

Bootstrapping the operator norm in high dimensions: Error estimation for covariance matrices and sketching

MILES E. LOPES^{1,a}, N. BENJAMIN ERICHSON^{2,b} and MICHAEL W. MAHONEY^{2,c}

¹*Department of Statistics, University of California, Davis, Mathematical Sciences Building, Davis, CA, 95616, USA, ^amelopes@ucdavis.edu*

²*Department of Statistics, University of California, Berkeley, Evans Hall, Berkeley, CA, 94720, USA and International Computer Science Institute, Berkeley, CA, 94704, USA, ^berichson@berkeley.edu,*

^c*mmahoney@stat.berkeley.edu*

Although the operator (spectral) norm is one of the most widely used metrics for covariance estimation, comparatively little is known about the fluctuations of error in this norm. To be specific, let $\widehat{\Sigma}$ denote the sample covariance matrix of n i.i.d. observations in \mathbb{R}^p that arise from a population matrix Σ , and let $T_n = \sqrt{n} \|\widehat{\Sigma} - \Sigma\|_{\text{op}}$. In the setting where the eigenvalues of Σ have a decay profile of the form $\lambda_j(\Sigma) \asymp j^{-2\beta}$, we analyze how well the bootstrap can approximate the distribution of T_n . Our main result shows that up to factors of $\log(n)$, the bootstrap can approximate the distribution of T_n with respect to the Kolmogorov metric at the rate of $n^{-\frac{\beta-1/2}{6\beta+4}}$, which does not depend on the ambient dimension p . In addition, we offer a supporting result of independent interest that establishes a high-probability upper bound for T_n based on flexible moment assumptions. More generally, we discuss the consequences of our work beyond covariance matrices, and show how the bootstrap can be used to estimate the errors of sketching algorithms in randomized numerical linear algebra (RandNLA). An illustration of these ideas is also provided with a climate data example.

Keywords: Bootstrap; error estimation; high-dimensional statistics; covariance estimation; randomized numerical linear algebra; sketching

1. Introduction

Within the areas of covariance estimation and principal components analysis, it is of central importance to understand how well a sample covariance matrix $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n X_i X_i^\top$ approximates its population version $\Sigma = \mathbb{E}[X_1 X_1^\top]$, where $X_1, \dots, X_n \in \mathbb{R}^p$ are centered i.i.d. observations. In particular, a major line of research in high-dimensional statistics has focused on the problem of deriving non-asymptotic bounds for the operator (spectral) norm error

$$T_n = \sqrt{n} \|\widehat{\Sigma} - \Sigma\|_{\text{op}},$$

where the norm is defined as $\|A\|_{\text{op}} = \sup_{\|u\|_2=1} \|Au\|_2$. A partial overview of work on this problem, as well as some of its extensions, may be found in the papers (Adamczak et al. (2011), Bickel and Levina (2008), Bunea and Xiao (2015), Cai, Zhang and Zhou (2010), Koltchinskii and Lounici (2017a), Lounici (2014), Minsker (2017), Rudelson (1999), among numerous others).

As a whole, this line of work offers many conceptual insights into the ways that error is influenced by model assumptions. However, the literature is less complete with regard to inference, and there are not many guarantees for the problem of constructing confidence intervals for T_n , which is equivalent

to constructing numerical bounds on the error of $\widehat{\Sigma}$, or confidence regions for Σ . Accordingly, the challenges of inference on high-dimensional covariance matrices have stimulated much recent activity, and there has been a particular interest to understand the limits of the bootstrap in this context (Johnstone and Paul (2018) §X.C, El Karoui and Purdom (2019), Han, Xu and Zhou (2018), Lopes, Blandino and Aue (2019), Naumov, Spokoiny and Ulyanov (2019), Yao and Lopes (2022)).

Simultaneously with these developments, the burgeoning field of randomized numerical linear algebra (RandNLA) has generated many other error estimation problems of a similar nature (Drineas and Mahoney, 2018, Halko, Martinsson and Tropp, 2011, Kannan and Vempala, 2017, Mahoney, 2011, Woodruff, 2014). A prototypical example deals with computing a fast randomized approximation of the product $A^T A$, where A is a very large matrix. Most commonly, the matrix A is randomly “sketched” into a much shorter matrix \tilde{A} , which can then be used to quickly compute $\tilde{A}^T \tilde{A}$ as an approximation to $A^T A$. In turn, it is necessary to assess the unknown error $\|\tilde{A}^T \tilde{A} - A^T A\|_{\text{op}}$, which leads to a notable parallel with the statistical literature: There are many existing theoretical error bounds, but very few tools for numerical error estimation (see Sections 1.2 and 4). Furthermore, the operator norm is of special importance, because it governs the accuracy of numerous matrix computations, and it frequently appears in numerical analysis.

Motivated by the challenges above, this paper aims to quantify how well the bootstrap can approximate the error distribution $\mathcal{L}(T_n)$ for sample covariance matrices, and likewise in the context of RandNLA. Specifically, we consider a setup where Σ has low “effective rank” and its ordered eigenvalues satisfy a decay profile of the form

$$\lambda_j(\Sigma) \asymp j^{-2\beta}, \tag{1.1}$$

for some parameter $\beta > 1/2$. Variations of this setting have drawn considerable attention in recent years, especially in connection with principal components analysis (e.g., Bunea and Xiao (2015), Koltchinskii, Löffler and Nickl (2020), Koltchinskii and Lounici (2017a,b), Lounici (2014), Reiss and Wahl (2020), and Naumov, Spokoiny and Ulyanov (2019), among others). Moreover, the condition (1.1) corresponds to problems where sketching algorithms can be highly effective.

1.1. Contributions

To briefly outline our main result, let the Kolmogorov metric be denoted as $d_K(\mathcal{L}(U), \mathcal{L}(V)) = \sup_{t \in \mathbb{R}} |\mathbb{P}(U \leq t) - \mathbb{P}(V \leq t)|$ for two generic random variables U and V , and let T_n^* denote the bootstrap version of T_n , obtained by sampling with replacement from (X_1, \dots, X_n) . Then, as long as (1.1) is satisfied and the observations have suitable tail behavior, it follows that the bound

$$d_K\left(\mathcal{L}(T_n), \mathcal{L}(T_n^*|X)\right) \leq c n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c \tag{1.2}$$

holds with probability at least $1 - \frac{c}{n}$, where $\mathcal{L}(T_n^*|X)$ is the conditional distribution of T_n^* given the observations. (Going forward, we use symbols such as c, c_0, c_1 , etc. to denote positive constants not depending on n whose values may change at each occurrence.) Most importantly, this non-asymptotic bound does not depend on the ambient dimension p , and explicitly accounts for the structural complexity parameter β .

From the standpoint of methodology, our work illustrates new possibilities for applying the bootstrap in the domains of computer science and applied mathematics. At this interface, the bootstrap has a largely untapped potential to make an impact, because error estimation allows randomized computations to be done *adaptively*, so that “just enough” work is done. More specifically, the estimated error

of a rough initial solution can be used to predict how much extra computation is needed to reach a high-quality solution—and this will be demonstrated numerically in Section 4. Lastly, to put this type of application into historical perspective, it is notable that the bootstrap has been traditionally labeled as “computationally intensive”, and so in this respect, it is relatively novel to use the bootstrap *in the service of computation*.

With regard to theoretical considerations, our work contributes to recent developments on bootstrap methods, as well as covariance estimation. For the bootstrap, we expand upon the progress achieved in the series of papers (Chernozhukov, Chetverikov and Kato, 2013, 2014, 2016, 2017), which address bootstrap approximations for “max statistics” of the form $M_n = \sup_{f \in \mathcal{F}} \mathbb{G}_n(f)$, where \mathcal{F} is a class of functions, and $\mathbb{G}_n(f) = \frac{1}{\sqrt{n}} \sum_{i=1}^n (f(X_i) - \mathbb{E}[f(X_i)])$. The basic similarity between M_n and T_n is that they can be represented in a common form, due to the variational representation of $\|\cdot\|_{\text{op}}$. Nevertheless, the statistic T_n seems to present certain technical obstructions with regard to previous results. First, in order to handle the metric d_K , the mentioned works typically require a “minimum variance condition” such as

$$\inf_{f \in \mathcal{F}} \text{var}(\mathbb{G}_n(f)) \geq c, \tag{1.3}$$

which poses a difficulty in our setting, because the minimum variance may decrease rapidly with n . As a result, a challenge arises in showing that our statistic is well approximated (in d_K) by the supremum $\sup_{f \in \mathcal{F}'_n} \mathbb{G}_n(f)$, where $\mathcal{F}'_n \subset \mathcal{F}$ is a “nice” subset for which $\inf_{f \in \mathcal{F}'_n} \text{var}(\mathbb{G}_n(f))$ decreases slowly with n . Second, further challenges are encountered when controlling the discretization error that comes from replacing \mathcal{F} with a discrete ϵ -net. More specifically, this error is significant in our analysis because the relevant class \mathcal{F} is exponentially larger than VC-type — in the sense that ϵ -covering numbers grow exponentially in $1/\epsilon$, rather than polynomially. By contrast, previous applications of bootstrap approximation results for max statistics have often been concerned with VC-type function classes, which allow for strong control of the discretization error.

Another segment of our work deals with tail bounds for $\|\widehat{\Sigma} - \Sigma\|_{\text{op}}$ that do not depend on the ambient dimension p , as studied in (Hsu, Kakade and Zhang, 2012, Koltchinskii and Lounici, 2017a, Minsker, 2017, Oliveira, 2010, Rudelson and Vershynin, 2007). In the setting of (1.1), this line of work shows that if the observations satisfy $\|X_i\|_2 \leq c$ almost surely, or $\|\langle u, X_i \rangle\|_{\psi_2} \asymp \|\langle u, X_i \rangle\|_2$ for all unit vectors u , then the operator norm error can be bounded as $\|\widehat{\Sigma} - \Sigma\|_{\text{op}} \leq c n^{-1/2} \log(n)^c$ with high probability. However, the ℓ_2 -boundedness condition is often restrictive, while the ψ_2 - L_2 equivalence condition is not well-suited to the discrete distributions that arise from resampling (Vershynin, 2018, §3.4.2). Consequently, as a way to streamline our analysis of both (X_1, \dots, X_n) and the bootstrap samples (X_1^*, \dots, X_n^*) , it is of interest to develop a bound that can be applied in a more general-purpose way. Indeed, an extension of this type is also suggested briefly in the paper (Rudelson and Vershynin, 2007), but to the best of our knowledge, such a result has not been available without involving dependence on the ambient dimension. Accordingly, one of our secondary main results (Theorem 2.2) serves this purpose by showing that if $q \geq 3$, and if ξ_1, \dots, ξ_n are i.i.d. random elements of a separable Hilbert space, then

$$\left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \otimes \xi_i - \mathbb{E}[\xi_i \otimes \xi_i] \right\|_{\text{op}}^q \right)^{1/q} \leq c \cdot \|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}} \cdot \left(\sqrt{\frac{\mathfrak{r}(q)}{n^{1-3/q}}} \vee \frac{\mathfrak{r}(q)}{n^{1-3/q}} \right).$$

where $\mathfrak{r}(q)$ is a parameter that plays the role of an effective rank, and satisfies $\mathfrak{r}(q) < \infty$ as long as $\|\xi_1\|$ has at least $2q$ moments. In particular, the proof of Theorem 2.2 extends the approach of Rudelson and Vershynin (2007) based on non-commutative Khintchine inequalities. For comparison, the paper (Chen, Gittens and Tropp, 2012) has also extends the approach of Rudelson and Vershynin (2007) to allow for sums of more general types of random matrices, but it differs from our work insofar as its results depend on the ambient dimension.

1.2. Related work

The most closely related work to ours is the recent paper (Han, Xu and Zhou, 2018), which studies bootstrap approximations for certain variants of T_n . To explain the connection, first recall that T_n may be written in terms of a supremum over the unit sphere $\mathbb{S}^{p-1} \subset \mathbb{R}^p$, namely $T_n = \sup_{u \in \mathbb{S}^{p-1}} \sqrt{n} |u^\top (\widehat{\Sigma} - \Sigma)u|$. As an alternative to this, the paper (Han, Xu and Zhou, 2018) analyzes “sparse versions” of T_n obtained by taking the supremum over $\{u \in \mathbb{S}^{p-1} \mid \|u\|_0 \leq s\}$, where $1 \leq s \leq p$. For these sparse versions of T_n , bootstrap approximation results are obtained in the Kolmogorov metric with rates of the form $s^{9/8}/n^{1/8}$, up to logarithmic factors. As this relates to our work, it should be emphasized that the setting in (Han, Xu and Zhou, 2018) is quite different, since the eigenvalues of Σ are not assumed to decay. The difference becomes most apparent when $s = p$, so that the sphere \mathbb{S}^{p-1} coincides with the set $\{u \in \mathbb{S}^{p-1} \mid \|u\|_0 \leq s\}$. In this case, the analysis without spectral decay requires $p \ll n^{1/9}$ for bootstrap consistency, whereas our setting places no constraints on p . Another work that looks at bootstrapping a structured variant of the operator norm is (Silin and Fan, 2020), but for the different purpose of analyzing the errors of empirical spectral projection matrices. Although this work appeared after our 2019 preprint (Lopes, Erichson and Mahoney, 2019), it is still worth noting that it shares the approach of using a variational representation of the norm and applying the techniques for bootstrapping suprema of empirical processes mentioned earlier.

Next, the recent papers (El Karoui and Purdom, 2019, Yao and Lopes, 2022) look at bootstrapping simple eigenvalues. These works show that if Σ is nearly low-rank, then the bootstrap can consistently approximate the distribution of $L_n = \sqrt{n}(\lambda_j(\widehat{\Sigma}) - \lambda_j(\Sigma))_{1 \leq j \leq j_0}$, provided that the eigenvalues $(\lambda_j(\Sigma))_{1 \leq j \leq j_0}$ each have multiplicity 1, and j_0 is held fixed as $(n, p) \rightarrow \infty$. This illustrates a key difference between the statistics T_n and L_n , since it is known that repeated eigenvalues can interfere with bootstrap approximations for L_n (e.g. Hall et al., 2009), whereas our work will show that the bootstrap can work for T_n even in the presence of repeated eigenvalues. Two more papers on bootstrap methods for high-dimensional sample covariance matrices are (Lopes, Blandino and Aue, 2019) and (Naumov, Spokoiny and Ulyanov, 2019). The first of these generalizes the parametric bootstrap for high-dimensional models without spectral decay, and it establishes consistency for linear spectral statistics, while the latter deals with bootstrapping the error of empirical spectral projection matrices. (We refer to (Koltchinskii and Lounici, 2017b,c) for other related distributional approximation results.)

At a more technical level, our work is related to the paper (Lopes, Lin and Müller, 2020), which analyzes rates of bootstrap approximation for max statistics of the form $M_n = \max_{1 \leq j \leq p} \sqrt{n} \bar{X}_j$, where the vector \bar{X} is the sample average of centered i.i.d. observations $X_1, \dots, X_n \in \mathbb{R}^p$ satisfying a “variance decay” condition. This condition has the form $\sigma_{(j)}^2 \asymp j^{-2\alpha}$ for some fixed parameter $\alpha > 0$, where $\sigma_{(1)}^2 \geq \dots \geq \sigma_{(p)}^2$ are the sorted versions of $\sigma_j^2 = \text{var}(X_{1j})$ for $j = 1, \dots, p$. One of the key steps in the analysis of (Lopes, Lin and Müller, 2020) is to “localize the maximizing index” for M_n . That is, if $\widehat{j} \in \{1, \dots, p\}$ is a random index such that $M_n = \sqrt{n} \bar{X}_{\widehat{j}}$, then the variance decay condition can be used to show that \widehat{j} is likely to fall into a small subset of $\{1, \dots, p\}$. In the present context, we also use this localization technique. Namely, the eigenvalue decay condition $\lambda_j(\Sigma) \asymp j^{-2\beta}$ implies that if $\widehat{u} \in \mathbb{S}^{p-1}$ denotes a maximizing index for $T_n = \sup_{u \in \mathbb{S}^{p-1}} \sqrt{n} |u^\top (\widehat{\Sigma} - \Sigma)u|$, then $\Sigma^{1/2} \widehat{u}$ is likely to fall into a suitable subset of an ellipsoid. In carrying out this localization argument, our Proposition A.1 and its supporting Lemmas A.1 and A.2 in this work are parallel versions of Proposition B.2 and Lemmas B.1 and B.2 from (Lopes, Lin and Müller, 2020). (The localization of the maximizer for the bootstrapped statistic T_n^* is similarly handled here in Appendix E.) Apart from this connection, the analyses in the two papers are essentially different. In particular, it should be emphasized that the high probability bound on $\|\widehat{\Sigma} - \Sigma\|_{\text{op}}$ in Theorem 2.2, the discretization error bound in Proposition 2.1, as well as the intermediate results in Appendices B, C, D, and F are distinct from the earlier paper. Furthermore,

the overall qualitative difference between the analyses is illustrated by the fact that in (Lopes, Lin and Müller, 2020), the rate of bootstrap approximation does not depend on the variance decay parameter α , whereas here, the rate does depend on the spectrum decay parameter β .

Finally, to conclude this section, we describe related work on the estimation of algorithmic error. Here, it is important to note that error estimation has a long history for *deterministic* algorithms, such as those in numerical partial differential equations and finite-element methods, where it is called a *posteriori error estimation* (Ainsworth and Oden, 2000, Babuška and Rheinboldt, 1978, Cangiani et al., 2017, Verfürth, 1994, among many others). However, in the literature on randomized algorithms, error estimation has received much less attention, and for certain types of computations there are only a few papers addressing error estimation: (low-rank approximation: Halko, Martinsson and Tropp, 2011, Liberty et al., 2007, Woolfe et al., 2008), (least-squares: Lopes, Wang and Mahoney, 2018), (classification: Lopes, 2019), (matrix multiplication: Ar et al., 1993, Lopes, Wang and Mahoney, 2019, Sarlós, 2006). Among these works, the only ones to address error estimation for the operator norm are (Halko, Martinsson and Tropp, 2011, Liberty et al., 2007, Woolfe et al., 2008), but this is done specifically for low-rank approximation, which is complementary to our applications. Also, the approach in these works is quite different from bootstrapping, and is based on the idea of bounding error in terms of random “test vectors”, which is rooted in the classical works (Dixon, 1983, Freivalds, 1979). In essence, the main difference between the test-vector approach and bootstrapping is that the former is inherently conservative, whereas the latter can be used to *directly estimate* the error distribution.

Outline. Section 2 presents the problem setup and main results, as well as the proofs for some of these results. Section 3 describes numerical results for inference tasks related to covariance matrices. Section 4 introduces the setting of sketching algorithms, and demonstrates the performance of the bootstrap in synthetic problems, as well as in a climate data example. Lastly, the proof of the main result (Theorem 2.1) and some additional numerical results are deferred to the supplementary material ((Lopes, Erichson and Mahoney, 2023)).

Notation and conventions. For a vector $v \in \mathbb{R}^m$, and a number $q \geq 1$, the ℓ_q -norm is $\|v\|_q = (\sum_{j=1}^m |v_j|^q)^{1/q}$. The unit sphere for the ℓ_2 -norm in \mathbb{R}^m is \mathbb{S}^{m-1} . For a real matrix M , its Frobenius norm is $\|M\|_F = \sqrt{\text{tr}(M^T M)}$, and its Schatten- q norm is $\|M\|_{S_q} = \text{tr}((M^T M)^{q/2})^{1/q}$. The identity matrix of size $m \times m$ is I_m , and the standard basis vectors in \mathbb{R}^m are $\{e_1, \dots, e_m\}$. The sorted singular values of a real matrix M are written as $\sigma_j(M) \geq \sigma_{j+1}(M)$, and similarly, if M is symmetric, then the sorted eigenvalues are written as $\lambda_j(M) \geq \lambda_{j+1}(M)$. For a random variable ξ , the L_q norm is $\|\xi\|_q = (\mathbb{E}[|\xi|^q])^{1/q}$. Also, if $\psi_q(x) = \exp(x^q) - 1$, then the ψ_q -Orlicz norm is given by $\|\xi\|_{\psi_q} = \inf\{r > 0 \mid \mathbb{E}[\psi_q(|\xi|/r)] \leq 1\}$. If ζ is another random variable, then the conditional distribution of ζ given ξ is denoted as $\mathcal{L}(\zeta|\xi)$. If a_n and b_n are sequences of non-negative real numbers, we write $a_n \lesssim b_n$ if there is a constant $c > 0$ not depending on n , and integer $n_0 \geq 1$ such that $a_n \leq cb_n$ for all $n \geq n_0$. Likewise, we write $a_n \asymp b_n$ if $a_n \lesssim b_n$ and $b_n \lesssim a_n$. Lastly, for maxima and minima, we use the notation $a_n \vee b_n = \max\{a_n, b_n\}$ and $a_n \wedge b_n = \min\{a_n, b_n\}$.

