

Covariance Estimation in High Dimensions via Kronecker Product Expansions

Theodoros Tsiligkaridis *, *Student Member, IEEE*, Alfred O. Hero III, *Fellow, IEEE*

Abstract

This paper presents a new method for estimating high dimensional covariance matrices. The method, permuted rank-penalized least-squares (PRLS), is based on a Kronecker product series expansion of the true covariance matrix. Assuming an i.i.d. Gaussian random sample, we establish high dimensional rates of convergence to the true covariance as both the number of samples and the number of variables go to infinity. For covariance matrices of low separation rank, our results establish that PRLS has significantly faster convergence than the standard sample covariance matrix (SCM) estimator. The convergence rate captures a fundamental tradeoff between estimation error and approximation error, thus providing a scalable covariance estimation framework in terms of separation rank, similar to low rank approximation of covariance matrices [1]. The MSE convergence rates generalize the high dimensional rates recently obtained for the ML Flip-flop algorithm [2], [3] for Kronecker product covariance estimation. We show that a class of block Toeplitz covariance matrices is approximatable by low separation rank and give bounds on the minimal separation rank r that ensures a given level of bias. Simulations are presented to validate the theoretical bounds. As a real world application, we illustrate the utility of the proposed Kronecker covariance estimator for spatio-temporal linear least squares prediction of multivariate wind speed measurements.

Index Terms

Structured covariance estimation, penalized least squares, Kronecker product decompositions, high dimensional convergence rates, mean-square error, multivariate prediction.

The research reported in this paper was supported in part by ARO grant W911NF-11-1-0391. Preliminary results in this paper have appeared at the 2013 IEEE International Symposium on Information Theory.

T. Tsiligkaridis and A. O. Hero, III, are with the Department of Electrical Engineering and Computer Science, University of Michigan, Ann Arbor, MI 48109 USA (e-mail: ttsili@umich.edu, hero@umich.edu).

I. INTRODUCTION

Covariance estimation is a fundamental problem in multivariate statistical analysis. It has received attention in diverse fields including economics and financial time series analysis (e.g., portfolio selection, risk management and asset pricing [4]), bioinformatics (e.g. gene microarray data [5], [6], functional MRI [7]) and machine learning (e.g., face recognition [8], recommendation systems [9]). In many modern applications, data sets are very large with both large number of samples n and large dimension d , often with $d \gg n$, leading to a number of covariance parameters that greatly exceeds the number of observations. The search for good low-dimensional representations of these data sets has led to much progress in their analysis. Recent examples include sparse covariance estimation [10], [11], [12], [13], low rank covariance estimation [14], [15], [16], [1], and Kronecker product estimation [17], [18], [19], [2], [3].

Kronecker product (KP) structure is a different covariance constraint from sparse or low rank constraints. KP represents a $pq \times pq$ covariance matrix Σ_0 as the Kronecker product of two lower dimensional covariance matrices. When the variables are multivariate Gaussian with covariance following the KP model, the variables are said to follow a matrix normal distribution [19], [17], [20]. This model has applications in channel modeling for MIMO wireless communications [21], geostatistics [22], genomics [23], multi-task learning [24], face recognition [8], recommendation systems [9] and collaborative filtering [25]. The main difficulty in maximum likelihood estimation of structured covariances is the nonconvex optimization problem that arises. Thus, an alternating optimization approach is usually adopted. In the case where there is no missing data, an extension of the alternating optimization algorithm of Werner *et al* [18], that the authors called the flip flop (FF) algorithm, can be applied to estimate the parameters of the Kronecker product model, called KGLasso in [2].

In this paper, we assume that the covariance can be represented as a sum of Kronecker products of two lower dimensional factor matrices, where the number of terms in the summation may depend on the factor dimensions. More concretely, we assume that there are $d = pq$ variables whose covariance Σ_0 has Kronecker product representation:

$$\Sigma_0 = \sum_{\gamma=1}^r \mathbf{A}_{0,\gamma} \otimes \mathbf{B}_{0,\gamma} \quad (1)$$

where $\{\mathbf{A}_{0,\gamma}\}$ are $p \times p$ linearly independent matrices and $\{\mathbf{B}_{0,\gamma}\}$ are $q \times q$ linearly independent matrices. We assume that the factor dimensions p, q are known. We note $1 \leq r \leq r_0 = \min(p^2, q^2)$ and refer to

Linear independence is with respect to the trace inner product defined in the space of symmetric matrices.

r as the *separation rank*. The model (1) is analogous to separable approximation of continuous functions [26]. It is evocative of a type of low rank principal component decomposition where the components are Kronecker products. However, the components in (1) are neither orthogonal nor normalized. The model (1) with separation rank 1 is relevant to channel modeling for MIMO wireless communications, where \mathbf{A}_0 is a transmit covariance matrix and \mathbf{B}_0 is a receive covariance matrix [21]. The rank 1 model is also relevant to other transposable models arising in recommendation systems like NetFlix and in gene expression analysis [9]. The model (1) with $r \geq 1$ has applications in spatiotemporal MEG/EEG covariance modeling [27], [28], [29], [30] and SAR data analysis [31]. We finally note that Van Loan and Pitsianis [32] have shown that any $pq \times pq$ matrix Σ_0 can be written as an orthogonal expansion of Kronecker products of the form (1), thus allowing any covariance matrix to be approximated by a bilinear decomposition of this form.

The main contribution of this paper is a convex optimization approach to estimating covariance matrices with KP structure of the form (1) and the derivation of tight high-dimensional MSE convergence rates as n , p and q go to infinity. We call our method the Permuted Rank-penalized Least Squares (PRLS) estimator. Similarly to other studies of high dimensional covariance estimation [33], [2], [13], [34], [35], we analyze the estimator convergence rate in Frobenius norm of PRLS, providing specific convergence rates holding with certain high probability. In other words, our analysis provides high probability guarantees up to absolute constants in all sample sizes and dimensions.

For estimating separation rank r covariance matrices of the form (1), we establish that PRLS achieves high dimensional consistency with a convergence rate of $O_P\left(\frac{r(p^2+q^2+\log \max(p,q,n))}{n}\right)$. This can be significantly faster than the convergence rate $O_P\left(\frac{p^2q^2}{n}\right)$ of the standard sample covariance matrix (SCM). For separation rank $r = 1$ this rate is identical to that of the FF algorithm, which fits the sample covariance matrix to a single Kronecker factor.

The PRLS method for estimating the Kronecker product expansion (1) generalizes previously proposed Kronecker product covariance models [17], [19] to the case of $r > 1$. This is a fundamentally different generalization than the $r = 1$ sparse KP models proposed in [9], [2], [3], [36]. Independently in [2], [3] and [36], it was established that the high dimensional convergence rate for these sparse KP models is of order $O_P\left(\frac{(p+q)\log \max(p,q,n)}{n}\right)$. While we do not pursue the the additional constraint of sparsity in this paper, we speculate that sparsity can be combined with the Kronecker sum model (1), achieving even better convergence.

Advantages of the proposed PRLS covariance estimator is illustrated on both simulated and real data. The application of PRLS to the NCEP wind dataset shows that a low order Kronecker sum provides

a remarkably good fit to the spatio-temporal sample covariance matrix: over 86% of all the energy is contained in the first Kronecker component of the Kronecker expansion as compared to only 41% in the principal component of the standard PCA eigen-expansion. Furthermore, by replacing the SCM in the standard linear predictor by our Kronecker sum estimator we demonstrate a 1.9 dB RMSE advantage for predicting next-day wind speeds from NCEP network past measurements.

The outline of the paper is as follows. Section II introduces the notation that will be used throughout the paper. Section III introduces the PRLS covariance estimation method. Section IV presents the high-dimensional MSE convergence rate of PRLS. Section V presents numerical experiments. The technical proofs are placed in the Appendix.

II. NOTATION

For a square matrix \mathbf{M} , define $\|\mathbf{M}\|_1 = \|\text{vec}(\mathbf{M})\|_1$ and $\|\mathbf{M}\|_\infty = \|\text{vec}(\mathbf{M})\|_\infty$, where $\text{vec}(\mathbf{M})$ denotes the vectorized form of \mathbf{M} (concatenation of columns of \mathbf{M} into a column vector). $\|\mathbf{M}\|_2$ is the spectral norm of \mathbf{M} . $\mathbf{M}_{i,j} = [\mathbf{M}]_{i,j}$ is the (i,j) th element of \mathbf{M} . Let the inverse transformation (from a vector to a matrix) be defined as: $\text{vec}^{-1}(\mathbf{x}) = \mathbf{X}$, where $\mathbf{x} = \text{vec}(\mathbf{X})$. Define the $pq \times pq$ permutation operator $\mathbf{K}_{p,q}$ such that $\mathbf{K}_{p,q}\text{vec}(\mathbf{N}) = \text{vec}(\mathbf{N}^T)$ for any $p \times q$ matrix \mathbf{N} . For a symmetric positive definite matrix \mathbf{M} , $\lambda(\mathbf{M})$ will denote the vector of real eigenvalues of \mathbf{M} and define $\lambda_{\max}(\mathbf{M}) = \|\mathbf{M}\|_2 = \max \lambda_i(\mathbf{M})$, and $\lambda_{\min}(\mathbf{M}) = \min \lambda_i(\mathbf{M})$. For any matrix \mathbf{M} , define the nuclear norm $\|\mathbf{M}\|_* = \sum_{l=1}^{r_M} |\sigma_l(\mathbf{M})|$, where $r_M = \text{rank}(\mathbf{M})$ and $\sigma_l(\mathbf{M})$ is the l th singular value of \mathbf{M} .

For a matrix \mathbf{M} of size $pq \times pq$, let $\{\mathbf{M}(i,j)\}_{i,j=1}^p$ denote its $q \times q$ block submatrices, where each block submatrix is $\mathbf{M}(i,j) = [\mathbf{M}]_{(i-1)q+1:iq, (j-1)q+1:jq}$. Also let $\{\bar{\mathbf{M}}(k,l)\}_{k,l=1}^q$ denote the $p \times p$ block submatrices of the permuted matrix $\bar{\mathbf{M}} = \mathbf{K}_{p,q}^T \mathbf{M} \mathbf{K}_{p,q}$. Define the permutation operator $\mathcal{R} : \mathbb{R}^{pq \times pq} \rightarrow \mathbb{R}^{p^2 \times q^2}$ by setting the $(i-1)p + j$ row of $\mathcal{R}(\mathbf{M})$ equal to $\text{vec}(\mathbf{M}(i,j))^T$. When \mathbf{M} is representable as the Kronecker product $\mathbf{M}_1 \otimes \mathbf{M}_2$, an illustration of this permutation operator is shown in Fig. 1.

Define the set of symmetric matrices $S^p = \{\mathbf{A} \in \mathbb{R}^{p \times p} : \mathbf{A} = \mathbf{A}^T\}$, the set of symmetric positive semidefinite (psd) matrices S_+^p , and the set of symmetric positive definite (pd) matrices S_{++}^p . \mathbf{I}_d is a $d \times d$ identity matrix. It can be shown that S_{++}^p is a convex set but is not closed [37]. Note that S_{++}^p is simply the interior of the closed convex cone S_+^p .

For a subspace U , define \mathbf{P}_U and \mathbf{P}_U^\perp as the orthogonal projection operators projecting onto U and U^\perp , respectively. The unit Euclidean sphere in \mathbb{R}^d is denoted by $\mathcal{S}_{d-1} = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}\|_2 = 1\}$. Let $(x)_+ = \max(x, 0)$.

Statistical convergence rates will be denoted by the $O_P(\cdot)$ notation, which is defined as follows.




Fig. 1. Original (top) and permuted covariance (bottom) matrix. The original covariance is $\Sigma_0 = \mathbf{A}_0 \otimes \mathbf{B}_0$, where \mathbf{A}_0 is a 10×10 Toeplitz matrix and \mathbf{B}_0 is a 20×20 unstructured p.d. matrix. Note that the permutation operator \mathcal{R} maps a symmetric p.s.d. matrix Σ_0 to a non-symmetric rank 1 matrix $\mathbf{R}_0 = \mathcal{R}(\Sigma_0)$.

Consider a sequence of real random variables $\{X_n\}_{n \in \mathbb{N}}$ defined on a probability space (Ω, \mathcal{F}, P) and a deterministic (positive) sequence of reals $\{b_n\}_{n \in \mathbb{N}}$. By $X_n = O_P(1)$ is meant: $\sup_{n \in \mathbb{N}} \Pr(|X_n| > K) \rightarrow 0$ as $K \rightarrow \infty$, where X_n is a sequence indexed by n , for fixed p, q . The notation $X_n = O_P(b_n)$ is equivalent to $\frac{X_n}{b_n} = O_P(1)$. By $X_n = o_p(1)$ is meant $\Pr(|X_n| > \epsilon) \rightarrow 0$ as $n \rightarrow \infty$ for any $\epsilon > 0$. By $\lambda_n \asymp b_n$ is meant $c_1 \leq \frac{\lambda_n}{b_n} \leq c_2$ for all n , where $c_1, c_2 > 0$ are finite constants.

III. PERMUTED RANK-PENALIZED LEAST-SQUARES

Available are n i.i.d. multivariate Gaussian observations $\{\mathbf{z}_t\}_{t=1}^n$, $\mathbf{z}_t \in \mathbb{R}^{pq}$, having zero-mean and covariance equal to (1). A sufficient statistic for estimating the covariance is the well-known sample covariance matrix (SCM):

$$\hat{\mathbf{S}}_n = \frac{1}{n} \sum_{t=1}^n \mathbf{z}_t \mathbf{z}_t^T \quad (2)$$

The SCM is an unbiased estimator of the true covariance matrix. However, when the number of samples n is smaller than the number of variables $d = pq$ the SCM suffers from high variance and a low rank approximation to the SCM is commonly used. The most common low rank approximation is to perform the eigendecomposition of $\hat{\mathbf{S}}_n$ and retain only the top r principal components resulting in an estimator, called the PCA estimator, of the form:

$$\hat{\mathbf{S}}_n^{PCA} = \sum_{i=1}^r \sigma_i^2 \nu_i \nu_i^T, \quad (3)$$

where $r < d$ is selected according to some heuristic. It is now well known [38], [39] that this PCA estimator suffers from high bias when n is smaller than $d = pq$.

An alternative approach to low rank covariance estimation was proposed in [1] specifying a low rank covariance estimator as the solution of the penalized least squares problem :

$$\hat{\Sigma}_n^\lambda \in \arg \min_{\mathbf{S} \in \mathcal{S}_{++}^d} \|\hat{\mathbf{S}}_n - \mathbf{S}\|_F^2 + \lambda \text{tr}(\mathbf{S}) \quad (4)$$

where $\lambda > 0$ is a regularization parameter.

