

Sufficient dimension reduction for longitudinally measured predictors

Ruth M. Pfeiffer,^{a,*†} Liliana Forzani^{b,c} and Efstathia Bura^d

We propose a method to combine several predictors (markers) that are measured repeatedly over time into a composite marker score without assuming a model and only requiring a mild condition on the predictor distribution. Assuming that the first and second moments of the predictors can be decomposed into a time and a marker component via a Kronecker product structure that accommodates the longitudinal nature of the predictors, we develop first-moment sufficient dimension reduction techniques to replace the original markers with linear transformations that contain sufficient information for the regression of the predictors on the outcome. These linear combinations can then be combined into a score that has better predictive performance than a score built under a general model that ignores the longitudinal structure of the data. Our methods can be applied to either continuous or categorical outcome measures. In simulations, we focus on binary outcomes and show that our method outperforms existing alternatives by using the AUC, the area under the receiver–operator characteristics (ROC) curve, as a summary measure of the discriminatory ability of a single continuous diagnostic marker for binary disease outcomes. Published 2011. This article is a US Government work and is in the public domain in the USA.

Keywords: discrimination; AUC; Kronecker product; sliced inverse regression (SIR)

1. Introduction

Much research effort is devoted to searching for predictors, also called markers, which may aid in diagnosis of disease or physical impairment. Ideally, one would obtain a single marker with very high specificity and sensitivity. However, such high-performance markers are yet to be found for many diseases. Strategies for combining information from multiple diagnostic predictors are needed, because a combination may provide a better tool for diagnosis or screening applications than any single marker on its own. In an earlier work [1], we proposed an approach to combine several biomarkers into a composite marker score without assuming a model for the distribution of the predictors. Using sufficient dimension reduction (SDR) techniques, we replaced the original markers with a lower-dimensional version, obtained through linear transformations of markers that contain sufficient information for regression of the predictors on the outcome. We focused on a first-moment method, sliced inverse regression (SIR) [2], and also on a second-moment method, sliced average variance estimation [3].

In this paper, we extend first-moment-based SDR (FMSDR) methods to combining several longitudinally measured biomarkers into a composite marker score under a mild distributional assumption and by exploiting their longitudinal structure. We assume that the means and the second moments of the markers can be separated into a marker-specific and a time-specific component via a Kronecker product structure. This substantially reduces the complexity of the first-moment-based dimension reduction subspace and results in better predictive accuracy of the resulting score compared to standard FMSDR methods.

Following the presentation of background material (Section 2), we present the new results (Section 3) for FMSDR methods for longitudinally measured predictors and an algorithm that extends SIR

^aBiostatistics Branch, National Cancer Institute, Bethesda, MD 20892-7244, U.S.A.

^bInstituto de Matemática Aplicada del Litoral, CONICET, Güemes 3450, (3000) Santa Fe, Argentina

^cFacultad de Ingeniería Química, UNL, Santiago del Estero 2654, (3000) Santa Fe, Argentina

^dDepartment of Statistics, George Washington University, Washington, DC 20052, U.S.A.

*Correspondence to: Ruth M. Pfeiffer, Biostatistics Branch, National Cancer Institute, Bethesda, MD 20892-7244, U.S.A.

†E-mail: pfeiffer@mail.nih.gov

(Section 4) to that setting. We carry out simulations to assess and compare the performance of the longitudinal extension of SIR (Section 5) and give a data example in Section 6 before concluding with a discussion in Section 7.

2. Background

We start with a brief overview of standard SDR methods and a review of results for the estimation of the first-moment-based dimension reduction subspace, before considering the setting of longitudinally measured markers.

2.1. Sufficient dimension reduction

Suppose we are interested in inferring the relationship between a response variable Y and a covariate vector $\mathbf{X} = (X_1, \dots, X_p)^T \in \mathbb{R}^p$. When the number of covariates p is large, it is very difficult to visualize how Y changes as a function of the components of \mathbf{X} , which makes modeling challenging. Dimension reduction aims to reduce the complexity of the regression or classification problem prior to model fitting. In particular, SDR [4] aims to find a function of \mathbf{X} , $\mathbf{Q}: \mathbb{R}^p \rightarrow \mathbb{R}^d$ with $d \leq p$, such that $\mathbf{Q}(\mathbf{X})$ contains the same information as \mathbf{X} about the response Y . This means that $F(Y|\mathbf{X} = \mathbf{x}) = F(Y|\mathbf{Q}(\mathbf{X}) = \mathbf{Q}(\mathbf{x}))$, where $F(\cdot|\cdot)$ is the conditional distribution function of Y given the second argument. As no information about Y is lost when \mathbf{X} is replaced by $\mathbf{Q}(\mathbf{X})$, $\mathbf{Q}(\mathbf{X})$ is called a sufficient reduction for the regression of Y on \mathbf{X} .

Although, in principle, \mathbf{Q} can be any function, researchers have used only linear sufficient transformations $\mathbf{Q}(\mathbf{X}) = \boldsymbol{\eta}^T \mathbf{X} = (\eta_1^T \mathbf{X}, \dots, \eta_d^T \mathbf{X})$ with $\boldsymbol{\eta} = (\eta_1, \dots, \eta_d) \in \mathbb{R}^{p \times d}$ in SDR methodology so far (e.g., [2–6]). One important fact about linear sufficient reductions is that if the columns of a matrix $\boldsymbol{\gamma}$ span the vector space spanned by the columns of $\boldsymbol{\eta}$, then $\boldsymbol{\eta}^T \mathbf{X}$ and $\boldsymbol{\gamma}^T \mathbf{X}$ contain the same information about Y , and therefore, inference needs to focus on the subspace $\mathcal{S} = \text{span}(\boldsymbol{\eta})$ spanned by the columns of $\boldsymbol{\eta}$ and not on the matrix $\boldsymbol{\eta}$ itself. We call such a subspace an SDR subspace.

Under mild conditions, the intersection of all SDR subspaces, called the central dimension reduction subspace and denoted by $\mathcal{S}_{Y|X}$ [4, Section 6.4, pp 108–112], is also sufficient. In particular, $\mathcal{S}_{Y|X}$ exists when the support of the distribution of \mathbf{X} is convex. The dimension $d = \dim(\mathcal{S}_{Y|X})$, referred to as the structural dimension of the regression of Y on \mathbf{X} , can take on any value in the set $\{0, 1, \dots, p\}$. When $d < p$, the structural dimension of the regression is smaller than the number of predictors, and the complexity of the regression is reduced.

If $\boldsymbol{\eta} = (\eta_1, \dots, \eta_d)$ is a basis for $\mathcal{S}_{Y|X}$, then the d linear combinations $\boldsymbol{\eta}^T \mathbf{X} = (\eta_1^T \mathbf{X}, \dots, \eta_d^T \mathbf{X})$ contain all the information in \mathbf{X} about Y . Moreover, the number of these linear combinations of \mathbf{X} , d , is the smallest number of linear combinations of \mathbf{X} with this property. In this sense, $\boldsymbol{\eta}^T \mathbf{X} = (\eta_1^T \mathbf{X}, \dots, \eta_d^T \mathbf{X})$ is a ‘minimal’ transformation of the covariate vector \mathbf{X} for the regression of Y on it. Thus, the goal of SDR methodology is to estimate d and a set of basis vectors $\boldsymbol{\eta} = (\eta_1, \dots, \eta_d)$ such that $\text{span}(\boldsymbol{\eta}) \subseteq \mathcal{S}_{Y|X}$. A detailed exposition of SDR methodology is provided in [4].

2.2. The first-moment-based sufficient dimension reduction subspace

If $\mathbf{Q}(\mathbf{X})$ is a sufficient reduction for the forward regression $Y|\mathbf{X}$, then it is also a sufficient statistic for the inverse regression $\mathbf{X}|Y$ [5]. This result implies that we can use the inverse regression of \mathbf{X} on Y to find sufficient reductions for the forward regression of Y on \mathbf{X} . The advantage of inverse regression is that a complex, p -dimensional multiple forward regression of Y on \mathbf{X} that cannot easily be visualized is translated to p univariate inverse regressions of each component of \mathbf{X} on Y . Thus, a p -dimensional problem is mapped to p one-dimensional ones so that the difficulty of inference in high dimensions is removed. Because of this, most SDR methods are based on inverse regression.

In general, the estimation of $\mathcal{S}_{Y|X}$ is based on finding a matrix $\boldsymbol{\Omega}$, called kernel, so that $\text{span}(\boldsymbol{\Omega}) \subseteq \mathcal{S}_{Y|X}$. First-moment methods such as SIR [2], principal fitted components (PFC) [6], and parametric inverse regression (PIR) [7] use kernel matrices computed from first moments of $\mathbf{X}|Y$, whereas second-moment methods such as sliced average variance estimation [3] use second-moment-based kernels. In this paper, we focus on first-moment methods and, in particular, on extensions of SIR to longitudinal settings.

The FMSDR subspace is the span of the centered mean of the inverse regression of \mathbf{X} on Y , $E(\mathbf{X}|Y) - E(\mathbf{X})$, scaled by the inverse of the marginal covariance of \mathbf{X} , Σ . That is,

$$\mathcal{S}_{\text{FMSDR}} = \text{span}(\Sigma^{-1}[E(\mathbf{X}|Y) - E(\mathbf{X})]). \quad (1)$$

This subspace is the estimation target of SIR with kernel matrix $\Omega_{\text{SIR}} = \Sigma^{-1} \text{cov}(E(\text{vec}(\mathbf{X}|Y)))$.

