



Dynamic Principal Component Analysis in High Dimensions

Xiaoyu Hu & Fang Yao

To cite this article: Xiaoyu Hu & Fang Yao (2024) Dynamic Principal Component Analysis in High Dimensions, Journal of the American Statistical Association, 119:545, 308-319, DOI: 10.1080/01621459.2022.2115917

To link to this article: <https://doi.org/10.1080/01621459.2022.2115917>

 View supplementary material [↗](#)

 Published online: 10 Oct 2022.

 Submit your article to this journal [↗](#)

 Article views: 1544

 View related articles [↗](#)

 View Crossmark data [↗](#)

 Citing articles: 4 View citing articles [↗](#)



Dynamic Principal Component Analysis in High Dimensions

Xiaoyu Hu^a and Fang Yao^b

^aDepartment of Statistics & Data Science, National University of Singapore, Singapore; ^bDepartment of Probability & Statistics, School of Mathematical Sciences, Center for Statistical Science, Peking University, Beijing, China

ABSTRACT

Principal component analysis is a versatile tool to reduce dimensionality which has wide applications in statistics and machine learning. It is particularly useful for modeling data in high-dimensional scenarios where the number of variables p is comparable to, or much larger than the sample size n . Despite an extensive literature on this topic, researchers have focused on modeling static principal eigenvectors, which are not suitable for stochastic processes that are dynamic in nature. To characterize the change in the entire course of high-dimensional data collection, we propose a unified framework to directly estimate dynamic eigenvectors of covariance matrices. Specifically, we formulate an optimization problem by combining the local linear smoothing and regularization penalty together with the orthogonality constraint, which can be effectively solved by manifold optimization algorithms. We show that our method is suitable for high-dimensional data observed under both common and irregular designs, and theoretical properties of the estimators are investigated under l_q ($0 \leq q \leq 1$) sparsity. Extensive experiments demonstrate the effectiveness of the proposed method in both simulated and real data examples.

ARTICLE HISTORY

Received June 2021
Accepted August 2022

KEYWORDS

Dimension reduction; Local linear smoothing; Manifold optimization; Sparsity

1. Introduction

Principal component analysis (PCA) has been widely used to reduce dimensionality and extract useful features by transforming the original variables into a few new uncorrelated variables while retaining most information in the data (Anderson 1963). It is an important tool in various applications, such as data compression and reconstruction (Sirovich and Kirby 1987; Turk and Pentland 1991). Despite its importance, existing works in this field mainly focus on modeling a static decomposition, where principal eigenvectors are invariant with respect to time. However, technological advances enable data collection in dynamic environments that often vary with time or other index variables. Such data are expected to possess dynamic structures with eigenvectors/eigenspaces varying with time, which makes existing methods less applicable. To tackle this issue, we aim to dynamically estimate leading eigenvectors of covariance matrices to capture the time-varying information, which is referred to as the dynamic PCA (DPCA). It has wide applications in signal processing, for example, subspace tracking (Delmas 2010).

A straightforward way to conduct the DPCA is to perform PCA on the sample covariance matrix at each observed grid point (Berrendero, Justel, and Svarc 2011). However, this has several limitations. First, one cannot obtain a smooth estimate over the whole time period in an integrative manner by their method. Second, under the irregularly/sparsely observed case for each subject, one cannot directly calculate the sample covariance matrix, which makes their method inapplicable.

More importantly, PCA is known to behave poorly in high-dimensional settings, where the number of variables p is comparable to or much larger than the sample size n (Johnstone and Lu 2009).

In the high-dimensional static case, various sparse PCA methods are developed and studied in the literature (Jolliffe, Trendafilov, and Uddin 2003; Zou, Hastie, and Tibshirani 2006; Shen and Huang 2008; Amini and Wainwright 2009; Witten, Tibshirani, and Hastie 2009; Berthet and Rigollet 2013; Brennan and Bresler 2019). For spiked covariance models, Johnstone and Lu (2009) proposed the diagonal thresholding algorithm by retaining variables with large sample variances, and Ma (2013), Deshpande and Montanari (2014), and Krauthgamer, Nadler, and Vilenchik (2015) further refined this estimation. Moreover, Ding et al. (2019) and Holtzman, Soffer, and Vilenchik (2020) proposed different algorithms for the sparse PCA problem in the spiked covariance model. In more general settings, Vu and Lei (2013) gave the nonasymptotic lower and upper bounds on the minimax subspace estimation error. In addition, Vu et al. (2013) considered a convex relaxation strategy based on the convex hull of low rank projection matrices. Yet, if one adopts the method in Berrendero, Justel, and Svarc (2011), the existing sparse PCA methods in the static case are not readily applicable for modeling time-varying principal eigenvectors. For instance, the irregular sampling scheme and the dependence among measurements from the same subject pose new challenges to methodological and theoretical developments.

Another possible way to obtain time-varying eigenvectors is through the eigendecomposition of estimated dynamic covariance matrices, and dynamic covariance models have been explored in the literature. In low-dimensional settings, the nonparametric or semiparametric estimators are constructed (Zhu et al. 2009; Yin et al. 2010; Yuan et al. 2012). For high-dimensional data, Chen and Leng (2016) proposed a sparse estimate using the kernel smoothing and thresholding. However, it is unclear about the quality of eigenvector estimates based on performing standard PCA on the sparse covariance matrices. Our numerical studies reveal that the principal eigenvectors deduced from dynamic covariance estimation in Chen and Leng (2016) perform sub-optimally, especially when the observational grids are sparse and the dimension is large. Moreover, these existing works about dynamic covariance matrices did not account for the dependence among observations from the same process which is an important nature of repeated measurements (Cai and Yuan 2011).

In this work, we propose a unified framework with theoretical guarantees for the DPCA. Specifically, to resolve the problems caused by the sample covariance matrix and the irregular sampling scheme, we use the local linear smoothing (Fan and Gijbels 1996) for estimation. To deal with high dimensionality, we restrict our attention to the eigenvectors with sparsity structures. Consequently, we formulate an optimization problem which combines the local linear smoothing and sparse regularization. The proposed method has some remarkable features. First, it is applicable to high-dimensional data under both common and irregular/sparse designs (Cai and Yuan 2011). Second, instead of adopting the convex relaxation strategy which is computationally expensive with the computational cost $O(p^3)$ per iteration (Vu et al. 2013), our optimization problem is directly defined on the Stiefel manifold, which can be solved by leveraging manifold optimization algorithms, for example, the proximal gradient method in Chen et al. (2020) operating with $O(p^2d)$, where d is the number of principal eigenvectors of interest. Third, our procedure consists of two steps: the first step generates an initial estimate from the optimization, and the second step refines the estimate by hard thresholding and re-optimization on the reduced set of variables. This two-step algorithm helps successfully identify significant features and enhance the interpretability, which leads to consistent estimators under the l_q ($0 \leq q \leq 1$) sparsity. Moreover, we show that the convergence rate of resulting estimators exhibits a phase transition phenomenon that attains either nonparametric or parametric rate, depending on the sampling frequency, that is, how sparse/dense the repeated measurements are observed, see Section 3. It is noteworthy that the convergence of estimated principal eigenvectors is faster than that in dynamic covariance estimation, which coincides with findings in the static case (Vu and Lei 2012; Bickel and Levina 2008; Cai and Zhou 2012).

While both DPCA and functional PCA (FPCA) (Ramsay and Silverman 2005) are tools to model random functions, they are essentially different frameworks, see Remark 1. Although FPCA is widely used to represent a single or a small number of functional processes, its performance is not guaranteed and can be unreliable in high dimensions due to error accumulation (Yao, Müller, and Wang 2005a,b; Chiou, Chen, and Yang 2014). In this regard, the DPCA is preferred to capture the dynamic

information with low-dimensional structures in applications such as data compression and reconstruction. This is illustrated in the real data example in terms of recovery errors in Section 5.

The remainder of the article is organized as follows. In Section 2, we first introduce the dynamic PCA, then we provide the l_q ($0 \leq q \leq 1$) sparsity assumption in dynamic settings and the formulation of our optimization problem, and describe procedures for practical implementation. In Section 3, we present theoretical results under suitable regularity conditions. Simulation results are included in Section 4, followed by an application to the heartbeat sound data in Section 5. The additional results and technical proofs are deferred to the Appendix and Supplementary Material.

2. Dynamic Principal Component Analysis with Sparsity

2.1. Dynamic Principal Component Analysis

We begin with some notations used in the sequel. For a matrix $A = (a_{ij})_{i,j=1}^p \in \mathbb{R}^{p \times p}$, $\text{vec}(A)$ denotes the vector with length p^2 obtained by stacking the columns of A . We define the Frobenius norm $\|A\|_F = (\sum_{i,j} a_{ij}^2)^{1/2}$, the elementwise l_∞ norm $\|A\|_\infty = \max_{1 \leq i,j \leq p} |a_{ij}|$ and the elementwise l_1 norm $\|A\|_1 = \sum_{i,j} |a_{ij}|$. For a vector $\mathbf{u} \in \mathbb{R}^p$, denote its l_q norm by $\|\mathbf{u}\|_q = (\sum_{j=1}^p |u_j|^q)^{1/q}$ with $\|\mathbf{u}\|_0$ defined as the number of nonzero elements. For two real numbers a and b , define $a \wedge b = \min(a, b)$ and $a \vee b = \max(a, b)$. We write $a \lesssim b$ if $a \leq Cb$ for some constant $C > 0$.

Let $\{\mathbf{X}(t) : t \in \mathcal{T}\}$ be a vector-valued stochastic process defined on a compact interval $\mathcal{T} = [0, 1]$, where $\mathbf{X}(t) = (X_1(t), \dots, X_p(t))^T$. The mean and diagonal covariance functions are assumed to be continuous and denoted by $\boldsymbol{\mu}(t) = (\mu_1(t), \dots, \mu_p(t))^T = E\mathbf{X}(t)$ and $\Sigma(t) = (\sigma_{jk}(t))_{j,k=1}^p = E\mathbf{X}(t)\mathbf{X}(t)^T - \boldsymbol{\mu}(t)\boldsymbol{\mu}(t)^T$, respectively. For each fixed t , applying multivariate PCA, we obtain

$$\mathbf{X}(t) = \boldsymbol{\mu}(t) + \sum_{k=1}^p \xi_k(t)\mathbf{u}_k(t), \tag{1}$$

where $\mathbf{u}_k(t)$ is the k th principal eigenvector and $\xi_k(t) = (\mathbf{X}(t) - \boldsymbol{\mu}(t))^T \mathbf{u}_k(t)$ is the k th principal component score with $E\xi_k(t) = 0$ and $\text{var}(\xi_k(t)) = \lambda_k(t)$. Without loss of generality, suppose that $\lambda_1(t) \geq \lambda_2(t) \geq \dots \geq \lambda_p(t) \geq 0$. Moreover, we have $\text{cov}(\xi_k(t), \xi_l(t)) = 0$ and $\mathbf{u}_k(t)^T \mathbf{u}_l(t) = 0$ for $k \neq l$ at each t . The time-varying version of PCA (1) is called dynamic PCA (DPCA).

