
Spiked Covariance Estimation from Modulo-Reduced Measurements

Elad Romanov

Hebrew University of Jerusalem

Or Ordentlich

Hebrew University of Jerusalem

Abstract

Consider the rank-1 spiked model: $\mathbf{X} = \sqrt{\nu}\xi\mathbf{u} + \mathbf{Z}$, where ν is the spike intensity, $\mathbf{u} \in \mathbb{S}^{k-1}$ is an unknown direction and $\xi \sim \mathcal{N}(0, 1)$, $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Motivated by recent advances in analog-to-digital conversion, we study the problem of recovering $\mathbf{u} \in \mathbb{S}^{k-1}$ from n i.i.d. *modulo-reduced* measurements $\mathbf{Y} = [\mathbf{X}] \bmod \Delta$, focusing on the high-dimensional regime ($k \gg 1$). We develop and analyze an algorithm that, for most directions \mathbf{u} and $\nu = \text{poly}(k)$, estimates \mathbf{u} to high accuracy using $n = \text{poly}(k)$ measurements, provided that $\Delta \gtrsim \sqrt{\log k}$. Up to constants, our algorithm accurately estimates \mathbf{u} at the smallest possible Δ that allows (in an information-theoretic sense) to recover \mathbf{X} from \mathbf{Y} . A key step in our analysis involves estimating the probability that a line segment of length $\approx \sqrt{\nu}$ in a random direction \mathbf{u} passes near a point in the lattice $\Delta\mathbb{Z}^k$. Numerical experiments show that the developed algorithm performs well even in a non-asymptotic setting.

1 Introduction

We consider the problem of estimating a spiked covariance matrix from Gaussian modulo-folded measurements. Let $\mathbf{u} \in \mathbb{S}^{k-1}$ be an unknown direction, and $\nu > 0$ be the signal-to-noise (SNR) ratio. Consider the spiked covariance matrix

$$\Sigma = \nu\mathbf{u}\mathbf{u}^\top + \mathbf{I}, \quad (1)$$

and denote $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \Sigma)$. Equivalently, one may write

$$\mathbf{X} = \sqrt{\nu}\xi\mathbf{u} + \mathbf{Z}, \quad (2)$$

where ξ, \mathbf{Z} have $\mathcal{N}(0, 1)$ entries; the one dimensional component $\sqrt{\nu}\xi\mathbf{u}$ is often thought of as the “signal”, whereas \mathbf{Z} is thought of as “noise”. In this paper, we consider the problem of estimating \mathbf{u} from n independent and modulo-reduced measurements of \mathbf{X} . Let $\Delta > 0$ be the *dynamic range*, and for $X \in \mathbb{R}$, denote the modulo operation by

$$Y = [X] \bmod \Delta \in \left[-\frac{1}{2}\Delta, \frac{1}{2}\Delta\right), \quad (3)$$

so that Y is the unique number in the half-open interval such that $X - Y \in \Delta\mathbb{Z}$. For a vector $\mathbf{X} \in \mathbb{R}^k$, $\mathbf{Y} = [\mathbf{X}] \bmod \Delta$ is defined by modulo-reducing each coordinate separately. In the setup we consider, one is given n independent copies of \mathbf{Y} , denoted by $\mathbf{y}_1, \dots, \mathbf{y}_n$, and wishes to estimate the unknown direction $\mathbf{u} \in \mathbb{S}^{k-1}$. Throughout, we denote by $\mathbf{x}_1, \dots, \mathbf{x}_n$ independent copies of \mathbf{X} , such that $\mathbf{y}_i = [\mathbf{x}_i] \bmod \Delta$. See Figure 1 for a graphical illustration, in $k = 2$ dimensions.

Motivation. Our motivation for considering this problem is driven by recent developments in signal processing. Analog-to-digital converters (ADCs), devices that convert analog (continuous) signals into digital (discrete, e.g. bits) signals, are an essential component in virtually all modern communication devices. From a mathematical perspective, ADCs are set out to solve *essentially* the following problem: Given a vector-valued random variable $\mathbf{X} \in \mathbb{R}^k$, find a quantizer (binning scheme), $\phi : \mathbb{R}^k \rightarrow \mathbb{R}^k$, with a finite range, $|\text{range}(\phi)| \leq 2^B$ (B being the allowable representation length in bits), so as to minimize the quantization error: $\mathbb{E}\|\mathbf{X} - \phi(\mathbf{X})\|^2$. Quantization is an extensively studied problem, and its fundamental limits (in information theory: “rate-distortion theory”), under setups of varying generality, are largely understood; see, e.g., [Cover and Thomas, 2012, Gersho and Gray, 2012].

Optimal vector quantizers (that achieve the fundamental limits) typically involve rather complicated constructions, that depend intricately on the *exact* statistics of \mathbf{X} . Since ADCs are implemented in mixed analog-digital circuits, sophisticated vector quantizers are prohibitive and the de-




Figure 1: A typical problem instance in $k = 2$ dimensions. Top: A point cloud, corresponding to $n = 5000$ i.i.d. samples from $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \nu \mathbf{u} \mathbf{u}^\top + \mathbf{I})$, for $\nu = 10^4$ and some fixed $\mathbf{u} \in \mathbb{S}^{k-1}$. Bottom: The modulo-reduced point cloud, with $\Delta = 80$.

sign is often restricted to architectures of a scalar product form: $\phi(\mathbf{X}) = (\phi_1(X_1), \dots, \phi_k(X_k))$ where $\phi_1, \dots, \phi_k : \mathbb{R} \rightarrow \mathbb{R}$. Perhaps the simplest – and most popular – architecture in practice is a uniform scalar quantizer, determined by its dynamic range $\Delta > 0$ and bit-rate $b = B/k$. The quantizer divides the interval $[-\Delta/2, \Delta/2]$ into 2^b intervals of equal size $\Delta 2^{-b}$, so that the scalar quantizer $\phi_1 = \dots = \phi_k = \phi$ maps $X \in \mathbb{R}$ to its closest interval center. Note that for this scheme to attain quantization error that vanishes as the quantization rate b increases, the dynamic range must be $\Delta \gtrsim \max_{1 \leq i \leq k} \sqrt{\text{Var}(X_i)}$.¹

Although simple to implement, the uniform quantizer can be pronouncedly sub-optimal for vector-valued signals, as it cannot leverage the cross-coordinate correlations that often occur in real-world applications. An important use-case in digital communications is Massive MIMO, where typically the number of users is much smaller than the number of receive antennas; this results in signals \mathbf{X} that have strong cross-entry correlations.² Thus, a quantization scheme that can ex-

¹If X falls inside $[-\Delta/2, \Delta/2]$, then $|X - \phi(X)| \leq \Delta 2^{-b}$. For the expected error to be of the same order, the probability of a *saturation*, namely $|X| > \Delta/2$, has to be small.

²A multi-user MIMO channel is modeled by

plot these statistical inter-dependencies, while retaining the simplicity of the uniform quantizer, is highly desired.

A recently proposed architecture, “modulo-ADCs” [Ordentlich et al., 2018], attempts to address these issues. Their idea is rather simple: do not *truncate* \mathbf{X} onto $[-\Delta/2, \Delta/2]^k$, as the uniform quantizer does; instead, apply **modulo-reduction** $\mathbf{Y} = [\mathbf{X}] \bmod \Delta \in [-\Delta/2, \Delta/2]^k$ and then quantize as before. For this idea to work, one clearly need some means of “unwrapping” \mathbf{X} from \mathbf{Y} (with high probability). When the coordinates of \mathbf{X} are independent and unimodal, with the mode at 0 (for example, a centered Gaussian), it is easy to see that the best estimator for \mathbf{X} from \mathbf{Y} (in the sense of error probability) is just $\tilde{\mathbf{X}} = \mathbf{Y}$. Thus, a coordinate X_i cannot be recovered once it saturates the ADC dynamic range, $|X_i| \geq \Delta/2$; so to consistently undo the modulo, one needs $\Delta \gtrsim \max_{1 \leq i \leq k} \sqrt{\text{Var}(X_i)}$, and the scheme has no advantage over the standard uniform quantizer. It turns out, however, that when \mathbf{X} has strong correlations, it is often possible to consistently unwrap at *substantially* smaller values of Δ ; see next section.

1.1 Related work

We start with very brief background on the spiked model, Eq. (2). In the high-dimensional statistics ($k \approx n$, $k, n \rightarrow \infty$) literature, the spiked model was popularized by [Johnstone, 2001], who studied the largest eigenvalue of the sample covariance matrix $\hat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$. Subsequent advances in random matrix theory [Baik et al., 2005, Paul, 2007] characterized the behavior of PCA (namely, the relation between the principal components of Σ and its empirical counterpart $\hat{\Sigma}$) rather precisely. Since then, a vast literature on the spiked model (and variations thereof) has emerged – which we make no pretense to survey here; as an entry point, geared towards statisticians, see [Wainwright, 2019, Chapter 8]. We cite the following minimax lower bound for the spike estimation problem (without modulo-reduction) [Wain-

$\mathbf{X} = \mathbf{H} \mathbf{S} + \mathbf{Z}$, where k the number of receive antennas, m is the number of transmitting users, each equipped with a single antenna, and $\mathbf{H} = [\mathbf{h}_1 \cdots \mathbf{h}_m] \in \mathbb{R}^{k \times m}$ is the channel matrix ($\mathbf{h}_i \in \mathbb{R}^k$ represents the channel gains from transmitter i to the receiver). The vector $\mathbf{S} = [S_1, \dots, S_m] \in \mathbb{R}^m$ represents the transmissions of the m users, and $\mathbf{Z} \in \mathbb{R}^k$ is white noise. The “signal” part, $\mathbf{H} \mathbf{S}$, lives inside an m -dimensional subspace in \mathbb{R}^k , where typically $m \ll k$. One example, among many, for the extreme case of rank $m = 1$ MIMO (corresponding to model (2) exactly), is in low-earth orbit (LEO) communications, where a phased array receiver is used to track a rapidly moving satellite.

wright, 2019, Example 15.19]:

$$\|\mathbf{u} - \hat{\mathbf{u}}\| \gtrsim \left(\frac{\sqrt{1+\nu}}{\nu} \sqrt{\frac{k}{n}} \right) \wedge 1. \quad (4)$$

In the regime $k = O(n)$, this rate is attained, up to prefactors, by PCA ($\hat{\mathbf{u}}$ taken to be the largest eigenvector of $\hat{\Sigma}$), see [Wainwright, 2019, Corollary 8.7].

Moving on, there has recently been a great deal of activity in the signal processing community around recovery from modulo-reduced measurements [Bhandari et al., 2017, Ordentlich et al., 2018, Bhandari et al., 2018, Graf et al., 2019, Romanov and Ordentlich, 2019, Bhandari and Kraemer, 2019, Bhandari et al., 2020, Bhandari et al., 2021, Weiss et al., 2021]. Most relevant to this paper is a line of works dealing with recovery from modulo-reduced measurements, and motivated by the modulo-ADC architecture described before [Ordentlich and Erez, 2017, Ordentlich et al., 2018]. The setting is this: the source is *Gaussian* $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, with Σ some covariance matrix (not necessarily spiked); one observes modulo-reduced measurements $\mathbf{Y} = [\mathbf{X}] \bmod \Delta$, and wishes to recover \mathbf{X} itself (with high probability). How large should Δ be so that consistent recovery is possible, in an information-theoretic sense? When it *is* possible, how could one do so *practically*? (Assuming Σ is known? And when it is *not*?) The answers, it turns out, depend rather intricately on the *diophantine* properties of the matrix Σ .

Let us start with the fundamental limits. A simple observation is that when $\Delta \gtrsim \sqrt{\log k} \cdot \max_{1 \leq \ell \leq k} \sqrt{\Sigma_{\ell, \ell}}$ ($\Sigma_{\ell, \ell}$ being the variance of the ℓ -th coordinate), one has $\mathbf{X} = \mathbf{Y}$ with high probability, so that consistent recovery is straightforward. It is easy to see that when \mathbf{X} is white, in other words $\Sigma \propto \mathbf{I}$, this requirement is in fact tight. For general Σ , one may readily show that the maximum a posteriori probability (MAP) estimator for \mathbf{X} given \mathbf{Y} is

$$\hat{\mathbf{X}}_{\text{MAP}}(\mathbf{Y} = \mathbf{y}) = \underset{\mathbf{x} : \mathbf{x} - \mathbf{y} \in \Delta \mathbb{Z}^k}{\operatorname{argmin}} \mathbf{x}^\top \Sigma^{-1} \mathbf{x}, \quad (5)$$

that is, one needs to minimize a quadratic form over the coset of \mathbf{y} . Searching over the coset directly (and consequently, computing $\hat{\mathbf{X}}_{\text{MAP}}$ exactly) is not computationally tractable, in all but the simplest cases; nonetheless, since $\hat{\mathbf{X}}_{\text{MAP}}$ is optimal in the sense of error probability, its performance characterizes the information-theoretic limits of the problem. The latter has a rather elegant geometric interpretation. Let $\mathcal{L} = \Delta \Sigma^{-1/2} \mathbb{Z}^k \subseteq \mathbb{R}^k$ be the lattice generated by $\Delta \Sigma^{-1/2} \in \mathbb{R}^{k \times k}$, and let $V_0 \subseteq \mathbb{R}^k$ be the Voronoi cell of $\mathbf{0} \in \mathcal{L}$. Then [Romanov and Ordentlich, 2021] the success probability of (5) is the Gaussian measure

of V_0 :

$$p_{\text{MAP}} = \Pr \left(\hat{\mathbf{X}}_{\text{MAP}} = \mathbf{X} \right) = \Pr_{\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} (\mathbf{Z} \in V_0).$$

As a corollary, it is not hard to show that³ $\Delta \lesssim |\Sigma|^{1/2k}$ implies that $p_{\text{MAP}} = o(1)$; see also Proposition 2. When the lattice \mathcal{L} is a *uniformly random lattice*, sampled, up to normalization, from the Haar measure over $\text{SL}_k(\mathbb{R})/\text{SL}_k(\mathbb{Z})$ (also called Haar-Siegel measure), one can show that V_0 is with high probability “sufficiently ball-like”, so that $\Delta \gtrsim |\Sigma|^{1/2k}$ is also a sufficient condition. Random lattices have played a prominent role in the lattice coding literature [Zamir, 2014], which is closely related to the present line of work. An important point is that “natural” random matrix ensembles, such as the spiked ensemble (1) with $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$, **do not** correspond to the Siegel-Haar measure on the space of lattices. In [Domanovitz and Erez, 2017] the authors demonstrate that certain orthogonally-invariant ensembles, that arise in channel coding theory, indeed allow for consistent recovery with Δ not much larger than $|\Sigma|^{1/2k}$. In particular, for the spiked ensemble (1), they show that with high probability over $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$, the error probability is small whenever $\Delta \gtrsim C(k) |\Sigma|^{1/2k} = C(k) (1 + \nu)^{1/2k}$, where $C(k)$ grows exponentially fast in k , but does not depend on the SNR ν .

