

A SPARSE DECOMPOSITION OF LOW RANK SYMMETRIC POSITIVE SEMIDEFINITE MATRICES*

THOMAS Y. HOU[†], QIN LI[‡], AND PENGCHUAN ZHANG[†]

Abstract. Suppose that $A \in \mathbb{R}^{N \times N}$ is symmetric positive semidefinite with rank $K \leq N$. Our goal is to decompose A into K rank-one matrices $\sum_{k=1}^K g_k g_k^T$ where the modes $\{g_k\}_{k=1}^K$ are required to be as sparse as possible. In contrast to eigendecomposition, these sparse modes are not required to be orthogonal. Such a problem arises in random field parametrization where A is the covariance function and is intractable to solve in general. In this paper, we partition the indices from 1 to N into several patches and propose to quantify the sparseness of a vector by the number of patches on which it is nonzero, which is called patchwise sparseness. Our aim is to find the decomposition which minimizes the total patchwise sparseness of the decomposed modes. We propose a domain-decomposition type method, called intrinsic sparse mode decomposition (ISMD), which follows the “local-modes-construction + patching-up” procedure. The key step in the ISMD is to construct local pieces of the intrinsic sparse modes by a joint diagonalization problem. Thereafter, a pivoted Cholesky decomposition is utilized to glue these local pieces together. Optimal sparse decomposition, consistency with different domain decomposition, and robustness to small perturbation are proved under the so-called regular-sparse assumption (see Definition 1.2). We provide simulation results to show the efficiency and robustness of the ISMD. We also compare the ISMD to other existing methods, e.g., eigendecomposition, pivoted Cholesky decomposition, and convex relaxation of sparse principal component analysis [R. LAI, J. LU, AND S. OSHER, *Comm. Math. Sci.*, to appear; V. Q. VU, J. CHO, J. LEI, AND K. ROHE, *Fantope projection and selection: A near-optimal convex relaxation of sparse PCA*, in Proceedings in Advances in Neural Information Processing Systems 26, 2013, pp. 2670–2678].

Key words. intrinsic sparse mode decomposition, principal component analysis, sparse PCA, joint diagonalization, pivoted Cholesky decomposition, matrix factorization

AMS subject classifications. 68Q25, 68R10, 68U05

DOI. 10.1137/16M107760X

1. Introduction. Many problems in science and engineering lead to huge *symmetric* and *positive semidefinite* (PSD) matrices. Often they arise from the discretization of self-adjoint PSD operators or their kernels, especially in the context of data science and partial differential equations.

Consider a symmetric PSD matrix of size $N \times N$, denoted as A . Since N is typically large, this causes serious obstructions when dealing numerically with such problems. Fortunately, in many applications the discretization A is low-rank or approximately low-rank, i.e., there exists $\{\psi_1, \dots, \psi_K\} \subset \mathbb{R}^N$ for $K \ll N$ such that

$$A = \sum_{k=1}^K \psi_k \psi_k^T \quad \text{or} \quad \left\| A - \sum_{k=1}^K \psi_k \psi_k^T \right\|_2 \leq \epsilon,$$

respectively. Here, $\epsilon > 0$ is some small number and $\|A\|_2 = \lambda_{max}(A)$ is the largest eigenvalue of A . To obtain such a low-rank decomposition/approximation of A , the

*Received by the editors May 31, 2016; accepted for publication (in revised form) November 22, 2016; published electronically March 16, 2017.

<http://www.siam.org/journals/mms/15-1/M107760.html>

Funding: This research was supported in part by Air Force MURI grant FA9550-09-1-0613, DOE grant DE-FG02-06ER257, and NSF grants DMS-1318377 and DMS-1159138.

[†]Applied and Computational Mathematics, Caltech, Pasadena, CA 91125 (hou@cms.caltech.edu, pzzhang@cms.caltech.edu).

[‡]Mathematics, University of Wisconsin-Madison, Madison, WI 53705 (qinli@math.wisc.edu).

most natural method is perhaps the eigendecomposition with $\{\psi_k\}_{k=1}^K$ as the eigenvectors corresponding to the largest K eigenvalues of A . An additional advantage of the eigendecomposition is the fact that eigenvectors are orthogonal to each other. However, eigenvectors are typically dense vectors, i.e., every entry is typically nonzero.

For a symmetric PSD matrix A with rank $K \ll N$, the aim of this paper is to find an alternative decomposition

$$(1.1) \quad A = \sum_{k=1}^K g_k g_k^T.$$

Here the number of components is still its rank K , which is optimal, and the modes $\{g_k\}_{k=1}^K$ are required to be as sparse as possible. In this paper, we work on the symmetric PSD matrices, which are typically the discretized self-adjoint PSD operators or their kernels. We could have just as well worked on the self-adjoint PSD operators. This would correspond to the case when $N = \infty$. Much of what will be discussed below applies equally well to this case.

Symmetric PSD matrices/operators/kernels appear in many science and engineering branches, and various efforts have been made to seek sparse modes. In statistics, sparse principal component analysis (PCA) and its convex relaxations [20, 45, 8, 39] are designed to sparsify the eigenvectors of data covariance matrices. In quantum chemistry, Wannier functions [41, 22] and other methods [33, 42, 32, 36, 24] have been developed to obtain a set of functions that approximately span the eigenspace of the Hamiltonian, but are spatially localized or sparse. In numerical homogenization of elliptic equations with rough coefficients [13, 14, 9, 35, 34], a set of multiscale basis functions is constructed to approximate the eigenspace of the elliptic operator and is used as the finite element basis to solve the equation. In most cases, sparse modes reduce the computational cost for further scientific experiments. Moreover, in some cases sparse modes have a better physical interpretation compared to the global eigenmodes. Therefore, it is of practical importance to obtain sparse (localized) modes.

1.1. Our results. The number of nonzero entries of a vector $\psi \in \mathbb{R}^N$ is called its l^0 norm, denoted by $\|\psi\|_0$. Since the modes in (1.1) are required to be as sparse as possible, the sparse decomposition problem is naturally formulated as the following optimization problem:

$$(1.2) \quad \boxed{\min_{\psi_1, \dots, \psi_K \in \mathbb{R}^N} \sum_{k=1}^K \|\psi_k\|_0 \quad \text{subject to (s.t.)} \quad A = \sum_{k=1}^K \psi_k \psi_k^T.}$$

However, this problem is rather difficult to solve because: first, minimizing l^0 norm results in a combinatorial problem and is computationally intractable in general; second, the number of unknown variables is $K \times N$, where N is typically a huge number. Therefore, we introduce the following patchwise sparseness as a surrogate of $\|\psi_k\|_0$ and make the problem computationally tractable.

DEFINITION 1.1 (patchwise sparseness). *Suppose that $\mathcal{P} = \{P_m\}_{m=1}^M$ is a disjoint partition of the N nodes, i.e., $[N] \equiv \{1, 2, 3, \dots, N\} = \sqcup_{m=1}^M P_m$. The patchwise sparseness of $\psi \in \mathbb{R}^N$ with respect to the partition \mathcal{P} , denoted by $s(\psi; \mathcal{P})$, is defined as*

$$s(\psi; \mathcal{P}) = \#\{P \in \mathcal{P} : \psi|_P \neq \mathbf{0}\}.$$

Throughout this paper, $[N]$ denotes the index set $\{1, 2, 3, \dots, N\}$; $\mathbf{0}$ denotes the vectors with all entries equal to 0; $|P|$ denotes the cardinality of a set P ; and $\psi|_P \in \mathbb{R}^{|P|}$ denotes the restriction of $\psi \in \mathbb{R}^N$ on patch P . Once the partition \mathcal{P} is fixed, smaller $s(\psi; \mathcal{P})$ means that ψ is nonzero on fewer patches, which implies a sparser vector. With patchwise sparseness as a surrogate of the l^0 norm, the sparse decomposition problem (1.2) is relaxed to

$$(1.3) \quad \min_{\psi_1, \dots, \psi_K \in \mathbb{R}^N} \sum_{k=1}^K s(\psi_k; \mathcal{P}) \quad \text{s.t.} \quad A = \sum_{k=1}^K \psi_k \psi_k^T.$$

If $\{g_k\}_{k=1}^K$ is an optimizer for (1.3), we call them a set of *intrinsic sparse modes* for A under partition \mathcal{P} . Since the objective function of problem (1.3) only takes nonnegative integer values, we know that for a symmetric PSD matrix A with rank K , there exists at least one set of intrinsic sparse modes.

It is obvious that the intrinsic sparse modes depend on the domain partition \mathcal{P} . Two extreme cases would be $M = N$ and $M = 1$. For $M = N$, $s(\psi; \mathcal{P})$ recovers $\|\psi\|_0$, and the patchwise sparseness minimization problem (1.3) recovers the original l^0 minimization problem (1.2). Unfortunately, it is computationally intractable. For $M = 1$, every nonzero vector has sparseness one, and thus the number of nonzero entries makes no difference. However, in this case problem (1.3) is computationally tractable. For instance, a set of (unnormalized) eigenvectors is one of the optimizers. We are interested in the sparseness defined in between, namely, a partition with a meso-scale patch size. Compared to $\|\psi\|_0$, the meso-scale partition sacrifices some resolution when measuring the support, but makes the optimization (1.3) efficiently solvable. Specifically, problem (1.3) with the following *regular-sparse* partitions enjoys many good properties. These properties enable us to design a very efficient algorithm to solve problem (1.3).

DEFINITION 1.2 (regular-sparse partition). *The partition \mathcal{P} is regular-sparse with respect to A if there exists a decomposition $A = \sum_{k=1}^K g_k g_k^T$ such that all nonzero modes on each patch P_m are linearly independent.*

If two intrinsic sparse modes are nonzero on exactly the same set of patches, which are called unidentifiable modes in Definition 3.6, it is easy to see that any rotation of these unidentifiable modes forms another set of intrinsic sparse modes. From a theoretical point of view, if a partition is regular-sparse with respect to A , the intrinsic sparse modes are unique up to rotations of unidentifiable modes; see Theorem 3.7. Moreover, as the partition gets refined, the original identifiable intrinsic sparse modes remain unchanged, while the original unidentifiable modes become identifiable and become sparser (in the sense of l^0 norm); see Theorem 3.9. In this sense, the intrinsic sparse modes are independent of the partition that we use. From a computational point of view, a regular-sparse partition ensures that the restrictions of the intrinsic sparse modes on each patch P_m can be constructed from rotations of local eigenvectors. Following this idea, we propose the intrinsic sparse mode decomposition (ISMD); see Algorithm 1. In Theorem 3.7, we have proved that the ISMD solves problem (1.3) exactly on regular-sparse partitions. We point out that even when the partition is not regular-sparse, numerical experiments show that the ISMD still generates a sparse decomposition of A .

The ISMD consists of three steps. In the first step, we perform eigendecomposition of A restricted on local patches $\{P_m\}_{m=1}^M$, denoted as $\{A_{mm}\}_{m=1}^M$, to get $A_{mm} = H_m H_m^T$. Here, columns of H_m are the unnormalized local eigenvectors of A on patch

P_m . In the second step, we recover the local pieces of intrinsic sparse modes, denoted by G_m , by rotating the local eigenvectors $G_m = H_m D_m$. The method to find the right local rotations $\{D_m\}_{m=1}^M$ is the core of the ISMD. All the local rotations are coupled by the decomposition constraint $A = \sum_{k=1}^K g_k g_k^T$ and it seems impossible to solve $\{D_m\}_{m=1}^M$ from this big coupled system. Surprisingly, when the partition is regular-sparse, this coupled system can be decoupled, and every local rotation D_m can be solved independently by a joint diagonalization problem (2.6). In the last “patch-up” step, we identify correlated local pieces across different patches by the pivoted Cholesky decomposition of a symmetric PSD matrix Ω and then glue them into a single intrinsic sparse mode. Here, Ω is the projection of A onto the subspace spanned by all the local pieces $\{G_m\}_{m=1}^M$; see (2.8). This step is necessary to reduce the number of decomposed modes to the optimal K , i.e., the rank of A . The last step also equips the ISMD with the power to identify long range correlation and to honor the intrinsic correlation structure hidden in A . The popular l^1 approach typically does not have this property.

The ISMD has very low computational complexity. There are two reasons for its efficiency: first, instead of computing the expensive global eigendecomposition, we compute only the local eigendecompositions of $\{A_{mm}\}_{m=1}^M$; second, there is an efficient algorithm to solve the joint diagonalization problems for the local rotations $\{D_m\}_{m=1}^M$. Moreover, because both performing the local eigendecompositions and solving the joint diagonalization problems can be done independently on each patch, the ISMD is embarrassingly parallelizable.

The stability of the ISMD is also explored when the input data A is mixed with noises. We study the small perturbation case, i.e., $\hat{A} = A + \epsilon \tilde{A}$. Here, A is the noiseless rank- K symmetric PSD matrix, \tilde{A} is the symmetric additive perturbation, and $\epsilon > 0$ quantifies the noise level. A simple thresholding step is introduced in the ISMD to achieve our aim: *to clean up the noise $\epsilon \tilde{A}$ and to recover the intrinsic sparse modes of A* . Under some assumptions, we can prove that sparse modes $\{\hat{g}_k\}_{k=1}^K$, produced by the ISMD with thresholding, exactly capture the supports of A 's intrinsic sparse modes $\{g_k\}_{k=1}^K$, and the error $\|\hat{g}_k - g_k\|$ is small; see section 4.1 for a precise description.

We have verified all the theoretical predictions with numerical experiments on several synthetic covariance matrices of high dimensional random vectors. Without parallel execution, for partitions with a large range of patch sizes, the computational cost of the ISMD is comparable to that of the partial eigendecomposition [37, 27]. For certain partitions, the ISMD could be ten times faster than the partial eigendecomposition. We have also implemented the convex relaxation of sparse PCA [24, 39] and compared these two methods. It turns out that the convex relaxation of sparse PCA fails to capture the long range correlation, needs to perform (partial) eigendecomposition on matrices repeatedly for many times, and is thus much slower than the ISMD. Moreover, we demonstrate the robustness of the ISMD on partitions which are not regular-sparse and on inputs which are polluted with small noises.

1.2. Applications. The ISMD leads to a sparse-orthogonal matrix factorization for any matrix. Given a matrix $X \in \mathbb{R}^{N \times M}$ of rank K and a partition \mathcal{P} of the index set $[N]$, the ISMD tries to solve the following optimization problem:

(1.4)

$$\min_{\substack{g_1, \dots, g_K \in \mathbb{R}^N \\ u_1, \dots, u_K \in \mathbb{R}^M}} \sum_{k=1}^K s(g_k; \mathcal{P}) \quad \text{s.t.} \quad X = \sum_{k=1}^K g_k u_k^T, \quad u_k^T u_{k'} = \delta_{k,k'} \quad \forall 1 \leq k, k' \leq K,$$

where $s(g_k; \mathcal{P})$ is the patchwise sparseness defined in Definition 1.1. Compared to the biorthogonal property of SVD, the ISMD requires orthogonality only in one dimension and requires sparsity in the other dimension. The method to obtain the decomposition (1.4) consists of three steps: first, compute $A = XX^T$; second, apply the ISMD to A to get $\{g_k\}_{k=1}^K$; third, project X onto $\{g_k\}_{k=1}^K$ to obtain $\{u_k\}_{k=1}^K$.

The sparse-orthogonal matrix factorization (1.4) has potential applications in statistics, machine learning, and uncertainty quantification. In statistics and machine learning, latent factor models with sparse loadings have found many applications ranging from DNA microarray analysis [11], facial and object recognition [40], web search models [1], etc. Specifically, latent factor models decompose a data matrix $X \in \mathbb{R}^{N \times M}$ by the product of the loading matrix $G \in \mathbb{R}^{N \times K}$ and the factor value matrix $U \in \mathbb{R}^{M \times K}$, with possibly small noise $E \in \mathbb{R}^{N \times M}$, i.e.,

$$(1.5) \quad X = GU^T + E.$$

The sparse-orthogonal matrix factorization (1.4) tries to find the optimal sparse loadings G under the condition that latent factors are *normalized* and *uncorrelated*, i.e., columns in U are orthonormal. In practice, the uncorrelated latent factors make lots of sense, but are not guaranteed by many existing matrix factorization methods, e.g., nonnegative matrix factorization (NMF) [25], sparse PCA [20, 45, 8], or structured sparse PCA [18].