If the predictors \mathbf{X} satisfy the *linearity condition* [4, p. 188], then $\mathcal{S}_{\text{FMSDR}}$ is contained in the central subspace,

$$\mathcal{S}_{\text{FMSDR}} \subseteq \mathcal{S}_{Y|X}. \quad (2)$$

This condition requires that there exists a linear sufficient transformation $\eta^T \mathbf{X}$ such that $E(\mathbf{X}|\eta^T \mathbf{X}) - E(\mathbf{X})$ is a linear function of $\eta^T \mathbf{X}$ [8]. An important feature of this condition is that it is required only for the marginal distribution of the predictors. The linearity condition holds when \mathbf{X} has an elliptical distribution [9], for example, a multivariate normal distribution, but normality is much stronger than needed. Coordinate-wise transformations or reweighting of the predictors can help achieve ellipticity of \mathbf{X} [10]. Moreover, the linearity condition holds approximately when p is large [11].

We show in Appendix A that for $\Delta_Y = \text{cov}(\mathbf{X}|Y)$ and $\Delta = E(\Delta_Y)$, the PFC subspace defined by $\mathcal{S}_{\text{PFC}} = \text{span}\{\Delta^{-1}[E(\mathbf{X}|Y) - E(\mathbf{X})]\}$ also equals the FMSDR subspace under the linearity condition. That is,

$$\mathcal{S}_{\text{FMSDR}} = \text{span}(\Sigma^{-1}[E(\mathbf{X}|Y) - E(\mathbf{X})]) = \text{span}(\Delta^{-1}[E(\mathbf{X}|Y) - E(\mathbf{X})]). \quad (3)$$

When the marginal distribution of the predictors \mathbf{X} satisfies the linearity condition, we can see from (2) and (3) that one can scale the conditional data, $\mathbf{X}|Y$, by using either the marginal covariance matrix Σ or the conditional Δ to obtain $\mathcal{S}_{\text{FMSDR}}$ and that

$$\mathcal{S}_{\text{FMSDR}} = \mathcal{S}_{\text{PFC}} = \mathcal{S}_{\text{SIR}} \subseteq \mathcal{S}_{Y|X}.$$

For $\mathcal{S}_{\text{FMSDR}}$ to recover all of $\mathcal{S}_{Y|X}$, and not just part of it, that is, to obtain $\mathcal{S}_{\text{FMSDR}} = \mathcal{S}_{Y|X}$, $\text{cov}(\eta^T \mathbf{X}|Y)$ must be positive definite for a linear sufficient transformation $\eta^T \mathbf{X}$ [12]. This is true, for example, when the predictors $\mathbf{X}|Y$ have a multivariate normal distribution and when the conditional covariance of \mathbf{X} does not depend on Y .

Under the linearity condition, any kernel matrix whose column space spans the same space as $\mathcal{S}_{\text{FMSDR}}$ can be used to estimate $\mathcal{S}_{Y|X}$ or a part of it. Different kernel matrices define different estimation methods, such as SIR, PIR, and PFC. Here, we focus on SIR and its extension to longitudinal data.

3. Longitudinal first-moment-based sufficient dimension reduction subspace

3.1. Notation and assumptions

We now estimate the first-moment-based dimension reduction subspace $\mathcal{S}_{\text{FMSDR}}$ for longitudinal data. Throughout this section, we assume that the linearity condition holds for \mathbf{X} .

In the longitudinal setting, the $p \times 1$ predictor vector \mathbf{X} is measured over T time points. To be specific, let $Y_i \in \mathbb{R}$ be a one-dimensional response variable for the i th individual in the study, with p -dimensional covariate vector $\mathbf{X}_{it} = (x_{i1t}, \dots, x_{ipt})^T \in \mathbb{R}^p$ that is measured at time points $t = 1, \dots, T_i$, for $i = 1, \dots, n$, where n is the total sample size. For notational simplicity, we assume that the markers are measured at the same time points for all individuals and that all individuals have the same number of observations over time; that is, $t_{ij} = t_j$, and $T_i = T$. The predictor vector over time can be represented as the $p \times T$ matrix

$$\mathbf{X}_i = (\mathbf{X}_{i1}, \dots, \mathbf{X}_{iT}) = \begin{bmatrix} X_{i11} & \cdots & X_{i1T} \\ X_{i21} & \cdots & X_{i2T} \\ \vdots & \vdots & \vdots \\ X_{ip1} & \vdots & X_{ipT} \end{bmatrix} \quad (4)$$

that corresponds to the $pT \times 1$ column vector $\text{vec}(\mathbf{X}_i)$ composed of the columns of \mathbf{X}_i stacked one after the other.

If one ignores the time structure of the data, standard dimension reduction can be applied to find $\eta \in \mathbb{R}^{pT \times d}$ such that $F(Y|\text{vec}(\mathbf{X})) = F(Y|\eta^T \text{vec}(\mathbf{X}))$. However, the time structure is integral to the nature of longitudinal data, and ignoring it could result in loss of accuracy in estimation for practically relevant sample sizes.

To accommodate the longitudinal structure of $\mathbf{X}|Y$, we assume that the first two moments of \mathbf{X} can be decomposed into a time and a marker component. In particular, we let both the mean and the covariance of the markers be Kronecker products of the two components to further improve classification or regression, as follows.

Assumption 1

The conditional mean of \mathbf{X} given Y has the following Kronecker product structure:

$$\text{vec}(E(\mathbf{X}|Y) - E(\mathbf{X})) = (\alpha \otimes \beta)\text{vec}(v_y), \tag{5}$$

for some $d \times s$ matrix v_y with $E(v_y) = 0$ and $\det(\text{cov}(\text{vec}(v_y))) > 0$. The matrix $\beta \in \mathbb{R}^{T \times s}$ captures the mean structure over time, and $\alpha \in \mathbb{R}^{p \times d}$ captures the mean structure of the markers regardless of time.

A practically important case is that of binary outcomes; that is, $Y = 0, 1$. For binary Y , the condition (5) is satisfied if, for example, $\text{vec}(E(\mathbf{X}|Y)) = \alpha_y \otimes \beta$, which implies that the means of the markers change over time only by a multiplicative factor that affects all markers equally and is the same for the two groups defined by Y ; that is, $\text{vec}(E(\mathbf{X}_t|Y)) = \beta_t \alpha_y$. Letting $p_y = P(Y = y)$ and using that $E(\mathbf{X}) = p_0 \alpha_0 \otimes \beta + (1 - p_0) \alpha_1 \otimes \beta$, we get that $\text{vec}(E(\mathbf{X}|Y = y) - E(\mathbf{X})) = (1 - p_y)(\alpha_0 - \alpha_1) \otimes \beta$, which means that the first-order moment condition is satisfied with $v_y = (1 - p_y)$.

Moreover, we assume that the second moments of \mathbf{X} , either unconditional or conditional on Y , have a Kronecker product structure (Assumption 2). In the next two sections, we study the second-moment assumptions and show that under this structure, our target estimation subspace $\mathcal{S}_{\text{FM\textsubscript{SDR}}}$ is itself spanned by Kronecker products of the marker and time components.

Adopting this moment structure accommodates the longitudinal nature of the data by focusing on the predictor and time aspects separately and results in a substantial reduction in the number of parameters to estimate. Researchers have previously used a Kronecker product structure of moments in linear discriminant analysis procedures for repeated measurements under the assumption of normally distributed data (e.g., [13] and [14]).

3.2. $\Sigma = \text{cov}(\text{vec}(\mathbf{X})) = \Sigma_1 \otimes \Sigma_2$

When the longitudinal data arise from a prospective cohort, it may be reasonable to assume that the $(Tp) \times (Tp)$ covariance matrix of $\text{vec}(\mathbf{X})$ can be written as $\Sigma = \Sigma_1 \otimes \Sigma_2$ (Assumption 2a), where Σ_1 captures the covariance structure of the p markers and Σ_2 models the temporal association. Thus, $\text{cov}(X_{it}, X_{js}) = \sigma_{ij}^{(1)} \sigma_{ts}^{(2)}$. In terms of correlation, this Kronecker product structure implies that

$$\text{cor}(X_{it}, X_{js}) = \frac{\sigma_{ij}^{(1)}}{\sqrt{\sigma_{ii}^{(1)} \sigma_{jj}^{(1)}}} \frac{\sigma_{ts}^{(2)}}{\sqrt{\sigma_{tt}^{(2)} \sigma_{ss}^{(2)}}}.$$

For the same marker measured at two different time points, $\text{cov}(X_{it}, X_{is}) = \sigma_{ii}^{(1)} \sigma_{ts}^{(2)}$, and for two markers measured at the same time point, $\text{cov}(X_{it}, X_{jt}) = \sigma_{ij}^{(1)} \sigma_{tt}^{(2)}$. This means that the correlation of two markers measured at the same time point does not depend on time, that is, $\text{cor}(X_{it}, X_{jt}) = \sigma_{ij}^{(1)} / (\sigma_{ii}^{(1)} \sigma_{jj}^{(1)})^{-1/2}$, and that the correlation of the same marker measured at two different time points is the same for all markers, $\text{cor}(X_{it}, X_{is}) = \sigma_{st}^{(2)} / (\sigma_{tt}^{(2)} \sigma_{ss}^{(2)})^{-1/2}$. As $\Sigma_1 \otimes \Sigma_2 = (c \Sigma_1) \otimes (c^{-1} \Sigma_2)$, the matrices Σ_1 and Σ_2 are determined up to a multiplicative constant, and we assume without loss of generality that $\sigma_{TT}^{(2)} = 1$.