As for practical recovery algorithms, let us start by assuming Σ is **known**. As already mentioned, computing the MAP estimator directly is intractable; instead [Ordentlich and Erez, 2017] proposed to use a sub-optimal estimator, the so-called *Integer Forcing* (IF) decoder. The idea is to find an invertible integer matrix $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_k]$ so as to minimize the maximal variance:

$$\mathbf{A}_{\text{IF}} = \underset{\substack{\mathbf{A} \in \mathbb{Z}^{k \times k}, \\ \text{rank}(\mathbf{A}) = k}}{\operatorname{argmin}} \max_{1 \leq \ell \leq k} \mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell. \quad (6)$$

The corresponding maximal deviation, $m_k(\Sigma) = \max_{1 \leq \ell \leq k} \sqrt{\mathbf{a}_\ell^\top \Sigma \mathbf{a}_\ell}$, is called the k th successive minimum of the lattice $\Sigma^{1/2} \mathbb{Z}^k$. Since $[\mathbf{A}\mathbf{Y}] \bmod \Delta = [\mathbf{A}\mathbf{X}] \bmod \Delta$, one can reliably recover \mathbf{X} from \mathbf{Y} , using $\hat{\mathbf{X}}_{\text{IF}} = \mathbf{A}_{\text{IF}}^{-1} ([\mathbf{A}_{\text{IF}}\mathbf{Y}] \bmod \Delta)$, whenever $\Delta \gtrsim m_k(\Sigma) \sqrt{\log k}$. Of course, to compute the IF decoder, Eq. (6), one clearly needs to know Σ .⁴ In

³ $|\Sigma|$ denotes the determinant of Σ .

⁴An important caveat is that Eq.(6) is actually a computationally hard problem, and may only be solved exactly for very small k . In practice, one usually solves this *approximately*, using a lattice reduction algorithm, like the Lenstra-Lenstra-Lovász (LLL) algorithm [Lenstra et al., 1982]. Observe that if one a priori restricts the minimization to $\mathbf{A} \in \text{SL}_k(\mathbb{Z})$, then Eq. (6) is equivalent to finding a shortest basis for the lattice $\Sigma^{1/2} \mathbb{Z}^k$.

some applications, for example in wireless communications (where Σ depends on the channel matrix, which rapidly changes over time) [Tse and Viswanath, 2005], this is *not* a reasonable assumption. In [Romanov and Ordentlich, 2021], the authors propose a **blind** unwrapping algorithm, that does not know Σ beforehand, in a setting where one needs to simultaneously unwrap many i.i.d. signals $\mathbf{y}_1, \dots, \mathbf{y}_n$. A natural step towards that end is to estimate Σ from the (modulo-reduced) data, from which the integer-forcing decoder (6) could be computed. Alas, directly computing the maximum likelihood estimator (MLE) of Σ from modulo-reduced measurements is not computationally feasible. Instead, they propose an algorithm which iteratively alternates between 1) A covariance estimation step, where a certain “proxy” of Σ is estimated; 2) An integer forcing decoder, computed from that proxy; the idea is to gradually “whiten” the entire dataset, effectively computing the IF decoder “in small steps”. They prove a result of the following form: when the error of the informed IF decoder (Eq. (6)) is *small enough*, then the error of the adaptive algorithm is *essentially comparable* to it, up to dimension-dependent prefactors. However, their algorithm is only suited to rather modest k , as seen both in the analysis (the prefactors are exponential in k) and the numerical experiments. The problem lies with their covariance estimation procedure, whose performance breaks down rapidly as the dimension increases. This is the starting point for the present paper.

Our contributions. We propose a computationally tractable algorithm to estimate the spike $\mathbf{u} \in \mathbb{S}^{k-1}$ from modulo-reduced measurements, under the spiked covariance model Eq. (1) and in high dimension k . We show that for *most* directions \mathbf{u} (formally: with high probability over $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$), estimation is possible with $n = \text{poly}(k)$ samples, under essentially the smallest Δ (up to constants) that allows for consistent unwrapping. Thus, in this setting, we *provably* overcome the curse of dimensionality suffered by the algorithm of [Romanov and Ordentlich, 2021]. While we do not directly tackle the unwrapping problem, note that in applications where the SNR ν is approximately known, the algorithm readily yields an estimate for Σ , from whence one could compute the IF decoder (6). Our numerical experiments below show that this method attains an unwrapping error probability that is not far from that of the informed IF decoder.

Notation. For sequences a_k, b_k we use the following standard notation: $a_k \vee b_k = \max\{a_k, b_k\}$, $a_k \wedge b_k = \min\{a_k, b_k\}$. By $a_k \lesssim b_k$ we mean that $a_k \leq Cb_k$ for some *universal* constant $C > 0$; we write $a_k \approx b_k$ whenever both $a_k \lesssim b_k$ and $b_k \lesssim a_k$. We also use big-O

notation; if M is a parameter, we use $a_k = O_M(b_k)$ to signify that the constants might depend on M . For a vector $\mathbf{v} \in \mathbb{R}^k$, $\|\mathbf{v}\|$ denotes its ℓ_2 (Euclidean) norm.

2 Proposed method

An observation. Our algorithm is based on the following observation: the eigen-structure of the covariance matrix of \mathbf{X} is preserved when truncated onto a ball. Set $R > 0$ a truncation radius. For $\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \Sigma)$, let \mathbf{X}_{Ball} be its spherically-truncated version: for $\mathcal{S} \subseteq \mathbb{R}^k$,

$$\begin{aligned} \Pr(\mathbf{X}_{\text{Ball}} \in \mathcal{S}) &= \Pr(\mathbf{X} \in \mathcal{S} | \mathbf{X} \in \mathcal{B}(\mathbf{0}, R)) \\ &= \frac{\Pr(\mathbf{X} \in \mathcal{S} \cap \mathcal{B}(\mathbf{0}, R))}{\Pr(\mathbf{X} \in \mathcal{B}(\mathbf{0}, R))}. \end{aligned}$$

Observe that $\mathbb{E}[\mathbf{X}_{\text{Ball}}] = \mathbf{0}$, since \mathbf{X}_{Ball} is symmetric. Denote the covariance by $\Sigma_{\text{Ball}} = \mathbb{E}[\mathbf{X}_{\text{Ball}}\mathbf{X}_{\text{Ball}}^\top]$.

Proposition 1 *Let Σ be any covariance matrix, with (orthonormal) eigenvectors $\mathbf{u}_1, \dots, \mathbf{u}_k \in \mathbb{S}^{k-1}$ and corresponding eigenvalues $\lambda_1 \geq \dots \geq \lambda_k \geq 0$. Then*

1. *The basis $\mathbf{u}_1, \dots, \mathbf{u}_k$ diagonalizes Σ_{Ball} . Denote $g_1, \dots, g_k \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$; the respective eigenvalues are:*

$$\mu_i = \mathbb{E} \left[\lambda_i g_i^2 \left| \sum_{i=1}^k \lambda_i g_i^2 \leq R^2 \right. \right]. \quad (7)$$

2. *The ordering is preserved: $\mu_1 \geq \dots \geq \mu_k$.*

Proposition 1 is not new by any means. It has appeared before in [Palombi et al., 2012], which considered covariance estimation from spherically-truncated Gaussian measurements (see also discussion later in this section). The proof of Item 1 is rather trivial; for completeness, we provide a short proof (Appendix A). Item 2 is considerably less so; we refer to [Palombi et al., 2012, Proposition 3.3] for the details.

The algorithm. We draw inspiration from Figure 1. For “most” directions $\mathbf{u} \in \mathbb{S}^{k-1}$, the points $\{\mathbf{y}_i\}_{i=1}^n$ are arranged, essentially, in parallel and separated stripes. The central stripe (that crosses the origin) consists of points that have not undergone folding, $\mathbf{y}_i = \mathbf{x}_i$. Picking only points inside a *small enough* ball $\mathcal{Y}_{\text{Ball}} = \{\mathbf{y}_i\}_{i=1}^n \cap \mathcal{B}(\mathbf{0}, R)$, we therefore obtain, approximately, an i.i.d. sample from \mathbf{X}_{Ball} . By Proposition 1, the leading eigenvector of Σ_{Ball} is \mathbf{u} , and therefore PCA with $\mathcal{Y}_{\text{Ball}}$ should yield a consistent estimator (as $n \rightarrow \infty$). See Figure 2 for a graphical illustration; and Algorithm 1 for a formal description.

Algorithm 1: The proposed algorithm

Input: Samples: $\{\mathbf{y}_i\}_{i=1}^n$; Radius: $R > 0$.

1 Pick samples inside ball:

$$\mathcal{K} = \{i \in [n] : \mathbf{y}_i \in \mathcal{B}(\mathbf{0}, R)\}.$$

 2 Form sample covariance: $\widehat{\Sigma} = \sum_{i \in \mathcal{K}} \mathbf{y}_i \mathbf{y}_i^\top$.

return $\widehat{\mathbf{u}}$, principal eigenvector of $\widehat{\Sigma}$.




Figure 2: An illustration of Algorithm 1. Red: points inside a small ball; Yellow: the central stripe ($\mathbf{x}_i = \mathbf{y}_i$).

The truncation radius. In spite of its seeming simplicity, the behavior of Algorithm 1 depends *drastically* on the truncation radius R . Its choice should balance between two opposing effects. On the one hand, the algorithm uses effectively $\mathbb{E}|\mathcal{K}| = n \cdot p_{\text{Ball}}$ measurements for estimation (where $p_{\text{Ball}} := \Pr(\mathbf{X} \in \mathcal{B}(\mathbf{0}, R))$), so R cannot be too small; on the other hand, we need to take only (or mostly) points from the central stripe, so R cannot be too large. Let us start with an observation: when $R/\sqrt{k} < 1$, p_{Ball} is exponentially small in k , so the algorithm requires $n \gtrsim 1/p_{\text{Ball}} = \exp(\Omega(k))$ measurements. Consequently, to (potentially) overcome the curse of dimensionality one must set $R/\sqrt{k} > 1$. In that case, $p_{\text{Ball}} \approx 1 \wedge (\sqrt{R}/\nu)$ (Lemma 2), so that for a large spike, $\nu = \omega(R)$, $n \gtrsim \sqrt{\nu/R}$; therefore, we shall henceforth restrict our attention to $\nu = \text{poly}(k)$. As we have said, R cannot be chosen too large, and in general there is a rather delicate **tradeoff** between the parameters Δ, R, ν and the direction $\mathbf{u} \in \mathbb{S}^{k-1}$ itself. Our main result, Theorem 1, says, roughly, the following: there *is* a choice $R = \Theta(\sqrt{k})$ such that for *most* directions, if $\Delta \gtrsim \sqrt{\log k}$ and $\nu = \text{poly}(k)$, then Algorithm 1 estimates \mathbf{u} with small error from only $n = \text{poly}(k)$ measurements.

On estimating ν . In this paper, we restrict our attention to estimating only the direction $\mathbf{u} \in \mathbb{S}^{k-1}$ (and not ν). We mention two potential strategies for estimating ν , not pursued here further due to space con-

straints:

- [Palombi et al., 2012] studies covariance estimation from spherically-truncated Gaussian measurements. Relying on Proposition 1, they prove that the mapping $\boldsymbol{\lambda} \mapsto \boldsymbol{\mu}$ between the true and truncated eigenvalues is invertible, and propose a fixed point iteration to recover $\boldsymbol{\lambda}$ from $\boldsymbol{\mu}$ (given *exactly*, without noise). Our proposed algorithm computes an estimate of Σ_{Ball} ; computing error bounds for the method of [Palombi et al., 2012], applied to $\widehat{\Sigma}$, is potentially challenging, especially in the regime where $\nu \gg R^2 = \Theta(k)$, where the mapping $\boldsymbol{\lambda} \mapsto \boldsymbol{\mu}$ is necessarily badly conditioned.
- Since only one eigenvalue of Σ is unknown, a more natural approach is to invert (numerically) the mapping $\nu \mapsto p_{\text{Ball}}(\nu) = \Pr(\mathbf{X} \in \mathcal{B}(\mathbf{0}, R))$, which is strictly decreasing.

2.1 Main results

The following is our main result. It shows that for most directions $\mathbf{u} \in \mathbb{S}^{k-1}$, the error attained by Algorithm 1 can be made quite small, with n reasonably controlled, and assuming only $\Delta \gtrsim \sqrt{\log k}$. To make the presentation lighter, we focus exclusively on the regime where the spike is not small, $\nu = \Omega(1)$ (which is also practically more interesting). Below, $\widehat{\mathbf{u}}$ denotes the largest eigenvector of $\widehat{\Sigma}$, with the sign ambiguity resolved by assuming that $\langle \mathbf{u}, \widehat{\mathbf{u}} \rangle \geq 0$.

Theorem 1 Fix a constant $M > 12$, and set $R = \Theta(\sqrt{k})$ as in (11). There is a universal constant C_* and a set $\mathcal{U}_M \subseteq \mathbb{S}^{k-1}$ with

$$\Pr_{\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})}(\mathbf{u} \in \mathcal{U}_M) = 1 - O_M(k^{-10}),$$

such that whenever $\Delta \geq C_*(M\sqrt{\log k} \vee \nu^{\frac{1}{2(k-1)}})$ and $\mathbf{u} \in \mathcal{U}_M$, the following error bounds hold (depending on the magnitude of ν), with probability $1 - O(k^{-10})$:

1. Assume that $1 \leq \nu \leq k$ and $n \gtrsim \log k$. Then

$$\|\mathbf{u} - \widehat{\mathbf{u}}\| \lesssim \frac{1}{\sqrt{\nu}} \left(\frac{k}{n} \vee \sqrt{\frac{k}{n}} \right) + \nu^{-1} k^{-M^2+12}. \quad (8)$$

2. Assume that $k \leq \nu \lesssim k^{2M^2-21}$. Then

$$\|\mathbf{u} - \widehat{\mathbf{u}}\| \lesssim \frac{\sqrt{\nu}}{n} + \sqrt{\frac{1}{n} \sqrt{\frac{\nu}{k}}} + \nu^{1/2} k^{-M^2+10.5}. \quad (9)$$

Discussion. Let us start with the small-spike regime, $1 \leq \nu \leq k$. The first term in Eq. (8) is, up

to prefactors, the error rate for PCA without modulo-folding, see e.g. [Wainwright, 2019, Chapter 8]; in particular, when $k = O(n)$, it matches the minimax lower bound Eq. (4). The second term corresponds to error incurred by erroneously taking “bad” points \mathbf{y}_i , that do not belong on the central stripe. By taking $\Delta \gtrsim \sqrt{\log k}$ large enough, this term can be made to decay arbitrarily (polynomially) fast as $k \rightarrow \infty$. We note that in this regime, the consequences of Theorem 1 are, in fact, rather unsurprising: if $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$, then with high probability, $\|\mathbf{u}\|_\infty \lesssim \sqrt{\log k/k}$. Since $\Delta \gtrsim \sqrt{\log k}$, the cube $[-\frac{1}{2}\Delta, \frac{1}{2}\Delta]^k$ contains a segment $\{t\mathbf{u} : t \in [-L, L]\}$ of length $2L \gtrsim \Delta/\|\mathbf{u}\|_\infty \gtrsim \sqrt{k}$; consequently, a large fraction of $\mathbf{x}_1, \dots, \mathbf{x}_n$ are actually themselves already inside the cube, since the “typical length” of the projection along \mathbf{u} , $|\langle \mathbf{u}, \mathbf{X} \rangle|$, is $\lesssim \sqrt{\nu}$ (the standard deviation).