Remark 1. As discussed in Section 1, the DPCA is essentially different from FPCA. In particular, the DPCA represents data in the Euclidean space, that is, it applies multivariate PCA at each t to obtain $\mathbf{X}(t) = \boldsymbol{\mu}(t) + \sum_{k=1}^p \xi_k(t)\mathbf{u}_k(t)$, where $\mathbf{u}_k(t), k = 1, \dots, p$ form an orthonormal basis in \mathbb{R}^p . By comparison, the FPCA represents $\mathbf{X}(t)$ in the infinite-dimensional function space, that is, $\mathbf{X}(t) = \boldsymbol{\mu}(t) + \sum_{k=1}^\infty \theta_k \boldsymbol{\phi}_k(t)$ where θ_k are uncorrelated functional principal scores and $\boldsymbol{\phi}_k(t)$ are orthonormal basis functions in the space of square integrable

functions $L^2(\mathcal{T})$. Note that, instead of the auto-covariance function $C(s, t) = E\mathbf{X}(s)\mathbf{X}(t)^\top - \boldsymbol{\mu}(s)\boldsymbol{\mu}(t)^\top$, the DPCA studies the much simpler diagonal covariance function $\Sigma(t) = C(t, t)$.

One advantage of DPCA is the ability to capture the dynamic information contained in data, which facilitates interpretation. Formally, the dynamic principal eigenvectors can be found by solving the following optimization problem,

$$\begin{aligned} \min_{V(t)} \int_{\mathcal{T}} E\|\mathbf{X}(t) - \boldsymbol{\mu}(t) - V(t)V(t)^\top\{\mathbf{X}(t) - \boldsymbol{\mu}(t)\}\|^2 dt \\ \text{s.t.} \quad V(t)^\top V(t) = I_d, \end{aligned} \quad (2)$$

where $V(t) \in \mathbb{R}^{p \times d}$, I_d is a $d \times d$ identity matrix and d is the number of principal eigenvectors of interest. The problem (2) is reduced to perform multivariate PCA at each t based on Lemma 2 in the Appendix, and a similar result can be found in Berrendero, Justel, and Svarc (2011). Note that if $\lambda_d(t) - \lambda_{d+1}(t) > 0$, then $U(t) = (\mathbf{u}_1(t), \dots, \mathbf{u}_d(t))$ is unique up to an orthogonal matrix, that is, $U(t)O$ is also an optimal solution of (2) for any $d \times d$ orthogonal matrix O . We refer to the subspace $\mathcal{S}(t)$ spanned by the column vectors of $U(t)$ as the dynamic principal subspace, and the corresponding projection matrix is given by $\Pi(t) = U(t)U(t)^\top$.

In reality, we observe noisy measurements at common or irregular design points, $y_{ijl} = x_{ij}(t_{il}) + \epsilon_{ijl}$, $t_{il} \in \mathcal{T}$, where ϵ_{ijl} are independent and identically distributed (iid) measurement errors independent of x_{ij} with mean zero and variance σ^2 , $i = 1, \dots, n$; $j = 1, \dots, p$ and $l = 1, \dots, m_i$, where m_i is the number of observations for each trajectory of the i th subject. We denote $N = \sum_{i=1}^n m_i$ and $\bar{m} = \sum_{i=1}^n m_i/n$. Under the common design, all the observations are sampled at the same locations, that is, $t_{1l} = t_{2l} = \dots = t_{nl} = t_l$ for all $l = 1, \dots, m$ where $m = \bar{m} = m_1 = \dots = m_n$, while the locations t_{il} are sampled independently from a compact interval \mathcal{T} under the irregular design (Cai and Yuan 2011).

An empirical version of (2) is formulated by substituting the expectation with its estimate. A naive estimate is to use the sample covariance matrix S , which however has some drawbacks as discussed in Section 1. First, since the data are collected at discrete grids, one can merely obtain estimates at observed times instead of the whole period \mathcal{T} . Second, the sample covariance matrices are infeasible under the irregular design. Third, the estimates may fluctuate significantly without considering smoothness. Therefore, a reliable and smooth estimate is desirable. To illustrate the main idea, we assume $\boldsymbol{\mu}(t) = \mathbf{0}$ for the moment. To obtain the estimate at the target time t , we borrow the information of the data observed at neighboring grids. Thus, motivated by the local linear smoothing, we propose an empirical optimization problem as follows,

$$\begin{aligned} \min_{V(t)} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il}(t) \|\mathbf{y}_{il} - V(t)V(t)^\top \mathbf{y}_{il}\|^2 \\ \text{s.t.} \quad V(t)^\top V(t) = I_d, \end{aligned}$$

where $w_{il}(t) = \{R_2 K_h(t_{il} - t) - R_1 K_h(t_{il} - t)(t_{il} - t)\} / \{R_0 R_2 - R_1^2\}$, $R_\ell = \sum_{i=1}^n \sum_{l=1}^{m_i} K_h(t_{il} - t)(t_{il} - t)^\ell$, $\ell = 0, 1, 2$, h is the bandwidth, $K_h(\cdot) = K(\cdot/h)/h$ and K is a kernel function (Fan

and Gijbels 1996). It can be equivalently formulated as

$$\begin{aligned} \max_{V(t)} \text{Tr}[V(t)^\top \hat{\Sigma}(t)V(t)] \\ \text{s.t.} \quad V(t)^\top V(t) = I_d, \end{aligned} \quad (3)$$

where $\hat{\Sigma}(t) = \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il}(t) \mathbf{y}_{il} \mathbf{y}_{il}^\top$ is the smoothed covariance matrix.

Note that our proposal readily adapts to both common and irregular designs using pooled data. More generally, incorporating the estimated mean function by the kernel smoothing, the estimator $\hat{U}(t)$ can be obtained by substituting $\hat{\Sigma}(t)$ in (3) with

$$\hat{\Sigma}(t) = \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il}(t) \mathbf{y}_{il} \mathbf{y}_{il}^\top - \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il}(t) \mathbf{y}_{il} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il}(t) \mathbf{y}_{il}^\top.$$

In addition, under the common design where the data are observed at regular grids, practitioners can adopt an alternative estimate of $\hat{\Sigma}(t)$ for simplified computation,

$$\hat{\Sigma}_{\text{common}}(t) = \sum_{l=1}^m w_l(t) \sum_{i=1}^n n^{-1} (\mathbf{y}_{il} - \bar{y}_l)(\mathbf{y}_{il} - \bar{y}_l)^\top, \quad (4)$$

where $\bar{y}_l = \sum_{i=1}^n y_{il}/n$, $w_l(t) = \{R_{2,c} K_h(t_l - t) - R_{1,c} K_h(t_l - t)(t_l - t)\} / (R_{2,c} R_{0,c} - R_{1,c}^2)$, $R_{\ell,c} = \sum_{i=1}^m K_h(t_l - t)(t_l - t)^\ell$, $\ell = 0, 1, 2$.

2.2. Sparsity and Estimation in High Dimensions

For high-dimensional data, the number of variables p is comparable to or even much larger than the sample size n . The estimator $\hat{U}(t)$ in Section 2.1 may become inconsistent without additional structures. The sparsity assumption is necessary to enhance the interpretability and improve the estimates in high dimensions. Assume that $U(t) \in \mathcal{U}(q, R_q; \mathcal{T})$, where

$$\begin{aligned} \mathcal{U}(q, R_q; \mathcal{T}) \\ = \left\{ U(t) \in \mathbb{R}^{p \times d}, t \in \mathcal{T} \mid U(t) \in \mathbb{V}_{p,d}, \sup_{t \in \mathcal{T}} \max_{1 \leq j \leq d} \|\mathbf{u}_j(t)\|_q \leq R_q \right\}, \end{aligned}$$

with $0 < q \leq 1$. When $q = 0$,

$$\begin{aligned} \mathcal{U}(0, R_0; \mathcal{T}) \\ = \left\{ U(t) \in \mathbb{R}^{p \times d}, t \in \mathcal{T} \mid U(t) \in \mathbb{V}_{p,d}, \sup_{t \in \mathcal{T}} \max_{1 \leq j \leq d} \|\mathbf{u}_j(t)\|_0 \leq R_0 \right\}. \end{aligned}$$

The set $\mathcal{U}(q, R_q; \mathcal{T})$ is nonempty and the l_q constraint is active only when $1 \leq R_q \leq p^{1-q/2}$. In Section 3, we consider bounded R_q to simplify the theoretical exposition. The family of leading eigenvectors over \mathcal{T} defined in $\mathcal{U}(q, R_q; \mathcal{T})$ generalizes the notion of static eigenvectors in Vu and Lei (2012). We stress that if the sparsity condition does not hold uniformly over \mathcal{T} , our method can still apply to the subregions of \mathcal{T} where this condition holds.

Recall that the projection matrix is $\Pi(t) = U(t)U(t)^\top$. By definition, $\Pi_{jj}(t) = 0$ holds if and only if each element of the j th row of $U(t)$ is zero. Further, it implies that if $\Pi_{jj}(t) = 0$, then all entries of the j th row/column of $\Pi(t)$ are 0. Denote the

support set $J(t) = \{j : \Pi_{jj}(t) > 0\}$. For notational convenience, we introduce the following block representation of $\Sigma(t)$:

$$\begin{pmatrix} \Sigma_{JJ}(t) & \Sigma_{JJ^c}(t) \\ \Sigma_{J^cJ}(t) & \Sigma_{J^cJ^c}(t) \end{pmatrix}.$$

Similar block representations can be defined for other matrices and vectors. Apparently, the principal eigenvectors at t does not depend on the variables outside of the set $J(t)$ in the sense that all elements of $U_{J^c}(t)$ are 0.

Note that $U(t) \in \mathbb{V}_{p,d}$, where $\mathbb{V}_{p,d} = \{V \in \mathbb{R}^{p \times d} | V^T V = I_d\}$ is the Stiefel manifold, which results in a non-convex problem which is hard to solve. Most existing algorithms for static sparse PCA require deflation steps (Shen and Huang 2008; Mackey 2009) or convex relaxation (d’Aspremont et al. 2007; Vu et al. 2013) to circumvent the orthogonality constraint, which either lack theoretical guarantees or are computational expensive. To avoid these issues, our optimization problem is defined directly on the Stiefel manifold $\mathbb{V}_{p,d}$ which can be solved by manifold optimization algorithms. The regularized manifold optimization problem is formulated as follows,

$$\begin{aligned} \min_{V(t)} & -\text{Tr}\{V(t)^T \hat{\Sigma}(t) V(t)\} + \rho_t \|V(t)\|_1 \\ \text{s.t.} & V(t)^T V(t) = I_d, \end{aligned} \tag{5}$$

where $\rho_t > 0$ is the regularization parameter at t . We allow the parameter ρ to depend on t , which makes our proposal fully adaptive to different sparsity levels varying with t . The optimization problem (5) deals with sparsity and orthogonality jointly, which can be solved effectively with recent developments for manifold optimization, for example, the proximal gradient method (Chen et al. 2020).