In Theorem 1, we show that $\mathcal{S}_{\text{FM\textsubscript{SDR}}}$ has a Kronecker structure induced by the Kronecker structure of Σ and the conditional mean of \mathbf{X} . The proof is given in Appendix A.

Theorem 1

Suppose that Assumption 1 holds and suppose that $\text{cov}(\text{vec}(\mathbf{X})) = \Sigma = \Sigma_1 \otimes \Sigma_2$. Then,

$$\mathcal{S}_{\text{FM\textsubscript{SDR}}} = \text{span}(\Sigma_1^{-1} \alpha \otimes \Sigma_2^{-1} \beta). \tag{6}$$

It follows immediately that $\mathcal{S}_{\text{SIR}} = \text{span}(\boldsymbol{\Omega}_{\text{SIR}}) = \text{span}(\boldsymbol{\Sigma}^{-1} \text{cov}(\text{vec}(\mathbf{E}(\mathbf{X}|Y)))) = \text{span}(\boldsymbol{\Sigma}_1^{-1} \boldsymbol{\alpha} \otimes \boldsymbol{\Sigma}_2^{-1} \boldsymbol{\beta})$.

$$3.3. \quad \boldsymbol{\Delta} = \mathbf{E}[\text{cov}(\text{vec}(\mathbf{X}|Y))] = \boldsymbol{\Delta}_1 \otimes \boldsymbol{\Delta}_2$$

Sometimes, assuming that the overall population covariance $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_1 \otimes \boldsymbol{\Sigma}_2$ may not be reasonable, for example, when the data arise from a retrospective case-control sample. A slightly less restrictive assumption is to impose a Kronecker product structure on the conditional covariances of the predictors given the response Y . That is, $\text{cov}(\text{vec}(\mathbf{X})|Y = y) = \boldsymbol{\Delta}_1^y \otimes \boldsymbol{\Delta}_2^y$. When either $\boldsymbol{\Delta}_1^y$ or $\boldsymbol{\Delta}_2^y$ does not depend on Y , this structure implies that $\boldsymbol{\Delta} = \mathbf{E}(\text{var}(\mathbf{X}|Y)) = \boldsymbol{\Delta}_1 \otimes \boldsymbol{\Delta}_2$ as well. If the markers are sampled over the same time points in the groups defined by Y , the assumption that the correlations of the markers over time $\boldsymbol{\Delta}_1$ do not depend on Y is reasonable, even though correlations of the markers differ for the groups defined by Y . If $\boldsymbol{\Delta} = \mathbf{E}(\text{cov}(\mathbf{X}|Y)) = \boldsymbol{\Delta}_1 \otimes \boldsymbol{\Delta}_2$ (Assumption 2b), then the following theorem, whose proof is analogous to that of Theorem 1, holds.

Theorem 2

Suppose that Assumption 1 holds and suppose that $\boldsymbol{\Delta} = \mathbf{E}(\text{cov}(\mathbf{X}|Y)) = \boldsymbol{\Delta}_1 \otimes \boldsymbol{\Delta}_2$. Then,

$$\mathcal{S}_{\text{PFC}} = \text{span}(\boldsymbol{\Delta}_1^{-1} \boldsymbol{\alpha} \otimes \boldsymbol{\Delta}_2^{-1} \boldsymbol{\beta}). \tag{7}$$

Theorems 1 and 2 imply that when the conditional first moment $\text{vec}(\mathbf{E}(\mathbf{X}|Y) - \mathbf{E}(\mathbf{X}))$ and either $\text{cov}(\text{vec}(\mathbf{X}))$ or $\mathbf{E}(\text{cov}(\text{vec}(\mathbf{X})|Y))$ have a Kronecker product structure, then the structure of $\mathcal{S}_{\text{FMSDR}}$ is also a Kronecker product. In particular, when the assumptions of Theorem 1 are satisfied,

$$\mathcal{S}_{\text{FMSDR}} = \mathcal{S}_{\text{PFC}} = \mathcal{S}_{\text{SIR}} = \text{span}(\boldsymbol{\Sigma}_1^{-1} \boldsymbol{\alpha} \otimes \boldsymbol{\Sigma}_2^{-1} \boldsymbol{\beta}) \subseteq \mathcal{S}_{Y|X}. \tag{8}$$

If the conditions of Theorem 2 are satisfied,

$$\mathcal{S}_{\text{FMSDR}} = \mathcal{S}_{\text{PFC}} = \mathcal{S}_{\text{SIR}} = \text{span}(\boldsymbol{\Delta}_1^{-1} \boldsymbol{\alpha} \otimes \boldsymbol{\Delta}_2^{-1} \boldsymbol{\beta}) \subseteq \mathcal{S}_{Y|X}. \tag{9}$$

On the basis of the results in (8) and (9), the estimation of $\mathcal{S}_{\text{FMSDR}}$ requires estimating many fewer parameters compared with the setting when no structure is imposed.

Remark

Under the additional assumption of multivariate normality of the markers, we can use likelihood ratio tests to assess if Assumptions 1 and 2a hold [15, 16].

4. The longitudinal sliced inverse regression algorithm

We estimate $\mathcal{S}_{\text{FMSDR}}$ by using the SIR kernel matrix $\boldsymbol{\Omega}_{\text{SIR}} = \boldsymbol{\Sigma}^{-1} \text{cov}(\mathbf{E}(\mathbf{X}|Y))$ based on repeated measurements of \mathbf{X} . In Section 3, we showed that under the conditions of either Theorem 1 or Theorem 2, the SIR subspace has a Kronecker product structure, that is, $\mathcal{S}_{\text{FMSDR}} = \text{span}(\boldsymbol{\Omega}_x) = \text{span}(\boldsymbol{\Omega}_{1x} \otimes \boldsymbol{\Omega}_{2x})$, where $\boldsymbol{\Omega}_{1x}$, a $p \times p$ matrix, captures the reduction with respect to the predictors and $\boldsymbol{\Omega}_{2x}$, a $T \times T$ matrix, captures the reduction of the markers with respect to time.

In the implementation of the algorithm, we standardize the predictors for numerical stability. Depending on whether the conditions of Theorem 1 or Theorem 2 are satisfied, we can standardize the predictors (centered and scaled) by using $\boldsymbol{\Sigma}$ or $\boldsymbol{\Delta}$, respectively, for their scaling (see (8) and (9)). However, it can be shown by direct calculation, as in the proof of Corollary 3.4 in [6], that if one replaces $\widehat{\boldsymbol{\Sigma}}_1$ and $\widehat{\boldsymbol{\Sigma}}_2$ by $\widehat{\boldsymbol{\Delta}}_1$ and $\widehat{\boldsymbol{\Delta}}_2$, respectively, in scaling the predictors, $\widehat{\mathcal{S}}_{\text{FMSDR}} (= \widehat{\mathcal{S}}_{\text{SIR}})$ remains the same. Hence, the algorithm uses the simpler $\widehat{\boldsymbol{\Sigma}}_1$ and $\widehat{\boldsymbol{\Sigma}}_2$ to scale the data.

We use the standardized predictors $\mathbf{Z} = \boldsymbol{\Sigma}^{-1/2}(\mathbf{X} - \mathbf{E}(\mathbf{X}))$ to compute the kernel matrix $\boldsymbol{\Omega}_z = \text{cov}(\mathbf{E}(\mathbf{Z}|Y))$. The latter relates to the kernel matrix for the original \mathbf{X} predictors through $\boldsymbol{\Omega}_x = \boldsymbol{\Sigma}^{1/2} \boldsymbol{\Omega}_z \boldsymbol{\Sigma}^{-1/2}$.

Suppose that a sample $(\mathbf{X}_i, Y_i), i = 1, \dots, n$ is available, where Y_i denotes the outcome variable and \mathbf{X}_i denotes the $(p \times T)$ matrix of predictors for individual i given in (4).

- The p -vector of covariate values for person i at time t is $\mathbf{X}_{it} = (x_{i1t}, \dots, x_{ipt})^T$, and $\bar{\mathbf{x}}_t = (\bar{x}_{1t}, \dots, \bar{x}_{pt})^T$ denotes the p -vector of predictor means over all subjects at time t . Similarly, $\mathbf{X}_{ij} = (x_{ij1}, \dots, x_{ijT})^T$ denotes the T -vector of the j th covariate values across all T time points for individual i , and $\bar{\mathbf{x}}_j$ is the T -vector of means of the j th predictor, $j = 1, \dots, p$. Let n_t be the

number of observations at time t across all p predictors. Similarly, n_j is the number of observations of the j th predictor across all time points.