The “interesting” regime is $\nu \gg k$. Note that unlike in the small-spike regime, here the error, Eq. (9), *increases* as ν grows. Moreover, the magnitude of ν has to be constrained by Δ : $\nu \lesssim k^{2M^2-21}$ (the constant 21 is itself not particularly important, and can be improved). Thus, to retain the scaling $\Delta \sim \sqrt{\log k}$, ν has to grow at most polynomially with k ; in that case, note that the term $\nu^{1/(2(k-1))}$ in the bound for Δ is always negligible. Furthermore, note that n has to scale at least as $n \gtrsim \sqrt{\nu}$, which anyhow precludes the practicality of the algorithm when ν is super-polynomial, regardless of the third term.

Let us try to get some intuition for the particular form of the bounds (8), (9), by considering a simplified setting, where one had direct access to all the measurements \mathbf{x}_i that lie inside the ball $\mathcal{B}(\mathbf{0}, R)$, and used them to perform PCA. There are roughly (Lemmas 1, 2)

$$\tilde{n} \approx p_{\text{Ball}} n \approx (1 \wedge \sqrt{k/\nu}) n$$

such measurements. By Proposition 1, the population covariance Σ_{Ball} is spiked, and one can show (Lemma 7) that the effective spike is


$$\tilde{\nu} \approx k \wedge \nu.$$

(Note that \mathbf{X}_{Ball} lies inside $\mathcal{B}(\mathbf{0}, R)$, and consequently $\lambda_1(\Sigma_{\text{Ball}}) \leq R^2 = \Theta(k)$). In particular, note that when $\nu = \omega(k)$, \tilde{n} decreases with ν but $\tilde{\nu}$ cannot grow further to compensate for this; this is the reason why the error in Eq. (9) degrades with ν . Now, assuming that \tilde{n} is large enough (this point is a little subtle, since we let $\tilde{\nu}$ grow as well), the error is bounded like


$$\|\mathbf{u} - \hat{\mathbf{u}}\| \lesssim \frac{1}{\sqrt{\tilde{\nu}}} \left(\frac{k}{\tilde{n}} \vee \sqrt{\frac{k}{\tilde{n}}} \right),$$

using “standard” bounds for PCA (e.g. [Wainwright, 2019, Chapter 8]). Plugging in the above estimates

for $\tilde{n}, \tilde{\nu}$ recovers the first terms in Eqs. (8), (9). The challenging part of the analysis (and our main technical contribution) is to control the last term: namely, show that for most directions \mathbf{u} , when $\Delta \gtrsim \sqrt{\log k}$, the contribution of the erroneously picked (“bad”) points is indeed very small with high probability.



(a) The estimation error, $\|\mathbf{u} - \hat{\mathbf{u}}\|$, as ν changes.



(b) Performance of the informed vs. blind integer forcing decoder, as $\Delta/\sqrt{\log k}$ changes.

Figure 3: Numerical results

Experiments. We demonstrate the validity and relevance of our results through numerical experiments:

- In Figure 3a we study the behavior of the error, $\|\mathbf{u} - \hat{\mathbf{u}}\|$, as the spike magnitude ν changes. For several values, $k = 50, 100, 150, 200$, we have set $\Delta = 16\sqrt{\log k}$, $n = k^2$ and varied $\nu = k^\alpha$ for an exponent $\alpha \in [0, 5.6]$; each point on the graph is the average error across $T = 2400$ repetition. We observe that as k increases, scaling $\Delta \propto \sqrt{\log k}$ indeed suffices for estimation. Moreover, we see that for small spikes, $\alpha < 1$, the error decreases as α increases, whereas when $\alpha > 1$ the error increases; this is consistent with Theorem 1.
- In Figure 3b we apply our algorithm as an in-

intermediate step for blind unwrapping. We set $k = 150$, $n = k^2 \approx 2 \cdot 10^4$, $\nu = k^3 \approx 3 \cdot 10^6$ and vary $\Delta = 2^\delta \sqrt{\log k}$. At every working point, we compute the error rate, namely the fraction of erroneously recovered samples $\hat{p}_e = \frac{1}{n} \sum_{i=1}^n \mathbb{1}_{\mathbf{x}_i \neq \hat{\mathbf{x}}_i}$, of the informed IF decoder (Eq. (6)), the blind IF decoder (computed from $\hat{\Sigma} = \nu \hat{\mathbf{u}} \hat{\mathbf{u}}^\top + \mathbf{I}$), and the trivial decoder $\hat{\mathbf{X}} = \mathbf{Y}$. For each method, δ is increased in jumps of 0.25, until the point where $p_e \leq 10^{-4}$; each point on the graph is the average of $T = 200$ repetitions (so, overall, $nT \approx 10^6$ single recovery trials). We see, for this particular setup, a gap of around $\delta = 2$ bits between the the informed and blind decoders, and of about $\delta = 5$ bits between the blind and trivial decoders. To put things in context, a hypothetical quantization scheme based around modulo-folding and the blind decoder could save up to 5 bits per coordinate (so $5k = 750$ bits overall) compared to the uniform quantizer (both designed so that the probability of a saturation is $\leq 10^{-4}$).

By Theorem 1, when $\nu = \text{poly}(k)$, the condition $\Delta \gtrsim \sqrt{\log k}$ ensure that one can estimate most directions $\mathbf{u} \in \mathbb{S}^{k-1}$ with $n = \text{poly}(k)$ measurements. Recall that the present problem was motivated by the modulo-unfolding problem (which is a *harder* problem). It turns out that for the latter, the condition $\Delta \gtrsim \sqrt{\log k}$ is actually necessary. Thus, if one's goal is to solve the unwrapping problem (e.g. for implementing modulo-ADCs), and to that end estimates the covariance as an *intermediate step*, then our algorithm succeeds with *essentially* the smallest allowable dynamic range. We show the following (see Appendix B for the proof):

Proposition 2 *Suppose that there exists $\hat{\mathbf{X}} = \hat{\mathbf{X}}(\mathbf{Y})$, with $\Pr(\mathbf{X} = \hat{\mathbf{X}}) \geq 0.1$. Then $\Delta \gtrsim (\sqrt{\log k} \vee \nu^{\frac{1}{2k}})$.*

We remark in passing that, once we have obtained an estimate for \mathbf{u} , and consequently for Σ , using Algorithm 1, we may use it to unwrap the measurements $\mathbf{x}_1, \dots, \mathbf{x}_n$ (using, e.g., the IF decoder). Having unwrapped the samples, we can use standard methods (e.g., PCA) to get an improved estimate of \mathbf{u} . We do not pursue this option here for two reasons: 1) The performance of such an algorithm depends on the unwrapping error probability, which is difficult to analyze. In particular, unwrapping errors could have a disastrous effect on the estimation error; 2) Our primary motivation for estimating \mathbf{u} in the first place was to perform unwrapping. To that end, once we obtain an estimate of \mathbf{u} with accuracy sufficient for unwrapping, further improvements are of limited interest.

3 Analysis

In this section, we give a proof outline for our main result, Theorem 1. In the interest of space, the proofs of most technical lemmas are relegated to the Appendix.

For $\delta \in (0, 1)$, set

$$\begin{aligned} z_2(\delta) &:= \left(k + 2\sqrt{k \log(1/\delta)} + 2 \log(1/\delta) \right)^{1/2}, \\ z_\infty(\delta) &:= \sqrt{2 \log k} + \sqrt{2 \log(2/\delta)}, \end{aligned} \quad (10)$$

so that for $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$ (see Appendix H, Lemma 17),

$$\Pr(\|\mathbf{Z}\| \geq z_2(\delta)) \leq \delta, \quad \Pr(\|\mathbf{Z}\|_\infty \geq z_\infty(\delta)) \leq \delta.$$

Going forward, we fix a truncation radius:

$$R = 2\sqrt{k} + z_2(0.1) = \Theta(\sqrt{k}). \quad (11)$$

This *particular* choice is rather arbitrary. One could carry out the analysis with any $R = C\sqrt{k}$ for $C > 1$; this would only change the constants in the bounds.

3.1 High level view

Divide the pairs $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i \in [n]}$ into groups. Denote by

$$\begin{aligned} \mathcal{K} &= \{i \in [n] : \mathbf{y}_i \in \mathcal{B}(\mathbf{0}, R)\}, \\ \mathcal{K}_{\text{Ball}} &= \{i \in [n] : \mathbf{x}_i \in \mathcal{B}(\mathbf{0}, R)\}, \end{aligned}$$

respectively the points that were picked by Algorithm 1, and those for which $\mathbf{x}_i \in \mathcal{B}(\mathbf{0}, R)$. Note that, conditioned on $i \in \mathcal{K}_{\text{Ball}}$, $\mathbf{x}_i \stackrel{d}{=} \mathbf{X}_{\text{Ball}}$; hence $\{\mathbf{x}_i\}_{i \in \mathcal{K}_{\text{Ball}}}$ is an i.i.d. sample from \mathbf{X}_{Ball} . Observe also that $\mathcal{K}_{\text{Ball}} \subseteq \mathcal{K}$, since $\|\mathbf{X} \bmod \Delta\| \leq \|\mathbf{X}\|$. We denote by $\mathcal{K}_{\text{Good}} \subseteq \mathcal{K}_{\text{Ball}}$ the subset of measurements for which $\mathbf{x}_i = \mathbf{y}_i$, in other words, such that $\mathbf{x}_i \in [-\frac{1}{2}\Delta, \frac{1}{2}\Delta]^k$ to begin with. The measurements in $\mathcal{K}_{\text{Bad}} = \mathcal{K} \setminus \mathcal{K}_{\text{Good}}$, will be called *bad*. We have

$$\mathcal{K}_{\text{Good}} \subseteq \mathcal{K}_{\text{Ball}} \subseteq \mathcal{K} = \mathcal{K}_{\text{Good}} \sqcup \mathcal{K}_{\text{Bad}} \subseteq [n].$$

Now, the sample covariance,

$$\hat{\Sigma}_{\mathbf{X}} := \frac{1}{|\mathcal{K}_{\text{Ball}}|} \sum_{i \in \mathcal{K}_{\text{Ball}}} \mathbf{x}_i \mathbf{x}_i^\top \quad (12)$$

is a consistent (as $|\mathcal{K}_{\text{Ball}}| \rightarrow \infty$) estimator for Σ_{Ball} , whose largest eigenvector is \mathbf{u} (Proposition 1). Consequently, PCA yields a consistent estimator for the unknown direction. Alas, the set $\mathcal{K}_{\text{Ball}}$ is not directly observable, so the algorithm uses $\mathcal{K} \supseteq \mathcal{K}_{\text{Ball}}$ instead.⁵

$$\hat{\Sigma} = \frac{1}{|\mathcal{K}_{\text{Ball}}|} \sum_{i \in \mathcal{K}} \mathbf{y}_i \mathbf{y}_i^\top. \quad (13)$$

⁵Note that we have normalized by $|\mathcal{K}_{\text{Ball}}|$, which is unknown. This is done for the sake of convenience in the analysis; the eigenvectors, of course, are not affected.

This injects *additional* error into the covariance estimation process, in two ways. First, the covariance is computed using the \mathbf{y}_i -s instead of the \mathbf{x}_i -s (the latter are unknown); we have $\mathbf{y}_i = \mathbf{x}_i$ only for $i \in \mathcal{K}_{\text{Good}}$, which may be a strict subset of \mathcal{K} . Second, we use additional samples, on top of $\mathcal{K}_{\text{Ball}}$: the points in $\mathcal{K} \setminus \mathcal{K}_{\text{Ball}}$ necessarily come from the wrong distribution. Set

$$\varepsilon_{\text{CovEst}} := \|\widehat{\boldsymbol{\Sigma}}_{\mathbf{X}} - \boldsymbol{\Sigma}_{\text{Ball}}\|, \quad \varepsilon_{\text{Pick}} := \|\widehat{\boldsymbol{\Sigma}}_{\mathbf{X}} - \widehat{\boldsymbol{\Sigma}}\|,$$

so that $\|\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_{\text{Ball}}\| \leq \varepsilon_{\text{CovEst}} + \varepsilon_{\text{Pick}}$. A bound on this operator norm yields, by standard eigenvector perturbation results, a bound on $\|\mathbf{u} - \widehat{\mathbf{u}}\|$. Note: $\varepsilon_{\text{CovEst}}$ is simply the statistical estimation error in estimating $\boldsymbol{\Sigma}_{\text{Ball}} = \text{Cov}(\mathbf{X}_{\text{Ball}})$ from $|\mathcal{K}_{\text{Ball}}|$ i.i.d. measurements; the other term, $\varepsilon_{\text{Pick}}$, is the error induced through picking *erroneous measurements*. We shall bound each error term separately.

3.2 The covariance estimation error

We start with $\varepsilon_{\text{CovEst}}$; the argument is quite standard. First, we show that with high probability, $|\mathcal{K}_{\text{Ball}}|$ is reasonably large. Denote

$$p_{\text{Ball}} = p_{\text{Ball}}(\nu; R) = \Pr(\mathbf{X} \in \mathcal{B}(\mathbf{0}, R)), \quad (14)$$

so that $|\mathcal{K}_{\text{Ball}}| \sim \text{Binomial}(p_{\text{Ball}}, n)$. Controlling $|\mathcal{K}_{\text{Ball}}|$ is straightforward using, e.g., Chernoff's inequality (Appendix H, Lemma 19):

Lemma 1 For any $\mathbf{u} \in \mathbb{S}^{k-1}$, for universal $c_1, c_2 > 0$,

$$\Pr(|\mathcal{K}_{\text{Ball}}| \leq c_1 p_{\text{Ball}} n \mid \mathbf{u}) \leq 2e^{-c_2 p_{\text{Ball}} n}.$$

Next is an (tight, for large ν) estimate for p_{Ball} ; the (short) proof is relegated to Appendix G.1:

Lemma 2

$$0.9 \operatorname{erf}\left(\sqrt{\frac{2k}{\nu}}\right) \leq p_{\text{Ball}} \leq \operatorname{erf}\left(\sqrt{\frac{2k + z_2(0.1)/2}{1 + \nu}}\right),$$

where, recall, the error function is defined by

$$\operatorname{erf}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{-t^2} dt = \Pr(|\xi| \leq \sqrt{2}x). \quad (15)$$

Note that Lemma 2 implies that $p_{\text{Ball}} = \Theta(1)$ when $\nu = O(k)$, whereas $p_{\text{Ball}} = \Theta(\sqrt{k/\nu})$ when $\nu = \omega(k)$; in other words, $p_{\text{Ball}} = \Theta(1 \wedge \sqrt{k/\nu})$.

Error bounds for covariance estimation rely on the concentration properties of the data. Thus, we need to show that \mathbf{X}_{Ball} “inherits” the favorable properties of the underlying Gaussian vector \mathbf{X} . We start with a general Lemma, whose proof appears in Appendix G.2:

Lemma 3 For every convex function $g : \mathbb{R}^k \rightarrow \mathbb{R}$,

$$\mathbb{E}[g(\mathbf{X}_{\text{Ball}})] \leq \mathbb{E}[g(\mathbf{X})].$$

The proof of Lemma 3 relies on the Gaussian correlation inequality. As an important corollary, it allows us to control the sub-Gaussian and sub-exponential norms of \mathbf{X}_{Ball} ; see Appendix C, Lemma 9.