2. Main results

Our setup is based on a sequence of models indexed by n , where all parameters may depend on n , unless stated otherwise. In particular, the dimensions $p = p(n)$ and $d = d(n)$ below may vary with n . If a parameter does not depend on n , then it is understood not to depend on p or d either.

Assumption 2.1 (Data-generating model).

- (i). There is a deterministic matrix $A \in \mathbb{R}^{d \times p}$ with $d \geq p$, and i.i.d. random vectors $Z_1, \dots, Z_n \in \mathbb{R}^d$, such that for each $i \in \{1, \dots, n\}$, the observation $X_i \in \mathbb{R}^p$ is generated as

$$X_i = A^\top Z_i. \tag{2.1}$$

- (ii). The random vector Z_1 has independent entries that satisfy $\mathbb{E}[Z_{1j}] = 0$, $\mathbb{E}[Z_{1j}^2] = 1$, and $\kappa := \mathbb{E}[Z_{1j}^4] > 1$ for all $j \in \{1, \dots, d\}$, where κ does not depend on n . In addition, the condition $\max_{1 \leq j \leq d} \|Z_{1j}\|_{\psi_2} \lesssim 1$ holds.
- (iii). There is a constant $\beta > 1/2$ not depending on n , such that for each $j \in \{1, \dots, p\}$, the singular value $\sigma_j(A)$ satisfies

$$\sigma_j(A) \asymp j^{-\beta}.$$

Remarks. In statistical applications, the matrix A is typically taken to be the square root $\Sigma^{1/2}$, with $p = d$. However, the extra generality of a rectangular matrix is needed for the application of our work to sketching algorithms in Section 4. To comment on two other aspects of Assumption 2.1, observe that it places no constraints on the relationship between n and p , and it allows for many eigenvalues of Σ to be repeated.

In order to state our main result, we need to precisely define the statistic T_n^* that arises from bootstrap sampling. Let (X_1^*, \dots, X_n^*) be drawn with replacement from (X_1, \dots, X_n) , and define the matrix

$$\widehat{\Sigma}^* = \frac{1}{n} \sum_{i=1}^n X_i^* (X_i^*)^\top.$$

Then, the bootstrapped counterpart of T_n is defined as

$$T_n^* = \sqrt{n} \|\widehat{\Sigma}^* - \widehat{\Sigma}\|_{\text{op}}.$$

The following is our main result.

Theorem 2.1. *Suppose that Assumption 2.1 holds. Then, there is a constant $c > 0$ not depending on n such that the event*

$$d_K(\mathcal{L}(T_n), \mathcal{L}(T_n^*|X)) \leq c n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c \tag{2.2}$$

occurs with probability at least $1 - \frac{c}{n}$.

Remarks. A high-level outline of the proof will be given in Section 2.3. To explain how the difference $\beta - 1/2$ arises in the rate of bootstrap approximation, we offer some informal discussion. As preparatory notation, define the ellipsoidal boundary set $\mathcal{E} = \{Au \mid u \in \mathbb{S}^{p-1}\}$, as well as its signed version $\Theta = \mathcal{E} \times \{\pm 1\}$, whose generic element is denoted by $\theta = (v, s)$. With these items in place, we will consider the following empirical process indexed by Θ ,

$$\mathbb{G}_n(\theta) = \frac{s}{\sqrt{n}} \sum_{i=1}^n \langle v, Z_i \rangle^2 - \mathbb{E}[\langle v, Z_i \rangle^2], \tag{2.3}$$

which allows T_n to be represented as

$$T_n = \sup_{\theta \in \Theta} \mathbb{G}_n(\theta). \tag{2.4}$$

Given that the set Θ is uncountable, a standard reduction is to approximate T_n with the supremum of \mathbb{G}_n over a discrete ϵ -net for Θ , where the metric is taken to be $\rho(\theta, \tilde{\theta}) = \|v - \tilde{v}\|_2 + |s - \tilde{s}|$. In turn, this requires us to control the discretization error, which leads to bounding the supremum of increments, denoted

$$\Delta_n(\epsilon) = \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})|.$$

In order for the discrete approximation to succeed, the quantity $\mathbb{E}[\Delta_n(\epsilon)]$ should vanish as $\epsilon \rightarrow 0$. However, the demonstration of this property depends on the complexity of Θ through the parameter β .

We can gain some intuition for the role of β by looking at how it affects $\mathbb{E}[\Delta_n(\epsilon)]$ in a much simpler case—where \mathbb{G}_n is replaced by a linear Gaussian process indexed by Θ . Namely, consider the process $\tilde{\mathbb{G}}_n(\theta) = n^{-1/2} \sum_{i=1}^n s \langle v, \zeta_i \rangle$, where ζ_1, \dots, ζ_n are independent standard Gaussian vectors. In this case, if $\tilde{\Delta}_n(\epsilon)$ denotes the analogue of $\Delta_n(\epsilon)$ for $\tilde{\mathbb{G}}_n$ and $p \gtrsim \epsilon^{-1/\beta}$, then the following lower bound can be shown using classical facts about Gaussian processes,

$$\mathbb{E}[\tilde{\Delta}_n(\epsilon)] \gtrsim \epsilon^{(\beta-1/2)/\beta}. \tag{2.5}$$

(See Lemma F.3 in Appendix F.) Thus, the main point to take away here is that even in the simpler case of a linear Gaussian process, the condition $\beta > 1/2$ is necessary in high dimensions for the discretization error to vanish as $\epsilon \rightarrow 0$.

Another benefit of looking at the linear Gaussian case is that the lower bound (2.5) provides a reference point for assessing our upper bound on the discretization error. For instance, it will follow from Proposition 2.1 that

$$\mathbb{E}[\Delta_n(\epsilon)] \lesssim \epsilon^{(\beta-1/2)/\beta} \log(n), \tag{2.6}$$

as shown in Section 2.2. Hence, in light of the quadratic nature of the process \mathbb{G}_n , it is notable that the dependence on ϵ does not change in comparison to the linear Gaussian case. Moreover, it also turns out that the dependence on ϵ even remains the same for L_q norms of $\Delta_n(\epsilon)$ when q is large.

One more point of theoretical interest is that the bound (2.6) arises in a situation where standard chaining seems to give a slower dependence on ϵ than a more problem-specific approach. As an example of a standard approach, one might try to show that \mathbb{G}_n is sub-exponential with respect to ρ , and then appeal to an entropy integral bound such as in (van der Vaart and Wellner, 2000, Theorem 2.2.4). However, this ultimately leads to an upper bound scaling like $\epsilon^{(\beta-1)/\beta}$, which would require the excessive condition $\beta > 1$ (as opposed to $\beta > 1/2$). Likewise, the development of new techniques for quadratic processes akin to \mathbb{G}_n has attracted interest in the literature, as surveyed in (Talagrand, 2014, §9.3-9.4). Nevertheless, it should also be noted that existing results in this direction do not seem to be directly applicable to our analysis of the bootstrap. For instance, the abstract approaches based on Talagrand’s γ_1 and γ_2 functionals lead to challenges in connection with the bootstrap, because the discrete process \mathbb{G}_n^* (arising from sampling with replacement) induces a random metric on Θ that does not lend itself to calculations. On the other hand, the approach taken here allows \mathbb{G}_n and \mathbb{G}_n^* to be treated on nearly equal footing (see Proposition D.1).

2.1. A general-purpose bound for sample covariance matrices

Below, we provide a general-purpose high-probability bound for the operator norm error of sample covariance matrices. Notably, the result only requires control on the moments of the norm of a random vector. Since the result may be of independent interest in other problems, we have stated it in the general context of a separable Hilbert space \mathcal{H} , with inner product and norm denoted as $\langle \cdot, \cdot \rangle$ and $\|\cdot\|$. In addition, if $x, y, z \in \mathcal{H}$, then $x \otimes y$ denotes the linear operator from \mathcal{H} to \mathcal{H} satisfying $(x \otimes y)z = \langle y, z \rangle x$.

Theorem 2.2. Let $q \geq 3$, and let $\xi_1, \dots, \xi_n \in \mathcal{H}$ be i.i.d. random elements. Also, let

$$r(q) = \frac{q \left(\mathbb{E}[\|\xi_1\|^{2q}] \right)^{\frac{1}{q}}}{\|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}}}. \tag{2.7}$$

Then, there is an absolute constant $c > 0$ such that

$$\left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \otimes \xi_i - \mathbb{E}[\xi_i \otimes \xi_i] \right\|_{\text{op}}^q \right)^{1/q} \leq c \|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}} \left(\sqrt{\frac{r(q)}{n^{1-3/q}}} \vee \frac{r(q)}{n^{1-3/q}} \right). \tag{2.8}$$

Remarks. It will sometimes be useful to consider the special case where the random variable $\|\xi_1\|$ can be described in terms of its ψ_2 -Orlicz norm. This gives

$$r(q) \leq \frac{c q^2 \|\|\xi_1\|\|_{\psi_2}^2}{\|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}}}, \tag{2.9}$$

for an absolute constant $c > 0$, which can be obtained from the facts about Orlicz norms summarized in Lemmas G.1 and G.2. Before proceeding directly to the proof of Theorem 2.2, we need a preparatory lemma, which is a slightly relaxed version of a result from (Rudelson, 1999, p.63). As a matter of notation, the Schatten- q norm of an operator M will be denoted as $\|M\|_{S_q}$.

Lemma 2.1. Let $x_1, \dots, x_n \in \mathcal{H}$ be fixed, and let $\varepsilon_1, \dots, \varepsilon_n$ be independent Rademacher random variables. Then, there is an absolute constant $c > 0$ such that for any $q \geq 2$,

$$\left(\mathbb{E} \left\| \sum_{i=1}^n \varepsilon_i x_i \otimes x_i \right\|_{S_q}^q \right)^{1/q} \leq c \cdot n^{1/q} \cdot \sqrt{q} \cdot \left(\max_{1 \leq i \leq n} \|x_i\| \right) \cdot \left\| \sum_{i=1}^n x_i \otimes x_i \right\|_{\text{op}}^{1/2}. \tag{2.10}$$

Proof. First, we make use of a non-commutative Khinchine inequality originating from (Lust-Piquard, 1986). Specifically, we use the version from (Pisier, 2016, Theorem 14.6), which implies

$$\left(\mathbb{E} \left\| \sum_{i=1}^n \varepsilon_i x_i \otimes x_i \right\|_{S_q}^q \right)^{1/q} \leq c \sqrt{q} \left\| \left(\sum_{i=1}^n \|x_i\|^2 x_i \otimes x_i \right)^{1/2} \right\|_{S_q}.$$

Next, observe that any operator M with rank at most r satisfies $\|M\|_{S_q} \leq r^{1/q} \|M\|_{\text{op}}$, and that the operator $\sum_{i=1}^n \|x_i\|_2^2 x_i \otimes x_i$ has rank at most n . Combining this with the previous bound leads to the stated result. \square

Proof of Theorem 2.2. The proof extends the approach developed in (Rudelson and Vershynin, 2007) to the case of unbounded random vectors. Using a standard symmetrization argument, we have

$$\left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \otimes \xi_i - \mathbb{E}[\xi_i \otimes \xi_i] \right\|_{\text{op}}^q \right)^{1/q} \leq c \left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \varepsilon_i \xi_i \otimes \xi_i \right\|_{\text{op}}^q \right)^{1/q}, \tag{2.11}$$

where $\varepsilon_1, \dots, \varepsilon_n$ are independent Rademacher variables that are also independent of ξ_1, \dots, ξ_n . Next, since $\|\cdot\|_{\text{op}} \leq \|\cdot\|_{S_q}$, it follows from Lemma 2.1 and the Cauchy-Schwarz inequality that

$$\left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \varepsilon_i \xi_i \otimes \xi_i \right\|_{\text{op}}^q \right)^{1/q} \leq c \cdot n^{1/q} \cdot \sqrt{\frac{q}{n}} \cdot \mathbb{E} \left[\max_{1 \leq i \leq n} \|\xi_i\|^{2q} \right]^{\frac{1}{2q}} \cdot \left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \otimes \xi_i \right\|_{\text{op}}^q \right)^{\frac{1}{2q}}. \tag{2.12}$$

To bound the L_{2q} norm of $\max_{1 \leq i \leq n} \|\xi_i\|$, we have

$$\mathbb{E} \left[\max_{1 \leq i \leq n} \|\xi_i\|^{2q} \right]^{\frac{1}{2q}} \leq n^{\frac{1}{2q}} \left(\mathbb{E}[\|\xi_1\|^{2q}] \right)^{\frac{1}{2q}}.$$

Also, to handle the last factor in the bound (2.12), we use

$$\left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \otimes \xi_i \right\|_{\text{op}}^q \right)^{1/q} \leq \left(\mathbb{E} \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \otimes \xi_i - \mathbb{E}[\xi_i \otimes \xi_i] \right\|_{\text{op}}^q \right)^{1/q} + \|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}}.$$

Hence, if L denotes the left side of (2.11), then

$$L \leq c \cdot n^{\frac{3}{2q}} \cdot \sqrt{\frac{q}{n}} \cdot \left(\mathbb{E}[\|\xi_1\|^{2q}] \right)^{\frac{1}{2q}} \cdot \sqrt{L + \|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}}}.$$

Finally, by putting

$$K = c \cdot n^{\frac{3}{2q}} \cdot \sqrt{\frac{q}{n}} \cdot \left(\mathbb{E}[\|\xi_1\|^{2q}] \right)^{\frac{1}{2q}} \cdot \|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}}^{1/2} \quad \text{and} \quad K' = 1/\|\mathbb{E}[\xi_1 \otimes \xi_1]\|_{\text{op}}$$

we may solve the quadratic inequality $L \leq K\sqrt{K'L + 1}$ with respect to L to reach the stated result. \square

2.2. An increment bound for \mathbb{G}_n

The next result plays a key role in the proof of Theorem 2.1, since it serves to control the discretization error that arises when approximating T_n with a supremum over an ϵ -net for Θ . (Recall that for any $\theta = (v, s)$ and $\tilde{\theta} = (\tilde{v}, \tilde{s})$ in Θ , we use the metric $\rho(\theta, \tilde{\theta}) = \|v - \tilde{v}\|_2 + |s - \tilde{s}|$.)

Proposition 2.1. *Let $q = \log(n) \vee 3$, let $\epsilon = \epsilon(n) \in (0, 1)$, and suppose that Assumption 2.1 holds. Then,*

$$\left\| \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| \right\|_q \lesssim \epsilon^{1 - \frac{1}{2\beta}} \log(n). \tag{2.13}$$

Proof. As an initial observation, note that the condition $\rho(\theta, \tilde{\theta}) \leq \epsilon < 1$ implies that the signs s and \tilde{s} must be equal. This leads to the algebraic identity

$$|\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| = \frac{1}{\sqrt{n}} \left| \sum_{i=1}^n \langle v + \tilde{v}, Z_i \rangle \langle v - \tilde{v}, Z_i \rangle - \mathbb{E}[\langle v + \tilde{v}, Z_i \rangle \langle v - \tilde{v}, Z_i \rangle] \right|.$$

To rewrite the quadratic forms in terms of a symmetric matrix, let

$$Q = \frac{1}{2} \left((v + \tilde{v})(v - \tilde{v})^\top + (v - \tilde{v})(v + \tilde{v})^\top \right),$$

so that

$$|\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| = \frac{1}{\sqrt{n}} \left| \sum_{i=1}^n Z_i^\top Q Z_i - \mathbb{E}[Z_i^\top Q Z_i] \right|.$$

Next, let $t > 0$ denote a free parameter to be chosen later, and define the vectors

$$\omega(t) = \frac{1}{2} \left[t(v + \tilde{v}) + \frac{1}{t}(v - \tilde{v}) \right] \quad \text{and} \quad \tilde{\omega}(t) = \frac{1}{2} \left[t(v + \tilde{v}) - \frac{1}{t}(v - \tilde{v}) \right].$$

In turn, it can be checked that these vectors give the following representation of \mathcal{Q} ,

$$\mathcal{Q} = \omega(t)\omega(t)^\top - \tilde{\omega}(t)\tilde{\omega}(t)^\top, \tag{2.14}$$

which has a certain invariance property, insofar as it holds for every $t > 0$, while \mathcal{Q} itself does not depend on t . The utility of this representation is that it will allow us to work with sums of squares, and also, to optimize with respect to the choice of t .

To proceed, we will define a particular ellipsoid that contains the vectors $\omega(t)$ and $\tilde{\omega}(t)$, and then take a supremum over this ellipsoid to derive a stochastic upper bound on $\sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})|$. For this purpose let $A(\epsilon) \in \mathbb{R}^{d \times p}$ be the matrix with the same s.v.d. as A , except that the singular value $\sigma_j(A)$ is replaced with $\sqrt{2} \min\{\sigma_j(A), \epsilon/2\}$ for every $j \in \{1, \dots, p\}$. Also, define $A(t, \epsilon) \in \mathbb{R}^{d \times 2p}$ as the column concatenation

$$A(t, \epsilon) = \left[tA, \frac{1}{t}A(\epsilon) \right].$$

With this matrix in hand, it can be shown that both vectors $\omega(t)$ and $\tilde{\omega}(t)$ lie in the ellipsoid $A(t, \epsilon)(\mathbb{B}_{2p}(2))$, where $\mathbb{B}_{2p}(2)$ denotes the ℓ_2 -ball of radius 2 in \mathbb{R}^{2p} . (For the details, see Lemma F.2.) In particular, this ellipsoid does not depend on the indices θ and $\tilde{\theta}$ underlying $\omega(t)$ and $\tilde{\omega}(t)$. As a result, we have

$$\sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| \leq \sup_{w \in \mathbb{B}_{2p}(2)} \frac{2}{\sqrt{n}} \left| \sum_{i=1}^n \langle Z_i, A(t, \epsilon)w \rangle^2 - \mathbb{E}[\langle Z_i, A(t, \epsilon)w \rangle^2] \right|.$$

We now apply Theorem 2.2 with $\xi_i = A(t, \epsilon)^\top Z_i$, which gives

$$\left\| \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| \right\|_q \lesssim \sqrt{n} \cdot \|A(t, \epsilon)\|_{\text{op}}^2 \cdot \left(\sqrt{\frac{r(q)}{n^{1-3/q}}} \vee \frac{r(q)}{n^{1-3/q}} \right). \tag{2.15}$$

Due to the choice $q = \log(n) \vee 3$, we have $n^{1-3/q} \asymp n$, and also, the bound (2.9) implies

$$r(q) \lesssim q^2 \frac{\|A(t, \epsilon)^\top Z_1\|_2^2}{\|A(t, \epsilon)\|_{\text{op}}^2}.$$

Furthermore, Lemma G.2 gives

$$\| \|A(t, \epsilon)^\top Z_1\|_2 \|_{\psi_2} \lesssim \|A(t, \epsilon)\|_F, \tag{2.16}$$

and then combining with (2.15) leads to

$$\left\| \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| \right\|_q \lesssim \left(q \cdot \|A(t, \epsilon)\|_{\text{op}} \cdot \|A(t, \epsilon)\|_F \right) \sqrt{\left(\frac{q^2}{\sqrt{n}} \|A(t, \epsilon)\|_F^2 \right)}.$$

Hence, to complete the proof, it remains to bound the norms of $A(t, \epsilon)$ and then specify a value of t . From the definition of $A(t, \epsilon)$ and a short calculation, we have

$$\|A(t, \epsilon)\|_F \leq t\|A\|_F + \frac{1}{t}\|A(\epsilon)\|_F \lesssim t + \frac{1}{t}\epsilon^{1-\frac{1}{2\beta}}, \tag{2.17}$$

as well as

$$\|A(t, \epsilon)\|_{\text{op}} \leq t\|A\|_{\text{op}} + \frac{1}{t}\|A(\epsilon)\|_{\text{op}} \lesssim t + \frac{1}{t}\epsilon. \tag{2.18}$$

Taking $t = \epsilon^{\frac{1}{2} - \frac{1}{4\beta}}$ leads to the stated result. □

2.3. Outline for the proof of Theorem 2.1

In this subsection, we define several objects that will recur in our arguments, and then explain how the main components of the proof fit together.