- We normalize the predictor matrix \mathbf{X}_i for the i th individual: $\text{vec}(\mathbf{Z}_i) = \widehat{\Sigma}^{-1/2}(\text{vec}(\mathbf{X}_i) - \text{vec}(\bar{\mathbf{x}}))$, where each entry of $\bar{\mathbf{x}}$ is the empirical mean for each predictor at each time point over the i observations. We find $\widehat{\Sigma} = \widehat{\Sigma}_1 \otimes \widehat{\Sigma}_2$ by computing

$$\widehat{\Sigma}_{1t} = \text{cov}(\mathbf{x}_{\cdot t}) = \frac{1}{n_t} \sum_{i=1}^{n_t} (\mathbf{x}_{it} - \bar{\mathbf{x}}_{\cdot t})(\mathbf{x}_{it} - \bar{\mathbf{x}}_{\cdot t})^T \quad (10)$$

$$\widehat{\Sigma}_1 = \frac{1}{T} \sum_{t=1}^T \widehat{\Sigma}_{1t}, \quad (11)$$

and

$$\widehat{\Sigma}_{2j} = \text{cov}(\mathbf{x}_{j \cdot}) = \frac{1}{n_j} \sum_{i=1}^{n_j} (\mathbf{x}_{ij} - \bar{\mathbf{x}}_{j \cdot})(\mathbf{x}_{ij} - \bar{\mathbf{x}}_{j \cdot})^T \quad (12)$$

$$\widehat{\Sigma}_2 = \frac{1}{p} \sum_{j=1}^p \widehat{\Sigma}_{2j}. \quad (13)$$

- As in standard SIR, we divide the support of Y into H slices. For continuous Y , this means that Y is replaced by a discrete version \tilde{Y} based on partitioning the observed range of Y into H fixed, non-overlapping slices. For categorical Y , the slices are the categories of Y . Let $\bar{\mathbf{z}}_t^{(h)}$ denote the p -vector of the standardized predictor means and $\bar{\mathbf{z}}_j^{(h)}$ be the T -vector of the means of the j th standardized predictor, $j = 1, \dots, p$, within slice h , $h = 1, \dots, H$. The proportion of observations in slice h at time t across all p predictors is $f_{\cdot t}^{(h)}$, and $f_{j \cdot}^{(h)}$ is the proportion of observations in slice h of the j th predictor across all time points. Let

$$\widehat{\Omega}_{1t} = \sum_{h=1}^H f_{\cdot t}^{(h)} \bar{\mathbf{z}}_t^{(h)} \bar{\mathbf{z}}_t^{(h)T} \quad (14)$$

$$\widehat{\Omega}_{1z} = \frac{1}{T} \sum_{t=1}^T \widehat{\Omega}_{1t}, \quad (15)$$

and

$$\widehat{\Omega}_{2j} = \sum_{h=1}^H f_{j \cdot}^{(h)} \bar{\mathbf{z}}_j^{(h)} \bar{\mathbf{z}}_j^{(h)T} \quad (16)$$

$$\widehat{\Omega}_{2z} = \frac{1}{p} \sum_{j=1}^p \widehat{\Omega}_{2j}. \quad (17)$$

Then,

$$\widehat{\text{cov}}(\mathbf{E}(\mathbf{Z}|Y)) = \widehat{\Omega}_{\text{LSIR}} = \widehat{\Omega}_{1z} \otimes \widehat{\Omega}_{2z}. \quad (18)$$

- We separately compute the singular value decompositions $\widehat{\Omega}_{1z} = \widehat{\mathbf{U}}_1 \widehat{\mathbf{D}}_1 \widehat{\mathbf{R}}_1^T$ and $\widehat{\Omega}_{2z} = \widehat{\mathbf{U}}_2 \widehat{\mathbf{D}}_2 \widehat{\mathbf{R}}_2^T$ to write $\widehat{\Omega}_{\text{LSIR}} = \widehat{\Omega}_{1z} \otimes \widehat{\Omega}_{2z} = (\widehat{\mathbf{U}}_1 \otimes \widehat{\mathbf{U}}_2)(\widehat{\mathbf{D}}_1 \otimes \widehat{\mathbf{D}}_2)(\widehat{\mathbf{R}}_1^T \otimes \widehat{\mathbf{R}}_2^T)$.
- To test for dimension, we estimate the rank of $\widehat{\Omega}_{1z}$ and $\widehat{\Omega}_{2z}$ separately by using the weighted chi-square test for dimension in [17], and the dimension of $\widehat{\Omega}_{\text{LSIR}}$ by the product of the two estimated ranks because $\text{rank}(\widehat{\Omega}_{\text{LSIR}}) = \text{rank}(\widehat{\Omega}_{1z}) \times \text{rank}(\widehat{\Omega}_{2z})$.

- Letting $\text{rank}(\widehat{\boldsymbol{\Omega}}_{iz}) = d_i$, an estimate of the central subspace is given by

$$\widehat{\boldsymbol{S}}_{\text{FMSDR}} = \widehat{\boldsymbol{S}}_{\text{SIR}} = \text{span}(\widehat{\boldsymbol{\Sigma}}_1^{-1}(\hat{U}_{11}, \dots, \hat{U}_{1d_1}) \otimes \widehat{\boldsymbol{\Sigma}}_2^{-1}(\hat{U}_{21}, \dots, \hat{U}_{2d_2})).$$

- To test for contributions of specific combinations of time points or predictors, we apply the test statistic proposed in Theorem 2 [18] separately to $\widehat{\boldsymbol{\Omega}}_{1z}$ and $\widehat{\boldsymbol{\Omega}}_{2z}$. That is, we use a Wald-type test to test hypotheses of the form $H_0 : \text{Cvec}(\mathbf{U}_{d_i}) = \mathbf{0}$ versus $H_1 : \text{Cvec}(\mathbf{U}_{d_i}) \neq \mathbf{0}$ where $\mathbf{U}_{d_i} = (U_{i1}, \dots, U_{id_i})$, $i = 1, 2$ for a prespecified $r \times kd_i$ matrix \mathbf{C} of zeroes and ones, where $k = p$ for testing contributions of specific marker combinations and $k = T$ for specific time points. The rank of \mathbf{C} , r , equals the number of the elements of \mathbf{U}_{d_i} set to zero. The test for marker contribution requires the computation of the asymptotic distribution of $\widehat{\boldsymbol{\Omega}}_z$, which is derived in Appendix A.

We focus on the longitudinal SIR (LSIR) algorithm that naturally accommodates categorical responses as the slices are by default the categories. For continuous responses, alternatively, one can use the longitudinal version of either PIR [7] or PFC [6], where the inverse regressions are fitted parametrically and the estimation is expected to be more accurate.

5. Simulations

In the simulations, we focus on binary outcomes, $Y = 0, 1$. We compare LSIR with standard SIR applied to the data ignoring their longitudinal structure; that is, we apply SIR to \mathbf{X} treated as the $pT \times 1$ vector, $\text{vec}(\mathbf{X})$. For binary Y , SIR and LSIR estimate at most a single direction in $S_{Y|\mathbf{X}}$, that is, $\mathcal{S}_{\text{FMSDR}}$ is given by a vector. The projection onto the space spanned by the SIR or LSIR kernel matrix can thus directly be used as a scalar diagnostic score. For categorical or continuous Y , if the dimension of the subspace is estimated to be larger than 1, the projections need to be further combined to obtain a scalar score, for example, by using the procedure in [1].

We quantify the discriminatory performance of the SIR and LSIR diagnostic scores with respect to the AUC, the area under the receiver-operator characteristics (ROC) curve. The ROC curve plots sensitivity against (1-specificity) (true versus false positivity) for all thresholds that can be used to define ‘test positive’ [19, p. 67]. The AUC values vary between 0.5 and 1, where 1 corresponds to perfect discriminatory accuracy of the test and 0.5 corresponds to no discriminatory ability. We can also express the AUC as the probability that the scalar diagnostic score for a randomly selected case S_1 exceeds that for a randomly selected control S_0 ; that is, $AUC = P(S_1 > S_0)$. We can compare two diagnostic tests by calculating the difference between their AUC values, with the larger values corresponding to the ‘better’ test.

We generated samples of $p = 5$ and $p = 10$ markers $\mathbf{X}^T = (X_1, \dots, X_p)$ measured over $T = 5$ and $T = 10$ time points from two normal populations with equal covariance matrices, $(\mathbf{X}|Y = i) \sim \text{MVN}(\boldsymbol{\alpha}_i \otimes \boldsymbol{\beta}, \boldsymbol{\Delta})$, $i = 0, 1$. Each $\boldsymbol{\alpha}_i$, $i = 0, 1$ is a vector of length p , and $\boldsymbol{\beta}$ is a vector of length T . We let $\boldsymbol{\alpha}_1 = p^{-1/2}(1, \dots, 1)$; $\boldsymbol{\alpha}_0$ has entries $\alpha_0(k) = 0$ for $k = 1, 3, 5$ and $\alpha_0(k) = p^{-1/2}$ for $k \neq 1, 3, 5$, and the entries of $\boldsymbol{\beta}$ are $\beta(k) = (T - k + 1)^{-1}$. We chose these values to reflect practically relevant AUC values. The choice of the $\boldsymbol{\beta}$ coefficients leads to larger differences in the group means for later time points; that is, measurements more proximal to Y contribute more to discriminating the two groups. We let the number of observations in the $Y = 0$ and $Y = 1$ groups be equal, $N = N_0 = N_1$. To obtain unbiased estimates of the AUC, we used an independently generated sample of 500 cases and 500 controls.

For reference, we also computed the AUC value by using 100,000 cases and 100,000 controls for the linear marker combination under the assumption that $\boldsymbol{\alpha}$, $\boldsymbol{\beta}$, and $\boldsymbol{\Sigma}$ are known exactly. We report this value as the ‘true AUC’.