To improve estimation, we treat the solution of (5) as an initial estimate which is denoted by $\hat{U}^0(t)$, and then propose a refined version. Specifically, we add a thresholding step to further filter out the variables irrelevant to the principal eigenvectors. Denote the set of remaining variables by $\hat{J}(t) = \{j : \hat{\Pi}_{jj}^0(t) \geq \gamma_t\}$, where $\hat{\Pi}^0(t) = \hat{U}^0(t) \hat{U}^0(t)^T$ and $\gamma_t > 0$ is the thresholding parameter at t . Since the estimate after thresholding may not belong to the Stiefel manifold $\mathbb{V}_{p,d}$, we reestimate the principal eigenvectors afterwards. Our refined estimate is given by

$$\hat{U}(t) = \begin{pmatrix} \hat{U}_{\hat{J}(t)}(t) \\ 0 \end{pmatrix},$$

where $\hat{U}_{\hat{J}(t)}(t)$ is the solution of the problem,

$$\begin{aligned} \min_{V(t)} & -\text{Tr}\{V(t)^T \hat{\Sigma}_{\hat{J}(t)\hat{J}(t)}(t) V(t)\} + \rho_t \|V(t)\|_{1,1} \\ \text{s.t.} & V(t)^T V(t) = I_d. \end{aligned} \tag{6}$$

Note that the estimated principal subspaces are readily obtained by spanning the columns of $\hat{U}(t)$ with projection matrices $\hat{\Pi}(t) = \hat{U}(t) \hat{U}(t)^T$. The two-step estimation procedure successfully identifies the significant variables and provides consistent estimators under general l_q sparsity, which is theoretically and empirically demonstrated in Sections 3 and 4.

2.3. Tuning Parameters

In this section, we discuss how to select parameters that are involved in the estimation procedure. Note that in the dynamic setting, the number of principal eigenvectors of interest d may be a constant or vary with t . There exists no consensus on the selection of d which depends on the specific application. For example, it could be selected based on the fraction of variance explained (FVE). In supervised problems such as regression or classification, it may be tuned by k -fold cross-validation to minimize the prediction/classification error. Here we mainly consider tuning three other parameters, the bandwidth h , the sparsity parameter ρ_t and the thresholding parameter γ_t . We suggest to select them in a sequential manner (Chen and Lei 2015; Chen and Leng 2016). For the bandwidth h , we use the leave-one-curve-out cross-validation approach (Rice and Silverman 1991; Yao, Müller, and Wang 2005a). Specifically, we tune the bandwidth h given $\rho_t = 0$ and $\gamma_t = 0$ by maximizing the cross-validated inner product,

$$h^* = \arg \max_{h \in \mathcal{A}_1} \frac{1}{nm} \sum_{i=1}^n \sum_{l=1}^{m_i} \text{Tr}\{\hat{U}_{h,0,0}^{-i}(t_{il})^T (\mathbf{y}_{il} - \hat{\boldsymbol{\mu}}(t_{il})) (\mathbf{y}_{il} - \hat{\boldsymbol{\mu}}(t_{il}))^T \hat{U}_{h,0,0}^{-i}(t_{il})\},$$

where $\hat{\boldsymbol{\mu}}$ is the estimated mean function which refers to the sample mean under the common case and the local linear estimate under the irregular design, \mathcal{A}_1 is a candidate set of h , $\hat{U}_{h,0,0}^{-i}$ is estimated by leaving out the i th subject with the bandwidth h , $\rho_t = 0$ and $\gamma_t = 0$. Next, the parameter ρ_t is determined by k -fold cross-validation. The data is divided into k -folds by subjects, denoted by $\mathcal{D}_1, \dots, \mathcal{D}_k$. Let $\hat{U}_{h,\rho_t,\gamma_t}^{-v}(t)$ be the estimator using data other than \mathcal{D}_v at time t with parameters h , ρ_t and γ_t . Let $\hat{\Sigma}_h^v(t)$ be the smoothed covariance matrix estimate at t using \mathcal{D}_v with the bandwidth h . Next, we choose ρ_t given the selected bandwidth h^* and $\gamma_t = 0$ by maximizing the cross-validated inner product,

$$\rho_t^* = \arg \max_{\rho_t \in \mathcal{A}_{2,t}} \frac{1}{k} \sum_{v=1}^k \text{Tr}\{[\hat{U}_{h^*,\rho_t,0}^{-v}(t)]^T \hat{\Sigma}_{h^*}^v(t) \hat{U}_{h^*,\rho_t,0}^{-v}(t)\},$$

where $\mathcal{A}_{2,t}$ is a candidate set for ρ_t . At last, we tune the thresholding parameter γ_t , given the selected bandwidth h^* and sparsity level ρ_t^* , by a trade-off between the explained variance $Ip(\gamma_t)$ and model complexity, that is, the number of retained variables, where

$$Ip(\gamma_t) = \frac{1}{k} \sum_{v=1}^k \text{Tr}\{[\hat{U}_{h^*,\rho_t^*,\gamma_t}^{-v}(t)]^T \hat{\Sigma}_{h^*}^v(t) \hat{U}_{h^*,\rho_t^*,\gamma_t}^{-v}(t)\},$$

where $\gamma_t \in \mathcal{A}_{3,t}$, $\mathcal{A}_{3,t}$ is a candidate set and $Ip(\gamma_t)$ is the cross-validated inner product when the threshold equals γ_t . The model complexity depicts the cardinality of the support set $\hat{J}(t)$. One can select the γ_t to achieve model parsimony without much information loss. We demonstrate the performance of selected parameters in Section 4.

3. Theoretical Results

In this section, we investigate the theoretical properties of the proposed estimator under both common and irregular designs.

To measure the performance of the estimator, we use the notion of the distance defined in Vu and Lei (2013). For $U, V \in \mathbb{V}_{p,d}$, the squared distance is defined by

$$d^2(U, V) = d^2(\mathcal{E}, \mathcal{F}) = \frac{1}{2} \|E - F\|_F^2, \quad (7)$$

where \mathcal{E} and \mathcal{F} are subspaces with projection matrices $E = UU^T$ and $F = VV^T$, respectively.

Some assumptions necessary for theoretical results are provided, concerning the properties of variables and kernel functions. **Assumption 1** ensures that the d -dimensional principal subspace is well-defined. In multivariate cases, the commonly used assumption for sparse PCA is that x_j^2 is sub-exponential, while **Assumption 2** is adapted to random processes, which holds rather generally, for example, Gaussian processes.

Assumption 1. Assume that $\lambda_d(t) - \lambda_{d+1}(t) > 0$ for all $t \in \mathcal{T}$.

Assumption 2. For each $j = 1, \dots, p$, $X_j^2(t)$ is sub-exponential uniformly in $t \in \mathcal{T}$, that is, there exists a positive constant λ_0 such that $\sup_{t \in \mathcal{T}} E e^{\lambda X_j^2(t)} < \infty$ for $|\lambda| < \lambda_0$. Also assume the measurement error ϵ^2 is sub-exponential.

Assumption 3. The mean functions $\mu_j(\cdot)$ and the diagonal covariance functions $\sigma_{jk}(t)$ are twice differentiable and the second derivative is bounded on \mathcal{T} for $j, k = 1, \dots, p$.

Assumption 4. Assume $\log p(n\bar{m}^{-1} + nh)^{-1} \rightarrow 0$ as $n \rightarrow \infty$.

Assumption 5. The kernel function $K(\cdot)$ is a bounded and symmetric probability density function on $[-1, 1]$ with $\int u^2 K(u) du < \infty$ and $\int K^2(u) du < \infty$.

The smoothness of mean and diagonal covariance functions is imposed in **Assumption 3**, while $\mathbf{X}(t)$ is not necessarily smooth. **Assumption 4** indicates $\log p = O(n^c)$ for some $c > 0$ since \bar{m} and h^{-1} typically grow at a fractional polynomial order of n . **Assumption 5** is standard in the kernel smoothing literature (Fan and Gijbels 1996; Chen and Leng 2016).

First, we quantify the performance of the thresholding step by investigating the false positive control and false negative control of $J(t)$. It is revealed in **Lemma 1** that, with a suitable parameter γ_t , we can recover the support set consistently. The condition $\min_{j \in J(t)} \Pi_{jj}(t) \geq 2\gamma_t$ assures that the important variables can be distinguished from the noise stochastically. Denote $\Gamma(t) = \Sigma(t) + \sigma^2 I_p$ where $I_p \in \mathbb{R}^{p \times p}$ is an identity matrix.

Lemma 1. Assume $U(t) \in \mathcal{U}(q, R_q; \mathcal{T})$ and recall that $\hat{\Pi}^0(t)$ is the initial estimator. Note that we have $\|\Pi(t) - \hat{\Pi}^0(t)\|_F^2 \leq C \|\hat{\Sigma}(t) - \Gamma(t)\|_\infty = o_p(1)$, $t \in \mathcal{T}$ for some positive constant $C > 0$. If $\min_{j \in J(t)} \Pi_{jj}(t) \geq 2\gamma_t$ and $\gamma_t > \|\Pi(t) - \hat{\Pi}^0(t)\|_F$, then the variable selection procedure $\hat{J}(t) := \{j : \hat{\Pi}_{jj}^0(t) \geq \gamma_t\}$ succeeds.

In the following, we state the theoretical properties of the eventually obtained estimators $\hat{U}(t)$. A theoretical challenge is how to carefully control the l_q ($0 \leq q \leq 1$) norm of our estimators obtained with a lasso-type penalty, which can be tackled by

the consistent variable selection. Moreover, we need to deal with the dependence between observations from the same trajectory with care to control the concentration bound of the local linear estimator, which is essential to the theoretical results. Next, we investigate the behavior of the resulting estimator under both irregular and common designs.

3.1. Rate of Convergence Under the Irregular Design

In this section, we provide a theoretical investigation of estimators under the irregular design. **Assumption 6** is about the sampling scheme under the irregular design (Cai and Yuan 2011). **Assumption 7** on the sampling frequency is assumed to quantify the within-subject dependence and facilitate the exposition of theoretical analysis. This is a standard condition used in functional data (Zhang and Wang 2016), which is fairly mild and holds for the common design and the irregular design with finite m_i or when m_i are not all vastly different, $1 \leq i \leq n$.