Recall that by Proposition 1, $\boldsymbol{\Sigma}_{\text{Ball}}$ is a spiked covariance matrix with largest eigenvector \mathbf{u} :

$$\lambda_1(\boldsymbol{\Sigma}_{\text{Ball}}) > \lambda_2(\boldsymbol{\Sigma}_{\text{Ball}}) = \dots = \lambda_k(\boldsymbol{\Sigma}_{\text{Ball}}),$$

so that applying Lemma 3 (with $g(\mathbf{X}) = \langle \mathbf{u}_i, \mathbf{X} \rangle^2$),

$$\lambda_1(\boldsymbol{\Sigma}_{\text{Ball}}) \leq R^2 \wedge (1 + \nu), \quad \lambda_2(\boldsymbol{\Sigma}_{\text{Ball}}) \leq 1.$$

The rest of the analysis proceeds along rather standard lines, as in e.g. [Wainwright, 2019, Section 8.2.2]; the full details are given in Appendix C. We prove:

Lemma 4 Suppose that $p_{\text{Ball}} n \gtrsim \log k$. Then, with probability $1 - O(k^{-10})$,

$$\begin{aligned} \varepsilon_{\text{CovEst}} &\lesssim (k \wedge (1 + \nu)) \sqrt{\frac{\log k}{p_{\text{Ball}} n}} \\ &\quad + \sqrt{(k \wedge (1 + \nu))} \left(\frac{k}{p_{\text{Ball}} n} \vee \sqrt{\frac{k}{p_{\text{Ball}} n}} \right). \end{aligned}$$

3.3 The sample picking error

Decompose

$$\sum_{i \in \mathcal{K}} \mathbf{y}_i \mathbf{y}_i^\top = \sum_{i \in \mathcal{K}_{\text{Good}}} \mathbf{x}_i \mathbf{x}_i^\top + \sum_{i \in \mathcal{K}_{\text{Bad}}} \mathbf{y}_i \mathbf{y}_i^\top,$$

so that

$$\begin{aligned} \varepsilon_{\text{Pick}} &= \frac{1}{|\mathcal{K}_{\text{Ball}}|} \cdot \left\| \sum_{i \in \mathcal{K}_{\text{Bad}}} \mathbf{y}_i \mathbf{y}_i^\top - \sum_{i \in \mathcal{K}_{\text{Ball}} \setminus \mathcal{K}_{\text{Good}}} \mathbf{x}_i \mathbf{x}_i^\top \right\| \\ &\leq \frac{2|\mathcal{K}_{\text{Bad}}| \cdot R^2}{|\mathcal{K}_{\text{Ball}}|} \lesssim k \cdot \frac{|\mathcal{K}_{\text{Bad}}|}{|\mathcal{K}_{\text{Ball}}|}. \end{aligned} \quad (16)$$

Above, we used: $\|\mathbf{x}_i\| \leq R$ for $i \in \mathcal{K}_{\text{Ball}}$; $\|\mathbf{y}_i\| \leq R$ for $i \in \mathcal{K}$; and $|\mathcal{K}_{\text{Ball}} \setminus \mathcal{K}_{\text{Good}}| \leq |\mathcal{K} \setminus \mathcal{K}_{\text{Good}}| \leq |\mathcal{K}_{\text{Bad}}|$.

The next Lemma is one of our main technical results. It states that for *most* directions $\mathbf{u} \in \mathbb{S}^{k-1}$, the probability that a pair (\mathbf{X}, \mathbf{Y}) is *bad*, meaning that $\mathbf{Y} \in \mathcal{B}(\mathbf{0}, R)$ but $\mathbf{X} \neq \mathbf{Y}$, is overwhelmingly small provided that $\Delta \gtrsim \sqrt{\log k} \vee \nu^{\frac{1}{2(k-1)}}$:

Lemma 5 Fix a constant $M \geq 1$. There is a universal $C_* > 0$ and a subset $\mathcal{U}_M \subseteq \mathbb{S}^{k-1}$ with

$$\Pr_{\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})}(\mathbf{u} \in \mathcal{U}_M) = 1 - O_M(k^{-10}),$$

such that if $\Delta \geq C_*(M\sqrt{\log k} \vee \nu^{\frac{1}{2(k-1)}})$, then for all $\mathbf{u} \in \mathcal{U}_M$,

$$\Pr((\mathbf{X}, \mathbf{Y}) \text{ is bad} \mid \mathbf{u}) \leq k^{-M^2}.$$

The proof appears in Appendix D. The key idea is to reduce the problem into a question in geometric probability: whether a randomly rotated line segment is far away from all non-zero lattice points.

Lemma 5 readily gives the following bound; the details are given in Appendix G.3:

Lemma 6 Assume the setup of Lemma 5, with $M > \sqrt{12}$, $\mathbf{u} \in \mathcal{U}_M$, $\Delta \geq C_*(M\sqrt{\log k} \vee \nu^{\frac{1}{2(k-1)}})$. Suppose that $p_{\text{Ball}} n \gtrsim \log k$. With probability $1 - O(k^{-10})$:

$$\varepsilon_{\text{Pick}} \lesssim \frac{k^{-M^2+12}}{p_{\text{Ball}}}.$$

3.4 Concluding the analysis

So far, we have shown that $\|\widehat{\Sigma} - \Sigma_{\text{Ball}}\|$ is small with high probability. To deduce that their largest eigenvectors are close as well (using eigenvector perturbation results), we first need to show that the spectral gap of Σ_{Ball} is large. We prove the following in Appendix E:

Lemma 7 There are universal C_1, C_2 such that for $\nu \geq e^{-C_1 k}$,

$$\lambda_1(\Sigma_{\text{Ball}}) \geq 1 + C_2(k \wedge \nu).$$

Consequently, $\lambda_1(\Sigma_{\text{Ball}}) - \lambda_2(\Sigma_{\text{Ball}}) \gtrsim (k \wedge \nu)$.

The proof of Theorem 1 follows by combining our bounds thus far. The details appear in Appendix F.

Acknowledgements

This work was supported in part by ISF under Grant 1791/17 and in part by the GENESIS Consortium via the Israel Ministry of Economy and Industry. The work of Elad Romanov was supported in part by an Einstein-Kaye fellowship from the Hebrew University of Jerusalem.

References

- [Artstein-Avidan et al., 2015] Artstein-Avidan, S., Giannopoulos, A., and Milman, V. D. (2015). *Asymptotic geometric analysis, Part I*, volume 202. American Mathematical Soc.
- [Baik et al., 2005] Baik, J., Arous, G. B., and Pécché, S. (2005). Phase transition of the largest eigenvalue for nonnull complex sample covariance matrices. *The Annals of Probability*, 33(5):1643–1697.
- [Bhandari and Krahmer, 2019] Bhandari, A. and Krahmer, F. (2019). On identifiability in unlimited sampling. In *2019 13th International conference on Sampling Theory and Applications (SampTA)*, pages 1–4.
- [Bhandari et al., 2021] Bhandari, A., Krahmer, F., and Poskitt, T. (2021). Unlimited sampling from theory to practice: Fourier-prony recovery and prototype adc. *arXiv preprint arXiv:2105.05818*.
- [Bhandari et al., 2017] Bhandari, A., Krahmer, F., and Raskar, R. (2017). On unlimited sampling. In *2017 International Conference on Sampling Theory and Applications (SampTA)*, pages 31–35. IEEE.
- [Bhandari et al., 2018] Bhandari, A., Krahmer, F., and Raskar, R. (2018). Unlimited sampling of sparse sinusoidal mixtures. In *2018 IEEE International Symposium on Information Theory (ISIT)*, pages 336–340.
- [Bhandari et al., 2020] Bhandari, A., Krahmer, F., and Raskar, R. (2020). On unlimited sampling and reconstruction. *IEEE Transactions on Signal Processing*.
- [Boucheron et al., 2013] Boucheron, S., Lugosi, G., and Massart, P. (2013). *Concentration inequalities: A nonasymptotic theory of independence*. Oxford university press.
- [Cover and Thomas, 2012] Cover, T. M. and Thomas, J. A. (2012). *Elements of Information Theory*. John Wiley & Sons.
- [Domanovitz and Erez, 2017] Domanovitz, E. and Erez, U. (2017). Outage behavior of integer forcing with random unitary pre-processing. *IEEE Transactions on Information Theory*, 64(4):2774–2790.
- [Gersho and Gray, 2012] Gersho, A. and Gray, R. M. (2012). *Vector quantization and signal compression*, volume 159. Springer Science & Business Media.
- [Graf et al., 2019] Graf, O., Bhandari, A., and Krahmer, F. (2019). One-bit unlimited sampling. In *ICASSP 2019 - 2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 5102–5106.
- [Johnstone, 2001] Johnstone, I. M. (2001). On the distribution of the largest eigenvalue in principal components analysis. *Annals of statistics*, pages 295–327.
- [Latała and Matlak, 2017] Latała, R. and Matlak, D. (2017). Royen’s proof of the gaussian correlation inequality. In *Geometric aspects of functional analysis*, pages 265–275. Springer.

- [Laurent and Massart, 2000] Laurent, B. and Massart, P. (2000). Adaptive estimation of a quadratic functional by model selection. *Annals of Statistics*, pages 1302–1338.
- [Lenstra et al., 1982] Lenstra, A. K., Lenstra, H. W., and Lovász, L. (1982). Factoring polynomials with rational coefficients. *Mathematische annalen*, 261(ARTICLE):515–534.
- [Ordentlich and Erez, 2017] Ordentlich, O. and Erez, U. (2017). Integer-forcing source coding. *IEEE Transactions on Information Theory*, 63(2):1253–1269.
- [Ordentlich et al., 2018] Ordentlich, O., Tabak, G., Hanumolu, P. K., Singer, A. C., and Wornell, G. W. (2018). A modulo-based architecture for analog-to-digital conversion. *IEEE journal of selected topics in signal processing*, 12(5):825–840.
- [Palombi et al., 2012] Palombi, F., Toti, S., and Filipini, R. (2012). Numerical reconstruction of the covariance matrix of a spherically truncated multinormal distribution. *Journal of Probability and Statistics*, 2017.
- [Paul, 2007] Paul, D. (2007). Asymptotics of sample eigenstructure for a large dimensional spiked covariance model. *Statistica Sinica*, pages 1617–1642.
- [Romanov and Ordentlich, 2019] Romanov, E. and Ordentlich, O. (2019). Above the Nyquist rate, modulo folding does not hurt. *IEEE Signal Processing Letters*, 26(8):1167–1171.
- [Romanov and Ordentlich, 2021] Romanov, E. and Ordentlich, O. (2021). Blind unwrapping of modulo reduced gaussian vectors: Recovering MSBs from LSBs. *IEEE Transactions on Information Theory*, 67(3):1897–1919.
- [Royen, 2014] Royen, T. (2014). A simple proof of the gaussian correlation conjecture extended to multivariate gamma distributions. *Far East Journal of Theoretical Statistics*, 48:139–145.
- [Tse and Viswanath, 2005] Tse, D. and Viswanath, P. (2005). *Fundamentals of wireless communication*. Cambridge university press.
- [Vershynin, 2018] Vershynin, R. (2018). *High-dimensional probability: An introduction with applications in data science*, volume 47. Cambridge university press.
- [Wainwright, 2019] Wainwright, M. J. (2019). *High-dimensional statistics: A non-asymptotic viewpoint*, volume 48. Cambridge University Press.
- [Weiss et al., 2021] Weiss, A., Huang, E., Ordentlich, O., and Wornell, G. W. (2021). Blind modulo analog-to-digital conversion. *arXiv preprint arXiv:2108.08937*.
- [Zamir, 2014] Zamir, R. (2014). *Lattice Coding for Signals and Networks: A Structured Coding Approach to Quantization, Modulation, and Multiuser Information Theory*. Cambridge University Press.

A Proof of Proposition 1, Item 1

Decompose \mathbf{X} along the principal components:

$$\mathbf{X} = \sum_{i=1}^k \sqrt{\lambda_i} g_i \mathbf{u}_i,$$

where $g_1, \dots, g_k \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$. Then,

$$\Sigma_{\text{Ball}} = \mathbb{E} [\mathbf{X}\mathbf{X}^\top \mid \|\mathbf{X}\|^2 \leq R^2] = \sum_{i,j=1}^k \mathbb{E} \left[\sqrt{\lambda_i \lambda_j} g_i g_j \mathbf{u}_i \mathbf{u}_j^\top \mid \sum_{i=1}^k \lambda_i g_i^2 \leq R^2 \right].$$

Now, observe that the cross terms, $i \neq j$, are zero, since conditioning onto the ball preserves the symmetry $(g_i, g_j) \mapsto (-g_i, g_j)$. Thus,

$$\Sigma_{\text{Ball}} = \sum_{i=1}^k \mathbb{E} \left[\lambda_i g_i^2 \mathbf{u}_i \mathbf{u}_i^\top \mid \sum_{i=1}^k \lambda_i g_i^2 \leq R^2 \right] =: \sum_{i=1}^k \mu_i \mathbf{u}_i \mathbf{u}_i^\top,$$

and so the claim is proved.

B Proof of Proposition 2

For brevity, define

$$p(\nu; \mathbf{u}) = \Pr_{\mathbf{X} \sim \mathcal{N}(\mathbf{0}, \nu \mathbf{u} \mathbf{u}^\top + \mathbf{I})} (\mathbf{X} = \widehat{\mathbf{X}}_{\text{MAP}}), \quad (17)$$

where $\widehat{\mathbf{X}}_{\text{MAP}}$ is the MAP estimator of \mathbf{X} from $\mathbf{Y} = [\mathbf{X}] \bmod \Delta$; in other words, $p(\nu; \mathbf{u})$ is the success probability of the MAP estimator at SNR ν with spike direction \mathbf{u} . Recalling that the MAP estimator is optimal in the sense of error probability, it is clear that to prove Proposition 2, it suffices to show that $p(\nu; \mathbf{u}) \geq 0.1$ implies that $\Delta \gtrsim \nu^{1/2k} \vee \sqrt{\log k}$.