The overall covariance matrix of the predictors $\boldsymbol{\Sigma}$ has the Kronecker structure, $\boldsymbol{\Sigma} = \boldsymbol{\Sigma}_1 \otimes \boldsymbol{\Sigma}_2$. We assumed an AR(1) structure for $\boldsymbol{\Sigma}_1$, that is, $\text{cor}(X_{ij}, X_{ik}) = \rho_p^{|k-j|}$, and similarly for $\boldsymbol{\Sigma}_2$, $\text{cor}(X_{ij}, X_{kj}) = \rho_T^{|k-j|}$, and for various choices of ρ_p and ρ_T . We computed the group-specific covariance matrix $\boldsymbol{\Delta}$ by using

$$\begin{aligned} \boldsymbol{\Sigma} &= \text{cov}(\text{vec}(\mathbf{X})) = E(\text{cov}(\text{vec}(\mathbf{X})|Y)) + \text{cov}(E(\text{vec}(\mathbf{X})|Y)) \\ &= \boldsymbol{\Delta} + E\{E(\text{vec}(\mathbf{X})|Y)E(\text{vec}(\mathbf{X})^T|Y)\} - E(\text{vec}(\mathbf{X}))E(\text{vec}(\mathbf{X})^T). \end{aligned}$$

For a model with $p = 5$ markers measured over $T = 10$ time points and the same mean and covariance matrices for both groups, $(\mathbf{X}|Y = i) \sim \text{MVN}(0, \boldsymbol{\Delta})$, $i = 0, 1$, the AUC estimates (with standard

errors) were 0.52 (0.01) for the truth, 0.51 (0.01) for LSIR, and 0.51 (0.01) for SIR for $N = 100$ cases and controls. Both methods thus performed well when there was no difference in the distribution of the predictors between the two groups.

Table I presents means over 100 simulations for various values of p , T and correlations ρ_T and ρ_p . For $p = T = 5$ with $\rho_p = \rho_T = 0.3$, the AUC values (with empirical standard errors in parentheses) for LSIR and SIR were 0.728 (0.021) and 0.699 (0.023) for $N = 100$ cases and controls, and 0.748 (0.013) and 0.741 (0.014) for $N = 500$ cases and controls. For the same setting with $p = T = 10$, the differences were more pronounced, with AUC values of LSIR and SIR 0.643 (0.034) and 0.577 (0.029) for $N = 100$ cases and controls and 0.699 (0.015) and 0.661 (0.017) for $N = 500$ cases and controls.

On the basis of the paired t -test, LSIR had significantly higher AUC values than SIR, even for $N = 500$ cases and $N = 500$ controls. For example, for $N = 500$ cases and controls, for $p = T = 5$ with $\rho_p = \rho_T = 0.3$, the paired t -test p -value was 0.0002, and for $p = T = 5$ with $\rho_p = 0, \rho_T = 0.4$, the paired t -test p -value was 0.0001.

Results were very similar when the markers were uncorrelated; that is, $\rho_T = \rho_p = 0$. In this setting, the AUC was higher than that of the correlated case, around 0.79 for the settings studied. SIR always resulted in lower AUC values for smaller sample sizes than LSIR. For $N = 100$, the difference was around 5%, a similar difference to the correlated case.

We also assessed the robustness of LSIR to violations of the moment assumptions. First, we violated the assumption of the Kronecker structure of the mean but let Σ have a Kronecker structure. The mean vector for the controls was equal to zero, and the mean vector for cases had values $(1/3pT, \dots, pT/3pT)$. For $p = 5$ and $T = 6$ with $\rho_t = 0.4$ and $\rho_p = 0.3$, the true AUC value was 0.73, and the AUCs (with empirical standard errors in parentheses) for LSIR and SIR were 0.64 (0.03) and 0.62 (0.03), respectively, for $N = 100$ cases and controls and 0.68 (0.02) and 0.67 (0.02) for $N = 500$ cases and controls. For $N = 1000$ cases and controls, LSIR and SIR both had $AUC = 0.68$ with standard error 0.02. We then assumed that the means had a Kronecker structure, with the same

Table I. Mean area under the curve (AUC) values over 100 simulations when Σ has a Kronecker structure and the correlations over time and markers follow an AR(1) structure for N cases and N controls.

p, T, ρ_p, ρ_T	AUC	Method	$N = 100$	$N = 200$	$N = 300$	$N = 400$	$N = 500$
5, 5, 0.3, 0.3	0.743	LSIR	0.728	0.743	0.746	0.747	0.748
			0.021	0.015	0.016	0.015	0.013
		SIR	0.699	0.728	0.735	0.737	0.741
			0.023	0.016	0.018	0.017	0.014
5, 10, 0.3, 0.3	0.758	LSIR	0.743	0.765	0.772	0.775	0.776
			0.023	0.017	0.015	0.016	0.015
		SIR	0.694	0.735	0.750	0.758	0.761
			0.024	0.023	0.016	0.018	0.017
10, 10, 0.3, 0.3	0.691	LSIR	0.643	0.677	0.690	0.695	0.699
			0.034	0.022	0.019	0.017	0.015
		SIR	0.577	0.619	0.641	0.654	0.661
			0.029	0.020	0.020	0.021	0.017
5, 5, 0, 0.4	0.724	LSIR	0.698	0.715	0.721	0.725	0.725
			0.026	0.019	0.016	0.015	0.017
		SIR	0.670	0.699	0.708	0.715	0.716
			0.025	0.019	0.018	0.016	0.018
5, 10, 0, 0.4	0.749	LSIR	0.747	0.766	0.772	0.777	0.778
			0.029	0.017	0.015	0.014	0.015
		SIR	0.700	0.737	0.750	0.760	0.765
			0.025	0.019	0.016	0.014	0.015
10, 10, 0, 0.4	0.698	LSIR	0.618	0.661	0.673	0.679	0.683
			0.039	0.022	0.019	0.018	0.019
		SIR	0.568	0.607	0.626	0.637	0.649
			0.021	0.024	0.021	0.020	0.018

Standard errors are given below the mean estimates. The size of the test set for the AUC computation was 500 cases and 500 controls.

choices of parameters as for the models in Table I, but the true covariance matrix Σ did not have a Kronecker structure. Instead, Σ had an AR(1) structure with $\rho = 0.3$. For $p = T = 5$, the AUC values (with empirical standard errors in parentheses) for LSIR and SIR were 0.718 (0.019) and 0.726 (0.021) for $N = 100$ cases and controls and 0.767 (0.015) and 0.770 (0.016) for $N = 500$ cases and controls. For the same setting with $p = 5$ and $T = 10$, the AUC values of LSIR and SIR were 0.769 (0.016) and 0.781 (0.015) for $N = 100$ cases and controls and 0.699 (0.015) and 0.661 (0.017) for $N = 500$ cases and controls. When the population means were zero and $(1/3pT, \dots, pT/3pT)$ for controls and cases, respectively, and the assumption on Σ was violated so that instead of a Kronecker product Σ had an AR(1) structure with $\rho = 0.4$, the true AUC value was 0.75, and the AUCs for LSIR and SIR were 0.681 (0.030) and 0.645 (0.024), respectively, for $N = 100$ cases and controls and 0.70 (0.02) and 0.69 (0.02), respectively, for $N = 500$ cases and controls. Thus, the LSIR procedure appears to be robust to violations of the assumption on the second moments.

R code for the LSIR algorithm, testing of dimension and testing of marker contributions is available from the first author upon request. R code that implements standard SIR is available in the package `dr` at <http://i.e.http://cran.r-project.org/cran.r-project.org/>.

6. Data analysis

We used data from the Vorarlberg Health Monitoring and Promotion Program [20], one of the world's largest ongoing population-based risk factor surveillance programs, to illustrate LSIR and compare it with SIR. The aim was to classify men into two groups: those who developed cancer at the end of follow-up and those who did not, by using three serum biomarkers. The markers were the log-transformed values of uric acid (*UA*), blood glucose (*GLUC*), and total serum cholesterol (*CHOL*). We used data from 100 male cancer patients who each had marker measurements at four time points prior to diagnosis, and we selected a sample of 100 controls who also had measurements at four time points and at the end of follow-up had not developed cancer. Thus, the $p = 3$ predictors $\mathbf{X}_t = (UA_t, GLUC_t, CHOL_t)^T$, $t = 1, \dots, 4$, were measured at $T = 4$ time points, thus the dimension of this classification problem is $3 \times 4 = 12$. The histograms of the log-transformed markers in the cases and the controls at time point $t = 1$, given in Figure 1, appeared reasonably symmetric.