In uncertainty quantification (UQ), we often need to parametrize a random field, denoted as $\kappa(x, \omega)$, with a finite number of random variables. Applying the ISMD to its covariance function, denoted by $\text{Cov}(x, y)$, we can get a parametrization with K random variables:

$$(1.6) \quad \kappa(x, \omega) = \bar{\kappa}(x) + \sum_{k=1}^K g_k(x) \eta_k(\omega),$$

where $\bar{\kappa}(x)$ is the mean field, the physical modes $\{g_k\}_{k=1}^K$ are sparse/localized, and the random variables $\{\eta_k\}_{k=1}^K$ are *centered*, *uncorrelated*, and have *unit variance*. The parametrization (1.6) has a form similar to the widely used Karhunen–Loève (KL) expansion [21, 28], but in the KL expansion the physical modes $\{g_k\}_{k=1}^K$ are eigenfunctions of the covariance function and are typically nonzero everywhere. Obtaining a sparse parametrization is important to uncover the intrinsic sparse feature in a random field and to achieve computational efficiency for further scientific experiments. In [15], such sparse parametrization methods are used to design efficient algorithms to solve partial differential equations with random inputs.

1.3. Connection with the sparse matrix factorization problem. Given a matrix $X \in \mathbb{R}^{N \times M}$ of M columns corresponding to M observations in \mathbb{R}^N , a sparse matrix factorization problem is to find a matrix $G = [g_1, \dots, g_r] \in \mathbb{R}^{N \times r}$, called the *dictionary*, and a matrix $U = [u_1, \dots, u_r] \in \mathbb{R}^{M \times r}$, called *decomposition coefficients*, such that GU^T approximates X well and the columns in G are sparse.

In [26, 43, 31], the authors formulated this problem as an optimization problem by penalizing the l1 norm of G , i.e., $\|G\|_1 := \sum_{k=1}^r \|g_k\|_1$, to enforce the sparsity of the dictionary. This can be written as

$$(1.7) \quad \boxed{\min_{G \in \mathbb{R}^{N \times r}, U \in \mathbb{R}^{M \times r}} \|X - GU^T\|_F^2 + \lambda \|G\|_1 \quad \text{s.t.} \quad \|u_k\|_2 \leq 1 \quad \forall 1 \leq k \leq r,}$$

where the parameter $\lambda > 0$ controls to what extent the dictionary G is regularized. We point out that the l1 penalty can be replaced by other penalties. For example,

the structured sparse PCA [18] uses certain l1/l2 norms of G to enforce sparsity with specific structures, e.g., rectangular structure on a grid. Problem (1.7) is not jointly convex in (G, U) . Certain specially designed algorithms have been developed to solve this optimization problem. We will discuss one of these methods in section 2.3.

There are two major differences between the optimization problem (1.4) and the optimization problem (1.7). First, the ISMD, which is designed to solve (1.4), requires that the decomposition coefficients U be orthonormal, while many other methods, including sparse PCA and structured sparse PCA, which are designed to solve (1.7), only normalize every column in U . One needs to decide whether the orthogonality in U is necessary in her application and choose the appropriate method. Second, the number of modes K in the ISMD must be the rank of the matrix, while the number of modes r in problem (1.7) is picked by users and can be any number. In other words, the ISMD is seeking an *exact matrix decomposition*, while other methods make a trade-off between the accuracy $\|X - GU^T\|_F$ and the sparsity $\|G\|_1$ by recovering the matrix approximately instead of obtaining an exact recovery. Although the ISMD can be modified to do matrix approximation (with the orthogonality constraint on U ; see Algorithm 3), the optimal sparsity of the dictionary G is no longer guaranteed. *Based on these two differences, we recommend the ISMD for sparse matrix factorization problems where the orthogonality in decomposition coefficients U is required and an exact (or nearly exact) decomposition is desired.* In our upcoming papers [16, 17], we will present our recent results on solving problem (1.7).

1.4. Outline. In section 2 we present our ISMD algorithm for low-rank matrices, analyze its computational complexity, and talk about its relation with other methods for sparse decomposition or approximation. In section 3 we present our main theoretical results, i.e., Theorem 3.7. In section 4, we discuss the stability of the ISMD by performing perturbation analysis. We also provide two modified ISMD algorithms: Algorithm 2 for low-rank matrix with small noise, and Algorithm 3 for sparse matrix approximation. Finally, we present a few numerical examples in section 5 to demonstrate the efficiency of the ISMD and compare its performance with other existing methods.

2. Intrinsic sparse mode decomposition. In this section, we present the algorithm of the ISMD and analyze its computational complexity. Its relation with other matrix decomposition methods is discussed at the end of this section. In the rest of the paper, $\mathbb{O}(n)$ denotes the set of real unitary matrices of size $n \times n$; \mathbb{I}_n denotes the identity matrix with size $n \times n$.

2.1. The algorithm of ISMD. Suppose that we have one symmetric positive symmetric matrix, denoted as $A \in \mathbb{R}^{N \times N}$, and a partition of the index set $[N]$, denoted as $\mathcal{P} = \{P_m\}_{m=1}^M$. The partition typically originates from the physical meaning of the matrix A . For example, if A is the discretized covariance function of a random field on domain $D \subset \mathbb{R}^d$, \mathcal{P} is constructed from certain domain partition of D . The submatrix of A , with row index in P_m and column index in P_n , is denoted as A_{mn} . To simplify our notation, we assume that indices in $[N]$ are rearranged such that A is written as below:

$$(2.1) \quad A = \begin{bmatrix} A_{11} & A_{12} & \cdots & A_{1M} \\ A_{21} & A_{22} & \cdots & A_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ A_{M1} & A_{M2} & \cdots & A_{MM} \end{bmatrix}.$$

Notice that when implementing the ISMD, there is no need to rearrange the indices as above. The ISMD tries to find the optimal sparse decomposition of A with respect to (w.r.t.) partition \mathcal{P} , defined as the minimizer of problem (1.3). The ISMD consists of three steps: local decomposition, local rotation, and global patch-up.

In the first step, we perform eigendecomposition

$$(2.2) \quad A_{mm} = \sum_{i=1}^{K_m} \gamma_{m,i} h_{m,i} h_{m,i}^T \equiv H_m H_m^T,$$

where K_m is the rank of A_{mm} and $H_m = [\gamma_{m,1}^{1/2} h_{m,1}, \gamma_{m,2}^{1/2} h_{m,2}, \dots, \gamma_{m,K_m}^{1/2} h_{m,K_m}]$. If A_{mm} is ill-conditioned, we truncate the small eigenvalues, and a truncated eigendecomposition is used as follows:

$$(2.3) \quad A_{mm} \approx \sum_{i=1}^{K_m} \gamma_{m,i} h_{m,i} h_{m,i}^T \equiv H_m H_m^T.$$

Let $K_{(t)} \equiv \sum_{m=1}^M K_m$ be the total local rank of A . We extend columns of H_m into \mathbb{R}^N by adding zeros, and get the block diagonal matrix

$$H_{ext} = \text{diag}\{H_1, H_2, \dots, H_M\}.$$

The correlation matrix with basis H_{ext} , denoted by $\Lambda \in \mathbb{R}^{K_{(t)} \times K_{(t)}}$, is such that

$$(2.4) \quad A = H_{ext} \Lambda H_{ext}^T.$$

Since columns of H_{ext} are orthogonal and span a space that contains $\text{range}(A)$, Λ exists and can be computed blockwisely as follows:

$$(2.5) \quad \Lambda = \begin{bmatrix} \Lambda_{11} & \Lambda_{12} & \cdots & \Lambda_{1M} \\ \Lambda_{21} & \Lambda_{22} & \cdots & \Lambda_{2M} \\ \vdots & \vdots & \ddots & \vdots \\ \Lambda_{M1} & \Lambda_{M2} & \cdots & \Lambda_{MM} \end{bmatrix}, \quad \Lambda_{mn} = H_m^\dagger A_{mn} (H_n^\dagger)^T \in \mathbb{R}^{K_m \times K_n},$$

where $H_m^\dagger \equiv (H_m^T H_m)^{-1} H_m^T$ is the (Moore–Penrose) pseudoinverse of H_m .

In the second step, on every patch P_m , we solve the following joint diagonalization problem to find a local rotation D_m :

$$(2.6) \quad \min_{V \in \mathbb{O}(K_m)} \sum_{n=1}^M \sum_{i \neq j} |(V^T \Sigma_{n;m} V)_{i,j}|^2,$$

in which

$$(2.7) \quad \Sigma_{n;m} \equiv \Lambda_{mn} \Lambda_{mn}^T.$$

We rotate the local eigenvectors with D_m and get $G_m = H_m D_m$. Again, we extend columns of G_m into \mathbb{R}^N by adding zeros, and get the block diagonal matrix

$$G_{ext} = \text{diag}\{G_1, G_2, \dots, G_M\}.$$

The correlation matrix with basis G , denoted by $\Omega \in \mathbb{R}^{K(t) \times K(t)}$, is such that

$$(2.8) \quad A = G_{ext} \Omega G_{ext}^T.$$

With Λ in hand, Ω can be obtained as follows:

$$(2.9) \quad \Omega = D^T \Lambda D, \quad D = \text{diag}\{D_1, D_2, \dots, D_M\}.$$

Joint diagonalization has been well studied in the blind source separation (BSS) community. We present some relevant theoretical results in supplementary material SM2. A Jacobi-like algorithm [5, 2], see Algorithm SM1 in the supplementary material, is used in our paper to solve problem (2.6). For most cases, we may want to normalize the columns of G_{ext} and put all the magnitude information in Ω , i.e.,

$$(2.10) \quad G_{ext} = \bar{G}_{ext} E, \quad \bar{\Omega} = E \Omega E^T,$$

where E is a diagonal matrix, with E_{ii} being the l^2 norm of the i th column of G_{ext} , \bar{G}_{ext} , and $\bar{\Omega}$ will substitute the roles of G and Ω in the rest of the algorithm.

In the third step, we use the pivoted Cholesky decomposition to patch up the local pieces G_m . Specifically, suppose the pivoted Cholesky decomposition of Ω is given as

$$(2.11) \quad \Omega = P L L^T P^T,$$

where $P \in \mathbb{R}^{K(t) \times K(t)}$ is a permutation matrix and $L \in \mathbb{R}^{K(t) \times K}$ is a lower triangular matrix with positive diagonal entries. Since A has rank K , both Λ and Ω have rank K . This is why L only has K nonzero columns. However, we point out that the rank K is automatically identified in the algorithm instead of given as an input parameter. Finally, A is decomposed as

$$(2.12) \quad A = G G^T \equiv G_{ext} P L (G_{ext} P L)^T.$$

The columns in G ($G_{ext} P L$) are our decomposed sparse modes.

The full algorithm is summarized in Algorithm 1. We point out that there are two extreme cases for the ISMD:

- The coarsest partition $\mathcal{P} = \{[N]\}$. In this case, the ISMD is equivalent to the standard eigendecomposition.
- The finest partition $\mathcal{P} = \{\{i\} : i \in [N]\}$. In this case, the ISMD is equivalent to the pivoted Cholesky factorization on \bar{A} where $\bar{A}_{ij} = \frac{A_{ij}}{\sqrt{A_{ii} A_{jj}}}$. If the normalization (2.10) is applied, the ISMD is equivalent to the pivoted Cholesky factorization of A in this case.

In these two extreme cases, there is no need to use the joint diagonalization step, and it is known that, in general, neither the ISMD nor the pivoted Cholesky decomposition generates sparse decomposition. When \mathcal{P} is neither of these two extreme cases, the joint diagonalization is applied to rotate the local eigenvectors, and thereafter the generated modes are patchwise sparse. Specifically, when the partition is regular-sparse, the ISMD generates the optimal patchwise sparse decomposition as stated in Theorem 3.7.

Remark 2.1. One can interpret H_m as the patchwise amplitude and D_m as the patchwise phase. The patchwise amplitude is easy to obtain using a local eigendecomposition (2.2), while the patchwise phase is obtained by the joint diagonalization (2.6).

Algorithm 1 Intrinsic sparse mode decomposition.

Require: $A \in \mathbb{R}^{N \times N}$: symmetric and PSD; $\mathcal{P} = \{P_m\}_{m=1}^M$: partition of index set $[N]$

Ensure: $G = [g_1, g_2, \dots, g_K]$: K is the rank of A , $A = GG^T$

- 1: \triangleright Local eigendecomposition
 - 2: **for** $m = 1, 2, \dots, M$ **do**
 - 3: Local eigendecomposition: $A_{mm} = H_m H_m^T$
 - 4: **end for**
 - 5: \triangleright Assemble correlation matrix Λ
 - 6: Assemble $\Lambda = H_{ext}^\dagger A (H_{ext}^\dagger)^T$ blockwisely as in (2.5)
 - 7: \triangleright Joint Diagonalization
 - 8: **for** $m = 1, 2, \dots, M$ **do**
 - 9: **for** $n = 1, 2, \dots, M$ **do**
 - 10: $\Sigma_{n;m} = \Lambda_{mn} \Lambda_{mn}^T$
 - 11: **end for**
 - 12: Solve the joint diagonalization problem (2.6) for D_m \triangleright Use Algorithm SM1 in the supplementary material
 - 13: **end for**
 - 14: \triangleright Assemble correlation matrix Ω and its pivoted Cholesky decomposition
 - 15: $\Omega = D^T \Lambda D$
 - 16: $\Omega = PLL^T P^T$
 - 17: \triangleright Assemble the intrinsic sparse modes G
 - 18: $G = H_{ext} DPL$
-

In fact, the ISMD solves the following optimization problem where we jointly diagonalize A_{mn} :

$$(2.13) \quad \begin{array}{l} \min_{G_m \in \mathbb{R}^{|P_m| \times K_m}} \sum_{n=1}^M \sum_{i \neq j} |B_{n;m}(i, j)|^2 \\ \text{s.t.} \quad G_m G_m^T = A_{mm}, \\ \quad \quad G_m B_{n;m} G_m^T = A_{mn} A_{nn}^\dagger A_{mn}^T, \end{array}$$

in which $A_{nn}^\dagger = \sum_{i=1}^{K_n} \gamma_{n,i}^{-1} h_{n,i} h_{n,i}^T$ is the (Moore–Penrose) pseudoinverse of A_{nn} . Equation (2.13) is not a unitary joint diagonalization problem, i.e., the variable G_m is not unitary. The ISMD solves this nonunitary joint diagonalization problem in two steps:

1. Perform a local eigendecomposition $A_{mm} = H_m H_m^T$. Then the feasible G_m can be written as $H_m D_m$ with a unitary matrix D_m .
2. Find the rotation D_m that solves the unitary joint diagonalization problem (2.6).

2.2. Computational complexity. The main computational cost of the ISMD comes from the local KL expansion, the joint diagonalization, and the pivoted Cholesky decomposition. To simplify the analysis, we assume that the partition \mathcal{P} is uniform, i.e., each group has $\frac{N}{M}$ nodes. On each patch, we perform eigendecomposition of A_{mm}

of size N/M and rank K_m . Then, the cost of the local eigendecomposition step is

$$\text{Cost}_1 = \sum_{m=1}^M \mathcal{O}((N/M)^2 K_m) = (N/M)^2 \mathcal{O}\left(\sum_{m=1}^M K_m\right).$$

For the joint diagonalization, the computational cost of Algorithm SM1 in the supplementary material is

$$\sum_{m=1}^M N_{corr,m} K_m^3 N_{iter,m}.$$

Here, $N_{corr,m}$ is the number of nonzero matrices in $\{\Sigma_{n;m}\}_{n=1}^M$. Notice that $\Sigma_{n;m} \equiv \Lambda_{mn} \Lambda_{mn}^T = 0$ if and only if $A_{mn} = 0$. Therefore, $N_{corr,m}$ may be much smaller than M if A is sparse. Nevertheless, we take an upper bound M to estimate the cost. $N_{corr,m} K_m^3$ is the computational cost for each sweeping in Algorithm SM1 in the supplementary material, and $N_{iter,m}$ is the number of iterations needed for the convergence. The asymptotic convergence rate is shown to be quadratic [2], and we see no more than six iterations needed in our numerical examples. Therefore, we can take $N_{iter,m} = \mathcal{O}(1)$ and in total we have

$$\text{Cost}_2 = \sum_{m=1}^M M \mathcal{O}(K_m^3) = M \mathcal{O}\left(\sum_{m=1}^M K_m^3\right).$$

Finally, the pivoted Cholesky decomposition of Ω , which is of size $\sum_{k=1}^M K_m$, has cost

$$\text{Cost}_3 = \mathcal{O}\left(\left(\sum_{k=1}^M K_m\right) K^2\right) = K^2 \mathcal{O}\left(\sum_{m=1}^M K_m\right).$$

Combining the computational costs in all three steps, we conclude that the total computational cost of the ISMD is

$$(2.14) \quad \text{Cost}_{\text{ISMD}} = ((N/M)^2 + K^2) \mathcal{O}\left(\sum_{m=1}^M K_m\right) + M \mathcal{O}\left(\sum_{m=1}^M K_m^3\right).$$

Making use of $K_m \leq K$, we have an upper bound for $\text{Cost}_{\text{ISMD}}$,

$$(2.15) \quad \text{Cost}_{\text{ISMD}} \leq \mathcal{O}(N^2 K/M) + \mathcal{O}(M^2 K^3).$$

When $M = \mathcal{O}((N/K)^{2/3})$, $\text{Cost}_{\text{ISMD}} \leq \mathcal{O}(N^{4/3} K^{5/3})$. Compared to the cost of partial eigendecomposition [37, 27], which is about $\mathcal{O}(N^2 K)$,¹ the ISMD is more efficient for low-rank matrices.