Bootstrap and Gaussian processes. Let (Z_1^*, \dots, Z_n^*) be sampled with replacement from (Z_1, \dots, Z_n) , and define the bootstrap counterpart of \mathbb{G}_n as

$$\mathbb{G}_n^*(\theta) = \frac{s}{\sqrt{n}} \sum_{i=1}^n \langle v, Z_i^* \rangle^2 - \mathbb{E}[\langle v, Z_i^* \rangle^2 | X],$$

where $\mathbb{E}[\cdot | X]$ refers to expectation that is conditional on X_1, \dots, X_n , and we note that $\mathbb{E}[\langle v, Z_i^* \rangle^2 | X] = \frac{1}{n} \sum_{i=1}^n \langle v, Z_i \rangle^2$. This definition of \mathbb{G}_n^* allows T_n^* to be expressed as

$$T_n^* = \sup_{\theta \in \Theta} \mathbb{G}_n^*(\theta).$$

In addition, we define G_n as the centered Gaussian process on Θ whose covariance structure matches that of \mathbb{G}_n ,

$$\text{cov}(G_n(\theta), G_n(\tilde{\theta})) = \text{cov}(\mathbb{G}_n(\theta), \mathbb{G}_n(\tilde{\theta})) \quad \text{for all } \theta, \tilde{\theta} \in \Theta.$$

Subsets of indices. In order to define some special subsets of \mathcal{E} and Θ , let

$$\ell_n = \left[(1 \vee \log(n)^3) \wedge p \right] \quad \text{and} \quad k_n = \left[\left(\ell_n \vee \log(n)^{\frac{6\beta+4}{\beta-1/2}} \right) \wedge p \right],$$

which always satisfy $1 \leq \ell_n \leq k_n \leq p$. Also, let the columns of $V_{k_n} \in \mathbb{R}^{p \times k_n}$ contain the leading k_n right singular vectors of A . Based on these items, we define \mathcal{E}_n^\uparrow as a subset of \mathcal{E} arising from vectors in \mathbb{S}^{p-1} that are “partially aligned” with the columns of V_{k_n} ,

$$\mathcal{E}_n^\uparrow = \left\{ Au \mid u \in \mathbb{S}^{p-1} \text{ and } \|V_{k_n}^\top u\|_2 > \frac{1}{2} k_n^{-\beta+1/2} \right\}.$$

Likewise, by analogy with the definition of Θ , let

$$\Theta_n^\uparrow = \mathcal{E}_n^\uparrow \times \{\pm 1\}.$$

The next piece of notation is an ϵ -net for Θ_n^\uparrow with respect to the metric ρ . This net is denoted as $\Theta_n^\uparrow(\epsilon) \subset \Theta_n^\uparrow$ and has the defining property that for any $\theta \in \Theta_n^\uparrow$, there is at least one point $\theta' \in \Theta_n^\uparrow(\epsilon)$ with $\rho(\theta, \theta') \leq \epsilon$. Throughout the proofs, we will mostly use the particular choice $\epsilon = \epsilon_n$ with

$$\epsilon_n = n^{-\beta/(6\beta+4)}. \tag{2.19}$$

Lastly, due to classical bounds on the metric entropy of ellipsoids (as recorded in Lemma G.3), it is possible to choose an ϵ_n -net $\Theta_n^\uparrow(\epsilon_n)$ with respect to ρ so that its cardinality satisfies $\log \text{card}(\Theta_n^\uparrow(\epsilon_n)) \lesssim \epsilon_n^{-1/\beta}$.

Decomposition into six main terms. We will bound the Kolmogorov distance between $\mathcal{L}(T_n)$ and $\mathcal{L}(T_n^*|X)$ with six terms,

$$d_K(\mathcal{L}(T_n), \mathcal{L}(T_n^*|X)) \leq \mathbf{I}_n + \mathbf{II}_n + \mathbf{III}_n + \widetilde{\mathbf{III}}_n + \widetilde{\mathbf{II}}_n + \widetilde{\mathbf{I}}_n,$$

which are defined below. The essential novelty of the proof deals with the four terms $(\mathbf{I}_n, \mathbf{II}_n, \widetilde{\mathbf{III}}_n, \widetilde{\mathbf{I}}_n)$, and almost all of the effort will be focused on these.

1. Localizing the maximizer of \mathbb{G}_n :

$$\mathbf{I}_n = d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta} \mathbb{G}_n(\theta)\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow} \mathbb{G}_n(\theta)\right)\right).$$

(We use the phrase ‘localizing the maximizer of \mathbb{G}_n ’, because the problem of showing that \mathbf{I}_n is small amounts to showing that the maximizing index for \mathbb{G}_n is likely to fall in Θ_n^\uparrow .)

2. Discrete approximation of \mathbb{G}_n :

$$\mathbf{II}_n = d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow} \mathbb{G}_n(\theta)\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow(\epsilon_n)} \mathbb{G}_n(\theta)\right)\right)$$

3. Gaussian approximation:

$$\mathbf{III}_n = d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow(\epsilon_n)} \mathbb{G}_n(\theta)\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow(\epsilon_n)} G_n(\theta)\right)\right)$$

4. Bootstrap approximation:

$$\widetilde{\mathbf{III}}_n = d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow(\epsilon_n)} G_n(\theta)\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow(\epsilon_n)} \mathbb{G}_n^*(\theta)|X\right)\right)$$

5. Discrete approximation of \mathbb{G}_n^* :

$$\widetilde{\mathbf{II}}_n = d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow(\epsilon_n)} \mathbb{G}_n^*(\theta)|X\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow} \mathbb{G}_n^*(\theta)|X\right)\right)$$

6. Localizing the maximizer of \mathbb{G}_n^* :

$$\widetilde{\mathbf{I}}_n = d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\uparrow} \mathbb{G}_n^*(\theta)|X\right), \mathcal{L}\left(\sup_{\theta \in \Theta} \mathbb{G}_n^*(\theta)|X\right)\right)$$

Altogether, the six terms are handled consecutively in Appendices A through E, with each appendix corresponding to a different term (except for \mathbf{III}_n and $\widetilde{\mathbf{III}}_n$, which are handled together).

To comment on some of the techniques used in the various stages of the proof, the bound on \mathbf{I}_n builds on a lower-tail bound for Gaussian maxima developed in (Lopes, Lin and Müller, 2020) and requires fine-grained control on the covariances $\text{cov}(\mathbb{G}_n(\theta), \mathbb{G}_n(\tilde{\theta}))$ for certain choices of θ and $\tilde{\theta}$. Next, the analysis of \mathbf{II}_n uses both the result on sample covariance matrices in Theorem 2.2 and the increment bound in Proposition 2.1. With regard to the terms \mathbf{III}_n and $\widetilde{\mathbf{III}}_n$, the localization and discrete approximation steps make it possible to invoke Gaussian and bootstrap approximation results from (Chernozhukov, Chetverikov and Kato, 2017). Lastly, the terms $\widetilde{\mathbf{II}}_n$ and $\widetilde{\mathbf{I}}_n$ are analyzed in correspondence with \mathbf{II}_n and \mathbf{I}_n , and here the general-purpose nature of Theorem 2.2 is especially helpful, since it allows the bootstrapped process \mathbb{G}_n^* to be treated in a similar manner to the original process \mathbb{G}_n .

3. Application to inference on covariance matrices

To illustrate the numerical performance of the bootstrap, this section looks at the coverage probabilities of bootstrap confidence regions for Σ . An additional set of numerical results dealing with simultaneous confidence intervals for the eigenvalues of Σ are presented in Appendix H of the supplementary material. It should also be noted that both sets of numerical results were obtained in a situation where the leading eigenvalue $\lambda_1(\Sigma)$ has high multiplicity.

Simulation settings. Simulations were based on the model described in Assumption 2.1, with $n \in \{300, 500, 700\}$ and $d = p = 1,000$, giving $n < p$ in every case. Also, the matrix A was constructed to be symmetric so that it can be interpreted as $A = \Sigma^{1/2}$. To specify A in more detail, its singular values (equivalently eigenvalues) were chosen as

$$\sigma_1(A) = \cdots = \sigma_5(A) = 1 \quad \text{and} \quad \sigma_j(A) = j^{-\beta} \quad \text{for } j \in \{6, \dots, p\},$$

with decay parameter values $\beta \in \{0.25, 0.50, 0.75, 1.0, 1.25\}$, and its eigenvectors were taken as the orthogonal factor from a QR decomposition of a $p \times p$ matrix with independent $N(0, 1)$ entries. Next, for each pair (n, β) , we conducted 5,000 trials in which the $n \times p$ data matrix $X = ZA$ was generated by filling $Z \in \mathbb{R}^{n \times p}$ with independent random variables drawn from $N(0, 1)$ or a standardized t_{20} distribution. Lastly, for each trial, we generated 500 bootstrap samples T_n^* by sampling the rows of X with replacement, as described in Section 2.

Error estimation and confidence regions. A natural way to formulate the problem of error estimation for $\widehat{\Sigma}$ is in terms of the $1 - \alpha$ quantile of T_n , denoted by $q_{1-\alpha}$. By definition, this quantity gives the tightest bound of the form

$$\|\widehat{\Sigma} - \Sigma\|_{\text{op}} \leq \frac{q_{1-\alpha}}{\sqrt{n}}$$

that holds with probability at least $1 - \alpha$. Likewise, if we let $\widehat{q}_{1-\alpha}$ denote the empirical $(1 - \alpha)$ -quantile of the bootstrap samples T_n^* , then we may regard $\widehat{q}_{1-\alpha}/\sqrt{n}$ as an error estimate for $\widehat{\Sigma}$.

Alternatively, the estimate $\widehat{q}_{1-\alpha}$ can be viewed as specifying an approximate $(1 - \alpha)$ -confidence region for Σ . That is, if we let $\mathbf{B}_{\text{op}}(r; \widehat{\Sigma}) \subset \mathbb{R}^{p \times p}$ denote the operator-norm ball of radius $r > 0$ centered at $\widehat{\Sigma}$, then $q_{1-\alpha}/\sqrt{n}$ is the smallest value of r such that

$$\mathbb{P}\left(\Sigma \in \mathbf{B}_{\text{op}}(r; \widehat{\Sigma})\right) \geq 1 - \alpha.$$

Hence, the ideal confidence region may be approximated with $\mathbf{B}_{\text{op}}\left(\frac{\widehat{q}_{1-\alpha}}{\sqrt{n}}; \widehat{\Sigma}\right)$.

The observed coverage probabilities based on $\widehat{q}_{1-\alpha}$ have been listed in Table 1, with the nominal coverage level always being set to 90%. (Note that these probabilities can be interpreted either with respect to the coverage of the error bound or the confidence region.) Looking at the results, we see that they generally conform with our theoretical analysis, in the sense that the performance of the bootstrap depends substantially on whether β is above or below 0.50. Namely, for $\beta > 0.50$ and sufficiently large sample sizes, the observed coverage comes within a percent or two of the desired level. Furthermore, this holds in spite of the fact that $\lambda_1(\Sigma)$ has high multiplicity. On the other hand, when $\beta < 0.50$, the bootstrap breaks down even at large sample sizes, with the observed coverage of 100% being far higher than the desired 90%. Apart from this dichotomy, it should be noted that the bootstrap has the favorable property that it errs reliably in the conservative direction, with the observed coverage never falling below the nominal level. More generally, we observed similar patterns in the context of simultaneous confidence intervals for the eigenvalues of Σ , as discussed in Appendix H.

Table 1. Observed coverage probabilities for $\widehat{q}_{1-\alpha}$, with $\alpha = 0.1$ and $p = 1,000$.

decay param. β	sample size n			decay param. β	sample size n		
	300	500	700		300	500	700
0.25	100%	100%	100%	0.25	100%	100%	100%
0.50	94.3%	92.0%	93.6%	0.50	95.4%	93.0%	92.2%
0.75	92.8%	92.3%	92.0%	0.75	92.9%	92.0%	91.9%
1.00	92.7%	91.7%	91.2%	1.00	92.5%	91.8%	91.6%
1.25	92.4%	91.5%	91.2%	1.25	92.5%	91.7%	91.5%

(a) $N(0, 1)$ distribution

(b) standardized t_{20} distribution

4. Application to randomized numerical linear algebra

Over the past decade, RandNLA has become the focus of intense activity in many fields related to large-scale computation (Drineas and Mahoney, 2018, Halko, Martinsson and Tropp, 2011, Kannan and Vempala, 2017, Mahoney, 2011, Woodruff, 2014). Broadly speaking, this new direction of research has stemmed from the principle that randomization is a very general mechanism for scaling up algorithms. However, in exchange for scalability, randomized sketching algorithms are typically less accurate than their deterministic predecessors. Therefore, in order to use sketching reliably, it is crucial to verify that the algorithmic error is small, which motivates new applications of the bootstrap beyond its traditional domains.

The purpose of this section is to illustrate how the bootstrap can be applied to estimate operator-norm error for randomized matrix multiplication, which has been a prominent topic in the RandNLA literature (e.g., Cohen, Nelson and Woodruff, 2016, Drineas and Kannan, 2001, Drineas, Kannan and Mahoney, 2006, Gupta et al., 2018, Holodnuk and Ipsen, 2015, Pagh, 2013). A related study of the bootstrap for this application can also be found in (Lopes, Wang and Mahoney, 2019), which differs from the current work insofar as it deals exclusively with the entrywise ℓ_∞ -norm and does not focus the role of spectrum decay.

To proceed, we will first provide a brief review of the algorithmic setting (Section 4.1), followed by an efficient implementation of the bootstrap using an extrapolation technique (Section 4.2). In turn, we will present numerical results for synthetic matrices (Section 4.3), as well as an example concerning spatial modes of temperature variation (Section 4.4).

4.1. Rudiments of sketching

Consider a situation involving a very large deterministic matrix $A \in \mathbb{R}^{d \times p}$ with $p \ll d$, where the product $A^T A$ is too expensive to compute to high precision. For instance, this often occurs when A must be stored on disk because it exceeds the constraints of fast memory. Alternatively, even when memory is not a bottleneck, matrix products can become too expensive if they must be computed frequently as a subroutine of a larger pipeline.

The basic idea of sketching is to work with a shorter version of A , referred to as a “sketch of A ”, and denoted as $\tilde{A} \in \mathbb{R}^{n \times p}$ where $n \ll d$. This matrix is defined as

$$\tilde{A} = SA,$$

where $S \in \mathbb{R}^{n \times d}$ is a random “sketching matrix” that is generated by the user. In particular, the user must choose the “sketch size” n . Intuitively, the matrix S is intended to shorten A in a way that retains most of the information, so that the inexpensive product $\tilde{A}^T \tilde{A}$ will provide a good approximation to $A^T A$.

The sketching matrix. Typically, the action of S upon A is interpreted in either of two ways: randomly projecting columns from \mathbb{R}^d into \mathbb{R}^n , or discretely sampling n among d rows. In addition, the matrix S is commonly generated by the user so that its rows are i.i.d., and that it satisfies $\mathbb{E}[S^\top S] = I_n$, which implies that $\tilde{A}^\top \tilde{A}$ is unbiased with respect to $A^\top A$. At a high level, these basic properties are sufficient to understand all of our work below, but numerous types of sketching matrices have been studied in the literature. For instance, two of the most well-known are the *Gaussian random projection* and *uniform row sampling* types, where the former has i.i.d. rows drawn from $N(0, \frac{1}{n}I_d)$, and the latter has i.i.d. rows drawn uniformly from $\{\sqrt{d/n}e_1, \dots, \sqrt{d/n}e_d\}$. More elaborate examples may be found in the references above.

Cost versus accuracy. Whenever sketching is implemented, the choice of the sketch size n plays a pivotal role in a tradeoff between computational cost and accuracy. To see this, note that on one hand, the cost to compute $\tilde{A}^\top \tilde{A}$ is generally proportional to n , with the number of operations being $O(np^2)$. On the other hand, the operator-norm error of $\tilde{A}^\top \tilde{A}$ tends to decrease stochastically like $1/\sqrt{n}$, because the difference $\tilde{A}^\top \tilde{A} - A^\top A$ can be expressed as a sample average of n centered random matrices (i.e., in the same way as $\widehat{\Sigma} - \Sigma$).

The need for error estimation. Although the choice of the sketch size n has clear importance, this choice also involves practical difficulties that expose a major gap between the theory and practice of sketching. Specifically, these difficulties arise because the value of the sketching error $\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}}$ is *unknown in practice*, as it depends on the unknown product $A^\top A$. Hence, it is hard for the user to know if any given choice of n will achieve a desired level of accuracy.

As a way to handle this dilemma, one option is to consult the RandNLA literature on theoretical error bounds for $\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}}$, as surveyed in the references above. However, much like in the setting of covariance estimation, these results usually only provide qualitative guidance, and they rarely offer an explicit numerical bound. Most often, this occurs because of unspecified theoretical constants, but there is also a second key limitation: Theoretical error bounds are generally formulated to hold in a worst-case sense, and so they often fail to account for special structure. Due to these issues, we propose instead to directly estimate the error via a computationally efficient bootstrap method. This has the twofold benefit of providing a numerical bound and adapting automatically to the structure of the problem at hand.

Comparison of sketching and covariance estimation. To clarify the relationship between the sketching error $\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}}$ and the covariance estimation error $\|\widehat{\Sigma} - \Sigma\|_{\text{op}}$, let A , Σ , and $\widehat{\Sigma}$ be understood as in the context of the model 2.1 with $Z \in \mathbb{R}^{n \times d}$ having rows Z_1, \dots, Z_n , and let

$$S = \frac{1}{\sqrt{n}} Z.$$

Under these conditions, the matrix S has the desired properties of a sketching matrix mentioned earlier, and furthermore $\tilde{A}^\top \tilde{A} - A^\top A = \widehat{\Sigma} - \Sigma$. However, it is worth highlighting that this formal similarity conceals some operational differences. For instance, the matrices Z and A are unobservable to the user in covariance estimation, whereas the user does have access to S and A in sketching. Secondly, in covariance estimation, the user often does not have the option to increase n , but in sketching, it is possible to construct a rough initial sketch of A for inspection, and then take a second sketch to improve performance. Later on, we will show how this second point has an important link with our error estimation method, because it will enable the user to *dynamically predict* the total sketch size needed to reach a given level of accuracy.

4.2. Error estimation with an extrapolated bootstrap

The intuition for applying the bootstrap to sketching comes from thinking of the matrix \tilde{A} as a “dataset” whose rows are “observations”. In particular, this interpretation is supported by the fact that many types of sketching matrices S cause the rows of \tilde{A} to be i.i.d. Therefore, we may expect that sampling from the rows of \tilde{A} with replacement will faithfully mimic the process that generated \tilde{A} .

To fix some notation for describing the bootstrap method, let $q_{1-\alpha}$ denote the $(1 - \alpha)$ -quantile of the sketching error variable $\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}}$, which is the minimal value such that the event

$$\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}} \leq q_{1-\alpha}$$

holds with probability at least $1 - \alpha$. Our main goal is to construct an estimate $\hat{q}_{1-\alpha}$ using only the sketch \tilde{A} as a source of information. Below, we state a basic version of the bootstrap method in Algorithm 1, which will later be accelerated via an extrapolation technique in Section 4.2.1.

Algorithm 1. (Bootstrap estimate of sketching error).

Input: The number of bootstrap samples B , and the sketch $\tilde{A} \in \mathbb{R}^{n \times p}$.

For $b = 1, \dots, B$ **do**

1. Form $\tilde{A}^* \in \mathbb{R}^{n \times p}$ by drawing n rows from \tilde{A} with replacement.
2. Compute the bootstrap sample $\varepsilon_b^* := \|(\tilde{A}^*)^\top (\tilde{A}^*) - \tilde{A}^\top \tilde{A}\|_{\text{op}}$.

Return: $\hat{q}_{1-\alpha} \leftarrow$ the $(1 - \alpha)$ -quantile of the values $\varepsilon_1^*, \dots, \varepsilon_B^*$.

Remark. Given that the construction of \tilde{A} is fully controlled by the user, one might ask why bootstrapping is preferable to carrying out many repetitions of the actual sketching process. The answer comes down to the fact that constructing \tilde{A} requires a computation involving the full matrix A , which often incurs high communication costs. In fact, this issue is one of the primary motivations for the whole subject of RandNLA, which usually deals with situations where it is only feasible to access A at most a handful of times. In contrast to the task of constructing \tilde{A} from A , Algorithm 1 only requires inexpensive access to the much smaller matrix \tilde{A} , and it requires no access to A whatsoever.

