

Polynomial Time and Sample Complexity for Non-Gaussian Component Analysis: Spectral Methods

Yan Shuo Tan ^{*} Roman Vershynin [†]

April 3, 2017

Abstract

The problem of Non-Gaussian Component Analysis (NGCA) is about finding a maximal low-dimensional subspace E in \mathbb{R}^n so that data points projected onto E follow a non-gaussian distribution. Although this is an appropriate model for some real world data analysis problems, there has been little progress on this problem over the last decade.

In this paper, we attempt to address this state of affairs in two ways. First, we give a new characterization of standard gaussian distributions in high-dimensions, which lead to effective tests for non-gaussianness. Second, we propose a simple algorithm, *Reweighted PCA*, as a method for solving the NGCA problem. We prove that for a general unknown non-gaussian distribution, this algorithm recovers at least one direction in E , with sample and time complexity depending polynomially on the dimension of the ambient space. We conjecture that the algorithm actually recovers the entire E .

^{*}Department of Mathematics, University of Michigan, yanshuo@umich.edu

[†]Department of Mathematics, University of Michigan, romanv@umich.edu

1 Introduction

1.1 Non-Gaussian Component Analysis

Dimension reduction is a necessary step for much of modern data analysis, the principle being that the structure or “interestingness” of a collection of data points is contained in a geometric structure which has much lower dimension than the ambient vector space. We consider the case where the geometric structure in question is a linear subspace. In other words, we are in the situation where the variation of the data points within this subspace contains some information which we would like to extract, while their variation in the complementary directions constitute mere noise.

In many cases, it is reasonable to think of the noise as being gaussian. Formally, we then have the following generative model. Let E be an unknown d -dimensional subspace of \mathbb{R}^n , and let E^\perp be the orthogonal complement of E . Let X be a random vector in \mathbb{R}^n , which we can decompose into two independent components: a non-gaussian component \tilde{X} that takes values in E , and a gaussian component g that takes values in E^\perp . In other words, we let $X = (\tilde{X}, g) \in E \oplus E^\perp$.¹

Our goal is to recover the subspace E from a sample of independent realizations of X . This is precisely the framework of the problem of Non-Gaussian Component Analysis (NGCA). We make no assumption on the relative magnitudes of \tilde{X} and g . When the noise component is much smaller, which is a reasonable assumption in some real world applications, E can be recovered using the standard Principal Component Analysis (PCA). However, PCA manifestly fails when the signal to noise ratio is small, i.e. when \tilde{X} has lower magnitude than g .

With mild distributional assumptions, applying a whitening transformation to the data points can be done efficiently with sample size linear in the dimension [21]. As such, we might as well assume that the distribution is already whitened (i.e. isotropic). In other words, for the rest of this paper, we work with the model:

Definition 1.1 (Isotropic NGCA model).

$$X = (\tilde{X}, g) \in E \oplus E^\perp, \quad \mathbb{E}X = 0, \quad \mathbb{E}XX^T = I_n. \quad (1.1)$$

The NGCA problem is closely related to the problem of Independent Component Analysis (ICA), but generalizes it in a crucial way. ICA assumes the existence of a latent variable s with independent coordinates, whereas in our case, the distribution of \tilde{X} is allowed to have any manner of dependencies amongst its entries.

1.2 Previous work on NGCA

As far as we know, the NGCA problem was first studied by Blanchard et al. in their 2006 paper [3]. Following this paper, there has been a small but growing body of work on the problem [15, 16, 14, 9, 8, 19, 22], most of which is based on adapting ideas from Projection Pursuit, and in particular, the iterative *FastICA* algorithm of Hyvärinen [13].

None of these works have provided an algorithm with provable polynomial time and sample complexity. Suggested algorithms either have scant theoretical justification, or have analyses that are unsatisfactory for two reasons. First, where there is a consistency guarantee for an estimator \hat{E} , the guarantee is either conditional on some uncheckable, idiosyncratic assumptions on how the

¹It is not necessary to assume that the gaussian and non-gaussian subspaces are perpendicular. They automatically become perpendicular if we apply a whitening transformation.

particular algorithm interacts with the “non-gaussianness” of \tilde{X} (for instance, see Assumption 1 in [8]), or states only that it converges to a certain subspace F that is related to E , but may not coincide with E itself. Second, rates of convergence either omit dependence on the dimension or are again given in terms of idiosyncratic assumptions.

Furthermore, the algorithms with better theoretical grounding, such as in [8, 19], have also become incredibly complex, and one hopes that simpler algorithms should suffice.

1.3 Reweighted PCA in other contexts

Our proposed algorithm, Reweighted PCA, will involve performing PCA on a weighted sample. The idea of doing this with weight functions that are non-linear in the sample points can be traced back at least as far as Brubaker and Vempala’s work on *Isotropic PCA* [4]. In that paper, the authors similarly use gaussian weights, but do so in order to handle clustering for gaussian mixture models that are highly non-spherical. In a later paper [11], Goyal, Vempala, and Xiao used Fourier weights to handle ICA. While our analysis is radically different, the idea for the algorithm was directly inspired by these two papers.

2 Main results

The principle that underlies our approach to NGCA is a new characterization of multi-dimensional gaussian distributions. Throughout this section, X denotes a random vector in \mathbb{R}^n and g is a standard gaussian random vector in \mathbb{R}^n . By X' we will always denote an independent copy of X .

Theorem 2.1 (First gaussian test). *Suppose X has the same radial distribution as g , i.e. $\|X\|_2$ and $\|g\|_2$ are identically distributed. If $\langle X, X' \rangle$ has the same distribution as $\langle g, g' \rangle$, then X has the same distribution as g , i.e. the standard gaussian distribution.*

We will prove this theorem in Section 3 via a decomposition of moment tensors and a resulting energy minimization property of the gaussian measure. The theorem guarantees that any non-gaussianness of X is always captured by either the norm $\|X\|_2$ or the dot product pairings $\langle X, X' \rangle$. It is clear that the norm condition on its own is not sufficient to guarantee that $X \stackrel{d}{=} g$. For instance, let θ have any non-uniform distribution on the sphere. Then $\|g\|_2 \theta$ has the same radial distribution as g , but is not itself gaussian.

This result by itself does not address the NGCA problem, in which we are looking for non-gaussian *directions* in the distribution of X . To this end, we propose a matrix version of the first gaussian test. Pick a parameter $\alpha > 0$ and consider the *test matrices*

$$\Phi_{X,\alpha} := \frac{1}{Z_\Phi} \mathbb{E} e^{-\alpha \|X\|_2^2} X X^T \quad \text{and} \quad \Psi_{X,\alpha} := \frac{1}{Z_\Psi} \mathbb{E} e^{-\alpha \langle X, X' \rangle} X (X')^T, \quad (2.1)$$

where the normalizing quantities $Z_\Phi = Z_{\Phi,X}(\alpha) := \mathbb{E} e^{-\alpha \|X\|_2^2}$ and $Z_\Psi = Z_{\Psi,X}(\alpha) := \mathbb{E} e^{-\alpha \langle X, X' \rangle}$ resemble partition functions in statistical mechanics.

For a standard gaussian random vector g , a straightforward computation (see Lemma C.3) shows that both test matrices are multiples of the identity, namely

$$\Phi_{g,\alpha} = (2\alpha + 1)^{-1} I_n \quad \text{and} \quad \Psi_{g,\alpha} = \alpha(\alpha^2 - 1)^{-1} I_n. \quad (2.2)$$

Our second test guarantees that the non-gaussianness of X is captured by one of the test matrices, and moreover that their eigenvectors reveal the non-gaussian directions of X .

Theorem 2.2 (Second gaussian test). *Consider a random vector X which follows the isotropic NGCA model (1.1), in which the projection \tilde{X} of X onto E is not gaussian. Then, for any $|\alpha|$ small enough, either $\Phi_{X,\alpha}$ has an eigenvalue not equal to $(2\alpha + 1)^{-1}$ or $\Psi_{X,\alpha}$ has an eigenvalue not equal to $\alpha(\alpha^2 - 1)^{-1}$. Furthermore, all eigenvectors corresponding to such eigenvalues lie in E .²*

In Section 4, we will show how to derive the second gaussian test from the first using a block diagonalization formula for each of the matrices $\Phi_{X,\alpha}$ and $\Psi_{X,\alpha}$. Again, it is easy to see that $\Phi_{X,\alpha}$ is not sufficient by itself to identify non-gaussian directions: Take $X = \|g\|_2 \theta$ as before, and this time that assume that θ is uniform on $\{\pm e_i\}_{i=1}^N$. The symmetry implies that $\Phi_{X,\alpha}$ is a scalar matrix, and computing its trace shows that it is equal to $(2\alpha + 1)^{-1} I_n$.