For matrix A , which has a sparse decomposition, the local ranks K_m are much smaller than its global rank K . An extreme case is $K_m = \mathcal{O}(1)$, which is, in fact, true for many random fields; see [7, 15]. In this case,

$$(2.16) \quad \text{Cost}_{\text{ISMD}} = \mathcal{O}(N^2/M) + \mathcal{O}(M^2) + \mathcal{O}(MK^2).$$

¹The cost can be reduced to $\mathcal{O}(N^2 \log(K))$ if a randomized SVD with some specific technique is applied.

When the partition gets finer (M increases), the computational cost first decreases due to the saving of local eigendecompositions. The computational cost achieves its minimum around $M = \mathcal{O}(N^{2/3})$ and then increases due to the increasing cost for the joint diagonalization. This trend is observed in our numerical examples; see Figure 4.

We point out that the M local eigendecompositions (2.2) and the joint diagonalization problems (2.6) are solved independently on different patches. Therefore, our algorithm is embarrassingly parallelizable. This will save the computational cost in the first two steps by a factor of M , which makes the ISMD even faster.

2.3. Connection with other matrix decomposition methods. Sparse decompositions of symmetric PSD matrices have been studied in different fields for a long time. There are, in general, two approaches to achieve sparsity: rotation and L^1 minimization.

The rotation approach begins with eigenvectors. Suppose that we have decided to retain and rotate K eigenvectors. Define $H = [h_1, h_2, \dots, h_K]$, with h_k being the k th eigenvector. We postmultiply H by a matrix $T \in \mathbb{R}^{K \times K}$ to obtain the rotated modes $G = [g_1, g_2, \dots, g_K] = HT$. The choice of T is determined by the rotation criterion we use. In data science, for the commonly used varimax rotation criterion [23, 19], T is an orthogonal matrix chosen to maximize the variance of squared modes within each column of G . This drives entries in G towards 0 or ± 1 . In quantum chemistry, every column in H and G corresponds to a function over a physical domain D , and certain specialized sparse modes—localized modes—are sought after. The most widely used criterion to achieve maximally localized modes is proposed in [33]. This criterion requires T to be unitary, and then minimizes the second moment:

$$(2.17) \quad \sum_{k=1}^K \int_D (x - x_k)^2 |g_k(x)|^2 dx,$$

where $x_k = \int_D x |g_k(x)|^2 dx$. More recently, a method weighted by higher degree polynomials is discussed in [42]. While these criteria work reasonably well for simple symmetric PSD functions/operators, they all suffer from nonconvex optimization—which requires a good starting point to converge to the global minimum. In addition, these methods only care about the eigenspace spanned by H instead of the specific matrix decomposition, and thus they cannot be directly applied to solve our problem (1.3).

The ISMD proposed in this paper follows the rotation approach. The ISMD implicitly finds a unitary matrix $T \in \mathbb{R}^{K \times K}$ to construct the intrinsic sparse modes

$$(2.18) \quad [g_1, g_2, \dots, g_K] = \left[\sqrt{\lambda_1} h_1, \sqrt{\lambda_2} h_2, \dots, \sqrt{\lambda_K} h_K \right] T.$$

Notice that we rotate the unnormalized eigenvector $\sqrt{\lambda_k} h_k$ to satisfy the decomposition constraint $A = \sum_{k=1}^K g_k g_k^T$. The criterion of the ISMD is to minimize the total patchwise sparseness as in (1.3). The success of the ISMD lies in the fact that as long as the domain partition is regular-sparse, the optimization problem (1.3) can be exactly and efficiently solved by Algorithm 1. Moreover, the intrinsic sparse modes produced by the ISMD are optimally localized because we are directly minimizing the total patchwise sparseness of $\{g_k\}_{k=1}^K$.

The L^1 minimization approach, pioneered by ScotLass [20], has a rich literature about solving the sparse matrix factorization problem (1.7); see [45, 8, 44, 39, 36, 24]. Problem (1.7) is highly nonconvex in (G, U) , and there has been a lot of effort (see, e.g., [8, 39, 24]) in relaxing it to a convex optimization. First, since there are no essential

constraints on U , one can get rid of U by considering the variational form [20, 45, 36]:

$$(2.19) \quad \boxed{\min_{G \in \mathbb{R}^{N \times K}} -\text{Tr}(G^T A G) + \mu \|G\|_1 \quad \text{s.t.} \quad G^T G = \mathbb{I}_K,}$$

where $A = X X^T$ is the covariance matrix as in the ISMD (1.3), and Tr is the trace operator on square matrices. Notice that the problem is still nonconvex due to the orthogonality constraint $G^T G = \mathbb{I}_K$. In the second step, the authors in [39] proposed the following semidefinite programming to obtain the sparse density matrix $W \in \mathbb{R}^{n \times n}$, which plays the same role as $G G^T$ in (2.19):

$$(2.20) \quad \boxed{\min_{W \in \mathbb{R}^{N \times N}} -\text{Tr}(A W) + \mu \|W\|_1 \quad \text{s.t.} \quad 0 \preceq W \preceq \mathbb{I}_N, \text{Tr}(W) = K.}$$

Here, $0 \preceq W \preceq \mathbb{I}_N$ means that both W and $\mathbb{I}_N - W$ are symmetric and positive semidefinite. Finally, the first K eigenvectors of W are used as the sparse modes G . An equivalent formulation was proposed in [24], and the authors proposed to pick K columns of W as the sparse modes G .

We will compare the advantages and disadvantages of the ISMD and the convex relaxation of sparse PCA in sections 5.2 and 5.6.

3. Theoretical results with regular-sparse partitions. In this section, we present the main theoretical results of the ISMD, i.e., Theorems 3.7 and 3.9 and its perturbation analysis. We first introduce a domain-decomposition type presentation of any feasible decomposition $A = \sum_{k=1}^K \psi_k \psi_k^T$. Then we discuss the regular-sparse property and use it to prove our main results. When no ambiguity arises, we denote patchwise sparseness $s(g_k; \mathcal{P})$ as s_k .

3.1. A domain-decomposition type presentation. For an arbitrary decomposition $A = \sum_{k=1}^K \psi_k \psi_k^T$, denote $\Psi \equiv [\psi_1, \dots, \psi_K]$ and $\Psi|_{P_m} \equiv [\psi_1|_{P_m}, \dots, \psi_K|_{P_m}]$. For a sparse decomposition, we expect that most columns in $\Psi|_{P_m}$ are zero, and thus we define the local dimension on patch P_m as follows.

DEFINITION 3.1 (local dimension). *The local dimension of a decomposition $A = \sum_{k=1}^K \psi_k \psi_k^T$ on patch P_m is the number of nonzero modes when restricted to this patch, i.e.,*

$$d(P_m; \Psi) = |S_m|, \quad S_m = \{k : \psi_k|_{P_m} \neq 0\}.$$

When no ambiguity arises, $d(P_m; \Psi)$ is written as d_m . We enumerate all the elements in S_m as $\{k_i^m\}_{i=1}^{d_m}$, and group together all the nonzero local pieces on patch P_m and obtain

$$(3.1) \quad \Psi_m \equiv [\psi_{m,1}, \dots, \psi_{m,d_m}], \quad \psi_{k_i^m}|_{P_m} = \psi_{m,i}.$$

Therefore, we have

$$(3.2) \quad \Psi|_{P_m} = \Psi_m L_m^{(\psi)},$$

where $L_m^{(\psi)}$ is a matrix of size $d_m \times K$ with the k_i^m th column being e_i for $i \in [d_m]$ and other columns being $\mathbf{0}$. Here, e_i is the i th column of \mathbb{I}_{d_m} . $L_m^{(\psi)}$ is called the *local indicator matrix* of Ψ on patch P_m . Restricting the decomposition constraint $A = \Psi \Psi^T$ to patch P_m , we have $A_{mm} = \Psi|_{P_m} (\Psi|_{P_m})^T$, where A_{mm} is the restriction

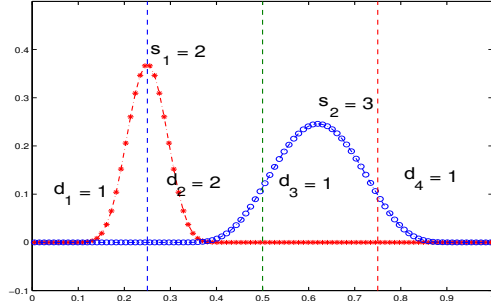


FIG. 1. Illustration of sparseness, local dimension, and $\Psi = \Psi_{ext}L^{(\psi)}$.

of A on patch P_m , as in (2.1). Since Ψ_m is obtained from $\Psi|_{P_m}$ by deleting zero columns, we have

$$(3.3) \quad A_{mm} = \Psi_m \Psi_m^T.$$

We stack up Ψ_m and $L_m^{(\psi)}$ as follows:

$$\Psi_{ext} \equiv \text{diag}\{\Psi_1, \Psi_2, \dots, \Psi_M\}, \quad L^{(\psi)} \equiv [L_1^{(\psi)}; L_2^{(\psi)}; \dots; L_M^{(\psi)}],$$

and then we have

$$(3.4) \quad \Psi = [\Psi|_{P_1}; \dots; \Psi|_{P_M}] = \Psi_{ext}L^{(\psi)}.$$

The intuition in (3.4) is that the local pieces Ψ_m are linked together by the indicator matrix $L^{(\psi)}$, and the modes Ψ on the entire domain $[N]$ can be recovered from Ψ_{ext} and $L^{(\psi)}$. We call $L^{(\psi)}$ the *indicator matrix* of Ψ .

We use a simple example to illustrate the patchwise sparseness, the local dimension, and (3.4). In this case, $\Psi \in \mathbb{R}^{N \times K}$ ($N = 100, K = 2$) is the discretized version of two functions on $[0, 1]$, and \mathcal{P} partitions $[0, 1]$ uniformly into four intervals as shown in Figure 1. ψ_1 , the red (online only) starred mode, is nonzero on the left two patches and ψ_2 , the blue (online only) circled mode, is nonzero on the right three patches. The sparseness of ψ_1 is 2, the sparseness of ψ_2 is 3, and the local dimensions of the four patches are 1, 2, 1, and 1, respectively, as we comment in Figure 1. Following the definitions above, we have $\Psi_1 = \psi_1|_{P_1}$, $L_1^{(\psi)} = [1, 0]$, $\Psi_2 = [\psi_1|_{P_2}, \psi_2|_{P_2}]$, $L_2^{(\psi)} = [1, 0; 0, 1]$, $\Psi_3 = \psi_2|_{P_3}$, $L_3^{(\psi)} = [0, 1]$, $\Psi_4 = \psi_2|_{P_4}$, and $L_4^{(\psi)} = [0, 1]$. Finally, we get

$$[\psi_1, \psi_2] = \Psi_{ext}L^{(\psi)} \equiv \begin{bmatrix} \psi_{1,1} & 0 & 0 & 0 \\ 0 & \psi_{1,2} & \psi_{2,2} & 0 \\ 0 & 0 & 0 & \psi_{2,3} \\ 0 & 0 & 0 & \psi_{2,4} \end{bmatrix} \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}.$$

With this domain-decomposition type representation of Ψ , the decomposition constraint is rewritten as

$$(3.5) \quad A = \Psi \Psi^T = \Psi_{ext} \Omega^{(\psi)} \Psi_{ext}^T, \quad \Omega^{(\psi)} \equiv L^{(\psi)} \left(L^{(\psi)} \right)^T.$$

Here, $\Omega^{(\psi)}$ has a role similar to that of Ω in the ISMD. It can be viewed as the correlation matrix of A under basis Ψ_{ext} , just like how Λ and Ω are defined.

Finally, we provide two useful properties of the local indicator matrices $L_m^{(\psi)}$, which are direct consequences of their definitions. Its proof is elementary and can be found in Appendix A.

PROPOSITION 3.2. *For an arbitrary decomposition $A = \Psi\Psi^T$:*

1. *The k th column of $L^{(\psi)}$, denoted as $l_k^{(\psi)}$, satisfies $\|l_k^{(\psi)}\|_1 = s_k$, where s_k is the patchwise sparseness of ψ_k , as in Definition 1.1. Moreover, different columns in $L^{(\psi)}$ have disjoint supports.*
2. *Define*

$$(3.6) \quad B_{n;m}^{(\psi)} \equiv \Omega_{mn}^{(\psi)} \left(\Omega_{mn}^{(\psi)} \right)^T,$$

where $\Omega_{mn}^{(\psi)} \equiv L_m^{(\psi)} (L_n^{(\psi)})^T$ is the (m, n) th block of $\Omega^{(\psi)}$. $B_{n;m}^{(\psi)}$ is diagonal with diagonal entries either 1 or 0. Moreover, $B_{n;m}^{(\psi)}(i, i) = 1$ if and only if there exists $k \in [K]$ such that $\psi_k|_{P_m} = \psi_{m,i}$ and $\psi_k|_{P_n} \neq \mathbf{0}$.

Since different columns in $L^{(\psi)}$ have disjoint supports, $\Omega^{(\psi)} \equiv L^{(\psi)} (L^{(\psi)})^T$ has a block-diagonal structure with K blocks. The k th diagonal block is the one contributed by $l_k^{(\psi)} (l_k^{(\psi)})^T$. Therefore, as long as we obtain $\Omega^{(\psi)}$, we can use the pivoted Cholesky decomposition to efficiently recover $L^{(\psi)}$. The ISMD follows this rationale: we first construct local pieces $\Psi_{ext} \equiv \text{diag}\{\Psi_1, \Psi_2, \dots, \Psi_M\}$ for a certain set of intrinsic sparse modes Ψ . Then from the decomposition constraint (3.5) we are able to compute $\Omega^{(\psi)}$. Finally, the pivoted Cholesky decomposition is applied to obtain $L^{(\psi)}$, and the modes are assembled by $\Psi = \Psi_{ext} L^{(\psi)}$. Obviously, the key step is to construct Ψ_{ext} , which are local pieces of a set of intrinsic sparse modes—this is exactly where the regular-sparse property and the joint diagonalization come into play.

3.2. Regular-sparse property and local modes construction. In this and the next subsection (sections 3.2 and 3.3), we assume that the submatrices A_{mm} are well conditioned and thus the exact local eigendecomposition (2.2) is used in the ISMD.

Combining the local eigendecomposition (2.2) and local decomposition constraint (3.3), there exists $D_m^{(\psi)} \in \mathbb{R}^{K_m \times d_m}$ such that

$$(3.7) \quad \Psi_m = H_m D_m^{(\psi)}.$$

Moreover, since the local eigenvectors are linearly independent, we have

$$(3.8) \quad d_m \geq K_m, \quad D_m^{(\psi)} \left(D_m^{(\psi)} \right)^T = \mathbb{I}_{K_m}.$$

We see that $d_m = K_m$ if and only if columns in Ψ_m are also linearly independent. In this case, $D_m^{(\psi)}$ is unitary, i.e., $D_m^{(\psi)} \in \mathbb{O}(K_m)$. This is exactly what is required by the regular-sparse property; see Definition 1.2. It is easy to see that we have the following equivalent definitions of regular-sparse property.

