

Estimating covariance and precision matrices along subspaces

Željko Kereta and Timo Klock

*Simula Research Laboratory
Machine Intelligence Department
Oslo, Norway*

e-mail: z.kereta@ucl.ac.uk; timo@simula.no

Abstract: We study the accuracy of estimating the covariance and the precision matrix of a D -variate sub-Gaussian distribution along a prescribed subspace or direction using the finite sample covariance. Our results show that the estimation accuracy depends almost exclusively on the components of the distribution that correspond to desired subspaces or directions. This is relevant and important for problems where the behavior of data along a lower-dimensional space is of specific interest, such as dimension reduction or structured regression problems. We also show that estimation of precision matrices is almost independent of the condition number of the covariance matrix. The presented applications include direction-sensitive eigenspace perturbation bounds, relative bounds for the smallest eigenvalue, and the estimation of the single-index model. For the latter, a new estimator, derived from the analysis, with strong theoretical guarantees and superior numerical performance is proposed.

MSC2020 subject classifications: Primary 62H12, 62J10; secondary 62G08, 62G05, 62J12.

Keywords and phrases: Covariance matrix, finite sample bounds, dimension reduction, rate of convergence, ordinary least squares, single-index model, precision matrix.

Received January 2020.

1. Introduction

Estimating the covariance $\Sigma = \mathbb{E}(X - \mathbb{E}X)(X - \mathbb{E}X)^\top$ and the precision matrix Σ^\dagger of a random vector $X \in \mathbb{R}^D$ is a standard and long standing problem in multivariate statistics with applications in a number of mathematical and applied fields. Notable examples include any form of dimension reduction, such as principal component analysis, nonlinear dimension reduction, manifold learning, but also problems ranging from classification, regression, and signal processing to econometrics, brain imaging and social networks.

Given independent copies X_1, \dots, X_N of X , the most widely used estimator is the sample covariance $\hat{\Sigma} := \frac{1}{N} \sum_{i=1}^N X_i X_i^\top$, and the inverse thereof. The crucial question in estimating covariance and precision matrices is to quantify the minimal number of samples N ensuring that for a desired accuracy $\varepsilon > 0$, and a confidence level $u > 0$, we have

$$\|\hat{\Sigma} - \Sigma\|_2 \leq \varepsilon S_\Sigma(X), \quad \text{respectively,} \quad \|\hat{\Sigma}^\dagger - \Sigma^\dagger\|_2 \leq \varepsilon S_{\Sigma^\dagger}(X), \quad (1.1)$$

with probability at least $1 - \exp(-u)$. Constants $S_{\Sigma}(X)$ and $S_{\Sigma^\dagger}(X)$ in (1.1) describe the dependence of the error with respect to the distribution of X and properties of Σ , respectively Σ^\dagger .

In practice however, we are often not directly interested in matrices Σ or Σ^\dagger themselves, but rather in how they, as (bi)-linear operators, act on specific vectors or matrices. In terms of concentration inequalities, this can be interpreted as developing bounds for $\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top$ and $\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top$, where \mathbf{A} and \mathbf{B} are a given pair of (rectangular) matrices. Bounds of this type in case of sub-Gaussian distributions are the principal subjects of this work.

For any submultiplicative matrix norm $\|\cdot\|$ perturbations $\|\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top\|$ and $\|\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top\|$ are bounded by $\|\mathbf{A}\|\|\hat{\Sigma} - \Sigma\|\|\mathbf{B}\|$ and $\|\mathbf{A}\|\|\hat{\Sigma}^\dagger - \Sigma^\dagger\|\|\mathbf{B}\|$. This suggests that standard bounds for $\|\hat{\Sigma} - \Sigma\|$ and $\|\hat{\Sigma}^\dagger - \Sigma^\dagger\|$ (see the overviews in Sections 1.1 and 1.2) suffice so long as the distribution is nearly isotropic, i.e. $\Sigma \approx \mathbf{Id}_D$. However, many modern data analysis tasks explicitly rely on anisotropic distributions because different spectral modalities of the covariance matrix provide crucial, and complementary, information about the task at hand. In this case, using norm submultiplicativity and standard bounds for $\|\hat{\Sigma} - \Sigma\|$ and $\|\hat{\Sigma}^\dagger - \Sigma^\dagger\|$ overestimates incurred errors because it decouples \mathbf{A} and \mathbf{B} from their effect on covariance and precision matrices. Thus, such bounds cannot capture the true behavior of the estimation error.

A typical example that leverages different modalities of (conditional) covariance matrices are problems that analyze the structure of point clouds, such as manifold learning. This is because such methods are often prefaced by a linearization step, where the globally non-linear geometry is locally approximated by tangential spaces. In such a case the conditional covariance of localized data points is anisotropic, i.e. eigenvalues in tangential directions are notably larger than those in non-tangential directions, and the degree of anisotropy increases as the data is more localized [37], facilitating the linearization.

Anisotropic distributions also play an important role in high-dimensional nonparametric regression problems that use structural assumptions. Denoting Y as the dependent output variable, a popular example is the multi-index model $\mathbb{E}[Y|X = \mathbf{x}] = g(\mathbf{A}^\top \mathbf{x})$, for a matrix \mathbf{A} with $\text{rank}(\mathbf{A}) \ll D$. The complexity of the underlying nonparametric regression problem can be significantly reduced by first identifying $\text{Im}(\mathbf{A})$ and then performing nonparametric regression in $\mathbb{R}^{\text{rank}(\mathbf{A})}$. Many methods in the structural dimension reduction literature [2, 39] propose estimators for \mathbf{A} that compare the global covariance matrix Σ (typically assumed to be isotropic), with conditional covariance matrices $\text{Cov}(X|Y \in \mathcal{R})$. Here $\mathcal{R} \subset \text{Im}(Y)$ represents a connected level set of the output. The rationale behind this approach is that conditioning breaks the isotropicity and induces different spectral modalities with respect to directions belonging to $\text{Im}(\mathbf{A})$ and its orthogonal complement. This is then leveraged to identify $\text{Im}(\mathbf{A})$ [15, 33].

We showcase the usability of bounds for $\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top$ on a concrete example in Section 4 by analyzing the ordinary least squares estimator $\Sigma^\dagger \text{Cov}(X, Y)$ as an estimator for the index vector \mathbf{a} in the single-index model $\mathbb{E}[Y|X = \mathbf{x}] = g(\mathbf{a}^\top \mathbf{x})$. Our analysis extends studies [10, 5] and shows how is the accuracy of

the estimator be affected by anisotropy of X . Furthermore, an examination of developed estimation bounds motivates a modification of the ordinary least squares estimator via an approach based on conditioning and averaging in the spirit of structural dimension reduction techniques [2, 39]. We validate the superiority of this modified estimator theoretically and numerically, and thereby show how a careful tracking of different modalities of the distribution X helps to develop improved methods for common data analysis tasks.

Bounds developed in this work also have a few immediate corollaries, which might be of independent interest. These include eigenspace perturbation bounds similar to [59, Theorem 1], but which are sensitive to the behavior of X in the direction corresponding to the eigenspace of interest, and a relative bound for the smallest eigenvalue of $\hat{\Sigma}$ comparable to [58, Theorem 2.2], but without the isotropy assumption.

1.1. State of the art: covariance matrix estimation

The most common bounds for estimating the covariance matrix from finitely many observations consider sub-Gaussian [54, 52] and bounded [6] random vectors. They state that with probability at least $1 - \exp(-u)$

$$\|\hat{\Sigma} - \Sigma\|_2 \lesssim \|X\|_{\psi_2}^2 \left(\sqrt{\frac{D+u}{N}} \vee \frac{D+u}{N} \right), \quad (1.2)$$

$$\|\hat{\Sigma} - \Sigma\|_F \lesssim C_X^2 \sqrt{\frac{u}{N}}, \quad \text{provided } \|X\|_2 \leq C_X \text{ a.s.}, \quad (1.3)$$

where $A \lesssim B$ means $A \leq CB$ for some universal constant C . Besides these two cases, researchers have over the years investigated minimal moment- or tail conditions on X such that bounds similar to (1.2) can be achieved. We refer to papers [52, 49, 1, 53] that consider more general classes of distributions. The most general setting we are aware of is [49] that considers distributions which for universal $C, \eta > 0$ satisfy the tail condition

$$\mathbb{P}(\|\mathbf{P}X\|_2^2 > t) \leq Ct^{-1-\eta}, \quad \text{for } t > C \text{rank}(\mathbf{P}),$$

for every orthogonal projection \mathbf{P} . Distributions satisfying this condition include log-concave random variables (e.g. uniform distributions on convex sets) and product distributions, where the marginals have uniformly bounded $4 + s$ moments for some $s > 0$.

Our bounds show that the sample covariance estimator $\hat{\Sigma}$ automatically and implicitly adapts to $\text{rank}(\hat{\Sigma})$, which serves as a complexity parameter of the estimation problem. However, in the regime $N < \text{rank}(\Sigma)$ the sample covariance is rank deficient and the estimation is in general not possible. Instead, structural assumptions, such as sparsity or bandedness, are needed to reduce the effective complexity of the problem and allow consistent estimation. These assumptions can be leveraged by regularized estimation techniques, which include banding

[9], thresholding [13, 8], or penalized likelihood estimation [21]. We refer to [17, 13] for detailed reviews of existing methods.

We also mention [29] that considers Gaussian random variables taking values in general Banach spaces, and [30] which analyze the concentration of spectral projectors of the covariance matrix, and (bi)-linear forms thereof, for Gaussian random variables in a separable Hilbert space. This is conceptually related to our work as the bounds take into account information about relevant spectral projectors, and their interplay with vectors to which they are applied. This has recently been extended to the estimation of smooth functionals of covariance operators [27, 28] for Gaussians in separable Hilbert spaces.

1.2. State of the art: precision matrix estimation

Estimation of the precision matrix is relevant for many problems, ranging from simple tasks such as data transformations (e.g. standardization $Z := \Sigma^{-1/2}X$), to applications that include linear discriminant analysis, graphical modeling, or complex data visualization. Furthermore, precision matrix encodes information about partial correlations of features of X . Namely, if X follows a Gaussian (or parnormal) distribution, the ij -th entry of Σ^\dagger is zero if the i -th and the j -th feature are conditionally independent.

The inverse $\hat{\Sigma}^\dagger$ of the sample covariance, constructed from N independent copies of a mean zero random vector $X \in \mathbb{R}^D$, is a well-behaved estimator of Σ^\dagger as $N \rightarrow \infty$ and D is considered fixed [3]. In such a case bounds for the precision matrix can be obtained by using general perturbation bounds for the Moore-Penrose inverse. One of the first such bounds [55] states that for $\mathbf{G} \in \mathbb{R}^{d_1 \times d_2}$, and an additively perturbed matrix $\mathbf{H} = \mathbf{G} + \Delta$, we have

$$\|\mathbf{H}^\dagger - \mathbf{G}^\dagger\| \leq \omega \max \{ \|\mathbf{G}^\dagger\|_2^2, \|\mathbf{H}^\dagger\|_2^2 \} \|\Delta\|, \quad \text{and} \quad (1.4)$$

$$\|\mathbf{H}^\dagger - \mathbf{G}^\dagger\| \leq \omega \|\mathbf{G}^\dagger\|_2 \|\mathbf{H}^\dagger\|_2 \|\Delta\|, \quad \text{if } \text{rank}(\mathbf{G}) = \text{rank}(\mathbf{H}), \quad (1.5)$$

where $\|\cdot\|$ is any unitarily invariant norm, and ω is a small universal constant [42]. Recent studies [36, 57] examine the influence of the perturbation in greater detail, implying the bound

$$\|\mathbf{H}^\dagger - \mathbf{G}^\dagger\|_F \leq \min \{ \|\mathbf{H}^\dagger\|_2 \|\mathbf{G}^\dagger \Delta\|_F, \|\mathbf{G}^\dagger\|_2 \|\mathbf{H}^\dagger \Delta\|_2 \},$$

if $\text{rank}(\mathbf{G}) = \text{rank}(\mathbf{H}) = \min\{d_1, d_2\}$.

Using these general perturbation bounds it is easy to derive first concentration bounds for the precision matrix. For example, assume X is sub-Gaussian and that the number of independent data samples satisfies $N = C4^k(D + u) \|X\|_{\psi_2}^4 / \lambda_{\text{rank}(\Sigma)}(\Sigma)^2$ for some universal $C > 0$ and $k \in \mathbb{N}$. Equation (1.2) and Weyl's bound [56] imply

$$\lambda_{\text{rank}(\Sigma)}(\hat{\Sigma}) \geq \lambda_{\text{rank}(\Sigma)}(\Sigma) - \|\hat{\Sigma} - \Sigma\|_2 \geq \lambda_{\text{rank}(\Sigma)}(\Sigma)(1 - 2^{-k}),$$

and consequently $\|\hat{\Sigma}^\dagger\|_2 \leq (1 - 2^{-k})^{-1} \|\Sigma^\dagger\|_2$. The perturbation bound (1.5) and covariance bound (1.2) give with probability $1 - \exp(-u)$

$$\|\hat{\Sigma}^\dagger - \Sigma^\dagger\|_2 \lesssim \|\Sigma^\dagger\|_2^2 \|\hat{\Sigma} - \Sigma\|_2 \lesssim \|\Sigma^\dagger\|_2^2 \|X\|_{\psi_2}^2 \sqrt{\frac{D+u}{N}}, \quad (1.6)$$

where the higher order term in (1.2) vanishes in the applicable regime $N \geq C(D+u) \|X\|_{\psi_2}^4 / \lambda_{\text{rank}(\Sigma)}^2(\Sigma)$. While this effectively provides bounds as soon as the covariance perturbation is bounded, in this work we show that (1.6) overestimates the error by assigning a quadratic scaling of $\|\Sigma^\dagger\|_2$ (Corollary 11).

Moreover, we are not aware of precision matrix bounds that take into account the specific nature of the perturbation nor of bounds for $\|\mathbf{A}(\mathbf{H}^\dagger - \mathbf{G}^\dagger)\mathbf{B}\|_2$ that are sensitive to objects of interest.

If $\text{rank}(\Sigma)$ grows with N , then $\hat{\Sigma}$ is not a consistent estimator of Σ and thus the precision matrix cannot be estimated well by inverting the sample covariance matrix. Various families of structured precision matrices have been studied to mitigate these issues, motivated by applications in genomics, finance, and other fields. Dominant assumptions are sparsity and bandedness, which are exploited through the use of regularized estimators. Algorithms for estimating Σ^\dagger based on regularization include computing Σ^\dagger column by column through entry-wise Lasso [41, 18], constrained ℓ_1 minimization [11], adaptive ℓ_1 minimization [12], ℓ_1 regularized score matching [38], or ridge regressors [51]. See [17, 13] for comprehensive overviews.

1.3. Overview and contributions

Throughout, $X \in \mathbb{R}^D$ is a sub-Gaussian random vector with $\tilde{X} := X - \mathbb{E}X$ and $\Sigma = \text{Cov}(X)$, and X_1, \dots, X_N are independent copies of X . Sub-Gaussians are a broad class of light-tailed distributions that have received increasing attention in recent years and are used in many branches of probability and statistics. We define finite sample estimators of $\mathbb{E}X$ and Σ by $\hat{\mu}_X := N^{-1} \sum_{i=1}^N X_i$ and $\hat{\Sigma} := N^{-1} \sum_{i=1}^N (X_i - \hat{\mu}_X)(X_i - \hat{\mu}_X)^\top$.

Let $\mathbf{A} \in \mathbb{R}^{d_1 \times D}$, $\mathbf{B} \in \mathbb{R}^{d_2 \times D}$ be matrices determining a direction, subspace, or generally an object of interest. We can summarize our findings as follows.

- (1) In Section 2 we show that with probability at least $1 - \exp(-u)$

$$\|\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top\|_2 \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2} \|\mathbf{B}\tilde{X}\|_{\psi_2} \mathfrak{m} \left(\sqrt{\frac{d_A + d_B + u}{N}} \right), \quad (1.7)$$

where $\mathfrak{m}(t) = t \vee t^2$, $d_A := \text{rank}(\mathbf{A}\Sigma)$, and $d_B := \text{rank}(\mathbf{B}\Sigma)$. This is similar to [52, Proposition 2.1] but replaces the sub-Gaussian norm $\|\tilde{X}\|_{\psi_2}$ by direction/subspace dependent quantities $\|\mathbf{A}\tilde{X}\|_{\psi_2}$ and $\|\mathbf{B}\tilde{X}\|_{\psi_2}$.

- (2) In Section 3 we show that with probability at least $1 - \exp(-u)$, we have

$$\|\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top\|_2 \lesssim \|\mathbf{A}\Sigma^\dagger\tilde{X}\|_{\psi_2} \|\mathbf{B}\Sigma^\dagger\tilde{X}\|_{\psi_2} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}, \quad (1.8)$$

provided $N \geq C \|\sqrt{\Sigma^\dagger} \tilde{X}\|_{\psi_2}^4$, for an absolute constant $C > 0$, where $\sqrt{\Sigma^\dagger} \tilde{X}$ is the standardization of a random variable X , with $\text{Cov}(\sqrt{\Sigma^\dagger} \tilde{X}) = \mathbf{Id}_D$.

- (3) In Section 3 we show stronger bounds for (1.7) and (1.8) in case of bounded random vectors.

Remark 1. *As we point out in Corollary 11, the bound (1.8) is interesting even when $\mathbf{A} = \mathbf{B} = \mathbf{Id}_D$, particularly in view of general perturbation bounds such as (1.6). Consider for example a random vector X for which the sub-Gaussian norm is a good proxy for the variance, i.e. so that $\|\Sigma^\dagger \tilde{X}\|_{\psi_2} \approx \sqrt{\|\Sigma^\dagger\|_2}$ holds. (This is the case for example if $X \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$, or more generally for strict sub-Gaussians [4]). The right hand side of (1.8) then scales linearly in $\|\Sigma^\dagger\|_2$, whereas (1.6) shows a quadratic behavior. This has a significant impact for ill-conditioned covariance matrices and implies that inverting the sample covariance exhibits better conditioning than inverting a general matrix perturbation. The same effect is observed if \mathbf{A} and \mathbf{B} are arbitrary.*

Two applications of bounds (1.7) and (1.8) are presented. In Section 2 we use the covariance bound (1.7) to establish a bound for perturbations of eigenspaces of the covariance matrix that is sensitive to the distribution of the random vector in the eigenspace of interest. This is relevant for example when estimating manifolds from unlabeled point cloud data, see [43, 44, 25].