For simplicity, we stated both gaussian tests for population rather than for finite samples; they involve taking expectations over the entire distribution of X which is typically unknown in practice. However, both tests are quite robust and work provably well on finite (polynomially large) samples. Robust versions of gaussian tests can be formulated in terms of *deviations from gaussian moments*

$$D_{X,r} := \sup_{v \in S^{n-1}} |\mathbb{E}\langle X, v \rangle^r - \mathbb{E}\langle g, v \rangle^r|. \quad (2.3)$$

Note that by the classical moment problem, if $D_{X,r} = 0$ for all positive integers r , then X has the same distribution as g , i.e. the standard gaussian distribution. The magnitude of $D_{X,r}$ can thus serve as a convenient quantitative measure of non-gaussianness of X .

Theorem 2.3 (First gaussian test, robust). *There is a universal constant $c > 0$ such that for each positive integer r , we have either*

$$|\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| \geq c\eta_r^2/\gamma_r \quad \text{or} \quad |\mathbb{E}\langle X, X' \rangle^r - \mathbb{E}\langle g, g' \rangle^r| \geq c\eta_r^2.$$

Here $\gamma_r = \mathbb{E}|\langle g, v \rangle|^r$ for an arbitrary vector $v \in S^{n-1}$ and $\eta_r = \min\{D_{X,r}, \gamma_r\}$.

As with the non-robust version, we will prove this theorem in Section 3 using a decomposition of moment tensors. There is a similar robust version of the second gaussian test, which we will skip here but state and prove in Section 4.

Robustness allows us to use finite sample averages instead of expectations in the gaussian tests, which is critical for practical applications. Indeed, consider a sample $X_1, \dots, X_N, X'_1, \dots, X'_N$ of $2N$ i.i.d. realizations of a random variable X . We can then define the sample versions of the test matrices in (2.1) in an obvious way:

$$\hat{\Phi}_{X,\alpha} = \frac{1}{\hat{Z}_\Phi} \sum_{i=1}^N e^{-\alpha\|X_i\|_2^2} X_i X_i^T \quad \text{and} \quad \hat{\Psi}_{X,\alpha} = \frac{1}{\hat{Z}_\Psi} \sum_{i=1}^N e^{-\alpha\langle X_i, X'_i \rangle} (X_i (X'_i)^T + X'_i X_i^T), \quad (2.4)$$

with the normalizing quantities $\hat{Z}_\Phi := \sum_{i=1}^N e^{-\alpha\|X_i\|_2^2}$ and $\hat{Z}_\Psi := 2 \sum_{i=1}^N e^{-\alpha\langle X_i, X'_i \rangle}$.

The second gaussian test leads to the following straightforward algorithm for solving NGCA problem based on a finite sample: Use the sample to compute the test matrices $\hat{\Phi}_{X,\alpha}$ and $\hat{\Psi}_{X,\alpha}$; select the eigenspaces corresponding to the eigenvalues that significantly deviate from the gaussian

²The matrix $\Phi_{X,\alpha}$ always exists, but when \tilde{X} is not subgaussian (i.e. can be rescaled so that marginals have tails lighter than a standard gaussian), $\Psi_{X,\alpha}$ may not be well-defined even for small α . In that case, $\|\tilde{X}\|_2$ has a different distribution from $\|g\|_2$, so that $\Phi_{X,\alpha}$ has non-gaussian eigenvalues. We can hence think of $\Phi_{X,\alpha}$ as the primary test matrix, and $\Psi_{X,\alpha}$ being an auxiliary that is only required in hard (effectively adversarial) cases.

eigenvalues. Then all vectors in both eigenspaces will be close to the non-gaussian subspace E which we are trying to find. Let us state this algorithm and its guarantees precisely.

Algorithm 1 REWEIGHTED PCA($X, \alpha_1, \alpha_2, \beta_1, \beta_2$)

Input: Data points $X_1, \dots, X_N, X'_1, \dots, X'_N$, scaling parameters $\alpha_1, \alpha_2 \in \mathbb{R}$, tolerance parameters $\beta_1, \beta_2 > 0$.

Output: Two estimates \hat{E}_Φ and \hat{E}_Ψ for E .

- 1: Compute test matrices $\hat{\Phi}_{X, \alpha_1}$ and $\hat{\Psi}_{X, \alpha_2}$.
 - 2: Compute the eigenspace \hat{E}_Φ of $\hat{\Phi}_{X, \alpha_1}$ corresponding to the eigenvalues that are farther than β_1 from the value $(2\alpha_1 + 1)^{-1}$.
 - 3: Compute the eigenspace \hat{E}_Ψ of $\hat{\Psi}_{X, \alpha_2}$ corresponding to the eigenvalues that are farther than β_2 from the value $\alpha_2(\alpha_2^2 - 1)^{-1}$.
-

Theorem 2.4 (Guarantee for Reweighted PCA). *Let X be a subgaussian³ random vector which follows the isotropic NGCA model (1.1), and with subgaussian norm bounded above by $K \geq 1$. Let r be the integer such that $D_{\tilde{X}, r} =: D > 0$ and $D_{\tilde{X}, r'} = 0$ for all $r' < r$. Then for any $\delta, \epsilon \in (0, 1)$, with probability at least $1 - \delta$, if we run REWEIGHTED PCA with a choice of parameters $\alpha_1, \alpha_2, \beta_1, \beta_2$ that is optimal up to constant multiples, at least one of \hat{E}_Φ and \hat{E}_Ψ is non-trivial, and any unit vector in their union is ϵ -close to one in E , so long as the sample size N is greater than $\text{poly}_r(n, 1/\epsilon, \log(1/\delta), 1/D, K)$. Here, poly_r is a polynomial whose total degree depends linearly on r .*

The idea of the proof is to use Davis-Kahan eigenvector perturbation theory [7]. The robust version of the second gaussian test exerts the existence of a gap between gaussian and non-gaussian eigenvalues. By bounding the deviation of the test matrix estimators $\hat{\Phi}_{X, \alpha}$ and $\hat{\Psi}_{X, \alpha}$ from their expectation, we can thus show that their eigenstructures are similar. We will prove this theorem formally in Section 5.

Note that REWEIGHTED PCA is a simple spectral algorithm, which obviously runs in polynomial time. The name of the algorithm stems from the first test matrix, which can be seen as a PCA matrix for the reweighted sample obtained when each point X_i is given the weight $e^{-\alpha \|X_i\|_2^2}$. As mentioned earlier in the section, $\Phi_{X, \alpha}$ reveals at least one non-gaussian direction in all but adversarial situations, and so can be considered the primary test matrix.

2.1 Organization of paper and notation

In Section 3, we will prove the first gaussian test and its robust version. In Section 4, we will prove the second gaussian test and state a robust version needed for proving our guarantee for Reweighted PCA. The guarantee is then proven in Section 5. For the sake of space, many technical details are deferred to the appendix. Throughout the paper, C and c denote absolute constants whose value may change from line to line. We let g_n denote the standard gaussian vector in \mathbb{R}^n . The subscript is omitted whenever the dimension is obvious.

³For a formal definition of subgaussian random vectors and an introduction to their properties, please see [21].

3 Proof of the first gaussian test

The first gaussian test is based on the first author's work on eccentricity tensors in [20]. For completeness, we will repeat the key arguments. The statements and proofs in this section are valid more generally for random variables X with finite moments of all orders (not necessarily subgaussian).

Recall the following fact from linear algebra. For any positive integer r , we may identify the r -th tensor product $T^r(\mathbb{R}^n) = \mathbb{R}^n \otimes \cdots \otimes \mathbb{R}^n$ with \mathbb{R}^{n^r} by picking as a basis the vectors $\{e_{i_1} \otimes e_{i_2} \otimes \cdots \otimes e_{i_r}\}_{1 \leq i_1, \dots, i_r \leq n}$. With this choice, the Euclidean inner product between any two pure tensors $u_1 \otimes \cdots \otimes u_r$ and $v_1 \otimes \cdots \otimes v_r$ can be written as

$$\langle u_1 \otimes \cdots \otimes u_r, v_1 \otimes \cdots \otimes v_r \rangle = \prod_{i=1}^r \langle u_i, v_i \rangle.$$

In particular, for power tensors $u^{\otimes r}$ and $v^{\otimes r}$, we have the formula $\langle u^{\otimes r}, v^{\otimes r} \rangle = \langle u, v \rangle^r$.

Now let X and Y be two independent random vectors. The above formula allows us to rewrite the r -th moment of their inner product as an inner product between their r -th moment tensors. Namely, we have

$$\mathbb{E}\langle X, Y \rangle^r = \mathbb{E}\langle X^{\otimes r}, Y^{\otimes r} \rangle = \langle M_X^r, M_Y^r \rangle, \quad (3.1)$$

where we define $M_X^r := \mathbb{E}X^{\otimes r}$. For independent copies X, X' of the same random vector having distribution μ , $M_X^r = M_{X'}^r$, so

$$\mathbb{E}\langle X, X' \rangle^r = \|M_X^r\|^2. \quad (3.2)$$

Next, for any random vector X , let X_{rot} denote a random vector that is independent of X , has the same radial distribution as X , and whose distribution is rotationally invariant. We call X_{rot} the *rotational symmetrization* of X . Comparing the moment tensors of a random vector and those of its rotational symmetrization give rise to what we shall call eccentricity tensors. Specifically, for any positive integer r , we define the r -th *eccentricity tensor* of X to be $E_X^r = M_X^r - M_{X_{rot}}^r$.