PROPOSITION 3.3. *The following assertions are equivalent.*

1. *The partition \mathcal{P} is regular-sparse w.r.t. A .*
2. *There exists a decomposition $A = \sum_{k=1}^K \psi_k \psi_k^T$ such that on every patch P_m its local dimension d_m is equal to the local rank K_m , i.e., $d_m = K_m$.*

3. The minimum of problem (1.3) is $\sum_{m=1}^M K_m$.

The proof is elementary and omitted here. By Proposition 3.3, for regular-sparse partitions local pieces of a set of intrinsic sparse modes can be constructed from rotating local eigenvectors, i.e., $\Psi_m = H_m D_m^{(\psi)}$. All of the local rotations $\{D_m^{(\psi)}\}_{m=1}^M$ are coupled by the decomposition constraint $A = \Psi \Psi^T$. At first glance, it seems impossible to find such D_m from this big coupled system. However, the following lemma gives a necessary condition that $D_m^{(\psi)}$ must satisfy so that $H_m D_m^{(\psi)}$ are local pieces of a set of intrinsic sparse modes. More importantly, this necessary condition turns out to be sufficient, and thus provides a criterion to find the local rotations.

LEMMA 3.4. *Suppose that \mathcal{P} is regular-sparse w.r.t. A and $\{\psi_k\}_{k=1}^K$ is an arbitrary set of intrinsic sparse modes. Denote the transformation from H_m to Ψ_m as $D_m^{(\psi)}$, i.e., $\Psi_m = H_m D_m^{(\psi)}$. Then $D_m^{(\psi)}$ is unitary and jointly diagonalizes $\{\Sigma_{n;m}\}_{n=1}^M$, which are defined in (2.7). Specifically, we have*

$$(3.9) \quad B_{n;m}^{(\psi)} = \left(D_m^{(\psi)}\right)^T \Sigma_{n;m} D_m^{(\psi)}, \quad m = 1, 2, \dots, M,$$

where $B_{n;m}^{(\psi)} \equiv \Omega_{mn}^{(\psi)} \left(\Omega_{mn}^{(\psi)}\right)^T$, defined in (3.6), is diagonal with diagonal entries either 0 or 1.

Proof. From item 3 in Proposition 3.3, any set of intrinsic sparse modes must have local dimension $d_m = K_m$ on patch P_m . Therefore, the transformation $D_m^{(\psi)}$ from H_m to Ψ_m must be unitary. Combining $\Psi_m = H_m D_m^{(\psi)}$ with the decomposition constraint (3.5), we get

$$A = H_{ext} D^{(\psi)} \Omega^{(\psi)} \left(D^{(\psi)}\right)^T H_{ext},$$

where $D^{(\psi)} = \text{diag}\{D_1^{(\psi)}, D_2^{(\psi)}, \dots, D_M^{(\psi)}\}$. Recall that $A = H_{ext} \Lambda H_{ext}$ and that H_{ext} has linearly independent columns. We obtain

$$(3.10) \quad \Lambda = D^{(\psi)} \Omega^{(\psi)} \left(D^{(\psi)}\right)^T,$$

or blockwisely,

$$(3.11) \quad \Lambda_{mn} = D_m^{(\psi)} \Omega_{mn}^{(\psi)} \left(D_n^{(\psi)}\right)^T.$$

Since $D_n^{(\psi)}$ is unitary, (3.9) naturally follows the definitions of $B_{n;m}^{(\psi)}$ and $\Sigma_{n;m}$. By item 2 in Proposition 3.2, we know that $B_{n;m}^{(\psi)}$ is diagonal with diagonal entries either 0 or 1. \square

Lemma 3.4 guarantees that $D_m^{(\psi)}$ for an arbitrary set of intrinsic sparse modes is the minimizer of the joint diagonalization problem (2.6). In the other direction, the following lemma guarantees that any minimizer of the joint diagonalization problem (2.6), denoted as D_m , transforms local eigenvectors H_m to G_m , which are the local pieces of certain intrinsic sparse modes.

LEMMA 3.5. *Suppose that \mathcal{P} is regular-sparse w.r.t. A and that D_m is a minimizer of the joint diagonalization problem (2.6). As in the ISMD, define $G_m = H_m D_m$. Then there exists a set of intrinsic sparse modes such that its local pieces on patch P_m are equal to G_m .*

Before we prove this lemma, we examine the uniqueness property of intrinsic sparse modes. It is easy to see that permutations and sign flips of a set of intrinsic sparse modes are still a set of intrinsic sparse modes. Specifically, if $\{\psi_k\}_{k=1}^K$ is a set of intrinsic sparse modes and $\sigma : [K] \rightarrow [K]$ is a permutation, $\{\pm\psi_{\sigma(k)}\}_{k=1}^K$ is another set of intrinsic sparse modes. Another kind of nonuniqueness comes from the following concept—identifiability.

DEFINITION 3.6 (identifiability). *For two modes $g_1, g_2 \in \mathbb{R}^N$, they are unidentifiable on partition \mathcal{P} if they are supported on the same patches, i.e., $\{P \in \mathcal{P} : g_1|_P \neq \mathbf{0}\} = \{P \in \mathcal{P} : g_2|_P \neq \mathbf{0}\}$. Otherwise, they are identifiable. For a collection of modes $\{g_i\}_{i=1}^k \subset \mathbb{R}^N$, they are unidentifiable if and only if any pair of them are unidentifiable. They are pairwise identifiable if and only if any pair of them are identifiable.*

It is important to point out that the identifiability above is based on the resolution of partition \mathcal{P} . Unidentifiable modes for partition \mathcal{P} may have different supports and become identifiable on a refined partition. Unidentifiable intrinsic sparse modes lead to another kind of nonuniqueness for intrinsic sparse modes. For instance, when two intrinsic sparse modes ψ_m and ψ_n are unidentifiable, then any rotation of $[\psi_m, \psi_n]$, while keeping other intrinsic sparse modes unchanged, is still a set of intrinsic sparse modes.

Local pieces of intrinsic sparse modes inherit this kind of nonuniqueness. Suppose $\Psi_m \equiv [\psi_{m,1}, \dots, \psi_{m,d_m}]$ are the local pieces of a set of intrinsic sparse modes Ψ on patch P_m . First, if $\sigma : [d_m] \rightarrow [d_m]$ is a permutation, $\{\pm\psi_{m,\sigma(i)}\}_{i=1}^{d_m}$ are local pieces of another set of intrinsic sparse modes. Second, if $\psi_{m,i}$ and $\psi_{m,j}$ are the local pieces of two unidentifiable intrinsic sparse modes, then any rotation of $[\psi_{m,i}, \psi_{m,j}]$, while keeping other local pieces unchanged, are local pieces of another set of intrinsic sparse modes. It turns out that this kind of nonuniqueness has a one-to-one correspondence with the nonuniqueness of joint diagonalizers for problem (2.6), which is characterized in Theorem SM2.1 in the supplementary material. Keeping this correspondence in mind, the proof of Lemma 3.5 is quite intuitive.

Proof of Lemma 3.5. Let $\Psi \equiv [\psi_1, \dots, \psi_K]$ be an arbitrary set of intrinsic sparse modes. We order columns in Ψ such that unidentifiable modes are grouped together, denoted as $\Psi = [\Psi_1, \dots, \Psi_Q]$, where Q is the number of unidentifiable groups. Accordingly on patch P_m , $\Psi_m = [\Psi_{m,1}, \dots, \Psi_{m,Q_m}]$, where Q_m is the number of nonzero unidentifiable groups. Denote the number of columns in each group as $n_{m,i}$, i.e., there are $n_{m,i}$ modes in $\{\psi_k\}_{k=1}^K$ that are nonzero and unidentifiable on patch P_m .

Making use of item 2 in Proposition 3.2, one can check that $\psi_{m,i}$ and $\psi_{m,j}$ are unidentifiable if and only if $B_{n;m}^{(\psi)}(i, i) = B_{n;m}^{(\psi)}(j, j)$ for all $n \in [M]$. Since unidentifiable pieces in Ψ_m are grouped together, the same diagonal entries in $\{B_{n;m}^{(\psi)}\}_{n=1}^M$ are grouped together as required in Theorem SM2.1 in the supplementary material. Now we apply Theorem SM2.1 in the supplementary material with M_k replaced by $\Sigma_{n;m}$, Λ_k replaced by $B_{n;m}^{(\psi)}$, D replaced by $D_m^{(\psi)}$, the number of distinct eigenvalues m replaced by Q_m , eigenvalue's multiplicity q_i replaced by $n_{m,i}$, and the diagonalizer V replaced by D_m . Therefore, there exists a permutation matrix Π_m and a block diagonal matrix V_m such that

$$(3.12) \quad D_m \Pi_m = D_m^{(\psi)} V_m, \quad V_m = \text{diag}\{V_{m,1}, \dots, V_{m,Q_m}\}.$$

Recalling that $G_m = H_m D_m$ and $\Psi_m = H_m D_m^{(\psi)}$, we obtain that

$$(3.13) \quad G_m \Pi_m = \Psi_m V_m = [\Psi_{m,1} V_{m,1}, \dots, \Psi_{m,Q_m} V_{m,Q_m}].$$

From (3.13), we can see that identifiable pieces are completely separated, and the small rotation matrices, $V_{m,i}$, only mix unidentifiable pieces $\Psi_{m,i}$. Π_m merely permutes the columns in G_m . From the nonuniqueness of local pieces of intrinsic sparse modes, we conclude that G_m are local pieces of another set of intrinsic sparse modes. \square

We point out that the local pieces $\{G_m\}_{m=1}^M$ constructed by the ISMD on different patches may correspond to different sets of intrinsic sparse modes. Therefore, the final “patch-up” step should further modify and connect them to build a set of intrinsic sparse modes. Fortunately, the pivoted Cholesky decomposition elegantly solves this problem.

3.3. Optimal sparse recovery and consistency of the ISMD. As defined in the ISMD, Ω is the correlation matrix of A with basis G_{ext} ; see (2.8). If Ω enjoys a block diagonal structure with each block corresponding to a single intrinsic sparse mode, just like $\Omega^{(\psi)} \equiv L^{(\psi)} (L^{(\psi)})^T$, the pivoted Cholesky decomposition can be utilized to recover the intrinsic sparse modes.

It is fairly easy to see that Ω indeed enjoys such a block diagonal structure when there is one set of intrinsic sparse modes that are pairwise identifiable. Denoting this identifiable set as $\{\psi_k\}_{k=1}^K$ (only its existence is needed), by (3.12), we know that on patch P_m there is a permutation matrix Π_m and a diagonal matrix V_m with diagonal entries either 1 or -1 such that $D_m \Pi_m = D_m^{(\psi)} V_m$. Recall that $\Lambda = D \Omega D^T = D^{(\psi)} \Omega^{(\psi)} (D^{(\psi)})^T$; see (2.9) and (3.11). We then have

$$(3.14) \quad \Omega = D^T D^{(\psi)} \Omega^{(\psi)} (D^{(\psi)})^T D = \Pi V^T \Omega^{(\psi)} V \Pi^T,$$

in which $V = \text{diag}\{V_1, \dots, V_m\}$ is diagonal with diagonal entries either 1 or -1 , and $\Pi = \text{diag}\{\Pi_1, \dots, \Pi_m\}$ is a permutation matrix. Since the action of ΠV^T does not change the block diagonal structure of $\Omega^{(\psi)}$, Ω still has such a structure, and the pivoted Cholesky decomposition can be readily applied. In fact, the action of ΠV^T exactly corresponds to the column permutation and sign flips of intrinsic sparse modes, which is the only kind of nonuniqueness of problem (1.3) that occurs when the intrinsic sparse modes are pairwise identifiable. For the general case when there are unidentifiable intrinsic sparse modes, Ω still has the block diagonal structure with each block corresponding to a group of unidentifiable modes, resulting in the following theorem.

THEOREM 3.7. *Suppose the domain partition \mathcal{P} is regular-sparse w.r.t. A . Let $A = GG^T$ be the decomposition given by the ISMD (2.12), and let $\Psi \equiv [\psi_1, \dots, \psi_K]$ be an arbitrary set of intrinsic sparse modes. Let columns in Ψ be ordered such that unidentifiable modes are grouped together, denoted as $\Psi = [\Psi_1, \dots, \Psi_Q]$, where Q is the number of unidentifiable groups and n_q is the number of modes in Ψ_q . Then there exists Q rotation matrices $U_q \in \mathbb{R}^{n_q \times n_q}$ ($1 \leq q \leq Q$) such that*

$$(3.15) \quad G = [\Psi_1 U_1, \dots, \Psi_Q U_Q],$$

with reordering of columns in G if necessary. It immediately follows that

- the ISMD generates one set of intrinsic sparse modes, and
- the intrinsic sparse modes are unique up to permutations and rotations within unidentifiable modes.

Proof. By (3.12), equation (3.14) still holds true with block diagonal V_m for $m \in [M]$. Without loss of generality, we assume that $\Pi = I$ since permutation does not

change the block diagonal structure that we desire. Then from (3.14) we have

$$(3.16) \quad \Omega = V^T \Omega^{(\psi)} V = V^T L^{(\psi)} \left(L^{(\psi)} \right)^T V.$$

In terms of blockwise formulation, we get

$$(3.17) \quad \Omega_{mn} = V_m^T \Omega_{mn}^{(\psi)} V_n = V_m^T L_m^{(\psi)} \left(L_n^{(\psi)} \right)^T V_n.$$

Correspondingly, by (3.13) the local pieces satisfy

$$G_m = [G_{m,1}, \dots, G_{m,Q_m}] = [\Psi_{m,1} V_{m,1}, \dots, \Psi_{m,Q_m} V_{m,Q_m}].$$

Now, we prove that Ω has the block diagonal structure in which each block corresponds to a group of unidentifiable modes. Specifically, $G_{m,i} = \Psi_{m,i} V_{m,i}$ and $G_{n,j} = \Psi_{n,j} V_{n,j}$ are two identifiable groups, i.e., $\Psi_{m,i}$ and $\Psi_{n,j}$ are from two identifiable groups, and we want to prove that the corresponding block in Ω , denoted as $\Omega_{m,i;n,j}$, is zero. From (3.17), one gets $\Omega_{m,i;n,j} = V_{m,i}^T L_{m,i}^{(\psi)} \left(L_{n,j}^{(\psi)} \right)^T V_{n,j}$, where $L_{m,i}^{(\psi)}$ are the rows in $L_m^{(\psi)}$ corresponding to $\Psi_{m,i}$. $L_{n,j}^{(\psi)}$ is defined similarly. Due to identifiability between $\Psi_{m,i}$ and $\Psi_{n,j}$, we know $L_{m,i}^{(\psi)} \left(L_{n,j}^{(\psi)} \right)^T = 0$, and thus we obtain the block diagonal structure of Ω .

In (2.11), the ISMD performs the pivoted Cholesky decomposition $\Omega = PLL^T P^T$ and generates sparse modes $G = G_{ext} PL$. Due to the block diagonal structure in Ω , every column in PL can only have nonzero entries on local pieces that are not identifiable. Therefore, columns in G have identifiable intrinsic sparse modes completely separated and unidentifiable intrinsic sparse modes rotated (including sign flip) by certain unitary matrices. Therefore, G is a set of intrinsic sparse modes.

Due to the arbitrary choice of Ψ , we know that the intrinsic sparse modes are unique to permutations and rotations within unidentifiable modes. \square

Remark 3.8. From the proof above, we can see that it is the block diagonal structure of Ω that leads to the recovery of intrinsic sparse modes. The pivoted Cholesky decomposition is one way to explore this structure. In fact, the pivoted Cholesky decomposition can be replaced by any other matrix decomposition that preserves this block diagonal structure, for instance, the eigendecomposition if there is no degeneracy.

Despite the fact that the intrinsic sparse modes depend on the partition \mathcal{P} , the following theorem guarantees that the solutions to problem (1.3) give consistent results as long as the partition is regular-sparse.