4.2.1. Extrapolation

Because the user has the option to increase the sketch size n by performing an extra round of sketching, it becomes possible to accelerate the bootstrap with an extrapolation technique that is often not applicable in covariance estimation. To develop the idea, we should first recall that the fluctuations of $\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}}$ tend to scale like $1/\sqrt{n}$ as a function of n , because the difference $\tilde{A}^\top \tilde{A} - A^\top A$ can be written as a centered sample average of n random matrices. Therefore, if we view the sketching error quantile as a function of n , say $q_{1-\alpha} = q_{1-\alpha}(n)$, then we may expect the following approximate relationship between a small “initial” sketch size n_0 , and a larger “final” sketch size n_1 ,

$$q_{1-\alpha}(n_1) \approx \sqrt{\frac{n_0}{n_1}} q_{1-\alpha}(n_0). \tag{4.1}$$

The significance of this approximation is that $q_{1-\alpha}(n_0)$ is computationally much easier to estimate than $q_{1-\alpha}(n_1)$, since the former involves bootstrapping a matrix of size $n_0 \times p$, rather than $n_1 \times p$. More general background on the connections between extrapolation and resampling methods can be found in (Bertail, 1997, Bertail and Politis, 2001, Bickel and Sakov, 2002, Bickel and Yahav, 1988, Lopes, 2019), among others.

Based on the heuristic approximation (4.1), we can obtain an inexpensive estimate of $\mathfrak{q}_{1-\alpha}(n_1)$ for any $n_1 > n_0$ by using

$$\widehat{\mathfrak{q}}_{1-\alpha}^{\text{ext}}(n_1) := \sqrt{\frac{n_0}{n_1}} \widehat{\mathfrak{q}}_{1-\alpha}(n_0), \quad (4.2)$$

where $\widehat{\mathfrak{q}}_{1-\alpha}(n_0)$ is obtained from Algorithm 1. More concretely, if the user has the ultimate intention of achieving $\|\widehat{A}^\top \tilde{A} - A^\top A\|_{\text{op}} \leq \epsilon_{\text{tol}}$ for some tolerance ϵ_{tol} , then extrapolation may be applied in the following way: First, the user should check the condition $\widehat{\mathfrak{q}}_{1-\alpha}(n_0) \leq \epsilon_{\text{tol}}$ to see if n_0 is already large enough. Second, if n_0 is too small, then the rule (4.2) instructs the user to obtain a final sketch size n_1 satisfying $\widehat{\mathfrak{q}}_{1-\alpha}^{\text{ext}}(n_1) \leq \epsilon_{\text{tol}}$, which is equivalent to $n_1 \geq \frac{n_0}{\epsilon_{\text{tol}}^2} \widehat{\mathfrak{q}}_{1-\alpha}(n_0)^2$. Furthermore, our numerical results will demonstrate that this simple technique remains highly effective even when n_1 is much larger than n_0 , such as by an *order of magnitude* (see Sections 4.3 and 4.4).

4.2.2. Assessment of cost

Since the overall purpose of sketching is to reduce computation, it is important to explain why the added cost of the bootstrap is manageable. In particular, the added cost should not be much higher than the cost of sketching itself. As a simple point of reference, the cost to construct \tilde{A} and then compute $\tilde{A}^\top \tilde{A}$ with most state-of-the-art sketching algorithms is at least $\underline{C}_{\text{sketch}} = \Omega(dp + n_1 p^2)$, where n_1 refers to the “final” sketch size described above. Next, to assess the cost of the bootstrap, we can take advantage of a small initial sketch size n_0 by using extrapolation, as well as the fact that the bootstrap samples can be trivially computed in parallel, with say m processors. When these basic factors are taken into account, the cost of the bootstrap turns out to be at most $\bar{C}_{\text{boot}} = O(Bn_0 p^2/m)$.

From this discussion of cost, perhaps the most essential point to emphasize is that $\underline{C}_{\text{sketch}}$ grows linearly with d , whereas \bar{C}_{boot} is *independent of d* . Indeed, this is of great importance for scalability, because randomized matrix multiplication is of primary interest in situations where d is extremely large. Beyond this general observation, we can also take a more detailed look to see that the condition $\bar{C}_{\text{boot}} = O(\underline{C}_{\text{sketch}})$ occurs when $B = O((\frac{n_1}{n_0} + \frac{d}{pn_0})m)$. Furthermore, such a condition on B can be considered realistic in light of our experiments, since the modest choice of $B = 50$ is shown to yield good results.

4.3. Numerical results for synthetic matrices

We now demonstrate the performance of the bootstrap estimate $\widehat{\mathfrak{q}}_{1-\alpha}$ in a range of conditions, both with and without extrapolation. Most notably, the numerical results for extrapolation are quite encouraging.

Simulation settings. The choices for the matrix $A \in \mathbb{R}^{d \times p}$ were developed in analogy with those in Section 3, except that in this context, the matrix is very tall with $d = 10,000$ and $p = 1,000$. If we let $A = UDV^\top$ denote the singular value decomposition, then the singular vectors were specified by taking $U \in \mathbb{R}^{d \times p}$ and $V \in \mathbb{R}^{p \times p}$ to be orthonormal factors from QR decompositions of matrices filled with independent $N(0, 1)$ entries. In addition, the singular values were chosen as $\sigma_1(A) = \dots = \sigma_5(A) = 1$ and $\sigma_j(A) = j^{-\beta}$ for $j \in \{6, \dots, p\}$, with decay parameters $\beta \in \{0.75, 1.0, 1.25\}$. These values were chosen in order to show that the bootstrap can work even when there are no gaps among the leading singular values.

Design of simulations. The design of the simulations can be understood in terms of Figure 1. For each value of β , and sketch size $n \in \{300, \dots, 2,100\}$, we performed 1,000 trials of sketching to compute independent copies $\tilde{A} \in \mathbb{R}^{n \times p}$ using two different types of sketching matrices: Gaussian random projection,




Figure 1: (Sketching with Gaussian random projections.) The plots illustrate bootstrap estimates for the 90% quantile of the error $\|\tilde{A}^T \tilde{A} - A^T A\|_{\text{op}}$.

and uniform row sampling, as defined in Section 4.1. In turn, the actual values of $\|\tilde{A}^T \tilde{A} - A^T A\|_{\text{op}}$ in these trials yielded a high quality surrogate for the true 90% quantile $q_{0.9} = q_{0.9}(n)$, plotted as a function of n with the black dashed line.

With regard to Algorithm 1, it was applied during each trial to compute $\hat{q}_{1-\alpha}$ using $B = 50$ bootstrap samples. The average of these estimates is plotted as a function of n with the solid blue line. In addition, the performance of the extrapolation rule (4.2) was studied by applying it to each estimate $\hat{q}_{0.9}(n_0)$ computed at $n_0 = 300$. The average of the extrapolated curves is plotted in solid red, with the pink envelope signifying ± 1 standard deviation.

Comments on results. Results for the case of Gaussian random projections are displayed in Figure 1, and similar results for the case of uniform row sampling are given in Appendix I. Figure 1 shows that on average, the bootstrap estimates are nearly equal to the true quantile over the entire range of sketch sizes $n \in \{300, \dots, 2,100\}$, both with and without extrapolation. Indeed, the performance of the extrapolated estimate is especially striking, because it shows that bootstrapping a rough initial sketch \tilde{A} of size $300 \times 1,000$ can be used to accurately predict the error of a much larger sketch of size $2,100 \times 1,000$. To put this into context, we should also remember that the original matrix A is of size $10,000 \times 1,000$, and hence the initial sketch is able to provide quite a bit of information about the sketching task for a small computational price. Moreover, the fact that the extrapolation works up to the larger sketch size of 2,100 means that a 7-fold speedup can be obtained in comparison to naively applying Algorithm 1 to the larger sketch. (In fact, the plots suggest that the extrapolation would remain accurate for sketch sizes beyond 2,100, and that even larger speedups are attainable.) Lastly, it is worth noting that even though a small choice of $B = 50$ bootstrap samples was used, the standard deviation of the extrapolated estimate is rather well-behaved, as indicated by the pink envelope.

4.4. Sea surface temperature measurements

Large-scale dynamical systems are ubiquitous in the physical sciences, and advances in technology for measuring these systems have led to rapidly increasing volumes of data. Consequently, it is often too costly to apply standard tools of exploratory data analysis in a direct manner, and there has been growing interest to use sketching as a data-reduction strategy that preserves the essential information (e.g. Bai et al., 2019, Brunton et al., 2015, Erichson et al., 2019, Saibaba, 2019, Tropp et al., 2019, among others).

This type of situation is especially common in fields such as climate science and fluid dynamics, where we may be presented with a very large matrix $A \in \mathbb{R}^{d \times p}$ whose rows form a long sequence of




Figure 2: (a) The relevant ENSO region, marked with a rectangle. (b) The true ENSO mode, obtained by exact computation with the full product $A^T A$.

“snapshots” that represent a dynamical system at time points $1, \dots, d$. As a concrete example, we consider satellite recordings of sea surface temperature that have been collected over the time period 1981–2018, and are available from the National Oceanic and Atmospheric Administration (NOAA) (Reynolds et al., 2002). More specifically, we deal with a particular subset of the data corresponding to $d = 13,271$ temporal snapshots at $p = 3,944$ spatial grid points in the eastern Pacific Ocean, shown in Figure (2a). From the standpoint of climate science, this region is important for studying the phenomenon known as the El Niño Southern Oscillation (ENSO).

This example is relevant to our discussion of sketching for several reasons. First, the matrix product $A^T A$ is of interest because it describes spatial modes of temperature variation through its eigenstructure. In particular, the fourth eigenvector (mode) of $A^T A$ identifies the intermittent El Niño and La Niña warming events that are influential global weather patterns, as displayed in Figure (2b) (Erichson et al., 2020). Second, the singular values of A have a natural decay profile, which is illustrated in Figure (4a).




Figure 3: The left and right panels show approximations to the ENSO mode based on the approximate product $\tilde{A}^T \tilde{A}$, obtained from Gaussian random projections with sketch sizes $n = 500$ and $n = 3,000$. A comparison with the exact ENSO mode in Figure (2b) above shows that an insufficient sketch size can lead to a substantial distortion.




Figure 4: The left panel displays the decaying eigenvalues of $A^T A$, where the x-axis is logarithmic. The right panel demonstrates that the extrapolated and non-extrapolated bootstrap methods accurately estimate the 90% quantile of the sketching error $\|\tilde{A}^T \tilde{A} - A^T A\|_{op}$ over a wide range of sketch sizes. The extrapolation rule gives accurate results at a final sketch size $n_1 = 5,000$ that is 10 times larger than the initial sketch size $n_0 = 500$.

Lastly, the example demonstrates the need for error estimation in order to guide the choice of sketch size. This can be seen in Figures (3a) and (3b) below, where it is shown that an insufficient sketch size can heavily distort the ENSO mode in comparison to the exact form given in Figure (2b).

To conclude this example, we present numerical results for the bootstrap error estimates. Analogously to Section 4.3, we consider the task of estimating the 90% quantile $q_{0.9}(n)$ of the sketching error, viewed as a function of n . The full matrix A is of size $13,271 \times 3,944$, as described earlier, except that it was normalized to satisfy $\sigma_1(A) = 1$, so that the results here can be easily compared on the same scale with the previous results in Section 4.3. Also, the results shown here in Figure (4b) are plotted in the same format, with the number of trials being 1,000, the number of bootstrap samples being $B = 50$, and the sketching matrices being Gaussian random projections.

From looking at Figure (4b), we see that the performance of the bootstrap in the case of the naturally generated matrix A is very similar to that in the previous cases of synthetic matrices. Namely, the averages of both the extrapolated and non-extrapolated estimates virtually overlap with the true curve, and furthermore, the fluctuations of the extrapolated estimates are well controlled. Lastly, the extrapolation rule accurately estimates the quantile value $q_{0.9}(n_1)$ at a final sketch size $n_1 = 5,000$ that is 10 times larger than the initial sketch size $n_0 = 500$, which shows the potential of this rule to accelerate computations without sacrificing the quality of estimation.

Funding

MEL was supported in part by NSF grants DMS-1613218 and DMS-1915786. NBE and MWM were supported in part by ARO, DARPA (FA8750-17-2-0122), NSF, and ONR.

Supplementary Material

Supplement to “Bootstrapping the operator norm in high dimensions: Error estimation for covariance matrices and sketching” (DOI: [10.3150/22-BEJ1463SUPP](https://doi.org/10.3150/22-BEJ1463SUPP); .pdf). The supplementary material (Lopes, Erichson and Mahoney, 2023) contains proofs of theoretical results, as well as additional numerical results.

References

- Adamczak, R., Litvak, A.E., Pajor, A. and Tomczak-Jaegermann, N. (2011). Sharp bounds on the rate of convergence of the empirical covariance matrix. *C. R. Math. Acad. Sci. Paris* **349** 195–200. [MR2769907](#) <https://doi.org/10.1016/j.crma.2010.12.014>
- Ainsworth, M. and Oden, J.T. (2000). *A Posteriori Error Estimation in Finite Element Analysis. Pure and Applied Mathematics (New York)*. New York: Wiley Interscience. [MR1885308](#) <https://doi.org/10.1002/9781118032824>
- Ar, S., Blum, M., Codenotti, B. and Gemmell, P. (1993). Checking approximate computations over the reals. In *Annual ACM Symposium on Theory of Computing (STOC)*.
- Babuška, I. and Rheinboldt, W.C. (1978). Error estimates for adaptive finite element computations. *SIAM J. Numer. Anal.* **15** 736–754. [MR0483395](#) <https://doi.org/10.1137/0715049>
- Bai, Z., Erichson, N.B., Meena, M.G., Taira, K. and Brunton, S.L. (2019). Randomized methods to characterize large-scale vortical flow networks. *PLoS ONE* **14** e0225265. <https://doi.org/10.1371/journal.pone.0225265>
- Bertail, P. (1997). Second-order properties of an extrapolated bootstrap without replacement under weak assumptions. *Bernoulli* **3** 149–179. [MR1466305](#) <https://doi.org/10.2307/3318585>
- Bertail, P. and Politis, D.N. (2001). Extrapolation of subsampling distribution estimators: The i.i.d. and strong mixing cases. *Canad. J. Statist.* **29** 667–680. [MR1888512](#) <https://doi.org/10.2307/3316014>
- Bickel, P.J. and Levina, E. (2008). Regularized estimation of large covariance matrices. *Ann. Statist.* **36** 199–227. [MR2387969](#) <https://doi.org/10.1214/009053607000000758>
- Bickel, P.J. and Sakov, A. (2002). Extrapolation and the bootstrap. *Sankhyā: The Indian Journal of Statistics, Series A* 640–652.
- Bickel, P.J. and Yahav, J.A. (1988). Richardson extrapolation and the bootstrap. *J. Amer. Statist. Assoc.* **83** 387–393. [MR0971364](#)
- Brunton, S.L., Proctor, J.L., Tu, J.H. and Kutz, J.N. (2015). Compressed sensing and dynamic mode decomposition. *J. Comput. Dyn.* **2** 165–191. [MR3507892](#) <https://doi.org/10.3934/jcd.2015002>
- Bunea, F. and Xiao, L. (2015). On the sample covariance matrix estimator of reduced effective rank population matrices, with applications to fPCA. *Bernoulli* **21** 1200–1230. [MR3338661](#) <https://doi.org/10.3150/14-BEJ602>
- Cai, T.T., Zhang, C.-H. and Zhou, H.H. (2010). Optimal rates of convergence for covariance matrix estimation. *Ann. Statist.* **38** 2118–2144. [MR2676885](#) <https://doi.org/10.1214/09-AOS752>
- Cangiani, A., Georgoulis, E.H., Pryer, T. and Sutton, O.J. (2017). A posteriori error estimates for the virtual element method. *Numer. Math.* **137** 857–893. [MR3719046](#) <https://doi.org/10.1007/s00211-017-0891-9>
- Chen, R.Y., Gittens, A. and Tropp, J.A. (2012). The masked sample covariance estimator: An analysis using matrix concentration inequalities. *Inf. Inference* **1** 2–20. [MR3311439](#) <https://doi.org/10.1093/imaiai/ias001>
- Chernozhukov, V., Chetverikov, D. and Kato, K. (2013). Gaussian approximations and multiplier bootstrap for maxima of sums of high-dimensional random vectors. *Ann. Statist.* **41** 2786–2819. [MR3161448](#) <https://doi.org/10.1214/13-AOS1161>
- Chernozhukov, V., Chetverikov, D. and Kato, K. (2014). Gaussian approximation of suprema of empirical processes. *Ann. Statist.* **42** 1564–1597. [MR3262461](#) <https://doi.org/10.1214/14-AOS1230>
- Chernozhukov, V., Chetverikov, D. and Kato, K. (2016). Empirical and multiplier bootstraps for suprema of empirical processes of increasing complexity, and related Gaussian couplings. *Stochastic Process. Appl.* **126** 3632–3651. [MR3565470](#) <https://doi.org/10.1016/j.spa.2016.04.009>
- Chernozhukov, V., Chetverikov, D. and Kato, K. (2017). Central limit theorems and bootstrap in high dimensions. *Ann. Probab.* **45** 2309–2352. [MR3693963](#) <https://doi.org/10.1214/16-AOP1113>
- Cohen, M.B., Nelson, J. and Woodruff, D.P. (2016). Optimal approximate matrix product in terms of stable rank. In *43rd International Colloquium on Automata, Languages, and Programming. LIPIcs. Leibniz Int. Proc. Inform.* **55** Art. No. 11. Wadern: Schloss Dagstuhl. Leibniz-Zent. Inform. [MR3577072](#)
- Dixon, J.D. (1983). Estimating extremal eigenvalues and condition numbers of matrices. *SIAM J. Numer. Anal.* **20** 812–814. [MR0708458](#) <https://doi.org/10.1137/0720053>
- Drineas, P. and Kannan, R. (2001). Fast Monte-Carlo algorithms for approximate matrix multiplication. In *42nd IEEE Symposium on Foundations of Computer Science (Las Vegas, NV, 2001)* 452–459. Los Alamitos, CA: IEEE Computer Soc. [MR1948734](#)
- Drineas, P., Kannan, R. and Mahoney, M.W. (2006). Fast Monte Carlo algorithms for matrices. I. Approximating matrix multiplication. *SIAM J. Comput.* **36** 132–157. [MR2231643](#) <https://doi.org/10.1137/S0097539704442684>