Since $X \stackrel{d}{=} X_{rot}$ if and only if X is rotationally invariant, we see that the eccentricity tensors of X are quantitative measures of how far its distribution is from being rotationally invariant. This interpretation is further supported by the following observation.

Lemma 3.1 (Orthogonality). *The eccentricity tensors of a random vector X are orthogonal to the moment tensors of its rotational symmetrization. In other words, for any positive integer r ,*

$$\langle E_X^r, M_{X_{rot}}^r \rangle = 0 \quad \text{and} \quad \|M_X^r\|_2^2 = \|M_{X_{rot}}^r\|_2^2 + \|E_X^r\|_2^2. \quad (3.3)$$

Proof. Let Q be a random orthogonal matrix chosen according to the Haar measure on $O(n)$. For any fixed vector $v \in \mathbb{R}^n$, Qv is uniformly distributed on the sphere of radius $\|v\|_2$, so if Y is any random vector independent of Q , applying Q to Y preserves its radial distribution but makes QY rotationally invariant.

Now choose Q to be independent of X and X_{rot} . Our previous discussion implies that $Q^T X \stackrel{d}{=} QX_{rot} \stackrel{d}{=} X_{rot}$. We use this to compute

$$\mathbb{E}\langle X, X_{rot} \rangle^r = \mathbb{E}\langle X, QX_{rot} \rangle^r = \mathbb{E}\langle Q^T X, X_{rot} \rangle^r = \mathbb{E}\langle X_{rot}, X_{rot} \rangle^r, \quad (3.4)$$

where X'_{rot} is an independent copy of X_{rot} . We may then apply identities (3.1) and (3.2) to rewrite the above equation as

$$\langle M_X^r, M_{X_{rot}}^r \rangle = \langle M_{X_{rot}}^r, M_{X_{rot}}^r \rangle. \quad (3.5)$$

Subtracting the right hand side from the left hand side gives (3.3). \square

Theorem 3.2. *Let X be a random vector in \mathbb{R}^n with finite moments of all orders. Then*

- a) (Minimization) *If X' is an independent copy of X , and X_{rot}, X'_{rot} are independent copies of its rotational symmetrization, we have*

$$\mathbb{E}\langle X, X' \rangle^r \geq \mathbb{E}\langle X_{rot}, X'_{rot} \rangle^r \quad (3.6)$$

for any positive integer r .

- b) (Uniqueness) *Furthermore, if equality holds in (3.6) for all r and we further assume that X has a subexponential distribution⁴, then X is rotationally invariant.*

Proof. Using identity (3.2), we rewrite the first claim as

$$\|M_X^r\|_2^2 \geq \|M_{X_{rot}}^r\|_2^2,$$

and this follows immediately from equation (3.3).

If equality holds for all positive integers r , then by (3.3), $E_X^r = 0$ for all r , implying that X and X_{rot} have the same moment tensors of all orders. Since exponential random variables are characterized by their moments (see Lemma A.1), X and X_{rot} have the same distribution. \square

Proof of Theorem 2.1. If X has the same radial distribution as g , then g is the rotational symmetrization of X . The claim is then a direct application of the uniqueness portion of Theorem 3.2. \square

We now move on to proving the robust version of the test, namely Theorem 2.3.

Lemma 3.3. *Let X be a random vector in \mathbb{R}^n . Let θ be uniformly distributed on the sphere S^{n-1} . Then the following hold for any positive integer r :*

- a) $M_{X_{rot}}^r = \mathbb{E}\|X\|_2^r M_\theta^r$.
b) $\|E_X^r\|_2^2 = (\mathbb{E}\langle X, X' \rangle^r) - (\mathbb{E}\|X\|_2^r)^2 (\mathbb{E}\langle \theta, \theta' \rangle^r)$.
c) *For any unit vector $v \in \mathbb{R}^n$,*

$$|\mathbb{E}\langle X, v \rangle^r - \mathbb{E}\langle g, v \rangle^r| \leq |\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| \mathbb{E}\langle \theta, \theta' \rangle^r + \left(\mathbb{E}\langle X, X' \rangle^r - (\mathbb{E}\|X\|_2^r)^2 \mathbb{E}\langle \theta, \theta' \rangle^r \right)^{1/2}. \quad (3.7)$$

- d) *In particular, when r is odd,*

$$|\mathbb{E}\langle X, v \rangle^r - \mathbb{E}\langle g, v \rangle^r| \leq (\mathbb{E}\langle X, X' \rangle^r)^{1/2} = |\mathbb{E}\langle X, X' \rangle^r - (\mathbb{E}\langle g, g' \rangle^r)|^{1/2}. \quad (3.8)$$

Proof. Deferred to Appendix A. \square

⁴For an introduction to the properties of subexponential distributions, we again refer the reader to [21].

By balancing the two terms on the right hand side in part c), we obtain the following lemma, whose proof is again deferred to Appendix A.

Lemma 3.4. *Let X be a random vector in \mathbb{R}^n for $n \geq 2$. Suppose there is a unit vector $v \in S^{n-1}$, an even integer $r \geq 2$, and a positive number $0 < \delta \leq 1$ such that $|\mathbb{E}\langle X, v \rangle^r - \mathbb{E}\langle g, v \rangle^r| \geq \delta \mathbb{E}\langle g, v \rangle^r$. Then either*

$$|\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| \geq \frac{\delta^2}{4} \mathbb{E}\langle g, v \rangle^r \quad \text{or} \quad |\mathbb{E}\langle X, X' \rangle^r - \mathbb{E}\langle g, g' \rangle^r| \geq \frac{15\delta^2}{64} (\mathbb{E}\langle g, v \rangle^r)^2.$$

Proof of Theorem 2.3. If r is odd, then the statement follows from (3.8). If r is even, set $\delta = \frac{\Delta}{\mathbb{E}\langle g, v \rangle^r}$ in the previous theorem. \square

4 Proof of the second gaussian test

In this section, we return to the setting where X follows the NGCA model (1.1). We further assume that the non-gaussian component \tilde{X} is a subgaussian random vector with subgaussian norm bounded by K . In order not to break the flow of the paper, most of the proofs are deferred to Appendix B.

The first step in proving the test is to notice that the independence of the gaussian and non-gaussian components allows us to block diagonalize the test matrices.

Lemma 4.1 (Block diagonalization for $\Phi_{X,\alpha}$ and $\Psi_{X,\alpha}$). *Assume E is spanned by the first d basis vectors. Then the test matrices $\Phi_{X,\alpha}$ and $\Psi_{X,\alpha}$ decompose into blocks in the following manner:*

$$\Phi_{X,\alpha} = \left(\begin{array}{c|c} \Phi_{\tilde{X},\alpha} & 0 \\ \hline 0 & \Phi_{g,\alpha} \end{array} \right), \quad \Psi_{X,\alpha} = \left(\begin{array}{c|c} \Psi_{\tilde{X},\alpha} & 0 \\ \hline 0 & \Psi_{g,\alpha} \end{array} \right). \quad (4.1)$$

We then observe that the trace of the test matrices are conveniently equal to the negated log derivatives of their respective partition functions.

Lemma 4.2 (Trace of $\Phi_{Y,\alpha}$ and $\Psi_{Y,\alpha}$). *Let Y be any random vector in \mathbb{R}^n . Then $\text{tr}(\Phi_{Y,\alpha}) = -(\log Z_{\Phi,Y})'(\alpha)$ and $\text{tr}(\Psi_{Y,\alpha}) = -(\log Z_{\Psi,Y})'(\alpha)$.*

Our next lemma shows that for α small enough, the partition functions themselves differentiate between gaussian and non-gaussian random vectors. This is obvious once we realize that they are just the moment generating functions of $\|X\|_2^2$ and $\langle X, X' \rangle$, and that these are analytic in a small neighborhood around 0.

Lemma 4.3 (Partition functions characterize gaussian distributions). *The following hold for any subgaussian random vector Y :*

- a) *If $Z_{\Phi,Y}(\alpha_k) = Z_{\Phi,g}(\alpha_k)$ for a sequence of values α_k converging to 0, then Y has the same radial distribution as g .*
- b) *If in addition, $Z_{\Psi,Y}(\beta_k) = Z_{\Psi,g}(\beta_k)$ for a sequence of values β_k converging to 0, then X has the standard gaussian distribution.*

We are now in a position to prove the second gaussian test.

Proof of Theorem 2.2. Let g_d denote the standard gaussian in \mathbb{R}^d . By Lemma 4.3, either $Z_{\Phi, \tilde{X}}(\alpha) \neq Z_{\Phi, g_d}(\alpha)$ for $|\alpha|$ small enough, or $Z_{\Psi, \tilde{X}}(\alpha) \neq Z_{\Psi, g_d}(\alpha)$ for $|\alpha|$ small enough. As such, either $(\log Z_{\Phi, \tilde{X}})'(\alpha) \neq (\log Z_{\Phi, g_d})'(\alpha)$ or $(\log Z_{\Psi, \tilde{X}})'(\alpha) \neq (\log Z_{\Psi, g_d})'(\alpha)$. Assume the former holds, and let $\lambda_1, \dots, \lambda_n$ denote the eigenvalues of $\Phi_{X, \alpha}$. Since we may write $\Phi_{X, \alpha}$ in a block form, these eigenvalues are either those of $\Phi_{\tilde{X}, \alpha}$ or $\Phi_{g, \alpha}$. Without loss of generality, we may assume that $\lambda_1, \dots, \lambda_d$ are the eigenvalues of $\Phi_{\tilde{X}, \alpha}$, and $\lambda_{d+1}, \dots, \lambda_n$ are those of $\Phi_{g, \alpha}$.