In Section 4.1 we use (1.8) to establish sharp concentration bounds for single-index model estimation. In this model a response $Y \in \mathbb{R}$ is assumed to follow the regression model $\mathbb{E}[Y|X] = f(\mathbf{a}^\top X)$, and the task is to estimate the unknown vector \mathbf{a} using a finite data set $\{(X_i, Y_i) : i = 1, \dots, N\}$. A common estimator is the normalized ordinary least squares vector, for which we provide direction-sensitive concentration bounds. Furthermore, our analysis yields an insight into how the estimator can be improved by a simple and straightforward procedure based on conditioning and averaging. This is presented in Section 4.2.

Most proofs are deferred to the Appendix for the sake of brevity and clarity.

1.4. General notation

We denote $[M] = \{1, \dots, M\}$, $a \wedge b = \min\{a, b\}$, $a \vee b = \max\{a, b\}$, and we may use the auxiliary function $\mathfrak{m}(t) = t \vee t^2$. For a real, symmetric matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$ we denote by $\lambda_1(\mathbf{A}) \geq \dots \geq \lambda_d(\mathbf{A})$ the ordered set of its eigenvalues, and by $\mathbf{u}_1(\mathbf{A}), \dots, \mathbf{u}_d(\mathbf{A})$ the corresponding eigenvectors. $\|\cdot\|_2$ denotes the spectral norm of a matrix, and the Euclidean norm of a vector, and $\langle \cdot, \cdot \rangle$ is the dot product. $\|\cdot\|_F$ is the Frobenius norm. $\mathbf{Id}_D \in \mathbb{R}^{D \times D}$ is the identity matrix. \mathbb{S}^{D-1} is the unit sphere in \mathbb{R}^D . For any random vector X we write $\tilde{X} := X - \mathbb{E}X$. For $p \geq 1$ and a random variable X we define the Orlicz norm

$$\|X\|_{\psi_p} := \inf\{s > 0 : \mathbb{E} \exp(|X/s|^p) \leq 2\}.$$

The definition extends to random vectors $X \in \mathbb{R}^D$ by

$$\|X\|_{\psi_p} := \sup_{\mathbf{v} \in \mathbb{S}^{D-1}} \|\mathbf{v}^\top X\|_{\psi_p} < \infty. \quad (1.9)$$

We only use $p = 1$ (sub-exponential) and $p = 2$ (sub-Gaussian). If Ω is a finite set, $|\Omega|$ denotes its cardinality. If Ω is an interval, $|\Omega|$ denotes its length. Throughout the paper C is a placeholder for a positive universal constant that may have a different value on each occurrence, even in the same line of text. Furthermore, we sometimes use $A \lesssim B$ instead of $A \leq CB$.

2. Covariance matrix estimation

In this section we present bounds for covariance and eigenspace estimation that are sensitive to the distribution of a given random vector in directions of interest. The following matrix concentration bound is the fundamental tool of our analysis.

Lemma 2. *Let $\mathbf{A} \in \mathbb{R}^{d_1 \times D}$, $\mathbf{B} \in \mathbb{R}^{d_2 \times D}$ and define $d_A = \text{rank}(\mathbf{A}\Sigma)$, and $d_B = \text{rank}(\mathbf{B}\Sigma)$. Then for all $u > 0$, with probability at least $1 - \exp(-u)$ we have for $\mathfrak{m}(t) = t \vee t^2$*

$$\|\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top\|_2 \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2} \|\mathbf{B}\tilde{X}\|_{\psi_2} \mathfrak{m} \left(\sqrt{\frac{d_A + d_B + u}{N}} \right). \quad (2.1)$$

Remark 3. *An analogous result holds for almost surely bounded random vectors by using a different concentration inequality. In that case $\|\cdot\|_{\psi_2}$ can be replaced by a bound for the Euclidean norm of X , and the dimensionality does not appear in the requirement on N . We will return to this point at the end of Section 3.*

The proof of Lemma 2 by and large follows along the lines of traditional concentration results. Indeed, if $\mathbf{A} = \mathbf{B}$, the result would follow by applying [52, Proposition 2.1] to the random vector $\mathbf{A}X$, along with some minor adjustments to account for the ranks. When $\mathbf{A} \neq \mathbf{B}$, a somewhat more careful tracking of the behavior of the random vector X , with respect directions induced by matrices \mathbf{A} and \mathbf{B} is needed. In the end, as (2.1) suggests, the payoff is that the error rate scales only with components of X along those directions.

Remark 4. *Some works that consider concentration inequalities for sub-Gaussian random variables do so with respect to the effective rank, defined as $r(\Sigma) = \text{tr}(\Sigma) / \|\Sigma\|_2$, see for instance [27] or [54, Remark 5.6.3]. Effective rank cannot exceed the true rank of a matrix, and unlike $\text{rank}(\Sigma)$, it is less affected by small eigenvalues. It can thus be a useful surrogate for approximately low dimensional distributions, and can in those cases be used to provide informative estimation bounds even if $N \ll \text{rank}(\Sigma)$.*

On the other hand, $r(\Sigma)$ is not as useful for precision matrix estimation, since inverting a matrix reverses the ordering of the eigenvalues. Thus, $r(\Sigma)$ should be replaced with $r(\Sigma^\dagger)$. Furthermore, our current proof technique for precision matrix estimation requires $\text{Im}(\hat{\Sigma}) = \text{Im}(\Sigma)$, which immediately implies $N \geq \text{rank}(\Sigma)$. To provide a unified framework, we decided to abstain from using the effective rank in our results.

Applying Lemma 2 we can reconstruct known error rates in case of low-rank distributions in high-dimensional ambient spaces.

Corollary 5. *For all $u > 0$ with probability at least $1 - \exp(-u)$ we have*

$$\|\hat{\Sigma} - \Sigma\|_2 \lesssim \|\tilde{X}\|_{\psi_2}^2 \mathfrak{m} \left(\sqrt{\frac{\text{rank}(\Sigma) + u}{N}} \right), \quad \text{where } \mathfrak{m}(t) = t \vee t^2. \quad (2.2)$$

Lemma 2 also has an immediate effect on the estimation of eigenvectors and eigenspaces. Denote by $\mathbf{P}_{i,l}(\Sigma) := \sum_{k=i}^l \mathbf{u}_k(\Sigma) \mathbf{u}_k(\Sigma)^\top$ the orthoprojector onto the space spanned by i -th to l -th eigenvectors of Σ , and let $\mathbf{P}_{i,l}(\hat{\Sigma})$ be the corresponding finite sample estimator. Moreover, denote $\mathbf{Q}_{i,l}(\hat{\Sigma}) := \mathbf{Id}_D - \mathbf{P}_{i,l}(\hat{\Sigma})$, and $\text{dist}(I_1; I_2) := \inf_{t \in I_1, t' \in I_2} |t - t'|$ for $I_1, I_2 \subset \mathbb{R}$.

Proposition 6. *Let $i \leq l \in \mathbb{N}$, and define*

$$\begin{aligned} \delta_{il} &= \text{dist}([\lambda_l(\Sigma), \lambda_i(\Sigma)]; [-\infty, \lambda_{l+1}(\hat{\Sigma})] \cup [\lambda_{i-1}(\hat{\Sigma}), +\infty]), \\ &\text{with } \lambda_0(\hat{\Sigma}) := \infty, \lambda_{D+1}(\hat{\Sigma}) = -\infty. \end{aligned} \quad (2.3)$$

For any $u > 0$, with probability at least $1 - \exp(-u)$ we have, with $\mathfrak{m}(t) = t \vee t^2$,

$$\|\mathbf{Q}_{i,l}(\hat{\Sigma}) \mathbf{P}_{i,l}(\Sigma)\|_2 \lesssim \frac{\|\mathbf{P}_{i,l}(\Sigma) \tilde{X}\|_{\psi_2} \|\tilde{X}\|_{\psi_2}}{\delta_{il}} \mathfrak{m} \left(\sqrt{\frac{\text{rank}(\Sigma) + u}{N}} \right). \quad (2.4)$$

Proof. Davis-Kahan Theorem in [7, Theorem 7.3.2] gives

$$\|\mathbf{Q}_{i,l}(\hat{\Sigma}) \mathbf{P}_{i,l}(\Sigma)\|_2 \leq \frac{\pi \|\mathbf{Q}_{i,l}(\hat{\Sigma})(\Sigma - \hat{\Sigma}) \mathbf{P}_{i,l}(\Sigma)\|_2}{\delta_{il}} \leq \frac{\pi \|(\Sigma - \hat{\Sigma}) \mathbf{P}_{i,l}(\Sigma)\|_2}{\delta_{il}}.$$

The claim now follows by applying Lemma 2 with $\mathbf{A} = \mathbf{Id}_D$ and $\mathbf{B} = \mathbf{P}_{i,l}(\Sigma)$. \square

Typical bounds for eigenspace perturbations $\|\mathbf{Q}_{i,l}(\hat{\Sigma}) \mathbf{P}_{i,l}(\Sigma)\|_2$ take the specific eigenspace into account only through the denominator, whereas the numerator relies on squared terms of the form $\|\tilde{X}\|_{\psi_2}^2$ in the sub-Gaussian case, or a bound for $\|X\|_2^2$ in the bounded case. Expression (2.4) is thus beneficial if $\|\mathbf{P}_{i,l}(\Sigma) \tilde{X}\|_{\psi_2}$ is smaller than $\|\tilde{X}\|_{\psi_2}$, as it provides a sharper estimate.

In order to ensure $\delta_{il} > 0$, the covariance matrix Σ must have a population eigengap, that is, $\delta_{il}^* := (\lambda_{i-1}(\Sigma) - \lambda_i(\Sigma)) \wedge (\lambda_l(\Sigma) - \lambda_{l+1}(\Sigma)) > 0$, and sufficiently many samples are required in order to stabilize δ_{il} around δ_{il}^* . The latter is typically achieved by first using Weyl's bound [56], giving $|\lambda_j(\hat{\Sigma}) - \lambda_j(\Sigma)| \leq \|\hat{\Sigma} - \Sigma\|_2$ for all $j \in [D]$, and then applying a concentration bound for $\|\hat{\Sigma} - \Sigma\|_2$. A consequence however is that Proposition 6 is only informative if we have sufficiently many samples with respect to δ_{il}^* and $\|\tilde{X}\|_{\psi_2}$. Thus, the estimation error is no longer sensitive to the eigenspace of interest.

Preserving the dependence on $\|\mathbf{P}_{i,l}(\Sigma) \tilde{X}\|_{\psi_2}$, instead of on $\|\tilde{X}\|_{\psi_2}$, requires the use of relative eigenvalue bounds. Such bounds have been recently provided

in [48] under the assumption that X is strictly sub-Gaussian [4], i.e. for some $K > 0$, and for arbitrary matrices $\mathbf{U} \in \mathbb{R}^{k \times D}$, X satisfies

$$\|\mathbf{U}\tilde{X}\|_{\psi_2}^2 \leq K \|\text{Cov}(\mathbf{U}\tilde{X})\|_2. \tag{2.5}$$

This asserts that the sub-Gaussian norm is a good proxy for variances of one-dimensional marginals of the random vector X , and is for instance satisfied for $X \sim \mathcal{N}(\mathbf{0}, \Sigma)$ with $K = 1$. For such distributions we obtain the following concentration of the sample eigengap.

Lemma 7. *Assume X satisfies (2.5) for some $K > 0$. Let $i, l \in \mathbb{N}$ with $1 \leq i \leq l \leq \text{rank}(\Sigma)$, and $\{i, l\} \neq \{1, \text{rank}(\Sigma)\}$. Let $\mathfrak{m}(t) = t \vee t^2$. There exists a constant $C_K > 0$, depending only on K , so that whenever*

$$N > C_K (\text{rank}(\Sigma) \vee u) \mathfrak{m} \left(\frac{\lambda_{l+1}(\Sigma)}{\lambda_l(\Sigma) - \lambda_{l+1}(\Sigma)} \vee \frac{\lambda_{i-1}(\Sigma)}{\lambda_{i-1}(\Sigma) - \lambda_i(\Sigma)} \right), \tag{2.6}$$

(with conventions $\lambda_0(\Sigma) = \infty$, $\lambda_{\text{rank}(\Sigma)+1}(\Sigma) = 0$, and $\infty/\infty = 0$) for any $u > 0$ with probability at least $1 - 2 \exp(-u)$ we have $\delta_{il} \geq \delta_{il}^*/2$.

The case $\{i, l\} = \{1, \text{rank}(\Sigma)\}$ is equivalent to asking whether $\text{Im}(\hat{\Sigma}) = \text{Im}(\Sigma)$. This holds for $N \geq \text{rank}(\Sigma)$ if the law of X has a density which is absolutely continuous with respect to the Lebesgue measure on $\text{Im}(\Sigma)$, because X_1, \dots, X_N are almost surely linearly independent. It also holds with probability at least $1 - \exp(-u)$ if X is sub-Gaussian as soon as $N > C_K(\text{rank}(\Sigma) + u)$, provided (2.5) holds, and if (2.5) does not hold then we need $N > C(\text{rank}(\Sigma) + u) \|\sqrt{\Sigma^\dagger} \tilde{X}\|_{\psi_2}^4$.

Lemma 7 now allows to refine Proposition 6 with a population eigengap.

Proposition 8. *Assume X satisfies (2.5) for some $K > 0$. Let $1 \leq i \leq l \leq \text{rank}(\Sigma)$ with $\{i, l\} \neq \{1, \text{rank}(\Sigma)\}$. Let $\mathfrak{m}(t) = t \vee t^2$. There exists $C_K > 0$, depending only on K , so that whenever N satisfies (2.6) for any $u > 0$, with probability at least $1 - \exp(-u)$ we have*

$$\begin{aligned} \|\mathbf{Q}_{i,l}(\hat{\Sigma})\mathbf{P}_{i,l}(\Sigma)\|_2 &\leq C_K \frac{\|\mathbf{P}_{i,l}(\Sigma)\tilde{X}\|_{\psi_2} \|\tilde{X}\|_{\psi_2}}{\delta_{il}^*} \mathfrak{m} \left(\sqrt{\frac{\text{rank}(\Sigma) + u}{N}} \right) \\ &\leq C_K \frac{\sqrt{K\lambda_i(\Sigma)} \|\tilde{X}\|_{\psi_2}}{\delta_{il}^*} \mathfrak{m} \left(\sqrt{\frac{\text{rank}(\Sigma) + u}{N}} \right) \end{aligned} \tag{2.7}$$

Proof. The first inequality follows immediately by first conditioning on events in Proposition 6, Lemma 7, and then using the union bound (probability $1 - 3 \exp(-u)$ can be adjusted by adjusting C_K). The second inequality in (2.7) comes from additionally using (2.5) and $\|\text{Cov}(\mathbf{P}_{i,l}(\Sigma)X)\|_2 = \lambda_i(\Sigma)$. \square

Recently, [59] showed a useful alternative

$$\|\mathbf{Q}_{i,l}(\hat{\Sigma})\mathbf{P}_{i,l}(\Sigma)\|_F \lesssim \frac{(D^{1/2} \|\hat{\Sigma} - \Sigma\|_2 \wedge \|\hat{\Sigma} - \Sigma\|_F)}{\delta_{il}^*}. \tag{2.8}$$

Thus, with (2.8) we do not have to stabilize the sample eigengap, and the eigenspace perturbation bound (2.8) can be used for arbitrary $N \geq 1$. A natural question to ask is whether Proposition 6 holds if δ_{il} is replaced with δ_{il}^* . The following example strongly suggests this is not the case.

Example 9. Assume that Proposition 6 holds with δ_{il}^* in place of δ_{il} . Let $X \sim \mathcal{N}(\mathbf{0}, \Sigma)$, where $\Sigma = \sum_{i=1}^{D-1} \mathbf{u}_i \mathbf{u}_i^\top + \eta^2 \mathbf{u}_D \mathbf{u}_D^\top$, for $\eta < 1$. Notice that in this case $\|\tilde{X}\|_{\psi_2} = 1$ and $\|\mathbf{u}_D^\top X\|_{\psi_2} = \eta$, by the definition of the sub-Gaussian norm. For any $N \geq 1$ with probability at least $1 - \exp(-u)$ we would have

$$\|\mathbf{Q}_{D,D}(\hat{\Sigma})\mathbf{P}_{D,D}(\Sigma)\|_2 \lesssim \frac{\eta}{1-\eta^2} \mathfrak{m}\left(\frac{D+u}{N}\right).$$

In particular, we have an increasingly small estimation error as $\eta \rightarrow 0$ by using only one data sample, i.e. by estimating $\mathbf{P}_{D,D}(\Sigma)$ based on the eigendecomposition of a rank one matrix XX^\top .

3. Precision matrix estimation

In this section we investigate directional estimates of the precision matrix Σ^\dagger through the empirical precision matrix $\hat{\Sigma}^\dagger$, analogously to results in Section 2.

Theorem 10. Let $\mathbf{A} \in \mathbb{R}^{d_1 \times D}$, $\mathbf{B} \in \mathbb{R}^{d_2 \times D}$. There exists a uniform constant $C > 0$ such that if $N > C(\text{rank}(\Sigma) + u)\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^4$ for any $u > 0$ with probability at least $1 - \exp(-u)$ we have

$$\|\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top\|_2 \lesssim \|\mathbf{A}\Sigma^\dagger\tilde{X}\|_{\psi_2}\|\mathbf{B}\Sigma^\dagger\tilde{X}\|_{\psi_2}\sqrt{\frac{\text{rank}(\Sigma) + u}{N}}. \quad (3.1)$$

Let us comment on the implications of Theorem 10. First, we note that $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}$ is the sub-Gaussian norm of the standardization of X . Provided X is strictly sub-Gaussian for some $K > 0$, see (2.5), we have $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^2 \leq K\|\text{Cov}(\sqrt{\Sigma^\dagger}\tilde{X})\|_2 = K$. It follows that for such distributions $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}$ has a negligible effect on estimating precision matrices.

Second, similar to Lemma 2, the bound (3.1) depends only on components of X induced by \mathbf{A} and \mathbf{B} . Since the sub-Gaussian norm can be interpreted as a proxy for the variance, this improves non-directional bounds whenever the eigenvalues of $\mathbf{A}\Sigma^\dagger\mathbf{A}^\top$ and $\mathbf{B}\Sigma^\dagger\mathbf{B}^\top$ are small compared to those of Σ^\dagger .