- Drineas, P. and Mahoney, M.W. (2018). Lectures on randomized numerical linear algebra. In *The Mathematics of Data. IAS/Park City Math. Ser.* **25** 1–48. Providence, RI: Amer. Math. Soc. [MR3839166](#)
- El Karoui, N. and Purdom, E. (2019). The non-parametric bootstrap and spectral analysis in moderate and high-dimension. In *AISTATS 2019. Proceedings of Machine Learning Research* **89** 2115–2124.
- Erichson, N.B., Mathelin, L., Kutz, J.N. and Brunton, S.L. (2019). Randomized dynamic mode decomposition. *SIAM J. Appl. Dyn. Syst.* **18** 1867–1891. [MR4021266](#) <https://doi.org/10.1137/18M1215013>
- Erichson, N.B., Zheng, P., Manohar, K., Brunton, S.L., Kutz, J.N. and Aravkin, A.Y. (2020). Sparse principal component analysis via variable projection. *SIAM J. Appl. Math.* **80** 977–1002. [MR4091178](#) <https://doi.org/10.1137/18M1211350>
- Freivalds, R. (1979). Fast probabilistic algorithms. In *Mathematical Foundations of Computer Science, 1979 (Proc. Eighth Sympos., Olomouc, 1979). Lecture Notes in Computer Science* **74** 57–69. Berlin: Springer. [MR0570975](#)
- Gupta, V., Wang, S., Courtade, T. and Ramchandran, K. (2018). Oversketch: Approximate matrix multiplication for the cloud. In *2018 IEEE International Conference on Big Data (Big Data)* 298–304. IEEE.
- Halko, N., Martinsson, P.G. and Tropp, J.A. (2011). Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Rev.* **53** 217–288. [MR2806637](#) <https://doi.org/10.1137/090771806>
- Hall, P., Lee, Y.K., Park, B.U. and Paul, D. (2009). Tie-respecting bootstrap methods for estimating distributions of sets and functions of eigenvalues. *Bernoulli* **15** 380–401. [MR2543867](#) <https://doi.org/10.3150/08-BEJ154>
- Han, F., Xu, S. and Zhou, W.-X. (2018). On Gaussian comparison inequality and its application to spectral analysis of large random matrices. *Bernoulli* **24** 1787–1833. [MR3757515](#) <https://doi.org/10.3150/16-BEJ912>
- Holodnak, J.T. and Ipsen, I.C.F. (2015). Randomized approximation of the Gram matrix: Exact computation and probabilistic bounds. *SIAM J. Matrix Anal. Appl.* **36** 110–137. [MR3306014](#) <https://doi.org/10.1137/130940116>
- Hsu, D., Kakade, S.M. and Zhang, T. (2012). Tail inequalities for sums of random matrices that depend on the intrinsic dimension. *Electron. Commun. Probab.* **17** no. 14. [MR2900355](#) <https://doi.org/10.1214/ECP.v17-1869>
- Johnstone, I.M. and Paul, D. (2018). PCA in high dimensions: An orientation. *Proc. IEEE Inst. Electr. Electron. Eng.* **106** 1277–1292. <https://doi.org/10.1109/JPROC.2018.2846730>
- Kannan, R. and Vempala, S. (2017). Randomized algorithms in numerical linear algebra. *Acta Numer.* **26** 95–135. [MR3653851](#) <https://doi.org/10.1017/S0962492917000058>
- Koltchinskii, V., Löffler, M. and Nickl, R. (2020). Efficient estimation of linear functionals of principal components. *Ann. Statist.* **48** 464–490. [MR4065170](#) <https://doi.org/10.1214/19-AOS1816>
- Koltchinskii, V. and Lounici, K. (2017a). Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli* **23** 110–133. [MR3556768](#) <https://doi.org/10.3150/15-BEJ730>
- Koltchinskii, V. and Lounici, K. (2017b). Normal approximation and concentration of spectral projectors of sample covariance. *Ann. Statist.* **45** 121–157. [MR3611488](#) <https://doi.org/10.1214/16-AOS1437>
- Koltchinskii, V. and Lounici, K. (2017c). New asymptotic results in principal component analysis. *Sankhya A* **79** 254–297. [MR3707422](#) <https://doi.org/10.1007/s13171-017-0106-6>
- Liberty, E., Woolfe, F., Martinsson, P.-G., Rokhlin, V. and Tygert, M. (2007). Randomized algorithms for the low-rank approximation of matrices. *Proc. Natl. Acad. Sci. USA* **104** 20167–20172. [MR2366406](#) <https://doi.org/10.1073/pnas.0709640104>
- Lopes, M.E. (2019). Estimating the algorithmic variance of randomized ensembles via the bootstrap. *Ann. Statist.* **47** 1088–1112. [MR3909961](#) <https://doi.org/10.1214/18-AOS1707>
- Lopes, M.E., Blandino, A. and Aue, A. (2019). Bootstrapping spectral statistics in high dimensions. *Biometrika* **106** 781–801. [MR4031199](#) <https://doi.org/10.1093/biomet/asz040>
- Lopes, M.E., Erichson, N.B. and Mahoney, M.W. (2019). Bootstrapping the operator norm in high dimensions: Error estimation for covariance matrices and sketching. Available at [arXiv:1909.06120](https://arxiv.org/abs/1909.06120).
- Lopes, M.E., Erichson, N.B. and Mahoney, M.W. (2023). Supplement to “Bootstrapping the Operator Norm in High Dimensions: Error Estimation for Covariance Matrices and Sketching.” <https://doi.org/10.3150/22-BEJ1463SUPP>
- Lopes, M.E., Lin, Z. and Müller, H.-G. (2020). Bootstrapping max statistics in high dimensions: Near-parametric rates under weak variance decay and application to functional and multinomial data. *Ann. Statist.* **48** 1214–1229. [MR4102694](#) <https://doi.org/10.1214/19-AOS1844>
- Lopes, M.E., Wang, S. and Mahoney, M.W. (2018). Error estimation for randomized least-squares algorithms via the bootstrap. In *International Conference on Machine Learning (ICML)* 3223–3232.

- Lopes, M.E., Wang, S. and Mahoney, M.W. (2019). A bootstrap method for error estimation in randomized matrix multiplication. *J. Mach. Learn. Res.* **20** Paper No. 39. [MR3948079](#)
- Lounici, K. (2014). High-dimensional covariance matrix estimation with missing observations. *Bernoulli* **20** 1029–1058. [MR3217437](#) <https://doi.org/10.3150/12-BEJ487>
- Lust-Piquard, F. (1986). Inégalités de Khintchine dans C_p ($1 < p < \infty$). *C. R. Acad. Sci. Paris Sér. I Math.* **303** 289–292. [MR0859804](#)
- Mahoney, M.W. (2011). Randomized algorithms for matrices and data. *Found. Trends Mach. Learn.* **3** 123–224.
- Minsker, S. (2017). On some extensions of Bernstein’s inequality for self-adjoint operators. *Statist. Probab. Lett.* **127** 111–119. [MR3648301](#) <https://doi.org/10.1016/j.spl.2017.03.020>
- Naumov, A., Spokoiny, V. and Ulyanov, V. (2019). Bootstrap confidence sets for spectral projectors of sample covariance. *Probab. Theory Related Fields* **174** 1091–1132. [MR3980312](#) <https://doi.org/10.1007/s00440-018-0877-2>
- NOAA (2019). National Oceanic and Atmospheric Administration.
- Oliveira, R.I. (2010). Sums of random Hermitian matrices and an inequality by Rudelson. *Electron. Commun. Probab.* **15** 203–212. [MR2653725](#) <https://doi.org/10.1214/ECP.v15-1544>
- Pagh, R. (2013). Compressed matrix multiplication. *ACM Trans. Comput. Theory* **5** Art. 9. [MR3100651](#) <https://doi.org/10.1145/2493252.2493254>
- Pisier, G. (2016). *Martingales in Banach Spaces. Cambridge Studies in Advanced Mathematics* **155**. Cambridge: Cambridge Univ. Press. [MR3617459](#)
- Reiss, M. and Wahl, M. (2020). Nonasymptotic upper bounds for the reconstruction error of PCA. *Ann. Statist.* **48** 1098–1123. [MR4102689](#) <https://doi.org/10.1214/19-AOS1839>
- Reynolds, R.W., Rayner, N.A., Smith, T.M., Stokes, D.C. and Wang, W. (2002). An improved in situ and satellite SST analysis for climate. *J. Climate* **15** 1609–1625.
- Rudelson, M. (1999). Random vectors in the isotropic position. *J. Funct. Anal.* **164** 60–72. [MR1694526](#) <https://doi.org/10.1006/jfan.1998.3384>
- Rudelson, M. and Vershynin, R. (2007). Sampling from large matrices: An approach through geometric functional analysis. *J. ACM* **54** Art. 21. [MR2351844](#) <https://doi.org/10.1145/1255443.1255449>
- Saibaba, A.K. (2019). Randomized subspace iteration: Analysis of canonical angles and unitarily invariant norms. *SIAM J. Matrix Anal. Appl.* **40** 23–48. [MR3896637](#) <https://doi.org/10.1137/18M1179432>
- Sarlós, T. (2006). Improved approximation algorithms for large matrices via random projections. In *Annual IEEE Symposium on Foundations of Computer Science (FOCS)*.
- Silin, I. and Fan, J. (2020). Hypothesis testing for eigenspaces of covariance matrix. Available at [arXiv:2002.09810](#).
- Talagrand, M. (2014). *Upper and Lower Bounds for Stochastic Processes: Modern Methods and Classical Problems. Ergebnisse der Mathematik und Ihrer Grenzgebiete. 3. Folge. a Series of Modern Surveys in Mathematics [Results in Mathematics and Related Areas. 3rd Series. a Series of Modern Surveys in Mathematics]* **60**. Heidelberg: Springer. [MR3184689](#) <https://doi.org/10.1007/978-3-642-54075-2>
- Tropp, J.A., Yurtsever, A., Udell, M. and Cevher, V. (2019). Streaming low-rank matrix approximation with an application to scientific simulation. *SIAM J. Sci. Comput.* **41** A2430–A2463. [MR3986561](#) <https://doi.org/10.1137/18M1201068>
- van der Vaart, A.W. and Wellner, J.A. (2000). *Weak Convergence and Empirical Processes*. Berlin: Springer.
- Verfürth, R. (1994). A posteriori error estimation and adaptive mesh-refinement techniques. *J. Comput. Appl. Math.* **50** 67–83.
- Vershynin, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics* **47**. Cambridge: Cambridge Univ. Press. [MR3837109](#) <https://doi.org/10.1017/9781108231596>
- Woodruff, D.P. (2014). Sketching as a tool for numerical linear algebra. *Found. Trends Theor. Comput. Sci.* **10** iv+157. [MR3285427](#) <https://doi.org/10.1561/04000000060>
- Woolfe, F., Liberty, E., Rokhlin, V. and Tygert, M. (2008). A fast randomized algorithm for the approximation of matrices. *Appl. Comput. Harmon. Anal.* **25** 335–366. [MR2455599](#) <https://doi.org/10.1016/j.acha.2007.12.002>
- Yao, J. and Lopes, M.E. (2022). Rates of bootstrap approximation for eigenvalues in high-dimensional PCA. *Statist. Sinica*. To appear. Available at [arXiv:2104.07328](#).

Supplementary Material

Bootstrapping the Operator Norm in High Dimensions: Error Estimation for Covariance Matrices and Sketching

MILES E. LOPES¹, N. BENJAMIN ERICHSON² and MICHAEL W. MAHONEY²

¹*Department of Statistics, University of California, Davis, Mathematical Sciences Building, Davis, CA, 95616, USA, melopes@ucdavis.edu*

²*Department of Statistics, University of California, Berkeley, Evans Hall, Berkeley, CA, 94720, USA and International Computer Science Institute, Berkeley, CA, 94704, USA erichson@berkeley.edu, mmahoney@stat.berkeley.edu*

Appendix A: The term \mathbf{I}_n : localizing the maximizer of \mathbb{G}_n

The following proposition is the main result of this section, and it will be established with several lemmas later on.

Proposition A.1. *Suppose that Assumption 2.1 holds. Then, there is a constant $c > 0$ not depending on n such that*

$$\mathbf{I}_n \lesssim n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c.$$

Proof. Here, we only explain the proof at a high level, with the details being given in the remainder of this section. Observe that for any $t \in \mathbb{R}$, we have

$$\left| \mathbb{P}\left(\sup_{\theta \in \Theta} \mathbb{G}_n(\theta) \leq t\right) - \mathbb{P}\left(\sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n(\theta) \leq t\right) \right| = \mathbb{P}\left(\mathcal{A}(t) \cap \mathcal{B}(t)\right),$$

where we define the events

$$\mathcal{A}(t) = \left\{ \sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n(\theta) \leq t \right\} \quad \text{and} \quad \mathcal{B}(t) = \left\{ \sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n(\theta) > t \right\}.$$

For any pair of real numbers $t_{1,n}$ and $t_{2,n}$ satisfying $t_{1,n} \leq t_{2,n}$, it is possible to show that the inclusion

$$(\mathcal{A}(t) \cap \mathcal{B}(t)) \subset (\mathcal{A}(t_{2,n}) \cup \mathcal{B}(t_{1,n})) \tag{A.1}$$

holds simultaneously for all $t \in \mathbb{R}$. To see this, first consider the case when $t > t_{1,n}$. Then, $\mathcal{B}(t) \subset \mathcal{B}(t_{1,n})$, and hence

$$(\mathcal{A}(t) \cap \mathcal{B}(t)) \subset \mathcal{B}(t) \subset \mathcal{B}(t_{1,n}) \subset (\mathcal{A}(t_{2,n}) \cup \mathcal{B}(t_{1,n})).$$

In the opposite case when $t \leq t_{1,n}$, we must also have $t \leq t_{2,n}$, which implies $\mathcal{A}(t) \subset \mathcal{A}(t_{2,n})$, and the steps in the previous case can be used in a similar way. Applying a union bound to (A.1), and then taking the supremum over $t \in \mathbb{R}$, we obtain

$$\mathbf{I}_n \leq \mathbb{P}(\mathcal{A}(t_{2,n})) + \mathbb{P}(\mathcal{B}(t_{1,n})).$$

The difficult part of the proof is carried out below in Lemmas A.1 and A.5. In those results, we will determine values of $t_{1,n}$ and $t_{2,n}$ for which the probabilities $\mathbb{P}(\mathcal{A}(t_{2,n}))$ and $\mathbb{P}(\mathcal{B}(t_{1,n}))$ are at most of order $n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$. Furthermore, the chosen values of k_n and ℓ_n will ensure that the inequality $t_{1,n} \leq t_{2,n}$ holds for all large n . \square

Remark A.1. Note that in the special case where $k_n = p$, the matrix V_{k_n} is a square orthogonal matrix, which implies $\Theta_n^\dagger = \Theta$, and as a result, the terms \mathbf{I} and $\tilde{\mathbf{I}}_n$ become exactly 0. Therefore, in the proofs that handle the terms \mathbf{I}_n and $\tilde{\mathbf{I}}_n$, we may assume without loss of generality that $k_n < p$. This small reduction will be needed for specifying how quickly ℓ_n and k_n grow as a function of n , namely $\ell_n \asymp \log(n)^3$ and $k_n \asymp \log(n)^{\frac{6\beta+4}{\beta-1/2}}$.

A.1. Bounding the probability $\mathbb{P}(\mathcal{A}(t_{2,n}))$

In this subsection, we will need to introduce another special subset of Θ . Namely, let $v_1, \dots, v_{\ell_n} \in \mathbb{R}^p$ denote the ℓ_n leading right singular vectors of A , and define the set

$$\Theta_{\ell_n} = \{(Av_1, 1), \dots, (Av_{\ell_n}, 1)\}, \quad (\text{A.2})$$

which satisfies $\Theta_{\ell_n} \subset \Theta_n^\dagger$. Furthermore, we may add the points in Θ_{ℓ_n} to the net $\Theta_n^\dagger(\epsilon_n)$ while preserving the condition $\log \text{card}(\Theta_n^\dagger(\epsilon_n)) \lesssim \epsilon_n^{-1/\beta}$, since ℓ_n grows logarithmically in n , whereas $\epsilon_n^{-1/\beta}$ grows algebraically in n . For this reason, we may assume without loss of generality that the condition $\Theta_{\ell_n} \subset \Theta_n^\dagger(\epsilon_n)$ holds in our work below.

Lemma A.1. *Suppose that Assumption 2.1 holds. Then, there are constants $c, c_2 > 0$ not depending on n such that the choice*

$$t_{2,n} = c_2 \ell_n^{-2\beta} \sqrt{\log(\ell_n)}$$

implies

$$\mathbb{P}(\mathcal{A}(t_{2,n})) \lesssim n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c.$$

Proof. By the definitions of $\mathcal{A}(t_{2,n})$ and \mathbf{III}_n , we have

$$\begin{aligned} \mathbb{P}(\mathcal{A}(t_{2,n})) &\leq \mathbb{P}\left(\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \mathbb{G}_n(\theta) \leq t_{2,n}\right) && \text{(since } \Theta_n^\dagger(\epsilon_n) \subset \Theta_n^\dagger) \\ &\leq \mathbb{P}\left(\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) \leq t_{2,n}\right) + \mathbf{III}_n \\ &\leq \mathbb{P}\left(\sup_{\theta \in \Theta_{\ell_n}} G_n(\theta) \leq t_{2,n}\right) + \mathbf{III}_n && \text{(since } \Theta_{\ell_n} \subset \Theta_n^\dagger(\epsilon_n)). \end{aligned}$$

Lemma C.1 will show that the term \mathbf{III}_n is at most of order $n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$, and so it remains to control supremum of G_n over Θ_{ℓ_n} . This is a substantial task, involving several ingredients that are developed in subsequent lemmas, and so we only explain how the ingredients are combined here. To proceed, define the standardized version of G_n as

$$\bar{G}_n(\theta) = G_n(\theta) / \varsigma_n(\theta) \quad \text{where} \quad \varsigma_n(\theta) = \sqrt{\text{var}(G_n(\theta))}.$$

Also, define the minimum standard deviation $\varsigma_n^\circ = \inf_{\theta \in \Theta_{\ell_n}} \varsigma_n(\theta)$, which is shown to satisfy the following lower bound in Lemma F.1,

$$\varsigma_n^\circ \gtrsim \ell_n^{-2\beta}.$$

Based on the definition of \bar{G}_n , it follows that

$$\begin{aligned} \mathbb{P}\left(\sup_{\theta \in \Theta_{\ell_n}} G_n(\theta) \leq t_{2,n}\right) &\leq \mathbb{P}\left(\sup_{\theta \in \Theta_{\ell_n}} \bar{G}_n(\theta) \leq \frac{t_{2,n}}{\varsigma_n^\circ}\right) \\ &\leq \mathbb{P}\left(\sup_{\theta \in \Theta_{\ell_n}} \bar{G}_n(\theta) \leq \sqrt{\log(\ell_n)}\right), \end{aligned}$$

where the last step can be arranged by the choice of c_2 in the definition of $t_{2,n}$. Next, in Lemmas A.2 and A.3 below, we show that the last probability satisfies the bound

$$\mathbb{P}\left(\sup_{\theta \in \Theta_{\ell_n}} \bar{G}_n(\theta) \leq \sqrt{\log(\ell_n)}\right) \lesssim \exp\left(-\frac{1}{2}\ell_n^{1/3}\right).$$

Finally, as explained in the remark above Lemma A.1, we may assume $k_n < p$, which implies $\ell_n \geq \log(n)^3$, and hence $\exp(-\frac{1}{2}\ell_n^{1/3}) \leq n^{-1/2} \leq n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$. This completes the proof. \square

Remark. The following lemma was developed in the paper (Lopes, Lin and Müller, 2020, Lemma B.2), and later extended in (Lopes and Yao, 2022). Although there are many upper-tail bounds for the maxima of Gaussian processes, there are relatively few lower-tail bounds, which is the notable aspect of this result.

Lemma A.2. *For each integer $m \geq 1$, let $\mathbf{R} = \mathbf{R}(m)$ be a correlation matrix in $\mathbb{R}^{m \times m}$, and let $\mathbf{R}^+ = \mathbf{R}^+(m)$ denote the matrix with (i, j) entry given by $\max\{\mathbf{R}_{ij}, 0\}$. Suppose the*

matrix R^+ is positive semidefinite for all m , and that there are constants $\epsilon_1 \in (0, 1)$ and $c > 0$, not depending on m , such that the inequalities

$$\sum_{i \neq j} R_{ij}^+ \leq cm$$

$$\max_{i \neq j} R_{ij}^+ \leq 1 - \epsilon_1$$

hold for all m . Lastly, let $(\zeta_1, \dots, \zeta_m)$ be a Gaussian vector drawn from $N(0, R)$. Then, there is a constant $C > 0$, not depending on m , such that the following inequality holds for all $m \geq 1$,

$$\mathbb{P}\left(\max_{1 \leq j \leq m} \zeta_j \leq \sqrt{\log(m)}\right) \leq C \exp\left(-\frac{1}{2}m^{1/3}\right). \quad (\text{A.3})$$

Remark. In essence, the next lemma shows that if we restrict \mathbb{G}_n to the finite set Θ_{ℓ_n} , then the correlation matrix of the resulting vector satisfies the conditions of Lemma A.2 (as needed for the completion of the proof of Lemma A.1).

Lemma A.3. Let the elements of Θ_{ℓ_n} be written as $\{\theta_1, \dots, \theta_{\ell_n}\}$. Also, let $R(\ell_n) \in \mathbb{R}^{\ell_n \times \ell_n}$ denote the correlation matrix of $(\mathbb{G}_n(\theta_1), \dots, \mathbb{G}_n(\theta_{\ell_n}))$, and define the matrix $R^+(\ell_n) \in \mathbb{R}^{\ell_n \times \ell_n}$ as

$$R_{ij}^+(\ell_n) = \max\{R_{ij}(\ell_n), 0\}.$$

Under these conditions, it follows that $R^+(\ell_n)$ is positive semidefinite, and there is a constant $c > 0$ not depending on n such that

$$\sum_{i \neq j} R_{ij}^+(\ell_n) \leq c\ell_n. \quad (\text{A.4})$$

Furthermore, there is a constant $\epsilon_1 \in (0, 1)$ not depending on n such that

$$\max_{i \neq j} R_{ij}^+(\ell_n) \leq 1 - \epsilon_1. \quad (\text{A.5})$$

Proof. Below, we will write $R_{ij} = R_{ij}(\ell_n)$ to ease notation, and likewise for R_{ij}^+ . By Lemma F.1 the following identity holds for all $i, j \in \{1, \dots, \ell_n\}$, where we let u_1, \dots, u_{ℓ_n} denote the leading ℓ_n left singular vectors of A ,

$$\text{cov}(\mathbb{G}_n(\theta_i), \mathbb{G}_n(\theta_j)) = 2\sigma_i(A)^2\sigma_j^2(A)\left(1\{i=j\} + \frac{(\kappa-3)}{2}\sum_{l=1}^d \langle e_l, u_i \rangle^2 \langle e_l, u_j \rangle^2\right).$$

For distinct i and j , this leads to

$$R_{ij} = \frac{(\kappa-3)\sum_{l=1}^d \langle e_l, u_i \rangle^2 \langle e_l, u_j \rangle^2}{\sqrt{2 + (\kappa-3)\sum_{l=1}^d \langle e_l, u_i \rangle^4} \sqrt{2 + (\kappa-3)\sum_{l=1}^d \langle e_l, u_j \rangle^4}}. \quad (\text{A.6})$$

In the case when $\kappa \leq 3$, we have $R_{ij}^+ = 0$ for $i \neq j$, and so the matrix R^+ is clearly positive semidefinite. Furthermore, both of the bounds (A.4) and (A.5) hold in this case.