Lemma C.3 tells us that $\lambda_{d+1} = \dots = \lambda_n = (2\alpha + 1)^{-1}$. On the other hand, by Lemma 4.2,

$$\sum_{i=1}^d \lambda_i = \text{tr}(\Phi_{\tilde{X}, \alpha}) = -(\log Z_{\Phi, \tilde{X}})'(\alpha).$$

By Lemma C.2, $-(\log Z_{\Phi, g_d})'(\alpha) = d(2\alpha + 1)^{-1}$, so we have $\sum_{i=1}^d \lambda_i \neq d(2\alpha + 1)^{-1}$. Dividing through by d , we get $\frac{1}{d} \sum_{i=1}^d \lambda_i \neq (2\alpha + 1)^{-1}$, which implies that at least one λ_i differs from this value for $1 \leq i \leq d$.

If it were the case that $(\log Z_{\Psi, \tilde{X}})'(\alpha) \neq (\log Z_{\Psi, g_d})'(\alpha)$, a similar argument involving $\Psi_{X, \alpha}$ gives the alternate conclusion. \square

It is tedious but not too difficult to make the second gaussian test quantitative. We do this by tracking how the non-gaussian moments for $\|\tilde{X}\|_2$ and $\langle \tilde{X}, \tilde{X}' \rangle$ contribute to the power series expansions for $-(\log Z_{\Phi, \tilde{X}})'$ and $-(\log Z_{\Psi, \tilde{X}})'$ around 0. This yields the following theorem.

Theorem 4.4 (Second gaussian test, robust). *Let r be the integer such that $D_{\tilde{X}, r} > 0$ and $D_{\tilde{X}, r'} = 0$ for all $r' < r$. Then either*

a) *for $|\alpha| \leq \eta_r^2 r / (CK^2)^r (d^{r+1} + (r+1)!)$, we have*

$$\left| \frac{1}{d} \sum_{i=1}^d \lambda_i(\Psi_{\tilde{X}, \alpha}) - \frac{\alpha}{\alpha^2 - 1} \right| \geq \frac{c\eta_r^2}{d(r-1)!} |\alpha|^{r-1}, \quad (4.2)$$

b) *or for $|\alpha| \leq \eta_r^2 r / (CK^2)^{r/2} \gamma_r (d^{r/2+1} + (r/2+1)!)$, we have*

$$\left| \frac{1}{d} \sum_{i=1}^d \lambda_i(\Phi_{\tilde{X}, \alpha}) - \frac{1}{2\alpha + 1} \right| \geq \frac{c\eta_r^2}{d(r/2-1)! \gamma_r} |\alpha|^{r/2-1}. \quad (4.3)$$

Here $\gamma_r = \mathbb{E}|\langle g, v \rangle|^r$ for an arbitrary vector $v \in S^{n-1}$ and $\eta_r = \min\{D_{X, r}, \gamma_r\}$.

5 Proof of guarantee for Reweighted PCA

The second gaussian test tells us how we can recover non-gaussian directions from $\Phi_{X, \alpha}$ and $\Psi_{X, \alpha}$. Our guarantee for Reweighted PCA algorithm shows that we can do the same with the plug-in estimators $\hat{\Phi}_{X, \alpha}$ and $\hat{\Psi}_{X, \alpha}$. To this end, we first provide concentration bounds for these estimators, whose proofs can be found in Appendix D.

Theorem 5.1 (Concentration for $\hat{\Phi}_{X, \alpha}$). *There is an absolute constant C such that for any $0 < \epsilon, \delta < 1$, and any $0 \leq \alpha < 1/CK^2 n$, we have $\mathbb{P}\left(\|\hat{\Phi}_{X, \alpha} - \Phi_{X, \alpha}\| > \epsilon\right) \leq \delta$ so long as $N \geq CK^2(n + \log(1/\delta))\epsilon^{-2}$.*

Theorem 5.2 (Concentration for $\hat{\Psi}_{X,\alpha}$). *There is an absolute constant C such that for any $0 < \epsilon, \delta < 1$, if $N \geq CK^2(n + \log(1/\delta))\epsilon^{-2}$ and $|\alpha| \leq 1/CK^2\tau(n + \tau)$, we have $\mathbb{P}(\|\hat{\Psi}_{X,\alpha} - \Psi_{X,\alpha}\| > \epsilon) \leq \delta$. Here, $\tau = \log^{1/2}(N/\min\{\delta, K\epsilon\})$.*

We use the following notion of distance between subspaces.

Definition 5.3 (Subspace distance). Let F and F' be subspaces of \mathbb{R}^n of dimensions m . Let U and U' be matrices whose columns form an orthonormal basis for F and F' respectively. The distance between F and F' is defined to be $d(F, F') := \|UU^T - U'(U')^T\|_F$.

Lemma 5.4 (Guarantee for \hat{E}_Φ). *Suppose the moments of $\|\tilde{X}\|_2^2$ and $\|g_d\|_2^2$ agree up to order $r-1$, but there is a number $\Delta > 0$ such that $|\mathbb{E}\|\tilde{X}\|_2^{2r} - \mathbb{E}\|g_d\|_2^{2r}| \geq \Delta$. For any $\delta, \epsilon \in (0, 1)$, pick α_1 such that $0 < \alpha_1 < \min\{\Delta r/(CK^2)^r(d^{r+1} + (r+1)!), 1/CK^2n\}$, and $\beta_1 = \Delta\alpha_1^{r-1}/4d(r-1)!$. Then with probability at least $1 - \delta$, REWEIGHTED PCA with $2N \geq CK^2d^{3/2}(n + \log(1/\delta))/\beta_1^2\epsilon^2$ samples together with this choice of α_1 and β_1 produces a nontrivial estimate \hat{E}_Φ of dimension $1 \leq \hat{d}_\Phi \leq d$, such that there is a \hat{d}_Φ -dimensional subspace $E_\Phi \subset E$ satisfying $d(\hat{E}_\Phi, E_\Phi) \leq \epsilon$.*

Proof. Combining Lemmas 4.1, 4.2, and C.3 tells us that in the right coordinates, $\Phi_{X,\alpha}$ block diagonalizes as

$$\Phi_{X,\alpha} = \left(\begin{array}{c|c} \Phi_{\tilde{X},\alpha} & 0 \\ \hline 0 & (2\alpha + 1)^{-1}I_{n-d} \end{array} \right). \quad (5.1)$$

Next, label the eigenvalues of Φ_{X,α_1} as $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$. We can find $0 \leq p \leq q \leq n$ such that the eigenvalues corresponding to the $\Phi_{\tilde{X},\alpha_1}$ block are $\lambda_1, \lambda_2, \dots, \lambda_p, \lambda_{q+1}, \dots, \lambda_n$. Using Theorem B.4, we then have

$$\left| \frac{1}{d} \left(\sum_{i=1}^p \lambda_i + \sum_{i=q+1}^n \lambda_i \right) - \frac{1}{2\alpha_1 + 1} \right| \geq \frac{\Delta}{2d(r-1)!} \alpha_1^{r-1} = 2\beta_1. \quad (5.2)$$

In particular, we have $\frac{1}{p} \sum_{i=1}^p \lambda_i - 1/(2\alpha_1 + 1) \geq 2\beta_1$, and $1/(2\alpha_1 + 1) - \frac{1}{n-q} \sum_{i=q+1}^n \lambda_i \geq 2\beta_1$. Since at least one of these sums of eigenvalues is non-empty, truncating the eigenvalues of Φ_{X,α_1} at the β_1 level gives us a non-trivial subspace of E .

In order to show that our empirical estimate $\hat{\Phi}_{X,\alpha_1}$ also has an approximation to this property, we will need to use the eigenvector perturbation theory explained in Appendix E. First, we need to bound from below the ‘‘eigengap’’ in Φ_{X,α_1} . Suppose first that $p \geq 1$, i.e. that there are eigenvalues larger than $(2\alpha_1 + 1)^{-1}$. Then by the pigeonhole principle, one can find i such that $(2\alpha_1 + 1)^{-1} + \beta_1/2 \geq \lambda_{i+1} \geq (2\alpha_1 + 1)^{-1}$ and $\lambda_i - \lambda_{i+1} \geq \beta_1/2d$. Similarly, if $q \leq n-1$, then we can find j such that $(2\alpha_1 + 1)^{-1} \geq \lambda_{j-1} \geq (2\alpha_1 + 1)^{-1} - \beta_1/2$ and $\lambda_{j-1} - \lambda_j \geq \beta_1/2d$.