We start with a simple observation:

Lemma 8 *The function $\nu \mapsto p(\nu; \mathbf{u})$ is decreasing.*

Proof. For any $\nu \geq 0$, denote $\mathbf{X}_\nu \sim \mathcal{N}(\mathbf{0}, \nu \mathbf{u} \mathbf{u}^\top + \mathbf{I})$ and $\mathbf{Y}_\nu = [\mathbf{X}_\nu] \bmod \Delta$, and let $g_\nu : [-\Delta/2, \Delta/2)^k \rightarrow \mathbb{R}^k$ be a deterministic function such that $g_\nu(\mathbf{Y}_\nu)$ is the MAP estimator for \mathbf{X}_ν . Fix any $\tau \geq 0$; we shall now construct a suboptimal estimator for \mathbf{X}_ν given \mathbf{Y}_ν , based on $g_{\nu+\tau}(\cdot)$. The idea is simple: we generate known noise $\mathbf{N} \sim \mathcal{N}(\mathbf{0}, \tau \mathbf{u} \mathbf{u}^\top)$ and set $\mathbf{Y}' = [\mathbf{Y}_\nu + \mathbf{N}] \bmod \Delta$, which also equals $\mathbf{Y}' = [\mathbf{X}_\nu + \mathbf{N}] \bmod \Delta$. Note that $(\mathbf{X}_{\nu+\tau}, \mathbf{Y}_{\nu+\tau}) \stackrel{d}{=} (\mathbf{X}_\nu + \mathbf{N}, \mathbf{Y}')$. Considering the sub-optimal estimator $\widehat{\mathbf{X}}' = g_{\nu+\tau}(\mathbf{Y}') - \mathbf{N}$ for \mathbf{X}_ν , we conclude,

$$p(\nu; \mathbf{u}) \geq \Pr(\mathbf{X}_\nu = \widehat{\mathbf{X}}') = \Pr(\mathbf{X}_\nu = g_{\nu+\tau}(\mathbf{Y}') - \mathbf{N}) = \Pr(\mathbf{X}_{\nu+\tau} = g_{\nu+\tau}(\mathbf{Y}_{\nu+\tau})) = p(\nu + \tau; \mathbf{u}).$$

■

Let us start by showing $\Delta \gtrsim \sqrt{\log k}$. By Lemma 8, the assumptions of Proposition 2 imply that $p(0; \mathbf{u}) \geq 0.1$. Now, it is easy to see that when $\nu = 0$, $\widehat{\mathbf{X}}_{\text{MAP}} = \mathbf{Y}$; in this case, the problem simply decouples across the different coordinates. Thus,

$$0.1 \leq p(0; \mathbf{u}) = \Pr(\mathbf{X} = \mathbf{Y}) = \Pr(|X_i| \leq \Delta/2 \text{ for all } 1 \leq i \leq k) = \left\{ \text{erf} \left(\Delta/2^{3/2} \right) \right\}^k,$$

and therefore $\text{erf}(\Delta/2^{3/2}) \geq 1 - ck^{-1}$ for some universal $c \geq 0$. Clearly, then, $\Delta = \omega(1)$ for large k , so by the standard estimate $1 - \text{erf}(x) \gtrsim e^{-x^2}/x$ (for large x), we get $e^{-c_1 \Delta^2}/\Delta \lesssim 1/k$ hence $\Delta \gtrsim \sqrt{\log k}$.

It remains to show $\Delta \gtrsim \nu^{1/2k}$. To that end, we will use a simple geometric characterization of the MAP estimator, following [Romanov and Ordentlich, 2021]. Let

$$\mathcal{L} = \Delta \Sigma^{-1/2} \mathbb{Z}^k = \left\{ \Delta \Sigma^{-1/2} \mathbf{t} : \mathbf{t} \in \mathbb{Z}^k \right\}$$

be the lattice generated by the matrix $\Delta \mathbf{\Sigma}^{-1/2}$, and denote by $V_0 \subset \mathbb{R}^k$ the Voronoi cell of $\mathbf{0} \in \mathcal{L}$ (that is, all points $\mathbf{a} \in \mathbb{R}^k$ whose closest lattice point is $\mathbf{0} \in \mathcal{L}$). By [Romanov and Ordentlich, 2021, Section III, Eq. (38)], the success probability of the MAP estimator is

$$\Pr \left(\widehat{\mathbf{X}}^{\text{MAP}}(\mathbf{Y}) = \mathbf{X} \right) = \Pr_{\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)} (\mathbf{Z} \in V_0).$$

Now, it is a well-known fact that V_0 is a convex symmetric set, with $\text{vol}_k(V_0) = |\Delta \mathbf{\Sigma}^{-1/2}| = \Delta^k / |\mathbf{\Sigma}|^{1/2}$. Let r_0 be the *effective radius* of \mathcal{L} , defined by

$$\text{vol}_k(\mathcal{B}(\mathbf{0}, r_0)) = \text{vol}_k(V_0) \implies r_0 = \frac{\Delta}{|\mathbf{\Sigma}|^{1/2k} \mathcal{V}_k^{1/k}},$$

(\mathcal{V}_k denotes the volume of the Euclidean unit ball). Recall that among all convex bodies with a given (finite) volume, a ball has the largest Gaussian measure. Thus,

$$0.1 \leq \Pr_{\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)} (\mathbf{Z} \in V_0) \leq \Pr_{\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)} (\mathbf{Z} \in \mathcal{B}(\mathbf{0}, r_0)) = \Pr_{\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)} (\|\mathbf{Z}\|^2 \leq r_0^2).$$

Note that $\|\mathbf{Z}\|^2 \sim \chi_{(k)}^2$, which concentrates around k with ‘‘typical’’ deviations of order $O(\sqrt{k})$ (see, e.g., Lemma 17). This gives $r_0 \gtrsim \sqrt{k}$, so

$$\Delta \gtrsim |\mathbf{\Sigma}|^{1/2k} \left(\mathcal{V}_k^{1/k} \sqrt{k} \right) = (1 + \nu)^{1/2k} \left(\mathcal{V}_k^{1/k} \sqrt{k} \right) \gtrsim (1 + \nu)^{1/2k},$$

where the last inequality follows from Stirling’s approximation: $\mathcal{V}_k \stackrel{k \rightarrow \infty}{\approx} \frac{1}{\sqrt{k\pi}} \left(\frac{2\pi e}{k} \right)^{k/2}$, and therefore $\mathcal{V}_k^{1/k} \sqrt{k} = \sqrt{2\pi e} + o(1) = \Theta(1)$.

C Proof of Lemma 4

Decompose \mathbf{X}_{Ball} along the principal components:

$$\mathbf{X}_{\text{Ball}} = w\mathbf{u} + \mathbf{p},$$

where w is the projection along \mathbf{u} and \mathbf{p} is the orthogonal complement. Note that, while w and \mathbf{p} are uncorrelated, they are **not** independent (as was the case without truncation, for a Gaussian vector) since we condition on $w^2 + \|\mathbf{p}\|^2 \leq R^2$. Also, recalling the ‘‘spiky’’ structure of $\mathbf{\Sigma}_{\text{Ball}}$,

$$\mathbb{E}(w^2) = \lambda_1(\mathbf{\Sigma}_{\text{Ball}}), \quad \mathbb{E}[\mathbf{p}\mathbf{p}^\top] = \mathbf{\Sigma}_{\text{Ball}} - \lambda_1(\mathbf{\Sigma}_{\text{Ball}}) \cdot \mathbf{u}\mathbf{u}^\top = \lambda_2(\mathbf{\Sigma}_{\text{Ball}}) \cdot (\mathbf{I} - \mathbf{u}\mathbf{u}^\top). \quad (18)$$

Condition on $N = |\mathcal{K}_{\text{Ball}}|$, and denote for convenience $\mathcal{K}_{\text{Ball}} = \{1, \dots, N\}$, so that $\{\mathbf{x}_i\}_{i=1}^N$ are i.i.d. measurements from \mathbf{X}_{Ball} . Write

$$\begin{aligned} \widehat{\mathbf{\Sigma}}_{\mathbf{X}} &= \frac{1}{N} \sum_{i=1}^N (w_i \mathbf{u} + \mathbf{p}_i) (w_i \mathbf{u} + \mathbf{p}_i)^\top \\ &= \left(\frac{1}{N} \sum_{i=1}^N w_i^2 \right) \mathbf{u}\mathbf{u}^\top + \left(\frac{1}{N} \sum_{i=1}^N w_i \mathbf{p}_i \right) \mathbf{u}^\top + \left(\frac{1}{N} \left(\sum_{i=1}^N w_i \mathbf{p}_i \right) \mathbf{u}^\top \right)^\top + \frac{1}{N} \sum_{i=1}^N \mathbf{p}_i \mathbf{p}_i^\top, \end{aligned} \quad (19)$$

so that the error can be decomposed as $\left\| \widehat{\mathbf{\Sigma}}_{\mathbf{X}} - \mathbf{\Sigma}_{\text{Ball}} \right\| \leq \varepsilon_1 + 2\varepsilon_2 + \varepsilon_3$, with

$$\begin{aligned} \varepsilon_1 &:= \left| \frac{1}{N} \sum_{i=1}^N w_i^2 - \mathbb{E}[w^2] \right|, \\ \varepsilon_2 &:= \left\| \frac{1}{N} \sum_{i=1}^N w_i \mathbf{p}_i \right\|, \\ \varepsilon_3 &:= \left\| \frac{1}{N} \sum_{i=1}^N \mathbf{p}_i \mathbf{p}_i^\top - \mathbb{E}[\mathbf{p}\mathbf{p}^\top] \right\|. \end{aligned} \quad (20)$$

We first show that \mathbf{X}_{Ball} inherits the sub-Gaussian concentration properties of \mathbf{X} . We denote, respectively, the sub-Gaussian and sub-exponential norms by $\|\cdot\|_{\psi_2}$ and $\|\cdot\|_{\psi_1}$. For a quick reminder on these norm (and Orlicz norms in general), see Definition 1 and Lemma 20.

Lemma 9 *We have*

$$\|\mathbf{p}\|_{\psi_2} \lesssim 1, \quad \|w^2\|_{\psi_1} \lesssim k \wedge (1 + \nu).$$

and

$$\|w\mathbf{p}\|_{\psi_1} \lesssim \sqrt{k \wedge (1 + \nu)}.$$

Proof. Let $\psi : [0, \infty) \rightarrow [0, \infty)$ be convex and increasing, and let $g : \mathbb{R} \rightarrow \mathbb{R}$ be such that $\mathbf{x} \mapsto |g(\mathbf{x})|$ is convex. Observe that $x \mapsto \psi(|g(x)|)$ is convex, and consequently, by Lemma 3, $\|g(\mathbf{X}_{\text{Ball}})\|_{\psi} \leq \|g(\mathbf{X})\|_{\psi}$, where $\|\cdot\|_{\psi}$ is the Orlicz ψ -norm (see Definition 1). Consequently,

$$\|w^2\|_{\psi_1} \leq \|\langle \mathbf{u}, \mathbf{X} \rangle^2\|_{\psi_1} \lesssim 1 + \nu, \quad \|\mathbf{p}\|_{\psi_2} \leq \|(\mathbf{I} - \mathbf{u}\mathbf{u}^\top)\mathbf{X}\|_{\psi_2} \lesssim 1.$$

Furthermore, using Lemma 20, Items 1 and 5,

$$\|w^2\|_{\psi_1} = \|w\|_{\psi_2}^2 \lesssim \|w\|_{\infty}^2 \leq R^2 \lesssim k.$$

This proves the first two bound. As for the last one,

$$\|w\mathbf{p}\|_{\psi_1} = \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \|w\langle \mathbf{v}, \mathbf{p} \rangle\|_{\psi_1} \stackrel{(\star)}{\leq} \|w\|_{\psi_2} \cdot \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \|\langle \mathbf{v}, \mathbf{p} \rangle\|_{\psi_2} = \|w\|_{\psi_2} \|\mathbf{p}\|_{\psi_2} \stackrel{(\star\star)}{\lesssim} \sqrt{k \wedge (1 + \nu)},$$

where (\star) follows from Lemma 20, Item 2, and $(\star\star)$ follows from Lemma 20, Item 1, and the first part of this proof. ■

We now bound the errors $\varepsilon_1, \varepsilon_2, \varepsilon_3$, again conditioned on $N = |\mathcal{K}_{\text{Ball}}|$:

Lemma 10 *Assume that $N \gtrsim \log k$. Then, with probability $1 - O(k^{-10})$,*

$$\varepsilon_1 \lesssim (k \wedge (1 + \nu)) \sqrt{\frac{\log k}{N}}.$$

Proof. By the centralization Lemma (Lemma 20, Item 3) and Lemma 9,

$$\|w_i^2 - \mathbb{E}[w_i^2]\|_{\psi_1} \leq \|w_i^2\|_{\psi_1} \lesssim k \wedge (1 + \nu).$$

By Bernstein's inequality (Lemma 21),

$$\Pr(\varepsilon_1 \geq t) \leq 2 \exp[-c_1 N (\delta \wedge \delta^2)], \quad \delta := \frac{t}{k \wedge (1 + \nu)}.$$

Set $t = \sqrt{\frac{10}{c_1}} (k \wedge (1 + \nu)) \sqrt{\frac{\log k}{N}}$. Then whenever $N \geq \frac{10}{c_1} \log k$, the probability is $\leq 2k^{-10}$. ■

Lemma 11 *With probability $1 - 2e^{-\Omega(k)}$,*

$$\varepsilon_2 \lesssim \sqrt{k \wedge (1 + \nu)} \left(\frac{k}{N} \vee \sqrt{\frac{k}{N}} \right).$$

Proof. Set $\mathbf{q} = \frac{1}{N} \sum_{i=1}^N w_i \mathbf{p}_i$, and observe that $\mathbb{E}[\mathbf{q}] = \mathbf{0}$, since w_i and \mathbf{p}_i are uncorrelated. We want to bound $\|\mathbf{q}\|$ with high probability; to that end, we use a standard ε -net argument, executed in detail for the sake of completeness. Using [Vershynin, 2018, Corollary 4.2.13], fix a $1/2$ -net \mathcal{N} of \mathbb{S}^{k-1} of size $|\mathcal{N}| \leq 5^k$. Let $\tilde{\mathbf{v}} \in \mathcal{N}$ be a member of the net, such that $\left\| \frac{\mathbf{q}}{\|\mathbf{q}\|} - \tilde{\mathbf{v}} \right\| \leq 1/2$. Now,

$$\|\mathbf{q}\| = \left\langle \mathbf{q}, \frac{\mathbf{q}}{\|\mathbf{q}\|} \right\rangle = \left\langle \mathbf{q}, \frac{\mathbf{q}}{\|\mathbf{q}\|} - \tilde{\mathbf{v}} \right\rangle + \langle \mathbf{q}, \tilde{\mathbf{v}} \rangle \leq \|\mathbf{q}\| \left\| \frac{\mathbf{q}}{\|\mathbf{q}\|} - \tilde{\mathbf{v}} \right\| + \langle \mathbf{q}, \tilde{\mathbf{v}} \rangle \leq \frac{1}{2} \|\mathbf{q}\| + \langle \mathbf{q}, \tilde{\mathbf{v}} \rangle,$$

which implies $\|\mathbf{q}\| \leq 2\langle \mathbf{q}, \tilde{\mathbf{v}} \rangle$. Consequently, $\|\mathbf{q}\| \leq 2 \max_{\mathbf{v} \in \mathcal{N}} \langle \mathbf{q}, \mathbf{v} \rangle$, so it suffices to bound the latter. Recalling, by Lemma 9, that $\|w_i \mathbf{p}_i\|_{\psi_1} \lesssim \sqrt{k \wedge (1 + \nu)}$, by Bernstein's inequality and a union bound over the net,

$$\Pr \left(\max_{\mathbf{v} \in \mathcal{N}} \langle \mathbf{q}, \mathbf{v} \rangle \geq t \right) \leq 2 \cdot 5^k \cdot \exp[-c_1 N(\delta \wedge \delta^2)], \quad \delta := \frac{t}{\sqrt{k \wedge (1 + \nu)}}.$$

Set

$$t = \sqrt{k \wedge (1 + \nu)} \left(\frac{10}{c_1} \frac{k}{N} \vee \sqrt{\frac{10}{c_1} \frac{k}{N}} \right),$$

so that the probability is bounded by $2 \cdot 5^k e^{-10k} = 2e^{-\Omega(k)}$. ■

Lemma 12 *With probability $1 - 2e^{-\Omega(k)}$,*

$$\varepsilon_3 \lesssim \left(\frac{k}{N} \vee \sqrt{\frac{k}{N}} \right).$$

Proof. By Lemma 9, the vectors \mathbf{p}_i are $O(1)$ -sub-Gaussian. The claim follows by Lemma 22, applied with $t \approx \sqrt{k}$, $\delta \approx \sqrt{\frac{k}{N}}$. ■

Proof of Lemma 4. The proof follows from Lemmas 10, 11 and 12, combined with $N \gtrsim p_{\text{Ball}} n$ from Lemma 1.