Third, as soon as $N > C(\text{rank}(\Sigma) + u)\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^4$, the estimation rate in (3.1) is similar to the covariance estimation rate in Lemma 2. That is, assume we are trying to estimate Σ^\dagger with the sample covariance matrix of the random vector $Z = \Sigma^\dagger X$ through iid. copies $Z_i = \Sigma^\dagger X_i$. In that case Lemma 2 gives precisely the bound (3.1). This should come as a bit of a surprise, since it implies that estimating the precision matrix through the inverse of the sample covariance has the same theoretical guarantees as if we had access to a random vector Z whose covariance is exactly Σ^\dagger . To further stress this point, we now compare

Theorem 10 for $\mathbf{A} = \mathbf{B} = \mathbf{Id}_D$ with the bound

$$\|\hat{\Sigma}^\dagger - \Sigma^\dagger\|_2 \lesssim \|\Sigma^\dagger\|_2^2 \|\Sigma\|_2 \sqrt{\frac{D+u}{N}}, \quad (3.2)$$

which was derived in Section 1.2 using general perturbation bounds for the matrix inverse.

Corollary 11. *There exists a uniform constant $C > 0$ such that if $N > C(\text{rank}(\Sigma) + u) \|\sqrt{\Sigma^\dagger} \tilde{X}\|_{\psi_2}^4$ for any $u > 0$ with probability at least $1 - \exp(-u)$ we have*

$$\|\hat{\Sigma}^\dagger - \Sigma^\dagger\|_2 \lesssim \|\Sigma^\dagger \tilde{X}\|_{\psi_2}^2 \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}. \quad (3.3)$$

Assuming the squared sub-Gaussian norm is a good proxy for the variance, i.e. if (2.5) holds, the right hand side in (3.3) becomes $K \|\Sigma^\dagger\|_2 \sqrt{(\text{rank}(\Sigma) + u)/N}$. Compared to (3.2), this shows that using a general perturbation bound overestimates the influence the matrix condition number $\|\Sigma\|_2 \|\Sigma^\dagger\|_2$ has on precision matrix estimation. The discrepancy between these two results suggests that the finite sample covariance estimator induces a specific type of a perturbation that performs a form of regularization when estimating the inverse.

An immediate corollary is a relative bound for the smallest eigenvalue of Σ .

Corollary 12. *Let $d := \text{rank}(\Sigma)$ and assume X satisfies (2.5) for some $K > 0$. There exists a constant $C > 0$ such that provided $N > C(d+u) \|\sqrt{\Sigma^\dagger} \tilde{X}\|_{\psi_2}^4$, we have with probability at least $1 - \exp(-u)$*

$$\left| \frac{\lambda_d(\Sigma)}{\lambda_d(\hat{\Sigma})} - 1 \right| \lesssim K \sqrt{\frac{d+u}{N}}. \quad (3.4)$$

If (2.5) does not hold the right hand side of (3.4) is, up to absolute constants, replaced by $\lambda_d(\Sigma) \|\Sigma^\dagger \tilde{X}\|_{\psi_2}^2 \sqrt{(d+u)/N}$.

The bound (3.4) is similar to [58, Theorem 2.2], which holds for isotropic random variables with finite fourth order moments. Moreover, it is possible to avoid the dependence on d if the number of samples N is large compared to $\sum_{i=1}^{d-1} \frac{\lambda_i(\Sigma)}{\lambda_i(\Sigma) - \lambda_d(\Sigma)}$, by using relative eigenvalue bounds from [48, Theorem 2.15].

Numerical validation To validate the results of Theorem 10 we consider $X \sim \mathcal{N}(\mathbf{0}, \Sigma)$ for two types of covariance matrices $\Sigma \in \mathbb{R}^{10 \times 10}$:

Setting 1: Set $\Sigma = \mathbf{U} \mathbf{S} \mathbf{U}^\top$, for $\mathbf{U} \in \mathbb{R}^{10 \times 10}$ sampled uniformly at random from the space of orthonormal matrices, and $\mathbf{S} = \text{Diag}(1, 1, 1, 1, 1, \nu, \nu, \nu, \nu, \nu)$. We consider $\nu = 10^{-j+1}$ for $j \in [10]$, which implies that the matrix condition number $\|\Sigma\|_2 \|\Sigma^\dagger\|_2$ ranges from 1 to 10^9 .

Setting 2: Set $\Sigma_{i,j} = \nu^{|i-j|}$, with $\nu \in \{0.5, 0.55, \dots, 0.9, 0.95\}$. This is a common model for distributions where entries of X correspond to values of a certain feature at different time stamps. It leads to correlated entries when the time stamps are close by, i.e. when $|i-j|$ is small.




FIG 1. Relative errors for directional and isotropic precision matrix estimation. Different colors correspond to different directions of evaluation (see legend). Per each color we plot 10 lines corresponding to the 10 parameter choices of ν in the construction of Σ . The plots show that the relative error does not depend on ν , despite the fact that changing ν changes the sub-Gaussian condition number by a factor of 10.

We choose matrices \mathbf{A} and \mathbf{B} by sampling an orthoprojector \mathbf{A} uniformly at random with $\text{rank}(\mathbf{A}) = 3$ and setting $\mathbf{B} = \mathbf{Id}_D - \mathbf{A}$. We repeat each experiment 100 times and report the averaged relative errors in Figure 1. Different lines of the same color correspond to different values of ν .

The estimation error shows the expected $N^{-1/2}$ rate. Moreover, directional errors show a clear dependence on the directional spectral norm. On the other hand, since we are sampling Gaussians the squared sub-Gaussian norm is a proxy for the variance of the given random vector, and the results confirm that the error term in Theorem 10 indeed scales with the corresponding (directional) sub-Gaussian norm. Lastly, although the condition number of Σ depends on ν , the specific value of ν does not affect how accurate $\hat{\Sigma}^\dagger$ is, as predicted by the theory.

Bounded random vectors As mentioned in Remark 3, a stronger form of direction dependent covariance and precision matrix estimation bounds hold for bounded random vectors. The proofs for the bounded case follow along similar lines as for the sub-Gaussian case, except for the use of a slightly different probabilistic argument. We now state only the results for the estimation of covariance and precision matrix, since the remaining bounds follow by analogy.

Theorem 13. Let $\mathbf{A} \in \mathbb{R}^{d_1 \times D}$, $\mathbf{B} \in \mathbb{R}^{d_2 \times D}$. Assume $\|\mathbf{A}\tilde{X}\|_2 \leq C_A$, $\|\mathbf{B}\tilde{X}\|_2 \leq C_B$ almost surely. Then for any $u > 0$ with probability at least $1 - \exp(-u)$ we have

$$\|\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top\|_F \lesssim C_A C_B \sqrt{\frac{1+u}{N}}. \quad (3.5)$$

Assume $\|\mathbf{A}\Sigma^\dagger\tilde{X}\|_2 \leq C_A^\dagger$, $\|\mathbf{B}\Sigma^\dagger\tilde{X}\|_2 \leq C_B^\dagger$, $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_2^2 \leq \Theta$ almost surely. There exists $C > 0$ such that provided $N > C(1+u)\Theta^2$, for any $u > 0$ with probability

at least $1 - \exp(-u)$ we have

$$\|\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top\|_F \lesssim C_A^\dagger C_B^\dagger \sqrt{\frac{1+u}{N}}. \quad (3.6)$$

4. Application to single-index model estimation

In this section we use the results of Sections 2 and 3 to establish concentration bounds for estimating the index vector \mathbf{a} in the single-index model $\mathbb{E}[Y|X = x] = g(\mathbf{a}^\top x)$. Moreover, directional terms that arise from using directional dependent matrix concentration bounds provide an insight into how to further improve the performance of the standard estimator. The second half of the section is thus devoted to describing this strategy (which is based on splitting up the data, conditioning and averaging), proving the error estimates, and providing numerical evidence to show and examine the claimed performance gains.

Throughout the section, we use a directional sub-Gaussian condition number

$$\kappa(\mathbf{P}, X) := \|\mathbf{P}\Sigma^\dagger \tilde{X}\|_{\psi_2}^2 \|\mathbf{P}\tilde{X}\|_{\psi_2}^2, \quad \mathbf{P} \text{ is an orthoprojector.}$$

This quantity can be seen as a restricted condition number $\|\mathbf{P}\Sigma^\dagger \mathbf{P}\|_2 \|\mathbf{P}\Sigma \mathbf{P}\|_2$ for the matrix $\Sigma = \text{Cov}(X)$. Indeed, for strict sub-Gaussians, i.e. those satisfying (2.5) for some K , the two are equal up to a factor depending on K .

4.1. Ordinary least squares for the single-index model

Single-index model (SIM) is a popular semi-parametric regression model that poses the relationship between the features $X \in \mathbb{R}^D$ and responses $Y \in \mathbb{R}$ as $\mathbb{E}[Y|X] = f(\mathbf{a}^\top X)$, where f is an unknown link function and $\mathbf{a} \in \mathbb{S}^{D-1}$ is an unknown index vector. SIM was developed in the 80s and 90s [10, 22] as an extension of generalized linear regression that does not specify the link function, and which could thus avoid errors incurred by model misspecification. Common applications are in econometrics [16, 40] and signal processing under sparsity assumptions on the index vector [46, 45]. It has been shown, e.g. in [19], that (in certain scenarios) the minimax estimation rate of SIM equals that of nonparametric univariate regression.

Methods for estimating the SIM from a finite data set $\{(X_i, Y_i) : i \in [N]\}$ often first construct an approximate index vector $\hat{\mathbf{a}}$, and then use nonparametric regression on $\{(\hat{\mathbf{a}}^\top X_i, Y_i) : i \in [N]\}$ to estimate the link function. With such an approach the generalization error of the resulting estimator depends largely on the error incurred by estimating the index vector. Thus, the construction of $\hat{\mathbf{a}}$ becomes the critical point.

An efficient approach, which first appeared in [35, 10], and later in modified forms in [20, 5], is to solve the ordinary least squares (OLS) problem

$$(\hat{c}, \hat{\mathbf{b}}) = \underset{c \in \mathbb{R}, \mathbf{b} \in \mathbb{R}^D}{\operatorname{argmin}} \sum_{i=1}^N (Y_i - c - \mathbf{b}^\top (X_i - \hat{\mu}_X))^2, \quad \text{where } \hat{\mu}_X = \sum_{i=1}^N \frac{X_i}{N}, \quad (4.1)$$

and then set $\hat{\mathbf{d}} := \hat{\mathbf{b}}/\|\hat{\mathbf{b}}\|_2$. It was shown in [5] that $\sqrt{N}(\hat{\mathbf{d}} - \mathbf{a})$ is asymptotically normal with mean zero, provided X has an elliptical distribution and f is a non-decreasing function that is strictly increasing on some non-empty sub-interval of the support of $\mathbf{a}^\top X$. Assuming only ellipticity of X and statistical independence of Y and X given $\mathbf{a}^\top X$, the population vector $\mathbf{b} := \Sigma^\dagger \text{Cov}(X, Y)$ is still contained in $\text{span}\{\mathbf{a}\}$, see e.g. [26, Proposition 3]. Provided $\mathbf{b} \neq 0$, in these cases the direction $\mathbf{d} := \mathbf{b}/\|\mathbf{b}\|_2$ equals the index vector \mathbf{a} up to sign. Thus, $\hat{\mathbf{d}}$ is in many cases a consistent estimator of the index vector \mathbf{a} , and under certain conditions it converges with an $N^{-1/2}$ rate.

The minimal $\|\cdot\|_2$ -norm solution of (4.1) admits a closed form

$$\hat{c} = \hat{\mu}_Y, \quad \hat{\mathbf{b}} = \hat{\Sigma}^\dagger \hat{\mathbf{r}}, \quad \text{with } \hat{\mu}_Y := \sum_{i=1}^N \frac{Y_i}{N}, \quad \hat{\mathbf{r}} := \sum_{i=1}^N \frac{(X_i - \hat{\mu}_X)(Y_i - \hat{\mu}_Y)}{N}. \quad (4.2)$$

Using the results of the previous two sections we can show a direction dependent concentration bound for the vector $\hat{\mathbf{b}}$.

Lemma 14. *Let $Y \in \mathbb{R}$ be sub-Gaussian. Denote $\mathbf{P} = \mathbf{d}\mathbf{d}^\top$, $\mathbf{Q} := \mathbf{Id}_D - \mathbf{P}$, and $\kappa_{\mathbf{P}\mathbf{Q}} = \kappa(\mathbf{P}, X) \vee \kappa(\mathbf{Q}, X)$. There exists $C > 0$ such that provided $N > C(\text{rank}(\Sigma) + u)\|\sqrt{\Sigma^\dagger} \tilde{X}\|_{\psi_2}^4$, for any $u > 0$, with probability at least $1 - \exp(-u)$ we have*

$$\|\mathbf{P}(\mathbf{b} - \hat{\mathbf{b}})\|_2 \lesssim \|\tilde{Y}\|_{\psi_2} \|\mathbf{P}\Sigma^\dagger \tilde{X}\|_{\psi_2} \sqrt{\kappa_{\mathbf{P}\mathbf{Q}}} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}, \quad (4.3)$$

$$\|\mathbf{Q}(\mathbf{b} - \hat{\mathbf{b}})\|_2 \lesssim \|\tilde{Y}\|_{\psi_2} \|\mathbf{Q}\Sigma^\dagger \tilde{X}\|_{\psi_2} \sqrt{\kappa_{\mathbf{P}\mathbf{Q}}} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}. \quad (4.4)$$

For the normalized vector, respectively the directions, the following bound holds.

Corollary 15. *Assume the setting of Lemma 14. There exists a constant $C > 0$ such that for any $u > 0$, with probability at least $1 - \exp(-u)$ we have*

$$\|\hat{\mathbf{d}} - \mathbf{d}\|_2 \lesssim \frac{\|\tilde{Y}\|_{\psi_2} \|\mathbf{Q}\Sigma^\dagger \tilde{X}\|_{\psi_2} \sqrt{\kappa_{\mathbf{P}\mathbf{Q}}} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}}{\|\mathbf{b}\|_2}, \quad \text{provided that} \quad (4.5)$$

$$N > C(\text{rank}(\Sigma) + u) \left(\|\sqrt{\Sigma^\dagger} \tilde{X}\|_{\psi_2}^4 \vee \frac{\|\tilde{Y}\|_{\psi_2}^2 \|\mathbf{P}\Sigma^\dagger \tilde{X}\|_{\psi_2}^2 \kappa_{\mathbf{P}\mathbf{Q}}}{\|\mathbf{b}\|_2^2} \right). \quad (4.6)$$

As mentioned before, under certain conditions we have $\mathbf{d} = \mathbf{a}$, where \mathbf{a} is the index vector in the given SIM. In such cases Corollary 15 confirms that $\hat{\mathbf{d}}$ is a consistent estimator of \mathbf{a} and achieves a $N^{-1/2}$ convergence rate, which has been observed in previous works. To provide some context for our result we give a brief description of the two most popular strategies for estimating \mathbf{a} .

The first group of methods can be distinguished by their simplicity and efficiency as they estimate the index vector using empirical estimates of first and second order moments of random variables X and Y . The studied OLS estimator (sometimes also referred to as the average derivative estimator [10, 35])

belongs to this group. Inverse regression based techniques also fall into this category, see e.g. [34, 15, 33] or the review [39], and their convergence rate usually equals $N^{-1/2}$ and is thus on par with OLS. A drawback caused essentially by the simplicity of these methods is that theoretical guarantees typically require X to be an elliptical distribution, and sometimes a form of monotonicity is needed to avoid issues that arise when dealing with symmetric functions such as $\mathbf{x} \mapsto (\mathbf{a}^\top \mathbf{x})^2$.

The second, more sophisticated but also computationally heavier, class of methods is based on solving more complicated optimal programs for the recovery of \mathbf{a} , to which a closed form solution does not exist. This includes non-parametric methods for sufficient dimension reduction, e.g. [14], aiming at estimating the gradients of $f(\mathbf{a}^\top X)$ at all training samples X_i ; methods based on combining index estimation and monotonic regression, e.g. [24, 23]; and methods that simultaneously estimate the index vector and use spline interpolation [47, 32], to name just a few. For these methods convergence rates up to $N^{-1/2}$ (sometimes slower) can typically be proven, and the theory requires less stringent assumptions than those exhibited by the first class of methods.

A new insight in Corollary 15, which has not been emphasized before, is the directionally-dependent influence of the spectrum of Σ on the estimation guarantee. To make this more precise, we now consider a special case when X is strictly sub-Gaussian, the link function f is Lipschitz-smooth, and the index \mathbf{a} is an eigenvector of Σ . We note that similar results hold if \mathbf{a} is an approximate eigenvector of Σ , in the sense that $\mathbf{a}^\top \Sigma^\dagger \mathbf{a} \approx (\mathbf{a}^\top \Sigma \mathbf{a})^{-1}$.

Corollary 16. *Let Y be sub-Gaussian, X be strictly sub-Gaussian (i.e. satisfying (2.5) for some $K > 0$) and assume $Y = f(\mathbf{a}^\top X) + \zeta$ for $\mathbb{E}\zeta = 0$ with $\sigma_\zeta := \|\zeta\|_{\psi_2} < \infty$. Assume consistent estimation, i.e. $\hat{\mathbf{d}} = \mathbf{a}$, and that $\Sigma \mathbf{a} = \sigma_{\mathbf{P}}^2 \mathbf{a}$. Then there exists $C_K > 0$, depending only on K , so that if*

$$N > C_K(\text{rank}(\Sigma) + u) \frac{(L + \sigma_\zeta \sigma_{\mathbf{P}}^{-1})^2 \kappa(\mathbf{Q}, X)}{\|\mathbf{b}\|_2^2},$$

for any $u > 0$ we have with probability at least $1 - \exp(-u)$

$$\|\hat{\mathbf{d}} - \mathbf{a}\|_2 \lesssim \frac{(L\sigma_{\mathbf{P}} + \sigma_\zeta) \sqrt{\|\mathbf{Q}\Sigma^\dagger \mathbf{Q}\|_2 \kappa(\mathbf{Q}, X)}}{\|\mathbf{b}\|_2} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}. \quad (4.7)$$

Corollary 16 shows that the variance in the direction of the index vector, $\sigma_{\mathbf{P}}^2 = \text{Var}(\mathbf{a}^\top X)$, and the variance in orthogonal directions, quantified through the spectrum of $\mathbf{Q}\Sigma^\dagger \mathbf{Q}$, influence index vector estimation in a significantly different manner. Namely, as long as $\sigma_{\mathbf{P}} \gg \sigma_\zeta$, smaller $\sigma_{\mathbf{P}}$ has a provably beneficial effect on estimation accuracy, whereas small non-zero eigenvalues of Σ (i.e. large eigenvalues of Σ^\dagger), corresponding to eigenvectors in $\text{Im}(\mathbf{Q})$, can only worsen the accuracy. A similar observation can be made when inspecting the asymptotic covariance of $\sqrt{N}(\hat{\mathbf{d}} - \mathbf{a})$, which after using $\Sigma \mathbf{a} = \sigma_{\mathbf{P}}^2 \mathbf{a}$ and [5, Theorem 1], is




FIG 2. (X, Y) sampled according to $X \sim \text{Uni}(\{X : \|X\|_2 \leq 1\})$ and $Y = f(\mathbf{a}^\top X) + \zeta$, where $\text{span}\{\mathbf{a}\}$ is the horizontal line, and $\zeta \sim \mathcal{N}(0, 0.01 \text{Var}(f(\mathbf{a}^\top X)))$. We use a dyadic level-set partitioning of the data into J sub-intervals. The color indicates the labeling.

given by

$$\underbrace{\frac{\sigma_{\mathbf{P}}^2}{\text{Cov}(f(\mathbf{a}^\top X), \mathbf{a}^\top X)}}_{=:T_1} \underbrace{\mathbf{Q}\Sigma^\dagger\mathbf{Q} \text{Cov}\left((\tilde{Y} - \tilde{X}^\top \mathbf{b})\tilde{X}\right)}_{=:T_2} \mathbf{Q}\Sigma^\dagger\mathbf{Q}.$$

The scalar factor T_1 is comparable to $\|\mathbf{b}\|_2^2$, and the matrix T_2 shows that the variance of the estimator grows with large eigenvalues of $\mathbf{Q}\Sigma^\dagger\mathbf{Q}$. On the other hand, the variance of the residual $\tilde{Y} - \tilde{X}^\top \mathbf{b}$ decreases if the accuracy of the linear fit of $\tilde{X}^\top \mathbf{b} \approx f(\mathbf{a}^\top X) - \mathbb{E}f(\mathbf{a}^\top X)$ improves, which can typically be observed under monotonic links f and decreasing $\sigma_{\mathbf{P}}^2 = \text{Var}(\mathbf{a}^\top X)$.