Assumption 6. Under the irregular design, $t_{i\ell}$, $i = 1, \dots, n$; $\ell = 1, \dots, m_i$ are independent and identically distributed from a density $f(\cdot)$ with compact support \mathcal{T} . In addition, the sampling density $f_{\mathcal{T}}$ is bounded away from zero and infinity and is twice continuously differentiable with a bound derivative on its support.

Assumption 7. Assume $\limsup_n \sum_{i=1}^n m_i^2 / n\bar{m}^2 < \infty$ and $\sup_n (n \max_i m_i / \sum_{i=1}^n m_i) < \infty$.

Theorem 1. Suppose that $U(t) \in \mathcal{U}(q, R_q; \mathcal{T})$ for $0 \leq q \leq 1$. Under **Assumptions 1–7**, for a fixed point $t \in [0, 1]$, if $\rho_t = O[\{\log p / (n\bar{m}h) + \log p / n\}^{1/2} + h^2]$ and $\min_{j \in J(t)} \Pi_{jj}(t) \geq 2\gamma_t$, where $\gamma_t^2 = O[\{\log p / (n\bar{m}h) + \log p / n\}^{1/2} + h^2]$, then

$$d\{U(t), \hat{U}(t)\} = O_p \left[\left\{ \left(\frac{\log p}{n\bar{m}h} + \frac{\log p}{n} \right)^{1/2} + h^2 \right\}^{1-q/2} \right].$$

From **Lemma 1**, the condition $\min_{j \in J(t)} \Pi_{jj}(t) \geq 2\gamma_t$ together with the choice of $\gamma(t)$ in **Theorem 1** ensures that signal variables can be distinguished from the noise, which leads to consistent variable selection. The parameter ρ_t is to balance the trade-off between bias and variance. The rate of convergence in **Theorem 1** consists of two parts, the variance term $\{\log p / (n\bar{m}h) + \log p / n\}^{1/2}$ and the bias term h^2 for $q = 0$, which is consistent with that of the mean estimation in Zhang and Wang (2016) up to the $\log p$ term accounting for high dimensionality. The convergence rate depends on \bar{m} through the total number of observations $n\bar{m}$. Thus, the magnitude of \bar{m} can be of any order of the sample size n as long as $\log p / (nh) \rightarrow 0$ and $h \rightarrow 0$, which demonstrates the advantage of our proposal in handling the sparsely observed data. A careful inspection shows that the convergence rate exhibits a phase transition phenomenon. When $\bar{m}h \rightarrow \infty$, the sampling frequency \bar{m} has no effect on the resulting rate, $(\log p / n)^{1/2-q/4}$, as if the whole curves are completely observed. Otherwise, the estimates attain the nonparametric rate $[\{\log p / (n\bar{m}h)\}^{1/2} + h^2]^{1-q/2}$ as if all $n\bar{m}$ observations are independently observed.

In contrast, the convergence rate is of the order $\{[\log p/(nh)]^{1/2} + h^2\}^{1-q}$ for the dynamic covariance estimation in Chen and Leng (2016) under the assumption that the columns of covariance matrices possess the l_q -type sparsity structure. There are two notable differences between the two rates. First, in our setting, the effective sample size is $(n\bar{m}h) \wedge n$ instead of nh which differs from the conventional nonparametric scheme. This is because we take the correlation among observations from the same subject into account which is an important nature of functional data or generally the repeated measurements data. Second, the dependence on the q for these two convergence rates is different. In the static case, it is known that the optimal rate for eigenvector estimation $(\log p/n)^{1/2-q/4}$ is faster than the rate obtained for covariance estimation $(\log p/n)^{1/2-q/2}$ (Bickel and Levina 2008; Vu and Lei 2012; Cai and Zhou 2012). Likewise in the dynamic setting, the convergence rate for eigenvectors in Theorem 1 is faster than the rate of the corresponding covariance estimation. The theoretical finding is corroborated in empirical studies that the eigenvector estimators based on the eigen-decomposition of dynamic covariance estimates perform sub-optimally.

Moreover, Lemma 3 reveals that $d\{U, \hat{U}\} \leq C_q \|\hat{\Sigma} - \Gamma\|_\infty^{1-q/2}$ where $C_q = Cd^2R_q$ for some positive constant $C > 0$. Therefore, the quantity R_0 may be allowed to grow to infinity, and the consistency of the estimator is guaranteed as long as $R_0\{(\log p/(nmh) + \log p/n)^{1/2} + h^2\} \rightarrow 0$. Given that the bandwidth is carefully tuned to balance the bias and variance, the quantity \bar{m} plays a crucial role in the convergence rate, which is illustrated in Corollary 1.

Corollary 1. Suppose that $U(t) \in \mathcal{U}(q, R_q)$ and conditions in Theorem 1 hold and t is a fixed point in $[0, 1]$.

1. When $\bar{m}/(n/\log p)^{1/4} \rightarrow 0$ and $h = O\{[\log p/(n\bar{m})]^{1/5}\}$,

$$d\{U(t), \hat{U}(t)\} = O_p \left[\left\{ \left(\frac{\log p}{n\bar{m}h} \right)^{1/2} + h^2 \right\}^{1-q/2} \right].$$

2. When $\bar{m}/(n/\log p)^{1/4} \rightarrow C$, where $C > 0$, and $h = O\{(\log p/n)^{1/4}\}$,

$$d\{U(t), \hat{U}(t)\} = O_p \left\{ \left(\frac{\log p}{n} \right)^{1/2-q/4} \right\}.$$

3. When $\bar{m}/(n/\log p)^{1/4} \rightarrow \infty$, $h = o\{(\log p/n)^{1/4}\}$ and $\bar{m}h \rightarrow \infty$,

$$d\{U(t), \hat{U}(t)\} = O_p \left\{ \left(\frac{\log p}{n} \right)^{1/2-q/4} \right\}.$$

As Corollary 1 reveals, the phase transition occurs when \bar{m} is of the order $(n/\log p)^{1/4}$. When \bar{m} is relatively small as in case (1), the nonparametric rate is determined jointly by quantities n and \bar{m} . With \bar{m} grows such that $\bar{m} \gtrsim (n/\log p)^{1/4}$, the rate achieves $(\log p/n)^{1/2-q/4}$ regardless of \bar{m} which coincides with the optimal rate for estimating static eigenvectors. Although the rates are of the same order in cases (2) and (3) which fall into the parametric paradigm, the bias in case (2) is nonvanishing (Zhang and Wang 2016). With the advantage of data pooling, the grids are allowed to be sparse under the irregular design as long as the sample size n suffices.

3.2. Rate of Convergence Under the Common Design

In this section, we focus on the common design where sampling locations $t_l, l = 1, \dots, m$ are deterministic.

Assumption 8. Under the common design, t_l 's are fixed and distinct, and $\max_{0 \leq l \leq m} |t_{l+1} - t_l| \leq Cm^{-1}$, where $t_0 = 0, t_{m+1} = 1$.

Assumption 9. The sampling frequency $m \rightarrow \infty$ and $1/(mh) = O(1), h \rightarrow 0$ as $n \rightarrow \infty$.

Under the common design, the data should be observed on sufficiently dense grids. To see this, if m is finite, no data is available in the suitably small neighboring region for some t , which causes large bias for the resulting estimates. Moreover, Assumption 9 guarantees that $h \geq \min_{j=1, \dots, m} |t - t_j| = O(1/m)$ for each $t \in \mathcal{T}$ to avoid the trivial estimator.

Theorem 2. Suppose that $U(t) \in \mathcal{U}(q, R_q; \mathcal{T})$ for $0 \leq q \leq 1$. Under Assumptions 1–5, 8, and 9, for a fixed point $t \in [0, 1]$, if $\rho_t = O\{[\log p/(nmh) + \log p/n]^{1/2} + h^2\}$ and $\min_{j \in J(t)} \Pi_{jj}(t) \geq 2\gamma_t$, where $\gamma_t^2 = O\{[\log p/(nmh) + \log p/n]^{1/2} + h^2\}$, then

$$d\{U(t), \hat{U}(t)\} = O_p \left[\left\{ \left(\frac{\log p}{nmh} + \frac{\log p}{n} \right)^{1/2} + h^2 \right\}^{1-q/2} \right].$$

At first glance, the convergence rates under common and irregular designs are similar. However, since the data are observed at common locations under this design, the number of locations where the data are used for estimation is of the order mh . Consequently, the sampling frequency m is required to be sufficiently large which differs from the case under the irregular design. The effect of m on the convergence rate is illustrated in Corollary 2.

Corollary 2. Suppose that $U(t) \in \mathcal{U}(q, R_q)$ and conditions in Theorem 2 hold and t is a fixed point in $[0, 1]$.

1. When $m/(n/\log p)^{1/4} \rightarrow 0$ and $h = O(1/m)$,

$$d\{U(t), \hat{U}(t)\} = O_p \left\{ \left(\frac{1}{m^2} \right)^{1-q/2} \right\}.$$

2. When $m/(n/\log p)^{1/4} \rightarrow C$, where $C > 0$, and $h = O\{(\log p/n)^{1/4}\} = O(1/m)$,

$$d\{U(t), \hat{U}(t)\} = O_p \left\{ \left(\frac{\log p}{n} \right)^{1/2-q/4} \right\}.$$

3. When $m/(n/\log p)^{1/4} \rightarrow \infty$, $h = o\{(\log p/n)^{1/4}\}$ and $mh \rightarrow \infty$,

$$d\{U(t), \hat{U}(t)\} = O_p \left\{ \left(\frac{\log p}{n} \right)^{1/2-q/4} \right\}.$$

The phase transition phenomenon appears more complex under the common design because of the interplay among the quantities h, n and m . The bandwidth is usually chosen to be of the order $\{\log p/(n\bar{m})\}^{1/5}$ to balance the associated bias and variance. Yet, restricted by the condition $1/(mh) = O(1)$

that guarantees observations available in the local window, the bandwidth is at least of the order $1/m$, see [Corollary 2](#). When m is relatively small, the sampling frequency m plays a dominant role in the convergence rate, that is, $(1/m^2)^{1-q/2}$, which is slower than that under the irregular design. When m grows such that $m \gtrsim (n/\log p)^{1/4}$, the parametric rate $(\log p/n)^{1/2-q/4}$ can be achieved. Again, it requires $m \rightarrow \infty$ to guarantee the consistency, otherwise, the bias is not negligible.

In summary, the phase transition occurs at the same order, $m = O\{(n/\log p)^{1/4}\}$, for both designs. When $m/(n/\log p)^{1/4} \rightarrow 0$, the rate depends on the total number of observations $n\bar{m}$ under the irregular case, while the rate is solely determined by the sampling frequency m under the common case. Otherwise, both designs achieve parametric rates. Under the irregular design, the estimator is consistent as long as $(\log p/n\bar{m}h) \vee (\log p/n) \rightarrow 0$. Thus, as the sample size permits, we can handle the extremely sparse case, that is, the sampling frequency is allowed to be very small. However, the quantity m should be sufficiently large to achieve reasonable estimates under the common design. Further, when the sampling frequency is small, the irregular design is preferable to the common design with faster convergence rates.