In the cases, the means of the markers for the four time points were $\bar{\mathbf{X}}_{t_1} = (4.45, 5.38, 3.48)$, $\bar{\mathbf{X}}_{t_2} = (4.47, 5.41, 3.50)$, $\bar{\mathbf{X}}_{t_3} = (4.50, 5.40, 3.43)$, $\bar{\mathbf{X}}_{t_4} = (4.54, 5.43, 3.46)$. In the controls $\bar{\mathbf{X}}_{t_1} = (4.47, 5.43, 3.64)$, $\bar{\mathbf{X}}_{t_2} = (4.45, 5.39, 3.55)$, $\bar{\mathbf{X}}_{t_3} = (4.49, 5.42, 3.60)$, $\bar{\mathbf{X}}_{t_4} = (4.52, 5.38, 3.56)$. To check the assumption of Kronecker structure of the means for binary outcomes Y , we have to assess

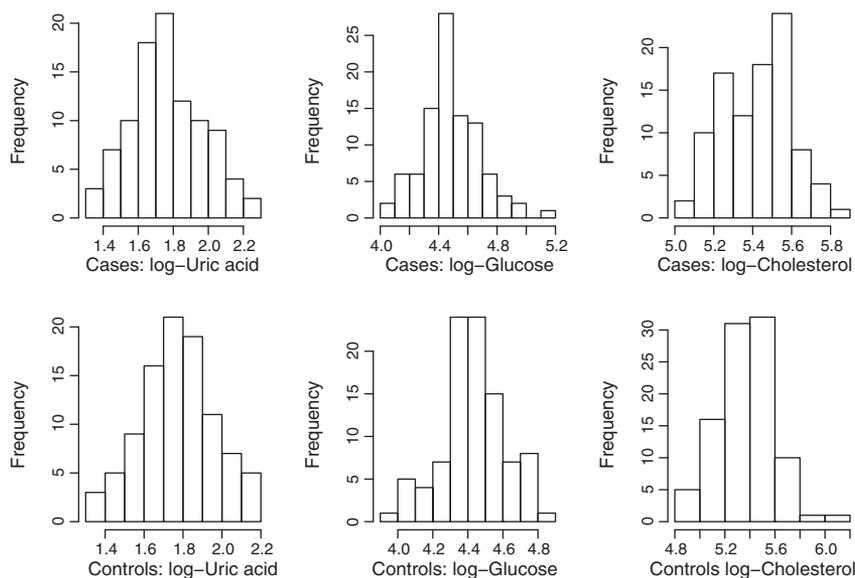


Figure 1. Distribution of the three biomarkers after log-transformation in cases (top panel) and controls (bottom panel).

if the means of the markers in the control and case groups vary only by a constant that may depend on time. Visual inspection of the ratios of the means for the different time points did not suggest violations of Assumption 1.

The correlations of the three markers across the four time points were similar, with

$$R_{UA} = \begin{bmatrix} 1.00 & 0.60 & 0.56 & 0.43 \\ 0.60 & 1.00 & 0.61 & 0.51 \\ 0.56 & 0.61 & 1.00 & 0.64 \\ 0.43 & 0.51 & 0.64 & 1.00 \end{bmatrix}, R_{GLUC} = \begin{bmatrix} 1.00 & 0.70 & 0.69 & 0.58 \\ 0.70 & 1.00 & 0.68 & 0.69 \\ 0.69 & 0.68 & 1.00 & 0.68 \\ 0.58 & 0.69 & 0.68 & 1.00 \end{bmatrix},$$

$$R_{CHOL} = \begin{bmatrix} 1.00 & 0.83 & 0.77 & 0.74 \\ 0.83 & 1.00 & 0.80 & 0.80 \\ 0.77 & 0.80 & 1.00 & 0.84 \\ 0.74 & 0.80 & 0.84 & 1.00 \end{bmatrix}.$$

Similarly, the correlations of the markers with each other for the four time points were similar:

$$R_{t1} = \begin{bmatrix} 1.00 & 0.04 & 0.23 \\ 0.04 & 1.00 & 0.25 \\ 0.23 & 0.25 & 1.00 \end{bmatrix}, R_{t2} = \begin{bmatrix} 1.00 & 0.05 & 0.12 \\ 0.05 & 1.00 & 0.10 \\ 0.12 & 0.10 & 1.00 \end{bmatrix},$$

$$R_{t3} = \begin{bmatrix} 1.00 & 0.16 & 0.11 \\ 0.16 & 1.00 & 0.11 \\ 0.11 & 0.11 & 1.00 \end{bmatrix}, R_{t4} = \begin{bmatrix} 1.00 & 0.13 & 0.02 \\ 0.13 & 1.00 & 0.17 \\ 0.02 & 0.17 & 1.00 \end{bmatrix}.$$

We used a likelihood ratio test based on assuming a multivariate normal distribution for the markers, to test the hypothesis of an unstructured covariance matrix against the alternative of a Kronecker product structure of the covariance matrix of the data [13]. Under the null hypothesis of a Kronecker structure, $-2 \log$ -likelihood = -1709.9 , and under the alternative hypothesis, $-2 \log$ -likelihood = -1781.60 , corresponding to a p -value 0.19 based on a chi-square distribution with $Tp(Tp + 1)/2 - T(T + 1)/2 - p(p + 1)/2 = 62$ degrees of freedom. Thus, there was no strong evidence in the data against the null hypothesis of a Kronecker product covariance structure.

The test for dimension estimated the dimension to be 1 for standard SIR ($p = 0.008$ for testing dimension 0 vs. 1) and 1 for LSIR for the time and marker parts separately.

For LSIR, the coefficients for the projection corresponding to time component were $\beta = (0.46, -0.31, 0.51, -0.66)$, and the coefficients for the marker component $\alpha = (UA, GLUC, CHOL) = (0.12, 0.99, 0.04)$. The resulting combined vector of coefficients for the projection was $(UA_1, GLUC_1, CHOL_1, \dots, UA_4, GLUC_4, CHOL_4) = (-0.05, -0.45, -0.02, 0.04, 0.30, 0.01, -0.06, -0.50, -0.02, 0.08, 0.66, 0.03)$. When we tested the contribution of individual markers using LSIR, there was evidence based on the Wald test that $GLUC$ significantly contributed to α ($p = 0.04$) whereas UA and $CHOL$ did not ($p = 0.73$ and $p = 0.54$, respectively). Applying the test to the time component showed that t_4 contributed significantly to β ($p = 0.05$).

On the basis of standard SIR, the coefficients for the projection were $(UA_1, GLUC_1, CHOL_1, \dots, UA_4, GLUC_4, CHOL_4) = (-0.16, -0.52, -0.02, 0.21, 0.25, 0.10, -0.01, -0.35, -0.12, 0.06, 0.67, -0.02)$. When we tested for significance of the contributions of individual markers at given time points, glucose was significant at the first time point ($p = 0.01$) and at the fourth time point ($p < 0.01$). However, we did not find the contribution of glucose at all four time points jointly to be significantly associated with outcome.

We also fitted a logistic regression model that included all predictors to the data. For binary outcomes, logistic regression is equivalent to standard SIR and yields the exact same projection. The corresponding log-odds ratios were $(UA_1, GLUC_1, CHOL_1, \dots, UA_4, GLUC_4, CHOL_4) = (-1.01, -3.21, -0.12, 1.18, 1.47, 0.62, -0.02, -2.18, -0.69, 0.32, 4.33, -0.14)$. On the basis of the logistic model, only glucose measured at the first and fourth time points was significantly associated with outcome ($p < 0.01$).

As we did not have an independent test set, we used twofold Monte Carlo cross validation to obtain unbiased estimates for the AUC. That is, we randomly assigned each of the observations in the case and control groups to one of two partitions of equal size. Subsequently, we used one of the partitions as the training set to build the LSIR and SIR predictors and evaluated the performance of our predictor on the remaining partition, labeled as the test set. We repeated this procedure 1000 times. For each of the test

sets, we computed the AUC and then averaged it. We estimated the dimension of both SIR and LSIR subspaces to be 1. The mean AUC value for the SIR predictor was 0.60 (0.04), whereas it was 0.63 (0.04) for the LSIR score. The 95% CI around the difference between the AUC values for SIR and LSIR was $(-0.04, -0.02)$, indicating that LSIR had better discriminatory performance than SIR. The difference in AUC values for this example agrees with what we saw in the simulations for a small number of predictors.

In summary, LSIR and SIR provided similar results. On the basis of LSIR, we identified glucose as a significant predictor and found that the fourth time point was important. This is plausible, as measurements taken at the fourth time point were most proximal to cancer diagnosis. SIR also found that the coefficient for glucose measured at times 1 and 4 was significant. The importance of glucose at the first time point is somewhat more difficult to explain, and could be a chance finding.

7. Discussion

In this paper, we show that under the assumption that the first two moments of repeatedly measured predictors have a Kronecker product structure, which reflects the longitudinal structure of the predictors, the first-moment-based dimension reduction subspace also has a Kronecker product structure. We propose an algorithm that utilizes this structure to implement a longitudinal version of SIR, which we call LSIR. This substantially reduces the complexity of the estimation of the first-moment-based dimension reduction subspace. The reduction in estimation burden leads to a noticeable improvement in the discriminatory ability of a score obtained by projecting the data onto that subspace for practically relevant sample sizes. For binary outcomes, the AUC for LSIR was approximately 5% higher than the AUC for a score computed from standard SIR, which ignored the repeated measurement structure. This seemingly small difference constitutes a substantial improvement in discriminatory power [21]. To see this, consider the simple case of a single continuous predictor that in both cases and controls arises from a normal population that differ only in their means; that is, $X_i \sim N(\mu_i, \sigma^2)$, $i = 0, 1$. The corresponding AUC is given by $AUC = \Phi((\mu_1 - \mu_0)/\sigma^2)$. The odds ratio (OR) associated with X in a logistic model based on a case-control sample is given by $\exp((\mu_1 - \mu_0)/\sigma^2)$. When $v = (\mu_1 - \mu_0)/\sigma^2 = 0.55$, we obtain $AUC = 0.71$. To improve the AUC by 5%, we need to increase v to $v = 0.70$, or equivalently, we need to increase the OR associated with X from 1.73 to 2.01, corresponding to an 18% increase in the OR.