D Proof of Lemma 5

The core of the argument is this: we reduce the question of whether $\Pr((\mathbf{X}, \mathbf{Y}) \text{ is bad} \mid \mathbf{u})$ is large to a *geometric question*; specifically, whether a randomly rotated line segment is close to any non-zero lattice point. The details proceed as follow.

Recall: the pair (\mathbf{X}, \mathbf{Y}) is bad when $\mathbf{Y} \in \mathcal{B}(\mathbf{0}, R)$ but $\mathbf{X} \notin \mathcal{Q}_\Delta := [-\frac{1}{2}\Delta, \frac{1}{2}\Delta]^k$. Our goal is to show that for most directions $\mathbf{u} \in \mathbb{S}^{k-1}$, the probability that (\mathbf{X}, \mathbf{Y}) is bad is small, specifically,

$$\Pr((\mathbf{X}, \mathbf{Y}) \text{ is bad} \mid \mathbf{u}) \leq k^{-M^2}.$$

We start by constraining ourselves to a set of “typical” vectors \mathbf{X} . As in Eq. (2), write, $\mathbf{X} = \sqrt{\nu} \xi \mathbf{u} + \mathbf{Z}$, for independent $\xi \sim \mathcal{N}(0, 1)$, $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$. Let $\delta > 0$ be a confidence parameter (we shall set $\delta = k^{-M^2}$ later), and consider the event

$$\mathcal{E}_{\mathbf{X}} = \{ \|\mathbf{Z}\|_2 \leq z_2(\delta/3), \|\mathbf{Z}\|_\infty \leq z_\infty(\delta/3), |\xi| \leq h(\delta/3) \} \quad (21)$$

where $h(\delta) = \sqrt{2 \log(2/\delta)}$ is such that $\Pr(|\xi| \geq h(\delta)) \leq \delta$, and $z_2(\delta), z_\infty(\delta)$ are as in Eq. (10). Clearly, $\Pr(\mathcal{E}_{\mathbf{X}}) \leq \delta$.

Operating under $\mathcal{E}_{\mathbf{X}}$, let us bound the event $\{(\mathbf{X}, \mathbf{Y}) \text{ is bad}\}$ by another, larger, event. To start, note that $\mathbf{X} \neq \mathbf{Y}$ implies that $\mathbf{X} = \mathbf{Y} + \Delta \mathbf{t}$ for some *non-zero* lattice vector $\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}$. Consequently, when $\mathbf{Y} \in \mathcal{B}(\mathbf{0}, R)$, $\mathbf{X} \neq \mathbf{Y}$ implies that $\mathbf{X} \in \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, R)$. Decomposing \mathbf{X} , this further implies that

$$\sqrt{\nu} \xi \mathbf{u} \in \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, R + \|\mathbf{Z}\|). \quad (22)$$

As for the condition $\mathbf{X} \notin \mathcal{Q}_\Delta$, equivalently $\|\mathbf{X}\|_\infty > \frac{1}{2}\Delta$, it follows from the triangle inequality that

$$\|\sqrt{\nu} \xi \mathbf{u}\|_\infty = |\sqrt{\nu} \xi| \|\mathbf{u}\|_\infty \geq \frac{1}{2}\Delta - \|\mathbf{Z}\|_\infty. \quad (23)$$

Let $\mathcal{U}_{1,M} \subseteq \mathbb{S}^{k-1}$ be the set of *incoherent* directions,

$$\mathcal{U}_{1,M} := \left\{ \mathbf{u} \in \mathbb{S}^{k-1} : \|\mathbf{u}\|_\infty \leq \frac{1}{C_1} \sqrt{\frac{\log k}{k}} \right\}, \quad (24)$$

with C_1 a universal constant such that $\Pr(\mathbf{u} \in \mathcal{U}_{1,M}) = 1 - O(k^{-10})$ for $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$ (see Lemma 18). Now, under $\mathcal{E}_{\mathbf{X}}$, and assuming that $\mathbf{u} \in \mathcal{U}_{1,M}$, Eqs. (22) and (23) imply that

$$\sqrt{\nu}\xi\mathbf{u} \in \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta\mathbf{t}, R_\delta), \quad |\sqrt{\nu}\xi| \geq C_1\Delta_\delta \sqrt{\frac{k}{\log k}}, \quad (25)$$

where we set

$$\Delta_\delta := \frac{1}{2}\Delta - z_\infty(\delta/3), \quad R_\delta := R + z_2(\delta/3). \quad (26)$$

Henceforth, we shall assume Δ to be large enough so that $\Delta_\delta > 0$. Consider the line segment, $L(\mathbf{u}) \subset \mathbb{R}^k$,

$$L(\mathbf{u}) = \{s\mathbf{u} : A \leq s \leq B\}, \quad \text{where } A := C_1\Delta_\delta \sqrt{\frac{k}{\log k}}, \quad B := \sqrt{\nu}h(\delta/3) = \sqrt{2\nu \log(6/\delta)}. \quad (27)$$

Observe that under $\mathcal{E}_{\mathbf{X}}$, the occurrence of the event in Eq. (25) implies, in particular, that

$$L(\mathbf{u}) \cap \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta\mathbf{t}, R_\delta) \neq \emptyset.$$

Note that given \mathbf{u} , this is a *deterministic* geometric condition. Set

$$\mathcal{U}_{2,M} := \left\{ \mathbf{u} \in \mathbb{S}^{k-1} : L(\mathbf{u}) \cap \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta\mathbf{t}, R_\delta) = \emptyset \right\}. \quad (28)$$

and

$$\mathcal{U}_M := \mathcal{U}_{1,M} \cap \mathcal{U}_{2,M}. \quad (29)$$

Summarizing the preceding discussion, we have argued that whenever $\mathbf{u} \in \mathcal{U}_M$, the event $\mathcal{E}_{\mathbf{X}}$ already implies that (\mathbf{X}, \mathbf{Y}) are *good*. Thus, for $\mathbf{u} \in \mathcal{U}_M$,

$$\Pr((\mathbf{X}, \mathbf{Y}) \text{ is bad} \mid \mathbf{u}) \leq \Pr(\mathcal{E}_{\mathbf{X}}^c \mid \mathbf{u}) = \Pr(\mathcal{E}_{\mathbf{X}}^c) \leq \delta. \quad (30)$$

The proof of Lemma 5 will follow from the following auxiliary result:

Lemma 13 *Fix a constant $M \geq 1$ and set $\delta = k^{-M^2}$. There is a universal constant $C_* > 0$, such that if $\Delta \geq C_* \left(M\sqrt{\log k} \vee \nu^{\frac{1}{2(k-1)}} \right)$ then, for $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$,*

$$\Pr(\mathbf{u} \notin \mathcal{U}_{2,M}) \leq O_M(k^{-10}).$$

Lemma 13 is purely a result in geometric probability. It states the following: take the 1D line segment $\tilde{L}(\mathbf{e}_1) \subseteq \mathbb{R}^k$, and rotate it uniformly in space (apply a random rotation $U \sim \text{Haar}(O(k))$). Then with high probability, the rotated segment will end up far away from all non-zero lattice points. The remainder of this section is devoted to proving Lemma 13.

Let us discretize the interval $[A, B]$ into disjoint sub-intervals of maximal length, such that the length of a sub-interval is $\leq R_\delta$; let $s_0 = A < s_1 < \dots < s_T = B$ be the corresponding end-points, and note that we may take $T \leq \left\lceil \frac{B-A}{R_\delta} \right\rceil + 1$. Clearly, any point in $L(\mathbf{u})$ must be $0.5R_\delta$ -close to some point in $\{s_1\mathbf{u}, \dots, s_T\mathbf{u}\}$. In particular,

$L(\mathbf{u}) \cap \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, R_\delta) \neq \emptyset$ implies that $s_\ell \mathbf{u} \in \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, 1.5R_\delta)$ for some $1 \leq \ell \leq T$. Consequently,

$$\begin{aligned} \Pr \left(L(\mathbf{u}) \cap \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, R_\delta) \neq \emptyset \right) &\leq \Pr \left(\{s_1 \mathbf{u}, \dots, s_T \mathbf{u}\} \cap \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, 1.5R_\delta) \neq \emptyset \right) \\ &\leq \sum_{\ell=1}^T \Pr \left(s_\ell \mathbf{u} \in \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B}(\Delta \mathbf{t}, 1.5R_\delta) \right) \\ &= \sum_{\ell=1}^T \Pr \left(\mathbf{u} \in \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right) \\ &=: \sum_{\ell=1}^T p_\ell. \end{aligned} \tag{31}$$

Since $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$, each term of Eq. (31) is, by definition,

$$p_\ell = \frac{\sigma_{k-1} \left(\mathbb{S}^{k-1} \cap \bigcup_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right)}{\sigma_{k-1}(\mathbb{S}^{k-1})} \leq \frac{\sum_{\mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\}} \sigma_{k-1} \left(\mathbb{S}^{k-1} \cap \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right)}{\sigma_{k-1}(\mathbb{S}^{k-1})},$$

where $\sigma_{k-1}(\cdot)$ denotes the surface area. Note that, on the one hand,

$$\begin{aligned} \sigma_{k-1} \left(\mathbb{S}^{k-1} \cap \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right) &\leq \sigma_{k-1} \left(\partial \left(\mathcal{B}(\mathbf{0}, 1) \cap \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right) \right) \\ &\stackrel{(*)}{\leq} \sigma_{k-1} \left(\partial \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right) \\ &= \left(\frac{1.5R_\delta}{s_\ell} \right)^{k-1} \sigma_{k-1}(\mathbb{S}^{k-1}), \end{aligned}$$

where $\partial(\cdot)$ denotes the boundary of a set, and $(*)$ follows from the well-known fact that for convex bodies $L \subset K$, $\sigma_{k-1}(\partial L) \leq \sigma_{k-1}(\partial K)$; see, e.g., [Artstein-Avidan et al., 2015, Theorem B.1.14]. On the other hand, clearly, $\sigma_{k-1} \left(\mathbb{S}^{k-1} \cap \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \right) = 0$ whenever $\mathbb{S}^{k-1} \cap \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) = \emptyset$. Setting

$$N_\ell = \left| \left\{ \mathbf{t} \in \mathbb{Z}^k \setminus \{\mathbf{0}\} : \mathbb{S}^{k-1} \cap \mathcal{B} \left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell} \right) \neq \emptyset \right\} \right|, \tag{32}$$

we conclude that

$$p_\ell \leq N_\ell \left(\frac{1.5R_\delta}{s_\ell} \right)^{k-1}. \tag{33}$$

Lemma 14 *We have*

$$N_\ell \leq \begin{cases} \mathcal{V}_k \cdot \left(\frac{s_\ell + 1.5R_\delta}{\Delta} + \sqrt{k} \right)^k & \text{if } s_\ell < \Delta\sqrt{k} + 1.5R_\delta \\ k \cdot \mathcal{V}_k \cdot \left(\frac{3R_\delta}{\Delta} + 2\sqrt{k} \right) \left(\frac{s_\ell + 1.5R_\delta}{\Delta} + \sqrt{k} \right)^{k-1} & \text{if } s_\ell \geq \Delta\sqrt{k} + 1.5R_\delta \end{cases},$$

where \mathcal{V}_k is the volume of the k -dimensional unit ball.

Proof. This is an essentially standard packing argument, made slightly more complicated (when s_ℓ is large) since we are considering intersections against a sphere rather than a ball. For radii $0 \leq r_1 \leq r_2$, denote the (closed) annulus by

$$\mathcal{A}(r_1, r_2) = \mathcal{B}(\mathbf{0}, r_2) \setminus \text{int}(\mathcal{B}(\mathbf{0}, r_1)).$$

Observe that $\mathbb{S}^{k-1} \cap \mathcal{B}\left(\frac{\Delta}{s_\ell} \mathbf{t}, \frac{1.5R_\delta}{s_\ell}\right) \neq \emptyset$ implies⁶ $\frac{\Delta}{s_\ell} \mathbf{t} \in \mathcal{A}\left(\left[1 - \frac{1.5R_\delta}{s_\ell}\right]_+, 1 + \frac{1.5R_\delta}{s_\ell}\right)$, so

$$N_\ell \leq \left| \mathbb{Z}^k \cap \mathcal{A}\left(\left[\frac{s_\ell - 1.5R_\delta}{\Delta}\right]_+, \frac{s_\ell + 1.5R_\delta}{\Delta}\right) \right|.$$

Next, we use the following packing argument: the sets $\mathbb{Z}^k + (-\frac{1}{2}, \frac{1}{2})^k$ are disjoint, so that if $\mathbf{t} \in \mathcal{A}(r_1, r_2)$ then $\mathbf{t} + (-\frac{1}{2}, \frac{1}{2})^k \subset \mathcal{A}\left(\left[r_1 - \sqrt{k}\right]_+, r_2 + \sqrt{k}\right)$. Therefore, by a volume comparison,

$$|\mathbb{Z}^k \cap \mathcal{A}(r_1, r_2)| \leq \frac{\text{Vol}_k\left(\mathcal{A}([r_1 - \sqrt{k}]_+, r_2 + \sqrt{k})\right)}{\text{Vol}_k\left((-1/2, 1/2)^k\right)} = \mathcal{V}_k \cdot \left((r_2 + \sqrt{k})^k - [r_1 - \sqrt{k}]_+^k\right).$$

Set $r_2 = \frac{s_\ell + 1.5R_\delta}{\Delta}$ and $r_1 = \left[\frac{s_\ell - 1.5R_\delta}{\Delta}\right]_+$, so that

$$N_\ell \leq \mathcal{V}_k \cdot \left(\left(\frac{s_\ell + 1.5R_\delta}{\Delta} + \sqrt{k} \right)^k - \left(\left[\frac{s_\ell - 1.5R_\delta}{\Delta} - \sqrt{k} \right]_+ \right)^k \right).$$

The second term is non-zero if and only if $s_\ell \geq \Delta\sqrt{k} + 1.5R_\delta$; the claimed bound follows from the inequality $|a^k - b^k| \leq k|b - a| \max\{|a|, |b|\}^{k-1}$. ■

We now conclude the proof of Lemma 13. Recall, by Eq. (31), that our goal is to bound $\sum_{\ell=1}^T p_\ell$, where p_ℓ is bounded in Eq. (33). We treat separately small and large terms in the sum.