We will now use this observation as a guiding principle for developing a modified estimator that splits the data into subsets with small $\sigma_{\mathbf{P}}$'s, computes corresponding OLS vectors on each subset, and then combines local estimators into a global estimator by weighted averaging.

4.2. Averaged conditional least squares for the single-index model

We now study the just discussed alternative procedure, where we first split the data into subsets, aiming to reduce the variance of the data distribution in the direction of the index vector, and then compute and average out the estimators from each subset. Since we have no a priori knowledge about the index vector, constructing such a partition seems challenging. However, when the link function is monotonic the partitioning is induced by a decomposition of $\text{Im}(Y)$ into equisized intervals, see Figure 2. For the sake of simplicity, in the following we assume $Y \in [0, 1]$ holds almost surely.

For $\ell \in [J]$ let $\mathcal{R}_{J,\ell} := [\frac{\ell-1}{J}, \frac{\ell}{J})$ denote equisized regions partitioning $[0, 1)$. Furthermore, define herein called level-sets $\mathcal{S}_{J,\ell} = \{(X_i, Y_i) : Y_i \in \mathcal{R}_{J,\ell}\}$, which induce a partition of the data set into J subsets based on the responses. Then estimate \mathbf{a} according to the following algorithm.

- Step 1 Solve (4.2) on each subset $\mathcal{S}_{J,\ell}$, by computing $\hat{c}_{J,\ell} \in \mathbb{R}$ and $\hat{\mathbf{b}}_{J,\ell} \in \mathbb{R}^D$.
- Step 2 Define the empirical density $\hat{\rho}_{J,\ell} := |\mathcal{S}_{J,\ell}|/N$, set the thresholding parameter $\alpha > 0$, and compute the averaged outer product matrix

$$\hat{\mathbf{M}}_J = \sum_{\ell=1}^J \mathbb{1}_{[\alpha J^{-1}, 1]}(\hat{\rho}_{J,\ell}) \hat{\rho}_{J,\ell} \hat{\mathbf{b}}_{J,\ell} \hat{\mathbf{b}}_{J,\ell}^\top. \tag{4.8}$$

Step 3 Use the eigenvector corresponding to the largest eigenvalue of \mathbf{M}_J , denoted as $\mathbf{u}_1(\hat{\mathbf{M}}_J)$, as an approximation of the index vector \mathbf{a} .

The parameter α is used to promote numerical stability by suppressing the contributions of sparsely populated subsets. In other words, we only keep those level-sets whose empirical mass behaves as if Y were uniformly distributed over $\text{Im}(Y)$ (which can in some problems be achieved by a suitable transformation of the responses). The parameter J on the other hand defines the number of sets we use in the partition of the given data set, and dictates the trade-off between $\text{Var}(\mathbf{a}^\top X | Y \in \mathcal{R}_{J,\ell})$ and the number of samples $|\mathcal{S}_{J,\ell}|$ in a given level set.

For a random vector Z we denote a conditional random vector $Z_{J,\ell} := Z | Y \in \mathcal{R}_{J,\ell}$. Note that $Z_{J,\ell}$ inherits sub-Gaussianity of Z provided $\mathbb{P}(Y \in \mathcal{R}_{J,\ell}) > 0$, see Lemma 23. We now analyze the approach under the following assumption:

(A) $\mathbf{b}_{J,\ell} := \text{Cov}(X | Y \in \mathcal{R}_{J,\ell})^\dagger \text{Cov}(X, Y | Y \in \mathcal{R}_{J,\ell}) \in \text{span}\{\mathbf{a}\}$ for all $\ell \in [J]$.

Assumption (A) is not particularly restrictive. For example, it can be shown that (A) holds if X is elliptically symmetric, which is a standard assumption when using the OLS functional (4.2), and if the function noise $Y - \mathbb{E}[Y|X]$ is statistically independent of X given $\mathbf{a}^\top X$ [26, Proposition 3].

Theorem 17. *Assume (A) holds and that $Y \in [0, 1]$ almost surely. Let $J > 0$, $\alpha > 0$, and assume we are given N iid. copies of (X, Y) . Denote $\mathbf{P} := \mathbf{a}\mathbf{a}^\top$, $\mathbf{Q} := \mathbf{Id}_D - \mathbf{P}$, $\Sigma_{J,\ell} := \text{Cov}(X | Y \in \mathcal{R}_{J,\ell})$, and*

$$\kappa_{J,\ell} := \kappa(\mathbf{P}, X_{J,\ell}) \vee \kappa(\mathbf{Q}, X_{J,\ell}), \text{ and } K_{J,\ell} := \|\sqrt{\Sigma_{J,\ell}^\dagger} \tilde{X}_{J,\ell}\|_{\psi_2}.$$

Let $I_J := \{\ell \in [J] : \hat{\rho}_{J,\ell} > \alpha J^{-1}\}$ be the index set containing the indices of active level-sets and set $K_J := \max_{\ell \in I_J} K_{J,\ell}$. There exists $C > 0$ such that, if

$$N > CK_J^4 \frac{J(\text{rank}(\Sigma) + \log(J) + u)}{\alpha}, \quad (4.9)$$

there exists a sign $s \in \{-1, 1\}$ so that for any $u > 0$, with probability at least $1 - \exp(-u)$ we have

$$\|s\mathbf{u}_1(\hat{\mathbf{M}}_J) - \mathbf{a}\|_2^2 \leq \varepsilon_{N,J,u} \frac{\sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \kappa_{J,\ell} \|\mathbf{Q}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2}{\sum_{\ell \in I_J} \hat{\rho}_{J,\ell} (\|b_{J,\ell}\|_2^2 - \varepsilon_{N,J,u} \kappa_{J,\ell} \|\mathbf{P}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2)}, \quad (4.10)$$

where

$$\varepsilon_{N,J,u} := C \frac{\text{rank}(\Sigma) + \log(J) + u}{\alpha J N}.$$

The same guarantee holds when replacing the denominator in (4.10) with $\lambda_1(\hat{\mathbf{M}}_J)$.

The error rate in Theorem 17 is dictated by $\varepsilon_{N,J,u}$ and there are two ways to interpret Theorem 17. First, for a fixed parameter J all terms in (4.9) and (4.10) (except N and $\varepsilon_{N,J,u}$) are constants, and we obtain a $N^{-1/2}$ convergence rate

for the non-squared error as soon as the sample size is sufficiently large. This is the same rate as the one achieved with the standard OLS estimator (4.2).

Parameter J however provides additional flexibility and an intriguing option is to select J as a function that grows with N while still complying with (4.9). Namely, using $u = \log(J)$ and $J \asymp N/\log(N)$, and plugging into (4.10), implies that a $\sqrt{\log(N/\log(N))\log(N)}N^{-1} \leq \log(N)N^{-1}$ rate is achievable, so long as the remaining terms involving parameter J are balanced.

To give an illustrative example of such a case, we consider the following idealized setting. We assume the noise-free regime $Y = f(\mathbf{a}^\top X)$ with strictly monotonic link in the sense that for some $L > 0$ we have

$$\text{Cov}(\mathbf{a}^\top X, f(\mathbf{a}^\top X)|Y \in \mathcal{R}_{J,\ell}) > L\text{Var}(\mathbf{a}^\top X|Y \in \mathcal{R}_{J,\ell}). \quad (4.11)$$

Furthermore, conditional random variables $X_{J,\ell}$ are assumed strictly sub-Gaussian, where the strict sub-Gaussianity constant $K > 0$ in (2.5) is independent of J and ℓ , and we model local covariance matrices as

$$\text{Cov}(\tilde{X}_{J,\ell}) = C_\ell J^{-2} \mathbf{P} + \mathbf{Q} \quad \text{with } c_1 \leq C_\ell \leq c_2, \quad (4.12)$$

where c_1, c_2 are universal constants independent of J . Condition (4.12) implies that partitioning the data into J level-sets results in a reduction of the variance along the direction of the index vector, while not affecting the variance in directions orthogonal to the index vector.

With strict sub-Gaussianity and (4.12), we have $K_J \leq \sqrt{K}$, $\kappa_{J,\ell} \leq (c_2/c_1)K^2$, $\|\mathbf{P}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \leq KJ^2/c_1$, and $\|\mathbf{Q}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \leq K$. Then, using $u = \log(J)$ and $J = \tau N/\log(N)$ for some $\tau > 0$, the result (4.10) can be written as

$$\|\mathbf{su}_1(\hat{\mathbf{M}}_J) - \mathbf{a}\|_2^2 \leq \frac{C_K \text{rank}(\Sigma)}{\alpha\tau \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \left(\|b_{J,\ell}\|_2^2 - \tau \frac{C_K(\text{rank}(\Sigma)+1)}{\alpha} \right)} \frac{\log^2(N)}{N^2}, \quad (4.13)$$

where $C_K > 0$ depends only on K , c_1 , and c_2 . Furthermore, using Assumption (A), and (4.12), we have $\text{Cov}(X, Y|Y \in \mathcal{R}_{J,\ell}) = \text{Cov}(a^\top X, Y|Y \in \mathcal{R}_{J,\ell}) a$, and with the asserted monotonicity (4.11) we get

$$\begin{aligned} \|b_{J,\ell}\|_2 &\geq \frac{\text{Cov}(X, Y|Y \in \mathcal{R}_{J,\ell})^\top \Sigma_{J,\ell}^\dagger \text{Cov}(X, Y|Y \in \mathcal{R}_{J,\ell})}{\|\text{Cov}(X, Y|Y \in \mathcal{R}_{J,\ell})\|_2} \\ &> L\text{Var}(a^\top X|Y \in \mathcal{R}_{J,\ell}) a^\top \Sigma_{J,\ell}^\dagger a = L. \end{aligned} \quad (4.14)$$

Plugging this bound into (4.13), and choosing $\tau \in (0, \alpha L^2/(2C_K(\text{rank}(\Sigma)+1)))$, we thus obtain a $\log^2(N)/N^2$ convergence rate for the squared error.

The described idealized setting can be a reasonable approximation of reality for strictly monotonic link functions f in the noise-free regime, since (4.12) can be motivated by the observation that in this case slicing the distribution of X based on Y reduces the variance of the data along the direction of the index vector, as shown in Figure 2. Thus, monotonicity is a key requirement for establishing a faster rate. In the case of noisy responses and strictly monotone links,

the matter is more delicate. Namely, by the minimax rate of linear regression, which equals $\text{Var}(Y - f(\mathbf{a}^\top X))/N$, result (4.13) can only hold for all $N \in \mathbb{N}$, if $\text{Var}(Y - f(\mathbf{a}^\top X)) \in \mathcal{O}(\log^2(N)/N)$ as $N \rightarrow \infty$.

The same observation carries over to more general strictly monotonic links, because, in the presence of noise, condition (4.12) and the lower bound (4.14) do not hold for all $J \in \mathbb{N}$. For instance, if J^2 exceeds $1/\text{Var}(Y - f(\mathbf{a}^\top X))$ or becomes even larger, the conditional distribution $Y|Y \in \mathcal{R}_{J,\ell}$ has roughly as much correlation with label noise as it does with the labels, $f(\mathbf{a}^\top X)$. This suggests $\|b_{J,\ell}\|_2 \approx 0$, and that further slicing the distribution of X based on Y will not decrease the variance in index vector direction, as required per (4.12). For $J^{-2} \gg \text{Var}(Y - f(\mathbf{a}^\top X))$ however, Condition (4.12) and (4.14) can still hold approximately, and our experiments below suggest that the proposed estimator benefits from an aggressive splitting of the data in the case of nonlinear, strictly monotonic links. We now confer the implications of Theorem 17, with synthetic examples, under a data-driven parameter choice rule for J .

Numerical setup and parameter choice We sample $X \sim \mathcal{N}(\mathbf{0}, \mathbf{Id}_D)$, with $D = 10$, and let $\mathbf{a} = (1, 0, \dots, 0)^\top$ (the specific choice of \mathbf{a} is irrelevant for the results due to the rotational invariance of X). Responses are generated by $Y = f(\mathbf{a}^\top X) + \zeta$, where $\zeta \sim \mathcal{N}(0, \sigma_\zeta^2 \text{Var}(f(\mathbf{a}^\top X)))$. We set $\alpha = 0.05$ and additionally exclude subsets with $|\mathcal{S}_{J,\ell}| < 2D$, for the sake of numerical stability. Our goal is twofold. First, to compare estimation of the index vector using the standard OLS approach (4.2) with the strategy proposed in this section, and second, to empirically confer the observations described after Theorem 17.

A critical step of the modified approach is the selection of J as a function of N . As mentioned in Theorem 17, the denominator effectively lower bounds $\lambda_1(\hat{\mathbf{M}}_J)$, and the corresponding concentration bound can be written as

$$\|\mathbf{s}\mathbf{u}_1(\hat{\mathbf{M}}_J) - \mathbf{a}\|_2^2 \lesssim \frac{\text{rank}(\mathbf{\Sigma}) + \log(J) + u \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \kappa_{J,\ell} \|\mathbf{Q}\mathbf{\Sigma}_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2}{\alpha N J \lambda_1(\hat{\mathbf{M}}_J)}.$$

Neglecting the dependence of $\kappa_{J,\ell}$ and $\|\mathbf{Q}\mathbf{\Sigma}_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2$ on J and $\log(J)$ -terms, originating from union bounds in the proofs, the error is minimized when maximizing $J\lambda_1(\hat{\mathbf{M}}_J)$. We thus propose to adaptively choose J by

$$J^* := \max\{J = \lceil (1.5)^k \rceil : k \in \mathbb{N}_0, J\lambda_1(\hat{\mathbf{M}}_J) > J'\lambda_1(\hat{\mathbf{M}}_{J'}) \forall J' < J\}, \quad (4.15)$$

where we use an exponential grid in \mathbb{N} to decrease the computational demand. Note also that $\lambda_1(\hat{\mathbf{M}}_J) \geq \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \|\hat{b}_{J,\ell}\|_2^2 \approx \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \|b_{J,\ell}\|_2^2$, so, using parameter choice (4.15), we expect no further increase of J if $\|b_{J,\ell}\|_2$'s decay rapidly.

Numerical experiments Figure 3 shows the errors $\|\hat{\mathbf{d}} - \mathbf{a}\|_2$ and the corresponding optimized choices J^* of the number of level-sets J , for different link functions f . Solid lines correspond to the estimator (4.2), the dashed lines to $\mathbf{u}_1(\hat{\mathbf{M}}_{J^*})$, and different colors represent different noise levels σ_ζ . In Figures 3(b)–3(d) we consider monotonic link functions and we can see that the

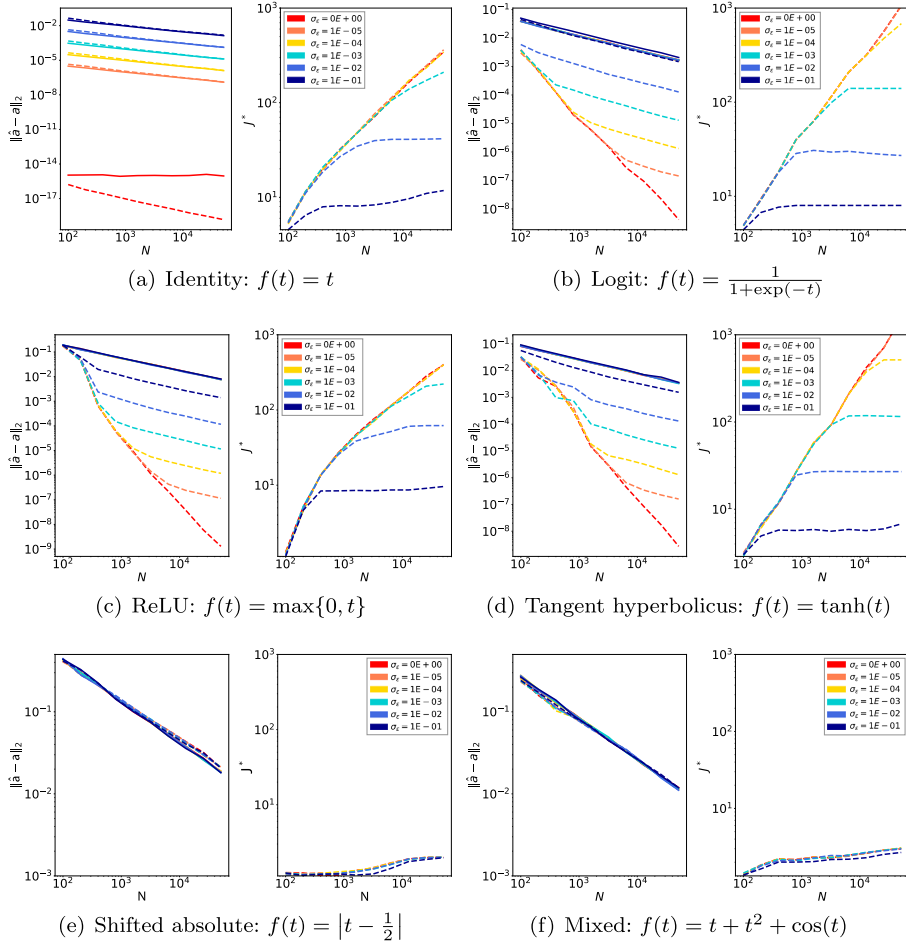


FIG 3. We plot the error $\|\hat{\mathbf{a}} - \mathbf{a}\|_2$ using (4.2) (solid lines), respectively $\hat{\mathbf{a}} = \mathbf{u}_1(\hat{\mathbf{M}}_{J^*})$ (dashed lines), for several link functions. The right plots in each subplot shows J^* that is chosen according to the rule (4.15). We see that in all cases where f is a nonlinear, monotonic function $\mathbf{u}_1(\hat{\mathbf{M}}_{J^*})$ improves upon (4.2), especially in scenarios with moderate noise levels σ_ζ . In the other cases the two estimators achieve similar accuracy.

approach presented in this section performs substantially better than the standard OLS approach. On the other hand, in Figures 3(a), 3(e) and 3(f) the two approaches achieve similar performance. In case of 3(a) this is because (4.2) is indeed optimal, according to the Gauss-Markov Theorem, whereas link functions in 3(e), 3(f), are not monotonic. In the latter case, the plots of J^* confirm that $\lambda_1(\hat{\mathbf{M}}_J)$ decays rapidly as a function of J , leaving J^* essentially constant as a function of N .