4. Simulation

In this section, several experiments are conducted to evaluate the numerical performance of our proposal under irregular and common designs. The observations are generated from the model, $\mathbf{y}_i(t_{il}) = \sum_{k=1}^{10} \xi_{ik}(t_{il})\mathbf{u}_k(t_{il}) + \epsilon_{il}$, $i = 1, \dots, n$; $l = 1, \dots, m_i$, where $\xi_{ik} \in \mathbb{R}$, $\mathbf{y}_i(t_{il}), \mathbf{u}_k(t_{il}) \in \mathbb{R}^p$ and $\epsilon_{il} \stackrel{\text{iid}}{\sim} N_p(0, \sigma^2 I_p)$. Moreover, we set $\xi_{ik}(t_{il}) \equiv \xi_{ik} \stackrel{\text{iid}}{\sim} N(0, \lambda_k)$ where $\lambda = (30, 18, 10, 5, 3, 2, 1, 0.5, 0.2, 0.1)^T$, which follows the classical generation mechanism of functional data. The sparse eigenvectors \mathbf{u}_k are obtained by applying Gram-Schmidt orthonormalization on $\mathbf{v}_k(t)$ defined later. Let $v_{k,(k-1) \times 5 + r}(t) = \phi_r(t)$, $k = 1, \dots, 10$, $r = 1, \dots, 5$, and other entries of \mathbf{v}_k be zero, where $v_{k,j}$ is the j -th element of the vector \mathbf{v}_k , $\phi_r(t)$ are functions in the Fourier basis, $\phi_r(t) = \sqrt{2} \sin(\pi(r+1)t)$ when r is odd, $\phi_r(t) = \sqrt{2} \cos(\pi r t)$ when r is even. The locations of nonzero elements are different for different eigenvectors.

We design simulation settings to demonstrate the effect of the sample size n and the sampling frequency under both common and irregular cases. Under the irregular design, we consider six settings for various combinations of m_i and n where m_i are iid from a discrete uniform distribution on the set \mathcal{M} . *Setting 1*: $n = 100$ and $\bar{m} = 100$, $\mathcal{M} = \{95, 100, 105\}$. *Setting 2*: $n = 100$ and $\bar{m} = 50$, $\mathcal{M} = \{45, 50, 55\}$. *Setting 3*: $n = 100$ and $\bar{m} = 20$, $\mathcal{M} = \{15, 20, 25\}$. *Setting 4*: $n = 500$ and $\bar{m} = 20$, $\mathcal{M} = \{19, 20, 21\}$. *Setting 5*: $n = 500$ and $\bar{m} = 10$, $\mathcal{M} = \{9, 10, 11\}$. *Setting 6*: $n = 500$ and $\bar{m} = 4$, $\mathcal{M} = \{3, 4, 5\}$. The time points t_{il} are iid sampled from the uniform distribution on $[0, 1]$. Under the common design, the data are sampled at $t_{il} = (2l)/(2m+1)$ with the sample size $n = 100$ and the sampling frequency $m = 20, 50, 100$, respectively. In each setting, we repeat 100 times independently for $p = 50, 100, 200$ and the noise level $\sigma^2 = 1, 3$, respectively.

For comparison purposes, we estimate eigenvectors by performing conventional PCA on sparse covariance matrices obtained by dynamic covariance models (DCM) ([Chen and Leng 2016](#)). Since the bandwidth selected by the leave-one-point-out cross-validation in the DCM might be inappropriate for repeated measurements, we use the leave-one-curve-out cross-validation instead and denote the resulting model by DCM+. Under the common design, we include the methods of [Berrendero, Justel, and Svarc \(2011\)](#) and [Johnstone and Lu \(2009\)](#), denoted by BJS and DT, respectively. Note that the parameters in DT are set to the recommended values in the original paper. We evaluate the performance of estimators by the mean integrated squared error (MISE) which is approximated by computing the average of squared errors, defined in (7), on a grid of 50 equally spaced points. Since the BJS and DT only obtain estimates under the sampling locations, to calculate the MISE for these two methods, we simply generate data at evaluated grids to obtain corresponding estimates. Note that the resulting errors MISE_{BJS} and MISE_{DT} are not relevant to the sampling frequency.

We begin with illustrating the selection and the performance of tuning parameters in our method. We set $d = 3$ under which the FVE is about 85%, and other parameters are chosen as discussed in [Section 2.3](#). Specifically, the bandwidth h is selected by leave-one-curve-out cross-validation, while the ρ_t and γ_t are determined by 5-fold cross-validation to save the computation time. Since the quantity m or \bar{m} might be large, to further reduce the computation, we randomly choose 10 observational points with equal probability for each curve to calculate the error in the validation step to tune the bandwidth. The effectiveness of the selection strategy for the parameters is illustrated in [Figure 1](#). As is shown, the selected parameters well depict the true smoothness of covariance matrices, the sparsity level and model complexity of the eigenvectors, respectively.

The results for $p=100$ and 200 with $\sigma^2 = 3$ are summarized in [Tables 1](#) and [2](#), while the results for $p = 50$ are qualitatively similar, thus not reported for space economy. Moreover, the results with $\sigma^2 = 1$ are provided in the Supplementary Material. The errors of the initial estimate \hat{U}^0 and the refined estimate \hat{U} are denoted by MISE_0 and MISE , respectively. As [Tables 1](#) and [2](#) show, the proposed method outperforms other methods in all settings, while the refined estimators perform slightly better than the initial estimates under both designs. The DCM and DCM+ methods perform decently when $m = 100$ and $p = 100$. However, their performance deteriorates significantly if the dimension increases or the sampling frequency becomes small. Under the irregular design, as noted in [Section 3.1](#), the grids can be sparse to obtain consistent estimation as the sample size permits. Thus, even in the very sparse case as $\bar{m} = 4$ or 10 ($n = 500$), the error is still well controlled and even smaller than that when $\bar{m} = 20$ ($n = 100$) due to a larger sample size and the advantage of data pooling. Moreover, when the total number of observations is comparable, a larger sample size usually leads to a better estimate. Under the common design, it is not surprising that the BJS method fails to obtain reasonable estimators as it did not accommodate high dimensionality. Although the DT yields sparse eigenvectors, it tends to select too few coordinates, which introduces larger bias. Further, as shown in the left of [Figure 2](#),

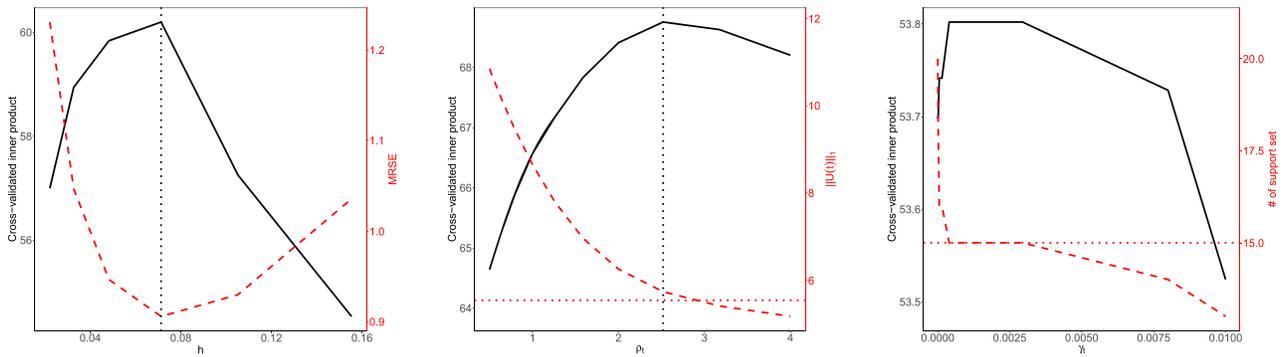


Figure 1. The performance of the cross-validation to select parameters under the common design with $p = 100, m = 50$ and $\sigma^2 = 3$. All the parameters are selected by maximizing the cross-validated inner product (black, solid, left y label). Left: bandwidth selection. The right y label indicates the $MRSE(h) = m^{-1} \sum_{l=1}^m \|\hat{\Sigma}_l(t) - \Sigma(t)\|_F^2 / \|\Sigma(t)\|_F^2$. An ideal h should be close to the $h_{opt} = \arg \min MRSE(h)$. The selected bandwidth matches the one attaining the minimal MRSE (treated as a benchmark; see red, dashed line, right y label). Middle: sparsity parameter selection. For a fixed t , the right y label represents $\|U(t)\|_1$. The $\|\cdot\|_1$ of estimated $U(t)$ with the selected ρ_t (red, dashed, right y label) approximately meets that of the true matrix (indicated by the horizontal dotted line). Right: thresholding parameter selection. The right y label represents the size of the support set. The γ_t is selected as the minimum value which maximizes the cross-validated inner product. The number of retained variables with the selected γ_t (red, dashed, right y label) is equal to the true number of relevant variables (horizontal, dotted).

Table 1. Average integrated squared errors and standard deviations over 100 replications for different settings under the irregular design and $\sigma^2 = 3$.

Model	MISE ₀	MISE	MISE _{DCM}	MISE _{DCM+}	
$p=100, n=100$	$\bar{m} = 100$	0.033 (0.014)	0.031(0.012)	0.040(0.037)	0.038(0.031)
	$\bar{m} = 50$	0.042(0.014)	0.041(0.016)	0.128(0.189)	0.121(0.157)
	$\bar{m} = 20$	0.112(0.056)	0.102(0.061)	0.488(0.237)	0.498(0.226)
$p=100, n=500$	$\bar{m} = 20$	0.023(0.003)	0.022(0.002)	0.021(0.006)	0.022(0.005)
	$\bar{m} = 10$	0.035(0.008)	0.034(0.007)	0.059(0.091)	0.080(0.130)
	$\bar{m} = 4$	0.078(0.016)	0.068(0.016)	0.465(0.158)	0.460(0.167)
$p=200, n=100$	$\bar{m} = 100$	0.038(0.023)	0.036(0.022)	0.127(0.168)	0.130(0.177)
	$\bar{m} = 50$	0.058(0.037)	0.059(0.044)	0.361(0.263)	0.322(0.275)
	$\bar{m} = 20$	0.148(0.071)	0.123(0.079)	0.577(0.193)	0.574(0.184)
$p=200, n=500$	$\bar{m} = 20$	0.024(0.002)	0.023(0.002)	0.095(0.154)	0.077(0.130)
	$\bar{m} = 10$	0.036(0.005)	0.035(0.005)	0.288(0.206)	0.269(0.214)
	$\bar{m} = 4$	0.104(0.025)	0.080(0.031)	0.515(0.087)	0.515(0.086)

our method performs significantly better than other methods uniformly over t . The results under other simulation settings lead to similar conclusions and are not reported.