- **Small terms:** ℓ -s such that $s_\ell < \Delta\sqrt{k} + 1.5R_\delta$. Note that there are $\leq \frac{\Delta\sqrt{k} + 1.5R_\delta}{R_\delta} + 1 \lesssim 1 + \frac{\Delta\sqrt{k}}{R_\delta}$ such terms. Bound

$$\mathcal{V}_k \leq (C/\sqrt{k})^k, \quad \frac{1.5R_\delta}{s_\ell} \leq \frac{1.5R_\delta}{A} \lesssim \frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}}$$

(recall $s_\ell \geq A$ and the definition of A in Eq. (27)). Assuming

$$\Delta \geq \frac{R_\delta}{\sqrt{k}}, \tag{34}$$

we have

$$\frac{s_\ell + 1.5R_\delta}{\Delta} + \sqrt{k} \leq \frac{(\Delta\sqrt{k} + 1.5R_\delta) + 1.5R_\delta}{\Delta} + \sqrt{k} \lesssim \sqrt{k} + \frac{R_\delta}{\Delta} \stackrel{(34)}{\lesssim} \sqrt{k}.$$

Plugging into Eq. (33) and Lemma 14,

$$p_\ell \lesssim C^k (1/\sqrt{k})^k \cdot (\sqrt{k})^k \left(\frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}} \right)^{k-1} \lesssim \left(C \frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}} \right)^{k-1},$$

for some universal C . Recalling again that there are $\lesssim 1 + \frac{\Delta\sqrt{k}}{R_\delta} \stackrel{(34)}{\lesssim} \frac{\Delta\sqrt{k}}{R_\delta}$ such terms, and that, by definition (Eq. (26)),

$$\Delta = 2\Delta_\delta + 2z_\infty(\delta/2) \lesssim \Delta_\delta + \sqrt{\log k \vee \log(1/\delta)}, \tag{35}$$

the total sum of the small terms is

$$\begin{aligned} &\lesssim \sqrt{\log k} \left(C \frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}} \right)^{k-2} + \sqrt{\log k \vee \log(1/\delta)} \cdot \frac{\sqrt{k}}{R_\delta} \cdot \left(C \frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}} \right)^{k-1} \\ &\lesssim \sqrt{\log k} \left(C \frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}} \right)^{k-2} + \sqrt{k \log k} \left(C \frac{R_\delta\sqrt{\log k}}{\Delta_\delta\sqrt{k}} \right)^{k-1}, \end{aligned}$$

⁶ $[x]_+$ denotes the positive part of x , namely, $[x]_+ = \max\{x, 0\}$.

where, for the second inequality, we used $R_\delta \geq z_2(\delta/3) \gtrsim \sqrt{\log(1/\delta)}$. Consequently, whenever

$$\Delta_\delta \gtrsim \frac{R_\delta \sqrt{\log k}}{\sqrt{k}},$$

the sum is exponentially decaying in k , and in particular $O(k^{-10})$. Again, recalling Eq. (35), the following condition on Δ is sufficient to get exponential decay:

$$\Delta \gtrsim \frac{R_\delta \sqrt{\log k}}{\sqrt{k}} + \sqrt{\log k \vee \log(1/\delta)}. \quad (36)$$

- **Large terms:** such that $s_\ell \geq \Delta\sqrt{k} + 1.5R_\delta$. Note that there are $\lesssim B/R_\delta \lesssim \frac{\sqrt{\nu \log(1/\delta)}}{R_\delta}$ such terms (recall the definition of B in Eq. (27)). Bounding $\frac{s_\ell + 1.5R_\delta}{\Delta} + \sqrt{k} \leq \frac{2s_\ell}{\Delta}$, we estimate, using $\mathcal{V}_k \leq (C/\sqrt{k})^k$ and assuming condition (34),

$$N_\ell \leq k \cdot \mathcal{V}_k \cdot \left(\frac{3R_\delta}{\Delta} + 2\sqrt{k} \right) \left(\frac{s_\ell + 1.5R_\delta}{\Delta} + \sqrt{k} \right)^{k-1} \lesssim k(C/\sqrt{k})^k \sqrt{k} \left(\frac{2s_\ell}{\Delta} \right)^{k-1},$$

so that, using Eq. (33),

$$p_\ell \leq N_\ell \left(\frac{1.5R_\delta}{s_\ell} \right)^{k-1} \lesssim k \left(3C \frac{R_\delta}{\Delta\sqrt{k}} \right)^{k-1}.$$

Again, since there are $\lesssim \frac{\sqrt{\nu \log(1/\delta)}}{R_\delta}$ such terms, the total contribution is

$$\lesssim \frac{k}{R_\delta} \sqrt{\nu \log(1/\delta)} \cdot \left(3C \frac{R_\delta}{\Delta\sqrt{k}} \right)^{k-1} \leq k \left(3C \frac{R_\delta \nu^{\frac{1}{2(k-1)}}}{\Delta\sqrt{k}} \right)^{k-1},$$

where, for the second inequality, we again used $R_\delta \gtrsim \sqrt{\log(1/\delta)}$. This is exponentially decreasing in k whenever

$$\Delta \gtrsim \frac{R_\delta}{\sqrt{k}} \nu^{\frac{1}{2(k-1)}}. \quad (37)$$

We finish by simplifying conditions (36) and (37) further. Setting $\delta = k^{-M^2}$, we may estimate

$$R_\delta \approx \sqrt{k} \vee \sqrt{\log(1/\delta)} = \sqrt{k} \vee (M\sqrt{\log k}),$$

so that for large $k \geq k_0(M)$, $R_\delta \approx \sqrt{k}$. Thus, (36) reads $\Delta \gtrsim M\sqrt{\log k}$, and (37) reads $\Delta \gtrsim \nu^{\frac{1}{2(k-1)}}$.

E Proof of Lemma 7

Recall the choice of R from Eq. (11). Decompose, rather arbitrarily, $R^2 = k + B$, so that $3k \leq B \leq C_3k$ for some C_3 . Note that if $g_1, \dots, g_k \sim \mathcal{N}(0, 1)$, then $\Pr\left(\sum_{i=2}^k g_i^2 \leq B\right) \geq 1 - e^{-C_4k}$ for some $C_4 > 0$. Following Proposition 1 Eq. (7),

$$\lambda_1(\Sigma_{\text{Ball}}) = \mathbb{E} \left[\lambda_i g_i^2 \mid \sum_{i=1}^k \lambda_i g_i^2 \leq R^2 \right] = \frac{\mathbb{E} \left[(1+\nu)g_1^2 \cdot \mathbf{1}_{(1+\nu)g_1^2 + \sum_{i=2}^k g_i^2 \leq k+B} \right]}{\Pr \left((1+\nu)g_1^2 + \sum_{i=2}^k g_i^2 \leq k+B \right)}.$$

Clearly, $\mathbf{1}_{(1+\nu)g_1^2 + \sum_{i=2}^k g_i^2 \leq k+B} \geq \mathbf{1}_{(1+\nu)g_1^2 \leq k} \cdot \mathbf{1}_{\sum_{i=2}^k g_i^2 \leq B}$, therefore,

$$\begin{aligned} \mathbb{E} \left[(1+\nu)g_1^2 \cdot \mathbf{1}_{(1+\nu)g_1^2 + \sum_{i=2}^k g_i^2 \leq k+B} \right] &\geq \mathbb{E} \left[(1+\nu)g_1^2 \cdot \mathbf{1}_{(1+\nu)g_1^2 \leq k} \cdot \mathbf{1}_{\sum_{i=2}^k g_i^2 \leq B} \right] \\ &\stackrel{(*)}{=} \mathbb{E} \left[(1+\nu)g_1^2 \cdot \mathbf{1}_{(1+\nu)g_1^2 \leq k} \right] \cdot \mathbb{E} \left[\mathbf{1}_{\sum_{i=2}^k g_i^2 \leq B} \right] \\ &\geq \mathbb{E} \left[(1+\nu)g_1^2 \cdot \mathbf{1}_{(1+\nu)g_1^2 \leq k} \right] (1 - e^{-C_4k}), \end{aligned}$$

where (\star) holds since this is the product of independent random variables. Furthermore, clearly,

$$\Pr \left((1 + \nu)g_1^2 + \sum_{i=1}^k g_i^2 \leq k + B \right) \leq \Pr \left((1 + \nu)g_1^2 \leq k + B \right) \leq \Pr \left((1 + \nu)g_1^2 \leq (1 + C_3)k \right).$$

Let $g_1^2 =: W \sim \chi^2(1)$, so that, finally,

$$\lambda_1(\Sigma_{\text{Ball}}) \geq (1 - e^{-C_4 k}) \frac{\mathbb{E} \left[(1 + \nu)W \cdot \mathbf{1}_{W \leq \frac{k}{1+\nu}} \right]}{\Pr \left(W \leq \frac{(1+C_3)k}{1+\nu} \right)}. \quad (38)$$

We continue case-by-case, depending on the magnitude of ν :

- (i) Suppose that $e^{-C_1 k} \leq \nu \leq 1$, where C_1 is a sufficiently small constant. Since W has an exponential tail and $\mathbb{E}[W] = 1$, there is some C_5 such that

$$\frac{\mathbb{E} \left[W \cdot \mathbf{1}_{W \leq \frac{k}{1+\nu}} \right]}{\Pr \left(W \leq \frac{(1+C_3)k}{1+\nu} \right)} \geq \frac{\mathbb{E} \left[W \cdot \mathbf{1}_{W \leq \frac{k}{2}} \right]}{\Pr \left(W \leq (1 + C_3)k \right)} \geq 1 - e^{-C_5 k},$$

therefore

$$\lambda_1(\Sigma_{\text{Ball}}) \geq (1 + \nu)(1 - e^{-C_4 k})(1 - e^{-C_5 k}) \geq 1 + C\nu$$

for small enough C , whenever C_1 is chosen sufficiently small compared to C_4, C_5 .

- (ii) Note that by [Palombi et al., 2012], $\lambda_1(\Sigma_{\text{Ball}})$ increases with ν . Consequently, for all $\nu \geq 1$, (i) implies that $\lambda_1(\Sigma_{\text{Ball}}) \geq 1 + C$. Now, suppose that $1 \leq \nu \leq Ak - 1$, where A is such that for all $A' \leq A$,

$$\frac{\mathbb{E} \left[W \cdot \mathbf{1}_{W \leq \frac{1}{A'}} \right]}{\Pr \left(W \leq \frac{(1+C_3)}{A'} \right)} \geq \frac{4}{5}.$$

Note that such A indeed exists, since the above ratio $\rightarrow 1$ as $A' \rightarrow 0$. Then

$$\lambda_1(\Sigma_{\text{Ball}}) \geq (1 + \nu)(1 - e^{-C_4 k}) \frac{\mathbb{E} \left[W \cdot \mathbf{1}_{W \leq \frac{k}{1+\nu}} \right]}{\Pr \left(W \leq \frac{(1+C_3)k}{1+\nu} \right)} \geq \frac{4}{5}(1 - e^{-C_4})(1 + \nu) \geq C'(1 + \nu),$$

so that $\lambda_1(\Sigma_{\text{Ball}}) \geq (1 + C) \vee C'(1 + \nu)$. Consequently, $\lambda_1(\Sigma_{\text{Ball}}) \geq 1 + C''\nu$ for some other C'' .

- (iii) $\nu \geq Ak - 1$. Consider the function

$$F(h) = \frac{1}{h} \cdot \frac{\mathbb{E} [W \cdot \mathbf{1}_{W \leq h}]}{\Pr (W \leq (1 + C_3)h)},$$

so that

$$\lambda_1(\Sigma_{\text{Ball}}) \geq (1 - e^{-C_4 k}) \cdot k \cdot F \left(\frac{k}{1 + \nu} \right) \geq k \cdot (1 - e^{-C_4}) \cdot \inf_{h \leq 1/A} F(h).$$

We are done if we show that the infimum is non-zero, and it clearly suffices to show that $\lim_{h \rightarrow 0+} F(h) > 0$. To do this, recall that $W \sim \chi^2(1)$ has a density $f_W(w) \propto w^{-1/2}e^{-w/2}$ supported on $w \geq 0$. Therefore, as $h \rightarrow 0+$,

$$\mathbb{E} [W \cdot \mathbf{1}_{W \leq h}] \sim Ch^{3/2}, \quad \Pr (W \leq (1 + C_3)h) \sim C'h^{1/2},$$

so that indeed $F(h) \sim C''$ as $h \rightarrow 0+$.

F Proof of Theorem 1

We shall use the following eigenvalue perturbation result [Wainwright, 2019, Theorem 8.5]:

Lemma 15 *Let \mathbf{A} be positive semidefinite, with a positive spectral gap: $\delta := \lambda_1(\mathbf{A}) - \lambda_2(\mathbf{A}) > 0$. Let $\mathbf{u} \in \mathbb{S}^{k-1}$ be its largest eigenvector. Suppose that $\widehat{\mathbf{A}}$ is positive semidefinite with $\|\mathbf{A} - \widehat{\mathbf{A}}\| \leq \delta/4$. Let $\widehat{\mathbf{u}}$ be its maximal eigenvector, with the sign chosen such that $\langle \mathbf{u}, \widehat{\mathbf{u}} \rangle \geq 0$ (part of the claim is that the largest eigenspace of $\widehat{\mathbf{A}}$ is 1-dimensional). Then*

$$\|\mathbf{u} - \widehat{\mathbf{u}}\| \leq \frac{4}{\delta} \cdot \left\| (\mathbf{I} - \mathbf{u}\mathbf{u}^\top) (\mathbf{A} - \widehat{\mathbf{A}})\mathbf{u} \right\|,$$

where $\mathbf{I} - \mathbf{u}\mathbf{u}^\top$ is the projection onto the orthogonal complement of \mathbf{u} .

We apply Lemma 15 with $\mathbf{A} = \boldsymbol{\Sigma}_{\text{Ball}}$ and $\widehat{\mathbf{A}} = \widehat{\boldsymbol{\Sigma}}$, using the error bounds developed so far. For brevity, denote

$$\delta := \lambda_1(\boldsymbol{\Sigma}_{\text{Ball}}) - \lambda_2(\boldsymbol{\Sigma}_{\text{Ball}}). \quad (39)$$

Lemma 16 *Assume the setup of Lemma 5, with Δ large, $M > \sqrt{12}$, and $\mathbf{u} \in \mathcal{U}_M$.*

Assume either of the following conditions hold:

- $1 \leq \nu \leq k$ and $n \gtrsim \frac{k}{\sqrt{\nu}} \vee \log k$.
- $k \leq \nu \lesssim k^{2M^2-21}$ and $n \gtrsim \sqrt{\nu}$.

Then with probability $1 - O(k^{-10})$, one has $\|\boldsymbol{\Sigma}_{\text{Ball}} - \widehat{\boldsymbol{\Sigma}}\| \leq \delta/4$

Proof. We consider two cases:

- Suppose that $1 \leq \nu \leq k$. By Lemma 2, $p_{\text{Ball}} \approx 1$, and by Lemma 7, $\delta \gtrsim \nu$. By Lemmas 4 and 6, it holds with probability $1 - O(k^{-10})$ that

$$\|\boldsymbol{\Sigma}_{\text{Ball}} - \widehat{\boldsymbol{\Sigma}}\| \lesssim \nu \sqrt{\frac{\log k}{n}} + \sqrt{\nu} \left(\frac{k}{n} \vee \sqrt{\frac{k}{n}} \right) + k^{-M^2+12}.$$

Consequently, when $n \gtrsim \log k$, $n \gtrsim \frac{k}{\sqrt{\nu}}$ and $n \gtrsim \frac{k}{\nu}$ (the last condition is redundant, since $\nu \geq 1$), it holds that, for large enough k , $\|\boldsymbol{\Sigma}_{\text{Ball}} - \widehat{\boldsymbol{\Sigma}}\| \leq \delta/4$.