Now let F be the span of the eigenvectors of Φ_{X,α_1} corresponding to $\lambda_1, \dots, \lambda_i, \lambda_j, \dots, \lambda_n$, and let \hat{F} be the eigenvectors of $\hat{\Phi}_{X,\alpha_1}$ corresponding to $\hat{\lambda}_1, \dots, \hat{\lambda}_i, \hat{\lambda}_j, \dots, \hat{\lambda}_n$. By Theorem 5.1, with probability at least $1 - \delta$, we have

$$\|\hat{\Phi}_{X,\alpha} - \Phi_{X,\alpha}\| \leq \frac{\beta_1\epsilon}{4\sqrt{2}d^{3/2}}. \quad (5.3)$$

We may then use Theorem E.4 to see that $d(\hat{F}, F) \leq \epsilon$.

We are not yet done, because we do not have access to \hat{F} . Nonetheless, we can show that \hat{F} contains \hat{E}_Φ . Using eigenvalue perturbation inequalities together with equation (5.3) tells us that we have

$$\hat{\lambda}_{i+1} \leq \lambda_{i+1} + \frac{\beta_1 \epsilon}{2d} \leq (2\alpha_1 + 1)^{-1} + \frac{\beta_1}{2} + \frac{\beta_1 \epsilon}{2d} \leq (2\alpha_1 + 1)^{-1} + \frac{2\beta_1}{3}, \quad (5.4)$$

and similarly that

$$\hat{\lambda}_{j-1} \leq \lambda_{j-1} - \frac{\beta_1 \epsilon}{2d} \leq (2\alpha_1 + 1)^{-1} - \frac{\beta_1}{2} - \frac{\beta_1 \epsilon}{2d} \leq (2\alpha_1 + 1)^{-1} - \frac{2\beta_1}{3}. \quad (5.5)$$

Let $\hat{I}_\Phi = \{i : |\hat{\lambda}_i - (1 - 2\alpha_1)^{-1}| > \beta_1\}$. We see that this set does not contain any index between $i + 1$ and $j - 1$, so \hat{E}_Φ , which comprises the span of the eigenvectors to these eigenvalues, does not contain any eigenvector that \hat{F} does not contain, as was to be shown. The inclusion then implies that we may find a subspace $E_\Phi \subset F$ such that $d(\hat{E}_\Phi, E_\Phi) \leq \epsilon$.

Finally, we observe that $\dim \hat{E}_\Phi \geq 1$, since

$$\frac{1}{p} \sum_{i=1}^p \hat{\lambda}_i - \frac{1}{2\alpha_1 + 1} \geq \frac{1}{p} \sum_{i=1}^p \lambda_i - \frac{\beta_1 \epsilon}{2d} - \frac{1}{2\alpha_1 + 1} > \beta_1, \quad (5.6)$$

and

$$\frac{1}{2\alpha_1 + 1} - \frac{1}{n - q} \sum_{i=q+1}^n \hat{\lambda}_i \geq \frac{1}{2\alpha_1 + 1} - \frac{1}{n - q} \sum_{i=q+1}^n \lambda_i - \frac{\beta_1 \epsilon}{2d} > \beta_1. \quad (5.7)$$

□

Lemma 5.5 (Guarantee for \hat{E}_Ψ). *Suppose the moments of $\langle X, X' \rangle$ and $\langle g, g' \rangle$ agree up to order $r - 1$ but $|\mathbb{E}\langle X, X' \rangle^r - \mathbb{E}\langle g, g' \rangle^r| \geq \Delta$. For any $\delta, \epsilon, \tau \in (0, 1)$, pick $0 < \alpha_2 < \min\{\Delta r / (CK^2)^r (d^{r+1} + (r + 1)!), 1 / CK^2 n^{1+\tau}\}$, and $\beta_2 = \Delta \alpha_2^{r-1} / 4d(r - 1)!$. Then with probability at least $1 - \delta$, REWEIGHTED PCA with sample size $2N$ satisfying $\exp(n^{2\tau}) \min\{\delta, K\epsilon\} \geq 2N \geq CK^2 d^{3/2} (n + \log(1/\delta)) / \beta_2^2 \epsilon^2$, together with this choice of α_2 and β_2 produces a nontrivial estimate \hat{E}_Ψ of dimension $1 \leq \hat{d}_\Psi \leq d$, such that there is a \hat{d}_Ψ -dimensional subspace $E_\Psi \subset E$ satisfying $d(\hat{E}_\Psi, E_\Psi) \leq \epsilon$.*

Proof. The proof is completely analogous to that for the previous theorem, except that we replace our estimates and identities for Φ_{X, α_1} with those for Ψ_{X, α_2} wherever necessary. □

Proof of Theorem 2.4. Combine the last two lemmas with Theorem 4.4 from the last section. □

Remark 5.6 (Selecting optimal parameters). If the problem parameters d, D, K and r were known before hand, then in principle, one could compute the optimal tuning parameters $\alpha_1, \alpha_2, \beta_1, \beta_2$. In practice, however, one rarely is in this situation, so one would have to estimate the problem parameters as a first step to solving the NGCA problem. Nonetheless, one can do this by the doubling/halving trick. In other words, we start with some fixed initial choice of α_1 and α_2 . Using Theorems 5.1 and 5.2, we can detect whether there are any outlier eigenvalues with high probability. If there are none, we halve α_1 and α_2 and try again, repeating this process until outliers show up. The number of iterations is then the base 2 logarithm of the final α_1 and α_2 , plus an additive constant. This is logarithmic in d, D , and K , and polynomial in r , so the algorithm remains efficient.

6 Discussion

We have presented and analyzed an algorithm that is guaranteed to return at least one non-gaussian direction efficiently, with sample and time complexity a polynomial in the problem parameters for a fixed r , where r is the smallest order at which the moments of any marginal of \tilde{X} differ from those of a standard gaussian.

Since the degree of the polynomial increases linearly in r , it would seem that the algorithm is practically useless if r is larger than a small constant. However, note that having all third and fourth moments equal those of a gaussian is a condition that is already stringent in one dimension, and which becomes even more so in higher dimensions. As such, unless \tilde{X} has some kind of adversarial distribution, r will be either 4 or 3, depending on whether \tilde{X} is centrally symmetric or not.

The algorithm also often delivers much more than is guaranteed for several reasons. First, in order to bound the subspace perturbation by ϵ , we used a very crude estimate of the eigengap, bounding it from below using the pigeonhole principle, which in the worst case assumes that the eigenvalues are spread out at regular intervals. This should not happen in practice, and we expect the non-gaussian eigenvalues to instead cluster relatively tightly around their average. If this happens, the sample complexity requirement can be relaxed by a factor of d .

Second, just as it is extremely unlikely for r to be higher than 4, for a general non-gaussian \tilde{X} and a small, random α , it is extremely unlikely for any of the non-gaussian values of $\Phi_{X,\alpha}$ to be equal to the gaussian one on the dot. This means that even though the guarantee is for one direction, in practice we most probably can recover the entire subspace E simultaneously with just $\hat{\Phi}_{X,\alpha}$ alone, albeit with a more sophisticated truncation technique. Indeed, we can provably recover all non-gaussian directions when \tilde{X} is an unconditional random vector, examples of which include the uniform distributions on spheres and hypercubes.

6.1 Conjectures and questions

We conjecture that Reweighted PCA actually recovers the entire non-gaussian subspace E , so long as \tilde{X} does not have any gaussian marginals. The first gaussian test for a random vector X using the distribution of its norm and dot product pairing also leads to further questions. For a fixed nonzero real number t , both of these appear in the formula for $\|Y_t\|_2^2$, where we set $Y_t := X + tX'$, so it is natural to ask whether Reweighted PCA works with $\Phi_{Y_t,\alpha}$ alone for some t . In particular, does it work for $t = -1$? It is also an open question whether $\langle X, X' \rangle$ alone is sufficient to test whether X is standard gaussian.

Acknowledgements

Both authors are partially supported by NSF Grant DMS 1265782 and U.S. Air Force Grant FA9550-14-1-0009.

References

- [1] ARORA, S., GE, R., MOITRA, A., AND SACHDEVA, S. Provable ICA with Unknown Gaussian Noise, and Implications for Gaussian Mixtures and Autoencoders. *Algorithmica* 72, 1 (2015), 215–236.