At last, to illustrate the performance of the refined estimates in achieving model parsimony, two criteria $TNR = TN/(TN+FP)$ and $TPR = TP/(TP+FN)$ are reported in Figure 2, where TP and TN are abbreviations for true positives and true negatives, respectively, that is, the number of significant or nonsignificant variables correctly identified by our method, similarly FP and FN stand for false positives and false negatives. The numerical values of $\hat{\Pi}_{jj}$ over 10^{-6} in magnitude are considered nonzero for the consideration of computation accuracy. As the principal eigenvectors vary with t , the signal $\min_{j \in J(t)} \hat{\Pi}_{jj}(t)$ may be close to zero which leads to a bit lower TPR at some t . The result reveals the fact that the refined estimate by thresholding yields better variable selection results, which is particularly useful for model interpretation.

5. Real Data Example

The heartbeat sound dataset from <http://www.timeseriesclassification.com/description.php?Dataset=Heartbeat> were sourced from several contributors around the world, collected at either a

clinical or nonclinical environment (Liu et al. 2016). The heart sound recordings were collected from different locations on the body. The typical four locations are the aortic area, pulmonic area, tricuspid area and mitral area, but could be one of nine different locations. Each recording was truncated to 5 seconds. A Spectrogram of each instance was then created with a window size of 0.061 seconds and an overlap of 70%. Each instance is arranged such that each dimension is a frequency band from the spectrogram, with $p = 61$ and $m = 405$. We focus on extracting the dynamic features of $n = 295$ pathological patients.

For the purpose of evaluation, we compute the mean squared recovery error on the held-out test sample. Specifically, we randomly choose 100 subjects as the test data, and treat the remaining data as the training set. The estimators are obtained using the training sample under different methods including the proposed method, DCM, DCM+, BJS, and DT. To better demonstrate the advantages of smoothing, we compare with a variant of our method with the bandwidth nearly 0, denoted by the rough estimator. More specifically, this variant is obtained by using the same optimization technique in our article and only replacing the smoothed covariance matrix with the sample covariance matrix. The mean squared recovery errors under a different number of eigenvectors d are calculated over all 405 grids, that is, $(100m)^{-1} \sum_{i,l} \|y_{il} - \hat{\mu}_l - \hat{U}_l \hat{U}_l^T (y_{il} - \hat{\mu}_l)\|^2$, where y_{il} are observations of the i -th test subject, $\hat{\mu}_l$ is the sample mean vector and \hat{U}_l is the estimator at the l -th observed locations. Moreover, to compare with the performance of FPCA (Ramsay and Silverman 2005) in terms of the low-dimensional representation, we implement FPCA for each functional variable and calculate the mean squared recovery error, $(100m)^{-1} \sum_{i,j,l} (y_{ijl} - \hat{\mu}_{jl} - \sum_{k=1}^d \hat{\xi}_{jk} \hat{\psi}_{jkl})^2$, where $\hat{\xi}_{jl} = m^{-1} \sum_{l=1}^m (y_{ijl} - \hat{\mu}_{jl}) \hat{\psi}_{jkl}$ and $\hat{\psi}_{jkl}$ is the value at the l th time point of the k th eigenfunction for the j th functional variable. Note that Proposed₀ and Proposed represent the proposed method without and with the refinement step, respectively. While the refinement step has little effect on the recovery errors, it in fact leads to a more parsimonious model with fewer retained variables, screening out about 20% insignificant variables. As Table 3 shows, the proposed method obtains favorable performance over other methods, suggesting

Table 2. Average integrated squared errors and standard deviations over 100 replications for different settings under the common design and $\sigma^2 = 3$.

Model		MISE ₀	MISE	MISE _{DCM}	MISE _{DCM+}	MISE _{BJS}	MISE _{DT}
$p=100$ $n=100$	$m=100$	0.032(0.016)	0.030(0.016)	0.064(0.096)	0.058(0.065)	0.643(0.075)	0.647(0.099)
	$m=50$	0.040(0.014)	0.037(0.013)	0.212(0.177)	0.199(0.159)		
	$m=20$	0.128(0.081)	0.114(0.085)	0.722(0.260)	0.706(0.276)		
$p=200$ $n=100$	$m=100$	0.037(0.028)	0.036(0.025)	0.242(0.275)	0.180(0.219)	1.021(0.098)	0.645 (0.118)
	$m=50$	0.053(0.028)	0.053(0.030)	0.581(0.299)	0.609(0.292)		
	$m=20$	0.157(0.063)	0.132(0.074)	1.026(0.125)	1.018(0.147)		

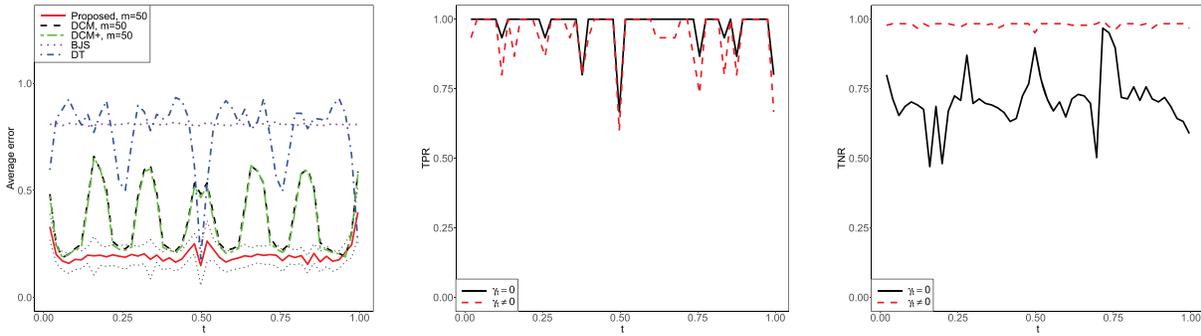


Figure 2. The common design: $p = 100$ and $\sigma^2 = 3$. Left: The errors, defined in (7), of different methods over 50 equally spaced points in $[0, 1]$. The two black dotted lines reflect uncertainty (standard deviation) of errors for the proposed estimate. Right: The performance of TPR (middle) and TNR (right) over t of the proposed estimates with $\gamma_t = 0$ (solid) or $\gamma_t \neq 0$ (dashed). It illustrates the advantage of the refinement step in screening out irrelevant variables and achieving desired model parsimony, which improves TNR without decreasing TPR.

Table 3. The mean squared recovery errors on the test data with a different number of principal eigenvectors/eigenfunctions for different methods.

	Proposed ₀	Proposed	DCM	DCM+	BJS	DT	Rough	FPCA
$d = 5$	0.201	0.201	0.274	0.222	0.279	0.302	0.273	3.156
$d = 6$	0.113	0.112	0.218	0.151	0.192	0.228	0.191	3.120
$d = 7$	0.072	0.072	0.187	0.114	0.132	0.184	0.131	3.078
$d = 8$	0.045	0.045	0.165	0.097	0.091	0.157	0.093	3.044

Table 4. The mean squared recovery errors on the test data for different methods under the irregular case with different \tilde{m} and $d = 6, 8$.

	$d = 6$				$d = 8$			
	Proposed ₀	Proposed	DCM	DCM+	Proposed ₀	Proposed	DCM	DCM+
$\tilde{m} = 30$	0.191	0.191	0.226	0.232	0.124	0.124	0.151	0.162
$\tilde{m} = 50$	0.150	0.150	0.174	0.202	0.072	0.072	0.120	0.150
$\tilde{m} = 80$	0.147	0.147	0.176	0.175	0.068	0.068	0.116	0.114

more accurate estimation from our approach for the dynamic PCA. In particular, our approach outperforms the rough estimator, showing the usefulness of the smoothing strategy. The DCM performs worse than DCM+ because it tends to select smaller bandwidth which is not satisfactory in this case. As seen in Table 3, the representation obtained by FPCA is not promising with large recovery errors.

To demonstrate the performance under the irregular design, we randomly sample $\tilde{m} = 30, 50, 80$ measurements, respectively, with equal probability from each subject of the training sample, and use the obtained irregular data for estimation. Since the methods BJS and DT are not feasible for the irregular design, we compare the recovery errors of the other three approaches for the dynamic PCA. Note that we report the results with $d = 6, 8$ in Table 4 for space economy, since the results exhibit a similar pattern for other values of d . It is demonstrated that the proposed method is capable of producing more desirable

estimates. Moreover, despite fewer observations, the recovery errors in Table 4 are still much lower than the errors of FPCA obtained under the common design.

6. Concluding Remarks

We propose a unified framework to estimate dynamic eigenvectors in high-dimensional settings by combining the local linear smoothing and the sparsity constraint under both common and irregular designs. The resulting estimators satisfy sparsity and orthogonality simultaneously. Different from the conventional nonparametric smoothing, the rates of convergence depend on the sampling frequency and the sample size jointly, exhibiting the phase transition phenomenon. When the sampling frequency is suitably large, the obtained rates are optimal as if the whole curves are available under both designs. Otherwise, the irregular design is preferred with a faster rate of convergence.

It is interesting to study other types of smoothing techniques, such as smoothing splines, for the problem of DPCA. Moreover, since the PCA is sensitive to outliers, it is also useful to develop a dynamic robust model. These topics are beyond the scope of the current article and deserve future study.

Appendix

Appendix A: Auxiliary Lemmas

In the sequel, we suppress the index t of pointwise results for convenience when no ambiguity arises. We write $a \asymp b$ if $a \lesssim b$ and $b \lesssim a$ hold simultaneously.

Lemma 2. Given the random function $\mathbf{X}(t) \in \mathbb{R}^p$ with the mean function $\boldsymbol{\mu}(t) = E\mathbf{X}(t)$, $t \in \mathcal{T}$. Denote $U(t) = (\mathbf{u}_1(t), \dots, \mathbf{u}_d(t))$, where $\mathbf{u}_1(t), \dots, \mathbf{u}_d(t)$ are the first d eigenvectors of $\mathbf{X}(t)$. Then, $U(t)$

is a solution of the optimization problem (2). Moreover, solving (2) is reduced to performing multivariate PCA at each t .