THEOREM 3.9. *Suppose that \mathcal{P}_c is a partition, \mathcal{P}_f is a refinement of \mathcal{P}_c , and \mathcal{P}_f is regular-sparse. Suppose $\{g_k^{(c)}\}_{k=1}^K$ and $\{g_k^{(f)}\}_{k=1}^K$ (with reordering if necessary) are the intrinsic sparse modes produced by the ISMD on \mathcal{P}_c and \mathcal{P}_f , respectively. Then for every $k \in \{1, 2, \dots, K\}$, $g_k^{(c)}$ and $g_k^{(f)}$ are supported on the same patches in the coarse partition \mathcal{P}_c , while the support patches of $g_k^{(f)}$ are contained in the support patches of $g_k^{(c)}$ in the fine partition \mathcal{P}_f , i.e.,*

$$\begin{aligned} \{P \in \mathcal{P}_c : g_k^{(f)}|_P \neq \mathbf{0}\} &= \{P \in \mathcal{P}_c : g_k^{(c)}|_P \neq \mathbf{0}\}, \\ \{P \in \mathcal{P}_f : g_k^{(f)}|_P \neq \mathbf{0}\} &\subset \{P \in \mathcal{P}_f : g_k^{(c)}|_P \neq \mathbf{0}\}. \end{aligned}$$

Moreover, if $g_k^{(c)}$ is identifiable on the coarse patch \mathcal{P}_c , it remains unchanged when the ISMD is performed on the refined partition \mathcal{P}_f , i.e., $g_k^{(f)} = \pm g_k^{(c)}$.

Proof. Given that the finer partition \mathcal{P}_f is regular-sparse, it is easy to prove the coarser partition \mathcal{P}_c is also regular-sparse.² Notice that if two modes are identifiable on the coarse partition \mathcal{P}_c , they must be identifiable on the fine partition \mathcal{P}_f . However, the opposite is not true, i.e., unidentifiable modes may become identifiable if the partition is refined. Based on this observation, Theorem 3.9 is a simple corollary of Theorem 3.7. \square

Finally, we provide a necessary condition for a partition to be regular-sparse as follows.

PROPOSITION 3.10. *If \mathcal{P} is regular-sparse w.r.t. A , all eigenvalues of Λ are integers. Here, Λ is computed in the ISMD by (2.5).*

Proof. Let $\{\psi_k\}_{k=1}^K$ be a set of intrinsic sparse modes. Since \mathcal{P} is regular-sparse, $D^{(\psi)}$ in (3.10) is unitary. Therefore, Λ and $\Omega^{(\psi)} \equiv L^{(\psi)} (L^{(\psi)})^T$ share the same eigenvalues. Due to the block-diagonal structure of $\Omega^{(\psi)}$, one can see that

$$\Omega^{(\psi)} \equiv L^{(\psi)} \left(L^{(\psi)} \right)^T = \sum_{k=1}^K l_k^{(\psi)} \left(l_k^{(\psi)} \right)^T$$

is, in fact, the eigendecomposition of $\Omega^{(\psi)}$. The eigenvalue corresponding to the eigenvector $l_k^{(\psi)}$ is $\|l_k^{(\psi)}\|_2^2$, which is also equal to $\|l_k^{(\psi)}\|_1$ because $L^{(\psi)}$ only elements 0 or 1. From item 1 in Proposition 3.2, $\|l_k^{(\psi)}\|_1 = s_k$, which is the patchwise sparseness of ψ_k . \square

Combining Theorems 3.7 and 3.9 and Proposition 3.10, we can develop a hierarchical process that gradually finds the finest regular-sparse partition and thus obtains the sparsest decomposition using the ISMD. This sparsest decomposition can be viewed as another definition of intrinsic sparse modes, which are independent of partitions. In our numerical examples, our partitions are all uniform but with different patch sizes. We see that even when the partition is not regular-sparse, the ISMD still produces a nearly optimal sparse decomposition.

4. Perturbation analysis and two modifications. In real applications, data are often contaminated by noises. For example, when measuring the covariance function of a random field, sample noise is inevitable if a Monte Carlo type sampling method is utilized. A basic requirement for a numerical algorithm is its stability w.r.t. small noises. In section 4.1, under several assumptions, we are able to prove that the ISMD is stable w.r.t. small perturbations in the input A . In section 4.2, we provide two modified ISMD algorithms that effectively handle noise in different situations.

4.1. Perturbation analysis of the ISMD. We consider the additive perturbation here, i.e., \hat{A} is an approximately low-rank symmetric PSD matrix that satisfies

$$(4.1) \quad \hat{A} = A + \epsilon \tilde{A}, \quad \|\tilde{A}\|_2 \leq 1.$$

Here, A is the noiseless rank- K symmetric PSD matrix and \tilde{A} is the symmetric additive perturbation, and $\epsilon > 0$ quantifies the noise level. We divide \tilde{A} into blocks that

²We provide the proof in supplementary materials; see Lemma SM1.1.

are conformal with blocks of A in (2.1) and thus $\widehat{A}_{mn} = A_{mn} + \epsilon \widetilde{A}_{mn}$. In this case, we need to apply the truncated local eigendecomposition (2.3) to capture the correct local rank K_m . Suppose the eigendecomposition of \widehat{A}_{mm} is

$$\widehat{A}_{mm} = \sum_{i=1}^{K_m} \widehat{\gamma}_{m,i} \widehat{h}_{n,i} \widehat{h}_{n,i}^T + \sum_{i>K_m} \widehat{\gamma}_{m,i} \widehat{h}_{n,i} \widehat{h}_{n,i}^T.$$

In this subsection, we assume that the noise level is very small with $\epsilon \ll 1$ such that there is an energy gap between $\widehat{\gamma}_{m,K_m}$ and $\widehat{\gamma}_{m,K_m+1}$. Therefore, the truncation (2.3) captures the correct local rank K_m , i.e.,

$$(4.2) \quad \widehat{A}_{mm} \approx \widehat{A}_{mm}^{(t)} \equiv \sum_{i=1}^{K_m} \widehat{\gamma}_{m,i} \widehat{h}_{n,i} \widehat{h}_{n,i}^T \equiv \widehat{H}_m \widehat{H}_m^T.$$

In the rest of the ISMD, the perturbed local eigenvectors \widehat{H}_m is used as H_m in the noiseless case. We expect that our ISMD is stable w.r.t. this small perturbation and generates slightly perturbed intrinsic sparse modes of A .

To carry out this perturbation analysis, we will restrict ourselves to the case when intrinsic sparse modes of A are pairwise identifiable, and thus it is possible to compare the error between the noisy output \widehat{g}_k with A 's intrinsic sparse mode g_k . When there are unidentifiable intrinsic sparse modes of A , it only makes sense to consider the perturbation of the subspace spanned by those unidentifiable modes, but we will not consider this case in this paper. The following lemma is a preliminary result on the perturbation analysis of local pieces G_m .

LEMMA 4.1. *Suppose that partition \mathcal{P} is regular-sparse w.r.t. A and all intrinsic modes are identifiable with each other. Furthermore, we assume that for all $m \in [M]$ there exists $E_m^{(eig)}$ such that*

$$(4.3) \quad \widehat{A}_{mm}^{(t)} = (I + \epsilon E_m^{(eig)}) A_{mm} (I + \epsilon (E_m^{(eig)})^T) \quad \text{and} \quad \|E_m^{(eig)}\|_2 \leq C_{eig}.$$

Here C_{eig} is a constant depending on A but not on ϵ or \widetilde{A} . Then there exists $E_m^{(jd)} \in \mathbb{R}^{K_m \times K_m}$ such that

$$(4.4) \quad \widehat{G}_m = (I + \epsilon E_m^{(eig)}) G_m (I + \epsilon E_m^{(jd)} + \mathcal{O}(\epsilon^2)) J_m \quad \text{and} \quad \|E_m^{(jd)}\|_F \leq C_{jd},$$

where G_m and \widehat{G}_m are local pieces constructed by the ISMD with input A and \widehat{A} , respectively, J_m is the product of a permutation matrix with a diagonal matrix having only ± 1 on its diagonal, and C_{jd} is a constant depending on A but not on ϵ or \widetilde{A} . Here, $\|\bullet\|_2$ and $\|\bullet\|_F$ are matrix spectral norm and Frobenius norm, respectively.

Lemma 4.1 ensures that local pieces of intrinsic sparse modes can be constructed with $\mathcal{O}(\epsilon)$ accuracy up to permutation and sign flips (characterized by J_m in (4.4)) under several assumptions. The identifiability assumption is necessary. Without such an assumption, these local pieces are not uniquely determined up to permutations and sign flips. The assumption (4.3) holds true when eigendecomposition of A_{mm} is well conditioned, i.e., both eigenvalues and eigenvectors are well conditioned. We expect that a stronger perturbation result is still true without making this assumption. The proof of Lemma 4.1 is an application of perturbation analysis for the joint diagonalization problem [4], and is presented in supplementary material SM3.

Finally, $\widehat{\Omega}$ is the correlation matrix of \widehat{A} with basis $\widehat{G}_{ext} = \text{diag}\{\widehat{G}_1, \widehat{G}_2, \dots, \widehat{G}_M\}$. Specifically, the (m, n) th block of $\widehat{\Omega}$ is given by

$$\widehat{\Omega}_{mn} = \widehat{G}_m^\dagger \widehat{A}_{mn} \left(\widehat{G}_n^\dagger\right)^T.$$

Without loss of generality, we can assume that $J_m = \mathbb{I}_{K_m}$ in (4.4).³ Based on the perturbation analysis of G_m in Lemma 4.1 and the standard perturbation analysis of pseudoinverse (for instance, see Theorem 3.4 in [38]), it is straightforward to get a bound of the perturbations in $\widehat{\Omega}$, i.e.,

$$(4.5) \quad \|\widehat{\Omega} - \Omega\|_2 \leq C_{ismd}\epsilon.$$

Here, C_{ismd} depends on the smallest singular value of G_m and the constants C_{eig} and C_{jd} in Lemma 4.1. Notice that when all intrinsic modes are identifiable with each other, the entries of Ω are either 0 or ± 1 . Therefore, when $C_{ismd}\epsilon$ is small enough, we can exactly recover Ω from $\widehat{\Omega}$ as below:

$$(4.6) \quad \Omega_{ij} = \begin{cases} -1 & \text{for } \widehat{\Omega}_{ij} < -0.5, \\ 0 & \text{for } \widehat{\Omega}_{ij} \in [-0.5, 0.5], \\ 1 & \text{for } \widehat{\Omega}_{ij} > 0.5. \end{cases}$$

Following Algorithm 1, we get the pivoted Cholesky decomposition $\Omega = PLL^T P^T$ and output the perturbed intrinsic sparse modes

$$\widehat{G} = \widehat{G}_{ext} PL.$$

Notice that when the patchwise sparseness information is all coded in L and we can reconstruct L exactly due to the thresholding step (4.6), then \widehat{G} has the same patchwise sparse structure as G . Moreover, because the local pieces \widehat{G}_{ext} are constructed with $\mathcal{O}(\epsilon)$ error, we have

$$(4.7) \quad \|\widehat{G} - G\|_2 \leq C_g \epsilon,$$

where the constant C_g only depends on the constants C_{eig} and C_{jd} in Lemma 4.1.

4.2. Two modified ISMD algorithms. In section 4.1, we have shown that the ISMD is robust to small noises under the assumption of regular sparsity and identifiability. In this section, we provide two modified versions of the ISMD to deal with the cases when these two assumptions fail. The first modification aims at constructing intrinsic sparse modes from noisy input \widehat{A} in the small noise region, as in section (4.1), but it does not require the regular sparsity and identifiability. The second modification aims at constructing a simultaneous low-rank and sparse approximation of \widehat{A} when the noise is big. Our numerical experiments demonstrate that these modified algorithms are quite effective in practice.

4.2.1. ISMD with thresholding. In the general case when unidentifiable pairs of intrinsic sparse modes exist, the thresholding idea (4.6) is still applicable, but the threshold ϵ_{th} should be learned from the data, i.e., the entries in $\widehat{\Omega}$. Specifically, there

³One can check that $\{J_m\}_{m=1}^M$ only affect the sign of recovered intrinsic sparse modes $[\widehat{g}_1, \widehat{g}_2, \dots, \widehat{g}_K]$ if pivoted Cholesky decomposition is applied on $\widehat{\Omega}$.

are $\mathcal{O}(1)$ entries in $\widehat{\Omega}$ corresponding to the slightly perturbed nonzero entries in Ω ; there are also many $\mathcal{O}(\epsilon)$ entries that are contributed by the noise $\epsilon\widetilde{A}$. If the noise level ϵ is small enough, we can see a gap between these two group of entries, and a threshold ϵ_{th} is chosen such that it separates these two groups. A simple 2-cluster algorithm is able to identify the threshold ϵ_{th} . In our numerical examples we draw the histogram of absolute values of entries in $\widehat{\Omega}$, which clearly shows the 2-cluster effect; see Figure 10. Finally, we set all the entries in $\widehat{\Omega}$ with absolute value less than ϵ_{th} to 0. In this approach we do not need to know the noise level ϵ a priori; we just learn the threshold from the data. To modify Algorithm 1 with this thresholding technique, we just need to add one line between assembling Ω (line 15) and the pivoted Cholesky decomposition (line 16); see Algorithm 2.

Algorithm 2 Intrinsic sparse mode decomposition with thresholding.

Require: $A \in \mathbb{R}^{N \times N}$: symmetric and PSD; $\mathcal{P} = \{P_m\}_{m=1}^M$: partition of index set $[N]$

Ensure: $G = [g_1, g_2, \dots, g_K]$: $A \approx GG^T$

- 1: The same as Algorithm 1 from line 1 to line 13
 - 2: \triangleright Assemble Ω , thresholding, and its pivoted Cholesky decomposition
 - 3: $\Omega = D^T \Lambda D$
 - 4: Learn a threshold ϵ_{th} from Ω and set all the entries in Ω with absolute value less than ϵ_{th} to 0
 - 5: $\Omega = PLL^T P^T$
 - 6: \triangleright Assemble the intrinsic sparse modes G
 - 7: $G = H_{ext} D P L$
-

It is important to point out that when the noise is large, the $\mathcal{O}(1)$ entries and $\mathcal{O}(\epsilon)$ entries mix together. In this case, we cannot identify such a threshold ϵ_{th} to separate them, and the assumption that there is an energy gap between $\widehat{\gamma}_{m,K_m}$ and $\widehat{\gamma}_{m,K_m+1}$ is invalid. In the next subsection, we will present the second modified version to overcome this difficulty.

4.2.2. Low-rank approximation with ISMD. In the case when there is no gap between $\widehat{\gamma}_{m,K_m}$ and $\widehat{\gamma}_{m,K_m+1}$ (i.e., no well-defined local ranks), or when the noise is so large that the threshold ϵ_{th} cannot be identified, we modify our ISMD to give a low-rank approximation of $A \approx GG^T$, in which G is observed to be patchwise sparse from our numerical examples.

In this modification, the normalization (2.10) is applied and thus we have

$$A \approx \widetilde{G}_{ext} \widetilde{\Omega} \widetilde{G}_{ext}^T.$$

It is important to point out that $\widetilde{\Omega}$ has the same block diagonal structure as Ω but has different eigenvalues. Specifically, for the case when there is no noise and the regular-sparse assumption holds true, $\widetilde{\Omega}$ has eigenvalues $\{\|g_k\|_2^2\}_{k=1}^K$ for a certain set of intrinsic sparse modes g_k , while Ω has eigenvalues $\{s_k\}_{k=1}^K$ (here s_k is the patchwise sparseness of the intrinsic sparse mode). We first perform eigendecomposition $\widetilde{\Omega} = \widetilde{L}\widetilde{L}^T$ and then assemble the final result by $G = \widetilde{G}_{ext}\widetilde{L}$. The modified algorithm is summarized in Algorithm 3.

Here we replace the pivoted Cholesky decomposition of Ω in Algorithm 1 by eigendecomposition of $\widetilde{\Omega}$. From Remark 3.8, this modified version generates exactly the same result as Algorithm 1 if all the intrinsic sparse modes have different l^2 norm

Algorithm 3 Intrinsic sparse mode decomposition for low-rank approximation.