To consider the opposite case when $\kappa > 3$, observe that $R^+ = R$, and so again, the matrix R^+ is positive semidefinite. In addition, when $i \neq j$, the formula (A.6) implies

$$R_{ij}^+ \leq \frac{\kappa-3}{2} \sum_{l=1}^d \langle e_l, u_i \rangle^2 \langle e_l, u_j \rangle^2,$$

and so

$$\begin{aligned} \sum_{1 \leq i \neq j \leq \ell_n} R_{ij}^+ &\leq \frac{\kappa-3}{2} \sum_{i=1}^{\ell_n} \sum_{l=1}^d \langle e_l, u_i \rangle^2 \sum_{j=1}^{\ell_n} \langle e_l, u_j \rangle^2 \\ &\leq \frac{\kappa-3}{2} \sum_{i=1}^{\ell_n} \sum_{l=1}^d \langle e_l, u_i \rangle^2 \\ &= \frac{\kappa-3}{2} \ell_n, \end{aligned}$$

where we have used the fact that $\sum_{j=1}^{\ell_n} \langle e_l, u_j \rangle^2 \leq 1$. This proves the bound (A.4). Turning to the second bound (A.5), we may again assume $\kappa > 3$. If we let $a = \sum_{l=1}^d \langle e_l, u_i \rangle^4$ and $b = \sum_{l=1}^d \langle e_l, u_j \rangle^4$, then an application of the Cauchy-Schwarz inequality to (A.6) gives

$$\begin{aligned} R_{ij}^+ &\leq \frac{\sqrt{ab}}{\sqrt{\left(\frac{2}{\kappa-3} + a\right)\left(\frac{2}{\kappa-3} + b\right)}}, \\ &\leq \frac{\sqrt{ab}}{\sqrt{\left(\frac{2}{\kappa-3}\right)^2 + ab}} \\ &\leq \frac{1}{\sqrt{\left(\frac{2}{\kappa-3}\right)^2 + 1}}, \end{aligned}$$

where the last step follows from the the fact that $ab \leq 1$. This proves (A.5). \square

A.2. Bounding the probability $\mathbb{P}(\mathcal{B}(t_{1,n}))$

Remark. Before handling the probability $\mathbb{P}(\mathcal{B}(t_{1,n}))$ in Lemma A.5 below, it is necessary to state a lemma involving a bit of matrix analysis. The proof is straightforward and is hence omitted. For notation, let $A = UDV^\top$ denote the s.v.d. of A , where $U \in \mathbb{R}^{d \times p}$ and $V \in \mathbb{R}^{p \times p}$ have orthonormal columns, and $D \in \mathbb{R}^{p \times p}$ is diagonal.

Lemma A.4. *Fix any $\delta \in (0, 1)$, and any symmetric matrix $M \in \mathbb{R}^{d \times d}$. Also, let $V_{k_n} \in \mathbb{R}^{p \times k_n}$ denote the first k_n columns of V , and define the map $\mathbb{T}_{k_n}^\delta : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$ that scales the first k_n diagonal entries of a matrix by δ , and leaves all other entries unchanged. Then, there is an absolute constant $c > 0$ such that*

$$\sup_{\|w\|_2 \leq 1, \|V_{k_n}^\top w\|_2 \leq \delta} |w^\top A^\top M A w| \leq c \left\| \mathbb{T}_{k_n}^\delta(D) U^\top M U \mathbb{T}_{k_n}^\delta(D) \right\|_{\text{op}}.$$

We now complete this section with the following bound on $\mathbb{P}(\mathcal{B}(t_{1,n}))$.

Lemma A.5. *Suppose that Assumption 2.1 holds. Then, there exists a constant $c_1 > 0$ not depending on n such that the choice*

$$t_{1,n} = c_1 k_n^{-2\beta+1} \log(n)$$

implies

$$\mathbb{P}(\mathcal{B}(t_{1,n})) \lesssim \frac{1}{n}.$$

Proof. Let $Z \in \mathbb{R}^{n \times d}$ be the matrix whose rows are Z_1, \dots, Z_n . By the definition of Θ_n^\dagger , if we take $\delta = \frac{1}{2} k_n^{-\beta+1/2}$, then

$$\sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n(\theta) \leq \sup_{\|w\|_2 \leq 1, \|V_{k_n}^\top w\|_2 \leq \delta} \sqrt{n} |w^\top A^\top \left(\frac{1}{n} Z^\top Z - I_d \right) A w|.$$

This bound allows us to apply Lemma A.4, which gives

$$\sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n(\theta) \leq c\sqrt{n} \left\| \mathbb{T}_{k_n}^\delta(D) U^\top \left(\frac{1}{n} Z^\top Z - I_d \right) U \mathbb{T}_{k_n}^\delta(D) \right\|_{\text{op}}. \quad (\text{A.7})$$

Next, we apply the form of Theorem 2.2 given by the bound (2.9), along with the choices $\xi_i = \mathbb{T}_{k_n}^\delta(D) U^\top Z_i$ and $q = \log(n) \vee 3$. This gives

$$\begin{aligned} \left\| \sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n(\theta) \right\|_q &\lesssim \left\| \mathbb{E}[\xi_1 \xi_1^\top] \right\|_{\text{op}} \cdot \left(\sqrt{\mathfrak{r}(q)} \vee \frac{\mathfrak{r}(q)}{\sqrt{n}} \right) \\ &\lesssim \left(q \cdot \left\| \mathbb{E}[\xi_1 \xi_1^\top] \right\|_{\text{op}}^{1/2} \cdot \left\| \xi_1 \right\|_2 \right) \vee \left(\frac{q^2}{\sqrt{n}} \cdot \left\| \xi_1 \right\|_2 \right)_{\psi_2}^2. \end{aligned} \quad (\text{A.8})$$

To simplify this bound, note that

$$\left\| \mathbb{E}[\xi_1 \xi_1^\top] \right\|_{\text{op}}^{1/2} = \left\| \mathbb{T}_{k_n}^\delta(D) U^\top \right\|_{\text{op}} \lesssim k_n^{-\beta+1/2}.$$

Also, a background fact in Lemma G.2 gives

$$\begin{aligned} \left\| \xi \right\|_2 \Big|_{\psi_2} &= \left\| \mathbb{T}_{k_n}^\delta(D) U^\top Z_1 \right\|_2 \Big|_{\psi_2} \\ &\lesssim \left\| \mathbb{T}_{k_n}^\delta(D) U^\top \right\|_F \\ &= \left(\sum_{j=1}^{k_n} \delta^2 \sigma_j^2(A) + \sum_{j=k_n+1}^p \sigma_j^2(A) \right)^{1/2} \\ &\lesssim k_n^{-\beta+1/2}. \end{aligned} \quad (\text{A.9})$$

So, combining with the earlier bound (A.8), we have

$$\left\| \sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n(\theta) \right\|_q \lesssim k_n^{-2\beta+1} \log(n),$$

which leads to the stated result by Chebyshev's inequality. \square

Appendix B: The term \mathbf{II}_n : discrete approximation of \mathbb{G}_n

Lemma B.1. *Suppose Assumption 2.1 holds. Then, there is a constant $c > 0$ not depending on n such that*

$$\mathbf{II}_n \lesssim n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c.$$

Proof. The approach is based on the fact that the Kolmogorov metric can always be bounded in two parts: a coupling term and an anti-concentration term. More specifically, for any two random variables ξ and ζ defined on the same probability space, the following inequality holds for any $r > 0$,

$$d_K(\mathcal{L}(\xi), \mathcal{L}(\zeta)) \leq \sup_{t \in \mathbb{R}} \mathbb{P}(|\zeta - t| \leq r) + \mathbb{P}(|\xi - \zeta| \geq r). \quad (\text{B.1})$$

For the present context, we will let ζ play the role of $\sup_{\theta \in \Theta_n^+(\epsilon_n)} \mathbb{G}_n(\theta)$ and let ξ play the role of $\sup_{\theta \in \Theta_n^+} \mathbb{G}_n(\theta)$.

With regard to the coupling inequality, first note that under the stated choices for ξ and ζ we have

$$\|\xi - \zeta\|_q \leq \left\| \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon_n} |\mathbb{G}_n(\theta) - \mathbb{G}_n(\tilde{\theta})| \right\|_q. \quad (\text{B.2})$$

Secondly, note that Proposition 2.1 shows that the right side of this bound is at most of order $\epsilon_n^{1-\frac{1}{2\beta}} \log(n)$ when $q = \log(n) \vee 3$. Hence, it follows from Chebyshev's inequality that if we take $r = c\epsilon_n^{1-\frac{1}{2\beta}} \log(n)$ for a sufficiently large constant $c > 0$ not depending on n , then

$$\mathbb{P}(|\xi - \zeta| \geq r) \leq \frac{c}{n}.$$

Next, with regard to the anti-concentration inequality, we will approximate ζ with another random variable, say χ , and then use an anti-concentration inequality for χ instead. To do this, it can be verified using the definition of the Kolmogorov metric that the following inequality holds

$$\sup_{t \in \mathbb{R}} \mathbb{P}(|\zeta - t| \leq r) \leq \sup_{t \in \mathbb{R}} \mathbb{P}(|\chi - t| \leq 2r) + 2d_K(\mathcal{L}(\zeta), \mathcal{L}(\chi)). \quad (\text{B.3})$$

Hence, if we choose $\chi = \sup_{\theta \in \Theta_n^+(\epsilon_n)} \mathbb{G}_n(\theta)$, then the distance $d_K(\mathcal{L}(\zeta), \mathcal{L}(\chi))$ is the same as \mathbf{III}_n , which is shown to be of order $cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$ in Lemma C.1. Furthermore, by using the stated choice of $r = c\epsilon_n^{1-\frac{1}{2\beta}} \log(n)$ and the fact that $\epsilon_n^{1-\frac{1}{2\beta}} = n^{-\frac{\beta-1/2}{6\beta+4}}$, it follows from Lemma B.2 that

$$\sup_{t \in \mathbb{R}} \mathbb{P}(|\chi - t| \leq 2r) \lesssim n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c,$$

which completes the proof. \square

B.1. Anti-concentration inequality for G_n

Lemma B.2. *Suppose that Assumption 2.1 holds, and let $\delta_n \in (0, 1)$ be any numerical sequence with $\log(1/\delta_n) \lesssim \log(n)$. Then, there is a constant $c > 0$ not depending on n such that*

$$\sup_{t \in \mathbb{R}} \mathbb{P} \left(\left| \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) - t \right| \leq \delta_n \right) \lesssim \delta_n \log(n)^c.$$

Proof. For each $\theta \in \Theta$, let

$$\varsigma_n(\theta) = \sqrt{\text{var}(G_n(\theta))}$$

as well as

$$\bar{\varsigma}_n = \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \varsigma_n(\theta) \quad \text{and} \quad \underline{\varsigma}_n = \inf_{\theta \in \Theta_n^\dagger(\epsilon_n)} \varsigma_n(\theta).$$

In addition, define the expected supremum

$$\mu_n = \mathbb{E} \left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) / \varsigma_n(\theta) \right].$$

As a consequence of the anti-concentration inequality in Theorem 3 of (Chernozhukov, Chetverikov and Kato, 2015), we have

$$\sup_{t \in \mathbb{R}} \mathbb{P} \left(\left| \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) - t \right| \leq \delta_n \right) \lesssim \frac{\bar{\varsigma}_n}{\underline{\varsigma}_n} \cdot \delta_n \cdot \left(\mu_n + \sqrt{1 \vee \log(\bar{\varsigma}_n / \delta_n)} \right),$$

(Note that in the paper (Chernozhukov, Chetverikov and Kato, 2015), the dependence of the bound on $\bar{\varsigma}_n$ and $\underline{\varsigma}_n$ is not given explicitly, but a scan through the proof shows that it is sufficient to use a prefactor of $\frac{\bar{\varsigma}_n}{\underline{\varsigma}_n}$.) To control the dependence on $\bar{\varsigma}_n$ and $\underline{\varsigma}_n$, we may use Lemma F.1 and the assumption $\log(1/\delta_n) \lesssim \log(n)$ to obtain

$$\frac{\bar{\varsigma}_n}{\underline{\varsigma}_n} \lesssim \log(n)^c \quad \text{and} \quad \log(\bar{\varsigma}_n / \delta_n) \lesssim \log(n),$$

for some constant $c > 0$ that does not depend on n .

To complete the proof, we must bound μ_n . An initial step is to work with the unstandardized process G_n by using the bound

$$\begin{aligned} \mu_n &\leq \frac{1}{\underline{\varsigma}_n} \mathbb{E} \left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} |G_n(\theta)| \right], \\ &\leq \frac{2}{\underline{\varsigma}_n} \mathbb{E} \left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) \right] + \frac{1}{\underline{\varsigma}_n} \mathbb{E} \left[|G_n(\theta_0)| \right] \quad (\text{some } \theta_0 \in \Theta_n^\dagger(\epsilon_n)), \\ &\leq \frac{2}{\underline{\varsigma}_n} \mathbb{E} \left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) \right] + \frac{\bar{\varsigma}_n}{\underline{\varsigma}_n}, \end{aligned}$$

where the second step is a general fact about processes that are symmetric about the origin (Talagrand, 2014, p.14). Next, we will compare G_n with a simpler Gaussian process,

whose expected supremum can be analyzed more easily. Due to the Sudakov-Fernique inequality (van der Vaart and Wellner, 2000, Proposition A.2.6), if we can construct a centered Gaussian process on $\Theta_n^\dagger(\epsilon_n)$, say $\Gamma_n(\theta)$, that satisfies the condition

$$\mathbb{E}\left[\left(G_n(\theta) - G_n(\tilde{\theta})\right)^2\right] \leq \mathbb{E}\left[\left(\Gamma_n(\theta) - \Gamma_n(\tilde{\theta})\right)^2\right] \quad (\text{B.4})$$

for all $\theta, \tilde{\theta} \in \Theta_n^\dagger(\epsilon_n)$, then the following bound will hold

$$\mathbb{E}\left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta)\right] \leq \mathbb{E}\left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \Gamma_n(\theta)\right].$$

For this purpose, way may apply Lemma F.1 to obtain the following formula any $\theta = (v, s)$ and $\tilde{\theta} = (\tilde{v}, \tilde{s})$,

$$\begin{aligned} \mathbb{E}\left[\left(G_n(\theta) - G_n(\tilde{\theta})\right)^2\right] &= 2(\|v\|_2^4 + \|\tilde{v}\|_2^4) + (\kappa - 3)(\|v\|_4^4 + \|\tilde{v}\|_4^4) \\ &\quad - 4s\tilde{s}\langle v, \tilde{v} \rangle^2 - 2(\kappa - 3)s\tilde{s} \sum_{l=1}^d \langle e_l, v \rangle^2 \langle e_l, \tilde{v} \rangle^2. \end{aligned}$$

To simplify this expression, let $w, \tilde{w} \in \mathbb{R}^d$ be vectors with respective l th coordinates equal to $\langle e_l, v \rangle^2$ and $\langle e_l, \tilde{v} \rangle^2$. It can then be checked that

$$\mathbb{E}\left[\left(G_n(\theta) - G_n(\tilde{\theta})\right)^2\right] = 2\left(s\|v\|_2^2 - \tilde{s}\|\tilde{v}\|_2^2\right)^2 + (\kappa - 3)\|sw - \tilde{s}\tilde{w}\|_2^2 + 4s\tilde{s}\left(\|v\|_2^2\|\tilde{v}\|_2^2 - \langle v, \tilde{v} \rangle^2\right).$$

Letting the three terms on the right be denoted as J_1, J_2 , and J_3 , we can obtain the following bounds by using the fact that all vectors $v \in \mathcal{E}$ satisfy $\|v\|_2^2 \leq \|\Sigma\|_{\text{op}} \lesssim 1$. For J_1 , we have

$$\begin{aligned} J_1 &\leq 4(s - \tilde{s})^2\|v\|_2^4 + 4(\|v\|_2 + \|\tilde{v}\|_2)^2(\|v\|_2 - \|\tilde{v}\|_2)^2 \\ &\lesssim (s - \tilde{s})^2 + \|v - \tilde{v}\|_2^2. \end{aligned}$$

Next, for J_2 , we have

$$\begin{aligned} J_2 &\leq 2(\kappa - 3)_+\|w\|_2^2(s - \tilde{s})^2 + 2(\kappa - 3)_+\|w - \tilde{w}\|_2^2 \\ &\lesssim (s - \tilde{s})^2 + \|v - \tilde{v}\|_2^2. \end{aligned}$$

Lastly, for the third term, it can be checked that $J_3 \lesssim \|v - \tilde{v}\|_2^2$, and then combining leads to

$$\mathbb{E}\left[\left(G_n(\theta) - G_n(\tilde{\theta})\right)^2\right] \leq c_0(s - \tilde{s})^2 + c_0\|v - \tilde{v}\|_2^2,$$

for some constant $c_0 > 0$ that does not depend on n . Next, we define a centered Gaussian process $\Gamma_n(\theta)$ for any $\theta = (v, s)$ according to

$$\Gamma_n(\theta) = \sqrt{c_0} s \zeta_0 + \sqrt{c_0} \langle v, \zeta \rangle,$$

where $\zeta \in \mathbb{R}^d$ is a standard Gaussian vector, and $\zeta_0 \in \mathbb{R}$ is an independent standard Gaussian variable. This yields

$$\mathbb{E}\left[(\Gamma_n(\theta) - \Gamma_n(\tilde{\theta}))^2\right] = c_0(s - \tilde{s})^2 + c_0\|v - \tilde{v}\|_2^2,$$

which shows that the condition (B.4) indeed holds. Finally, the expected supremum of Γ_n over $\Theta_n^\dagger(\epsilon_n)$ satisfies

$$\begin{aligned} \mathbb{E}\left[\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \Gamma_n(\theta)\right] &\lesssim \mathbb{E}[|\zeta_0|] + \mathbb{E}\left[\sup_{\|u\|_2=1} \langle Au, \zeta \rangle\right] \\ &\lesssim 1 + \mathbb{E}\left[\|A^\top \zeta\|_2\right] \\ &\leq 1 + \|A\|_F \\ &\lesssim 1, \end{aligned}$$

which completes the proof. \square

Appendix C: The terms \mathbf{III}_n and $\widetilde{\mathbf{III}}_n$ (Gaussian and bootstrap approximation)

The following lemma is obtained as an application of the Gaussian and bootstrap approximation results in (Chernozhukov, Chetverikov and Kato, 2017).

Lemma C.1. *Suppose that Assumption 2.1 holds. Then, there is a constant $c > 0$ not depending on n such that*

$$\mathbf{III}_n \leq cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c, \quad (\text{C.1})$$

and the event

$$\widetilde{\mathbf{III}}_n \leq cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c \quad (\text{C.2})$$

holds with probability at least $1 - \frac{c}{n}$.

Proof. We first establish (C.1), and then turn to (C.2) at the end of the proof. Let $m = \text{card}(\Theta_n^\dagger(\epsilon_n))$, and define i.i.d. vectors $\xi_1, \dots, \xi_n \in \mathbb{R}^m$ as follows. Let $\{\theta_1, \dots, \theta_m\}$ be an enumeration of $\Theta_n^\dagger(\epsilon_n)$, with j th element represented as $\theta_j = (v_j, s_j)$. Next, for each $i \in \{1, \dots, n\}$ and $j \in \{1, \dots, m\}$, define the random variable

$$\xi_{ij} = k_n^{4\beta-1} s_j \left(\langle Z_i, v_j \rangle^2 - \mathbb{E}[\langle Z_i, v_j \rangle^2] \right).$$

(The scale factor $k_n^{4\beta-1}$ will only play a technical role in order to prevent the variance of ξ_{ij} from becoming too small.) This definition gives the relation

$$\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} k_n^{4\beta-1} \mathbb{G}_n(\theta) = \max_{1 \leq j \leq m} \sqrt{n} \bar{\xi}_j \quad (\text{C.3})$$

where $\bar{\xi}_j = \frac{1}{n} \sum_{i=1}^n \xi_{ij}$. Although the left side of this relation is a scaled version of $\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \mathbb{G}_n(\theta)$, it is important to note that the Kolmogorov metric is scale-invariant, and so the distance between the suprema of $k_n^{4\beta-1} \mathbb{G}_n$ and $k_n^{4\beta-1} G_n$ is equivalent to the distance between the suprema of \mathbb{G}_n and G_n .