Proof. Let $h(V(t)) = E\|\mathbf{X}(t) - \boldsymbol{\mu}(t) - V(t)V(t)^T\{\mathbf{X}(t) - \boldsymbol{\mu}(t)\}\|^2$. We first show that $U(t)$ is a solution of (2), and define the optimization problem for multivariate PCA at each t as,

$$\begin{aligned} \min_{V(t)} \quad & h(V(t)) \\ \text{s.t.} \quad & V(t)^T V(t) = I_d. \end{aligned} \quad (\text{A.1})$$

Denote by $U^*(t)$ the solution of (A.1), and let $h^*(t) = h(U^*(t))$. Moreover, we denote the solution of (2) by $U^{**}(t)$, and let $h^{**}(t) = h(U^{**}(t))$. Thus, we have $\int_{\mathcal{T}} \{h^*(t) - h^{**}(t)\} dt \leq 0$. Since $U^*(t)$ is feasible for (2), then $\int_{\mathcal{T}} h^{**}(t) dt \leq \int_{\mathcal{T}} h^*(t) dt$. Due to the fact that $\int_{\mathcal{T}} h^*(t) dt = \int_{\mathcal{T}} h^{**}(t) dt$, we conclude that $U^*(t)$ is a solution of (2). Note that (A.1) seeks an orthonormal matrix $V(t) \in \mathbb{R}^{p \times d}$ to minimize $E\|\mathbf{X}(t) - \boldsymbol{\mu}(t) - V(t)V(t)^T\{\mathbf{X}(t) - \boldsymbol{\mu}(t)\}\|^2$, which is equivalent to finding an orthonormal matrix to maximize $\text{Tr}(V(t)^T \Sigma(t) V(t))$. Thus, $U(t)$ is a solution of (A.1), then it is also a solution of (2).

Next, we show that (2) is reduced to the multivariate PCA at each t . If $U^{**}(t)$ is not the solution of (A.1), then there exists some $\tilde{t} \in \mathcal{T}$ such that the columns of $U^{**}(\tilde{t})$ does not correspond to the first d eigenvectors of $\mathbf{X}(\tilde{t})$. It contradicts the argument that $U^{**}(t)$ is a solution of (2) since replacing $U^{**}(\tilde{t})$ with $U(\tilde{t})$ leads to a smaller objective value for (2). \square

Lemma 3. Recall that U and \hat{U} are true and estimated principal eigenvectors with projection matrices $\Pi = UU^T$ and $\hat{\Pi} = \hat{U}\hat{U}^T$, respectively. Assume $U \in \mathcal{U}(q, R_q)$, $0 \leq q \leq 1$. If conditions in Lemma 1 hold, with appropriate choice of parameters, we have $d\{U, \hat{U}\} \leq C_q \|\hat{\Sigma} - \Gamma\|_{\infty}^{1-q/2}$ where $C_q = Cd^2 R_q$ for some positive constant $C > 0$.

Lemma 4. If $h \rightarrow 0$ and $n\bar{m}h \rightarrow \infty$, then under the irregular design, we have

- (a) $R_{\ell} \asymp n\bar{m}h^{\ell}(1 + o_p(1))$, $\ell = 0, 1, 2$. Moreover, $R_2 R_0 - R_1^2 \asymp n^2 \bar{m}^2 h^2 (1 + o_p(1))$.
- (b) $E\left[\left\{R_2 K_h(t_{i\ell} - t) - R_1 K_h(t_{i\ell} - t)(t_{i\ell} - t)\right\} x_{ijl} x_{ikl}\right]^2 = O(n^2 \bar{m}^2 h^3)$.
- (c) $E\left(\tilde{w}_{il} x_{ijl} x_{ikl} \tilde{w}_{i'l'} x_{i'j'l'} x_{i'k'l'}\right) = O(n^2 \bar{m}^2 h^4)$ for $(i, l) \neq (i', l')$, where $\tilde{w}_{il} = R_2 K_h(t_{i\ell} - t) - R_1 K_h(t_{i\ell} - t)(t_{i\ell} - t)$.

Lemma 5. Under Assumptions 2–7, we have for each $t \in \mathcal{T}$,

$$\begin{aligned} \max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ \tilde{w}_{il} x_{ijl} x_{ikl} - E\left(\tilde{w}_{il} x_{ijl} x_{ikl}\right) \right\} \right| \\ = O_p\{(\log p)^{1/2} (n^3 \bar{m}^3 h^3 + n^3 \bar{m}^4 h^4)^{1/2}\}. \end{aligned}$$

The proofs of Lemmas 3–5 are deferred to the Supplementary Material. In the following, we provide the proof of Theorem 1, while the proof of Theorem 2 is analogous which could be found in the Supplementary Material.

Appendix B: Proofs of Main Results

Proof of Theorem 1. From Lemma 3, we have $d\{\hat{U}(t), U(t)\} \leq C_q \|\hat{\Sigma}(t) - \Sigma(t) - \sigma^2 I_p\|_{\infty}^{1-q/2}$. Thus, it suffices to quantify the error $\|\hat{\Sigma}(t) - \Sigma(t) - \sigma^2 I_p\|_{\infty}$.

Let $\Omega(t) = E(\mathbf{X}(t)\mathbf{X}(t)^T)$, we have $\Sigma(t) = \Omega(t) - \boldsymbol{\mu}(t)\boldsymbol{\mu}(t)^T$. Under the irregular design, using the triangle inequality,

$$\|\hat{\Sigma}(t) - \Sigma(t) - \sigma^2 I_p\|_{\infty}$$

$$\begin{aligned} &= \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{Y}_{il} \mathbf{Y}_{il}^T - \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{Y}_{il} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{Y}_{il}^T \right. \\ &\quad \left. - \{\Omega(t) - \boldsymbol{\mu}(t)\boldsymbol{\mu}(t)^T\} - \sigma^2 I_p \right\|_{\infty} \\ &\leq \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{Y}_{il} \mathbf{Y}_{il}^T - \Omega(t) - \sigma^2 I_p \right\|_{\infty} \\ &\quad + \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{Y}_{il} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{Y}_{il}^T - \boldsymbol{\mu}(t)\boldsymbol{\mu}(t)^T \right\|_{\infty} \\ &= M_1(t) + M_2(t). \end{aligned} \quad (\text{B.2})$$

Note that

$$\begin{aligned} M_1(t) &= \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} (\mathbf{x}_{il} + \boldsymbol{\epsilon}_{il}) (\mathbf{x}_{il} + \boldsymbol{\epsilon}_{il})^T - \Omega(t) - \sigma^2 I_p \right\|_{\infty} \\ &\leq \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il} \mathbf{x}_{il}^T - \Omega(t) \right\|_{\infty} + 2 \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il} \boldsymbol{\epsilon}_{il}^T \right\|_{\infty} \\ &\quad + \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \boldsymbol{\epsilon}_{il} \boldsymbol{\epsilon}_{il}^T - \sigma^2 I_p \right\|_{\infty}, \end{aligned} \quad (\text{B.3})$$

where $\boldsymbol{\epsilon}_{il} = (\epsilon_{i1l}, \dots, \epsilon_{ipl})^T$. Similarly,

$$\begin{aligned} M_2(t) &\leq \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il}^T - \boldsymbol{\mu}(t)\boldsymbol{\mu}(t)^T \right\|_{\infty} \\ &\quad + 2 \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \boldsymbol{\epsilon}_{il}^T \right\|_{\infty} \\ &\quad + \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \boldsymbol{\epsilon}_{il} \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \boldsymbol{\epsilon}_{il}^T \right\|_{\infty}. \end{aligned} \quad (\text{B.4})$$

To bound the term $\Delta_n = \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il} \mathbf{x}_{il}^T - \Omega(t) \right\|_{\infty}$, we have

$$\begin{aligned} \Delta_n &= \left\| \sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} \mathbf{x}_{il} \mathbf{x}_{il}^T - \Omega(t) \right\|_{\infty} \\ &= \max_{j,k} \left| \frac{\sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ \tilde{w}_{il} x_{ijl} x_{ikl} - E(\tilde{w}_{il} x_{ijl} x_{ikl}) \right\}}{R_0 R_2 - R_1^2} \right| \\ &\quad + \frac{\sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ E(\tilde{w}_{il} x_{ijl} x_{ikl}) - \tilde{w}_{il} \omega_{jk}(t) \right\}}{R_0 R_2 - R_1^2} \\ &\leq \frac{\max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ \tilde{w}_{il} x_{ijl} x_{ikl} - E(\tilde{w}_{il} x_{ijl} x_{ikl}) \right\} \right|}{\left| R_0 R_2 - R_1^2 \right|} \\ &\quad + \frac{\max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ E(\tilde{w}_{il} x_{ijl} x_{ikl}) - \tilde{w}_{il} \omega_{jk}(t) \right\} \right|}{\left| R_0 R_2 - R_1^2 \right|} \\ &= I + II, \end{aligned}$$

where $\tilde{w}_{il} = R_2 K_h(t_{i\ell} - t) - R_1 K_h(t_{i\ell} - t)(t_{i\ell} - t)$ and $\omega_{jk}(t) = E\{x_{ij}(t)x_{ik}(t)\}$. The second equality holds due to the fact that $\sum_{i=1}^n \sum_{l=1}^{m_i} w_{il} = 1$.

Denote $\zeta_1 = \max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ \tilde{w}_{il} x_{ijl} x_{ikl} - E(\tilde{w}_{il} x_{ijl} x_{ikl}) \right\} \right|$. By Lemma 5, we conclude $\zeta_1 = O_p\{(\log p)^{1/2} (n^3 \bar{m}^3 h^3 + n^3 \bar{m}^4 h^4)^{1/2}\}$.

From (a) of Lemma 4, we have $R_0R_2 - R_1^2 \asymp n^2\bar{m}^2h^2(1 + o_p(1))$. Consequently,

$$I = O_p \left\{ \left(\frac{\log p}{n\bar{m}h} + \frac{\log p}{n} \right)^{1/2} \right\}.$$

Next we bound the term II. Denote that $\zeta_2 = \max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ E(\tilde{w}_{il}x_{ijl}x_{ikl}) - \tilde{w}_{il}\omega_{jk}(t) \right\} \right|$. Notice that

$$\begin{aligned} \zeta_2 &= \max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left\{ E(\tilde{w}_{il}\omega_{jk}(t_{il})) - \tilde{w}_{il}\omega_{jk}(t) \right\} \right| \\ &= \max_{j,k} \left| \sum_{i=1}^n \sum_{l=1}^{m_i} \left(E \left[\tilde{w}_{il} \left\{ \omega_{jk}(t) + \omega_{jk}^{(1)}(t)(t_{il} - t) + \frac{\omega_{jk}^{(2)}(\xi_{il})}{2}(t_{il} - t)^2 \right\} \right] - \tilde{w}_{il}\omega_{jk}(t) \right) \right| \\ &\leq \max_{j,k} |\omega_{jk}(t)| \left| \sum_{i=1}^n \sum_{l=1}^{m_i} (\tilde{w}_{il} - E\tilde{w}_{il}) \right| + \left| \sum_{i=1}^n \sum_{l=1}^{m_i} E \left\{ \frac{\tilde{w}_{il}\omega_{jk}^{(2)}(\xi_{il})}{2}(t_{il} - t)^2 \right\} \right| \\ &= II_1 + II_2, \end{aligned}$$

where ξ_{il} is between t and t_{il} , and the inequality holds since $\sum_{i=1}^n \sum_{l=1}^{m_i} \tilde{w}_{il}(t_{il} - t) = 0$. Using similar arguments for the proof of Lemma 4(b), we obtain $II_1 = O_p\{(n\bar{m}h)^{3/2}\}$. To bound the term II_2 , notice that $\omega_{jk}^{(2)}(\xi_{il})$ is bounded by Assumption 3 and

$$\begin{aligned} E\{\tilde{w}_{il}(t_{il} - t)^2\} &= E \left[\left\{ R_2K_h(t_{il} - t) - R_1K_h(t_{il} - t)(t_{il} - t) \right\} (t_{il} - t)^2 \right]. \end{aligned}$$