An important step before any data analysis is to test the assumption about the Kronecker structure of the moments. Under the additional assumption of multivariate normally distributed markers, we can employ likelihood ratio tests following [16]. Nevertheless, we found in simulations that even under violations of the Kronecker product structure, fitting a more parsimonious covariance matrix did not result in a noticeable or even statistically significant loss of discriminatory power.

Moment structures similar to the ones that we assumed in this paper appeared for the first time in growth curve models [22] that are used to analyze longitudinal and repeated measures data. Roy and Khattree [13] and Srivastava *et al.* [14], for example, used these models in discrimination and classification of multivariate repeated measures data. Although all these approaches study variations of such structures on the mean and covariance of $\mathbf{X}|Y$, they also require normally distributed data. In contrast, we do not require normality or a specific distribution, and our approach is computationally tractable even for large numbers of markers. In addition, the LSIR algorithm can easily be adapted to handle missing data and a variable number of observations for each individual.

One of the reviewers of this paper alerted us to the fact that a very similar approach to our method, currently unpublished, was independently proposed as part of a PhD thesis [23]. In particular, Theorem 1 is a common result in our paper and in [23] (see Chapter 8), but our other methodological results, such as Theorem 2 and testing for dimension and marker contribution when the moments have a Kronecker product structure, are new and appear only in our paper, to the best of our knowledge.

Li and Yin [24], who applied the dimension reduction method of Li *et al.* [25] to settings where both the outcome and the predictors are measured repeatedly over time, first introduced SDR into the analysis of longitudinal data. Time is considered as a categorical covariate, dimension reduction subspaces are estimated separately for each time point, and then combined into a single matrix. However, this approach ignores the correlation across time, an integral feature of longitudinal data.

In summary, we present a new SDR approach for longitudinally measured predictors. By exploiting the structure of moments of the longitudinal predictors, we obtain more accurate estimates of the dimension reduction subspace and hence more accurate marker projections for classification. A limitation of the proposed work is that first-moment methods for binary outcomes can detect at most one dimension

in the central dimension reduction subspace and thus may miss important information of the predictors. In future work, we plan to extend this approach to second-moment-based methods.

Appendix A.

Proof of (3)

By the linearity condition using (2), we have

$$\Sigma^{-1}(\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})) = \eta \mathbf{v}_y \tag{19}$$

with $\eta \in \mathcal{S}_{Y|X}$ and $\text{cov}(\mathbf{v}_y) > 0$ and $\eta^T \eta$ the identity matrix. Then, $\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X}) = \Sigma \eta \mathbf{v}_y$, $\text{span}(\eta) = \Sigma^{-1} \text{span}((\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})))$ and

$$\Sigma = \mathbb{E}(\text{cov}(\mathbf{X}|Y)) + \text{cov}(\mathbb{E}(\mathbf{X}|Y)) = \Delta + \text{cov}(\mathbb{E}(\mathbf{X}|Y)) = \Delta + \Sigma \eta \text{cov}(\mathbf{v}_y) \eta^T \Sigma.$$

With $(\mathbf{A} + \mathbf{BCD})^{-1} = \mathbf{A}^{-1} - \mathbf{A}^{-1} \mathbf{B}(\mathbf{C}^{-1} + \mathbf{DA}^{-1} \mathbf{B})^{-1} \mathbf{DA}^{-1}$ [26], we obtain

$$\Delta^{-1} = \Sigma^{-1} + \eta (\text{cov}(\mathbf{v}_y)^{-1} + \eta^T \Sigma^{-1} \eta)^{-1} \eta^T,$$

and, therefore, from (19),

$$\begin{aligned} \Delta^{-1} \text{span}(\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})) &= \Sigma^{-1} \text{span}(\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})) + \text{span}(\eta (\text{cov}(\mathbf{v}_y)^{-1} + \eta^T \Sigma^{-1} \eta)^{-1} \mathbf{v}_y) \\ &= \Sigma^{-1} \text{span}(\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})) \end{aligned}$$

because $\text{span}(\eta) = \Sigma^{-1} \text{span}((\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})))$. □

Proof of Theorem 1

From the definition of the FMSSDR subspace, we have

$$\begin{aligned} \mathcal{S}_{\text{FMSSDR}} &= \Sigma^{-1} \text{span}(\text{vec}(\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X}))) = (\Sigma_1^{-1} \otimes \Sigma_2^{-1}) \text{span}(\text{cov}(\text{vec}(\mathbb{E}(\mathbf{X}|Y)))) \\ &= (\Sigma_1^{-1} \otimes \Sigma_2^{-1}) \text{span}(\alpha \otimes \beta) \text{cov}(\mathbf{v}_y) (\alpha^T \otimes \beta^T) = (\Sigma_1^{-1} \otimes \Sigma_2^{-1}) \text{span}(\alpha \otimes \beta) \\ &= \text{span}(\Sigma_1^{-1} \alpha \otimes \Sigma_2^{-1} \beta). \end{aligned}$$

The second equality follows from Proposition 2.6 in [9, p. 75]. □

Asymptotic distribution of $\widehat{\Omega}_{\text{LSIR}}$

First, notice that if $\text{vec}(\mathbb{E}(\mathbf{X}|Y) - \mathbb{E}(\mathbf{X})) = (\alpha \otimes \beta) \text{vec}(\mathbf{v}_y)$, the same moment condition holds for the standardized predictors \mathbf{Z} ,

$$\begin{aligned} \text{vec}(\mathbb{E}(\mathbf{Z}|Y)) &= \Sigma^{-1/2} (\alpha \otimes \beta) \text{vec}(\mathbf{v}_y) = (\Sigma_1 \otimes \Sigma_2)^{-1/2} (\alpha \otimes \beta) \text{vec}(\mathbf{v}_y) \\ &= [(\Sigma_1^{-1/2} \alpha) \otimes (\Sigma_2^{-1/2} \beta)] \text{vec}(\mathbf{v}_y) = (\alpha_z \otimes \beta_z) \text{vec}(\mathbf{v}_y), \end{aligned}$$

where $\alpha_z = \Sigma_1^{-1/2} \alpha$ and $\beta_z = \Sigma_2^{-1/2} \beta$. We derive the asymptotic distribution of $\widehat{\Omega}_{1t}$ given by (14). The proof for $\widehat{\Omega}_{2t}$ is analogous. Recall that $\bar{\mathbf{z}}_t^{(h)}$ denotes the p -vector of the standardized predictor means within slice h , $h = 1, \dots, H$. The proportion of observations in slice h at time t across all p predictors is $f_t^{(h)}$. Let $\hat{\mathbf{Z}}_t = (\bar{\mathbf{z}}_t^{(1)} \sqrt{f_t^{(1)}}, \dots, \bar{\mathbf{z}}_t^{(H)} \sqrt{f_t^{(H)}})$ to obtain $\widehat{\Omega}_{1t} = \hat{\mathbf{Z}}_t \hat{\mathbf{Z}}_t^T$. Because $\widehat{\Omega}_{1t}$ is a sample covariance matrix, for each $t = 1, \dots, T$,

$$n^{1/2} \text{vec}(\widehat{\Omega}_{1t} - \Omega_{1t}) \longrightarrow N_{pd}(\mathbf{0}, \mathbf{V}_{1t}) \tag{20}$$

where $\Omega_{1t} = \text{cov}(\mathbb{E}(\mathbf{Z}_t|Y)) = (\alpha_z \otimes \beta_{z,t}) \text{cov}(\mathbf{v}_y) (\alpha_z \otimes \beta_{z,t})^T$ with $\beta_{z,t}$ the t th column of the matrix β_z , and $\mathbf{V}_{1t} = \sum_h p_t^{(h)} \mathbf{V}_{1t}^{(h)}$ is the $p^2 \times p^2$ matrix with $\mathbf{V}_{1t}^{(h)} = \text{var}[\text{vec}(\mathbf{Z}_t^{(h)} - \mathbb{E}(\mathbf{Z}_t^{(h)}|Y = h))(\mathbf{Z}_t^{(h)} - \mathbb{E}(\mathbf{Z}_t^{(h)}|Y = h))^T | Y = h]$ and $p_t^{(h)} = P(Y \text{ falls in slice } h)$.