The estimator (4) has several useful interpretations. First, it can be interpreted as a convex relaxation of the non-convex rank constrained Frobenius norm minimization problem

$$\arg \min_{\mathbf{S} \in \mathcal{S}_{++}^d, \text{rank}(\mathbf{S}) \leq r} \|\hat{\mathbf{S}}_n - \mathbf{S}\|_F^2,$$

whose solution, by the Eckhart-Young theorem, is the PCA estimator (3). Second, it can be interpreted as a covariance version of the lasso regression problem, i.e., finding a low rank psd ℓ_2 approximation to the sample covariance matrix. The term $\text{tr}(\mathbf{S})$ in 4 is equivalent to the ℓ_1 norm on the eigenvalues of the psd matrix \mathbf{S} . As shown in [1] the solution to the convex minimization in (4) converges to the ensemble covariance $\Sigma_0 = \mathbb{E}[\mathbf{z}_t \mathbf{z}_t^T]$ at the minimax optimal rate. Corollary 1 in [1] establishes that, for $\lambda = C' \|\Sigma_0\|_2 \sqrt{\frac{r(\Sigma_0) \log(2d)}{n}}$, $n \geq cr(\Sigma_0) \log^2(\max(2d, n))$ and $C', c > 0$ sufficiently large, establishes a tight Frobenius norm error bound, which states that with probability $1 - \frac{1}{2d}$:

$$\|\hat{\Sigma}_n^\lambda - \Sigma_0\|_F^2 \leq \inf_{\mathbf{S} \succ 0} \|\Sigma_0 - \mathbf{S}\|_F^2 + C \|\Sigma_0\|_2^2 \text{rank}(\mathbf{S}) \frac{r(\Sigma_0) \log(2d)}{n}$$

where $r(\Sigma_0) = \frac{\text{tr}(\Sigma_0)}{\|\Sigma_0\|_2} \leq \min\{\text{rank}(\Sigma_0), d\}$ is the effective rank [1]. The absolute constant C is given by $\frac{(1+\sqrt{2})^2}{8} (C')^2$.

The estimator (4) was developed in [1] for the more general problem where there could be missing data.

Here we propose a similar nuclear norm penalization approach to estimate low separation-rank covariance matrices of form (1). Motivated by Van Loan and Pitsianis's work [32], we propose:

$$\hat{\mathbf{R}}_n^\lambda \in \arg \min_{\mathbf{R} \in \mathbb{R}^{p^2 \times q^2}} \|\hat{\mathbf{R}}_n - \mathbf{R}\|_F^2 + \lambda \|\mathbf{R}\|_* \quad (5)$$

where $\hat{\mathbf{R}}_n = \mathcal{R}(\hat{\mathbf{S}}_n)$ is the permuted SCM of size $p^2 \times q^2$ (see Notation section). The minimum-norm problem considered in [32] is:

$$\min_{\mathbf{R} \in \mathbb{R}^{p^2 \times q^2} : \text{rank}(\mathbf{R}) \leq r} \|\hat{\mathbf{R}}_n - \mathbf{R}\|_F^2 \quad (6)$$

Specifically, let $\mathbf{S} = \sum_{i=1}^r \mathbf{A}_i \otimes \mathbf{B}_i$ where for all i the dimensions of the matrices \mathbf{A}_i and \mathbf{B}_i are fixed. Then, as the Frobenius norm of a matrix is invariant to permutation of its elements, it follows that $\|\mathbf{S}_n - \mathbf{S}\|_F = \|\mathbf{R}_n - \mathbf{R}\|_F$ where $\mathbf{R}_n = \mathcal{R}(\mathbf{S}_n)$ and $\mathbf{R} = \mathcal{R}(\mathbf{S})$ (which is a matrix of algebraic rank r).

We note that (5) is a convex relaxation of (6) and is more amenable to numerical optimization. Furthermore, we show a tradeoff between approximation error (i.e., the error induced by model mismatch between the true covariance and the model) and estimation error (i.e., the error due to finite sample size) by analyzing the solution of (5). We also note that (5) is a strictly convex problem, so there exists a unique solution that can be efficiently found using well established numerical methods [37].

The solution of (5) is closed form and is given by a thresholded singular value decomposition:

$$\hat{\mathbf{R}}_n^\lambda = \sum_{j=1}^{\min(p^2, q^2)} \left(\sigma_j(\hat{\mathbf{R}}_n) - \frac{\lambda}{2} \right)_+ \mathbf{u}_j \mathbf{v}_j^T \quad (7)$$

where \mathbf{u}_j and \mathbf{v}_j are the left and right singular vectors of $\hat{\mathbf{R}}_n$. This is converted back to a square $pq \times pq$ matrix $\hat{\Sigma}_n^\lambda$ by applying the inverse permutation operator \mathcal{R}^{-1} to $\hat{\mathbf{R}}_n$ (see Notation section).

Efficient methods for numerically evaluating penalized objectives like (5) have been recently proposed [40], [41] and do not require computing the full SVD. Although empirically observed to be fast, the computational complexity of the algorithms presented in [40] and [41] is unknown. The rank- r SVD can be computed with $O(p^2 q^2 r)$ floating point operations. There exist faster randomized methods for truncated SVD requiring only $O(p^2 q^2 \log(r))$ floating point operations [42]. Thus, the computational complexity of solving (5) scales well with respect to the desired separation rank r .

The next theorem shows that the de-permuted version of (7) is symmetric and positive definite.

Theorem 1. *Consider the de-permuted solution $\hat{\Sigma}_n^\lambda = \mathcal{R}^{-1}(\hat{\mathbf{R}}_n^\lambda)$. The following are true:*

- 1) *The solution $\hat{\Sigma}_n^\lambda$ is symmetric with probability 1.*
- 2) *If $n \geq pq$, then the solution $\hat{\Sigma}_n^\lambda$ is positive definite with probability 1.*

Proof: See Appendix A. ■

We believe that the PRLS estimate $\hat{\Sigma}_n^\lambda$ is positive definite even if $n < pq$ for appropriately selected $\lambda > 0$. In our simulations, we always found $\hat{\Sigma}_n^\lambda$ to be positive definite. We have also found that the condition number of the PRLS estimate is orders of magnitude smaller than that of the SCM.

IV. HIGH DIMENSIONAL CONSISTENCY OF PRLS

In this section, we show that RPLS achieves the MSE statistical convergence rate of $O_P\left(\frac{r(p^2+q^2+\log M)}{n}\right)$. This result is clearly superior to the statistical convergence rate of the naive SCM estimator [35],

$$\|\hat{\mathbf{S}}_n - \Sigma_0\|_F^2 = O_P\left(\frac{p^2q^2}{n}\right), \quad (8)$$

particularly when $p, q \rightarrow \infty$.

The next result provides a relation between the spectral norm of $\hat{\mathbf{R}}_n - \mathbf{R}_0$, the Frobenius norm of $\mathbf{R} - \mathbf{R}_0$ and the Frobenius norm of the the estimation error $\hat{\mathbf{R}}_n^\lambda - \mathbf{R}_0$.

Theorem 2. *Consider the convex optimization problem (5). When $\lambda \geq 2\|\hat{\mathbf{R}}_n - \mathbf{R}_0\|_2$, the following holds:*

$$\|\hat{\mathbf{R}}_n^\lambda - \mathbf{R}_0\|_F^2 \leq \inf_{\mathbf{R}} \left\{ \|\mathbf{R} - \mathbf{R}_0\|_F^2 + \frac{(1 + \sqrt{2})^2}{4} \lambda^2 \text{rank}(\mathbf{R}) \right\} \quad (9)$$

Proof: See Appendix B. ■

A. High Dimensional Operator Norm Bound for the Permuted Sample Covariance Matrix

In this subsection, we establish a tight bound on the spectral norm of the error matrix

$$\Delta_n = \hat{\mathbf{R}}_n - \mathbf{R}_0 = \mathcal{R}(\hat{\mathbf{S}}_n - \Sigma_0). \quad (10)$$

The standard strong law of large numbers implies that for fixed dimensions p, q , we have $\Delta_n \rightarrow 0$ almost surely as $n \rightarrow \infty$. The next result will characterize the finite sample fluctuations of this convergence (in probability) measured by the spectral norm as a function of the sample size n and Kronecker factor dimensions p, q . This result will be useful for establishing a tight bound on the Frobenius norm convergence rate of PRLS and can guide the selection of the regularization paramater in (5).

Theorem 3. *(Operator Norm Bound on Permuted SCM) Assume $\|\Sigma_0\|_2 < \infty$ for all p, q and define $M = \max(p, q, n)$. Fix $\epsilon' < \frac{1}{2}$. Assume $t \geq \max(\sqrt{4C_1 \ln(1 + \frac{2}{\epsilon'})}, 4C_2 \ln(1 + \frac{2}{\epsilon'}))$ and $C = \max(C_1, C_2) > 0$*

. Then, with probability at least $1 - 2M^{-\frac{t}{4C}}$,

$$\|\Delta_n\|_2 \leq \frac{C_0 t}{1 - 2\epsilon'} \max \left\{ \frac{p^2 + q^2 + \log M}{n}, \sqrt{\frac{p^2 + q^2 + \log M}{n}} \right\} \quad (11)$$

where $C_0 = \|\Sigma_0\|_2 > 0$.

Proof: See Appendix D. ■

The proof technique is based on a large deviation inequality, derived in Lemma 2 in Appendix C. This inequality characterizes the tail behavior of the quadratic form $\mathbf{x}^T \Delta_n \mathbf{y}$ over the spheres $\mathbf{x} \in \mathcal{S}_{p^2-1}$ and $\mathbf{y} \in \mathcal{S}_{q^2-1}$. Using Lemma 2 and a sphere covering argument, the result of Theorem 3 follows (see Appendix E). Fig. 2 empirically validates the tightness of the bound (11) under the trivial separation rank 1 covariance $\Sigma_0 = \mathbf{I}_p \otimes \mathbf{I}_q$.




Fig. 2. Monte Carlo simulation for growth of spectral norm $\|\Delta_n\|_2^2$ as a function of p for fixed $n = 10$ and $q = 5$. The predicted curve is a least-square fit of a quadratic model $y = ax^2 + b$ to the empirical curve, and is a great fit. This example shows the tightness of the probabilistic bound (11).

The constants C_1, C_2 are defined in Lemma 2 in Appendix B.

The constant $\frac{C_0 t}{1 - 2\epsilon'}$ in front of the rate can be optimized by minimizing it as a function of ϵ' over the interval $(0, 1/2)$.

B. High Dimensional MSE Convergence Rate for PRLS

Using the result in Thm. 3 and the bound in Thm. 2, we next provide a tight bound on the MSE estimation error.

Theorem 4. Define $M = \max(p, q, n)$. Set $\lambda = \lambda_n = \frac{2C_0 t}{1-2\epsilon'} \max \left\{ \frac{p^2 + q^2 + \log M}{n}, \sqrt{\frac{p^2 + q^2 + \log M}{n}} \right\}$ with t satisfying the conditions of Thm. 3. Then, with probability at least $1 - 2M^{-\frac{t}{4C}}$:

$$\begin{aligned} \|\hat{\Sigma}_n^\lambda - \Sigma_0\|_F^2 &\leq \inf_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r} \|\mathbf{R} - \mathbf{R}_0\|_F^2 \\ &\quad + C' r \max \left\{ \left(\frac{p^2 + q^2 + \log M}{n} \right)^2, \frac{p^2 + q^2 + \log M}{n} \right\} \end{aligned} \quad (12)$$

where $C' = \left(C_0 t \frac{1+\sqrt{2}}{1-2\epsilon'} \right)^2 = (3(1+\sqrt{2})C_0 t)^2 > 0$.

Proof: See Appendix E. ■

When Σ_0 is truly a sum of r Kronecker products with factor dimensions p and q , there is no model mismatch and the approximation error $\inf_{\{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r\}} \|\mathbf{R} - \mathbf{R}_0\|_F^2$ is zero. In this case, in the large- p, q, n asymptotic regime where $p^2 + q^2 + \log M = o(n)$, it follows that $\|\hat{\Sigma}_n^\lambda - \Sigma_0\|_F = O_P(\sqrt{\frac{r(p^2 + q^2 + \log M)}{n}}) = o_p(1)$. This asymptotic MSE convergence rate of the estimated covariance to the true covariance reflects the number of degrees of freedom of the model, which is on the order of the total number $r(p^2 + q^2)$ of unknown parameters. This result extends the recent high-dimensional results obtained in [2], [3], [43] for the single Kronecker product model (i.e., $r = 1$).

Recall that $r \leq r_0 = \min(p^2, q^2)$. For the case when $p \sim q$, and $r \sim r_0$, we have a fully saturated Kronecker product model and the number of model parameters are of the order $p^4 \sim d^2$, and the SCM convergence rate (8) coincides with the rate obtained in Thm. 4.

For covariance models of low separation rank-i.e., $r \ll r_0$, Thm. 4 asserts that the high dimensional MSE convergence rate of PRLS can be much lower than the naive SCM convergence rate. Thus PRLS is an attractive alternative to rank-based series expansions like principal component analysis (PCA). We note that each term in the expansion $\mathbf{A}_{0,\gamma} \otimes \mathbf{B}_{0,\gamma}$ can be full-rank, while each term in the standard PCA expansion is rank 1.

Finally, we observe that Thm. 4 captures the tradeoff between estimation error and approximation error. In other words, choosing a smaller r than the true separation rank would incur a larger approximation error $\inf_{\{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r\}} \|\mathbf{R} - \mathbf{R}_0\|_F^2 > 0$, but smaller estimation error on the order of $O_P(\frac{r(p^2 + q^2 + \log M)}{n})$.

C. Approximation Error

It is well known from least-squares approximation theory that the residual error can be rewritten as:

$$\inf_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r} \|\mathbf{R} - \mathbf{R}_0\|_F^2 = \sum_{k=r+1}^{r_0} \sigma_k^2(\mathbf{R}_0), \quad (13)$$

where $\{\sigma_k(\mathbf{R}_0)\}$ are the singular values of \mathbf{R}_0 . In the high dimensional setting, the sample size n grows with the dimensions p, q so that the maximum separation rank r_0 also grows to infinity, and the approximation error (13) may not be finite. In this case the bound in Theorem 4 will not be finite. Hence, an additional condition will be needed to ensure that the sum (13) remains finite as $p, q \rightarrow \infty$: the singular values of \mathbf{R}_0 need to decay faster than $O(1/k)$.

We show next that the class of block-Toeplitz covariance matrices have bounded approximation error if the separation rank scales like $\log(\max(p, q))$. To show this, we first provide a tight variational bound on the singular value spectrum of any $p^2 \times q^2$ matrix \mathbf{R} . Note that the work on high dimensional Toeplitz covariance estimation under operator and Frobenius norms [33], [34] are not applicable to the block-Toeplitz case. To establish Thm. 5 on block Toeplitz matrices we first need the following Lemma.

Lemma 1. (*Variational Bound on Singular Value Spectrum*) Let \mathbf{R} be an arbitrary matrix of size $p^2 \times q^2$. Let \mathbf{P}_k be an orthogonal projection of \mathbb{R}^{q^2} onto \mathbb{R}^k . Then, for $k = 1, \dots, r_0 - 1$ we have:

$$\sigma_{k+1}^2(\mathbf{R}) \leq \|(\mathbf{I}_{q^2} - \mathbf{P}_k)\mathbf{R}^T\|_2^2 \quad (14)$$

with equality iff $\mathbf{P}_k = \mathbf{V}_k \mathbf{V}_k^T$. Also, $\mathbf{V}_k = [\mathbf{v}_1, \dots, \mathbf{v}_k]$, where \mathbf{v}_i is the i th column of \mathbf{V} and $\mathbf{R} = \mathbf{U} \mathbf{\Sigma} \mathbf{V}^T$ is the singular value decomposition.

Proof: See Appendix F. ■

Using this fundamental lemma, we can characterize the approximation error for estimating block-Toeplitz matrices with exponentially decaying off-diagonal norms. Such matrices arise, for example, as covariance matrices of multivariate stationary random processes of dimension m (see (17)) and take the block Toeplitz form:

$$\underbrace{\mathbf{\Sigma}_0}_{(N+1)m \times (N+1)m} = \begin{bmatrix} \mathbf{\Sigma}(0) & \mathbf{\Sigma}(1) & \dots & \mathbf{\Sigma}(N) \\ \mathbf{\Sigma}(-1) & \mathbf{\Sigma}(0) & \dots & \mathbf{\Sigma}(N-1) \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{\Sigma}(-N) & \mathbf{\Sigma}(-N+1) & \dots & \mathbf{\Sigma}(0) \end{bmatrix} \quad (15)$$

where each submatrix is of size $m \times m$. For a zero-mean vector process $\mathbf{y} = \{\mathbf{y}(0), \dots, \mathbf{y}(N)\}$, the submatrices are given by $\mathbf{\Sigma}(\tau) = \mathbb{E}[\mathbf{y}(0)\mathbf{y}(\tau)^T]$.