Require: $A \in \mathbb{R}^{N \times N}$: symmetric and PSD; $\mathcal{P} = \{P_m\}_{m=1}^M$: partition of index set $[N]$

Ensure: $G = [g_1, g_2, \dots, g_K]$: $A \approx GG^T$

- 1: The same as Algorithm 1 from line 1 to line 13
 - 2: \triangleright Assemble Ω , normalization, and its eigendecomposition
 - 3: $\Omega = D^T \Lambda D$
 - 4: $G_{ext} = \bar{G}_{ext} E$, $\bar{\Omega} = E \Omega E^T$ as in (2.10)
 - 5: $\bar{\Omega} = \bar{L} \bar{L}^T$
 - 6: \triangleright Assemble the intrinsic sparse modes G
 - 7: $G = \bar{G}_{ext} \bar{L}$
-

(there are no repeated eigenvalues in $\bar{\Omega}$). The advantage of the pivoted Cholesky decomposition is its low computational cost and the fact that it always exploits the (unordered) block diagonal structure of Ω . However, it is more sensitive to noise compared to eigendecomposition. In contrast, eigendecomposition is much more robust to noise. Moreover, eigendecomposition gives the optimal low-rank approximation of $\bar{\Omega}$. Thus Algorithm 3 gives a more accurate low-rank approximation for A compared to Algorithms 1 and 2 that use the pivoted Cholesky decomposition.

5. Numerical experiments. In this section, we demonstrate the robustness of our intrinsic sparse mode decomposition method and compare its performance with that of the eigendecomposition, the pivoted Cholesky decomposition, and the convex relaxation of sparse PCA. All of our computations are performed using MATLAB R2015a (64-bit) on an Intel Core i7-3770 (3.40 GHz). The pivoted Cholesky decomposition is implemented in MATLAB according to Algorithm 3.1 in [29].

We will use synthetic covariance matrices of a random permeability field, which models the underground porous media, as the symmetric PSD input A . This random permeability model is adapted from the porous media problem [12, 10] where the physical domain D is two dimensional. The basic model has a constant background and several localized features to model the subsurface channels and inclusions, i.e.,

$$(5.1) \quad \kappa(x, \omega) = \kappa_0 + \sum_{k=1}^K \eta_k(\omega) g_k(x), \quad x \in [0, 1]^2,$$

where κ_0 is the constant background, $\{g_k\}_{k=1}^K$ are characteristic functions of channels and inclusions, and η_k are the associated uncorrelated latent variables controlling the permeability of each feature. Here, we have $K = 35$, including 16 channels and 18 inclusions. Among these modes, there is one artificial smiling face mode that has disjoint branches. It is used here to demonstrate that the ISMD is able to capture long range correlation. For this random medium, the covariance function is

$$(5.2) \quad a(x, y) = \sum_{k=1}^K g_k(x) g_k(y), \quad x, y \in [0, 1]^2.$$

Since the length scales of channels and inclusions are very small, with width about $1/32$, we need a fine grid to resolve these small features. Such a fine grid is also needed when we do further scientific experiments [12, 10, 15]. In this paper, the physical domain $D = [0, 1]^2$ is discretized using a uniform grid with $h_x = h_y = 1/96$,

resulting in $A \in \mathbb{R}^{N \times N}$ with $N = 96^2$. One sample of the random field (shown also from a bird's-eye view) and the covariance matrix are plotted in Figure 2. It can be seen that the covariance matrix is sparse and concentrates along the diagonal since modes in the ground-truth media are all localized functions.

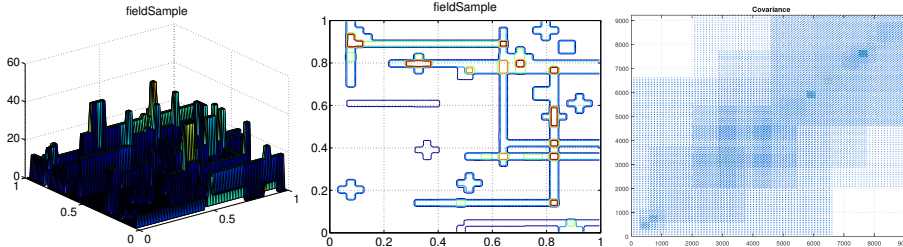


FIG. 2. One sample and the bird's-eye view. The covariance matrix is plotted on the right.

Note that this example is synthetic because we construct A from a sparse decomposition (5.2). We would like to test whether different matrix factorization methods, like eigendecomposition, the Cholesky decomposition, and the ISMD, are able to recover this sparse decomposition, or even find a sparser decomposition for A .

5.1. Numerical results of ISMD. The partitions we take for this example are all uniform domain partition with $H_x = H_y = H$. We run the ISMD with patch sizes $H \in \{1, 1/2, 1/3, 1/4, 1/6, 1/8, 1/12, 1/16, 1/24, 1/32, 1/48, 1/96\}$ in this section. For the coarsest partition $H = 1$, the ISMD is exactly the eigendecomposition of A . For the finest partition $H = 1/96$, the ISMD is equivalent to the pivoted Cholesky factorization on \bar{A} where $\bar{A}_{ij} = \frac{A_{ij}}{\sqrt{A_{ii}A_{jj}}}$. The pivoted Cholesky factorization on A is also implemented. It is no surprise that all of the above methods produce 35 modes. The number of modes is exactly the rank of A . We plot the first six modes for each method in Figure 3. We can see that both the eigendecomposition (ISMD with $H = 1$) and the pivoted Cholesky factorization on A generate modes which mix different localized feathers together. On the other hand, the ISMD with $H = 1/8$ and $H = 1/32$ exactly recover the localized feathers, including the smiling face.

We use Lemma 3.4 to check when the regular-sparse property fails. It turns out that for $H \geq 1/16$ the regular-sparse property holds and for $H \leq 1/24$ it fails. The eigenvalues of Λ 's for $H = 1, 1/8$ and $1/32$ are plotted in Figure 4 on the left side. The eigenvalues of Λ when $H = 1$ are all 1's, since every eigenvector has patchwise sparseness 1 in this trivial case. The eigenvalues of Λ when $H = 1/16$ are all integers, corresponding to patchwise sparseness of the intrinsic sparse modes. The eigenvalues of Λ when $H = 1/32$ are not all integers any more, which indicates that this partition is not regular-sparse w.r.t. A according to Lemma 3.4.

The consistency of the ISMD (Theorem 3.9) manifests from $H = 1$ to $H = 1/8$ in Figure 3. As Theorem 3.9 states, the supports of the intrinsic sparse modes on a coarser partition contain those on a finer partition. In other words, we get sparser modes when we refine the partition as long as the partition is regular-sparse. After checking all the 35 recovered modes, we see that the intrinsic sparse modes get sparser and sparser from $H = 1$ to $H = 1/6$. When $H \leq 1/6$, all of the 35 intrinsic sparse modes are identifiable with each other, and these intrinsic modes remain the same for $H = 1/8, 1/12, 1/16$. When $H \leq 1/24$, the regular-sparse property fails, but we still get the sparsest decomposition (the same decomposition with $H = 1/8$). For

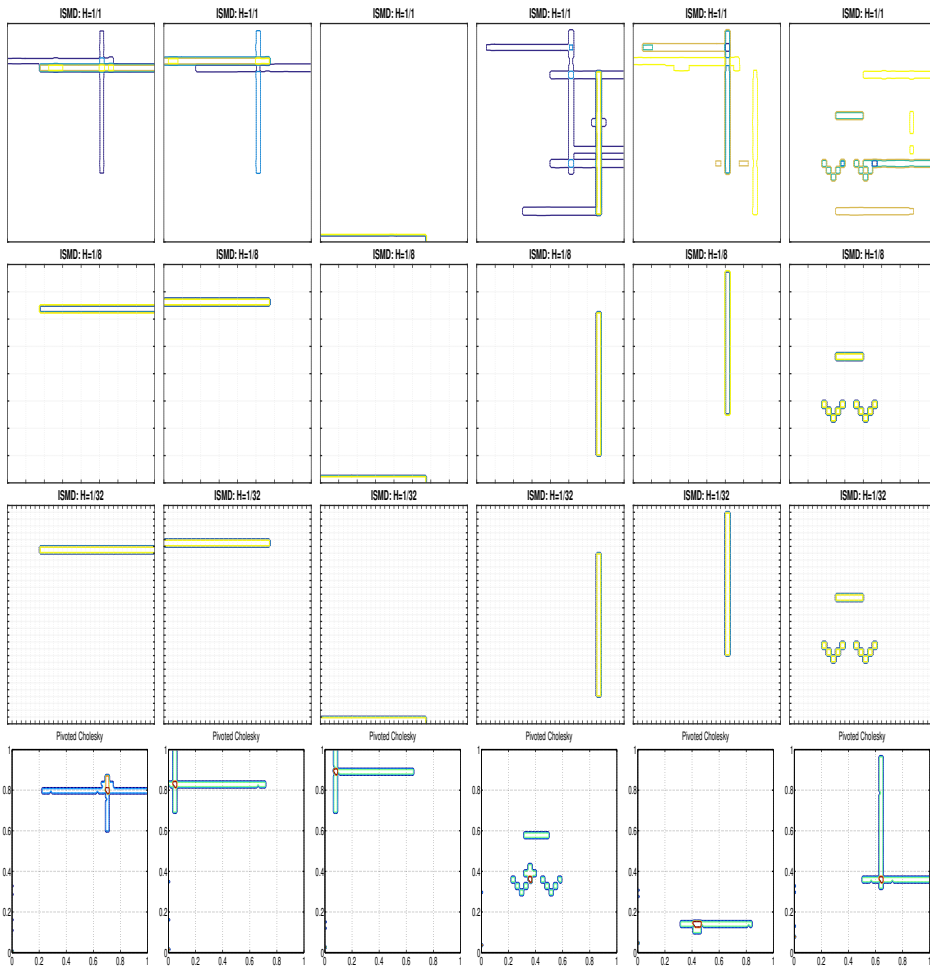


FIG. 3. First six eigenvectors ($H = 1$). First six intrinsic sparse modes ($H = 1/8$, regular-sparse). First six intrinsic sparse modes ($H = 1/32$; not regular-sparse). First six modes from the pivoted Cholesky decomposition of A .

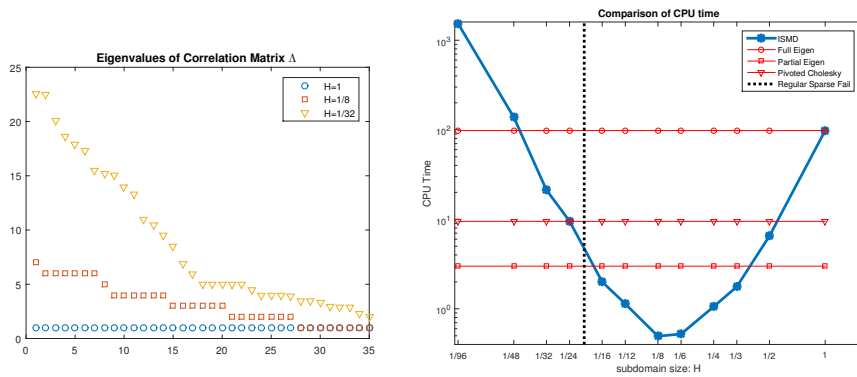


FIG. 4. Left: Eigenvalues of Λ for $H = 1, 1/8, 1/32$. By Lemma 3.4, the partition with $H = 1/32$ is not regular-sparse. Right: CPU time (unit: second) for different partition sizes H .

$H = 1/32$, we recover exactly 33 intrinsic sparse modes but get the other two mixed together. This is not surprising since the partition is not regular-sparse any more. For $H = 1/48$, we exactly recover all of the 35 intrinsic sparse modes again. Table 1 lists the cases when we exactly recover the sparse decomposition (5.2) from which we construct A . From Theorem 3.7, this decomposition is the optimal sparse decomposition (defined by problem (1.3)) for $H \geq 1/16$. We suspect that this decomposition is also optimal in the L^0 sense (defined by problem (1.2)).

TABLE 1
Cases when the ISMD gets exact recovery of the sparse decomposition (5.2).

H	1	1/2	1/3	1/4	1/6	1/8	1/12	1/16	1/24	1/32	1/48	1/96
regular-sparse	✓	✓	✓	✓	✓	✓	✓	✓	✗	✗	✗	✗
exact recovery	✗	✗	✗	✗	✓	✓	✓	✓	✓	✗	✓	✗

The CPU time of the ISMD for different H 's is shown in Figure 4 on the right side. We compare the CPU time for the full eigendecomposition $\mathbf{eig}(A)$, the partial eigendecomposition $\mathbf{eigs}(A, 35)$, and the pivoted Cholesky decomposition. For $1/16 \leq H \leq 1/3$, the ISMD is even faster than the partial eigendecomposition. Specifically, the ISMD is ten times faster for the case $H = 1/8$. Notice that the ISMD performs the local eigendecomposition by \mathbf{eig} in MATLAB, and thus does not need any prior information about the rank K . If we also assume prior information on the local rank K_m , the ISMD would be even faster. The CPU time curve has a V-shape as predicted by our computational estimation (2.16). The cost first decreases as we refine the mesh because the cost of local eigendecompositions decreases. Then it increases as we refine further because there are M joint diagonalization problem (2.6) to be solved. When M is very large, i.e., $H = 1/48$ or $H = 1/96$, the 2 layer for-loops from line 5 to line 10 in Algorithm 1 become extremely slow in MATLAB. When implemented in other languages that have little overhead cost for multiple for-loops, e.g., C or C++, the actual CPU time for $H = 1/96$ would be roughly the same as the CPU time for the pivoted Cholesky decomposition.

5.2. Comparison with the semidefinite relaxation of sparse PCA. In comparison, the semidefinite relaxation of sparse PCA (problem (2.20)) gives poor results in this example. We have tested several values of μ , and found that parameter $\mu = 0.0278$ gives the best performance in the sense that the first 35 eigenvectors of W capture the most variance in A . The first 35 eigenvectors of W , shown in Figure 5, explain 95% of the variance, but all of them mix several intrinsic modes the same way the eigendecomposition does in Figure 3. For this example, it is not clear how to choose the best 35 columns out of all the 9216 columns in W , as proposed in [24]. If columns of W are ordered by l^2 norm in descending order, the first 35 columns can only explain 31.46% of the total variance, although they are indeed localized. Figure 6 shows the first six columns of W with the largest norms.

We also compare the CPU time of the ISMD with that of the semidefinite relaxation of sparse PCA (2.20). The sparse PCA is computed using the split Bregman iteration. Each split Bregman iteration requires an eigendecomposition of a matrix of size $N \times N$. In comparison, the ISMD is cheaper than a single eigendecomposition, as shown in Figure 4. It has been observed that the split Bregman iteration converges linearly. If we set the error tolerance to be $O(\delta)$, the number of iterations needed is about $O(1/\delta)$. In our implementation, we set the error tolerance to be 10^{-3} , and we need to perform 852 iterations. Overall, to solve the convex optimization prob-

lem (2.20) with split Bregman iteration takes over 1000 times more CPU time than the ISMD with $H = 1/8$.

It is expected that the ISMD is much faster than sparse PCA since the sparse PCA needs to perform partial eigendecomposition many times to solve problem (2.20), but the ISMD has computational cost comparable to one single partial eigendecomposition. As we discussed in section 1.3, sparse PCA is designed and works reasonably well for problem (1.7). When sparse PCA is applied to our sparse decomposition problem (1.3), it does not work well. However, it is not always the case that the ISMD gives a sparser and more accurate decomposition of A than sparse PCA. In subsection 5.6, we will present another example in which sparse PCA gives a better performance than the ISMD.