Let us examine the results for monotonic functions in more detail. The plots for J^* in Figures 3(b)–3(d) show that the number of level-sets J^* indeed grows

as a function of N , and it does so up to a level dictated by the noise level σ_ζ . This shows that $\lambda_1(\hat{\mathbf{M}}_J)$ does not decay faster than J^{-1} over a reasonably large range of N values, i.e. until $J^{-1} \approx \sigma_\zeta$. As a consequence, our approach achieves a faster, N^{-1} estimation rate for the index vector. Specifically, in the noise-free case this holds asymptotically as $N \rightarrow \infty$. On the other hand, in the case of corrupted Y 's, we first have a faster N^{-1} convergence, and then a sharp transition into the usual $N^{-1/2}$ rate. The number of points N at which this transition occurs depends inversely on the level of noise.

Acknowledgements

We like to thank the anonymous referees for their feedback that helped to greatly improve earlier versions of the manuscript. This work is supported by the grant *Function-driven Data Learning in High Dimensions* by the Research Council of Norway.

Appendix

A.1. Additional technical results

Lemma 18 (Properties of the sub-Gaussian norm). *Let X, Z be sub-Gaussian random vectors in \mathbb{R}^D , let $\Sigma = \text{Cov}(X)$, and $\mathbf{A} \in \mathbb{R}^{D \times D}$. Then we have*

- (1) $\|X - \mathbb{E}X\|_{\psi_2} \lesssim \|X\|_{\psi_2}$ (holds also for sub-exponential random variables),
- (2) $\|\mathbf{A}\tilde{X}\|_{\psi_2} \leq \|\mathbf{A}\|_2 \|\mathbf{P}_{\text{Im}(\mathbf{A}^\top)}\tilde{X}\|_{\psi_2}$,
- (3) $\|\mathbf{A}\Sigma\mathbf{A}^\top\|_2 \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2}^2$,
- (4) $X^\top Z$ is sub-exponential with $\|X^\top Z - \mathbb{E}[X^\top Z]\|_{\psi_1} \lesssim \|X\|_{\psi_2} \|\tilde{Z}\|_{\psi_2}$.

Proof. Property (1) is shown in [54, Lemma 2.6.8] for sub-Gaussian random variables. Applying the definition of the sub-Gaussian norm (1.9) the claim follows, since $\mathbf{v}^\top X$ is a sub-Gaussian random variable for every $\mathbf{v} \in \mathbb{S}^{D-1}$. The same line of arguments holds for sub-exponential vectors.

For (2) we compute for arbitrary $\mathbf{v} \in \mathbb{S}^{D-1}$ and $\mathbf{A}^\top \mathbf{v} \in \text{Im}(\mathbf{A}^\top)$

$$\begin{aligned} \|\mathbf{v}^\top \mathbf{A}\tilde{X}\|_{\psi_2} &= \|\mathbf{A}^\top \mathbf{v}\|_2 \left(\frac{\mathbf{A}^\top \mathbf{v}}{\|\mathbf{A}^\top \mathbf{v}\|_2} \right)^\top \tilde{X} \|_{\psi_2} \\ &\leq \|\mathbf{A}^\top\|_2 \sup_{\mathbf{u} \in \text{Im}(\mathbf{A}^\top) \cap \mathbb{S}^{D-1}} \|\mathbf{u}^\top \tilde{X}\|_{\psi_2} \leq \|\mathbf{A}\|_2 \|\mathbf{P}_{\text{Im}(\mathbf{A}^\top)}\tilde{X}\|_{\psi_2}. \end{aligned}$$

For (3) we first note that [54, Proposition 2.5.2] implies $\text{Var}(u) \lesssim \|\tilde{u}\|_{\psi_2}^2$ for any sub-Gaussian u , where $\tilde{u} = u - \mathbb{E}[u]$. Thus, for every $\mathbf{v} \in \mathbb{S}^{D-1}$ we have

$$\mathbf{v}^\top \text{Cov}(\mathbf{A}X)\mathbf{v} = \text{Var}(\mathbf{v}^\top \mathbf{A}X) \lesssim \|\mathbf{v}^\top \mathbf{A}\tilde{X}\|_{\psi_2}^2.$$

Taking the supremum with respect to $\mathbf{v} \in \mathbb{S}^{D-1}$, the result follows since $\text{Cov}(\mathbf{A}X)$ is positive semidefinite.

Property (4) follows from the centering property (1) and [54, Lemma 2.7.6]. \square

Lemma 19. Let $\mathbf{A} \in \mathbb{R}^{D \times D}$ be positive semidefinite, $\mathbf{B}_1 \in \mathbb{R}^{d_1 \times D}$, $\mathbf{B}_2 \in \mathbb{R}^{d_2 \times D}$. For $\mathbf{u} \in \mathbb{R}^{d_1}$, $\mathbf{v} \in \mathbb{R}^{d_2}$ we have $\mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_2^\top \mathbf{v} \leq \sqrt{\mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top \mathbf{u}} \sqrt{\mathbf{v}^\top \mathbf{B}_2 \mathbf{A} \mathbf{B}_2^\top \mathbf{v}}$. Moreover, $\|\mathbf{B}_1 \mathbf{A} \mathbf{B}_2^\top\|_2^2 \leq \|\mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top\|_2 \|\mathbf{B}_2 \mathbf{A} \mathbf{B}_2^\top\|_2$.

Proof. Applying the Cauchy-Schwartz inequality we have

$$\mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_2^\top \mathbf{v} = \langle \mathbf{A}^{1/2} \mathbf{B}_1^\top \mathbf{u}, \mathbf{A}^{1/2} \mathbf{B}_2^\top \mathbf{v} \rangle \leq \|\mathbf{A}^{1/2} \mathbf{B}_1^\top \mathbf{u}\|_2 \|\mathbf{A}^{1/2} \mathbf{B}_2^\top \mathbf{v}\|_2. \quad (\text{A.1})$$

By the same line of argument we have $\|\mathbf{A}^{1/2} \mathbf{B}_1^\top \mathbf{u}\|_2^2 = \mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top \mathbf{u}$ and $\|\mathbf{A}^{1/2} \mathbf{B}_2^\top \mathbf{v}\|_2^2 = \mathbf{v}^\top \mathbf{B}_2 \mathbf{A} \mathbf{B}_2^\top \mathbf{v}$, giving the first statement. Considering now $\|\mathbf{u}\|_2 = \|\mathbf{v}\|_2 = 1$, we have

$$\sup_{\substack{\|\mathbf{u}\|_2=1, \\ \|\mathbf{v}\|_2=1}} \mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_2^\top \mathbf{v} \leq \sqrt{\sup_{\|\mathbf{u}\|_2=1} \mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top \mathbf{u}} \sqrt{\sup_{\|\mathbf{v}\|_2=1} \mathbf{v}^\top \mathbf{B}_2 \mathbf{A} \mathbf{B}_2^\top \mathbf{v}}. \quad (\text{A.2})$$

Notice that since \mathbf{A} is positive semidefinite, then $\mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top$ and $\mathbf{B}_2 \mathbf{A} \mathbf{B}_2^\top$ are positive semidefinite. Therefore,

$$\sup_{\|\mathbf{u}\|_2=1} \mathbf{u}^\top \mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top \mathbf{u} = \sup_{\|\mathbf{u}\|_2=1} \|(\mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top)^{1/2} \mathbf{u}\|_2^2 = \|\mathbf{B}_1 \mathbf{A} \mathbf{B}_1^\top\|_2.$$

An analogous expression holds for the other term. Identifying the quadratic form on the left hand side in (A.2) as the operator norm of $\mathbf{B}_1 \mathbf{A} \mathbf{B}_2^\top$, the conclusion follows. \square

The following is a standard concentration bound for sub-Gaussian and sub-exponential random vectors around their mean.

Lemma 20. Let $\{X_i : i \in [N]\}$ be independent copies of a centered random vector $X \in \mathbb{R}^D$. Denote $\hat{\boldsymbol{\mu}} := N^{-1} \sum_{i=1}^N X_i$ and $\mathbf{m}(t) = t \vee t^2$. For any $u > 0$ the following hold with probability at least $1 - \exp(-u)$.

- (1) If $\|X\|_{\psi_2} < \infty$ we have $\|\hat{\boldsymbol{\mu}}\|_2 \lesssim \|X\|_{\psi_2} \sqrt{\frac{D+u}{N}}$.
- (2) If $\|X\|_{\psi_1} < \infty$ we have $\|\hat{\boldsymbol{\mu}}\|_2 \lesssim \|X\|_{\psi_1} \mathbf{m}\left(\sqrt{\frac{D+u}{N}}\right)$.

Proof. The argument for the two bounds follows along analogous lines. Let $\delta < 1/4$, and \mathcal{N} be a δ -net of \mathbb{S}^{D-1} . We first use [54, Exercise 4.4.3] to rewrite

$$\|\hat{\boldsymbol{\mu}}\|_2 = \sup_{\mathbf{v} \in \mathbb{S}^{D-1}} \mathbf{v}^\top \hat{\boldsymbol{\mu}} = \sup_{\mathbf{v} \in \mathbb{S}^{D-1}} N^{-1} \sum_{i=1}^N \mathbf{v}^\top X_i \leq 2 \sup_{\mathbf{v} \in \mathcal{N}} N^{-1} \sum_{i=1}^N \mathbf{v}^\top X_i.$$

The term $N^{-1} \sum_{i=1}^N \mathbf{v}^\top X_i$ is a sum of either sub-Gaussian or sub-exponential random variables, for any $\mathbf{v} \in \mathcal{N}$. In the former case we have $\|\mathbf{v}^\top X_i\|_{\psi_1} \leq \|X\|_{\psi_1}$, and $\|\mathbf{v}^\top X_i\|_{\psi_2} \leq \|X\|_{\psi_2}$ in the latter. Hoeffding's inequality [54, Theorem 2.6.2], in the sub-Gaussian case, or Bernstein's inequality [54, Theorem 2.8.1], in the sub-exponential case, now yield concentration bounds for the sums. The claim follows by applying the union bound over $\mathbf{v} \in \mathcal{N}$, where the number of events is bounded by $|\mathcal{N}| \leq 12^D$, see for instance [54, Corollary 4.2.13]. \square

A.2. Proofs for Section 2

Proof of Lemma 2. Let $\tilde{X} = X - \mathbb{E}X$. Let $\mathbf{U}_A \in \mathbb{R}^{d_A \times D}$ and $\mathbf{U}_B \in \mathbb{R}^{d_B \times D}$ be matrices whose rows contain the orthonormal basis for $\text{Im}(\mathbf{A}\Sigma)$ and $\text{Im}(\mathbf{B}\Sigma)$, respectively. Since Σ and $\hat{\Sigma}$ are symmetric, and $\text{Im}(\hat{\Sigma}) \subseteq \text{Im}(\Sigma)$, we have $\hat{\Sigma} - \Sigma = \mathbf{P}_\Sigma(\hat{\Sigma} - \Sigma)\mathbf{P}_\Sigma$, where \mathbf{P}_Σ is the orthogonal projection onto $\text{Im}(\Sigma)$. We thus have

$$\|\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top\|_2 = \|\mathbf{A}_\Sigma(\hat{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_2,$$

for $\mathbf{A}_\Sigma = \mathbf{U}_A\mathbf{A}\mathbf{P}_\Sigma \in \mathbb{R}^{d_A \times D}$ and $\mathbf{B}_\Sigma = \mathbf{U}_B\mathbf{B}\mathbf{P}_\Sigma \in \mathbb{R}^{d_B \times D}$. Denote now $\tilde{\Sigma} = N^{-1} \sum_{i=1}^N \tilde{X}_i \tilde{X}_i^\top$. Notice that $\tilde{\Sigma}$, compared to the empirical covariance $\hat{\Sigma}$, uses the true, instead of the empirical mean of \tilde{X} . We then have

$$\|\mathbf{A}_\Sigma(\hat{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_2 \leq \|\mathbf{A}_\Sigma(\tilde{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_2 + \left\| \sum_{i=1}^N \frac{\mathbf{A}_\Sigma \tilde{X}_i}{N} \right\|_2 \left\| \sum_{i=1}^N \frac{\mathbf{B}_\Sigma \tilde{X}_i}{N} \right\|_2.$$

By Lemma 20 the second term is always of higher order, and the resulting error can be absorbed into a corresponding upper bound for the first term. Thus, in the following we focus on $\|\mathbf{A}_\Sigma(\tilde{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_2$.

We closely follow the proof of [52, Proposition 2.1]. Let $\delta < 1/4$, and by \mathcal{N} , \mathcal{M} denote δ -nets of \mathbb{S}^{d_A-1} and \mathbb{S}^{d_B-1} . From [54, Exercise 4.4.3] we have

$$\|\mathbf{A}_\Sigma(\tilde{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_2 \leq (1 - 2\delta)^{-1} \sup_{\substack{\mathbf{x} \in \mathcal{N} \\ \mathbf{y} \in \mathcal{M}}} \langle \mathbf{A}_\Sigma(\tilde{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top \mathbf{x}, \mathbf{y} \rangle \quad (\text{A.3})$$

$$\leq 2 \sup_{\substack{\mathbf{x} \in \mathcal{N} \\ \mathbf{y} \in \mathcal{M}}} \langle (\tilde{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top \mathbf{x}, \mathbf{A}_\Sigma^\top \mathbf{y} \rangle. \quad (\text{A.4})$$

Consider now any pair $(\mathbf{x}, \mathbf{y}) \in \mathcal{N} \times \mathcal{M}$, and write

$$\langle \tilde{\Sigma}\mathbf{B}_\Sigma^\top \mathbf{x}, \mathbf{A}_\Sigma^\top \mathbf{y} \rangle = \sum_{i=1}^N \frac{\langle \tilde{X}_i (\tilde{X}_i^\top \mathbf{B}_\Sigma^\top \mathbf{x}), \mathbf{A}_\Sigma^\top \mathbf{y} \rangle}{N} = \sum_{i=1}^N \frac{\langle \mathbf{A}_\Sigma \tilde{X}_i, \mathbf{y} \rangle \langle \mathbf{B}_\Sigma \tilde{X}_i, \mathbf{x} \rangle}{N}.$$

Since $\langle \mathbf{A}_\Sigma \tilde{X}_i, \mathbf{y} \rangle$ and $\langle \mathbf{B}_\Sigma \tilde{X}_i, \mathbf{x} \rangle$ are sub-Gaussian, their product is sub-exponential, and from [54, Lemma 2.7.7] we have

$$\|\langle \mathbf{A}_\Sigma \tilde{X}_i, \mathbf{y} \rangle \langle \mathbf{B}_\Sigma \tilde{X}_i, \mathbf{x} \rangle\|_{\psi_1} \leq \|\langle \mathbf{A}_\Sigma \tilde{X}_i, \mathbf{y} \rangle\|_{\psi_2} \|\langle \mathbf{B}_\Sigma \tilde{X}_i, \mathbf{x} \rangle\|_{\psi_2} \leq \|\mathbf{A}_\Sigma \tilde{X}\|_{\psi_2} \|\mathbf{B}_\Sigma \tilde{X}\|_{\psi_2},$$

which we denote by $\sigma_{\mathbf{AB}} := \|\mathbf{A}_\Sigma \tilde{X}\|_{\psi_2} \|\mathbf{B}_\Sigma \tilde{X}\|_{\psi_2}$ for short. Since $\mathbb{E}[\tilde{\Sigma}] = \Sigma$, by Lemma 20 we have for any $u > 0$, with $\mathbf{m}(t) = t \vee t^2$,

$$\mathbb{P} \left(\left| \langle \tilde{\Sigma}\mathbf{B}_\Sigma^\top \mathbf{x}, \mathbf{A}_\Sigma^\top \mathbf{y} \rangle - \langle \Sigma\mathbf{B}_\Sigma^\top \mathbf{x}, \mathbf{A}_\Sigma^\top \mathbf{y} \rangle \right| \geq \sigma_{\mathbf{AB}} \mathbf{m} \left(\sqrt{\frac{1+u}{N}} \right) \right) \leq \exp(-u).$$

The size of the nets can be bounded as $|\mathcal{N}| \leq 12^{d_A}$ and $|\mathcal{M}| \leq 12^{d_B}$, see [54, Corollary 4.2.13]. Thus, considering all pair $(\mathbf{x}, \mathbf{y}) \in \mathcal{N} \times \mathcal{M}$ and using the union

bound we get, for some universal constant C (bounded e.g. by $\log(12) + 1$),

$$\begin{aligned} \mathbb{P} \left(\left| \sup_{\mathbf{x} \in \mathcal{N}, \mathbf{y} \in \mathcal{M}} \langle \tilde{\Sigma} \mathbf{B}_{\Sigma}^{\top} \mathbf{x}, \mathbf{A}_{\Sigma}^{\top} \mathbf{y} \rangle - \langle \Sigma \mathbf{B}_{\Sigma}^{\top} \mathbf{x}, \mathbf{A}_{\Sigma}^{\top} \mathbf{y} \rangle \right| \geq C \sigma_{\mathbf{AB}} \mathfrak{m} \left(\sqrt{\frac{u + (d_A + d_B)}{N}} \right) \right) \\ \leq |\mathcal{N}| |\mathcal{M}| \exp(-u - (d_A + d_B) \log(12)) \\ \leq \exp((d_A + d_B) \log(12) - u - (d_A + d_B) \log(12)). \end{aligned}$$