Theorem 5. Consider a block-Toeplitz p.d. matrix Σ_0 of size $(N+1)m \times (N+1)m$, with $\|\Sigma(\tau)\|_F^2 \leq C'u^{2|\tau|}q$ for all $\tau = -N, \dots, N$ and constant $u \in (0, 1)$. Let $\hat{\Sigma}_n^\lambda$ be the de-permuted matrix $\mathcal{R}^{-1}(\hat{\mathbf{R}}_n^\lambda)$, where $\hat{\mathbf{R}}_n^\lambda$ is given in (7). Using the minimal separation rank r :

$$r \geq \frac{\log(pq/\epsilon)}{\log(1/u)}.$$

Then, the PRLS algorithm estimates Σ_0 up to an absolute tolerance $\epsilon \in (0, 1)$ with convergence rate guarantee:

$$\|\hat{\Sigma}_n^\lambda - \Sigma_0\|_F^2 \leq \epsilon + C'r \frac{p^2 + q^2 + \log M}{n} \quad (16)$$

holding with probability at least $1 - \max(p, q, n)^{-t/4C}$ for λ chosen as perscribed in Thm. 4. Here, $t > 1$ is constant and $C, C' > 0$ are constants specified in Thm. 4.

Proof: See Appendix G. ■

The exponential norm decay condition of Thm. 5 is satisfied by a first-order vector autoregressive process:

$$\mathbf{Z}_t = \Phi \mathbf{Z}_{t-1} + \mathcal{E}_t \quad (17)$$

with $u = \|\Phi\|_2 \in (0, 1)$, where $\mathbf{Z}_t \in \mathbb{R}^m$. For $\mathcal{E}_t \sim N(0, \Sigma_\epsilon)$, this is a multivariate Gaussian process. Collecting data over a time horizon of size $N+1$, we concatenate these observations into a large random vector \mathbf{z} of dimension $(N+1)m$, where m is the process dimension. The resulting covariance matrix has the block-Toeplitz form assumed in Thm. 5. Figure 3 shows bounds constructed using the Frobenius upper bound on the spectral norm in (14) and using the projection matrix \mathbf{P}_k as discussed in the proof of Thm. 5. The bound given in the proof of Thm. 5 (in black) is shown to be linear in log-scale, thus justifying the exponential decay of the Kronecker spectrum.

V. SIMULATION RESULTS

We consider dense positive definite matrices Σ_0 of dimension $d = 625$. Taking $p = q = 25$, we note that the number of free parameters that describe each Kronecker product is of the order $p^2 + q^2 \sim p^2$, which is essentially of the same order as the number of unknown parameters required to specify each eigenvector of Σ_0 , i.e., $pq \sim p^2$.

A. Sum of Kronecker Product Covariance

The covariance matrix shown in Fig. 4 was constructed using (1) with $r = 3$, with each p.d. factor chosen as $\mathbf{C}\mathbf{C}^T$, where \mathbf{C} is a square Gaussian random matrix. Fig. 5 shows the empirical performance




Fig. 3. Kronecker spectrum and bounds based on Lemma 1. The upper bound ‘Bound - frob’ (in green) is obtained using the bound (14) using the basis associated with the minimum ℓ_2 approximation error (i.e., the optimal basis computed by SVD as outlined in the equality condition of Lemma 1). The upper bound ‘Bound GS - frob’ (in magenta) is constructed using the variational bound (14) with projection matrix \mathbf{P}_k having columns drawn from the orthonormal basis constructed in the proof of Thm. 5. The upper bound ‘Bound GS - frob 2’ (in black) is constructed from the bound (47) in the proof of Thm. 5.

of covariance matching (CM) (i.e., solution of (6) with $r = 3$), PRLS and SVT (i.e., solution of (4)). We note that the Kronecker spectrum contains only three nonzero terms while the true covariance is full rank. The PRLS spectrum is more concentrated than the eigenspectrum and, from Fig. 5, we observe PRLS outperforms covariance matching (CM), SVT and SCM across all n .




Fig. 4. Simulation A. True dense covariance is constructed using the sum of KP model (1), with $r = 3$. Left panel: True positive definite covariance matrix Σ_0 . Middle panel: Kronecker spectrum (eigenspectrum of Σ_0 in permuted domain). Right panel: Eigenspectrum (Eigenvalues of Σ_0). Note that the Kronecker spectrum is much more concentrated than the eigenspectrum.




Fig. 5. Simulation A. Normalized MSE performance for true covariance matrix in Fig. 4 as a function of sample size n . PRLS outperforms CM, SVT (i.e., solution of (4)) and the standard SCM estimator. Here, $p = q = 25$ and $N_{MC} = 80$. For $n = 20$, PRLS achieves a 7.91 dB MSE reduction over SCM and SVT achieves a 1.80 dB MSE reduction over SCM.

B. Block Toeplitz Covariance

The covariance matrix shown in Fig. 6 was constructed by first generating a Gaussian random square matrix Φ of spectral norm $0.95 < 1$, and then simulating the block Toeplitz covariance for the process shown in (17). Fig. 7 compares the empirical performance of PRLS and SVT (i.e., the solution of (4) with appropriate scaling for the regularization parameter). We observe that the Kronecker product estimator performs much better than both SVT (i.e., the solution of (4)) and naive SCM estimator. This is most likely due to the fact that the repetitive block structure of Kronecker products better summarizes the covariance structure. We observe from Fig. 6 that for this block Toeplitz covariance, the Kronecker spectrum decays more rapidly (exponentially) than the eigenspectrum.

VI. APPLICATION TO WIND SPEED PREDICTION

In this section, we demonstrate the performance of PRLS in a real world application: wind speed prediction. We apply our methods to the Irish wind speed dataset and the NCEP dataset.

A. Irish Wind Speed Data

We use data consisting of time series consisting of daily average wind speed recordings during the period 1961 – 1978 at $q = 11$ meteorological stations. This data set has many temporal coordinates, spanning a total of $n_{total} = 365 \cdot 8 = 2920$ daily average recordings of wind speed at each station. More details on this data set can be found in [44], [45], [46], [47] and it can be downloaded from Statlib <http://lib.stat.cmu.edu/datasets>. We used the same square root transformation, estimated seasonal effect offset and station-specific mean offset as in [44], yielding the multiple (11) velocity measures. We used the data from years 1969 – 1970 for training and the data from 1971 – 1978 for testing.

The task is to predict the average velocity for the next day using the average wind velocity in each of the $p - 1$ previous days. The full dimension of each observation vector is $d = pq$, and each d -dimensional observation vector is formed by concatenating the p time-consecutive q -dimensional vectors (each entry containing the velocity measure for each station) without overlapping the time segments. The SCM was estimated using data from the training period consisting of years 1969 – 1970. Linear predictors over the time series were constructed by using these estimated covariance matrices in an ordinary least squares predictor. Specifically, we constructed the SCM linear predictor of all stations' wind velocity from the $p - 1$ previous samples of the $q = 11$ stations' time series:

$$\hat{\mathbf{v}}_t = \Sigma_{2,1} \Sigma_{1,1}^{-1} \mathbf{v}_{t-1:t-(p-1)} \quad (18)$$




Fig. 6. Simulation B. True dense block-Toeplitz covariance matrix. Left panel: True positive definite covariance matrix Σ_0 . Middle panel: Kronecker spectrum (eigenspectrum of Σ_0 in permuted domain). Right panel: Eigenspectrum (Eigenvalues of Σ_0). Note that the Kronecker spectrum is much more concentrated than the eigenspectrum.




Fig. 7. Simulation B. Normalized MSE performance for covariance matrix in Fig. 6 as a function of sample size n . PRLS outperforms SVT (i.e., solution of (4)) and the standard SCM estimator. Here, $p = q = 25$ and $N_{MC} = 80$. For $n = 108$, PRLS achieves a 6.88 dB MSE reduction over SCM and SVT achieves a 0.37 dB MSE reduction over SCM. Note again that the Kronecker spectrum is much more concentrated than the eigenspectrum.

where $\mathbf{v}_{t-1:t-(p-1)} \in \mathbb{R}^{(p-1)q}$ is the stacked wind velocities from the previous $p - 1$ time instants and $\Sigma_{2,1} \in \mathbb{R}^{q \times q(p-1)}$ and $\Sigma_{1,1} \in \mathbb{R}^{q(p-1) \times q(p-1)}$ are submatrices of the $qp \times qp$ standard SCM:

$$\hat{\mathbf{S}}_n = \begin{bmatrix} \Sigma_{1,1} & \Sigma_{1,2} \\ \Sigma_{2,1} & \Sigma_{2,2} \end{bmatrix}$$

The PRLS predictor was similarly constructed using our proposed estimator of the $qp \times qp$ Kronecker sum covariance matrix instead of the SCM. The coefficients of each of these predictors, $\Sigma_{2,1}\Sigma_{1,1}^{-1}$, were subsequently applied to predict over the test set.

The predictors were tested on the data from years 1971–1978, corresponding to $n_{test} = 365 \cdot 8 = 2920$ days, as the ground truth. Using non-overlapping samples and $p = 8$, we have a total of $n = \lceil \frac{365 \cdot 2}{p} \rceil = 91$ training samples of full dimension $d = 88$.

Fig. 8 shows the Kronecker product factors that make up the solution of Eq. (6) with $r = 1$ and the PRLS estimate. The PRLS estimate contains $r_{eff} = 6$ nonzero terms in the KP expansion. It is observed that the first order temporal factor gives a decay in correlations over time, and spatial correlations between weather stations are present. The second order temporal and spatial factors can potentially give insight into long range dependencies.

Fig. 10 shows the root mean squared error (RMSE) prediction performance over the testing period of 2920 days for the forecasts based on the standard SCM, PRLS estimator, Lounici's SVT estimator [1], and regularized Tyler [48]. The PRLS estimator was implemented using a regularization parameter $\lambda_n = C \|\hat{\mathbf{S}}_n\|_2 \sqrt{\frac{p^2 + q^2 + \log(\max(p, q, n))}{n}}$ with $C = 0.13$. The constant C was chosen by optimizing the prediction RMSE on the training set over a range of regularization parameters λ parameterized by C . The SVT estimator proposed by Lounici [1] was implemented using a regularization parameter $\lambda = C \sqrt{\text{tr}(\hat{\mathbf{S}}_n) \|\hat{\mathbf{S}}_n\|_2} \sqrt{\frac{\log(2pq)}{n}}$ with constant $C = 1.9$ optimized in a similar manner. The regularized Tyler estimator was implemented using the data-dependent shrinkage coefficient suggested in Eqn. (13) in [48]. Fig. 11 shows a sample period of 150 days. We observe that PRLS tracks the actual wind speed better than the SCM-based predictor does.

B. NCEP Wind Speed Data

We use data representative of the wind conditions in the lower troposphere (surface data at .995 sigma level) for the global grid (90°N - 90°S, 0°E - 357.5°E). We obtained the data from the National Centers for Environmental Prediction reanalysis project (Kalnay et al. [49]), which is available online at the NOAA website <ftp://ftp.cdc.noaa.gov/Datasets/ncep.reanalysis.dailyavgs/surface>. Daily averages of U (east-west) and V (north-south) wind components were collected using a station grid of size 144×73 (2.5 degree

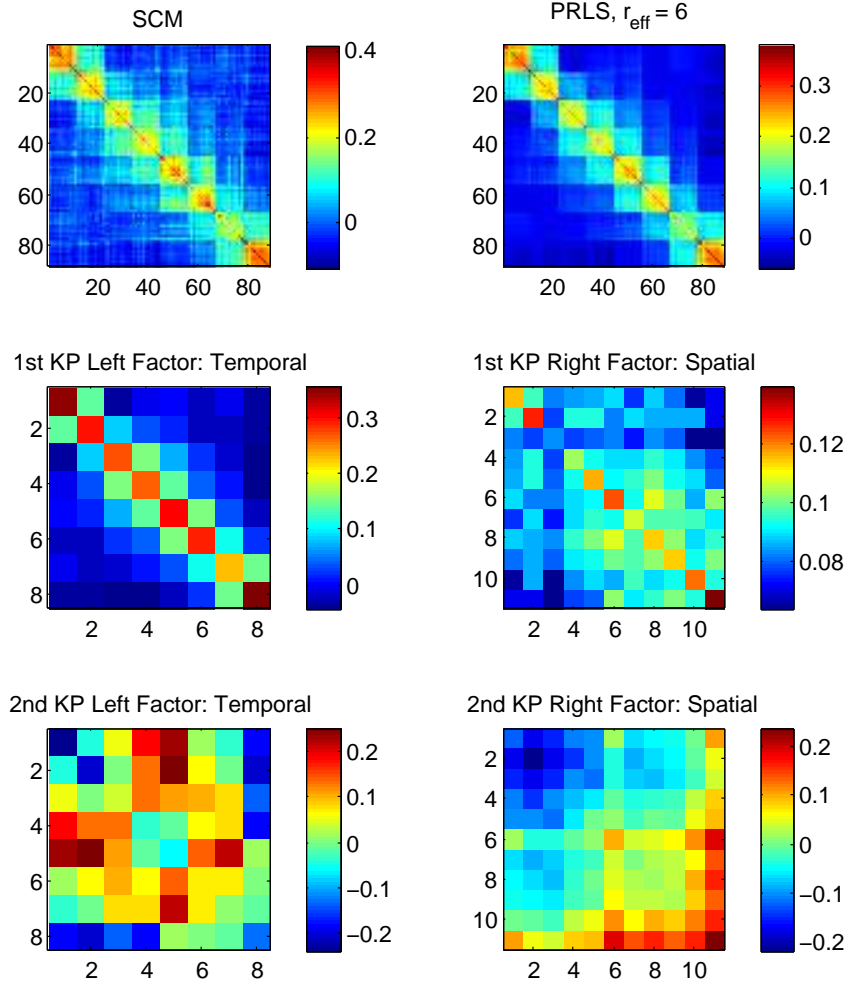


Fig. 8. Irish wind speed data: Sample covariance matrix (SCM) (top left), PRLS covariance estimate (top right), temporal Kronecker factor for first KP component (middle left) and spatial Kronecker factor for first KP component (middle right), temporal Kronecker factor for second KP component (bottom left) and spatial Kronecker factor for second KP component (bottom right). Note that the second order factors are not necessarily positive definite, although the sum of the components (i.e., the PRLS solution) is positive definite for large enough n . Each KP factor has unit Frobenius norm. Note that the plotting scales the image data to the full range of the current colormap to increase visual contrast.

latitude \times 2.5 degree longitude global grid) over the years 1948 – 2012. The wind speed is computed by taking the magnitude of the wind vector.

