Statistics*, **17**, 590–610.
31. Wu, Q., Liang, F., and Mukherjee, S. (2008). Regularized sliced inverse regression for kernel models. Technical report, Duke University.
32. 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.
33. Yin, X., Li, B. and Cook, R. D. (2008) Successive direction extraction for estimating the central subspace in a multiple-index regression. *Journal of Multivariate Analysis*, **99**, 1733-175.
34. Zhu, H. and Li, L. (2011). Biological pathway selection through nonlinear dimension reduction. *Biostatistics*, **12**, 429–444.