Using the δ -net approximation bound (A.3) we thus get

$$\begin{aligned} \mathbb{P} \left(\|\mathbf{A}_{\Sigma}(\tilde{\Sigma} - \Sigma) \mathbf{B}_{\Sigma}^{\top}\|_2 \geq 2C \sigma_{\mathbf{AB}} \mathfrak{m} \left(\sqrt{\frac{u + (d_A + d_B)}{N}} \right) \right) \\ \leq \mathbb{P} \left(\left| \sup_{\substack{\mathbf{x} \in \mathcal{N} \\ \mathbf{y} \in \mathcal{M}}} \langle \tilde{\Sigma} \mathbf{B}_{\Sigma}^{\top} \mathbf{x}, \mathbf{A}_{\Sigma}^{\top} \mathbf{y} \rangle - \langle \Sigma \mathbf{B}_{\Sigma}^{\top} \mathbf{x}, \mathbf{A}_{\Sigma}^{\top} \mathbf{y} \rangle \right| \geq C \sigma_{\mathbf{AB}} \mathfrak{m} \left(\sqrt{\frac{u + (d_A + d_B)}{N}} \right) \right). \end{aligned}$$

The result now follows from $\|\mathbf{A}_{\Sigma} \tilde{X}\|_{\psi_2} = \|\mathbf{A} \tilde{X}\|_{\psi_2}$ and $\|\mathbf{B}_{\Sigma} \tilde{X}\|_{\psi_2} = \|\mathbf{B} \tilde{X}\|_{\psi_2}$ and thus $\sigma_{\mathbf{AB}} = \|\mathbf{A} \tilde{X}\|_{\psi_2} \|\mathbf{B} \tilde{X}\|_{\psi_2}$. \square

Proof of Lemma 7. We use the shorthand notation $\lambda_i := \lambda_i(\Sigma)$ and $\hat{\lambda}_i := \hat{\lambda}_i(\Sigma)$. By the definition of δ_{il} we have

$$\begin{aligned} \delta_{il} &= |\hat{\lambda}_{i-1} - \lambda_i| \wedge |\lambda_l - \hat{\lambda}_{l+1}| \geq (\hat{\lambda}_{i-1} - \lambda_i) \wedge (\lambda_l - \hat{\lambda}_{l+1}) \\ &\geq \left((\lambda_{i-1} - \lambda_i) + (\hat{\lambda}_{i-1} - \lambda_{i-1}) \right) \wedge \left((\lambda_l - \lambda_{l+1}) + (\lambda_{l+1} - \hat{\lambda}_{l+1}) \right), \end{aligned}$$

so it suffices to show $\hat{\lambda}_{i-1} - \lambda_{i-1} \geq -(\lambda_{i-1} - \lambda_i)/2$ and $\lambda_{l+1} - \hat{\lambda}_{l+1} \geq -(\lambda_l - \lambda_{l+1})/2$. Note that if $i = 1$ or $l = \text{rank}(\Sigma)$, the corresponding gap is trivial and no concentration is required. Hence, we can focus on the case $i > 1$ and $l < \text{rank}(\Sigma)$. Using [48, Proposition 3.13], we obtain that for any $u > 0$ a relative eigenvalue bound

$$\hat{\lambda}_{i-1} - \lambda_{i-1} < -\frac{\lambda_{i-1} - \lambda_i}{2}, \quad (\text{A.5})$$

holds with probability at most $\exp(-u)$, provided

$$N > C_K \left(\frac{\lambda_{i-1}}{\lambda_{i-1} - \lambda_i} \vee 1 \right) \left(\sum_{j=1}^{i-1} \frac{\lambda_j}{\lambda_j - \lambda_{i-1} + \frac{\lambda_{i-1} - \lambda_i}{2}} \vee u \frac{\lambda_{i-1}}{\lambda_{i-1} - \lambda_i} \right). \quad (\text{A.6})$$

Furthermore, since $\lambda_{i-1} - (\lambda_{i-1} - \lambda_i)/2 < \lambda_j$ for all $j \geq i-1$, the summands in the sum in (A.6) are increasing, and we thus get

$$\sum_{j=1}^{i-1} \frac{\lambda_j}{\lambda_j - \lambda_{i-1} + \frac{\lambda_{i-1} - \lambda_i}{2}} \leq (i-1) \frac{2\lambda_{i-1}}{\lambda_{i-1} - \lambda_i} \lesssim \text{rank}(\Sigma) \frac{\lambda_{i-1}}{\lambda_{i-1} - \lambda_i}.$$

Thus, (A.6) is implied by the requirement on N in the statement. Similarly, using [48, Proposition 3.10] we get that

$$\hat{\lambda}_{l+1} - \lambda_{l+1} > \frac{\lambda_l - \lambda_{l+1}}{2} \quad (\text{A.7})$$

holds with probability at most $\exp(-u)$, provided

$$N > C_K \left(\frac{\lambda_{l+1}}{\lambda_l - \lambda_{l+1}} \vee 1 \right) \left(\sum_{j=l+1}^D \frac{\lambda_j}{\lambda_{l+1} - \lambda_j + \frac{\lambda_l - \lambda_{l+1}}{2}} \vee u \frac{\lambda_{l+1}}{\lambda_l - \lambda_{l+1}} \right). \quad (\text{A.8})$$

As before, we have $\lambda_j - (\lambda_l - \lambda_{l+1})/2 < \lambda_{l+1}$ for all $j \geq l+1$, so that

$$\sum_{j=l+1}^D \frac{\lambda_j}{\lambda_{l+1} - \lambda_j + \frac{\lambda_l - \lambda_{l+1}}{2}} \leq (\text{rank}(\mathbf{\Sigma}) - l) \frac{2\lambda_{l+1}}{\lambda_l - \lambda_{l+1}} \lesssim \text{rank}(\mathbf{\Sigma}) \frac{\lambda_{l+1}}{\lambda_l - \lambda_{l+1}}.$$

It follows that (A.8) is implied by the requirement on N in the statement. Taking the complementary event of the union of events (A.5) and (A.7), which holds with probability $1 - 2\exp(-u)$, we get $\hat{\lambda}_{i-1} - \lambda_{i-1} \geq -(\lambda_{i-1} - \lambda_i)/2$ and $\lambda_{l+1} - \hat{\lambda}_{l+1} \geq -(\lambda_l - \lambda_{l+1})/2$ and thus the result is proven. \square

A.3. Proofs for Section 3

Lemma 21. *If $\|\sqrt{\mathbf{\Sigma}^\dagger}(\mathbf{\Sigma} - \hat{\mathbf{\Sigma}})\sqrt{\mathbf{\Sigma}^\dagger}\|_2 < 1$, then $\ker(\hat{\mathbf{\Sigma}}) \cap \text{Im}(\mathbf{\Sigma}) = \{0\}$.*

Proof. We prove the statement by contraction. Assume there exists a nonzero $\mathbf{v} \in \ker(\hat{\mathbf{\Sigma}}) \cap \text{Im}(\mathbf{\Sigma})$. Since $\mathbf{v} \in \text{Im}(\mathbf{\Sigma})$ we have $\mathbf{v} \in \text{Im}(\sqrt{\mathbf{\Sigma}^\dagger})$ implying the existence of a $\mathbf{u} \in \mathbb{S}^{D-1}$ with $\mathbf{v} = \sqrt{\mathbf{\Sigma}^\dagger}\mathbf{u}$ (possibly after scaling \mathbf{v}). Using the properties of \mathbf{u} and \mathbf{v} , we get a contradiction by

$$1 = \mathbf{v}^\top \mathbf{\Sigma} \mathbf{v} = \mathbf{v}^\top (\mathbf{\Sigma} - \hat{\mathbf{\Sigma}}) \mathbf{v} = \mathbf{u}^\top \sqrt{\mathbf{\Sigma}^\dagger} (\mathbf{\Sigma} - \hat{\mathbf{\Sigma}}) \sqrt{\mathbf{\Sigma}^\dagger} \mathbf{u} \leq \|\sqrt{\mathbf{\Sigma}^\dagger} (\mathbf{\Sigma} - \hat{\mathbf{\Sigma}}) \sqrt{\mathbf{\Sigma}^\dagger}\|_2 < 1.$$

\square

Proof of Theorem 10. Assume for now $\|\sqrt{\mathbf{\Sigma}^\dagger}(\mathbf{\Sigma} - \hat{\mathbf{\Sigma}})\sqrt{\mathbf{\Sigma}^\dagger}\|_2 < \frac{1}{2}$ which will be ensured later by concentration arguments. Conditioned on this event, Lemma 21 implies $\ker(\hat{\mathbf{\Sigma}}) \cap \text{Im}(\mathbf{\Sigma}) = \{0\}$. Since $\text{Im}(\hat{\mathbf{\Sigma}}) \subset \text{Im}(\mathbf{\Sigma})$ holds almost surely, this implies $\text{Im}(\mathbf{\Sigma}) = \text{Im}(\hat{\mathbf{\Sigma}})$ almost surely. Denote now $\mathbf{\Delta} := \hat{\mathbf{\Sigma}} - \mathbf{\Sigma}$. We want to use a von Neumann series argument to express $\hat{\mathbf{\Sigma}}^\dagger$ in terms of $\mathbf{\Sigma}^\dagger$ and the corresponding finite sample perturbation. Recall that for general symmetric matrices \mathbf{U} and \mathbf{V} with $\text{Im}(\mathbf{U}) = \text{Im}(\mathbf{V})$, we have $(\mathbf{U}\mathbf{V})^\dagger = \mathbf{V}^\dagger \mathbf{U}^\dagger$ [50, Corollary 1.2]. Letting \mathbf{P}_Σ be the orthogonal projection onto $\text{Im}(\mathbf{\Sigma})$, and $\mathbf{Q}_\Sigma = \text{Id}_D - \mathbf{P}_\Sigma$, we can thus write

$$\begin{aligned} \hat{\mathbf{\Sigma}}^\dagger &= (\mathbf{\Sigma} + \mathbf{\Delta})^\dagger = \left(\mathbf{\Sigma}^{1/2} (\mathbf{P}_\Sigma + \sqrt{\mathbf{\Sigma}^\dagger} \mathbf{\Delta} \sqrt{\mathbf{\Sigma}^\dagger}) \mathbf{\Sigma}^{1/2} \right)^\dagger \\ &= \sqrt{\mathbf{\Sigma}^\dagger} (\mathbf{P}_\Sigma + \sqrt{\mathbf{\Sigma}^\dagger} \mathbf{\Delta} \sqrt{\mathbf{\Sigma}^\dagger})^\dagger \sqrt{\mathbf{\Sigma}^\dagger} \end{aligned}$$

$$\begin{aligned}
&= \sqrt{\Sigma^\dagger}(\mathbf{I}_D - \mathbf{Q}_\Sigma + \sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^\dagger\sqrt{\Sigma^\dagger} \\
&= \sqrt{\Sigma^\dagger}(\mathbf{I}_D + \sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^{-1}\sqrt{\Sigma^\dagger},
\end{aligned}$$

where we used the additivity of the Moore-Penrose inverse for matrices with orthogonal non-trivial eigenspaces in the last equality. Using a von Neumann series expansion for the second term, we obtain

$$\begin{aligned}
\hat{\Sigma}^\dagger &= \sqrt{\Sigma^\dagger}(\mathbf{I}_D + \sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^\dagger\sqrt{\Sigma^\dagger} = \sqrt{\Sigma^\dagger} \sum_{k=0}^{\infty} (-\sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^k \sqrt{\Sigma^\dagger} \\
&= \Sigma^\dagger + \sqrt{\Sigma^\dagger} \sum_{k=1}^{\infty} (-\sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^k \sqrt{\Sigma^\dagger},
\end{aligned}$$

where we used $\sqrt{\Sigma^\dagger}\sqrt{\Sigma^\dagger} = \Sigma^\dagger$ in the last equality. Subtracting Σ^\dagger , and multiplying by \mathbf{A} from the right and \mathbf{B}^\top from the left, it follows that

$$\begin{aligned}
\mathbf{A}(\hat{\Sigma}^\dagger - \Sigma^\dagger)\mathbf{B}^\top &= \mathbf{A}\sqrt{\Sigma^\dagger} \sum_{k=1}^{\infty} (-\sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^k \sqrt{\Sigma^\dagger}\mathbf{B}^\top \\
&= -\mathbf{A}\Sigma^\dagger\Delta\Sigma^\dagger\mathbf{B}^\top + \mathbf{A}\Sigma^\dagger\Delta\sqrt{\Sigma^\dagger} \sum_{k=0}^{\infty} (-\sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger})^k \sqrt{\Sigma^\dagger}\Delta\Sigma^\dagger\mathbf{B}^\top. \quad (\text{A.9})
\end{aligned}$$

The rest of the proof requires to apply Lemma 2 multiple times to concentrate Δ when multiplied by different matrices from left and right. More specifically, since $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^2 \gtrsim \|\text{Cov}(\sqrt{\Sigma^\dagger}X)\|_2 \gtrsim 1$, we can assume the regime $N > \text{rank}(\Sigma) + u$ by the assumption on N in Theorem 10. Then, we have by Lemma 2 with probability $1 - 4\exp(-u)$

$$\begin{aligned}
\|\mathbf{A}\Sigma^\dagger\Delta\Sigma^\dagger\mathbf{B}^\top\|_2 &\lesssim \|\mathbf{A}\Sigma^\dagger\tilde{X}\|_{\psi_2} \|\mathbf{B}\Sigma^\dagger\tilde{X}\|_{\psi_2} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}, \\
\|\mathbf{A}\Sigma^\dagger\Delta\sqrt{\Sigma^\dagger}\|_2 &\lesssim \|\mathbf{A}\Sigma^\dagger\tilde{X}\|_{\psi_2} \|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}, \quad (\text{A.10}) \\
\|\sqrt{\Sigma^\dagger}\Delta\Sigma^\dagger\mathbf{B}^\top\|_2 &\lesssim \frac{\|\mathbf{B}\Sigma^\dagger\tilde{X}\|_{\psi_2}}{\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}}, \text{ and } \|\sqrt{\Sigma^\dagger}\Delta\sqrt{\Sigma^\dagger}\| \leq \frac{1}{2},
\end{aligned}$$

whenever $N > C(\text{rank}(\Sigma) + u)\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^4$ (we used properties of the sub-Gaussian norm in Lemma 18 to get $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^2 \geq C\|\text{Cov}(\sqrt{\Sigma^\dagger}X)\|_2 = C$ and thus $\|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^2 \lesssim \|\sqrt{\Sigma^\dagger}\tilde{X}\|_{\psi_2}^4$). Using identity (A.9), combining the above bounds and adjusting the uniform constant C , the desired result holds with probability $1 - \exp(-u)$. \square

Proof of Theorem 13. For the covariance estimation bound denote first $\tilde{\Sigma} = N^{-1} \sum_{i=1}^N \tilde{X}_i \tilde{X}_i^\top$, where $\tilde{X}_i = X_i - \mathbb{E}X$, and decompose the error, as in the

proof of Lemma 2 into

$$\begin{aligned} \|\mathbf{A}(\hat{\Sigma} - \Sigma)\mathbf{B}^\top\|_F &= \|\mathbf{A}_\Sigma(\hat{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_F \\ &\leq \|\mathbf{A}_\Sigma(\tilde{\Sigma} - \Sigma)\mathbf{B}_\Sigma^\top\|_F + \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{A}_\Sigma \tilde{X}_i \right\|_F \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{B}_\Sigma \tilde{X}_i \right\|_F. \end{aligned}$$

The second term is again of higher order, with high probability, and can thus be disregarded. For the first term, define random matrices $\xi_i = \mathbf{A}_\Sigma \tilde{X}_i \tilde{X}_i^\top \mathbf{B}_\Sigma^\top - \mathbf{A}_\Sigma \Sigma \mathbf{B}_\Sigma^\top$. First note $\mathbb{E}[\xi_i] = \mathbf{0}$, and

$$\begin{aligned} \|\xi_i\|_F &\leq \|\mathbf{A}_\Sigma \tilde{X}_i \tilde{X}_i^\top \mathbf{B}_\Sigma^\top\|_F + \|\mathbf{A}_\Sigma^\top \Sigma \mathbf{B}_\Sigma^\top\|_F \\ &= \|\mathbf{A}_\Sigma \tilde{X}_i\|_2 \|\mathbf{B}_\Sigma \tilde{X}_i\|_2 + \mathbb{E} \|\mathbf{A}_\Sigma \tilde{X} \tilde{X}^\top \mathbf{B}_\Sigma^\top\|_F \leq 2C_A C_B. \end{aligned}$$

Moreover, ξ_i take values in the Hilbert space $(\mathbb{R}^{d_A \times d_B}, \|\cdot\|_F)$. Thus, by [6, Section 2.4], for every $u > 0$ we have

$$\left\| \frac{1}{N} \sum_{i=1}^N \xi_i \right\|_F \leq \frac{2\sqrt{2}C_A C_B u}{\sqrt{N}},$$

with probability at least $1 - \exp(-u^2)$. The conclusion follows after adjusting the confidence level u to account for the higher order term and the union bound.

The first steps for establishing the bound for precision matrices are the same as in the proof of Theorem 10. Indeed, the only differences are in the following lines of inequalities. Instead of (A.10), we have

$$\begin{aligned} \|\mathbf{A}\Sigma^\dagger \Delta \Sigma^\dagger \mathbf{B}^\top\|_F &\lesssim C_A^\dagger C_B^\dagger \sqrt{\frac{1+u}{N}}, \text{ and } \|\sqrt{\Sigma^\dagger} \Delta \Sigma^\dagger \mathbf{B}^\top\|_F \lesssim C_B^\dagger \sqrt{\Theta} \sqrt{\frac{1+u}{N}}; \\ \|\mathbf{A}\Sigma^\dagger \Delta \sqrt{\Sigma^\dagger}\|_F &\lesssim \frac{C_A^\dagger}{\sqrt{\Theta}}, \text{ and } \|\sqrt{\Sigma^\dagger} \Delta \sqrt{\Sigma^\dagger}\|_F \leq \frac{1}{2}, \end{aligned}$$

using the just proven covariance bounds. The claim now follows by using this together with (A.9), as in the proof of Theorem 10. \square

A.4. Proofs for Section 4

We first need a bound for $\mathbf{r} = \text{Cov}(X, Y)$, and a concentration around the finite sample estimate $\hat{\mathbf{r}} = N^{-1} \sum_{i=1}^N (X_i - \hat{\mu}_X)(Y_i - \hat{\mu}_Y)$.