The proof of (C.1) is completed by applying Proposition 2.1 in (Chernozhukov, Chetverikov and Kato, 2017) to $\max_{1 \leq j \leq m} \sqrt{n} \bar{\xi}_j$. In order to apply this result, it is enough to note that the vectors ξ_1, \dots, ξ_n are centered, i.i.d., and satisfy the following conditions, which can be verified using Lemmas F.1, G.1, and G.2:

$$\min_{1 \leq j \leq m} \text{var}(\xi_{1j}) \gtrsim 1 \quad (\text{C.4})$$

$$\max_{1 \leq j \leq m} \|\xi_{1j}\|_{\psi_1} \lesssim k_n^{4\beta-1}. \quad (\text{C.5})$$

$$\max_{1 \leq j \leq m} \mathbb{E}[|\xi_{1j}|^{2+l}] \lesssim k_n^{l(12\beta-3)} \quad \text{for } l \in \{1, 2\}. \quad (\text{C.6})$$

Based on these conditions, as well as $k_n \lesssim \log(n)^c$ and $\log(m) \lesssim \epsilon_n^{-1/\beta}$ (by Lemma G.3), it follows from Proposition 2.1 in (Chernozhukov, Chetverikov and Kato, 2017) that there is a constant $c > 0$ not depending on n such that

$$d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} k_n^{4\beta-1} \mathbb{G}_n(\theta)\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} k_n^{4\beta-1} G_n(\theta)\right)\right) \lesssim n^{-\frac{1}{6}} \epsilon_n^{-\frac{7}{6\beta}} \log(n)^c.$$

Substituting in the choice $\epsilon_n = n^{-\frac{\beta}{6\beta+4}}$ leads to (C.1).

Finally, to prove (C.2), let $(\xi_1^*, \dots, \xi_n^*)$ be drawn with replacement from (ξ_1, \dots, ξ_n) , and let $\bar{\xi}_j^* = \frac{1}{n} \sum_{i=1}^n \xi_{ij}^*$ for any $j \in \{1, \dots, m\}$. This gives the bootstrap counterpart of the relation (C.3),

$$\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} k_n^{4\beta-1} \mathbb{G}_n^*(\theta) = \max_{1 \leq j \leq m} \sqrt{n} (\bar{\xi}_j^* - \bar{\xi}_j).$$

Due to this relation, Proposition 4.3 in the aforementioned paper shows that under the conditions (C.4)-(C.6), there is a constant $c > 0$ not depending on n such that the event

$$d_K\left(\mathcal{L}\left(\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} k_n^{4\beta-1} \mathbb{G}_n^*(\theta) \mid X\right), \mathcal{L}\left(\sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} k_n^{4\beta-1} G_n(\theta)\right)\right) \lesssim n^{-\frac{1}{6}} \epsilon_n^{-\frac{7}{6\beta}} \log(n)^c,$$

holds with probability at least $1 - \frac{c}{n}$. As before, substituting in the choice $\epsilon_n = n^{-\frac{\beta}{6\beta+4}}$ leads to (C.2). \square

Appendix D: The term $\tilde{\Pi}_n$: discrete approximation of \mathbb{G}_n^*

Lemma D.1. *Suppose that Assumption 2.1 holds. Then, there is a constant $c > 0$ not depending on n such that the event*

$$\tilde{\Pi}_n \leq cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$$

occurs with probability at least $1 - \frac{c}{n}$.

Proof. Recall the Kolmogorov distance can always be bounded in terms of an anti-concentration term and a coupling term, as in (B.1). Using such an approach, we have

$$\tilde{\mathbf{I}}_n \leq \tilde{\mathbf{I}}'_n + \tilde{\mathbf{I}}''_n$$

where we define the following terms for a fixed number $\delta > 0$,

$$\tilde{\mathbf{I}}'_n = \sup_{t \in \mathbb{R}} \mathbb{P} \left(\left| \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \mathbb{G}_n^*(\theta) - t \right| \leq \delta \mid X \right)$$

and

$$\tilde{\mathbf{I}}''_n = \mathbb{P} \left(\left| \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \mathbb{G}_n^*(\theta) - \sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n^*(\theta) \right| \geq \delta \mid X \right).$$

When $\delta = c\epsilon_n^{1-\frac{1}{2\beta}} \log(n)^c$ for a sufficiently large constant $c > 0$ not depending on n , we will show in Proposition D.1 below that $\tilde{\mathbf{I}}''_n$ is at most c/n with probability at least $1 - c/n$.

To address the anti-concentration term, recall the inequality (B.3), which implies

$$\tilde{\mathbf{I}}'_n \leq \sup_{t \in \mathbb{R}} \mathbb{P} \left(\left| \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} G_n(\theta) - t \right| \leq 2\delta \right) + 2\tilde{\mathbf{I}}\mathbf{I}_n.$$

For the stated choice of δ , it is shown in Lemma B.2 that the first term on the right side is at most of order $\epsilon_n^{1-\frac{1}{2\beta}} \log(n)^c = n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$. Finally, Lemma C.1 shows that the event $\tilde{\mathbf{I}}\mathbf{I}_n \leq cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$ holds with probability at least $1 - c/n$, which completes the proof. \square

Remark. To introduce another piece of notation, for any $q \geq 1$ and random variable ξ , define the conditional norms

$$\|\xi\|_{q|X} = \mathbb{E} \left[|\xi|^q \mid X \right]^{1/q} \quad \text{and} \quad \|\xi\|_{\psi_2|X} = \inf \left\{ r > 0 \mid \mathbb{E}[\psi_2(|\xi|/r) \mid X] \leq 1 \right\}. \quad (\text{D.1})$$

Proposition D.1. *Let $q = \log(n) \vee 3$, and suppose that Assumption 2.1 holds. Then, there is a constant $c > 0$ not depending on n such that the event*

$$\left\| \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon_n} |\mathbb{G}_n^*(\theta) - \mathbb{G}_n^*(\tilde{\theta})| \right\|_{q|X} \leq c\epsilon_n^{1-\frac{1}{2\beta}} \log(n)^{5/2} \quad (\text{D.2})$$

holds with probability at least $1 - \frac{c}{n}$, and the event

$$\mathbb{P} \left(\left| \sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n^*(\theta) - \sup_{\theta \in \Theta_n^\dagger(\epsilon_n)} \mathbb{G}_n^*(\theta) \right| \geq c\epsilon_n^{1-\frac{1}{2\beta}} \log(n)^{5/2} \mid X \right) \leq \frac{c}{n} \quad (\text{D.3})$$

also holds with probability at least $1 - \frac{c}{n}$.

Proof. We only prove (D.2), since (D.3) is essentially a direct consequence, using the same reasoning as in (B.2). To begin, note that the first half of the of the proof of Proposition 2.1 can be repeated to show that

$$\sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon_n} |\mathbb{G}_n^*(\theta) - \mathbb{G}_n^*(\tilde{\theta})| \leq \sup_{w \in \mathbb{B}_{2p}(2)} \frac{2}{\sqrt{n}} \left| \sum_{i=1}^n \langle Z_i^*, A(t, \epsilon_n) w \rangle^2 - \mathbb{E}[\langle Z_i^*, A(t, \epsilon_n) w \rangle^2 | X] \right|,$$

where (Z_1^*, \dots, Z_n^*) are i.i.d. samples with replacement from (Z_1, \dots, Z_n) , and we retain the definition of $A(t, \epsilon_n)$ from that proof. Next, we use the shorthand $\xi_i^* = A(t, \epsilon_n)^\top Z_i^*$ and apply Theorem 2.2 to the right side above, yielding

$$\left\| \sup_{\rho(\theta, \tilde{\theta}) \leq \epsilon_n} |\mathbb{G}_n^*(\theta) - \mathbb{G}_n^*(\tilde{\theta})| \right\|_{q|X} \leq c \left(\sqrt{q} \cdot \|\mathbb{E}[\xi_1^* (\xi_1^*)^\top | X]\|_{\text{op}}^{1/2} \cdot \mathbb{E}[\|\xi_1^*\|_2^{2q} | X]^{\frac{1}{2q}} \right) \vee \left(\frac{q}{\sqrt{n}} \mathbb{E}[\|\xi_1^*\|_2^{2q} | X]^{\frac{1}{q}} \right). \quad (\text{D.4})$$

To simplify the previous bound, note that since Z_1^* is drawn uniformly from (Z_1, \dots, Z_n) , it follows that the inequality

$$\mathbb{E}[\|\xi_1^*\|_2^{2q} | X]^{\frac{1}{q}} \leq \max_{1 \leq i \leq n} \|A(t, \epsilon_n)^\top Z_i\|_2^2$$

holds almost surely, and similarly

$$\begin{aligned} \|\mathbb{E}[\xi_1^* (\xi_1^*)^\top | X]\|_{\text{op}} &= \left\| \frac{1}{n} \sum_{i=1}^n \xi_i \xi_i^\top \right\|_{\text{op}} \\ &\leq \max_{1 \leq i \leq n} \|A(t, \epsilon_n)^\top Z_i\|_2^2. \end{aligned}$$

Hence, to complete the proof, it suffices to derive a high-probability bound on $\max_{1 \leq i \leq n} \|A(t, \epsilon_n)^\top Z_i\|_2^2$. Using the facts in Lemmas G.1 and G.2, as well as the earlier bounds (2.16) and (2.17), we have

$$\begin{aligned} \left\| \max_{1 \leq i \leq n} \|A(t, \epsilon_n)^\top Z_i\|_2^2 \right\|_{\psi_1} &\lesssim \log(n) \left\| \|A(t, \epsilon_n)^\top Z_1\|_2 \right\|_{\psi_2}^2 \\ &\lesssim \log(n) \|A(t, \epsilon_n)\|_F^2 \\ &\lesssim \log(n) \left(t + \frac{1}{t} \epsilon_n^{1 - \frac{1}{2\beta}} \right)^2. \end{aligned}$$

Therefore, taking $t = \epsilon_n^{\frac{1}{2} - \frac{1}{4\beta}}$ implies

$$\mathbb{P} \left(\max_{1 \leq i \leq n} \|A(t, \epsilon_n)^\top Z_i\|_2^2 \geq c \log(n)^2 \epsilon_n^{1 - \frac{1}{2\beta}} \right) \leq \frac{c}{n},$$

which leads to (D.2) after combining with (D.4). \square

Appendix E: The term $\tilde{\mathbf{I}}_n$: localizing the maximizer of \mathbb{G}_n^*

Lemma E.1. *Suppose that Assumption 2.1 holds. Then, there is a constant $c > 0$, not depending on n , such that the event*

$$\tilde{\mathbf{I}}_n \leq cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$$

holds with probability at least $1 - c/n$.

Proof. Observe that

$$\begin{aligned} \tilde{\mathbf{I}}_n &= \sup_{t \in \mathbb{R}} \left| \mathbb{P}\left(\sup_{\theta \in \Theta_n^*} \mathbb{G}_n^*(\theta) \leq t \mid X\right) - \mathbb{P}\left(\sup_{\theta \in \Theta} \mathbb{G}_n^*(\theta) \leq t \mid X\right) \right| \\ &= \sup_{t \in \mathbb{R}} \mathbb{P}\left(\mathcal{A}'(t) \cap \mathcal{B}'(t) \mid X\right), \end{aligned}$$

where we define the events

$$\mathcal{A}'(t) = \left\{ \sup_{\theta \in \Theta_n^*} \mathbb{G}_n^*(\theta) \leq t \right\}$$

and

$$\mathcal{B}'(t) = \left\{ \sup_{\theta \in \Theta \setminus \Theta_n^*} \mathbb{G}_n^*(\theta) > t \right\}.$$

By repeating the argument in the proof of Proposition A.1, the following inequality holds for any real numbers $t'_{n,1}$ and $t'_{n,2}$ satisfying $t'_{1,n} \leq t'_{2,n}$,

$$\tilde{\mathbf{I}}_n \leq \mathbb{P}(\mathcal{A}'(t'_{2,n}) \mid X) + \mathbb{P}(\mathcal{B}'(t'_{1,n}) \mid X).$$

To complete the proof, it remains to show there are choices of $t'_{1,n}$ and $t'_{2,n}$ such that $t'_{1,n} \leq t'_{2,n}$ for all large n , and the quantities $\mathbb{P}(\mathcal{A}'(t'_{2,n}) \mid X)$ and $\mathbb{P}(\mathcal{B}'(t'_{1,n}) \mid X)$ are at most of order $n^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c$ with probability at least $1 - c/n$. Such choices of $t'_{1,n}$ and $t'_{2,n}$ are established below in Lemma E.2. Note also that the condition $t'_{1,n} \leq t'_{2,n}$ only needs to be established in the case when $k_n < p$, due to the considerations in Remark A.1. \square

Lemma E.2. *Suppose that Assumption 2.1 holds. Then, there are positive constants c_1 , c_2 , and c , not depending on n , for which the following statement is true:*

If $t'_{1,n}$ and $t'_{2,n}$ are chosen as

$$t'_{1,n} = c_1 k_n^{-2\beta+1} \log(n)^3 \tag{E.1}$$

$$t'_{2,n} = c_2 \ell_n^{-2\beta} \sqrt{\log(\ell_n)}, \tag{E.2}$$

then the events

$$\mathbb{P}(\mathcal{A}'(t'_{2,n}) \mid X) \leq cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c \tag{E.3}$$

and

$$\mathbb{P}(\mathcal{B}'(t'_{1,n}) \mid X) \leq \frac{c}{n}$$

both occur with probability at least $1 - \frac{c}{n}$.

Proof. Based on the definitions of $\widetilde{\mathbf{II}}_n$, $\widetilde{\mathbf{III}}_n$, and \mathbf{III}_n , we have

$$\begin{aligned} \mathbb{P}(\mathcal{A}'(t'_{n,2})|X) &= \mathbb{P}(\sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n^*(\theta) \leq t'_{n,2} | X) \\ &\leq \mathbb{P}(\sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n(\theta) \leq t'_{n,2}) + \widetilde{\mathbf{II}}_n + \widetilde{\mathbf{III}}_n + \mathbf{III}_n + \mathbf{II}_n. \end{aligned}$$

With regard to the first term of the last line, Lemma A.1 shows that the following holds for a suitable choice of c_2 in (E.2),

$$\mathbb{P}(\sup_{\theta \in \Theta_n^\dagger} \mathbb{G}_n(\theta) \leq t'_{n,2}) \lesssim cn^{-\frac{\beta-1/2}{6\beta+4}} \log(n)^c.$$

Combining this with the bounds on $\widetilde{\mathbf{II}}_n$, $\widetilde{\mathbf{III}}_n$, \mathbf{III}_n , and \mathbf{II}_n in Lemmas D.1, C.1, and B.1, we reach the stated result in (E.3).

We now turn to controlling $\mathbb{P}(\mathcal{B}'(t'_{1,n})|X)$. Letting $q = \log(n) \vee 3$, the basic goal is to identify a number b_n that satisfies

$$\left\| \sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n^*(\theta) \right\|_{q|X} \leq b_n \quad (\text{E.4})$$

with probability at least $1 - \frac{c}{n}$. If this can be established, then Chebyshev's inequality will imply that the bound

$$\mathbb{P}(\sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n^*(\theta) \geq eb_n | X) \leq e^{-q}$$

holds with probability at least $1 - \frac{c}{n}$. Hence, the number eb_n corresponds to $t'_{1,n}$, and also, our choice of q gives $e^{-q} \leq 1/n$.

To proceed with the details, recall that the s.v.d. of A is written as $A = UDV^\top$, and for any $\delta > 0$, the map $\mathbf{T}_{k_n}^\delta : \mathbb{R}^{p \times p} \rightarrow \mathbb{R}^{p \times p}$ is defined to act on a matrix by scaling the first k_n diagonal entries by δ and leaving all other entries unchanged. Also, let $\delta = \frac{1}{2}k_n^{-\beta+1/2}$ and let $\xi_i^* = \mathbf{T}_{k_n}^\delta(D)U^\top Z_i^*$, where (Z_1^*, \dots, Z_n^*) are i.i.d. samples with replacement from (Z_1, \dots, Z_n) .

With this notation in place, the argument leading up to (A.8) in the proof of Lemma A.5 can be repeated for the process \mathbb{G}_n^* to obtain

$$\left\| \sup_{\theta \in \Theta \setminus \Theta_n^\dagger} \mathbb{G}_n^*(\theta) \right\|_{q|X} \leq c \left(q \cdot \left\| \mathbb{E}[\xi_1^* (\xi_1^*)^\top | X] \right\|_{\text{op}}^{1/2} \cdot \|\xi_1^*\|_2 \right) \vee \left(\frac{q^2}{\sqrt{n}} \cdot \|\xi_1^*\|_2^2 \right), \quad (\text{E.5})$$

for some constant $c > 0$ not depending on n . To simplify this bound, first notice that if we let $\xi_i = \mathbf{T}_{k_n}^\delta(D)U^\top Z_i$, then

$$\mathbb{E}[\xi_1^* (\xi_1^*)^\top | X] = \frac{1}{n} \sum_{i=1}^n \xi_i \xi_i^\top,$$

which leads to

$$\left\| \mathbb{E}[\xi_1^* (\xi_1^*)^\top | X] \right\|_{\text{op}}^{1/2} \leq \max_{1 \leq i \leq n} \|\mathbf{T}_{k_n}^\delta(D)U^\top Z_i\|_2.$$

Similarly, we have

$$\|\xi_1^*\|_2 \left\| \right\|_{\psi_2|X} \leq c \max_{1 \leq i \leq n} \|\mathbf{T}_{k_n}^\delta(D)U^\top Z_i\|_2,$$

for some constant $c > 0$ that does not depend on n . In turn, we may use the facts about Orlicz norms given in Lemmas G.1 and G.2 to obtain

$$\begin{aligned} \left\| \max_{1 \leq i \leq n} \|\mathbb{T}_{k_n}^\delta(D)U^\top Z_i\|_2 \right\|_{\psi_2} &\lesssim \sqrt{\log(n)} \|\mathbb{T}_{k_n}^\delta(D)U^\top Z_1\|_2 \Big\|_{\psi_2} \\ &\lesssim \sqrt{\log(n)} \|\mathbb{T}_{k_n}^\delta(D)\|_F \\ &\lesssim \sqrt{\log(n)} k_n^{-\beta+1/2}, \end{aligned}$$

where the last step re-uses the calculation from (A.9). This implies that the bounds

$$\left\| \mathbb{E}[\xi_1^*(\xi_1^*)^\top | X] \right\|_{\text{op}}^{1/2} \leq c \log(n) k_n^{-\beta+1/2}$$

and

$$\left\| \|\xi_1^*\|_2 \right\|_{\psi_2|X} \leq c \log(n) k_n^{-\beta+1/2}$$

simultaneously hold with probability at least $1 - c/n$. Therefore, combining with the bound (E.5) shows that the event

$$\left\| \sup_{\theta \in \Theta \setminus \Theta_n^*} \mathbb{G}_n^*(\theta) \right\|_{q|X} \leq c \log(n)^3 k_n^{-2\beta+1}$$

holds with probability at least $1 - c/n$. Hence, the number b_n in (E.4) may be taken proportional to $\log(n)^3 k_n^{-2\beta+1}$, which completes the proof. \square

Appendix F: Supporting results and proofs

This section contains a lemma summarizing facts about the covariance structure of the process \mathbb{G}_n (Lemma F.1), as well as some supporting details involved in the proof of Proposition 2.1, and a proof of the lower bound (2.5) from the main text.