- [2] BILLINGSLEY, P. *Probability and Measure - Third Edition*. 1995.
- [3] BLANCHARD, G., KAWANABE, M., SUGIYAMA, M., SPOKOINY, V., AND MÜLLER, K.-R. In Search of Non-Gaussian Components of a High-Dimensional Distribution. *Journal of Machine Learning Research* 7, 2 (2006), 247–282.
- [4] BRUBAKER, S. C., AND VEMPALA, S. S. Isotropic PCA and affine-invariant clustering. *Proceedings - Annual IEEE Symposium on Foundations of Computer Science, FOCS* (2008), 551–560.
- [5] ÇINLAR, E. *Probability and stochastics*, vol. 230. 2011.
- [6] COMON, P. Independent component analysis, A new concept? *Signal Processing* 36, 3 (1994), 287–314.
- [7] DAVIS, C., AND KAHAN, W. M. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis* 7, 1 (1970), 1–46.
- [8] DIEDERICH, E., JUDITSKY, A., NEMIROVSKI, A., AND SPOKOINY, V. Sparse non Gaussian component analysis by semidefinite programming. *Machine Learning* 91, 2 (2013), 211–238.
- [9] DIEDERICH, E., JUDITSKY, A., SPOKOINY, V., AND SCHÜTTE, C. Sparse non-Gaussian component analysis. *IEEE Transactions on Information Theory* 56, 6 (2010), 3033–3047.
- [10] FRIEZE, A., JERRUM, M., AND KANNAN, R. Learning linear transformations. *Proceedings of 37th Conference on Foundations of Computer Science* (1996), 359–368.
- [11] GOYAL, N., VEMPALA, S., AND XIAO, Y. Fourier PCA and robust tensor decomposition. In *Proceedings of the 46th Annual ACM Symposium on Theory of Computing - STOC '14* (New York, New York, USA, 2014), no. c, ACM Press, pp. 584–593.
- [12] HUBER, P. J. Projection Pursuit. *The Annals of Statistics* 13, 2 (1985), 435–475.
- [13] HYVARINEN, A. Fast and robust fixed-point algorithms for independent component analysis. *IEEE Transactions on Neural Networks* 10, 3 (May 1999), 626–634.
- [14] KAWANABE, M., SUGIYAMA, M., BLANCHARD, G., AND MÜLLER, K.-R. A new algorithm of non-Gaussian component analysis with radial kernel functions. *Annals of the Institute of Statistical Mathematics* 59, 1 (2006), 57–75.
- [15] KAWANABE, M., AND THEIS, F. J. Estimating non-Gaussian subspaces by characteristic functions. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 3889 LNCS (2006), 157–164.
- [16] KAWANABE, M., AND THEIS, F. J. Joint low-rank approximation for extracting non-Gaussian subspaces. *Signal Processing* 87, 8 (2007), 1890–1903.
- [17] M. KAWANABE. Linear dimension reduction based on the fourth-order cumulant tensor. *Proc. of Artificial Neural Networks – ICANN 2005* (2005), 151–156.

- [18] SASAKI, H., HYVÄRINEN, A., AND SUGIYAMA, M. Clustering via mode seeking by direct estimation of the gradient of a log-density. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* 8726 LNAI, PART 3 (2014), 19–34.
- [19] SASAKI, H., NIU, G., AND SUGIYAMA, M. Non-Gaussian Component Analysis with Log-Density Gradient Estimation. *International Conference on Artificial Intelligence and Statistics* 51 (2016), 1–20.
- [20] TAN, Y. S. Energy optimization for distributions on the sphere and improvement to the Welch bounds. Available at <http://arxiv.org/abs/1612.06343>.
- [21] VERSHYNIN, R. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge University Press, Cambridge, 2011, pp. 210–268.
- [22] VIRTÄ, J., NORDHAUSEN, K., AND OJA, H. Projection Pursuit for non-Gaussian Independent Components. Available at <https://arxiv.org/abs/1612.05445>.
- [23] YU, Y., WANG, T., AND SAMWORTH, R. J. A useful variant of the Davis-Kahan theorem for statisticians. *Biometrika* 102, 2 (2015), 315–323.

A Details for Section 3

Lemma A.1. *Let X be a subexponential random vector in \mathbb{R}^n . Then the distribution of X is determined by its moment tensors.*

Proof. Let $\phi_X(v) = \mathbb{E}e^{i\langle X, v \rangle}$ denote the characteristic function of X , and let $K = \|X\|_{\psi_1}$ denote the subexponential norm of X . We then have the following moment growth condition [21]:

$$\sup_{v \in S^{n-1}} \limsup_{r \rightarrow \infty} \frac{(\mathbb{E}|\langle X, v \rangle|^r)^{1/r}}{r} \lesssim K. \quad (\text{A.1})$$

This condition implies that for each $v \in S^{n-1}$, the function $t \mapsto \mathbb{E}e^{it\langle X, v \rangle}$ can be written as a power series with coefficients $\frac{\mathbb{E}\langle X, v \rangle^r}{r!}$ [2], so $\phi_X(v)$ is determined by the moments $\mathbb{E}\langle X, v \rangle^r$. By (3.1), $\mathbb{E}\langle X, v \rangle^r = \langle M_X^r, v^{\otimes r} \rangle$, so these are functions of the moment tensors. Finally, it is a fact from elementary probability that X is determined by its characteristic function [5]. \square

Proof of Lemma 3.3. For the first statement, observe that $X_{rot} = \|X\|_2 \theta$, with $\|X\|_2$ and θ independent. We thus have

$$M_{X_{rot}}^r = \mathbb{E}(\|X\|_2 \theta)^{\otimes r} = \mathbb{E}\|X\|_2^r \mathbb{E}\theta^{\otimes r} = \mathbb{E}\|X\|_2^r M_\theta^r.$$

Next, rewrite (3.3) as $\|E_X^r\|_2^2 = \|M_X^r\|_2^2 - \|M_{X_{rot}}^r\|_2^2$. By definition, we have $\|M_X^r\|_2^2 = \mathbb{E}\langle X, X' \rangle^r$ and using a), we get $\|M_{X_{rot}}^r\|_2^2 = (\mathbb{E}\|X\|_2^r)^2 \mathbb{E}\langle \theta, \theta' \rangle^r$.

To prove part c), fix v and write

$$\mathbb{E}\langle X, v \rangle^r - \mathbb{E}\langle g, v \rangle^r = \langle M_X^r - M_g^r, v^{\otimes r} \rangle = \langle M_{X_{rot}}^r - M_g^r, v^{\otimes r} \rangle + \langle E_X^r, v^{\otimes r} \rangle.$$

We use a) to write

$$\langle M_{X_{rot}}^r - M_g^r, v^{\otimes r} \rangle = \langle \mathbb{E}\|X\|_2^r M_\theta^r - \mathbb{E}\|g\|_2^r M_\theta^r, v^{\otimes r} \rangle = (\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r) \mathbb{E}\langle \theta, v \rangle^r.$$

Notice that $\mathbb{E}\langle \theta, v \rangle^r = \mathbb{E}\langle \theta, \theta' \rangle^r$. We then combine the last two equations with b) and Cauchy-Schwarz to get (3.7). Finally, to get the last claim, we use the fact that $\mathbb{E}\langle \theta, \theta' \rangle^r = \mathbb{E}\langle g, g' \rangle^r = 0$ whenever r is odd. \square

Proof of Lemma 3.4. Observe that (3.7) gives the bound

$$\delta \mathbb{E}\langle g, v \rangle^r \leq |\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| \mathbb{E}\langle \theta, \theta' \rangle^r + \left(\mathbb{E}\langle X, X' \rangle^r - (\mathbb{E}\|X\|_2^r)^2 \mathbb{E}\langle \theta, \theta' \rangle^r \right)^{1/2}. \quad (\text{A.2})$$

Suppose $|\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| \leq \frac{\delta^2}{4} \mathbb{E}\langle g, v \rangle^r$. Then the second term on the right in equation (A.2) has to be large. Indeed, since $\delta \leq 1$ and $\mathbb{E}\langle \theta, \theta' \rangle^r \leq 1/2$ for $r, n \geq 2$, we have

$$\begin{aligned} \left(\mathbb{E}\langle X, X' \rangle^r - (\mathbb{E}\|X\|_2^r)^2 \mathbb{E}\langle \theta, \theta' \rangle^r \right)^{1/2} &\geq \delta \mathbb{E}\langle g, v \rangle^r - \frac{\delta^2}{4} \mathbb{E}\langle \theta, \theta' \rangle^r \mathbb{E}\langle g, v \rangle^r \\ &\geq \frac{7\delta}{8} \mathbb{E}\langle g, v \rangle^r. \end{aligned}$$

Now, applying the fact that $\mathbb{E}\langle g, g' \rangle^r = (\mathbb{E}\|g\|_2^r)^2 \mathbb{E}\langle \theta, \theta' \rangle^r$, we use the reverse triangle inequality and the above bound to write

$$\begin{aligned} |\mathbb{E}\langle X, X' \rangle^r - \mathbb{E}\langle g, g' \rangle^r| &\geq \left| \mathbb{E}\langle X, X' \rangle^r - (\mathbb{E}\|X\|_2^r)^2 \mathbb{E}\langle \theta, \theta' \rangle^r \right| - \left| (\mathbb{E}\|X\|_2^r)^2 - (\mathbb{E}\|g\|_2^r)^2 \right| \mathbb{E}\langle \theta, \theta' \rangle^r \\ &\geq \left(\frac{7\delta}{8} \mathbb{E}\langle g, v \rangle^r \right)^2 - \left| (\mathbb{E}\|X\|_2^r)^2 - (\mathbb{E}\|g\|_2^r)^2 \right| \mathbb{E}\langle \theta, \theta' \rangle^r. \end{aligned} \quad (\text{A.3})$$