1) *Continental US Region*: We considered a 10×10 grid of stations, corresponding to latitude range 25°N – 47.5°N and longitude range 125°W – 97.5°W . For this selection of variables, $q = 10 \cdot 10 = 100$ is the total number of stations and $p - 1 = 7$ is the prediction time lag. We preprocessed the raw data using




Fig. 9. Irish wind speed data: Kronecker spectrum of SCM (left) and Eigenspectrum of SCM (right). The first and second KP components contain 94.60% and 1.07% of the spectrum energy. The first and second eigenvectors contain 36.28% and 28.76% of the spectrum energy. The KP spectrum is more compact than the eigenspectrum. Here, the eigenspectrum is truncated at $\min(p^2, q^2) = 8^2 = 64$ to match the Kronecker spectrum. Each spectrum was normalized such that each component has height equal to the percentage of energy associated with it.

the detrending procedure outlined in Haslett et al. [44]. More specifically, we first performed a square root transformation, then estimated and subtracted the station-specific means from the data and finally estimated and subtracted the seasonal effect (see Fig. 12). The resulting features/observations are called the velocity measures [44]. The SCM was estimated using data from the training period consisting of years 2003 – 2007. Since the SCM is not full rank, the linear predictor (18) was implemented with the Moore-Penrose pseudo-inverse of $\Sigma_{1,1}$. The predictors were tested on the data from years 2008 – 2012 as the ground truth. Using non-overlapping samples and $p = 8$, we have a total of $n = \lceil \frac{365 \cdot 5}{p} \rceil = 228$ training samples of full dimension $d = 800$.

Fig. 13 shows the Kronecker product factors that make up the solution of Eq. (6) with $r = 2$ and the PRLS covariance estimate. The PRLS estimate contains $r_{eff} = 6$ nonzero terms in the KP expansion. It is observed that the first order temporal factor gives a decay in correlations over time, and spatial correlations between weather stations are present. The second order temporal and spatial factors give some insight into longer range dependencies.




Fig. 10. Irish wind speed data: RMSE prediction performance across q stations for linear estimators using SCM (blue), PRLS (green), SVT (red) and regularized Tyler (magenta). PRLS, SVT and regularized Tyler respectively achieve an average reduction in RMSE of 3.32, 2.50 and 2.79 dB as compared to SCM (averaged across stations).

Fig. 15 shows the root mean squared error (RMSE) prediction performance over the testing period of 1825 days for the forecasts based on the standard SCM, PRLS, SVT [1] and regularized Tyler [48]. The PRLS estimator was implemented using a regularization parameter $\lambda_n = C \|\hat{\mathbf{S}}_n\|_2 \sqrt{\frac{p^2 + q^2 + \log(\max(p, q, n))}{n}}$ with $C = 0.036$. The constant C was chosen by optimizing the prediction RMSE on the training set over a range of regularization parameters λ parameterized by C (as in Irish wind speed data set). The SVT estimator proposed by Lounici [1] was implemented using a regularization parameter $\lambda = C \sqrt{\text{tr}(\hat{\mathbf{S}}_n) \|\hat{\mathbf{S}}_n\|_2} \sqrt{\frac{\log(2pq)}{n}}$ with constant $C = 0.31$ optimized in a similar manner. Fig. 16 shows a sample period of 150 days. It is observed that SCM has unstable performance, while the Kronecker product estimator offers better tracking of the wind speeds.




Fig. 11. Irish wind speed data: Prediction performance for linear estimators using SCM (blue), SVT (red) and PRLS (green) for a time interval of 150 days. The actual (ground truth) wind speeds are shown in black. PRLS offers better tracking performance as compared to SVT and SCM.




Fig. 12. NCEP wind speed data (Continental US): Seasonal effect as a function of day of the year. A 14th order polynomial is fit by the least squares method to the average of the square root of the daily mean wind speeds over all stations and over all training years.

2) *Arctic Ocean Region*: We considered a 10×10 grid of stations, corresponding to latitude range 90°N - 67.5°N and longitude range 0°E - 22.5°E . For this selection of variables, $q = 10 \cdot 10 = 100$ is the total number of stations and $p - 1 = 7$ is the prediction time lag. We preprocessed the raw data using




Fig. 13. NCEP wind speed data (Continental US): Sample covariance matrix (SCM) (top left), PRLS covariance estimate (top right), temporal Kronecker factor for first KP component (middle left) and spatial Kronecker factor for first KP component (middle right), temporal Kronecker factor for second KP component (bottom left) and spatial Kronecker factor for second KP component (bottom right). Note that the second order factors are not necessarily positive definite, although the sum of the components (i.e., the PRLS solution) is positive definite for large enough n . Each KP factor has unit Frobenius norm. Note that the plotting scales the image data to the full range of the current colormap to increase visual contrast.

the detrending procedure outlined in Haslett et al. [44]. More specifically, we first performed a square root transformation, then estimated and subtracted the station-specific means from the data and finally estimated and subtracted the seasonal effect (see Fig. 17). The resulting features/observations are called the velocity measures [44]. The SCM was estimated using data from the training period consisting of years 2003 – 2007. Since the SCM is not full rank, the linear predictor (18) was implemented with the




Fig. 14. NCEP wind speed data (Continental US): Kronecker spectrum of SCM (left) and Eigenspectrum of SCM (right). The first and second KP components contain 85.88% and 3.48% of the spectrum energy. The first and second eigenvectors contain 40.93% and 23.82% of the spectrum energy. The KP spectrum is more compact than the eigenspectrum. Here, the eigenspectrum is truncated at $\min(p^2, q^2) = 8^2 = 64$ to match the Kronecker spectrum. Each spectrum was normalized such that each component has height equal to the percentage of energy associated with it.

Moore-Penrose pseudo-inverse of $\Sigma_{1,1}$. The predictors were tested on the data from years 2008 – 2012 as the ground truth. Using non-overlapping samples and $p = 8$, we have a total of $n = \lceil \frac{365 \cdot 5}{p} \rceil = 228$ training samples of full dimension $d = 800$.

Fig. 18 shows the Kronecker product factors that make up the solution of Eq. (6) with $r = 2$ and the PRLS covariance estimate. The PRLS estimate contains $r_{eff} = 2$ nonzero terms in the KP expansion. It is observed that the first order temporal factor gives a decay in correlations over time, and spatial correlations between weather stations are present. The second order temporal and spatial factors give some insight into longer range dependencies.

Fig. 20 shows the root mean squared error (RMSE) prediction performance over the testing period of 1825 days for the forecasts based on the standard SCM, PRLS, and regularized Tyler [48]. The PRLS estimator was implemented using a regularization parameter $\lambda_n = C \|\hat{\mathbf{S}}_n\|_2 \sqrt{\frac{p^2 + q^2 + \log(\max(p, q, n))}{n}}$ with $C = 0.073$. The constant C was chosen by optimizing the prediction RMSE on the training set over a range of regularization parameters λ parameterized by C (as in Irish wind speed data set). The SVT estimator proposed by Lounici [1] was implemented using a regularization parameter $\lambda =$




Fig. 15. NCEP wind speed data (Continental US): RMSE prediction performance across q stations for linear estimators using SCM (blue), SVT (red), PRLS (green) and regularized Tyler (magenta). The estimators PRLS, SVT, and regularized Tyler respectively achieve an average reduction in RMSE of 1.90, 1.59, and 0.66 dB as compared to SCM (averaged across stations).

$C\sqrt{\text{tr}(\hat{S}_n)\|\hat{S}_n\|_2}\sqrt{\frac{\log(2pq)}{n}}$ with constant $C = 0.47$ optimized in a similar manner. Fig. 21 shows a sample period of 150 days. It is observed that SCM has unstable performance, while the Kronecker product estimator offers better tracking of the wind speeds.

VII. CONCLUSION

We have introduced a framework for covariance estimation based on separation rank decompositions using a series of Kronecker product factors. We proposed a least-squares estimator in a permuted linear space with nuclear norm penalization, named PRLS. We established high dimensional consistency for PRLS with guaranteed rates of convergence. The analysis shows that for low separation rank covariance models, our proposed method outperforms the standard SCM estimator. For the class of block-Toeplitz matrices with exponentially decaying off-diagonal norms, we showed that the separation rank is small,




Fig. 16. NCEP wind speed data (Continental US): Prediction performance for linear estimators using SCM (blue), SVT (red) and PRLS (green) for a time interval of 150 days. The actual (ground truth) wind speeds are shown in black. PRLS offers better tracking performance as compared to SCM and SVT.

and specialized our convergence bounds to this class. We also presented synthetic simulations that showed the benefits of our methods.

As a real world application we demonstrated the performance of the proposed Kronecker product-based estimator in wind speed prediction using an Irish wind speed dataset and a recent US NCEP dataset. Implementation of a standard covariance-based prediction scheme using our Kronecker product estimator achieved performance gains as compared to standard with respect to previously proposed covariance-based predictors.

There are several questions that remain open and are worthy of additional study. First, while the proposed penalized least squares Kronecker sum approximation yields a unique solution, the solution requires specification of the parameter λ , which specifies both the separation rank, and the amount of spectral shrinkage in the approximation. It would be worthwhile to investigate optimal or consistent methods of choosing this regularization parameter, e.g. using Stein's theory of unbiased risk minimization. Second, while we have proven positive definiteness of the Kronecker sum approximation when the number of samples is greater than the variable dimension in our experiments we have observed that positive definiteness is preserved more generally. Maximum likelihood estimation of Kronecker sum covariance and inverse covariance matrices is a worthwhile open problem. Finally, extensions of the low separation rank estimation method (PRLS) developed here to missing data follow naturally through the methodology of low rank covariance estimation studied in [1].




Fig. 17. NCEP wind speed data (Arctic Ocean): Seasonal effect as a function of day of the year. A 14th order polynomial is fit by the least squares method to the average of the square root of the daily mean wind speeds over all stations and over all training years.

ACKNOWLEDGEMENT

The research reported in this paper was supported in part by ARO grant W911NF-11-1-0391.

APPENDIX A

PROOF OF THEOREM 1

Proof:

1) Symmetry

Recall the permuted version of the sample covariance $\hat{\mathbf{S}}_n$, i.e., $\hat{\mathbf{R}}_n = \mathcal{R}(\hat{\mathbf{S}}_n)$. The SVD of $\hat{\mathbf{R}}_n$ can be obtained as a solution to the minimum norm problem (Thm. 1 and Cor. 2 in [32], Sec. 3 in [50]):

$$\min_{\{\mathbf{A}_k, \mathbf{B}_k\}_k} \left\| \hat{\mathbf{S}}_n - \sum_{k=1}^r \mathbf{A}_k \otimes \mathbf{B}_k \right\|_F^2 \quad (19)$$




Fig. 18. NCEP wind speed data (Arctic Ocean): Sample covariance matrix (SCM) (top left), PRLS covariance estimate (top right), temporal Kronecker factor for first KP component (middle left) and spatial Kronecker factor for first KP component (middle right), temporal Kronecker factor for second KP component (bottom left) and spatial Kronecker factor for second KP component (bottom right). Note that the second order factors are not necessarily positive definite, although the sum of the components (i.e., the PRLS solution) is positive definite for large enough n . Each KP factor has unit Frobenius norm. Note that the plotting scales the image data to the full range of the current colormap to increase visual contrast.

subject to the orthogonality constraints $\text{tr}(\mathbf{A}_k^T \mathbf{A}_l) = \text{tr}(\mathbf{B}_k^T \mathbf{B}_l) = 0$ for $k \neq l$. Since the Frobenius norm is invariant to permutations, we have the equivalent optimization problem:

$$\min_{\{\mathbf{u}_k, \mathbf{v}_k\}_k} \left\| \hat{\mathbf{R}}_n - \sum_{k=1}^r \sigma_k \mathbf{u}_k \mathbf{v}_k^T \right\|_F^2 \quad (20)$$

subject to the orthonormality conditions $\mathbf{u}_k^T \mathbf{u}_l = \mathbf{v}_k^T \mathbf{v}_l = 1$ for $k = l$ and 0 if $k \neq l$. The correspondence of (19) with (20) is given by the mapping $\mathbf{u}_k = \text{vec}(\mathbf{A}_k)$ and $\mathbf{v}_k = \sigma_k \text{vec}(\mathbf{B}_k)$. The SVD of $\hat{\mathbf{R}}_n$ can be




Fig. 19. NCEP wind speed data (Arctic Ocean): Kronecker spectrum of SCM (left) and Eigenspectrum of SCM (right). The first and second KP components contain 91.12% and 3.28% of the spectrum energy. The first and second eigenvectors contain 47.99% and 19.68% of the spectrum energy. The KP spectrum is more compact than the eigenspectrum. Here, the eigenspectrum is truncated at $\min(p^2, q^2) = 8^2 = 64$ to match the Kronecker spectrum. Each spectrum was normalized such that each component has height equal to the percentage of energy associated with it.

written in matrix form as $\mathbf{U}\Sigma\mathbf{V}^T$.

We next show that the symmetry of $\hat{\mathbf{S}}_n$ implies that the PRLS solution is symmetric by showing that the reshaped singular vectors \mathbf{u}_k and \mathbf{v}_k correspond to symmetric matrices. From the SVD definition [51], the right singular vectors \mathbf{v}_k are eigenvectors of $\mathbf{M}_n = \hat{\mathbf{R}}_n^T \hat{\mathbf{R}}_n$ and thus satisfy the eigenrelation:

$$\mathbf{M}_n \mathbf{v}_k = \sigma_k^2 \mathbf{v}_k \quad (21)$$

where $\sigma_k = [\Sigma]_{k,k}$. Expressing (21) in terms of the permutation operator \mathcal{R} , we obtain:

$$\sum_{i,j=1}^p \left\langle \mathbf{v}_k, \text{vec}(\hat{\mathbf{S}}_n(i, j)) \right\rangle \text{vec}(\hat{\mathbf{S}}_n(i, j)) = \sigma_k^2 \mathbf{v}_k \quad (22)$$

Define the $q \times q$ matrix \mathbf{V}_k such that $\mathbf{v}_k = \text{vec}(\mathbf{V}_k)$. Rewriting (22) by reshaping vectors into matrices,




Fig. 20. NCEP wind speed data (Arctic Ocean): RMSE prediction performance across q stations for linear estimators using SCM (blue) and PRLS (green). The estimators PRLS, SVT and regularized Tyler respectively achieve an average reduction in RMSE of 4.64, 3.91 and 3.41 dB as compared to SCM (averaged across stations).

we have after some algebra:

$$\begin{aligned}
 \sigma_k^2 \mathbf{V}_k &= \sum_{i,j=1}^p \text{tr}(\mathbf{V}_k^T \hat{\mathbf{S}}_n(i,j)) \hat{\mathbf{S}}_n(i,j) \\
 &= \underbrace{\sum_{i=1}^p \text{tr}(\mathbf{V}_k^T \hat{\mathbf{S}}_n(i,i)) \hat{\mathbf{S}}_n(i,i)}_{\mathbf{K}_1} \\
 &\quad + \underbrace{\sum_{i < j} \text{tr}(\mathbf{V}_k^T \hat{\mathbf{S}}_n(i,j)) (\hat{\mathbf{S}}_n(i,j) + \hat{\mathbf{S}}_n(j,i))}_{\mathbf{K}_2} \\
 &\quad + \underbrace{\sum_{i < j} \text{tr}(\mathbf{V}_k^T (\hat{\mathbf{S}}_n(j,i) - \hat{\mathbf{S}}_n(i,j))) \hat{\mathbf{S}}_n(j,i)}_{\mathbf{E}}
 \end{aligned} \tag{23}$$




Fig. 21. NCEP wind speed data (Arctic Ocean): Prediction performance for linear estimators using SCM (blue), SVT (red) and PRLS (green) for a time interval of 150 days. The actual (ground truth) wind speeds are shown in black. PRLS offers better tracking performance as compared to SCM and SVT.