Note that

$$\begin{aligned} E \left\{ R_2K_h(t_{il} - t)(t_{il} - t)^2 \right\} &= E \left[\left\{ \sum_{i=1}^n \sum_{l=1}^{m_i} K_h(t_{il} - t)(t_{il} - t)^2 \right\} K_h(t_{il} - t)(t_{il} - t)^2 \right] \\ &= O(n\bar{m}h^4), \end{aligned}$$

by the change of variables. Analogously, we show that $E \left[\left\{ R_1K_h(t_{il} - t)(t_{il} - t) \right\} (t_{il} - t)^2 \right] = O(n\bar{m}h^4)$. Thus, $II_2 = O(n^2\bar{m}^2h^4)$. According to Lemma 4(a), we have $R_0R_2 - R_1^2 = n^2\bar{m}^2h^2(1 + o_p(1))$. Combining these pieces together leads to the fact that $II = O_p\{h^2 + 1/(n\bar{m}h)^{1/2}\}$.

The rates of other terms are proved using similar arguments which are omitted here to save space. By (B.2), (B.3), and (B.4), we obtain

$$\|\hat{\Sigma}(t) - \Gamma(t)\|_\infty = O_p \left\{ \left(\frac{\log p}{n\bar{m}h} + \frac{\log p}{n} \right)^{1/2} + h^2 \right\},$$

which completes the proof together with Lemma 3. □

Supplementary Materials

The supplementary material contains the algorithm details, additional simulation and technical proofs.

Acknowledgments

Fang Yao is the corresponding author. Xiaoyu Hu is the first author, and conducted this research when she was a PhD student at Peking University. The authors thank the Editor, Associate Editor and two anonymous referees for their many helpful comments that have resulted in significant improvements in the article.

Funding

This research is supported by National Key R&D Program of China (No. 2020YFE0204200, 2022YFA1003801), National Natural Science Foundation of China (No. 11931001, 12292981, 11871080), the LMAM, and the LMEQF.

References

Amini, A. A., and Wainwright, M. J. (2009), "High-Dimensional Analysis of Semidefinite Relaxations for Sparse Principal Components," *The Annals of Statistics*, 37, 2877–2921. [308]

Anderson, T. W. (1963), "Asymptotic Theory for Principal Component Analysis," *The Annals of Mathematical Statistics*, 34, 122–148. [308]

Berrendero, J. R., Justel, A., and Svarc, M. (2011), "Principal Components for Multivariate Functional Data," *Computational Statistics & Data Analysis*, 55, 2619–2634. [308,310,314]

Berthet, Q., and Rigollet, P. (2013), "Computational Lower Bounds for Sparse PCA," arXiv preprint arXiv:1304.0828. [308]

Bickel, P. J., and Levina, E. (2008), "Covariance Regularization by Thresholding," *The Annals of Statistics*, 36, 2577–2604. [309,313]

Brennan, M., and Bresler, G. (2019), "Optimal Average-Case Reductions to Sparse PCA: From Weak Assumptions to Strong Hardness," in *Conference on Learning Theory*, PMLR, pp. 469–470. [308]

Cai, T. T., and Yuan, M. (2011), "Optimal Estimation of the Mean Function Based on Discretely Sampled Functional Data: Phase Transition," *The Annals of Statistics*, 39, 2330–2355. [309,310,312]

Cai, T. T., and Zhou, H. H. (2012), "Minimax Estimation of Large Covariance Matrices Under ℓ_1 -norm," *Statistica Sinica*, 22, 1319–1349. [309,313]

Chen, K., and Lei, J. (2015), "Localized Functional Principal Component Analysis," *Journal of the American Statistical Association*, 110, 1266–1275. [311]

Chen, S., Ma, S., Man-Cho So, A., and Zhang, T. (2020), "Proximal Gradient Method for Nonsmooth Optimization Over the Stiefel Manifold," *SIAM Journal on Optimization*, 30, 210–239. [309,311]

Chen, Z., and Leng, C. (2016), "Dynamic Covariance Models," *Journal of the American Statistical Association*, 111, 1196–1207. [309,311,312,313,314]

Chiou, J.-M., Chen, Y.-T., and Yang, Y.-F. (2014), "Multivariate Functional Principal Component Analysis: A Normalization Approach," *Statistica Sinica*, 24, 1571–1596. [309]

d’Aspremont, A., Ghaoui, L. E., Jordan, M. I., and Lanckriet, G. R. (2007), "A Direct Formulation for Sparse PCA Using Semidefinite Programming," *SIAM Review*, 49, 434–448. [311]

Delmas, J.-P. (2010), "Subspace Tracking for Signal Processing," in *Adaptive Signal Processing: Next Generation Solutions*, Hoboken, New Jersey: Wiley, 211–270. [308]

- Deshpande, Y., and Montanari, A. (2014), "Sparse PCA via Covariance Thresholding," in *Advances in Neural Information Processing Systems*, Red Hook, New York: Curran Associates Inc. pp. 334–342. [308]
- Ding, Y., Kunisky, D., Wein, A. S., and Bandeira, A. S. (2019), "Subexponential-Time Algorithms for Sparse PCA," arXiv preprint arXiv:1907.11635. [308]
- Fan, J., and Gijbels, I. (1996), "Local Polynomial Modelling and Its Applications," *Monographs on Statistics and Applied Probability*, Boca Raton, London, New York: Chapman & Hall/CRC. [309,310,312]
- Holtzman, G., Soffer, A., and Vilenchik, D. (2020), "A Greedy Anytime Algorithm for Sparse PCA," in *Conference on Learning Theory*, pp. 1939–1956. [308]
- Johnstone, I. M., and Lu, A. Y. (2009), "On Consistency and Sparsity for Principal Components Analysis in High Dimensions," *Journal of the American Statistical Association*, 104, 682–693. [308,314]
- Jolliffe, I. T., Trendafilov, N. T., and Uddin, M. (2003), "A Modified Principal Component Technique Based on the LASSO," *Journal of Computational and Graphical Statistics*, 12, 531–547. [308]
- Krauthgamer, R., Nadler, B., and Vilenchik, D. (2015), "Do Semidefinite Relaxations Solve Sparse PCA up to the Information Limit?" *The Annals of Statistics*, 43, 1300–1322. [308]
- Liu, C., Springer, D., Li, Q., Moody, B., Juan, R. A., Chorro, F. J., Castells, F., Roig, J. M., Silva, I., Johnson, A. E., et al. (2016), "An Open ACCESS DATABASE for the evaluation of Heart Sound Algorithms," *Physiological Measurement*, 37, 2181. [315]
- Ma, Z. (2013), "Sparse Principal Component Analysis and Iterative Thresholding," *The Annals of Statistics*, 41, 772–801. [308]
- Mackey, L. W. (2009), "Deflation Methods for Sparse PCA," in *Advances in Neural Information Processing Systems*, Red Hook, New York: Curran Associates Inc. pp. 1017–1024. [311]
- Ramsay, J. O., and Silverman, B. W. (2005), *Functional Data Analysis* (2nd ed.), New York: Springer. [309,315]
- Rice, J. A., and Silverman, B. W. (1991), "Estimating the Mean and Covariance Structure Nonparametrically When the Data Are Curves," *Journal of the Royal Statistical Society, Series B*, 53, 233–243. [311]
- Shen, H. and Huang, J. Z. (2008), "Sparse Principal Component Analysis via Regularized Low Rank Matrix Approximation," *Journal of Multivariate Analysis*, 99, 1015–1034. [308,311]
- Sirovich, L. and Kirby, M. (1987), "Low-Dimensional Procedure for the Characterization of Human Faces," *Journal of the Optical Society of America*, 4, 519–524. [308]
- Turk, M. A. and Pentland, A. P. (1991), "Face Recognition using Eigenfaces," in *IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–591. [308]
- Vu, V. and Lei, J. (2012), "Minimax Rates of Estimation for Sparse PCA in High Dimensions," in *International Conference on Artificial Intelligence and Statistics*, pp. 1278–1286. [309,310,313]
- Vu, V. Q., Cho, J., Lei, J., and Rohe, K. (2013), "Fantope Projection and Selection: A Near-Optimal Convex Relaxation of Sparse PCA," in *Advances in Neural Information Processing Systems*, Red Hook, New York: Curran Associates Inc. pp. 2670–2678. [308,309,311]
- Vu, V. Q., and Lei, J. (2013), "Minimax Sparse Principal Subspace Estimation in High Dimensions," *The Annals of Statistics*, 41, 2905–2947. [308,312]
- Witten, D. M., Tibshirani, R., and Hastie, T. (2009), "A Penalized Matrix Decomposition, With Applications to Sparse Principal Components and Canonical Correlation Analysis," *Biostatistics*, 10, 515–534. [308]
- Yao, F., Müller, H.-G., and Wang, J.-L. (2005a), "Functional Data Analysis for Sparse Longitudinal Data," *Journal of the American Statistical Association*, 100, 577–590. [309,311]
- (2005b), "Functional Linear Regression Analysis for Longitudinal Data," *The Annals of Statistics*, 33, 2873–2903. [309]
- Yin, J., Geng, Z., Li, R., and Wang, H. (2010), "Nonparametric Covariance Model," *Statistica Sinica*, 20, 469–479. [309]
- Yuan, Y., Zhu, H., Lin, W., and Marron, J. (2012), "Local Polynomial Regression for Symmetric Positive Definite Matrices," *Journal of the Royal Statistical Society, Series B*, 74, 697–719. [309]
- Zhang, X., and Wang, J.-L. (2016), "From Sparse to Dense Functional Data and Beyond," *The Annals of Statistics*, 44, 2281–2321. [312,313]
- Zhu, H., Chen, Y., Ibrahim, J. G., Li, Y., Hall, C., and Lin, W. (2009), "Intrinsic Regression Models for Positive-Definite Matrices with Applications to Diffusion Tensor Imaging," *Journal of the American Statistical Association*, 104, 1203–1212. [309]
- Zou, H., Hastie, T., and Tibshirani, R. (2006), "Sparse Principal Component Analysis," *Journal of Computational & Graphical Statistics*, 15, 265–286. [308]