Next, notice that

$$\begin{aligned} \left| (\mathbb{E}\|X\|_2^r)^2 - (\mathbb{E}\|g\|_2^r)^2 \right| &= |\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| (\mathbb{E}\|X\|_2^r + \mathbb{E}\|g\|_2^r) \\ &= |\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r| \cdot 2\mathbb{E}\|g\|_2^r + (\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r)^2, \end{aligned}$$

so by the assumption on $|\mathbb{E}\|X\|_2^r - \mathbb{E}\|g\|_2^r|$, we have

$$\begin{aligned} \left| (\mathbb{E}\|X\|_2^r)^2 - (\mathbb{E}\|g\|_2^r)^2 \right| \mathbb{E}\langle \theta, \theta' \rangle^r &\leq \frac{\delta^2}{4} \mathbb{E}\langle g, v \rangle^r \cdot 2\mathbb{E}\|g\|_2^r \cdot \mathbb{E}\langle \theta, \theta' \rangle^r + \left(\frac{\delta^2}{4} \mathbb{E}\langle g, v \rangle^r \right)^2 \mathbb{E}\langle \theta, \theta' \rangle^r \\ &= \frac{\delta^2}{2} (\mathbb{E}\langle g, v \rangle^r)^2 + \left(\frac{\delta^2}{4} \mathbb{E}\langle g, v \rangle^r \right)^2 \mathbb{E}\langle \theta, \theta' \rangle^r \\ &\leq \frac{17\delta^2}{32} (\mathbb{E}\langle g, v \rangle^r)^2. \end{aligned} \quad (\text{A.4})$$

We can now substitute (A.4) into (A.3) to get

$$|\mathbb{E}\langle X, X' \rangle^r - \mathbb{E}\langle g, g' \rangle^r| \geq \frac{15\delta^2}{64} (\mathbb{E}\langle g, v \rangle^r)^2. \quad \square$$

B Details for Section 4

Proof of Lemma 4.1. The decompositions follow easily from the independence of the two components of the mixed vector, \tilde{X} and g , as well as the unconditional symmetry of the gaussian component. Let us illustrate this by proving the decomposition for $\Phi_{X,\alpha}$. First, note that $e^{-\alpha\|X\|_2^2} = e^{-\alpha\|\tilde{X}\|_2^2}e^{-\alpha\|g\|_2^2}$, so that $Z_{\Phi,X}(\alpha) = Z_{\Phi,\tilde{X}}(\alpha)Z_{\Phi,g}(\alpha)$. The top left d by d block is hence given by

$$\frac{\mathbb{E}e^{-\alpha\|X\|_2^2}\tilde{X}\tilde{X}^T}{Z_{\Phi,X}(\alpha)} = \frac{Z_{\Phi,g}(\alpha)\mathbb{E}e^{-\alpha\|\tilde{X}\|_2^2}\tilde{X}\tilde{X}^T}{Z_{\Phi,X}(\alpha)} = \frac{\mathbb{E}e^{-\alpha\|\tilde{X}\|_2^2}\tilde{X}\tilde{X}^T}{Z_{\Phi,\tilde{X}}(\alpha)} = \Phi_{\tilde{X},\alpha}.$$

The bottom right d' by d' block is also computed similarly. Finally, any entry outside these two blocks is of the form

$$\frac{\mathbb{E}e^{-\alpha\|X\|_2^2}\tilde{X}_i g_j}{Z_{\Phi,X}(\alpha)} = \frac{\mathbb{E}e^{-\alpha\|X\|_2^2}\tilde{X}_i(-g_j)}{Z_{\Phi,X}(\alpha)} = -\frac{\mathbb{E}e^{-\alpha\|X\|_2^2}\tilde{X}_i g_j}{Z_{\Phi,X}(\alpha)} = 0. \quad \square$$

Proof of Lemma 4.2. We have

$$\text{tr}(\Phi_{X,\alpha}) = \frac{\mathbb{E}\|X\|_2^2 e^{-\alpha\|X\|_2^2}}{\mathbb{E}e^{-\alpha\|X\|_2^2}} = \frac{-Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} = -(\log Z_{\Phi,X})'(\alpha).$$

The calculation for $\Psi_{X,\alpha}$ is similar. \square

In order to prove Lemma 4.3, we first need to establish the analyticity for the two partition functions.

Lemma B.1 (Analyticity for $Z_{\Phi,X}$ and $Z_{\Psi,X}$). *Let X be a subgaussian random vector in \mathbb{R}^n with subgaussian norm bounded by $K \geq 1$. The functions $Z_{\Phi,X}$ and $Z_{\Psi,X}$ are both analytic on $(-1/CK^2, 1/CK^2)$. They are given by the formulae $Z_{\Phi,X}(\alpha) = \sum_{r=0}^{\infty} \mathbb{E}\|X\|_2^{2r} (-\alpha)^r / r!$ and $Z_{\Psi,X}(\alpha) = \sum_{r=0}^{\infty} \mathbb{E}\langle X, X' \rangle^r (-\alpha)^r / r!$. Furthermore, by choosing C sufficiently large, on this interval they satisfy the bounds*

$$|Z_{\Phi,X}(\alpha)|, |Z_{\Psi,X}(\alpha)| \leq e^{CK^2 n |\alpha|} + \frac{CK^2 |\alpha|}{1 - CK^2 |\alpha|}. \quad (\text{B.1})$$

Proof. Let us first prove the bounds in (B.1). Observe that

$$\mathbb{E}e^{-\alpha\|X\|_2^2} \leq \mathbb{E}e^{|\alpha|\|X\|_2^2} = \sum_{n=0}^{\infty} \frac{\mathbb{E}\|X\|_2^{2n}}{n!} |\alpha|^n. \quad (\text{B.2})$$

Here, Tonelli allows us to interchange the sum and expectation. We next use Lemma D.1 to bound the terms of this series. Indeed, using the equivalent estimate (D.5), we have

$$\mathbb{E}\|X\|_2^{2r} \leq C^r K^{2r} (n^r + r!)$$

for some universal constant C . Substituting this into (B.2) and using $|\alpha| \leq 1/CK^2$, we have

$$\begin{aligned}\mathbb{E}e^{-\alpha\|X\|_2^2} &\leq \sum_{r=0}^{\infty} \frac{(CK^2)^r (n^r + r!)}{r!} |\alpha|^r \\ &= \sum_{r=0}^{\infty} \frac{(CK^2 n |\alpha|)^r}{r!} + \sum_{r=1}^{\infty} (CK^2 |\alpha|)^r \\ &= e^{CK^2 n |\alpha|} + \frac{CK^2 |\alpha|}{1 - CK^2 |\alpha|}.\end{aligned}$$

One may prove the bound for $Z_{\Psi,X}$ by doing the same computation but using (D.2) instead of (D.1).

We next handle analyticity of $Z_{\Phi,X}$. We shall prove by induction on r that we may differentiate under the integral sign to get the formula

$$Z_{\Phi,X}^{(r)}(\alpha) = (-1)^r \mathbb{E} \|X\|_2^{2r} e^{-\alpha\|X\|_2^2}. \quad (\text{B.3})$$

Assume the formula is true for all $r' < r$. Then

$$Z_{\Phi,X}^{(r)}(\alpha) = (-1)^{r-1} \lim_{h \rightarrow 0} \mathbb{E} \frac{\|X\|_2^{2r-2} e^{-(\alpha+h)\|X\|_2^2} - \|X\|_2^{2r-2} e^{-\alpha\|X\|_2^2}}{h} \quad (\text{B.4})$$

$$= (-1)^r \lim_{h \rightarrow 0} \mathbb{E} \|X\|_2^{2r-2} e^{-\alpha\|X\|_2^2} \frac{1 - e^{-h\|X\|_2^2}}{h} \quad (\text{B.5})$$

Next, note that the integrand is positive and by the mean value theorem, for a fixed value of $\|X\|_2^2$, we have

$$\frac{1 - e^{-h\|X\|_2^2}}{h} = \|X\|_2^2 e^{-h'\|X\|_2^2}$$

for some $h' \in [0, h]$ if $h > 0$ and $h' \in [h, 0]$ otherwise. As such, we have

$$\|X\|_2^{2r-2} e^{-\alpha\|X\|_2^2} \frac{1 - e^{-h\|X\|_2^2}}{h} \leq \|X\|_2^{2r} e^{(|h|-\alpha)\|X\|_2^2}$$

For $|h| - \alpha \leq 1/CK^2$, one can easily show that this is integrable by expanding this as a power series in $\|X\|_2^2$ and bounding the growth of the coefficients as above. As such, we may apply the Dominated Convergence Theorem to push the limit inside the expectation in (B.4), thereby yielding (B.3).