Lemma 22. *Let $\mathbf{A} \in \mathbb{R}^{k \times D}$. If $X \in \mathbb{R}^D, Y \in \mathbb{R}$ are sub-Gaussian, we have $\|\mathbf{A}\mathbf{r}\|_2 \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2}$. Furthermore, with probability at least $1 - \exp(-u)$ we have*

$$\|\mathbf{A}(\mathbf{r} - \hat{\mathbf{r}})\|_2 \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2} \left(\sqrt{\frac{k+u}{N}} \vee \frac{k+u}{N} \right).$$

Proof. By Lemma 18 we have $\text{Var}(\mathbf{v}^\top \mathbf{A}X) \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2}^2$, and $\text{Var}(Y) \lesssim \|\tilde{Y}\|_{\psi_2}^2$, which implies

$$\begin{aligned} \|\mathbf{A}\mathbf{r}\|_2 &= \|\text{Cov}(\mathbf{A}X, Y)\|_2 = \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \mathbf{v}^\top \text{Cov}(\mathbf{A}X, Y) = \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \text{Var}(\mathbf{v}^\top \mathbf{A}X, Y) \\ &\leq \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \sqrt{\text{Var}(\mathbf{v}^\top \mathbf{A}X)} \sqrt{\text{Var}(Y)} \lesssim \|\mathbf{A}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2}. \end{aligned}$$

Define the random variable $Z_i := \tilde{X}_i \tilde{Y}_i - \text{Cov}(X, Y)$ with $\mathbb{E}Z_i = 0$. We write

$$\mathbf{A}(\mathbf{r} - \hat{\mathbf{r}}) = \frac{1}{N} \sum_{i=1}^N \mathbf{A}Z_i - \left(\frac{1}{N} \sum_{i=1}^N \mathbf{A}\tilde{X}_i \right) \left(\frac{1}{N} \sum_{i=1}^N \tilde{Y}_i \right).$$

As in the proof of Lemma 2, by applying Lemma 20 it follows that the second term is of higher order. On the other hand, the first term is an empirical mean of a sub-exponential centered variable $\mathbf{A}Z_i$,

$$\begin{aligned} \|\mathbf{A}Z_i\|_{\psi_1} &= \|\mathbf{A}\tilde{X}_i \tilde{Y}_i - \text{Cov}(\mathbf{A}X, Y)\|_{\psi_1} \lesssim \|\mathbf{A}\tilde{X}_i \tilde{Y}_i\|_{\psi_1} \\ &\lesssim \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \|\mathbf{v}^\top \mathbf{A}\tilde{X}_i \tilde{Y}_i\|_{\psi_1} \lesssim \|\mathbf{A}\tilde{X}_i\|_{\psi_2} \|\tilde{Y}_i\|_{\psi_2}, \end{aligned}$$

where we use the centering property of the sub-exponential norm, and the bound for the sub-exponential norm by the product of sub-Gaussian norms (Lemma 18). Applying now Lemma 20, the statement follows. \square

Proof of Lemma 14. We first note that $\|\sqrt{\Sigma^\dagger}X\|_{\psi_2}^2 \gtrsim \|\text{Cov}(\sqrt{\Sigma^\dagger}X)\|_2 \geq 1$, so by choosing C in the statement large enough, we assume the regime $N > (\text{rank}(\Sigma) + u)$ in the following. We begin with a bound for $\|\mathbf{P}(\Sigma^\dagger \mathbf{r} - \hat{\Sigma}^\dagger \hat{\mathbf{r}})\|_2$.

$$\begin{aligned} \|\mathbf{P}(\Sigma^\dagger \mathbf{r} - \hat{\Sigma}^\dagger \hat{\mathbf{r}})\|_2 &\leq \|\mathbf{P}(\Sigma^\dagger - \hat{\Sigma}^\dagger)\mathbf{P}\mathbf{r}\|_2 + \|\mathbf{P}(\Sigma^\dagger - \hat{\Sigma}^\dagger)\mathbf{Q}\mathbf{r}\|_2 \\ &\quad + \|\mathbf{P}\hat{\Sigma}^\dagger \mathbf{P}(\hat{\mathbf{r}} - \mathbf{r})\|_2 + \|\mathbf{P}\hat{\Sigma}^\dagger \mathbf{Q}(\hat{\mathbf{r}} - \mathbf{r})\|_2 \\ &\leq \|\mathbf{P}(\Sigma^\dagger - \hat{\Sigma}^\dagger)\mathbf{P}\|_2 \|\mathbf{P}\mathbf{r}\|_2 + \|\mathbf{P}(\Sigma^\dagger - \hat{\Sigma}^\dagger)\mathbf{Q}\|_2 \|\mathbf{Q}\mathbf{r}\|_2 \\ &\quad + \|\mathbf{P}\hat{\Sigma}^\dagger \mathbf{P}\|_2 \|\mathbf{P}(\hat{\mathbf{r}} - \mathbf{r})\|_2 + \|\mathbf{P}\hat{\Sigma}^\dagger \mathbf{Q}\|_2 \|\mathbf{Q}(\hat{\mathbf{r}} - \mathbf{r})\|_2. \end{aligned}$$

By Lemma 22 we have $\|\mathbf{Q}\mathbf{r}\|_2 \leq \|\mathbf{Q}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2}$, $\|\mathbf{P}\mathbf{r}\|_2 \leq \|\mathbf{P}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2}$. Furthermore, since $\mathbf{r} \in \text{Im}(\Sigma)$, we can rewrite $\|\mathbf{Q}(\mathbf{r} - \hat{\mathbf{r}})\|_2 = \|\mathbf{U}_\mathbf{Q}(\mathbf{r} - \hat{\mathbf{r}})\|_2$ and $\|\mathbf{P}(\mathbf{r} - \hat{\mathbf{r}})\|_2 = \|\mathbf{U}_\mathbf{P}(\mathbf{r} - \hat{\mathbf{r}})\|_2$, where the rows of $\mathbf{U}_\mathbf{Q} \in \mathbb{R}^{d_Q \times D}$ and $\mathbf{U}_\mathbf{P} \in \mathbb{R}^{d_P \times D}$ contain orthonormal bases for $\text{Im}(\mathbf{Q}\Sigma)$ and $\text{Im}(\mathbf{P}\Sigma)$, respectively. Using Lemma 22 and $d_Q \vee d_P \leq \text{rank}(\Sigma)$, we have with probability at least $1 - 2\exp(-u)$

$$\begin{aligned} \|\mathbf{Q}(\mathbf{r} - \hat{\mathbf{r}})\|_2 &\lesssim \|\mathbf{Q}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}, \\ \|\mathbf{P}(\mathbf{r} - \hat{\mathbf{r}})\|_2 &\lesssim \|\mathbf{P}\tilde{X}\|_{\psi_2} \|\tilde{Y}\|_{\psi_2} \sqrt{\frac{\text{rank}(\Sigma) + u}{N}}. \end{aligned} \tag{A.11}$$

Furthermore, by Theorem 10 whenever $N > C(\text{rank}(\mathbf{\Sigma}) + u)\|\sqrt{\mathbf{\Sigma}^\dagger}\tilde{X}\|_{\psi_2}^4$ we get with probability $1 - 2\exp(-u)$ for each $\mathbf{B} \in \{\mathbf{P}, \mathbf{Q}\}$

$$\|\mathbf{P}(\mathbf{\Sigma}^\dagger - \hat{\mathbf{\Sigma}}^\dagger)\mathbf{B}\|_2 \lesssim \|\mathbf{P}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\|\mathbf{B}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\sqrt{\frac{\text{rank}(\mathbf{\Sigma}) + u}{N}}. \quad (\text{A.12})$$

Conditioned on all four events we now have

$$\begin{aligned} \|\mathbf{P}(\mathbf{\Sigma}^\dagger - \hat{\mathbf{\Sigma}}^\dagger)\mathbf{P}\|_2\|\mathbf{P}\mathbf{r}\|_2 &\lesssim \|\tilde{Y}\|_{\psi_2}\|\mathbf{P}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\sqrt{\kappa(\mathbf{P}, X)}\sqrt{\frac{\text{rank}(\mathbf{\Sigma}) + u}{N}}, \\ \|\mathbf{P}(\mathbf{\Sigma}^\dagger - \hat{\mathbf{\Sigma}}^\dagger)\mathbf{Q}\|_2\|\mathbf{Q}\mathbf{r}\|_2 &\lesssim \|\tilde{Y}\|_{\psi_2}\|\mathbf{P}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\sqrt{\kappa(\mathbf{Q}, X)}\sqrt{\frac{\text{rank}(\mathbf{\Sigma}) + u}{N}}. \end{aligned}$$

Moreover, since $N > \text{rank}(\mathbf{\Sigma}) + u$ as stated in the beginning, we get

$$\begin{aligned} \|\mathbf{P}\hat{\mathbf{\Sigma}}^\dagger\mathbf{P}\|_2\|\mathbf{P}(\hat{\mathbf{r}} - \mathbf{r})\|_2 &\leq (\|\mathbf{P}\mathbf{\Sigma}^\dagger\mathbf{P}\|_2 + \|\mathbf{P}(\hat{\mathbf{\Sigma}}^\dagger - \mathbf{\Sigma}^\dagger)\mathbf{P}\|_2)\|\mathbf{P}(\hat{\mathbf{r}} - \mathbf{r})\|_2 \\ &\lesssim \|\tilde{Y}\|_{\psi_2}\|\mathbf{P}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\sqrt{\kappa(\mathbf{P}, X)}\sqrt{\frac{\text{rank}(\mathbf{\Sigma}) + u}{N}}, \end{aligned}$$

and

$$\begin{aligned} \|\mathbf{P}\hat{\mathbf{\Sigma}}^\dagger\mathbf{Q}\|_2\|\mathbf{Q}(\hat{\mathbf{r}} - \mathbf{r})\|_2 &\leq (\|\mathbf{P}\mathbf{\Sigma}^\dagger\mathbf{Q}\|_2 + \|\mathbf{P}(\hat{\mathbf{\Sigma}}^\dagger - \mathbf{\Sigma}^\dagger)\mathbf{Q}\|_2)\|\mathbf{Q}(\hat{\mathbf{r}} - \mathbf{r})\|_2 \\ &\lesssim \|\tilde{Y}\|_{\psi_2}\|\mathbf{P}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\sqrt{\kappa(\mathbf{Q}, X)}\sqrt{\frac{\text{rank}(\mathbf{\Sigma}) + u}{N}}, \end{aligned}$$

where we use property (3) in Lemma 18, exploiting the identity $\mathbf{\Sigma}^\dagger = \mathbf{\Sigma}^\dagger\mathbf{\Sigma}\mathbf{\Sigma}^\dagger$, and using Lemma 19 in the second line. Combining the bounds and conditioning on the events above, we have that with probability at least $1 - 4\exp(-u)$

$$\|\mathbf{P}(\mathbf{\Sigma}^\dagger\mathbf{r} - \hat{\mathbf{\Sigma}}^\dagger\hat{\mathbf{r}})\|_2 \lesssim \|\tilde{Y}\|_{\psi_2}\|\mathbf{P}\mathbf{\Sigma}^\dagger\tilde{X}\|_{\psi_2}\sqrt{\kappa(\mathbf{P}, X) \vee \kappa(\mathbf{Q}, X)}\sqrt{\frac{\text{rank}(\mathbf{\Sigma}) + u}{N}},$$

if $N > C(\text{rank}(\mathbf{\Sigma}) + u)(\|\sqrt{\mathbf{\Sigma}^\dagger}\tilde{X}\|_{\psi_2}^4)$. Adjusting the constant C to account for the change in the probability constant and the requirement on N , the claim follows. The proof for the bound on $\|\mathbf{Q}\hat{\mathbf{\Sigma}}^\dagger\mathbf{r}\|_2$ follows analogous lines or argument. \square

Proof of Corollary 15. Assume for the moment $\hat{\mathbf{b}}^\top\mathbf{b} > 0$. In this case \mathbf{b} and $\mathbf{P}\hat{\mathbf{b}}$ are co-linear since $\mathbf{P}\hat{\mathbf{b}} = \langle \mathbf{d}, \hat{\mathbf{b}} \rangle \mathbf{d}$. Therefore, $\mathbf{d} = \frac{\mathbf{b}}{\|\mathbf{b}\|_2} = \frac{\mathbf{P}\hat{\mathbf{b}}}{\|\mathbf{P}\hat{\mathbf{b}}\|_2}$, which implies

$$\|\mathbf{P}(\hat{\mathbf{d}} - \mathbf{d})\|_2 = \|\mathbf{P}\hat{\mathbf{b}}\|_2 \left| \frac{1}{\|\hat{\mathbf{b}}\|_2} - \frac{1}{\|\mathbf{P}\hat{\mathbf{b}}\|_2} \right| \leq \left| \frac{\|\mathbf{P}\hat{\mathbf{b}}\|_2 - \|\hat{\mathbf{b}}\|_2}{\|\hat{\mathbf{b}}\|_2} \right| \leq \frac{\|\mathbf{Q}\hat{\mathbf{b}}\|_2}{\|\hat{\mathbf{b}}\|_2}.$$

Since $\mathbf{Q}\mathbf{d} = \mathbf{0}$ we now have

$$\|\hat{\mathbf{d}} - \mathbf{d}\|_2 = \sqrt{\|\mathbf{P}(\hat{\mathbf{d}} - \mathbf{d})\|_2^2 + \|\mathbf{Q}\hat{\mathbf{d}}\|_2^2} \leq \sqrt{2} \frac{\|\mathbf{Q}\hat{\mathbf{b}}\|_2}{\|\hat{\mathbf{b}}\|_2} \leq \sqrt{2} \frac{\|\mathbf{Q}\hat{\mathbf{b}}\|_2}{\|\mathbf{b}\|_2 - \|\mathbf{P}(\hat{\mathbf{b}} - \mathbf{b})\|_2}.$$

Therefore, it suffices to ensure $\|\mathbf{P}(\hat{\mathbf{b}} - \mathbf{b})\|_2 < \frac{1}{2}\|\mathbf{b}\|_2$, which would also give $\hat{\mathbf{b}}^\top \mathbf{b} > 0$. To that end, Lemma 14 gives that $\|\mathbf{P}(\hat{\mathbf{b}} - \mathbf{b})\|_2 \leq \frac{1}{2}\|\mathbf{b}\|_2$ holds with probability at least $1 - \exp(-u)$ provided

$$N > C(\text{rank}(\boldsymbol{\Sigma}) + u) \left(\|\sqrt{\boldsymbol{\Sigma}^\dagger} \tilde{X}\|_{\psi_2}^4 \vee \frac{\|\tilde{Y}\|_{\psi_2}^2 \|\mathbf{P}\boldsymbol{\Sigma}^\dagger \tilde{X}\|_{\psi_2}^2 \kappa_{\mathbf{P}\mathbf{Q}}}{\|\mathbf{b}\|_2^2} \right).$$

The result follows by bounding $\|\mathbf{Q}\hat{\mathbf{b}}\|_2$ through Lemma 14. \square

Proof of Corollary 16. We first note

$$\|\tilde{Y}\|_{\psi_2} \leq \|f(\mathbf{a}^\top X) - \mathbb{E}f(\mathbf{a}^\top X)\|_{\psi_2} + \sigma_\zeta.$$

Furthermore, we can use a bounded difference inequality for unbounded spaces, see [31, Theorem 1], to get $\|f(\mathbf{a}^\top X) - \mathbb{E}f(\mathbf{a}^\top X)\|_{\psi_2} \lesssim L\|\mathbf{a}^\top X\|_{\psi_2} \lesssim \sqrt{K}L\sigma_{\mathbf{P}}$, where we used the strict sub-Gaussianity in (2.5) and $\mathbf{a}^\top \boldsymbol{\Sigma} \mathbf{a} = \sigma_{\mathbf{P}}^2$. Therefore,

$$\|\tilde{Y}\|_{\psi_2} \lesssim KL\sigma_{\mathbf{P}} + \sigma_\zeta.$$

Using $\kappa(\mathbf{P}, X) \lesssim K^2$, $\|\sqrt{\boldsymbol{\Sigma}^\dagger} \tilde{X}\|_{\psi_2} = \sqrt{K}$, and plugging in the bounds on $\|\tilde{Y}\|_{\psi_2}$, as well as $\|\mathbf{P}\boldsymbol{\Sigma}^\dagger \tilde{X}\|_{\psi_2} = \frac{\sqrt{K}}{\sigma_{\mathbf{P}}}$ and $\|\mathbf{Q}\boldsymbol{\Sigma}^\dagger \tilde{X}\|_{\psi_2} = \sqrt{K}\|\mathbf{Q}\boldsymbol{\Sigma}^\dagger \mathbf{Q}\|_2$, the claim follows. \square

Lemma 23. *If Z is sub-Gaussian and E an event with $\mathbb{P}(E) > 0$, then $Z|E$ is also sub-Gaussian.*

Proof. Assume without loss of generality that $Z \in \mathbb{R}$. The argument for vectors then follows by the definition. We use the characterization of sub-Gaussianity by the moment bound [54, Proposition 2.5.2, b)]. Let $p \geq 1$. By the law of total expectation it follows

$$\mathbb{E}[|Z|^p] = \mathbb{E}[|Z|^p|E]\mathbb{P}(E) + \mathbb{E}[|Z|^p|E^c]\mathbb{P}(E^c) \geq \mathbb{E}[|Z|^p|E]\mathbb{P}(E).$$

Dividing by $\mathbb{P}(E)$ and using the monotonicity of the p -th root yields

$$(\mathbb{E}[|Z|^p|E])^{1/p} \leq \frac{(\mathbb{E}[|Z|^p])^{1/p}}{\mathbb{P}(E)^{1/p}} \lesssim \frac{\|Z\|_{\psi_2} \sqrt{p}}{\mathbb{P}(E)},$$

where we used $\mathbb{P}(E) \leq 1$ and the sub-Gaussianity of Z in the last inequality. \square

Proof of Theorem 17. Using $0 \leq \max_{s=\pm 1} \langle \mathbf{su}_1(\hat{\mathbf{M}}_J), \mathbf{a} \rangle \leq 1$, we first compute

$$\begin{aligned} \min_{s=\pm 1} \|\mathbf{su}_1(\hat{\mathbf{M}}_J) - \mathbf{a}\|_2^2 &= \min_{s=\pm 1} 2(1 - \langle \mathbf{su}_1(\hat{\mathbf{M}}_J), \mathbf{a} \rangle) \\ &\leq 2(1 - \langle \mathbf{u}_1(\hat{\mathbf{M}}_J), \mathbf{a} \rangle^2) \leq 2\|\mathbf{Q}\mathbf{u}_1(\hat{\mathbf{M}}_J)\|_2^2. \end{aligned}$$