Because $\text{cov}(\mathbf{v}_y)$ is positive definite, $\text{cov}(\mathbf{v}_y) = \mathbf{U}\mathbf{U}^T$. Thus,

$$\begin{aligned}\boldsymbol{\Omega}_1 &= \frac{1}{T} \sum_{t=1}^T \boldsymbol{\Omega}_{1t} = \frac{1}{T} \sum_{t=1}^T (\boldsymbol{\alpha}_z \otimes \boldsymbol{\beta}_{z,t}) \mathbf{U}\mathbf{U}^T (\boldsymbol{\alpha}_z \otimes \boldsymbol{\beta}_{z,t})^T \\ &= \frac{1}{T} \sum_{t=1}^T (\boldsymbol{\alpha}_z \mathbf{U} \boldsymbol{\beta}_{z,t}^T \boldsymbol{\beta}_{z,t} \mathbf{U}^T \boldsymbol{\alpha}_z^T) = \boldsymbol{\alpha}_z \mathbf{U} \left(\frac{1}{T} \sum_{t=1}^T \boldsymbol{\beta}_{z,t}^T \boldsymbol{\beta}_{z,t} \right) \mathbf{U}^T \boldsymbol{\alpha}_z^T.\end{aligned}$$

With the eigendecomposition of the $s \times s$ matrix $\sum_{t=1}^T \boldsymbol{\beta}_{z,t}^T \boldsymbol{\beta}_{z,t} / T = \mathbf{V}\mathbf{D}\mathbf{V}^T$, $\boldsymbol{\Omega}_1 = \boldsymbol{\alpha}_z \mathbf{U}\mathbf{V}\mathbf{D}\mathbf{V}^T \mathbf{U}^T \boldsymbol{\alpha}_z^T = (\boldsymbol{\alpha}_z \otimes (\mathbf{V}\mathbf{D}^{1/2})) \text{cov}(\mathbf{v}_y) (\boldsymbol{\alpha}_z \otimes (\mathbf{V}\mathbf{D}^{1/2}))^T$. Therefore, from (20), we obtain

$$n^{1/2} \text{vec} \left(\frac{1}{T} \sum_{t=1}^T \widehat{\boldsymbol{\Omega}}_{1t} - \boldsymbol{\Omega}_1 \right) \longrightarrow N_{pd}(\mathbf{0}, \mathbf{V}_1)$$

with $\mathbf{V}_1 = \frac{1}{T^2} \sum_{t=1}^T \mathbf{V}_{1t} + \frac{1}{T^2} \sum_t \sum_{s,s \neq t} \text{cov}(\text{vec}(\widehat{\boldsymbol{\Omega}}_{1t}), \text{vec}(\widehat{\boldsymbol{\Omega}}_{1s}))$. Now,

$$\begin{aligned}\text{cov}(\text{vec}(\widehat{\boldsymbol{\Omega}}_{1t}), \text{vec}(\widehat{\boldsymbol{\Omega}}_{1s})) &= \text{cov}(\text{vec}(\sum_{h=1}^H f_t^{(h)} \bar{\mathbf{z}}_t^{(h)} \bar{\mathbf{z}}_t^{(h)T}), \text{vec}(\sum_{h=1}^H f_s^{(h)} \bar{\mathbf{z}}_s^{(h)} \bar{\mathbf{z}}_s^{(h)T})) \\ &= \sum_{h=1}^H f_t^{(h)} f_s^{(h)} \text{cov}(\text{vec}(\bar{\mathbf{z}}_t^{(h)} \bar{\mathbf{z}}_t^{(h)T}), \text{vec}(\bar{\mathbf{z}}_s^{(h)} \bar{\mathbf{z}}_s^{(h)T}))\end{aligned}$$

For $t = 1, \dots, T$, let $\tilde{\mathbf{Z}}_t^{(h)} = \bar{\mathbf{z}}_t^{(h)} \bar{\mathbf{z}}_t^{(h)T}$, with elements $[\tilde{\mathbf{Z}}_t^{(h)}]_{kl} = (f_t^{(h)})^2 \sum_{i=1}^{n_t} Z_{ikt}^{(h)} \sum_{i=1}^{n_t} Z_{ilt}^{(h)}$. The p^2 diagonal elements of $\text{cov}(\text{vec}(\tilde{\mathbf{Z}}_t^{(h)}), \text{vec}(\tilde{\mathbf{Z}}_s^{(h)}))$ are given by $\text{var}([\tilde{\mathbf{Z}}_t^{(h)}]_{kl})$ and $\text{var}([\tilde{\mathbf{Z}}_s^{(h)}]_{qr})$, and the off diagonal elements are $\text{cov}(\text{vec}([\tilde{\mathbf{Z}}_t^{(h)}]_{kl}), \text{vec}([\tilde{\mathbf{Z}}_s^{(h)}]_{qr}))$, where $k, l, q, r = 1, \dots, p$.

We can derive the asymptotic behavior of $\widehat{\boldsymbol{\Omega}}_2$ in a similar manner, by using \mathbf{Z}^T instead of \mathbf{Z} and $\text{vec}(\mathbf{E}(\mathbf{Z}^T | Y)) = (\boldsymbol{\beta}_z \otimes \boldsymbol{\alpha}_z) \text{vec}(\mathbf{v}_y^T)$ in the aforementioned calculations, leading to

$$n^{1/2} \text{vec} \left(\frac{1}{p} \sum_{j=1}^p \widehat{\boldsymbol{\Omega}}_{2j} - \boldsymbol{\Omega}_2 \right) \longrightarrow N_{pd}(\mathbf{0}, \mathbf{V}_2),$$

where $\boldsymbol{\Omega}_2 = \boldsymbol{\beta}_z \mathbf{U} \left(\frac{1}{p} \sum_{j=1}^p \boldsymbol{\alpha}_{z,j}^T \boldsymbol{\alpha}_{z,j} \right) \mathbf{U}^T \boldsymbol{\beta}_z^T$. From the expressions for $\boldsymbol{\Omega}_1$ and $\boldsymbol{\Omega}_2$, it can be seen that $\text{span}(\boldsymbol{\Omega}_1 \otimes \boldsymbol{\Omega}_2) = \text{span}(\boldsymbol{\alpha}_z \otimes \boldsymbol{\beta}_z)$.

Acknowledgements

We thank the associate editor and the reviewers for their helpful comments and suggestions.

References

1. Pfeiffer RM, Bura E. A model free approach to combining biomarkers. *Biometrical Journal* 2008; **50**:558–570.
2. Li KC. Sliced inverse regression for dimension reduction (with discussion). *Journal of the American Statistical Association* 1991; **86**:316–342.
3. Cook RD, Weisberg S. Discussion of “Sliced inverse regression for dimension reduction”. *Journal of the American Statistical Association* 1991; **86**:328–332.
4. Cook RD. *Regression Graphics: Ideas for Studying Regressions through Graphics*. Wiley: New York, 1998.
5. Cook RD. Fisher lecture. *Dimension reduction in regression statistical science* 2007; **22**:1–26.
6. Cook RD, Forzani L. Principal fitted components for dimension reduction in regression. *Statistical Science* 2008; **23**:485–501.
7. Bura E, Cook RD. Estimating the structural dimension of regressions via parametric inverse regression. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 2001; **63**:393–410.
8. Li KC, Duan N. Regression analysis under link violations. *Annals of Statistics* 1989; **17**:1009–1052.
9. Eaton ML. *Multivariate Statistics: A Vector Space Approach*. John Wiley & Sons, Inc.: New York, 1983.
10. Cook RD, Nachtsheim CJ. Reweighting to achieve elliptically contoured covariates in regression. *Journal of the American Statistical Association* 1994; **89**:592–599.

11. Hall P, Li KC. On almost linearity of low-dimensional projections from high-dimensional data. *Annals of Statistics* 1993; **21**:867–889.
12. Shao Y, Cook RD, Weisberg S. Marginal tests with sliced average variance estimation. *Biometrika* 2007; **94**:285–296.
13. Roy A, Khattree R. Discrimination and classification with repeated measures data under different covariance structures. *Communications in Statistics—Simulation and Computation* 2005; **34**:167–178.
14. Srivastava MS, Von Rosen T, Von Rosen D. Estimation and testing in general multivariate linear models with Kronecker product covariance structure. *Sankhya, Series A* 2009; **71**:137–163.
15. Roy A, Khattree R. Tests for mean and covariance structures relevant in repeated measures based discriminant analysis. *Journal of Applied Statistical Science* 2003; **12**:91–104.
16. Roy A, Khattree R. On implementation of a test for Kronecker product covariance structure for multivariate repeated measures data in. *Statistical Methodology* 2005; **2**:297–306.
17. Bura E, Yang J. Dimension estimation in sufficient dimension reduction: a unifying approach. *Journal of Multivariate Analysis* 2011; **102**:130–142.
18. Bura E, Pfeiffer RM. On the distribution of the left singular vectors of a random matrix and its applications. *Statistics & Probability Letters* 2008; **15**:2275–2280.
19. Pepe MS. *The Statistical Evaluation of Medical Tests for Classification and Prediction*, Oxford Statistical Science Series. Oxford University Press: New York, 2003.
20. Ulmer H, Kelleher C, Diem G, Concin H. Long-term tracking of cardiovascular risk factors among men and women in a large population-based health system: the Vorarlberg Health Monitoring & Promotion Programme. *European Heart Journal* 2003; **24**:1004–1013.
21. Pepe MS, Janes H, Longton G, Leisenring W, Newcomb P. Limitations of the odds ratio in gauging the performance of a diagnostic, prognostic, or screening marker. *American Journal of Epidemiology* 2004; **159**:882–890.
22. Potthoff RF, Roy SN. A generalized multivariate analysis of variance model useful especially for growth curve problems. *Biometrika* 1964; **51**:313–326.
23. Kim MK. On dimension folding of matrix or array valued statistical objects. *PhD dissertation*, Pennsylvania State University, 2010.
24. Li LX, Yin XR. Longitudinal data analysis using sufficient dimension reduction method. *Computational Statistics & Data Analysis* 2009; **53**:4106–4115.
25. Li B, Cook RD, Chiaromonte F. Dimension reduction for the conditional mean in regressions with categorical predictors. *Annals of Statistics* 2003; **31**:1636–1668.
26. Henderson HV, Searle SR. On deriving the inverse of a sum of matrices. *SIAM Review* 1981; **23**:53–60.