Clearly, \mathbf{K}_1 is symmetric since all submatrices $\hat{\mathbf{S}}_n(i, i)$ are symmetric. Since $\hat{\mathbf{S}}_n(j, i) = \hat{\mathbf{S}}_n(i, j)^T$, it follows that \mathbf{K}_2 is also symmetric. To finish the proof, we show $\mathbf{E} = 0$. Define the set

$$\mathcal{L} = \left\{ (i, j) : i < j, \hat{\mathbf{S}}_n(i, j) \neq 0, \hat{\mathbf{S}}_n(i, j) \neq \hat{\mathbf{S}}_n(j, i), \right. \\ \left. \hat{\mathbf{S}}_n(i, j) \neq \hat{\mathbf{S}}_n(i', j') \forall i' \neq i, j' \neq j \right\}$$

The set \mathcal{L} is nonempty with probability 1 for any sample size. Let $l = \text{card}(\mathcal{L})$. Then, we can rewrite:

$$\mathbf{E} = \sum_{(i,j) \in \mathcal{L}} \text{tr}(\mathbf{V}_k^T (\hat{\mathbf{S}}_n(j, i) - \hat{\mathbf{S}}_n(i, j))) \hat{\mathbf{S}}_n(j, i) \quad (24)$$

Since $\hat{\mathbf{S}}_n(j, i) \neq 0$ with probability 1, $\mathbf{E} = 0$ iff $\text{tr}(\mathbf{V}_k^T (\hat{\mathbf{S}}_n(j, i) - \hat{\mathbf{S}}_n(i, j))) = 0$ for all $i < j$. Using the properties of the trace operator, rewriting $\text{tr}(\mathbf{V}_k^T (\hat{\mathbf{S}}_n(j, i) - \hat{\mathbf{S}}_n(i, j))) = \text{tr}((\mathbf{V}_k^T - \mathbf{V}_k) \hat{\mathbf{S}}_n(j, i))$, we conclude from the decomposition $\sigma_k^2 \mathbf{V}_k = \mathbf{K}_1 + \mathbf{K}_2 + \mathbf{E}$ that $\mathbf{V}_k = \mathbf{V}_k^T$ if $\mathbf{E} = 0$. To finish the proof, we show that $\mathbf{E} = 0$ with probability 1. Taking the $\text{vec}(\cdot)$ of (24), we conclude that $\mathbf{E} = 0$ is equivalent to

$$0 = \sum_{(i,j) \in \mathcal{L}} a_{i,j} \hat{\mathbf{S}}_n(j, i) \quad (25)$$

where $a_{i,j} = \text{tr}((\mathbf{V}_k^T - \mathbf{V}_k) \hat{\mathbf{S}}_n(j, i))$. The equation (25) can be rewritten as the linear equations:

$$\mathbf{D} \mathbf{a} = \mathbf{0} \quad (26)$$

where $\mathbf{a} = \{a_{i,j}\}_{(i,j) \in \mathcal{L}} \in \mathbb{R}^l$ and the columns of the $q^2 \times l$ matrix \mathbf{D} are given by $\mathbf{d}_{i,j} = \text{vec}(\hat{\mathbf{S}}_n(j, i)) \in \mathbb{R}^{q^2}$. Solutions of (26) are given by $\mathbf{a} \in \text{Nul}(\mathbf{D})$. Since the matrix \mathbf{D} is full-rank, $\mathbf{a} = \mathbf{0}$ is the only

solution of (25). This implies $\mathbf{E} = 0$, and therefore, $\mathbf{V}_k = \mathbf{V}_k^T$. Since k is arbitrary, all reshaped right singular vectors of $\hat{\mathbf{R}}_n$ are symmetric. A similar argument holds for all reshaped left singular vectors \mathbf{u}_k . The proof is complete.

2) Positive Definiteness

The sample covariance matrix $\hat{\mathbf{S}}_n$ is positive definite with probability 1 if $n \geq pq$. First, consider the minimum norm problem (19). The factors \mathbf{A}_k and \mathbf{B}_k are symmetric by part (1). If we show that a solution to (19) has p.d. Kronecker factors, then the weighted sum with positive scalars is also p.d. and as a result, the PRLS solution given by $\hat{\Sigma}_n^\lambda = \sum_{k=1}^{r_0} \left(\sigma_k(\hat{\mathbf{R}}_n) - \frac{\lambda}{2} \right)_+ \mathbf{U}_k \otimes \mathbf{V}_k$ is positive definite (see (7)).

Fix $l \in \{1, \dots, r_0\}$. We will show that in (19) \mathbf{A}_k and \mathbf{B}_k can be restricted to be p.d. matrices. Define the eigendecompositions of \mathbf{A}_l and \mathbf{B}_l :

$$\begin{aligned}\mathbf{A}_l &= \Psi_l \mathbf{D}_l \Psi_l^T \\ \mathbf{B}_l &= \Xi_l \mathbf{\Lambda}_l \Xi_l^T\end{aligned}$$

where $\{\Psi_l\}_l, \{\Xi_l\}_l$ are sets of orthonormal matrices and $\mathbf{D}_l, \mathbf{\Lambda}_l$ are diagonal matrices. Let $\mathbf{D}_l = \text{diag}(d_l^1, \dots, d_l^p)$ and $\mathbf{\Lambda}_l = \text{diag}(\lambda_l^1, \dots, \lambda_l^q)$. Set $\mathbf{Q}_l = \Psi_l \otimes \Xi_l$. Define $\mathbf{F}_l = \mathbf{Q}_l^T \hat{\mathbf{S}}_n \mathbf{Q}_l$. The objective

function (19) can be rewritten as:

$$\| \hat{\mathbf{S}}_n - \sum_{k=1}^r \mathbf{A}_k \otimes \mathbf{B}_k \|_F^2 \quad (27)$$

$$\begin{aligned} &= \| \mathbf{Q}_l^T \left(\hat{\mathbf{S}}_n - \sum_{k=1}^r \mathbf{A}_k \otimes \mathbf{B}_k \right) \mathbf{Q}_l \|_F^2 \\ &= \| \mathbf{F}_l - \sum_{k=1}^r \mathbf{Q}_l^T (\mathbf{A}_k \otimes \mathbf{B}_k) \mathbf{Q}_l \|_F^2 \\ &= \| \mathbf{F}_l - \underbrace{\sum_{k \neq l} (\Psi_l^T \mathbf{A}_k \Psi_l) \otimes (\Xi_l^T \mathbf{B}_k \Xi_l) - (\Psi_l^T \mathbf{A}_l \Psi_l) \otimes (\Xi_l^T \mathbf{B}_l \Xi_l)}_{\mathbf{M}_l} \|_F^2 \\ &= \| \mathbf{M}_l - \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 \\ &= \| \mathbf{M}_l \|_F^2 + \| \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 - 2\text{tr}(\mathbf{F}_l(\mathbf{D}_l \otimes \mathbf{\Lambda}_l)) + 2 \sum_{k \neq l} \text{tr}((\Psi_l^T \mathbf{A}_k \Psi_l \otimes \Xi_l^T \mathbf{B}_k \Xi_l)(\mathbf{D}_l \otimes \mathbf{\Lambda}_l)) \\ &= \| \mathbf{M}_l \|_F^2 - \| \mathbf{F}_l \|_F^2 + \| \mathbf{F}_l - \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 + 2 \sum_{k \neq l} \text{tr}(\mathbf{B}_k \mathbf{B}_l) \text{tr}(\mathbf{A}_k \mathbf{A}_l) \\ &= \| \mathbf{M}_l \|_F^2 - \| \mathbf{F}_l \|_F^2 \\ &\quad + \| \mathbf{F}_l - \text{diag}(\mathbf{F}_l) + \text{diag}(\mathbf{F}_l) - \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 \end{aligned} \quad (28)$$

$$\begin{aligned} &= \| \mathbf{M}_l \|_F^2 - \| \mathbf{F}_l \|_F^2 + \| \mathbf{F}_l - \text{diag}(\mathbf{F}_l) \|_F^2 + \| \text{diag}(\mathbf{F}_l) - \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 \\ &\quad + 2\text{tr}((\mathbf{F}_l - \text{diag}(\mathbf{F}_l))(\text{diag}(\mathbf{F}_l) - \mathbf{D}_l \otimes \mathbf{\Lambda}_l)) \\ &= \| \mathbf{M}_l \|_F^2 - \| \mathbf{F}_l \|_F^2 + \| \mathbf{F}_l - \text{diag}(\mathbf{F}_l) \|_F^2 \\ &\quad + \| \text{diag}(\mathbf{F}_l) - \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 \end{aligned} \quad (29)$$

where in equality (28) we used the orthogonality of Kronecker factors in the SVD. In equality (29), we used the fact that the matrices $\mathbf{F}_l - \text{diag}(\mathbf{F}_l)$ and $\text{diag}(\mathbf{F}_l) - \mathbf{D}_l \otimes \mathbf{\Lambda}_l$ have disjoint support. We note that the term $\| \mathbf{M}_l \|_F^2 - \| \mathbf{F}_l \|_F^2 + \| \mathbf{F}_l - \text{diag}(\mathbf{F}_l) \|_F^2$ is independent of $\mathbf{D}_l, \mathbf{\Lambda}_l$. The positive definiteness of $\hat{\mathbf{S}}_n$ implies that the diagonal elements of \mathbf{F}_l are all positive. Let $\text{diag}(\mathbf{F}_l) = \text{diag}(\{f_{(i-1)q+j}\}_{i,j}) > 0$. Simple algebra yields:

$$\begin{aligned} &\| \text{diag}(\mathbf{F}_l) - \mathbf{D}_l \otimes \mathbf{\Lambda}_l \|_F^2 \\ &= \sum_{i=1}^p \sum_{j=1}^q (f_{(i-1)q+j} - d_l^i \lambda_l^j)^2 = a_l + b_l \end{aligned}$$

where

$$a_l = \sum_{i=1}^p \sum_{j=1}^q (f_{(i-1)q+j} - |d_l^i| |\lambda_l^j|)^2$$

$$b_l = 2 \sum_{i=1}^p \sum_{j=1}^q f_{(i-1)q+j} (|d_l^i| |\lambda_l^j| - d_l^i \lambda_l^j)$$

We note that the term a_l is invariant to any sign changes of the eigenvalues $\{d_l^i, \lambda_l^j\}_{i,j}$ and the term b_l is non-negative and equals zero iff d_l^i, λ_l^j have the same sign for all i, j . By contradiction, it follows that the eigenvalues $\{d_l^i\}_{i=1}^p$ and $\{\lambda_l^j\}_{j=1}^q$ must all have the same sign (if not, then the minimum norm is not achieved by $(\mathbf{A}_l, \mathbf{B}_l)$). Without loss of generality (since $\mathbf{A}_l \otimes \mathbf{B}_l = (-\mathbf{A}_l) \otimes (-\mathbf{B}_l)$), the signs can be assumed to be positive. We conclude that there exist p.d. matrices $(\mathbf{A}_l, \mathbf{B}_l)$ that achieve the minimum norm of (27). This holds for any l so the proof is complete. ■

APPENDIX B

PROOF OF THEOREM 2

Proof: The proof generalizes Thm. 1 in [1] to nonsquare matrices. A necessary and sufficient condition for the minimizer of (5) is that there exists a $\hat{\mathbf{V}} \in \partial \|\hat{\mathbf{R}}^\lambda\|_*$ such that:

$$\left\langle 2(\hat{\mathbf{R}}^\lambda - \hat{\mathbf{R}}_n) + \lambda \hat{\mathbf{V}}, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \leq 0 \quad (30)$$

for all \mathbf{R} . From (30), we obtain for any $\mathbf{V} \in \partial \|\mathbf{R}\|_1$:

$$\begin{aligned} & 2 \left\langle \hat{\mathbf{R}}^\lambda - \mathbf{R}_0, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle + \lambda \left\langle \hat{\mathbf{V}} - \mathbf{V}, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \\ & \leq -\lambda \left\langle \mathbf{V}, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle + 2 \left\langle \hat{\mathbf{R}}_n - \mathbf{R}_0, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \end{aligned} \quad (31)$$

The monotonicity of subdifferentials of convex functions implies:

$$\left\langle \hat{\mathbf{V}} - \mathbf{V}, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \geq 0 \quad (32)$$

From Example 2 in [52], we have the characterization of the subdifferential of a nuclear norm of a nonsquare matrix:

$$\partial \|\mathbf{R}\|_* = \left\{ \sum_{j=1}^r \mathbf{u}_j(\mathbf{R}) \mathbf{v}_j(\mathbf{R})^T + \mathbf{P}_U^\perp \mathbf{W} \mathbf{P}_V^\perp : \|\mathbf{W}\|_2 \leq 1 \right\}$$

where $r = \text{rank}(\mathbf{R})$, $U = \text{span}\{\mathbf{u}_j\}$ and $V = \text{span}\{\mathbf{v}_j\}$. Thus, for $\mathbf{R} = \sum_{j=1}^r \sigma_j(\mathbf{R}) \mathbf{u}_j \mathbf{v}_j^T$, $r = \text{rank}(\mathbf{R})$, we can write:

$$\mathbf{V} = \sum_{j=1}^r \mathbf{u}_j \mathbf{v}_j^T + \mathbf{P}_U^\perp \mathbf{W} \mathbf{P}_V^\perp \quad (33)$$

where \mathbf{W} can be chosen such that $\|\mathbf{W}\|_2 \leq 1$ and

$$\left\langle \mathbf{P}_U^\perp \mathbf{W} \mathbf{P}_V^\perp, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle = \|\mathbf{P}_U^\perp \hat{\mathbf{R}}^\lambda \mathbf{P}_V^\perp\|_* \quad (34)$$

Next, note the equality:

$$\begin{aligned} \|\hat{\mathbf{R}}^\lambda - \mathbf{R}_0\|_F^2 + \|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F^2 - \|\mathbf{R} - \mathbf{R}_0\|_F^2 \\ = 2 \left\langle \hat{\mathbf{R}}^\lambda - \mathbf{R}_0, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \end{aligned} \quad (35)$$

Using (32), (34) and (35) in (31), we obtain:

$$\begin{aligned} \|\hat{\mathbf{R}}^\lambda - \mathbf{R}_0\|_F^2 + \|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F^2 + \lambda \|\mathbf{P}_U^\perp \hat{\mathbf{R}}^\lambda \mathbf{P}_V^\perp\|_* \\ \leq \|\mathbf{R} - \mathbf{R}_0\|_F^2 + \lambda \left\langle \sum_{j=1}^r \mathbf{u}_j \mathbf{v}_j^T, -(\hat{\mathbf{R}}^\lambda - \mathbf{R}) \right\rangle \\ + 2 \left\langle \hat{\mathbf{R}}^\lambda - \mathbf{R}_0, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \end{aligned} \quad (36)$$