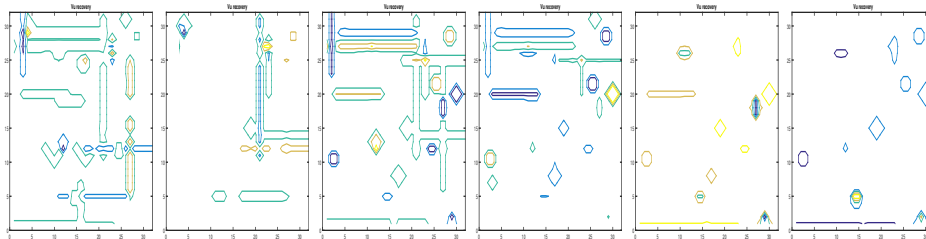


FIG. 5. *Sparse PCA: The first six eigenvectors of W . The first 35 eigenvectors of W explain 95% of the variance.*

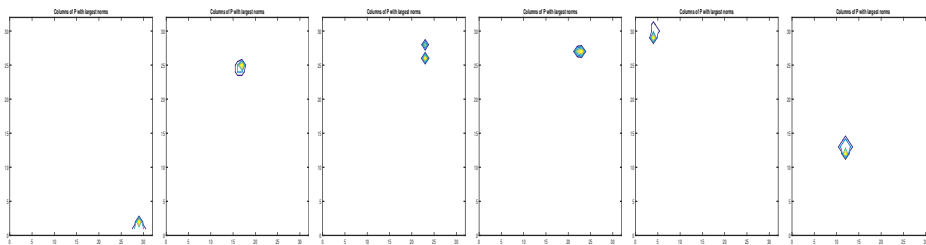


FIG. 6. *Sparse PCA: six columns of W with largest norms. The first 35 columns with largest norms only explain 31.46% of the variance.*

We point out that unlike the structured sparse PCA [18], the ISMD does not take advantage of the specific (rectangular) structure of the physical modes. The “smiling face” mode shows that the ISMD can recover nonconvex and nonlocal sparse modes. Therefore, the ISMD is expected to perform equally well even when there is no such structure known.

5.3. ISMD with small noises. In this subsection we report the test on the robustness of the ISMD. In the following, we perturb the rank-35 covariance matrix $A \in \mathbb{R}^{9216 \times 9216}$ with a random matrix:

$$\hat{A} = A + \epsilon \tilde{A},$$

where ϵ is the noise level and \tilde{A} is a random matrix with independently and identically distributed elements uniformly distributed in $[-1, 1]$. Notice that all elements in A are uniformly bounded by 1, and thus ϵ is a relative noise level. Since all the intrinsic sparse modes are identifiable with each other for the partition with patch

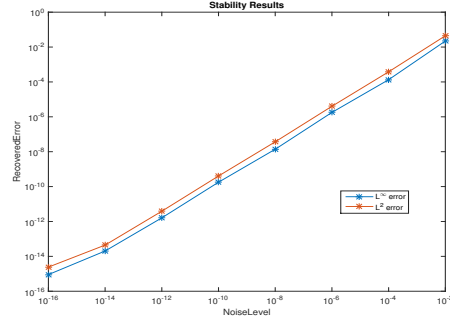


FIG. 7. L^∞ and L^2 error increases linearly as the noise level increases.

size $H = 1/16$, we perform ISMD with simple thresholding (4.6) on \hat{A} to get the perturbed intrinsic sparse modes $\hat{G} \equiv [\hat{g}_1, \dots, \hat{g}_K]$. The l^∞ and l^2 error are defined as below:

$$Err_\infty = \max_{k=1,2,\dots,K} \frac{\|\hat{g}_k - g_k\|_2}{\|g_k\|_2}, \quad Err_2 = \sqrt{\sum_{k=1}^K \frac{\|\hat{g}_k - g_k\|_2^2}{\|g_k\|_2^2}}.$$

Figure 7 shows that Err_∞ and Err_2 depend linearly on the noise level ϵ , which validates our stability analysis in section 4.1.

5.4. Separate global and localized modes with ISMD. In this example, we consider a more sophisticated model in which the media contain several global modes, i.e.,

$$(5.3) \quad \kappa(x, \omega) = \sum_{k=1}^{K_1} \xi_k(\omega) f_k(x) + \sum_{k=1}^{K_2} \eta_k(\omega) g_k(x), \quad x \in [0, 1]^2,$$

where $\{g_k\}_{k=1}^{K_2}$, and η_k models the localized features like channels and inclusions as above, $\{f_k\}_{k=1}^{K_1}$ are functions with support on the entire domain $D = [0, 1]^2$, and ξ_k are the associated latent variables with global influence on the entire domain. Here, we keep the 35 localized features as before, but add two global features with $f_1(x) = \sin(2\pi x_1 + 4\pi x_2)/2$, $f_2(x) = \sin(4\pi x_1 + 2\pi x_2)/2$. ξ_1 and ξ_2 are set to be uncorrelated and have variance 1. For this random medium, the covariance function is

$$(5.4) \quad a(x, y) = \sum_{k=1}^{K_1} f_k(x) f_k(y) + \sum_{k=1}^{K_2} g_k(x) g_k(y), \quad x, y \in [0, 1]^2.$$

As before, we discretize the covariance function with $h_x = h_y = 1/96$ and represent A by a matrix of size 9216×9216 . One sample of the random field (and the bird's-eye view) and the covariance matrix are plotted in Figure 8. It can be seen that the covariance matrix is dense now because we have two global modes.

We apply the ISMD with patch size $H = 1/16$ on A and get 37 intrinsic sparse modes as expected. Moreover, two of them are rotations of $[f_1, f_2]$ and the other 35 are exactly the 35 localized modes in the construction (5.4). We plot the first six intrinsic sparse modes in Figure 9. As we can see, the ISMD separates the global

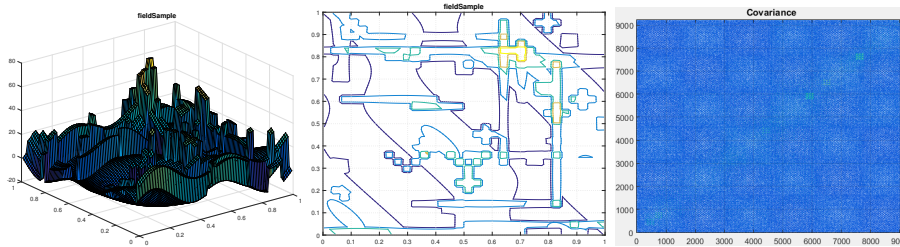


FIG. 8. One sample and the bird's-eye view. The covariance matrix is plotted on the right.

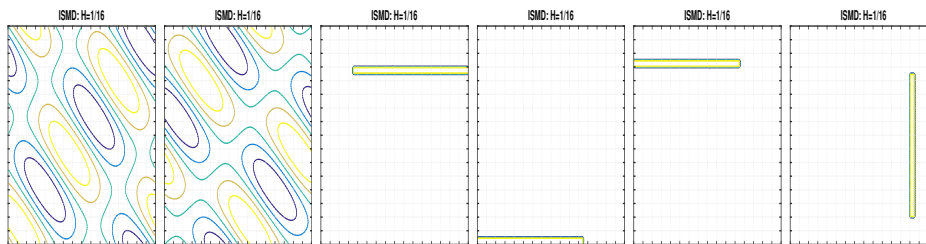


FIG. 9. First six intrinsic sparse modes ($H = 1/16$, regular-sparse)

modes and localized modes in A , or equivalently we separate the low-rank dense part and sparse part of A . The reason why we can achieve this separation is that the representation (5.4) in fact solves the patchwise sparseness minimization problem (1.3). The low-rank-plus-sparse decomposition (also known as Robust PCA; see [6, 3, 30]) can also separate the low-rank dense part and the sparse part in A . However, the computational cost of robust PCA is much more expensive than the ISMD.

5.5. Application of Algorithm 2. When A is constructed from model (5.4) but is mixed with small noises as in section 5.3, we cannot simply apply the thresholding (4.6) any more. In this case, we have unidentifiable modes f_1 and f_2 , and thus Ω may contain nonzero values other than ± 1 . For the noise level $\epsilon = 10^{-6}$, Figure 10 (left) shows the histogram of absolute values of entries in $\hat{\Omega}$. We can clearly see a gap between $\mathcal{O}(\epsilon)$ entries and $\mathcal{O}(1)$ entries from Figure 10(left). Therefore, we choose a threshold $\epsilon_{th} = 10^{-3}$ and apply the modified ISMD algorithm 2 on \hat{A} . The first six perturbed intrinsic sparse modes \hat{g}_k are shown in Figure 11. We can see that their supports are exactly the same as those of the unperturbed intrinsic sparse modes g_k in Figure 9. In fact, the first 37 perturbed intrinsic sparse modes $\{\hat{g}_k\}_{k=1}^{37}$ capture exactly the supports of the unperturbed intrinsic sparse modes $\{g_k\}_{k=1}^{37}$. However, we have several extra perturbed intrinsic sparse modes with very small l^2 error since $\hat{\Omega}$ has rank more than 37.

When we raise the noise level ϵ to 10^{-4} , the histogram of the absolute values in $\hat{\Omega}$ is shown in Figure 10(right). In this case, we cannot identify a gap any more. From Figure 10(left), we see that the exact Ω has entries in the order of 10^{-3} . Therefore, the noise level $\epsilon = 10^{-4}$ is large enough to mix together the true nonzero values and noisy null values in $\hat{\Omega}$. In Figure 10 the total counts are different because only values between $10^{-16.5}$ and $10^{0.5}$ are counted.

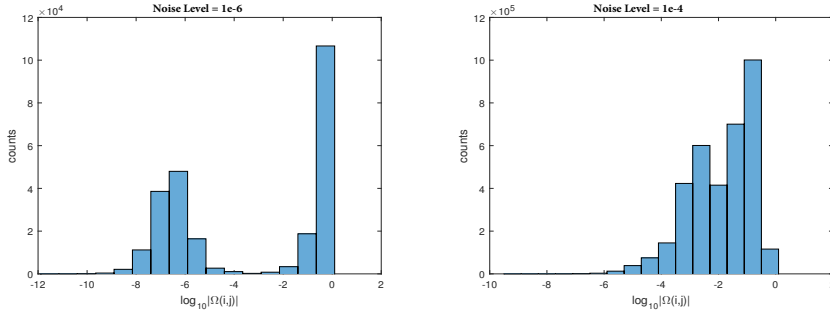


FIG. 10. Histogram of absolute values of entries in $\hat{\Omega}$.

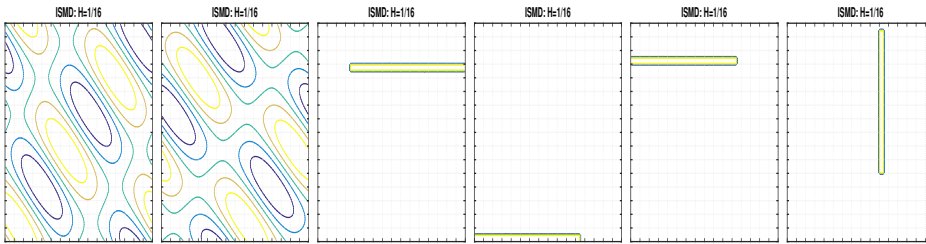


FIG. 11. Application of Algorithm 2 ($H = 1/16$, approximately regular-sparse): First six intrinsic sparse modes

5.6. Application of Algorithm 3. In this section, we consider the one-dimensional Poisson kernel

$$a(x, y) = e^{-\frac{|x-y|}{l}}, \quad x, y \in [-1, 1],$$

where $l = 1/16$. To refine the small scale, $a(x, y)$ is discretized by a uniform grid with $h = 1/512$, resulting in $A \in \mathbb{R}^{1024 \times 1024}$. In Figure 12 we plot the covariance matrix. By truncating the eigendecomposition with 45 modes, we can approximate A with spectral norm error 5%, and these 45 KL modes are plotted on the right panel of the figure. As one can see, they are all global functions.

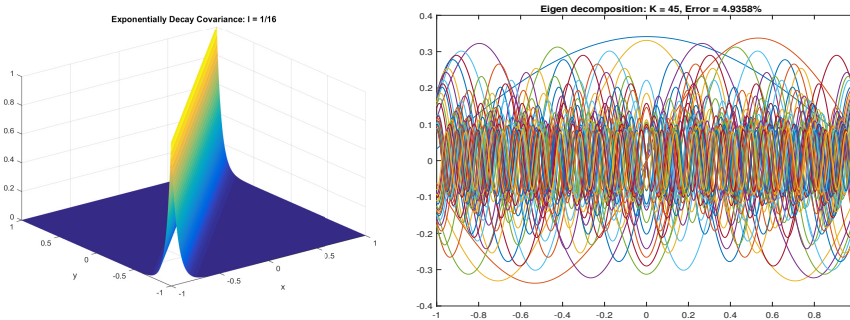


FIG. 12. Eigendecomposition: Covariance function and its first 45 KL modes. Error is 4.936%. Both local and global dimensions are 45.

We decompose the domain into 2, 4, and 8 patches, respectively, and apply Al-

gorithm 3 with thresholding (4.6) to each case. For all three cases, every mode has patchwise sparseness of either 1 or 2. In Figure 13, the left panels show the modes that are nonzero on more than one patch, and the right panels collect the modes that are nonzero on only one patch. To achieve the same accuracy with the eigen-decomposition, the numbers of modes needed are 45, 47, and 49, respectively. The total number is slightly larger than the number of eigenmodes, but most modes are localized. For the two-patch case, each patch contains 23 nonzero modes, and for the four-patch case, each patch contains either 12 or 13 nonzero modes, and for the eight-patch case, each patch contains only seven nonzero modes.

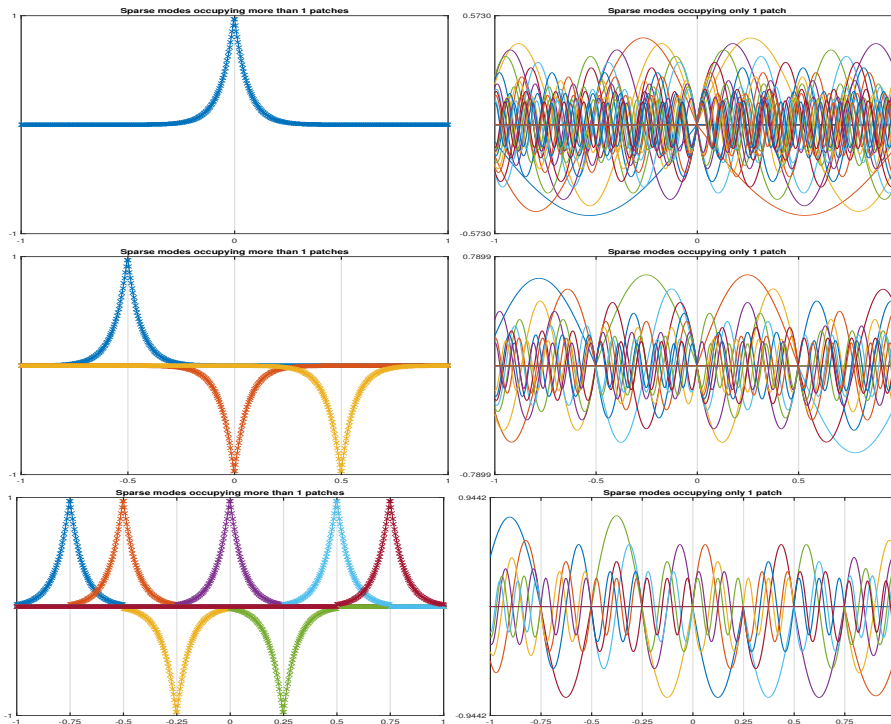


FIG. 13. Upper: Two patches case. Error is 4.95%. Global dimension is 45 and the local dimension is 23 for both patches. Middle: Four patches case. Error is 4.76%. Global dimension is 47 and the local dimension is 12, 13, 13, 12, respectively. Bottom: Eight patches case. Error is 4.42%. Global dimension is 49 and the local dimension is 7 for all patches.

For this translational invariant Poisson kernel, the semidefinite relaxation of sparse PCA (problem (2.20)) also gives satisfactory sparse approximation in the sense of problem (2.19). Numerical tests show that when $\mu < 2$, sparse PCA tends to put too much weight on the sparsity and it leads to poor approximation of A (over 90% error). In Figure 14 we plot 47 physical modes selected out of 513 columns of W , with $\mu = 2.7826$. The error is 4.94%. We also show five of them on the right panel. Note that we have used the translation invariance property in selecting the columns of W .