The Davis-Kahan Theorem [7, Theorem 7.3.1] for $\hat{\mathbf{M}}_J$ and $\mathbf{P} = \mathbf{a}\mathbf{a}^\top$ then gives

$$\|\mathbf{Q}\mathbf{u}_1(\hat{\mathbf{M}}_J)\|_2^2 = \|\mathbf{Q}\mathbf{u}_1(\hat{\mathbf{M}}_J)\mathbf{u}_1(\hat{\mathbf{M}}_J)^\top\|_F^2 \leq \frac{\|\mathbf{Q}(\mathbf{P} - \hat{\mathbf{M}}_J)\|_F^2}{\lambda_1(\hat{\mathbf{M}}_J)^2} = \frac{\|\mathbf{Q}\hat{\mathbf{M}}_J\|_F^2}{\lambda_1(\hat{\mathbf{M}}_J)^2}. \quad (\text{A.13})$$

It remains to find a concentration bound for the matrix $\mathbf{Q}\hat{\mathbf{M}}_J$ around zero and a lower bound for the leading eigenvalue $\lambda_1(\hat{\mathbf{M}}_J)$. Now let $\pi : [|I_J|] \rightarrow I_J$ be bijective and introduce the matrix

$$\hat{\mathbf{B}}_J = \left[\sqrt{\hat{\rho}_{J,\pi(1)}} \hat{\mathbf{b}}_{J,\pi(1)} \mid \cdots \mid \sqrt{\hat{\rho}_{J,\pi(|I_J|)}} \hat{\mathbf{b}}_{J,\pi(|I_J|)} \right] \in \mathbb{R}^{D \times |I_J|},$$

satisfying $\hat{\mathbf{M}}_J = \hat{\mathbf{B}}_J \hat{\mathbf{B}}_J^\top$, and thus $\lambda_1(\hat{\mathbf{M}}_J) = \lambda_1(\hat{\mathbf{B}}_J \hat{\mathbf{B}}_J^\top)$. Using $\|\mathbf{G}\mathbf{H}\|_F \leq \|\mathbf{G}\|_2 \|\mathbf{H}\|_F$ and $\|\mathbf{G}\|_F = \|\mathbf{G}^\top\|_F$, which hold for arbitrary \mathbf{G} and \mathbf{H} , yields

$$\|\mathbf{Q}\hat{\mathbf{M}}_J\|_F^2 = \|\mathbf{Q}\hat{\mathbf{B}}_J \hat{\mathbf{B}}_J^\top\|_F^2 \leq \|\hat{\mathbf{B}}_J\|_2^2 \|\mathbf{Q}\hat{\mathbf{B}}_J\|_F^2 \leq \lambda_1(\hat{\mathbf{M}}_J) \|\mathbf{Q}\hat{\mathbf{B}}_J\|_F^2.$$

By Lemma 23, conditioning on a sub-Gaussian random vector gives a sub-Gaussian random vector. Thus, by (4.4) we have simultaneously for all $\ell \in I_J$

$$\|\mathbf{Q}\hat{\mathbf{b}}_{J,\ell}\|_2^2 \lesssim \|\tilde{Y}_{J,\ell}\|_{\psi_2}^2 \|\mathbf{Q}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \kappa_{J,\ell} \frac{\text{rank}(\Sigma_{J,\ell}) + \log(|I_J|) + u}{N_{J,\ell}}, \quad (\text{A.14})$$

with probability at least $1 - \exp(-u)$, provided $N_{J,\ell} > C(\text{rank}(\Sigma_{J,\ell}) + \log(|I_J|) + u)K_{J,\ell}^4$ for all $\ell \in I_J$. By definition of I_J , we have $N_{J,\ell} > \alpha N J^{-1}$, and thus the previous condition is satisfied whenever for all $\ell \in I_J$

$$\frac{N}{J} > \frac{C(\text{rank}(\Sigma_{J,\ell}) + \log(J) + u)K_{J,\ell}^4}{\alpha},$$

which is implied by $N > CJK^4(\text{rank}(\Sigma) + \log(J) + u)/\alpha$. Under the same conditions and with similar probability we obtain

$$\begin{aligned} \|\mathbf{Q}\hat{\mathbf{M}}_J\|_F^2 &\leq \lambda_1(\hat{\mathbf{M}}_J) \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \|\mathbf{Q}\hat{\mathbf{b}}_{J,\ell}\|_2^2 \\ &\lesssim \lambda_1(\hat{\mathbf{M}}_J) \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \|\tilde{Y}_{J,\ell}\|_{\psi_2}^2 \|\mathbf{Q}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \kappa_{J,\ell} \frac{\text{rank}(\Sigma) + \log(J) + u}{N_{J,\ell}} \\ &\lesssim \frac{\lambda_1(\hat{\mathbf{M}}_J)(\text{rank}(\Sigma) + \log(J) + u)}{\alpha N J} \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \|\mathbf{Q}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \kappa_{J,\ell}, \end{aligned}$$

where we used $N_{J,\ell} > \alpha N J^{-1}$ and $\|\tilde{Y}_{J,\ell}\|_{\psi_2} \lesssim |\mathcal{R}_{J,\ell}| = J^{-1}$ (definition of $\mathcal{R}_{J,\ell}$) in the last inequality. To bound $\lambda_1(\hat{\mathbf{M}}_J)$ from below, we note that on the same event where bound (A.14) holds, we also have for all $\ell \in I_J$

$$\begin{aligned} \|\mathbf{P}(\hat{\mathbf{b}}_{J,\ell} - \mathbf{b}_{J,\ell})\|_2^2 &\lesssim \|\tilde{Y}_{J,\ell}\|_{\psi_2}^2 \|\mathbf{P}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \kappa_{J,\ell} \frac{\text{rank}(\Sigma_{J,\ell}) + \log(J) + u}{N_{J,\ell}} \\ &\lesssim \frac{\text{rank}(\Sigma) + \log(J) + u}{\alpha N J} \|\mathbf{P}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell}\|_{\psi_2}^2 \kappa_{J,\ell} \end{aligned} \quad (\text{A.15})$$

Thus, multiplying $\hat{\mathbf{M}}_J$ from left and right by \mathbf{a} and using (A), we obtain

$$\begin{aligned} \lambda_1(\hat{\mathbf{M}}_J) &\geq \mathbf{a}^\top \hat{\mathbf{M}}_J \mathbf{a} = \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \langle \mathbf{a}, \hat{\mathbf{b}}_{J,\ell} \rangle^2 = \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \left(\|b_{J,\ell}\|_2^2 - \left\| \mathbf{P}(b_{J,\ell} - \hat{b}_{J,\ell}) \right\|_2^2 \right) \\ &\geq \sum_{\ell \in I_J} \hat{\rho}_{J,\ell} \left(\|b_{J,\ell}\|_2^2 - \frac{C(\text{rank}(\Sigma) + \log(J) + u)}{\alpha N J} \left\| \mathbf{P}\Sigma_{J,\ell}^\dagger \tilde{X}_{J,\ell} \right\|_{\psi_2}^2 \kappa_{J,\ell} \right) \end{aligned}$$

for some universal constant $C > 0$. The result follows by plugging the bounds on $\|\widehat{\mathbf{Q}\mathbf{M}_J}\|_F^2$ and $\lambda_1(\widehat{\mathbf{M}}_J)$ into (A.13), and using the definition of $\varepsilon_{N,J,u}$. \square

References

- [1] ADAMCZAK, R., LITVAK, A., PAJOR, A., AND TOMCZAK-JAEGERMANN, N. (2010). Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of the American Mathematical Society* **23**, 2, 535–561. [MR2601042](#)
- [2] ADRAGNI, K. P. AND COOK, R. D. (2009). Sufficient dimension reduction and prediction in regression. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **367**, 1906, 4385–4405. [MR2546393](#)
- [3] ANDERSON, T. W. (2003). *An Introduction to Multivariate Statistical Analysis*, 3rd ed. John Wiley & Sons. [MR1990662](#)
- [4] ARBEL, J., MARCHAL, O., AND NGUYEN, H. D. (2019). On strict sub-gaussianity, optimal proxy variance and symmetry for bounded random variables. *to appear in ESAIM: Probability and statistics arxiv:1901.09188*. [MR4053001](#)
- [5] BALABDAOUI, F., GROENEBOOM, P., AND HENDRICKX, K. (2019). Score estimation in the monotone single-index model. *Scandinavian Journal of Statistics* **46**, 2, 517–544. [MR3948566](#)
- [6] BELKIN, M., ROSASCO, L., AND VITO, E. D. (2010). On learning with integral operators. *Journal of Machine Learning Research* **11**, 905–934. [MR2600634](#)
- [7] BHATIA, R. (2013). *Matrix analysis*. Vol. **169**. Springer Science & Business Media. [MR1477662](#)
- [8] BICKEL, P. AND LEVINA, E. (2008a). Covariance Regularization by thresholding. *The Annals of Statistics* **36**, 6, 2577–2604. [MR2485008](#)
- [9] BICKEL, P. AND LEVINA, E. (2008b). Regularized estimation of large covariance matrices. *The Annals of Statistics* **36**, 1, 199–227. [MR2387969](#)
- [10] BRILLINGER, D. (1982). A generalized linear model with “gaussian” regressor variables. *A Festschrift For Erich L. Lehmann*, 97. [MR0689741](#)
- [11] CAI, T., LIU, W., AND LUO, X. (2011). A constrained ℓ_1 minimization approach to sparse precision matrix estimation. *Journal of the American Statistical Association* **106**, 594–607. [MR2847973](#)
- [12] CAI, T., LIU, W., AND ZHOU, H. (2016). Estimating sparse precision matrix: optimal rates of convergence and adaptive estimation. *Annals of Statistics* **44**, 455–488. [MR3476606](#)
- [13] CAI, T., REN, Z., AND ZHOU, H. H. (2016). Estimating structured high-dimensional covariance and precision matrices: Optimal rates and adaptive estimation. *Electronic Journal of Statistics* **10**, 1, 1–59. [MR3466172](#)
- [14] DALALYAN, A. S., JUDITSKY, A., AND SPOKOINY, V. (2008). A new algorithm for estimating the effective dimension-reduction subspace. *Journal of Machine Learning Research* **9**, Aug, 1647–1678. [MR2438819](#)

- [15] DENNIS COOK, R. (2000). SAVE: A method for dimension reduction and graphics in regression. *Communications in Statistics – Theory and Methods* **29**, 9-10, 2109–2121. [MR1645673](#)
- [16] DOBSON, A. AND BARNETT, A. (2008). *An introduction to generalized linear models*. Chapman and Hall/CRC. [MR2068205](#)
- [17] FAN, J., LIAO, Y., AND LIU, H. (2016). An overview of the estimation of large covariance and precision matrices. *The Econometrics Journal* **19**, 1, C1–C32. [MR3501529](#)
- [18] FRIEDMAN, J., HASTIE, T., AND TIBSHIRANI, R. (2008). Sparse inverse covariance estimation with the graphical lasso. *Biostatistics* **9**, 3, 432–441.
- [19] GYÖRFI, L., KOHLER, M., KRZYŻAK, A., AND WALK, H. (2006). *A distribution-free theory of nonparametric regression*. Springer Science & Business Media. [MR1920390](#)
- [20] HRISTACHE, M., JUDITSKY, A., AND SPOKOINY, V. (2001). Direct estimation of the index coefficient in a single-index model. *The Annals of Statistics* **29**, 3, 595–623. [MR1865333](#)
- [21] HUANG, J., LIU, N., POURAHMADI, M., AND LIU, L. (2006). Covariance matrix selection and estimation via penalised normal likelihood. *Biometrika* **93**, 1, 85–98. [MR2277742](#)
- [22] ICHIMURA, H. (1993). Semiparametric least squares (SLS) and weighted SLS estimation of single-index models. *Journal of Econometrics* **58**, 1-2, 71–120. [MR1230981](#)
- [23] KAKADE, S. M., KANADE, V., SHAMIR, O., AND KALAI, A. (2011). Efficient learning of generalized linear and single index models with isotonic regression. In *Advances in Neural Information Processing Systems*. 927–935.
- [24] KALAI, A. T. AND SASTRY, R. (2009). The isotron algorithm: High-dimensional isotonic regression. In *COLT*. Citeseer.
- [25] KLASING, K., ALTHOFF, D., WOLLHERR, D., AND BUSS, M. (2009). Comparison of surface normal estimation methods for range sensing applications. In *2009 IEEE International Conference on Robotics and Automation*. IEEE, 3206–3211.
- [26] KLOCK, T., LANTERI, A., AND VIGOGNA, S. (2020). Estimating multi-index models with response-conditional least squares. *arXiv preprint arXiv:2003.04788*.
- [27] KOLTCHINSKII, V. (2017). Asymptotically efficient estimation of smooth functionals of covariance operators. *arXiv preprint arXiv:1710.09072*.
- [28] KOLTCHINSKII, V., LÖFFLER, M., NICKL, R., AND OTHERS. (2020). Efficient estimation of linear functionals of principal components. *The Annals of Statistics* **48**, 1, 464–490. [MR4065170](#)
- [29] KOLTCHINSKII, V., LOUNICI, K., AND OTHERS. (2017a). Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli* **23**, 1, 110–133. [MR3556768](#)
- [30] KOLTCHINSKII, V., LOUNICI, K., AND OTHERS. (2017b). Normal approximation and concentration of spectral projectors of sample covariance. *The Annals of Statistics* **45**, 1, 121–157. [MR3611488](#)

- [31] KONTOROVICH, A. (2014). Concentration in unbounded metric spaces and algorithmic stability. *Proceedings of Machine Learning Research*, Vol. **32**. *Proceedings of Machine Learning Research*, 28–36.
- [32] KUCHIBHOTLA, A. K., PATRA, R. K., AND OTHERS. (2020). Efficient estimation in single index models through smoothing splines. *Bernoulli* **26**, 2, 1587–1618. [MR4058379](#)
- [33] LI, B. AND WANG, S. (2007). On directional regression for dimension reduction. *Journal of the American Statistical Association* **102**, 479, 997–1008. [MR2354409](#)
- [34] LI, K.-C. (1991). Sliced inverse regression for dimension reduction. *Journal of the American Statistical Association* **86**, 414, 316–327. [MR1137117](#)
- [35] LI, K.-C. AND DUAN, N. (1989). Regression analysis under link violation. *The Annals of Statistics* **17**, 3, 1009–1052. [MR1015136](#)
- [36] LI, W., CHEN, Y., VONG, S., AND LUO, Q. (2018). Some refined bounds for the perturbation of the orthogonal projection and the generalized inverse. *Numerical Algorithms* **79**, 2, 1–21. [MR3857930](#)
- [37] LITTLE, A. V., MAGGIONI, M., AND ROSASCO, L. (2017). Multiscale geometric methods for data sets I: Multiscale SVD, noise and curvature. *Applied and Computational Harmonic Analysis* **43**, 3, 504–567. [MR3683673](#)
- [38] LIU, W. AND LUO, X. (2015). Fast and adaptive sparse precision matrix estimation in high dimensions. *Journal of Multivariate Analysis* **135**, 153–162. [MR3306432](#)
- [39] MA, Y. AND ZHU, L. (2013). A review on dimension reduction. *International Statistical Review* **81**, 1, 134–150. [MR3047506](#)
- [40] MCALEER, M. AND DA VEIGA, B. (2008). Single-index and portfolio models for forecasting value-at-risk thresholds. *Journal of Forecasting* **27**, 3, 217–235. [MR2420302](#)
- [41] MEINSHAUSEN, N. AND BÜHLMANN, P. (2006). High-dimensional graphs and variable selection with the lasso. *The Annals of Statistics* **34**, 3, 1436–1462. [MR2278363](#)
- [42] MENG, L. AND ZHENG, B. (2010). The optimal perturbation bounds of the Moore-Penrose inverse under the Frobenius norm. *Linear Algebra and its Applications* **432**, 4, 956–963. [MR2577640](#)
- [43] MITRA, N. J. AND NGUYEN, A. (2003). Estimating surface normals in noisy point cloud data. In *Proceedings of the nineteenth annual symposium on Computational geometry*. ACM, 322–328.
- [44] OUYANG, D. AND FENG, H.-Y. (2005). On the normal vector estimation for point cloud data from smooth surfaces. *Computer-Aided Design* **37**, 10, 1071–1079.
- [45] PLAN, Y. AND VERSHYNIN, R. (2016). The generalized lasso with non-linear observations. *IEEE Transactions on information theory* **62**, 3, 1528–1537. [MR3472264](#)
- [46] PLAN, Y., VERSHYNIN, R., AND YUDOVINA, E. (2016). High-dimensional estimation with geometric constraints. *Information and Inference: A Journal of the IMA* **6**, 1, 1–40. [MR3636866](#)
- [47] RADCHENKO, P. (2015). High dimensional single index models. *Journal of*

- Multivariate Analysis* 139, 266–282. [MR3349492](#)
- [48] REISS, M. AND WAHL, M. (2020). Nonasymptotic upper bounds for the reconstruction error of PCA. *The Annals of Statistics* **48**, 2, 1098–1123. [MR4102689](#)
- [49] SRIVASTAVA, N. AND VERSHYNIN, R. (2013). Covariance estimation for distributions with $2 + \varepsilon$ moments. *The Annals of Probability* **41**, 5, 3081–3111. [MR3127875](#)
- [50] TIAN, Y. AND CHENG, S. (2004). Some identities for Moore–Penrose inverses of matrix products. *Linear and Multilinear Algebra* **52**, 6, 405–420. [MR2102196](#)
- [51] VAN WIERINGEN, W. N. (2019). The Generalized Ridge Estimator of the Inverse Covariance Matrix. *Journal of Computational and Graphical Statistics* **28**, 4, 932–942. [MR4045859](#)
- [52] VERSHYNIN, R. (2012a). How close is the sample covariance matrix to the actual covariance matrix? *Journal of Theoretical Probability* **25**, 3, 655–686. [MR2956207](#)
- [53] VERSHYNIN, R. (2012b). Introduction to the non-asymptotic analysis of random matrices. 210–268. [MR2963170](#)
- [54] VERSHYNIN, R. (2018). *High-dimensional probability: An introduction with applications in data science*. Vol. **47**. Cambridge University Press. [MR3837109](#)
- [55] WEDIN, P.-Å. (1973). Perturbation theory for pseudo-inverses. *BIT Numerical Mathematics* **13**, 2, 217–232. [MR0336982](#)
- [56] WEYL, H. (1912). Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Mathematische Annalen* **71**, 4, 441–479. [MR1511670](#)
- [57] XU, X. (2018). On the perturbation of the Moore-Penrose inverse of a matrix. *arXiv preprint arXiv:1809.00203*. [MR4060045](#)
- [58] YASKOV, P. (2014). Lower bounds on the smallest eigenvalue of a sample covariance matrix. *Electronic Communications in Probability* **19**. [MR3291620](#)
- [59] YU, Y., WANG, T., AND SAMWORTH, R. (2014). A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika* **102**, 2, 315–323. [MR3371006](#)