From trace duality, we have:

$$\begin{aligned} \left\langle \sum_{j=1}^r \mathbf{u}_j \mathbf{v}_j^T, -(\hat{\mathbf{R}}^\lambda - \mathbf{R}) \right\rangle \\ = \left\langle \mathbf{P}_U \sum_{j=1}^r \mathbf{u}_j \mathbf{v}_j^T \mathbf{P}_V, -(\hat{\mathbf{R}}^\lambda - \mathbf{R}) \right\rangle \\ \leq \left\| \sum_{j=1}^r \mathbf{u}_j \mathbf{v}_j^T \right\|_2 \|\mathbf{P}_U^T (\hat{\mathbf{R}}^\lambda - \mathbf{R}) \mathbf{P}_V^T\|_* \\ = \|\mathbf{P}_U (\hat{\mathbf{R}}^\lambda - \mathbf{R}) \mathbf{P}_V\|_* \end{aligned}$$

where we used the symmetry of projection matrices. Using this bound in (36), we obtain:

$$\begin{aligned} \|\hat{\mathbf{R}}^\lambda - \mathbf{R}_0\|_F^2 + \|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F^2 + \lambda \|\mathbf{P}_U^\perp \hat{\mathbf{R}}^\lambda \mathbf{P}_V^\perp\|_* \\ \leq \|\mathbf{R} - \mathbf{R}_0\|_F^2 + \lambda \|\mathbf{P}_U (\hat{\mathbf{R}}^\lambda - \mathbf{R}) \mathbf{P}_V\|_* \\ + 2 \left\langle \Delta_n, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle \end{aligned} \quad (37)$$

where $\Delta_n = \hat{\mathbf{R}}^\lambda - \mathbf{R}_0$. Define the orthogonal projection of \mathbf{R} onto the outer product span of U and V as $\mathcal{P}_{U,V}(\mathbf{R}) = \mathbf{R} - \mathbf{P}_U^\perp \mathbf{R} \mathbf{P}_V^\perp$. Then, we decompose:

$$\begin{aligned} \left\langle \Delta_n, \hat{\mathbf{R}}^\lambda - \mathbf{R} \right\rangle &= \left\langle \Delta_n, \mathcal{P}_{U,V}(\hat{\mathbf{R}}^\lambda - \mathbf{R}) \right\rangle \\ &\quad + \left\langle \Delta_n, \mathbf{P}_U^\perp (\hat{\mathbf{R}}^\lambda - \mathbf{R}) \mathbf{P}_V^\perp \right\rangle \end{aligned}$$

By the Cauchy-Schwarz inequality and trace-duality:

$$\begin{aligned}
\|\mathbf{P}_U(\hat{\mathbf{R}}^\lambda - \mathbf{R})\mathbf{P}_V\|_* &\leq \sqrt{\text{rank}(\mathbf{R})}\|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F \\
|\langle \Delta_n, \mathcal{P}_{U,V}(\hat{\mathbf{R}}^\lambda - \mathbf{R}) \rangle| &\leq \|\Delta_n\|_2 \|\mathcal{P}_{U,V}(\hat{\mathbf{R}}^\lambda - \mathbf{R})\|_* \\
&\leq \|\Delta_n\|_2 \sqrt{2\text{rank}(\mathbf{R})}\|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F \\
|\langle \Delta_n, \mathbf{P}_U^\perp(\hat{\mathbf{R}}^\lambda - \mathbf{R})\mathbf{P}_V^\perp \rangle| &\leq \|\Delta_n\|_2 \|\mathbf{P}_U^\perp \hat{\mathbf{R}}^\lambda \mathbf{P}_V^\perp\|_*
\end{aligned}$$

where we used $\mathbf{P}_U^\perp \mathbf{R} \mathbf{P}_V^\perp = 0$. Using these bounds in (37), we further obtain:

$$\begin{aligned}
&\|\hat{\mathbf{R}}^\lambda - \mathbf{R}_0\|_F^2 + \|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F^2 + (\lambda - 2\|\Delta_n\|_2)\|\mathbf{P}_U^\perp \hat{\mathbf{R}}^\lambda \mathbf{P}_V^\perp\|_* \\
&\leq \|\mathbf{R} - \mathbf{R}_0\|_F^2 + ((2\sqrt{2}\|\Delta_n\|_2 + \lambda)\sqrt{r})(\sqrt{\|\hat{\mathbf{R}}^\lambda - \mathbf{R}\|_F^2})
\end{aligned} \tag{38}$$

Using the arithmetic-mean geometric-mean inequality in the RHS of (38) and the assumption $\lambda \geq 2\|\Delta_n\|_2$, we obtain:

$$\|\hat{\mathbf{R}}^\lambda - \mathbf{R}_0\|_F^2 \leq \|\mathbf{R} - \mathbf{R}_0\|_F^2 + \frac{\lambda^2(1 + \sqrt{2})^2}{4}r$$

This concludes the proof. ■

APPENDIX C

LEMMA 2

Lemma 2. (*Concentration of Measure for Coupled Gaussian Chaos*) Let $\mathbf{x} = [x_1, \dots, x_{p^2}]^T \in \mathcal{S}_{p^2-1}$ and $\mathbf{y} = [y_1, \dots, y_{q^2}]^T \in \mathcal{S}_{q^2-1}$. In the SCM (2) assume that $\{\mathbf{z}_t\}$ are i.i.d. multivariate normal $\mathbf{z}_t \sim N(0, \Sigma_0)$. Recall Δ_n in (10). For all $\tau \geq 0$:

$$\mathbb{P}(|\mathbf{x}^T \Delta_n \mathbf{y}| \geq \tau) \leq 2 \exp \left(\frac{-n\tau^2/2}{C_1 \|\Sigma_0\|_2^2 + C_2 \|\Sigma_0\|_2 \tau} \right) \tag{39}$$

where $C_1 = \frac{4e}{\sqrt{6\pi}} \approx 2.5044$ and $C_2 = e\sqrt{2} \approx 3.8442$ are absolute constants.

Proof: This proof is based on concentration of measure for Gaussian matrices and is similar to proof techniques used in compressed sensing (see Appendix A in [53]) and in estimation of matrix variate normal models (see Appendix C in [54]). Note that by the definition of the reshaping permutation operator $\mathcal{R}(\cdot)$, we have:

$$\Delta_n = \frac{1}{n} \sum_{t=1}^n \begin{bmatrix} \text{vec}(\mathbf{z}_t(1)\mathbf{z}_t(1)^T)^T - \mathbb{E}[\text{vec}(\mathbf{z}_t(1)\mathbf{z}_t(1)^T)^T] \\ \vdots \\ \text{vec}(\mathbf{z}_t(p)\mathbf{z}_t(p)^T)^T - \mathbb{E}[\text{vec}(\mathbf{z}_t(p)\mathbf{z}_t(p)^T)^T] \end{bmatrix}$$

where $\mathbf{z}_t(i) = [\mathbf{z}_t]_{(i-1)q+1:iq}$ is the i th subvector of the t th observation \mathbf{z}_t . Thus, we can write:

$$\mathbf{x}^T \Delta_n \mathbf{y} = \frac{1}{n} \sum_{t=1}^n \psi_t$$

where

$$\begin{aligned} \psi_t &= \sum_{i,j=1}^p \sum_{k,l=1}^q \mathbf{X}_{i,j} \mathbf{Y}_{k,l} \\ &\quad \times ([\mathbf{z}_t]_{(i-1)q+k} [\mathbf{z}_t]_{(j-1)q+l} - \mathbb{E}[[\mathbf{z}_t]_{(i-1)q+k} [\mathbf{z}_t]_{(j-1)q+l}]) \end{aligned} \quad (40)$$

and $\mathbf{X} \in \mathbb{R}^{p \times p}$ and $\mathbf{Y} \in \mathbb{R}^{q \times q}$ are reshaped versions of \mathbf{x} and \mathbf{y} . Defining $\mathbf{M} = \mathbf{X} \otimes \mathbf{Y}$, we can write (40) as:

$$\psi_t = \mathbf{z}_t^T \mathbf{M} \mathbf{z}_t - \mathbb{E}[\mathbf{z}_t^T \mathbf{M} \mathbf{z}_t]$$

The statistic (40) has the form of Gaussian chaos of order 2 [55]. Many of the random variables involved in the summation (40) are correlated, which makes the analysis difficult. To simplify the concentration of measure derivation, using the joint Gaussian property of the data, we note that a stochastic equivalent of $\mathbf{z}_t^T \mathbf{M} \mathbf{z}_t$ is $\beta_t^T \tilde{\mathbf{M}} \beta_t$, where $\tilde{\mathbf{M}} = \Sigma_0^{1/2} \mathbf{M} \Sigma_0^{1/2}$, and $\beta_t \sim N(\mathbf{0}, \mathbf{I}_{pq})$ is a random vector with i.i.d. standard normal components. With this decoupling, we have:

$$\begin{aligned} \mathbb{E}|\psi_t|^2 &= \mathbb{E} \left| \beta_t^T \tilde{\mathbf{M}} \beta_t - \mathbb{E}[\beta_t^T \tilde{\mathbf{M}} \beta_t] \right|^2 \\ &= \mathbb{E} \left| \sum_{i_1 \neq i_2} [\beta_t]_{i_1} [\beta_t]_{i_2} \tilde{\mathbf{M}}_{i_1, i_2} + \sum_{i_1=1}^d ([\beta_t]_{i_1}^2 - 1) \tilde{\mathbf{M}}_{i_1, i_1} \right|^2 \\ &= \sum_{i_1 \neq i_2} \sum_{i'_1 \neq i'_2} \mathbb{E}[[\beta_t]_{i_1} [\beta_t]_{i_2} [\beta_t]_{i'_1} [\beta_t]_{i'_2}] \tilde{\mathbf{M}}_{i_1, i_2} \tilde{\mathbf{M}}_{i'_1, i'_2} \\ &\quad + \sum_{i_1} \sum_{i'_1} \mathbb{E}([[\beta_t]_{i_1}^2 - 1][\beta_t]_{i'_1}^2 - 1)] \tilde{\mathbf{M}}_{i_1, i_1} \tilde{\mathbf{M}}_{i'_1, i'_1} \\ &= \sum_{i_1 \neq i_2} \tilde{\mathbf{M}}_{i_1, i_2}^2 + 2 \sum_{i_1} \tilde{\mathbf{M}}_{i_1, i_1}^2 \\ &= \|\tilde{\mathbf{M}}\|_F^2 + \|\text{diag}(\tilde{\mathbf{M}})\|_F^2 \\ &\leq 2\|\tilde{\mathbf{M}}\|_F^2 \leq 2\|\Sigma_0\|_2^2 \|\mathbf{M}\|_F^2 = 2\|\Sigma_0\|_2^2 \end{aligned}$$

where in the last step we used $\|\mathbf{M}\|_F = \|\mathbf{X}\|_F \|\mathbf{Y}\|_F = 1$.

Using a well known moment bound on Gaussian chaos (see p. 65 in [55]) and Stirling's formula, it can be shown (see, for example, Appendix A in [53]) that for all $m \geq 3$:

$$\mathbb{E}|\psi_t|^m \leq m! W^{m-2} v_t / 2 \quad (41)$$

where

$$W = e\sqrt{\mathbb{E}|\psi_t|^2} \leq e\sqrt{2}\|\Sigma_0\|_2$$

$$v_t = \frac{2e}{\sqrt{6\pi}}\mathbb{E}|\psi_t|^2 \leq \frac{4e}{\sqrt{6\pi}}\|\Sigma_0\|_2^2$$

From Bernstein's inequality (see Thm. 1.1 in [53]), we obtain:

$$\mathbb{P}\left(\left|\frac{1}{n}\sum_{t=1}^n\psi_t\right| \geq \tau\right) \leq 2\exp\left(\frac{-n^2\tau^2/2}{nv_1 + Wn\tau}\right)$$

$$\leq 2\exp\left(\frac{-n\tau^2/2}{C_1\|\Sigma_0\|_2^2 + C_2\|\Sigma_0\|_2\tau}\right)$$

This concludes the proof. ■

APPENDIX D

PROOF OF THEOREM 3

Proof: Let $\mathcal{N}(\mathcal{S}_{d'-1}, \epsilon')$ denote an ϵ' -net on the d' -dimensional sphere $\mathcal{S}_{d'-1}$. Let $\mathbf{x}_1 \in \mathcal{S}_{p^2-1}$ and $\mathbf{y}_1 \in \mathcal{S}_{q^2-1}$ be such that $|\mathbf{x}_1^T \Delta_n \mathbf{y}_1| = \|\Delta_n\|_2$. By the definition of ϵ' -net, there exists $\mathbf{x}_2 \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon')$ and $\mathbf{y}_2 \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon')$ such that $\|\mathbf{x}_1 - \mathbf{x}_2\|_2 \leq \epsilon'$ and $\|\mathbf{y}_1 - \mathbf{y}_2\|_2 \leq \epsilon'$. Then, by the Cauchy-Schwarz inequality:

$$|\mathbf{x}_1^T \Delta_n \mathbf{y}_1| - |\mathbf{x}_2^T \Delta_n \mathbf{y}_2| \leq |\mathbf{x}_1^T \Delta_n \mathbf{y}_1 - \mathbf{x}_2^T \Delta_n \mathbf{y}_2|$$

$$= |\mathbf{x}_1^T \Delta_n (\mathbf{y}_1 - \mathbf{y}_2) + (\mathbf{x}_1 - \mathbf{x}_2)^T \Delta_n \mathbf{y}_2|$$

$$\leq 2\epsilon' \|\Delta_n\|_2$$

Since $\|\Delta_n\|_2 = |\mathbf{x}_1^T \Delta_n \mathbf{y}_1|$, this implies:

$$\|\Delta_n\|_2(1 - 2\epsilon') \leq \max \left\{ |\mathbf{x}_2^T \Delta_n \mathbf{y}_2| : \mathbf{x}_2 \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon'), \mathbf{y}_2 \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon'), \right.$$

$$\left. \|\mathbf{x}_1 - \mathbf{x}_2\|_2 \leq \epsilon', \|\mathbf{y}_1 - \mathbf{y}_2\|_2 \leq \epsilon' \right\}$$

$$\leq \max \left\{ |\mathbf{x}^T \Delta_n \mathbf{y}| : \mathbf{x} \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon'), \mathbf{y} \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon') \right\}$$

As a result,

$$\|\Delta_n\|_2 \leq (1 - 2\epsilon')^{-1} \max_{\mathbf{x} \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon'), \mathbf{y} \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon')} |\mathbf{x}^T \Delta_n \mathbf{y}| \quad (42)$$

From Lemma 5.2 in [35], we have the bound on the cardinality of the ϵ' -net:

$$\text{card}(\mathcal{N}(\mathcal{S}_{d'-1}, \epsilon')) \leq \left(1 + \frac{2}{\epsilon'}\right)^{d'}. \quad (43)$$

From (42), (43) and the union bound:

$$\begin{aligned} & \mathbb{P}(\|\Delta_n\|_2 \geq \epsilon) \\ & \leq \mathbb{P}\left(\max_{\mathbf{x} \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon'), \mathbf{y} \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon')} |\mathbf{x}^T \Delta_n \mathbf{y}| \geq \epsilon(1 - 2\epsilon')\right) \\ & \leq \mathbb{P}\left(\bigcup_{\mathbf{x} \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon'), \mathbf{y} \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon')} |\mathbf{x}^T \Delta_n \mathbf{y}| \geq \epsilon(1 - 2\epsilon')\right) \\ & \leq \text{card}(\mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon')) \text{card}(\mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon')) \\ & \quad \times \max_{\mathbf{x} \in \mathcal{N}(\mathcal{S}_{p^2-1}, \epsilon'), \mathbf{y} \in \mathcal{N}(\mathcal{S}_{q^2-1}, \epsilon')} \mathbb{P}(|\mathbf{x}^T \Delta_n \mathbf{y}| \geq \epsilon(1 - 2\epsilon')) \\ & \leq \left(1 + \frac{2}{\epsilon'}\right)^{p^2+q^2} \mathbb{P}(|\mathbf{x}^T \Delta_n \mathbf{y}| \geq \epsilon(1 - 2\epsilon')) \end{aligned}$$

Using Lemma 2, we further obtain:

$$\begin{aligned} & \mathbb{P}(\|\Delta_n\|_2 \geq \epsilon) \\ & \leq 2 \left(1 + \frac{2}{\epsilon'}\right)^{p^2+q^2} \exp\left(\frac{-n\epsilon^2(1 - 2\epsilon')^2/2}{C_1\|\Sigma_0\|_2^2 + C_2\|\Sigma_0\|_2\epsilon(1 - 2\epsilon')}\right) \end{aligned} \quad (44)$$