Lemma F.1. *Suppose that Assumption 2.1 holds, and let two generic elements of Θ be denoted as $\theta = (v, s)$ and $\tilde{\theta} = (\tilde{v}, \tilde{s})$. Then,*

$$\text{cov}(\mathbb{G}_n(\theta), \mathbb{G}_n(\tilde{\theta})) = 2s\tilde{s}\langle v, \tilde{v} \rangle^2 + (\kappa - 3)s\tilde{s} \sum_{l=1}^d \langle e_l, v \rangle^2 \langle e_l, \tilde{v} \rangle^2, \quad (\text{F.1})$$

as well as

$$\inf_{\theta \in \Theta_n^*} \sqrt{\text{var}(\mathbb{G}_n(\theta))} \gtrsim k_n^{-4\beta+1} \quad \text{and} \quad \sup_{\theta \in \Theta_n^*} \sqrt{\text{var}(\mathbb{G}_n(\theta))} \lesssim 1. \quad (\text{F.2})$$

Furthermore, if u_i and v_i are the left and right singular vectors of A corresponding to the singular value $\sigma_i(A)$, and we let $\theta_i = (Av_i, 1)$, then

$$\text{cov}(\mathbb{G}_n(\theta_i), \mathbb{G}_n(\theta_j)) = 2\sigma_i^2(A)\sigma_j^2(A) \left(1\{i=j\} + \frac{(\kappa-3)}{2} \sum_{l=1}^d \langle e_l, u_i \rangle^2 \langle e_l, u_j \rangle^2 \right). \quad (\text{F.3})$$

Lastly, if we let Θ_{ℓ_n} be as defined in (A.2), then

$$\inf_{\theta \in \Theta_{\ell_n}} \sqrt{\text{var}(\mathbb{G}_n(\theta))} \gtrsim \ell_n^{-2\beta}. \quad (\text{F.4})$$

Proof. We start with the basic identity

$$\text{cov}(\mathbb{G}_n(\theta), \mathbb{G}_n(\tilde{\theta})) = \text{cov}(sZ_1^\top vv^\top Z_1, \tilde{s}Z_1^\top \tilde{v}\tilde{v}^\top Z_1).$$

Since the entries of Z_1 are standardized and independent with kurtosis κ , it follows from (Bai and Silverstein, 2010, eqn. 9.8.6) that

$$\text{cov}(sZ_1^\top vv^\top Z_1, \tilde{s}Z_1^\top \tilde{v}\tilde{v}^\top Z_1) = 2s\tilde{s} \text{tr}(vv^\top \tilde{v}\tilde{v}^\top) + (\kappa - 3)s\tilde{s} \sum_{l=1}^d \langle e_l, v \rangle^2 \langle e_l, \tilde{v} \rangle^2,$$

which implies both (F.1) and (F.3).

To establish the lower bound in (F.2), observe that the previous paragraph gives

$$\text{var}(\mathbb{G}_n(\theta)) = 2\|v\|_2^4 + (\kappa - 3)\|v\|_4^4. \quad (\text{F.5})$$

Due to the assumption $\kappa > 1$, there is some fixed $\epsilon_0 \in (0, 1)$ not depending n such that $\kappa - 3 \geq -(2 - \epsilon_0)$. Consequently, the basic inequality $\|v\|_2 \geq \|v\|_4$ implies $(\kappa - 3)\|v\|_4^4 \geq -(2 - \epsilon_0)\|v\|_2^4$, and hence

$$\text{var}(\mathbb{G}_n(\theta)) \geq \epsilon_0 \|v\|_2^4. \quad (\text{F.6})$$

Next, observe that for any $v \in \Theta_n^\dagger$, there is some $w \in \mathbb{S}^{p-1}$ satisfying $v = Aw$ and $\|V_{k_n}^\top w\|_2 > \frac{1}{2}k_n^{-\beta+1/2}$. Hence, the spectral decomposition $A^\top A = VD^2V^\top$, with v_l denoting the l th column of V leads to

$$\begin{aligned} \|v\|_2^2 &= w^\top A^\top Aw \\ &\geq \sum_{l=1}^{k_n} \sigma_l^2(A) \langle w, v_l \rangle^2 \\ &\geq \sigma_{k_n}^2(A) \|V_{k_n}^\top w\|_2^2 \\ &\gtrsim k_n^{-4\beta+1}. \end{aligned}$$

This implies the lower bound in (F.2). Meanwhile, the upper bound in (F.2) follows from (F.5) and the fact that $\|v\|_2^2 \leq \|A^\top A\|_{\text{op}} \lesssim 1$. Finally, the lower bound (F.4) follows from (F.6). \square

Lemma F.2. *Let the vectors $\omega(t)$ and $\tilde{\omega}(t)$ as well as the matrix $A(t, \epsilon)$ be as defined in the proof of Proposition 2.1. Then, the vectors $\omega(t)$ and $\tilde{\omega}(t)$ both lie in the ellipsoid $A(t, \epsilon)(\mathbb{B}_{2p}(2))$.*

Proof. The proof amounts to showing that the vector $\frac{1}{2}(v + \tilde{v})$ lies in $A(\mathbb{B}_p(1))$ and the vector $\frac{1}{2}(v - \tilde{v})$ lies in $A(\epsilon)(\mathbb{B}_p(1))$. Since the first of these items is straightforward to check, we focus only on handling $\frac{1}{2}(v - \tilde{v})$. From the context of the proof of Proposition 2.1, note that $\|\frac{1}{2}(v - \tilde{v})\|_2 \leq \frac{1}{2}\epsilon$. Also, let the s.v.d. of A be written as $A = UDV^\top$, where $U \in \mathbb{R}^{d \times p}$, $D \in \mathbb{R}^{p \times p}$, and $V \in \mathbb{R}^{p \times p}$. Hence, if we let u_1, \dots, u_p denote the columns of U , then

$$\sum_{l=1}^p \frac{(u_l, (v - \tilde{v})/2)^2}{\epsilon^2/4} \leq 1.$$

Meanwhile, considering the expression $v - \tilde{v} = A(w - \tilde{w}) = UDV^\top(w - \tilde{w})$ gives

$$\sum_{l=1}^p \frac{(u_l, (v - \tilde{v})/2)^2}{\sigma_l^2(A)} = \sum_{l=1}^p \frac{\sigma_l^2(A)(e_l, V^\top(w - \tilde{w})/2)^2}{\sigma_l^2(A)} = \|\frac{1}{2}(w - \tilde{w})\|_2^2 \leq 1.$$

Hence, if we let $\check{\sigma}_l(A) = \sqrt{2} \min\{\sigma_l(A), \epsilon/2\}$, then combining leads to

$$\sum_{l=1}^p \frac{(u_l, (v - \tilde{v})/2)^2}{\check{\sigma}_l^2(A)} \leq 1.$$

Likewise, if we let $\check{D} = \text{diag}(\check{\sigma}_1(A), \dots, \check{\sigma}_p(A))$, then the previous display shows that the vector $\check{D}^{-1}U^\top(v - \tilde{v})/2$ lies in the ball $\mathbb{B}_p(1)$, and since V is orthogonal, the vector $x := V\check{D}^{-1}U^\top(v - \tilde{v})/2$ also lies in $\mathbb{B}_p(1)$. In turn, we have

$$\begin{aligned} A(\epsilon)x &= (U\check{D}V^\top)V\check{D}^{-1}U^\top(v - \tilde{v})/2 \\ &= UU^\top(v - \tilde{v})/2 \\ &= (v - \tilde{v})/2 \end{aligned}$$

where the last step follows from the fact that $v - \tilde{v}$ lies in the image of U . Altogether, this means that $(v - \tilde{v})/2$ lies in the ellipsoid $A(\epsilon)(\mathbb{B}_p(1))$. \square

Lemma F.3. *Let $A \in \mathbb{R}^{d \times p}$ satisfy the conditions in Assumption 2.1, and let $\tilde{\mathbb{G}}_n$ be the Gaussian process defined in Section 2 of the main text. Also, let $\epsilon = \epsilon(n) \in (0, 1)$ be a sequence of numbers converging to 0 as $n \rightarrow \infty$, and suppose that $p \gtrsim \epsilon^{-1/\beta}$. Then,*

$$\mathbb{E} \left[\sup_{\rho(\theta, \theta') \leq \epsilon} |\tilde{\mathbb{G}}_n(\theta) - \tilde{\mathbb{G}}_n(\theta')| \right] \gtrsim \epsilon^{(\beta-1/2)/\beta}. \quad (\text{F.7})$$

Proof. Let $\zeta \in \mathbb{R}^d$ be a Gaussian vector drawn from $N(0, I_d)$, and observe that

$$\mathbb{E} \left[\sup_{\rho(\theta, \theta') \leq \epsilon} |\tilde{\mathbb{G}}_n(\theta) - \tilde{\mathbb{G}}_n(\theta')| \right] = \mathbb{E} \left[\sup_{v, v' \in \mathcal{E}, \|v - v'\|_2 \leq \epsilon} \langle \zeta, v - v' \rangle \right], \quad (\text{F.8})$$

where we have used the fact that if $\theta = (v, s)$ and $\theta' = (s', v')$ satisfy $\rho(\theta, \theta') < 1$ then the signs s and s' must be equal, as well as the fact that $\mathcal{E} = -\mathcal{E}$. By the definition of \mathcal{E} , we may express $v - v'$ as

$$v - v' = A(u - u'),$$

for some $u, u' \in \mathbb{S}^{p-1}$. It is a basic fact that when $p \geq 2$, the Minkowski sum $\mathbb{S}^{p-1} + \mathbb{S}^{p-1}$ is equal to $\mathbb{B}_p(2)$, and so

$$\mathbb{E} \left[\sup_{\rho(\theta, \theta') \leq \epsilon} |\tilde{\mathbb{G}}_n(\theta) - \tilde{\mathbb{G}}(\theta')| \right] = \mathbb{E} \left[\sup_{\|w\|_2 \leq 2, \|Aw\|_2 \leq \epsilon} \langle \zeta, Aw \rangle \right]. \quad (\text{F.9})$$

Again, let the s.v.d. of A be written as $A = UDV^\top$, where $U \in \mathbb{R}^{d \times p}$, $D \in \mathbb{R}^{p \times p}$, and $V \in \mathbb{R}^{p \times p}$. Since the p -dimensional random vector $z := U^\top \zeta$ has the Gaussian distribution $N(0, I_p)$ and $V^\top(\mathbb{B}_p(2)) = \mathbb{B}_p(2)$, it follows that

$$\mathbb{E} \left[\sup_{\rho(\theta, \theta') \leq \epsilon} |\tilde{\mathbb{G}}_n(\theta) - \tilde{\mathbb{G}}(\theta')| \right] = \mathbb{E} \left[\sup_{\|w\|_2 \leq 2, \|Dw\|_2 \leq \epsilon} \langle z, Dw \rangle \right]. \quad (\text{F.10})$$

Next, define the ellipsoid

$$\mathcal{E} = \left\{ x \in \mathbb{R}^p \mid \sum_{j=1}^p \frac{x_j^2}{\sigma_j(A)^2 \wedge \epsilon^2} \leq 1 \right\}. \quad (\text{F.11})$$

which can be verified to satisfy the inclusion $\mathcal{E} \subset \{Dw \mid \|w\|_2 \leq 2, \|Dw\|_2 \leq \epsilon\}$, yielding the lower bound

$$\mathbb{E} \left[\sup_{\rho(\theta, \theta') \leq \epsilon} |\tilde{\mathbb{G}}_n(\theta) - \tilde{\mathbb{G}}(\theta')| \right] \geq \mathbb{E} \left[\sup_{x \in \mathcal{E}} \langle z, x \rangle \right]. \quad (\text{F.12})$$

Using the argument in the proof of Proposition 2.5.1 in (Talagrand, 2014), the right hand side can be further lower bounded as

$$\mathbb{E} \left[\sup_{x \in \mathcal{E}} \langle z, x \rangle \right] \gtrsim \left(\sum_{j=1}^p \sigma_j(A)^2 \wedge \epsilon^2 \right)^{1/2}. \quad (\text{F.13})$$

Now consider two cases. If $\sigma_p(A) \geq \epsilon$, then

$$\left(\sum_{j=1}^p \sigma_j(A)^2 \wedge \epsilon^2 \right)^{1/2} \geq (p\epsilon^2)^{1/2} \gtrsim \epsilon^{(\beta-1/2)/\beta},$$

where the second step uses the assumption that $p \gtrsim \epsilon^{-1/\beta}$. Alternatively, in the case when $\sigma_p(A) < \epsilon$, we may let k denote the smallest index $j \in \{1, \dots, p\}$ such that $\sigma_j(A) < \epsilon$, which leads to a lower bound of the form $((k-1)\epsilon^2)^{1/2}$. Due to the assumption that $\sigma_j(A) \asymp j^{-\beta}$, we must have $k \gtrsim \epsilon^{-1/\beta}$, which implies $((k-1)\epsilon^2)^{1/2} \gtrsim \epsilon^{(\beta-1/2)/\beta}$, as needed. \square

Appendix G: Background results

Lemma G.1 (Facts about Orlicz norms). *Let ξ, ξ_1, \dots, ξ_m be any sequence of random variables, and let $q \geq 1$, $x > 0$, and $r \in \{1, 2\}$. Then, there are absolute constants $c, c_0 > 0$ such that the following hold*

$$\|\xi^2\|_{\psi_1} = \|\xi\|_{\psi_2}^2, \quad (\text{G.1})$$

$$\|\xi\|_{\psi_1} \leq c \|\xi\|_{\psi_2}, \quad (\text{G.2})$$

$$\|\xi\|_q \leq c q^{\frac{1}{r}} \|\xi\|_{\psi_r}, \quad (\text{G.3})$$

$$\mathbb{P}(|\xi| \geq x) \leq c \exp \left\{ -c_0 \frac{x^r}{\|\xi\|_{\psi_r}^r} \right\}, \quad (\text{G.4})$$

and

$$\left\| \max_{1 \leq j \leq m} \xi_j \right\|_{\psi_r} \leq c \log(m+1)^{\frac{1}{r}} \max_{1 \leq j \leq m} \|\xi_j\|_{\psi_r}. \quad (\text{G.5})$$

Proof. The first four statements follow from Lemmas 2.7.6 and 2.7.7, as well as Propositions 2.5.2 and 2.7.1 in (Vershynin, 2018). The fifth statement can be found in Lemma 2.2.2 of (van der Vaart and Wellner, 2000). \square

Lemma G.2. Fix any matrix $M \in \mathbb{R}^{p \times d}$ and vector $v \in \mathbb{R}^d$, and let the random vector $Z_1 \in \mathbb{R}^d$ be as in Assumption 2.1. Then, there is a constant $c > 0$ not depending on n such that

$$\| \|MZ_1\|_2 \|_{\psi_2} \leq c \|M\|_F,$$

and

$$\left\| \langle v, Z_1 \rangle^2 - \mathbb{E}[\langle v, Z_1 \rangle^2] \right\|_{\psi_1} \leq c \|v\|_2^2.$$

Proof. The first statement is a slight reformulation of (Vershynin, 2018, Theorem 6.3.2), while the second statement is a special case of (Lopes, Wang and Mahoney, 2019, Lemma 14). \square

Lemma G.3. Suppose that Assumption 2.1 holds. Let $\delta = \delta(n) \in (0, 1)$ be a sequence of numbers converging to 0 as $n \rightarrow \infty$, and let $\Theta(\delta)$ be a minimal δ -net for Θ with respect to the metric ρ . Then,

$$\log \text{card}(\Theta(\delta)) \lesssim \delta^{-1/\beta}.$$

Proof. This result is a consequence of (Tikhomirov (1993) Theorem XVI), (see also Kolmogorov and Tikhomirov (1959)). \square

Appendix H: Simultaneous confidence intervals for population eigenvalues

This appendix is a continuation of Section 3 from the main text. Here, we consider the problem of approximating a collection of random intervals $\mathcal{I}_1, \dots, \mathcal{I}_p$ that satisfy

$$\mathbb{P} \left(\bigcap_{j=1}^p \{ \lambda_j(\Sigma) \in \mathcal{I}_j \} \right) \geq 1 - \alpha. \quad (\text{H.1})$$

Our approach is based on Weyl’s inequality, which ensures that the condition

$$|\lambda_j(\widehat{\Sigma}) - \lambda_j(\Sigma)| \leq \|\widehat{\Sigma} - \Sigma\|_{\text{op}}$$

holds simultaneously for all $j \in \{1, \dots, p\}$, with probability 1. To proceed, let $q_{1-\alpha}$ again denote the $(1-\alpha)$ -quantile of T_n , and let $\mathcal{I}_j = [\lambda_j(\widehat{\Sigma}) \pm q_{1-\alpha}/\sqrt{n}]$. Then, Weyl’s inequality implies that the condition (H.1) must hold. In turn, we may use the bootstrap estimate $\widehat{q}_{1-\alpha}$ to form the approximate intervals defined by $\widehat{\mathcal{I}}_j = [\lambda_j(\widehat{\Sigma}) \pm \widehat{q}_{1-\alpha}/\sqrt{n}]$.

As a way to gain robustness against the effects of eigenvalue multiplicity, the papers (Hall and Hosseini-Nasab, 2006; Hall et al., 2009) also considered an approach of this type — but instead using the Frobenius norm, which can lead to potentially much wider intervals than the operator norm.

The simulation results for the intervals $\widehat{\mathcal{I}}_1, \dots, \widehat{\mathcal{I}}_p$ are given in Table 2. The entries of the table are the observed simultaneous coverage probabilities for a nominal level of 90%, and the simulation settings are the same as those used in Section 3. Although the intervals are conservative due to Weyl’s inequality, they are still close enough to the nominal level to be of practical interest when $\beta > 0.50$. By contrast, when $\beta < 0.50$, the bootstrap again breaks down in the same way that it did for confidence regions (as discussed in Section 3). Lastly, given that these results are obtained in a situation where $\lambda_1(\Sigma)$ has high multiplicity, it is important to reiterate that a naive application of the bootstrap to the individual sample eigenvalues (rather than the operator norm error) is known to work poorly in the presence of such multiplicity. Hence, the user may be willing to tolerate some conservatism from bootstrapping the operator norm, in order to avoid the harms of closely spaced population eigenvalues.

Table 2. Observed simultaneous coverage probabilities for $\widehat{\mathcal{I}}_1, \dots, \widehat{\mathcal{I}}_p$, with $\alpha = 0.1$ and $p = 1,000$.

decay param. β	sample size n			decay param. β	sample size n		
	300	500	700		300	500	700
0.25	100%	100%	100%	0.25	100%	100%	100%
0.50	96.9%	96.5%	97.0%	0.50	97.2%	96.7%	96.5%
0.75	94.5%	94.3%	93.3%	0.75	94.0%	93.9%	93.8%
1.00	93.1%	92.1%	91.5%	1.00	92.9%	91.7%	91.5%
1.25	92.6%	91.7%	91.2%	1.25	92.6%	91.4%	91.4%

(a) $N(0, 1)$ distribution

(b) standardized t_{20} distribution

Appendix I: Additional numerical results on sketching




Figure 1: (Sketching with uniform row sampling.) The plots illustrate bootstrap estimates for the 90% quantile of the error $\|\tilde{A}^\top \tilde{A} - A^\top A\|_{\text{op}}$. The curves are labeled as in the legend of Figure 1 in the main text.

References

- BAI, Z. D. and SILVERSTEIN, J. W. (2010). *Spectral Analysis of Large Dimensional Random Matrices* **20**. Springer.
- CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2015). Comparison and anti-concentration bounds for maxima of Gaussian random vectors. *Probability Theory and Related Fields* **162** 47–70.
- CHERNOZHUKOV, V., CHETVERIKOV, D. and KATO, K. (2017). Central limit theorems and bootstrap in high dimensions. *The Annals of Probability* **45** 2309–2352.
- HALL, P. and HOSSEINI-NASAB, M. (2006). On properties of functional principal components analysis. *Journal of the Royal Statistical Society: Series B* **68** 109–126.
- HALL, P., LEE, Y. K., PARK, B. U. and PAUL, D. (2009). Tie-respecting bootstrap methods for estimating distributions of sets and functions of eigenvalues. *Bernoulli* **15** 380–401.
- KOLMOGOROV, A. N. and TIKHOMIROV, V. M. (1959). ε -entropy and ε -capacity of sets in function spaces. *Uspekhi Matematicheskikh Nauk* **14** 3–86.
- LOPES, M. E., LIN, Z. and MÜLLER, H. G. (2020). Bootstrapping max statistics in high dimensions: Near-parametric rates under weak variance decay and application to functional and multinomial data. *Annals of Statistics* **48** 1214–1229.
- LOPES, M. E., WANG, S. and MAHONEY, M. W. (2019). A bootstrap method for error estimation in randomized matrix multiplication. *Journal of Machine Learning Research* **20** 1–40.
- LOPES, M. E. and YAO, J. (2022). A sharp lower-tail bound for Gaussian maxima with application to bootstrap methods in high dimensions. *Electronic Journal of Statistics* **16** 58–83.
- TALAGRAND, M. (2014). *Upper and Lower Bounds for Stochastic Processes: Modern Methods and Classical Problems*. Springer.

- TIKHOMIROV, V. M. (1993). ϵ -Entropy and ϵ -Capacity of Sets In Functional Spaces. In *Selected Works of A. N. Kolmogorov: Volume III: Information Theory and the Theory of Algorithms* 86–170. Springer.
- VAN DER VAART, A. W. and WELLNER, J. A. (2000). *Weak Convergence and Empirical Processes*. Springer.
- VERSHYNIN, R. (2018). *High-Dimensional Probability: An Introduction with Applications in Data Science*. Cambridge.