- Suppose that $k \leq \nu \leq k^{2M^2-21}$. By Lemma 2, $p_{\text{Ball}} \approx \sqrt{k/\nu}$, and by Lemma 7, $\delta \gtrsim k$. By Lemmas 4 and 6, it holds with probability $1 - O(k^{-10})$, provided that $p_{\text{Ball}}n \gtrsim \log k \implies n \gtrsim \sqrt{\frac{\nu}{k}} \log k$, that

$$\|\boldsymbol{\Sigma}_{\text{Ball}} - \widehat{\boldsymbol{\Sigma}}\| \lesssim k \sqrt{\frac{\sqrt{\nu} \log k}{n\sqrt{k}}} + k \left(\frac{\sqrt{\nu}}{n} \vee \sqrt{\frac{\sqrt{\nu}}{n\sqrt{k}}} \right) + k \cdot \nu^{1/2} k^{-M^2+10.5}.$$

Consequently, whenever $n \gtrsim \sqrt{\frac{\nu}{k}} \log k$, $n \gtrsim \sqrt{\nu}$ (the first condition is redundant), and $\nu \lesssim k^{-2M^2+21}$, it holds that $\|\boldsymbol{\Sigma}_{\text{Ball}} - \widehat{\boldsymbol{\Sigma}}\| \leq \delta/4$.

■

Proof of Theorem 1. We apply Lemma 15. Write, as before, $\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_{\text{Ball}} = (\widehat{\boldsymbol{\Sigma}} - \widehat{\boldsymbol{\Sigma}}_{\mathbf{X}}) + (\widehat{\boldsymbol{\Sigma}}_{\mathbf{X}} - \boldsymbol{\Sigma}_{\text{Ball}})$, so $\left\| (\mathbf{I} - \mathbf{u}\mathbf{u}^\top) (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_{\text{Ball}})\mathbf{u} \right\| \leq \varepsilon_{\text{Pick}} + \left\| (\mathbf{I} - \mathbf{u}\mathbf{u}^\top) (\widehat{\boldsymbol{\Sigma}}_{\mathbf{X}} - \boldsymbol{\Sigma}_{\text{Ball}})\mathbf{u} \right\|$. Using the decomposition Eq. (19), and recalling the notation in Eq. (20), we conclude that under the conditions of Lemma 16, with probability $1 - O(k^{-10})$,

$$\|\mathbf{u} - \widehat{\mathbf{u}}\| \leq \frac{4}{\delta} \left\| (\mathbf{I} - \mathbf{u}\mathbf{u}^\top) (\widehat{\boldsymbol{\Sigma}} - \boldsymbol{\Sigma}_{\text{Ball}})\mathbf{u} \right\| \leq \frac{4}{\delta} (\varepsilon_{\text{Pick}} + 2\varepsilon_2 + \varepsilon_3).$$

Note that the term ε_1 does not appear, since it corresponds to a components of the difference $\widehat{\boldsymbol{\Sigma}}_{\mathbf{X}} - \boldsymbol{\Sigma}_{\text{Ball}}$ which is parallel to \mathbf{u} . Using Lemmas 6, 11 and 12, $\delta \gtrsim k \vee \nu$ (Lemma 7), and $p_{\text{Ball}} \approx 1 \wedge \sqrt{k/\nu}$, we conclude that the following holds with probability $1 - O(k^{-10})$:

- Suppose that $1 \leq \nu \leq k$ and $n \gtrsim \frac{k}{\sqrt{\nu}} \vee \log k$. Then

$$\|\mathbf{u} - \hat{\mathbf{u}}\| \lesssim \frac{1}{\sqrt{\nu}} \left(\frac{k}{n} \vee \sqrt{\frac{k}{n}} \right) + \frac{1}{\nu} \cdot k^{-M^2+12}.$$

Note that the requirement $n \gtrsim \frac{k}{\sqrt{\nu}}$ may effectively be omitted from the statement of the Theorem. The reason is that a bound of the form $\|\mathbf{u} - \hat{\mathbf{u}}\| \leq B$, for any $B \geq 2$, is completely vacuous (since $\mathbf{u}, \hat{\mathbf{u}}$ are unit vectors). As we are not keeping track of the exact constants, it suffices to note that the first term in the upper bound becomes meaningful only when $n \gtrsim \frac{k}{\sqrt{\nu}}$.

- Suppose that $k \leq \nu \lesssim k^{2M^2-21}$ and $n \gtrsim \sqrt{\nu}$. Then

$$\|\mathbf{u} - \hat{\mathbf{u}}\| \lesssim \frac{\sqrt{\nu}}{n} + \sqrt{\frac{1}{n} \frac{\nu}{k}} + \nu^{1/2} k^{-M^2+10.5}.$$

The requirement $n \gtrsim \sqrt{\nu}$ is omitted from the statement of the Theorem, for the same reason as in the previous case.

G Proof of additional lemmas

In this section, we provide several short proofs, that were omitted from the main text due to space constraints.

G.1 Proof of Lemma 2

Upper bound: $\langle \mathbf{u}, \mathbf{X} \rangle \sim \mathcal{N}(0, 1 + \nu)$; clearly, $\mathbf{X} \in \mathcal{B}(\mathbf{0}, R)$ implies $|\langle \mathbf{u}, \mathbf{X} \rangle| \leq R$, so $p_{\text{Ball}} \leq \Pr(|\langle \mathbf{u}, \mathbf{X} \rangle| \leq R) = \text{erf}(R/\sqrt{2(1+\nu)})$.

Lower bound: Writing $\mathbf{X} = \sqrt{\nu}\xi\mathbf{u} + \mathbf{Z}$, the event $\{\sqrt{\nu}|\xi| \leq 2\sqrt{k}\} \cap \{\|\mathbf{Z}\| \leq z_2(0.1)\}$ implies $\mathbf{X} \in \mathcal{B}(\mathbf{0}, R)$. Thus,

$$\begin{aligned} p_{\text{Ball}} &\geq \Pr\left(\{\|\mathbf{Z}\| \leq z_2(0.1)\} \cap \{\sqrt{\nu}|\xi| \leq 2\sqrt{k}\}\right) \\ &\stackrel{(\star)}{=} \Pr(\|\mathbf{Z}\| \leq z_2(0.1)) \cdot \Pr(\sqrt{\nu}|\xi| \leq 2\sqrt{k}) \\ &\geq 0.9 \cdot \text{erf}\left(\sqrt{\frac{2k}{\nu}}\right), \end{aligned}$$

where (\star) follows since these are independent events.

G.2 Proof of Lemma 3

By the Gaussian correlation inequality [Royen, 2014, Latała and Matlak, 2017], for $f, h: \mathbb{R}^k \rightarrow \mathbb{R}$ quasiconcave⁷, and *one* of whom symmetric, $\mathbb{E}[f(\mathbf{X})h(\mathbf{X})] \geq \mathbb{E}[f(\mathbf{X})]\mathbb{E}[h(\mathbf{X})]$. Consequently, for g convex, $\mathbb{E}[g(\mathbf{X})h(\mathbf{X})] \leq \mathbb{E}[g(\mathbf{X})]\mathbb{E}[h(\mathbf{X})]$. The Lemma follows by taking $h(\mathbf{X}) = \mathbb{1}_{\mathbf{X} \in \mathcal{B}(\mathbf{0}, R)}$.

G.3 Proof of Lemma 6

First, suppose that $n \leq k^{M^2-10}$. Then by Lemma 5 and Markov's inequality, with probability $1 - O(nk^{-M^2}) = 1 - O(k^{-10})$, it holds that $|\mathcal{K}_{\text{Bad}}| = 0$, and consequently (Eq. (16)), $\varepsilon_{\text{Pick}} = 0$. Next, suppose that $n \geq k^{M^2-10}$. By Chernoff's inequality (Lemma 19), Lemma 5 and Lemma 1 it hold with probability $1 - O(k^{-10})$ that

$$\varepsilon_{\text{Pick}} \lesssim \frac{k}{np_{\text{Ball}}} |\mathcal{K}_{\text{Bad}}| \lesssim \frac{k}{np_{\text{Ball}}} \left(nk^{-M^2} \vee \log k \right) \leq \frac{k^{-M^2+1}}{p_{\text{Ball}}} + \frac{k \log k}{np_{\text{Ball}}} \leq \frac{2k^{-M^2+12}}{p_{\text{Ball}}}.$$

⁷ h is quasiconcave if $h(tx + (1-t)y) \geq \min\{h(x), h(y)\}$ for all $t \in [0, 1]$. Note that: (i) A concave function is quasiconcave; (ii) The indicator function of a convex set is quasiconcave (but *not* concave).

H Auxiliary technical lemmas

The following are tail bounds for some norms of a Gaussian random vector:

Lemma 17 *Let $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$. Then*

(i) ℓ_2 :

$$\Pr\left(\|\mathbf{Z}\|_2^2 \geq k + 2\sqrt{kx} + 2x\right) \leq e^{-x}, \quad \Pr\left(\|\mathbf{Z}\|_2^2 \leq k - 2\sqrt{kx}\right) \leq e^{-x}.$$

(ii) ℓ_∞ :

$$\Pr\left(\|\mathbf{Z}\|_\infty \geq \sqrt{2 \log k} + x\right) \leq 2e^{-\frac{1}{2}x^2}.$$

Proof. (i) is the well-known inequality of Laurent and Massart [Laurent and Massart, 2000, Lemma 1]. (ii) is a special case of the Borell-TIS inequality; alternatively, it follows from the Gaussian Lipschitz concentration inequality, e.g. [Boucheron et al., 2013, Theorem 5.6], with the elementary bound $\mathbb{E}[\max_{1 \leq i \leq k} Z_i] \leq \sqrt{2 \log k}$. ■

The following is an immediate corollary:

Lemma 18 (ℓ_∞ bound for a uniform vector in \mathbb{S}^{k-1}) *Suppose that $\mathbf{u} \sim \text{Unif}(\mathbb{S}^{k-1})$. There are absolute constants C, c such that*

$$\Pr\left(\|\mathbf{u}\|_\infty \geq c\sqrt{\frac{\log k}{k}}\right) \leq Ck^{-10}.$$

Proof. Let $\mathbf{Z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_k)$, so that $\mathbf{u} \stackrel{d}{=} \mathbf{Z}/\|\mathbf{Z}\|$. Choosing c large enough and using Lemma 17,

$$\Pr\left(\|\mathbf{u}\|_\infty \geq c\sqrt{\frac{\log k}{k}}\right) \leq \Pr\left(\|\mathbf{Z}\| \leq \frac{1}{2}\sqrt{k}\right) + \Pr\left(\|\mathbf{Z}\|_\infty \geq \frac{1}{2}c\sqrt{\log k}\right) \leq e^{-\Omega(k)} + O(k^{-10}) = O(k^{-10}).$$

■

Next is Chernoff's inequality for Bernoulli random variables [Vershynin, 2018, Theorem 2.3.1, Exercise 2.3.2]:

Lemma 19 (Chernoff's inequality) *Let X_1, \dots, X_n be independent Bernoulli random variables. Set $S_n = \sum_{i=1}^n X_i$ and $\mu = \mathbb{E}[S_n]$. Then for all $\alpha > 1$,*

$$\Pr(S_n \geq \alpha\mu) \leq e^{-\mu(\alpha/e)^{-\alpha}}, \quad \Pr(S_n \leq \mu/\alpha) \leq (e\alpha)^{\mu/\alpha} e^{-\mu}.$$

In particular, there are some absolute constants c, C such that

$$\Pr(S_n \geq C\mu) \leq e^{-c\mu}, \quad \Pr(S_n \leq \mu/C) \leq e^{-c\mu}.$$

We recall some properties of the sub-Gaussian and sub-exponential norms. The following is taken from [Vershynin, 2018, Chapter 2]:

Definition 1 (Orlicz norm) *Let $\psi : [0, \infty) \rightarrow [0, \infty)$ be convex, strictly increasing such that*

$$\psi(0) = 0, \quad \psi(x) \rightarrow \infty \text{ as } x \rightarrow \infty.$$

For a random variable X , define its ψ -Orlicz norm by

$$\|X\|_\psi = \inf\{t : \mathbb{E}\psi(|X|/t) \leq 1\}. \quad (40)$$

For a random vector \mathbf{X} , its ψ -Orlicz norm is $\|\mathbf{X}\|_\psi = \sup_{\mathbf{v} \in \mathbb{S}^{k-1}} \|\langle \mathbf{X}, \mathbf{v} \rangle\|_\psi$.

It is not hard to show that $\|\cdot\|_\psi$ is indeed a norm. The choices

$$\psi_2(x) = e^{x^2} - 1, \quad \psi_1(x) = e^x - 1$$

correspond to the sub-Gaussian and sub-exponential norms respectively. X is sub-Gaussian (resp. sub-exponential) in the “usual” sense if and only if $\|X\|_{\psi_2} < \infty$ (resp. $\|X\|_{\psi_1} < \infty$); see [Vershynin, 2018, Chapter 2] for more background. We briefly mention some properties of these norms that are used in the paper:

Lemma 20 *The following holds:*

1. $\|X^2\|_{\psi_1} = \|X\|_{\psi_2}^2$.
2. $\|XY\|_{\psi_1} \leq \|X\|_{\psi_2}\|Y\|_{\psi_2}$ (X, Y **do not** need to be independent).
3. *Centralization lemma:* $\|X - \mathbb{E}[X]\|_{\psi_i} \leq \|X\|_{\psi_i}$ for $i = 1, 2$.
4. *For independent X_1, \dots, X_n :* $\|\sum_{i=1}^n X_i\|_{\psi_2}^2 \leq C \sum_{i=1}^n \|X_i\|_{\psi_2}^2$, for some $C > 0$ universal.
5. *Hoeffding’s lemma:* for a bounded random variable, $\|\mathbf{X}\|_{\psi_2} \leq C\|\mathbf{X}\|_\infty$.

The following is Bernstein’s inequality for sums of independent sub-exponential random variables [Vershynin, 2018, Theorem 2.8.1]:

Lemma 21 (Bernstein’s inequality) *Let X_1, \dots, X_n be independent and sub-exponential. Set $S_n = \sum_{i=1}^n X_i$. Then for all $t \geq 0$,*

$$\Pr(|S_n - \mathbb{E}[S_n]| \geq t) \leq 2 \exp \left[-c \min \left(\frac{t^2}{\sum_{i=1}^n \|X_i\|_{\psi_1}^2}, \frac{t}{\max_{1 \leq i \leq n} \|X_i\|_{\psi_1}} \right) \right],$$

where $c > 0$ is an absolute constant.

Lastly, we cite a concentration inequality for sample covariance matrices with sub-Gaussian measurements [Vershynin, 2018, Theorem 4.6.1]:

Lemma 22 *Let $\mathbf{x}_1, \dots, \mathbf{x}_n \in \mathbb{R}^k$ be independent, centered and sub-Gaussian. Denote $K = \max_{1 \leq i \leq n} \|\mathbf{x}_i\|_{\psi_2}$. Let $\widehat{\Sigma} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$ be the sample covariance. There is $C > 0$ such that with probability at least $1 - 2e^{-t^2}$,*

$$\left\| \widehat{\Sigma} - \mathbb{E}[\widehat{\Sigma}] \right\| \leq K^2 \max\{\delta, \delta^2\}, \quad \text{where } \delta = C \left(\sqrt{\frac{k}{n}} + \frac{t}{\sqrt{n}} \right).$$

(Note that in [Vershynin, 2018, Theorem 4.6.1], the result is stated for isotropic vectors, meaning $\text{Cov}(\mathbf{x}_i) = \mathbf{I}$. However, the proof goes through, verbatim, also without this assumption.)