We finish the proof by considering the two separate regimes. First, let us consider the Gaussian tail regime which occurs when $\epsilon \leq \frac{C_1\|\Sigma_0\|_2}{C_2(1-2\epsilon')}$. For this regime, the bound (44) can be relaxed to:

$$\begin{aligned} & \mathbb{P}(\|\Delta_n\|_2 \geq \epsilon) \\ & \leq 2 \left(1 + \frac{2}{\epsilon'}\right)^{p^2+q^2} \exp\left(\frac{-n\epsilon^2(1 - 2\epsilon')^2/2}{2C_1\|\Sigma_0\|_2^2}\right) \end{aligned} \quad (45)$$

Let us choose:

$$\epsilon = \frac{t\|\Sigma_0\|_2}{1 - 2\epsilon'} \sqrt{\frac{p^2 + q^2 + \log M}{n}}$$

Then, from (45), we have:

$$\begin{aligned} & \mathbb{P}\left(\|\Delta_n\|_2 \geq \frac{t\|\Sigma_0\|_2}{1 - 2\epsilon'} \sqrt{\frac{p^2 + q^2 + \log M}{n}}\right) \\ & \leq 2 \left(1 + \frac{2}{\epsilon'}\right)^{p^2+q^2} \exp\left(\frac{-t^2(p^2 + q^2 + \log M)}{4C_1}\right) \\ & \leq 2 \left(\left(1 + \frac{2}{\epsilon'}\right) e^{-t^2/(4C_1)}\right)^{p^2+q^2} M^{-t^2/(4C_1)} \\ & \leq 2M^{-t^2/(4C_1)} \end{aligned}$$

This concludes the bound for the Gaussian tail regime. The exponential tail regime follows by similar arguments. Assuming $\epsilon \geq \frac{C_1 \|\Sigma_0\|_2}{C_2(1-2\epsilon')}$, and setting $\epsilon = \frac{t\|\Sigma_0\|_2}{1-2\epsilon'} \frac{p^2+q^2+\log M}{n}$, we obtain from (44):

$$\begin{aligned} & \mathbb{P} \left(\|\Delta_n\|_2 \geq \frac{t\|\Sigma_0\|_2}{1-2\epsilon'} \frac{p^2+q^2+\log M}{n} \right) \\ & \leq 2 \left(1 + \frac{2}{\epsilon'} \right)^{p^2+q^2} \exp \left(\frac{-t(p^2+q^2+\log M)}{4C_2} \right) \\ & \leq 2 \left(\left(1 + \frac{2}{\epsilon'} \right) e^{-t/(4C_2)} \right)^{p^2+q^2} M^{-t/(4C_2)} \\ & \leq 2M^{-t/(4C_2)} \end{aligned}$$

where we used the assumption $t \geq 4C_2 \ln(1 + \frac{2}{\epsilon'})$. The proof is completed by combining both regimes and letting $C_0 = \|\Sigma_0\|_2$ and noting that $t > 1$, along with $\frac{tC_2}{C_1} > 1$. ■

APPENDIX E

PROOF OF THEOREM 4

Proof: Define the event

$$\mathcal{E}_r = \left\{ \|\hat{\mathbf{R}}_n^\lambda - \mathbf{R}_0\|_F^2 > \inf_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r} \|\mathbf{R} - \mathbf{R}_0\|_F^2 + \frac{(1+\sqrt{2})^2}{4} \lambda_n^2 r \right\}$$

where λ_n is chosen as in the statement of the theorem.

Theorem 2 implies that on the event $\lambda_n \geq 2\|\Delta_n\|_2$, with probability 1, we have for any $1 \leq r \leq r_0$:

$$\|\hat{\mathbf{R}}_n^\lambda - \mathbf{R}_0\|_F^2 \leq \inf_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r} \|\mathbf{R} - \mathbf{R}_0\|_F^2 + \frac{(1+\sqrt{2})^2}{4} \lambda_n^2 r$$

Using this and Theorem 3, we obtain:

$$\begin{aligned} \mathbb{P}(\mathcal{E}_r) &= \mathbb{P}(\mathcal{E}_r \cap \{\lambda_n \geq 2\|\Delta_n\|_2\}) + \mathbb{P}(\mathcal{E}_r \cap \{\lambda_n < 2\|\Delta_n\|_2\}) \\ &\leq \mathbb{P}(\mathcal{E}_r | \lambda_n \geq 2\|\Delta_n\|_2) \overset{0}{\mathbb{P}(\lambda_n \geq 2\|\Delta_n\|_2)} \\ &\quad + \mathbb{P}(\lambda_n < 2\|\Delta_n\|_2) \\ &= \mathbb{P} \left(\|\Delta_n\|_2 > \frac{C_0 t}{1-2\epsilon'} \right. \\ &\quad \left. \times \max \left\{ \frac{p^2+q^2+\log M}{n}, \sqrt{\frac{p^2+q^2+\log M}{n}} \right\} \right) \\ &\leq 2M^{-t/4C} \end{aligned}$$

This concludes the proof. ■

APPENDIX F

PROOF OF LEMMA 1

Proof: From the min-max theorem of Courant-Fischer-Weyl [56]:

$$\begin{aligned}\sigma_{k+1}^2(\mathbf{R}) &= \lambda_{k+1}(\mathbf{R}\mathbf{R}^T) \\ &= \min_{\mathcal{V}: \dim(\mathcal{V}^\perp) \leq k} \max_{\|\mathbf{v}\|_2=1, \mathbf{v} \in \mathcal{V}} \langle \mathbf{R}\mathbf{R}^T \mathbf{v}, \mathbf{v} \rangle\end{aligned}$$

Define the set

$$\mathcal{V}_k = \{\mathbf{v} \in \mathbb{R}^{p^2} : \|\mathbf{v}\|_2 = 1, \mathbf{v} \perp \text{Col}(\mathbf{R}\mathbf{P}_k\mathbf{R}^T)\} \subset S^{p^2-1}.$$

Choosing $\mathcal{V} = \text{Col}(\mathbf{R}\mathbf{P}_k\mathbf{R}^T)^\perp$, we have the upper bound:

$$\sigma_{k+1}^2(\mathbf{R}) \leq \max_{\mathbf{v} \in \mathcal{V}_k} \langle \mathbf{R}\mathbf{R}^T \mathbf{v}, \mathbf{v} \rangle$$

Using the definition of \mathcal{V}_k and the orthogonality principle, we have:

$$\begin{aligned}\langle \mathbf{R}\mathbf{R}^T \mathbf{v}, \mathbf{v} \rangle &= \langle \mathbf{R}(\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v}, \mathbf{v} \rangle \\ &= \langle (\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v}, \mathbf{R}^T \mathbf{v} \rangle \\ &= \langle (\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v}, (\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v} \rangle \\ &= \|(\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v}\|_2^2\end{aligned}$$

Using this equality and the definition of the spectral norm [56]:

$$\begin{aligned}\sigma_{k+1}^2(\mathbf{R}) &\leq \max_{\mathbf{v} \in \mathcal{V}_k} \|(\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v}\|_2^2 \\ &\leq \max_{\mathbf{v} \in S_{p^2-1}} \|(\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T \mathbf{v}\|_2^2 \\ &= \|(\mathbf{I} - \mathbf{P}_k)\mathbf{R}^T\|_2^2\end{aligned}$$

Equality follows when choosing $\mathbf{P}_k = \mathbf{V}_k\mathbf{V}_k^T$. This is seen by writing $\mathbf{I} = \mathbf{V}\mathbf{V}^T$ and using the definition of the spectral norm and the sorting of the singular values. The proof is complete. \blacksquare

APPENDIX G

PROOF OF THEOREM 5

Proof: Note that (λ, \mathbf{u}) is an eigenvalue-eigenvector pair of the square symmetric matrix $\mathbf{R}_0^T \mathbf{R}_0$ if:

$$\sum_{i,j} \text{vec}(\boldsymbol{\Sigma}_0(i,j)) \langle \mathbf{u}, \text{vec}(\boldsymbol{\Sigma}_0(i,j)) \rangle = \lambda \mathbf{u} \quad (46)$$

So for $\lambda > 0$, the eigenvector \mathbf{u} must lie in the span of the vectorized submatrices $\{\text{vec}(\Sigma_0(i, j))\}_{i,j}$. Motivated by this result, we use the Gram-Schmidt procedure to construct a basis that incrementally spans more and more of the subspace $\text{span}(\{\text{vec}(\Sigma_0(i, j))\}_{i,j})$. For the special case of the block-Toeplitz matrix, we have:

$$\text{span}(\{\text{vec}(\Sigma_0(i, j))\}_{i,j}) = \text{span}(\{\text{vec}(\Sigma(\tau))\}_{\tau=-N}^N)$$

where the mapping is given by $\Sigma_0(i, j) = \Sigma(j - i)$. Note that $\Sigma(-\tau) = \Sigma(\tau)^T$.

For simplicity, consider the case $k = 2k' + 1$ for some $k' \geq 0$. From Lemma 1, we are free to choose an orthonormal basis set $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ and form the projection matrix $\mathbf{P}_k = \mathbf{V}_k \mathbf{V}_k^T$, where the columns of \mathbf{V}_k are the vectors $\{\mathbf{v}_j\}$. We form the orthonormal basis using the Gram-Schmidt procedure [56]:

$$\begin{aligned} \tilde{\mathbf{v}}_0 &= \text{vec}(\Sigma(0)), \\ \mathbf{v}_0 &= \frac{\tilde{\mathbf{v}}_0}{\|\tilde{\mathbf{v}}_0\|_2} \\ \tilde{\mathbf{v}}_1 &= \text{vec}(\Sigma(1)) - \frac{\langle \text{vec}(\Sigma(1)), \tilde{\mathbf{v}}_0 \rangle}{\|\tilde{\mathbf{v}}_0\|_2^2} \tilde{\mathbf{v}}_0, \\ \mathbf{v}_1 &= \frac{\tilde{\mathbf{v}}_1}{\|\tilde{\mathbf{v}}_1\|_2} \\ \tilde{\mathbf{v}}_{-1} &= \text{vec}(\Sigma(-1)) - \frac{\langle \text{vec}(\Sigma(-1)), \tilde{\mathbf{v}}_0 \rangle}{\|\tilde{\mathbf{v}}_0\|_2^2} \tilde{\mathbf{v}}_0 \\ &\quad - \frac{\langle \text{vec}(\Sigma(-1)), \tilde{\mathbf{v}}_1 \rangle}{\|\tilde{\mathbf{v}}_1\|_2^2} \tilde{\mathbf{v}}_1, \\ \mathbf{v}_{-1} &= \frac{\tilde{\mathbf{v}}_{-1}}{\|\tilde{\mathbf{v}}_{-1}\|_2} \\ &\text{etc.} \end{aligned}$$

With this choice of orthonormal basis, it follows that for every $k = 2k' + 1$, we have the orthogonal projector:

$$\mathbf{P}_k = \mathbf{v}_0 \mathbf{v}_0^T + \sum_{l=1}^{k'} (\mathbf{v}_l \mathbf{v}_l^T + \mathbf{v}_{-l} \mathbf{v}_{-l}^T)$$

This corresponds to a variant of a sequence of Householder transformations [56]. Using Lemma 1:

$$\begin{aligned}
\sigma_{k+1}^2(\mathbf{R}_0) &\leq \|\mathbf{R}_0(\mathbf{I} - \mathbf{P}_k)\|_2^2 \\
&\leq \|\mathbf{R}_0 - \mathbf{R}_0\mathbf{P}_k\|_F^2 \\
&\leq p \sum_{l=k'+1}^N \|\boldsymbol{\Sigma}(l)\|_F^2 + \|\boldsymbol{\Sigma}(-l)\|_F^2 \\
&\leq 2C'pq \sum_{l=k'+1}^N u^{2l} \\
&\leq 2C'pq \frac{u^{2k'+2}}{1-u^2} \\
&\leq 2C'pq \frac{u^k}{1-u^2}
\end{aligned} \tag{47}$$

where we used Lemma 3 to obtain (47). To finish the proof, using the bound above and (13):

$$\begin{aligned}
\inf_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq r} \|\mathbf{R} - \mathbf{R}_0\|_F^2 &= \sum_{k=r}^{r_0-1} \sigma_{k+1}^2(\mathbf{R}_0) \\
&\leq \frac{2C'pq}{1-u^2} \sum_{k=r}^{r_0-1} u^k \\
&\leq 2C'pq \frac{u^r}{(1-u)^2}
\end{aligned}$$

The proof is complete. ■

APPENDIX H

LEMMA 3

Lemma 3. *Consider the notation and setting of proof of Thm. 5. Then, for the projection matrix \mathbf{P}_k chosen, we have for $k = 2k' + 1, k' \geq 1$:*

$$\sigma_{k+1}^2(\mathbf{R}_0) \leq \|\mathbf{R}_0 - \mathbf{R}_0\mathbf{P}_k\|_F^2 \leq p \sum_{l=k'+1}^N \|\boldsymbol{\Sigma}(l)\|_F^2 + \|\boldsymbol{\Sigma}(-l)\|_F^2$$

Proof: To illustrate the row-subtraction technique, we consider the simplified scenario $k' = 1$. The proof can be easily generalized to all $k' \geq 1$. Without loss of generality, we write the permuted covariance

$$\boldsymbol{\Sigma}_0 = \begin{bmatrix} \boldsymbol{\Sigma}(0) & \boldsymbol{\Sigma}(1) \\ \boldsymbol{\Sigma}(-1) & \boldsymbol{\Sigma}(0) \end{bmatrix} \tag{48}$$

as:

$$\mathbf{R}_0 = \mathcal{R}(\mathbf{\Sigma}_0) = \begin{bmatrix} \text{vec}(\mathbf{\Sigma}(0))^T \\ \text{vec}(\mathbf{\Sigma}(1))^T \\ \text{vec}(\mathbf{\Sigma}(-1))^T \\ \text{vec}(\mathbf{\Sigma}(0))^T \end{bmatrix}$$

Using the Gram-Schmidt submatrix basis construction of the proof of Thm. 5, the sequence of projection matrices can be written as:

$$\mathbf{P}_1 = \mathbf{v}_0 \mathbf{v}_0^T$$

$$\mathbf{P}_2 = \mathbf{v}_0 \mathbf{v}_0^T + \mathbf{v}_1 \mathbf{v}_1^T$$

$$\mathbf{P}_3 = \mathbf{v}_0 \mathbf{v}_0^T + \mathbf{v}_1 \mathbf{v}_1^T + \mathbf{v}_{-1} \mathbf{v}_{-1}^T$$

where \mathbf{v}_i is the orthonormal basis constructed in the proof of Thm. 5. The singular value bound $\sigma_1^2(\mathbf{R}_0) \leq \|\mathbf{R}_0\|_F^2 = 2\|\mathbf{\Sigma}(0)\|_F^2 + \|\mathbf{\Sigma}(1)\|_F^2 + \|\mathbf{\Sigma}(-1)\|_F^2$ is trivial [56].