6. Conclusions and future work. In this paper, we introduced a new matrix factorization method, the intrinsic sparse mode decomposition (ISMD), to obtain a sparse decomposition of low-rank symmetric positive semidefinite matrices. Instead of minimizing the total number of nonzero entries of the decomposed modes, the ISMD

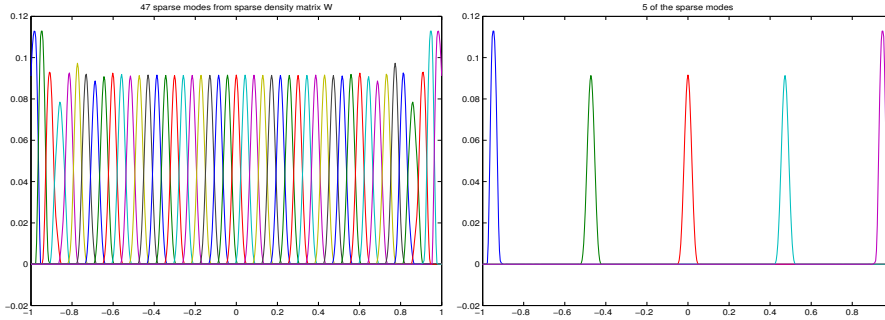


FIG. 14. *Sparse PCA: $\mu = 2.7826$. We specifically choose 47 columns out of W and show five of them.*

minimizes the total patchwise sparseness with a prescribed partition of index set $[N]$. The decomposed modes from the ISMD are called intrinsic sparse modes for the decomposed matrix w.r.t. the partition. The ISMD is equivalent to the eigendecomposition for the coarsest partition and recovers the pivoted Cholesky decomposition for the finest partition. If the partition is regular-sparse w.r.t. the matrix to be decomposed, we prove that the ISMD gives the optimal patchwise sparse decomposition. We also prove that as long as the partition is regular-sparse, the decomposed modes get sparser (in the sense of l^0 norm) as the partition is refined. Finally, we provide a preliminary result on perturbation analysis of the ISMD based on the assumption that the partition is regular-sparse, and the intrinsic sparse modes are identifiable with each other. Numerical examples on synthetic data demonstrate the robustness and efficiency of the ISMD.

Currently, the perturbation analysis is based on an extra assumption that roughly requires that the local eigendecomposition be well conditioned; see (4.3). It would be desirable to perform a perturbation analysis without such an assumption or propose a more stable version of the ISMD. In the paper, we also discussed the differences between the sparse-orthogonal matrix factorization problem (1.4) and the general sparse matrix factorization problem (1.7). We pointed out that the ISMD is not designed to solve the general matrix factorization problem. The ISMD is recommended as a sparse matrix factorization method only if the orthogonality in decomposition coefficients U is required, and an exact (or nearly exact) decomposition is desired.

Finally, we have provided a heuristic algorithm (e.g., Algorithm 3) to solve problem (1.3) for matrix factorization with large noise. Ultimately, the complete resolution of this matrix factorization problem in the presence of large noise requires a better formulation and a more robust algorithm.

Appendix A. Proof of Proposition 3.2.

1. $l_k^{(\psi)}$, divided into patches, can be written as $l_k^{(\psi)} = [l_{1,k}; l_{2,k}; \dots; l_{M,k}]$. From definition (3.2), we have $\|l_{m,k}\|_1 = 1$ if $\psi_k|_{P_m} \neq \mathbf{0}$ and 0 otherwise. Therefore, we obtain

$$\|l_k^{(\psi)}\|_1 = \sum_{m=1}^M \|l_{m,k}\|_1 = s_k(\psi_k; \mathcal{P}).$$

Moreover, on patch P_m different ψ_k 's correspond to different local pieces in Ψ_m (when they are identical, we keep both when constructing Ψ_m), and thus different columns in $L_m^{(\psi)}$ have disjoint supports. Therefore, different columns

in $L^{(\psi)}$ have disjoint supports.

2. From definition (3.2), the j th row of $L_n^{(\psi)}$ is equal to $\mathbf{e}_{k_j^m}^T$, where $\mathbf{e}_{k_j^m}$ is the k_j^m th column of \mathbb{I}_K . Then we have $(L_n^{(\psi)})^T L_n^{(\psi)} = \sum_{j=1}^{d_n} \mathbf{e}_{k_j^n} \mathbf{e}_{k_j^n}^T$. Therefore, we obtain

(A.1)

$$B_{n;m}^{(\psi)} \equiv L_m^{(\psi)} (L_n^{(\psi)})^T L_n^{(\psi)} (L_m^{(\psi)})^T = \sum_{j=1}^{d_n} L_m^{(\psi)} \mathbf{e}_{k_j^n} (L_m^{(\psi)} \mathbf{e}_{k_j^n})^T = \sum_{j=1}^{d_n} l_{m,k_j^n} l_{m,k_j^n}^T,$$

where l_{m,k_j^n} is the k_j^n th column of $L_m^{(\psi)}$.

From definition (3.2), l_{m,k_i^m} , the k_i^m th column of $L_m^{(\psi)}$ is equal to \mathbf{e}_i for $i \in [d_m]$ and all other columns are $\mathbf{0}$. Therefore,

$$(A.2) \quad \sum_{k=1}^K l_{m,k} l_{m,k}^T = \sum_{i=1}^{d_m} l_{m,k_i^m} l_{m,k_i^m}^T = \sum_{i=1}^{d_m} \mathbf{e}_i \mathbf{e}_i^T = \mathbb{I}_{d_m}.$$

Equation (A.1) sums over $k \in \{k_j^n\}_{j=1}^{d_n} \subset [K]$, and then we conclude that $B_{n;m}^{(\psi)}$ is diagonal with diagonal entries either 1 or 0. Moreover, if $B_{n;m}^{(\psi)}(i, i) = 1$, the term $\mathbf{e}_i \mathbf{e}_i^T$ has to be included in the summation in (A.1). Among all terms $\{l_{m,k} l_{m,k}^T\}_{k=1}^K$, only $l_{m,k_i^m} l_{m,k_i^m}^T$ is equal to $\mathbf{e}_i \mathbf{e}_i^T$ due to the definition of $L_m^{(\psi)}$. Therefore, the term $l_{m,k_i^m} l_{m,k_i^m}^T$ has to be included in the summation in (A.1). Therefore, there exists $j \in [d_n]$ such that $k_j^n = k_i^m$. In other words, there exist $k \in [K]$ and $j \in [d_n]$ such that $\psi_k|_{P_m} = \psi_{m,i}$ and $\psi_k|_{P_n} = \psi_{n,j}$.

REFERENCES

- [1] D. AGARWAL AND B.-C. CHEN, *Regression-based latent factor models*, in Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, ACM, New York, 2009, pp. 19–28.
- [2] A. BUNSE-GERSTNER, R. BYERS, AND V. MEHRMANN, *Numerical methods for simultaneous diagonalization*, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 927–949, <https://doi.org/10.1137/0614062>.
- [3] E. J. CANDÈS, X. LI, Y. MA, AND J. WRIGHT, *Robust principal component analysis?*, J. ACM, 58 (2011), 11, <https://doi.org/10.1145/1970392.1970395>.
- [4] J.-F. CARDOSO, *Perturbation of Joint Diagonalizers*, Tech. Report 94D023, Signal Department, Telecom Paris, Paris, 1994, <http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.45.3&rep=rep1&type=pdf>.
- [5] J. F. CARDOSO AND A. SOULOUMIAC, *Blind beamforming for non-Gaussian signals*, IEE Proceedings F—Radar and Signal Processing, 140 (1993), pp. 362–370, <https://doi.org/10.1049/ip-f-2.1993.0054>.
- [6] V. CHANDRASEKARAN, S. SANGHAVI, P. A. PARRILO, AND A. S. WILLSKY, *Rank-sparsity incoherence for matrix decomposition*, SIAM J. Optim., 21 (2011), pp. 572–596, <https://doi.org/10.1137/090761793>.
- [7] Y. CHEN, J. JAKEMAN, C. GITTELSON, AND D. XIU, *Local polynomial chaos expansion for linear differential equations with high dimensional random inputs*, SIAM J. Sci. Comput., 37 (2015), pp. A79–A102, <https://doi.org/10.1137/140970100>.
- [8] A. D’ASPRESMONT, L. EL GHAOU, M. JORDAN, AND G. LANCKRIET, *A direct formulation for sparse PCA using semidefinite programming*, SIAM Rev., 49 (2007), pp. 434–448, <https://doi.org/10.1137/050645506>.
- [9] Y. EFENDIEV, J. GALVIS, AND X.-H. WU, *Multiscale finite element methods for high-contrast problems using local spectral basis functions*, J. Comput. Phys., 230 (2011), pp. 937–955.
- [10] J. GALVIS AND Y. EFENDIEV, *Domain decomposition preconditioners for multiscale flows in high contrast media: Reduced dimension coarse spaces*, Multiscale Model. Simul., 8 (2010), pp. 1621–1644, <https://doi.org/10.1137/100790112>.

- [11] C. GAO AND B. E. ENGELHARDT, *A sparse factor analysis model for high dimensional latent spaces*, in NIPS: Workshop on Analysis Operator Learning vs. Dictionary Learning: Fraternal Twins in Sparse Modeling, 2012.
- [12] M. GHOMMEM, M. PRESHO, V. M. CALO, AND Y. EFENDIEV, *Mode decomposition methods for flows in high-contrast porous media. Global-local approach*, J. Comput. Phys., 253 (2013), pp. 226–238, <https://doi.org/10.1016/j.jcp.2013.06.033>.
- [13] T. Y. HOU AND X.-H. WU, *A multiscale finite element method for elliptic problems in composite materials and porous media*, J. Comput. Phys., 134 (1997), pp. 169–189, <https://doi.org/10.1006/jcph.1997.5682>.
- [14] T. Y. HOU, X.-H. WU, AND Y. ZHANG, *Removing the cell resonance error in the multiscale finite element method via a Petrov-Galerkin formulation*, Commun. Math. Sci., 2 (2004), pp. 185–205, <https://doi.org/10.4310/CMS.2004.v2.n2.a3>.
- [15] Y. T. HOU, Q. LI, AND P. ZHANG, *Exploring the locally low dimensional structure in solving random elliptic pdes*, Multiscale Model. Simul., to appear.
- [16] Y. T. HOU AND P. ZHANG, *Sparse operator compression of elliptic operators – part I: Second order elliptic operators*, Res. Math. Sci., preprint.
- [17] Y. T. HOU AND P. ZHANG, *Sparse operator compression of elliptic operators – part II: High order elliptic operators*, Res. Math. Sci., preprint.
- [18] R. JENATTON, G. OBOZINSKI, AND F. R. BACH, *Structured sparse principal component analysis*, in Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics (AISTATS-10), 2010, pp. 366–373.
- [19] I. JOLLIFFE, M. UDDIN, AND S. VINES, *Simplified EOFs—Three alternatives to rotation*, Climate Research, 20 (2002), pp. 271–279, <https://doi.org/10.3354/cr020271>.
- [20] I. T. JOLLIFFE, N. T. TRENDALOV, AND M. UDDIN, *A modified principal component technique based on the lasso*, J. Comput. Graph. Stat., 12 (2003), pp. 531–547.
- [21] K. KARHUNEN, *Über lineare Methoden in der Wahrscheinlichkeitsrechnung*, Annales Academiae scientiarum Fennicae: Mathematica - Physica, Universitat Helsinki, Helsinki, Finland, 1947.
- [22] W. KOHN, *Image of the Fermi surface in the vibration spectrum of a metal*, Phys. Rev. Lett., 2 (1959), 393, <https://doi.org/10.1103/PhysRevLett.2.393>.
- [23] W. J. KRZANOWSKI AND F. H. C. MARRIOTT, *Multivariate Analysis: Kendall's Library of Statistics*, Volume 2, Kendall's Advanced Theory of Statistics, John Wiley, New York, 1995.
- [24] R. LAI, J. LU, AND S. OSHER, *Density matrix minimization with L_1 regularization*, Comm. Math. Sci., to appear.
- [25] D. D. LEE AND H. S. SEUNG, *Learning the parts of objects by non-negative matrix factorization*, Nature, 401 (1999), pp. 788–791, <https://doi.org/10.1038/44565>.
- [26] H. LEE, A. BATTLE, R. RAINA, AND A. Y. NG, *Efficient sparse coding algorithms*, in Proceedings of the 19th International Conference on Neural Information Processing Systems, 2006, pp. 801–808.
- [27] R. B. LEHOUCQ AND D. C. SORENSEN, *Deflation techniques for an implicitly restarted Arnoldi iteration*, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 789–821, <https://doi.org/10.1137/S0895479895281484>.
- [28] M. LOÈVE, *Probability Theory I*, 4th ed., Grad. Texts Math. 45, Springer-Verlag, New York, Heidelberg, 1977.
- [29] C. LUCAS, *Lapack-style codes for level 2 and 3 pivoted Cholesky factorizations*, LAPACK Working Note 161, 2004.
- [30] X. LUO, *High dimensional low rank and sparse covariance matrix estimation via convex minimization*, preprint, 2011, <http://arxiv.org/abs/1111.1133>.
- [31] J. MAIRAL, F. BACH, J. PONCE, AND G. SAPIRO, *Online learning for matrix factorization and sparse coding*, J. Mach. Learn. Res., 11 (2010), pp. 19–60.
- [32] N. MARZARI, A. A. MOSTOFI, J. R. YATES, I. SOUZA, AND D. VANDERBILT, *Maximally localized Wannier functions: Theory and applications*, Rev. Mod. Phys., 84 (2012), 1419, <https://doi.org/10.1103/RevModPhys.84.1419>.
- [33] N. MARZARI AND D. VANDERBILT, *Maximally localized generalized Wannier functions for composite energy bands*, Phys. Rev. B, 56 (1997), 12847, <https://doi.org/10.1103/PhysRevB.56.12847>.
- [34] H. OWHADI, *Multi-grid with rough coefficients and multiresolution operator decomposition from hierarchical information games*, preprint, 2015, <http://arxiv.org/abs/1503.03467>.
- [35] H. OWHADI, L. ZHANG, AND L. BERLYAND, *Polyharmonic homogenization, rough polyharmonic splines and sparse super-localization*, ESAIM Math. Model. Numer. Anal., 48 (2014), pp. 517–552, <https://doi.org/10.1051/m2an/2013118>.

- [36] V. OZOLIŅŠ, R. LAI, R. CAFLISCH, AND S. OSHER, *Compressed modes for variational problems in mathematics and physics*, Proc. Natl. Acad. Sci. USA, 110 (2013), pp. 18368–18373, <http://www.pnas.org/content/110/46/18368.short>.
- [37] D. C. SORENSEN, *Implicit application of polynomial filters in a k-step Arnoldi method*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 357–385, <https://dx.doi.org/10.1137/0613025>.
- [38] G. W. STEWART, *On the perturbation of pseudo-inverses, projections and linear least squares problems*, SIAM Rev., 19 (1977), pp. 634–662, <https://doi.org/10.1137/1019104>.
- [39] V. Q. VU, J. CHO, J. LEI, AND K. ROHE, *Fantope projection and selection: A near-optimal convex relaxation of sparse PCA*, in Proceedings in Advances in Neural Information Processing Systems 26, 2013, pp. 2670–2678.
- [40] A. WAGNER, J. WRIGHT, A. GANESH, Z. ZHOU, H. MOBAHI, AND Y. MA, *Toward a practical face recognition system: Robust alignment and illumination by sparse representation*, IEEE Trans. Pattern Anal. Mach. Intell., 34 (2012), pp. 372–386, <https://doi.org/10.1109/TPAMI.2011.112>.
- [41] G. H. WANNIER, *The structure of electronic excitation levels in insulating crystals*, Phys. Rev., 52 (1937), 191, <https://doi.org/10.1063/1.1744206>.
- [42] E. WEINAN, T. LI, AND J. LU, *Localized bases of eigensubspaces and operator compression*, Proc. Natl. Acad. Sci. USA, 107 (2010), pp. 1273–1278.
- [43] D. M. WITTEN, R. TIBSHIRANI, AND T. HASTIE, *A penalized matrix decomposition, with applications to sparse principal components and canonical correlation analysis*, Biostatistics, 10 (2009), pp. 515–534, <https://doi.org/10.1093/biostatistics/kxp008>.
- [44] Y. ZHANG, A. D’ASPROMONT, AND L. EL GHAOUI, *Sparse PCA: Convex relaxations, algorithms and applications*, in Handbook on Semidefinite, Conic and Polynomial Optimization, Springer, New York, 2012, pp. 915–940.
- [45] H. ZOU, T. HASTIE, AND R. TIBSHIRANI, *Sparse principal component analysis*, J. Comput. Graph. Statist., 15 (2004), pp. 265–286.