For the second singular value, we want to prove the bound:

$$\sigma_2^2(\mathbf{R}_0) \leq \|\mathbf{\Sigma}(1)\|_F^2 + \|\mathbf{\Sigma}(-1)\|_F^2 \quad (49)$$

To show this, we use the variational bound of Lemma 1:

$$\begin{aligned} \sigma_2^2(\mathbf{R}_0) &\leq \|\mathbf{R}_0 - \mathbf{R}_0 \mathbf{P}_1\|_F^2 \\ &= \left\| \begin{bmatrix} \text{vec}(\mathbf{\Sigma}(0))^T - \langle \text{vec}(\mathbf{\Sigma}(0)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \\ \text{vec}(\mathbf{\Sigma}(1))^T - \langle \text{vec}(\mathbf{\Sigma}(1)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \\ \text{vec}(\mathbf{\Sigma}(-1))^T - \langle \text{vec}(\mathbf{\Sigma}(-1)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \\ \text{vec}(\mathbf{\Sigma}(0))^T - \langle \text{vec}(\mathbf{\Sigma}(0)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \end{bmatrix} \right\|_F^2 \\ &= \left\| \begin{bmatrix} \mathbf{0}^T \\ \text{vec}(\mathbf{\Sigma}(1))^T - \langle \text{vec}(\mathbf{\Sigma}(1)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \\ \text{vec}(\mathbf{\Sigma}(-1))^T - \langle \text{vec}(\mathbf{\Sigma}(-1)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \\ \mathbf{0}^T \end{bmatrix} \right\|_F^2 \\ &= \|\text{vec}(\mathbf{\Sigma}(1)) - \langle \text{vec}(\mathbf{\Sigma}(1)), \mathbf{v}_0 \rangle \mathbf{v}_0\|_2^2 \\ &\quad + \|\text{vec}(\mathbf{\Sigma}(-1)) - \langle \text{vec}(\mathbf{\Sigma}(-1)), \mathbf{v}_0 \rangle \mathbf{v}_0\|_2^2 \\ &\leq \|\mathbf{\Sigma}(1)\|_F^2 + \|\mathbf{\Sigma}(-1)\|_F^2 \end{aligned}$$

where in the last step, we used the Pythagorean principle from least-squares theory [57]-i.e. $\|\mathbf{A} - \frac{\langle \mathbf{A}, \mathbf{B} \rangle}{\|\mathbf{B}\|_F} \mathbf{B}\|_F^2 \leq \|\mathbf{A}\|_F^2$ for any matrices \mathbf{A}, \mathbf{B} of the same order. Next, we want to show

$$\sigma_3^2(\mathbf{R}_0) \leq \|\boldsymbol{\Sigma}(-1)\|_F^2 \quad (50)$$

Define $\gamma(j) = \text{vec}(\boldsymbol{\Sigma}(j)) - \langle \text{vec}(\boldsymbol{\Sigma}(j)), \mathbf{v}_0 \rangle \mathbf{v}_0$. Using similar bounds and the above, after some algebra:

$$\begin{aligned} \sigma_3^2(\mathbf{R}_0) &\leq \|\mathbf{R}_0 - \mathbf{R}_0 \mathbf{P}_2\|_F^2 \\ &= \left\| \begin{bmatrix} \mathbf{0}^T \\ \gamma(1)^T - \langle \text{vec}(\boldsymbol{\Sigma}(1)), \mathbf{v}_1 \rangle \mathbf{v}_1^T \\ \gamma(-1)^T - \langle \text{vec}(\boldsymbol{\Sigma}(-1)), \mathbf{v}_1 \rangle \mathbf{v}_1^T \\ \mathbf{0}^T \end{bmatrix} \right\|_F^2 \\ &= \|\text{vec}(\boldsymbol{\Sigma}(-1))^T - \langle \text{vec}(\boldsymbol{\Sigma}(-1)), \mathbf{v}_0 \rangle \mathbf{v}_0^T \\ &\quad - \langle \text{vec}(\boldsymbol{\Sigma}(-1)), \mathbf{v}_1 \rangle \mathbf{v}_1^T\|_2^2 \\ &= \|\text{vec}(\boldsymbol{\Sigma}(-1))^T\|_2^2 - |\langle \text{vec}(\boldsymbol{\Sigma}(-1)), \mathbf{v}_0 \rangle|^2 \\ &\quad - |\langle \text{vec}(\boldsymbol{\Sigma}(-1)), \mathbf{v}_1 \rangle|^2 \\ &\leq \|\boldsymbol{\Sigma}(-1)\|_F^2 \end{aligned}$$

where we observed that $\gamma(1) = \langle \text{vec}(\boldsymbol{\Sigma}(1)), \mathbf{v}_1 \rangle \mathbf{v}_1$ and used the Pythagorean principle again.

Using \mathbf{P}_3 and similar bounds, it follows that $\sigma_4^2(\mathbf{R}_0) = 0$, which makes sense since the separation rank of (48) is at most 3. Generalizing to $k' \geq 1$ and noting that $\|\boldsymbol{\Sigma}_0\|_F^2 = p\|\boldsymbol{\Sigma}(0)\|_F^2 + \sum_{l=1}^{p-1} (p-l)\|\boldsymbol{\Sigma}(l)\|_F^2 + \|\boldsymbol{\Sigma}(l)\|_F^2 \leq p\|\boldsymbol{\Sigma}(0)\|_F^2 + p \sum_{l=1}^{p-1} \|\boldsymbol{\Sigma}(l)\|_F^2 + \|\boldsymbol{\Sigma}(-l)\|_F^2$, we conclude the proof. \blacksquare

REFERENCES

- [1] K. Lounici, “High-dimensional covariance matrix estimation with missing observations,” *arXiv:1201.2577v5*, May 2012.
- [2] T. Tsiligkaridis, A. O. Hero, and S. Zhou, “On Convergence of Kronecker Graphical Lasso Algorithms,” *IEEE Transactions on Signal Processing*, vol. 61, no. 7, pp. 1743–1755, April 2013.
- [3] —, “Convergence Properties of Kronecker Graphical Lasso Algorithms,” *arXiv:1204.0585*, July 2012.
- [4] J. Bai and S. Shi, “Estimating high dimensional covariance matrices and its applications,” *Annals of Economics and Finance*, vol. 12, no. 2, pp. 199–215, 2011.
- [5] J. Xie and P. M. Bentler, “Covariance structure models for gene expression microarray data,” *Structural Equation Modeling: A Multidisciplinary Journal*, vol. 10, no. 4, pp. 556–582, 2003.
- [6] A. Hero and B. Rajaratnam, “Hub discovery in partial correlation graphs,” *IEEE Transactions on Information Theory*, vol. 58, no. 9, pp. 6064–6078, September 2012.

- [7] G. Derado, F. D. Bowman, and C. D. Kilts, "Modeling the spatial and temporal dependence in fMRI data," *Biometrics*, vol. 66, no. 3, pp. 949–957, September 2010.
- [8] Y. Zhang and J. Schneider, "Learning multiple tasks with a sparse matrix-normal penalty," *Advances in Neural Information Processing Systems*, vol. 23, pp. 2550–2558, 2010.
- [9] G. I. Allen and R. Tibshirani, "Transposable regularized covariance models with an application to missing data imputation," *The Annals of Applied Statistics*, vol. 4, no. 2, pp. 764–790, 2010.
- [10] M. Yuan and Y. Lin, "Model selection and estimation in the gaussian graphical model," *Biometrika*, vol. 94, pp. 19–35, 2007.
- [11] O. Banerjee, L. E. Ghaoui, and A. d'Aspremont, "Model selection through sparse maximum likelihood estimation for multivariate Gaussian or binary data," *Journal of Machine Learning Research*, vol. 9, pp. 485–516, March 2008.
- [12] P. Ravikumar, M. Wainwright, G. Raskutti, and B. Yu, "High-dimensional covariance estimation by minimizing ℓ_1 -penalized log-determinant divergence," *Electronic Journal of Statistics*, vol. 5, pp. 935–980, 2011.
- [13] A. Rothman, P. Bickel, E. Levina, and J. Zhu, "Sparse permutation invariant covariance estimation," *Electronic Journal of Statistics*, vol. 2, pp. 494–515, 2008.
- [14] J. Fan, Y. Fan, and J. Lv, "High dimensional covariance matrix estimation using a factor model," *Journal of Econometrics*, vol. 147, no. 1, pp. 1348–1360, 2008.
- [15] G. Fitzmaurice, N. Laird, and J. Ware, *Applied longitudinal analysis*. Wiley-Interscience, 2004.
- [16] I. Johnstone and A. Lu, "On consistency and sparsity for principal component analysis in high dimensions," *Journal of the American Statistical Association*, vol. 104, no. 486, pp. 682–693, 2009.
- [17] P. Dutilleul, "The mle algorithm for the matrix normal distribution," *J. Statist. Comput. Simul.*, vol. 64, pp. 105–123, 1999.
- [18] K. Werner, M. Jansson, and P. Stoica, "On estimation of covariance matrices with Kronecker product structure," *IEEE Transactions on Signal Processing*, vol. 56, no. 2, February 2008.
- [19] A. Dawid, "Some matrix-variate distribution theory: notational considerations and a bayesian application," *Biometrika*, vol. 68, pp. 265–274, 1981.
- [20] A. K. Gupta and D. K. Nagar, *Matrix Variate Distributions*. Chapman Hill, 1999.
- [21] K. Werner and M. Jansson, "Estimation of kronecker structured channel covariances using training data," in *Proceedings of EUSIPCO*, 2007.
- [22] N. Cressie, *Statistics for Spatial Data*. Wiley, New York, 1993.
- [23] J. Yin and H. Li, "Model selection and estimation in the matrix normal graphical model," *Journal of Multivariate Analysis*, vol. 107, pp. 119–140, 2012.
- [24] E. Bonilla, K. M. Chai, and C. Williams, "Multi-task gaussian process prediction," *Advances in Neural Information Processing Systems*, pp. 153–160, 2008.
- [25] K. Yu, J. Lafferty, S. Zhu, and Y. Gong, "Large-scale collaborative prediction using a nonparametric random effects model," *ICML*, pp. 1185–1192, 2009.
- [26] G. Beylkin and M. J. Mohlenkamp, "Algorithms for numerical analysis in high dimensions," *SIAM Journal on Scientific Computing*, vol. 26, no. 6, pp. 2133–2159, 2005.
- [27] J. C. de Munck, H. M. Huizenga, L. J. Waldorp, and R. M. Heethaar, "Estimating stationary dipoles from meg/eeeg data contaminated with spatially and temporally correlated background noise," *IEEE Transactions on Signal Processing*, vol. 50, no. 7, July 2002.
- [28] J. C. de Munck, F. Bijma, P. Gaura, C. A. Sieluzycski, M. I. Branco, and R. M. Heethaar, "A maximum-likelihood estimator

- for trial-to-trial variations in noisy meg/eeeg data sets,” *IEEE Transactions on Biomedical Engineering*, vol. 51, no. 12, 2004.
- [29] F. Bijma, J. de Munck, and R. Heethaar, “The spatiotemporal meg covariance matrix modeled as a sum of kronecker products,” *NeuroImage*, vol. 27, pp. 402–415, 2005.
 - [30] S. C. Jun, S. M. Plis, D. M. Ranken, and D. M. Schmidt, “Spatiotemporal noise covariance estimation from limited empirical magnetoencephalographic data,” *Physics in Medicine and Biology*, vol. 51, pp. 5549–5564, 2006.
 - [31] A. Rucci, S. Tebaldini, and F. Rocca, “SKP-shrinkage estimator for SAR multi-baselines applications,” in *Proceedings of IEEE Radar Conference*, 2010.
 - [32] C. V. Loan and N. Pitsianis, “Approximation with Kronecker Products,” in *Linear Algebra for Large Scale and Real Time Applications*. Kluwer Publications, 1993, pp. 293–314.
 - [33] T. Cai, C. Zhang, and H. Zhou, “Optimal rates of convergence for covariance matrix estimation,” *The Annals of Statistics*, vol. 38, no. 4, pp. 2118–2144, 2010.
 - [34] P. Bickel and E. Levina, “Regularized estimation of large covariance matrices,” *The Annals of Statistics*, vol. 36, no. 1, pp. 199–227, 2008.
 - [35] R. Vershynin, “Introduction to the non-asymptotic analysis of random matrices,” *arXiv:1011.3027v7*, November 2011.
 - [36] C. Leng and C. Y. Tang, “Sparse matrix graphical models,” *Journal of the American Statistical Association*, vol. 107, pp. 1187–1200, October 2012.
 - [37] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2004.
 - [38] S. Lee, F. Zou, and F. A. Wright, “Convergence and prediction of principal component scores in high-dimensional settings,” *The Annals of Statistics*, vol. 38, no. 6, pp. 3605–3629, 2010.
 - [39] N. R. Rao, J. A. Mingo, R. Speicher, and A. Edelman, “Statistical eigen-inference from large wishart matrices,” *The Annals of Statistics*, pp. 2850–2885, 2008.
 - [40] J.-F. Cai, E. J. Candes, and Z. Shen, “A singular value thresholding algorithm for matrix completion,” *SIAM Journal of Optimization*, vol. 20, no. 4, pp. 1956–1982, 2010.
 - [41] J.-F. Cai and S. Osher, “Fast singular value thresholding without singular value decomposition,” UCLA, Tech. Rep., 2010.
 - [42] N. Halko, P. G. Martinsson, and J. A. Tropp, “Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions,” *SIAM Review*, vol. 53, no. 2, pp. 217–288, May 2011.
 - [43] T. Tsiligkaridis, A. O. Hero, and S. Zhou, “Kronecker Graphical Lasso,” in *Proceedings of IEEE Statistical Signal Processing (SSP) Workshop*, 2012.
 - [44] J. Haslett and A. E. Raftery, “Space-time modeling with long-memory dependence: assessing ireland’s wind power resource,” *Applied Statistics*, vol. 38, no. 1, pp. 1–50, 1989.
 - [45] T. Gneiting, “Nonseparable, stationary covariance functions for space-time data,” *Journal of the American Statistical Association (JASA)*, vol. 97, no. 458, pp. 590–600, 2002.
 - [46] X. de Luna and M. Genton, “Predictive spatio-temporal models for spatially sparse environmental data,” *Statistica Sinica*, vol. 15, pp. 547–568, 2005.
 - [47] M. Stein, “Space-time covariance functions,” *Journal of the American Statistical Association (JASA)*, vol. 100, pp. 310–321, 2005.
 - [48] Y. Chen, A. Wiesel, and A. Hero, “Robust shrinkage estimation of high dimensional covariance matrices,” *IEEE Transactions on Signal Processing*, vol. 59, no. 9, pp. 4097–4107, September 2011.
 - [49] E. Kalnay, M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, Y. Zhu,

- M. Chelliah, W. Ebisuzaki, W. Higgins, J. Janowiak, K. C. Mo, C. Ropelewski, J. Wang, A. Leetmaa, R. Reynolds, R. Jenne, and D. Joseph, "The ncep/ncar 40-year reanalysis project," *Bulletin of the American Meteorological Society*, vol. 77, no. 3, p. 437471, 1996.
- [50] E. Tyrtyshnikov, "Kronecker-product approximations for some function-related matrices," *Linear Algebra and its Applications*, vol. 379, pp. 423–437, 2004.
- [51] G. H. Golub and C. V. Loan, *Matrix Computations*. JHU Press, 1996.
- [52] G. A. Watson, "Characterization of the subdifferential of some matrix norms," *Linear Algebra and Applications*, vol. 170, pp. 33–45, 1992.
- [53] H. Rauhut, K. Schnass, and P. Vandergheynst, "Compressed sensing and redundant dictionaries," *IEEE Transactions on Information Theory*, vol. 54, no. 5, pp. 2210–2219, May 2008.
- [54] S. Zhou, "Gemini: Graph estimation with matrix variate normal instances," *arXiv 1209.5075*, September 2012.
- [55] M. Ledoux and M. Talagrand, *Probability in Banach spaces. Isoperimetry and processes*. Springer-Verlag, Berlin, Heidelberg, New York, 1991.
- [56] R. A. Horn and C. R. Johnson, *Matrix Analysis*, 1st ed. Cambridge University Press, 1990.
- [57] A. Papoulis and S. U. Pillai, *Probability, Random Variables and Stochastic Processes*. Mc Graw Hill, 2002.