In particular, differentiating $Z_{\Phi,X}$ at 0, we see that its Taylor series at 0 is given by

$$Z_{\Phi,X}(\alpha) \sim \sum_{r=0}^{\infty} \frac{\mathbb{E} \|X\|_2^{2r}}{r!} (-\alpha)^r. \quad (\text{B.6})$$

The formula above shows that the Taylor series is absolutely convergent on our chosen interval. We next need to show that $Z_{\Phi,X}$ agrees with its Taylor series on this interval, meaning we have to show that the remainder term for the r -th Taylor polynomial goes to zero pointwise. The Lagrange form of the remainder term is written as

$$R_{Z_{\Phi,X},r}(\alpha) = \frac{Z_{\Phi,X}^{(r+1)}(\alpha')}{(r+1)!} \alpha^{r+1}$$

where $0 < |\alpha'| < |\alpha|$. Applying Cauchy-Schwarz to the formula (B.3), we get

$$|Z_{\Phi,X}^{(r+1)}(\alpha)| \leq \left(\mathbb{E} \|X\|_2^{4r+2} \right)^{1/2} \left(\mathbb{E} e^{-2\alpha \|X\|_2^2} \right)^{1/2}. \quad (\text{B.7})$$

Lemma D.1 again allows us to compute

$$\left(\mathbb{E} \|X\|_2^{4r+2} \right)^{1/2} \leq (CK^2)^{r+1} (n^{r+1} + (r+1)!).$$

This implies that for any $C' > 2C$,

$$\begin{aligned} \|R_{Z_{\Phi,X},r}\|_{L_\infty([-1/C'K^2, 1/C'K^2])} &\leq \frac{\|Z_{\Phi,X}^{(r+1)}\|_{L_\infty([-1/C'K^2, 1/C'K^2])}}{(C'K^2)^{r+1}(r+1)!} \\ &\leq \frac{(CK^2)^{r+1}(n^{r+1} + (r+1)!)}{(C'K^2)^{r+1}(r+1)!} \left(\mathbb{E} e^{2\|X\|_2^2/C'K^2} \right)^{1/2} \\ &\leq \left(\frac{C}{C'} \right)^{r+1} \left(\frac{n^{r+1}}{(r+1)!} + 1 \right) \left(e^{2nC/C'} + \frac{2C/C'}{1 - 2C/C'} \right). \end{aligned}$$

Using the fact that $r! \sim (\frac{r}{e})^r$, this last expression decays to zero as r tends to ∞ . Finally, to prove the claim for $Z_{\Psi,X}$, we repeat the same arguments. \square

Note that in the course of proving the last lemma, we have also proved the following result to be used elsewhere in the paper.

Lemma B.2 (Taylor remainder terms for $Z_{\Phi,X}$ and $Z_{\Psi,X}$). *Let X be a subgaussian random vector in \mathbb{R}^n with subgaussian norm bounded above by $K \geq 1$. There is an absolute constant C such that for all $0 < \alpha < 1/CK^2$, on the interval $[-\alpha, \alpha]$, the remainder terms for the r -th degree Taylor polynomials for $Z_{\Phi,X}$ and $Z_{\Psi,X}$ at 0 satisfy the uniform bound*

$$\|R_{Z_{\Phi,X},r}\|_\infty, \|R_{Z_{\Psi,X},r}\|_\infty \leq (CK^2)^{r+1} \alpha^{r+1} \left(\frac{n^{r+1}}{(r+1)!} + 1 \right) \left(e^{CK^2\alpha n} + \frac{CK^2\alpha}{1 - CK^2\alpha} \right) \quad (\text{B.8})$$

Proof of Lemma 4.3. By Lemma B.1, all four functions are analytic in a neighborhood of 0. Now recall that two different analytic functions cannot agree on a sequence with an accumulation point. \square

We now move on to proving Theorem 4.4. This requires the following technical lemma.

Lemma B.3. *Let X be subgaussian random vector in \mathbb{R}^n with subgaussian norm bounded above by $K \geq 1$. Suppose the moments of $\|X\|_2^2$ and $\|g\|_2^2$ agree up to order $r-1$, but there is a number $\Delta > 0$ such that $\left| \mathbb{E} \|X\|_2^{2r} - \mathbb{E} \|g\|_2^{2r} \right| \geq \Delta$, then there is an absolute constant C such that for $|\alpha| \leq \Delta r / (CK^2)^r (n^{r+1} + (r+1)!)$, we have*

$$|(\log Z_{\Phi,X})'(\alpha) - (\log Z_{\Phi,g})'(\alpha)| \geq \frac{\Delta}{2(r-1)!} |\alpha|^{r-1}. \quad (\text{B.9})$$

Similarly, suppose the moments of $\langle X, X' \rangle$ and $\langle g, g' \rangle$ agree up to order $r-1$ but $|\mathbb{E} \langle X, X' \rangle^r - \mathbb{E} \langle g, g' \rangle^r| \geq \Delta$, then for $|\alpha| \leq \Delta r / (CK^2)^r (n^{r+1} + (r+1)!)$, we have

$$|(\log Z_{\Psi,X})'(\alpha) - (\log Z_{\Psi,g})'(\alpha)| \geq \frac{\Delta}{2(r-1)!} |\alpha|^{r-1}. \quad (\text{B.10})$$

Proof. Let us first prove (B.9). For every positive integer k , let $p_{X,k}(\alpha) = \sum_{j=0}^k \mathbb{E}\|X\|_2^{2j} \alpha^j / j!$ denote the k -th Taylor polynomial of $Z_{\Phi,X}$, and define $p_{g,k}$ analogously. For convenience, also denote the k -th Taylor remainder term as $R_{X,k} := R_{Z_{\Phi,X},k}$. For any α , we then have

$$(\log Z_{\Phi,X})'(\alpha) - (\log Z_{\Phi,g})'(\alpha) = \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{Z'_{\Phi,g}(\alpha)}{Z_{\Phi,g}(\alpha)}, \quad (\text{B.11})$$

which we can then bound using

$$\left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{Z'_{\Phi,g}(\alpha)}{Z_{\Phi,g}(\alpha)} \right| \geq \left| \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} - \frac{p'_{g,r}(\alpha)}{p_{g,r-1}(\alpha)} \right| - \left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} \right| - \left| \frac{p'_{g,r}(\alpha)}{p_{g,r-1}(\alpha)} - \frac{Z'_{\Phi,g}(\alpha)}{Z_{\Phi,g}(\alpha)} \right|. \quad (\text{B.12})$$

We now bound each of these three terms individually. First, we need upper and lower bounds for $p_{X,k}(\alpha)$. Using the $\|X\|_2^2$ moment bound (D.5), we have

$$\begin{aligned} |p_{X,k}(\alpha) - 1| &\leq \sum_{j=1}^k \frac{\mathbb{E}\|X\|_2^{2j} |\alpha|^j}{j!} \\ &\leq \sum_{j=1}^k \frac{(CK^2)^j (n^j + j!) |\alpha|^j}{j!} \\ &= \sum_{j=1}^k \frac{(CK^2 n |\alpha|)^j}{j!} + \sum_{j=1}^k (CK^2 |\alpha|)^j \\ &\leq e^{CK^2 n |\alpha|} + \frac{CK^2 |\alpha|}{1 - CK^2 |\alpha|}. \end{aligned}$$

By sharpening the constant C in our assumption on $|\alpha|$ if necessary, we may thus ensure that

$$|p_{X,k}(\alpha) - 1| \leq \frac{1}{2} \quad (\text{B.13})$$

By the same argument, we can also ensure that

$$\left| p'_{X,k}(\alpha) - \mathbb{E}\|X\|_2^2 \right| \leq \frac{1}{2}. \quad (\text{B.14})$$

By our assumptions on the moments of $\|X\|_2^2$ and $\|g\|_2^2$, we have $p_{X,r-1} \equiv p_{g,r-1}$. Furthermore, only the leading terms of $p'_{X,r}$ and $p'_{g,r}$ differ. This, together with (B.13) implies that

$$\begin{aligned} \left| \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} - \frac{p'_{g,r}(\alpha)}{p_{g,r-1}(\alpha)} \right| &\geq \frac{2}{3} |p'_{X,r}(\alpha) - p'_{g,r}(\alpha)| \\ &\geq \frac{2\Delta |\alpha|^{r-1}}{3(r-1)!}. \end{aligned} \quad (\text{B.15})$$

Next, we have

$$\left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} \right| \leq \left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| + \left| \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} \right|. \quad (\text{B.16})$$

Again we bound these two terms individually. Using the identity $p_{X,r}(\alpha) = p_{X,r-1}(\alpha) + \mathbb{E}\|X\|_2^{2r}(-\alpha)^r/r!$, we get

$$\begin{aligned} \left| \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} \right| &= \left| \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| \left| 1 - \frac{p_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} \right| \\ &= \left| \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)p_{X,r-1}(\alpha)} \right| \frac{\mathbb{E}\|X\|_2^{2r}|\alpha|^r}{r!} \end{aligned} \quad (\text{B.17})$$

Using the bounds on $p_{X,r}$ and $p'_{X,r}$ (B.13) and (B.14), together with the $\|X\|_2^2$ moment bound (D.5), we get

$$\left| \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)p_{X,r-1}(\alpha)} \right| \frac{\mathbb{E}\|X\|_2^{2r}|\alpha|^r}{r!} \leq \frac{3}{8} \frac{(CK^2)^r(n^r + r!)|\alpha|^r}{r!}. \quad (\text{B.18})$$

For the first term in (B.16), we write

$$\begin{aligned} \left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| &= |(\log Z_{\Phi,X}(\alpha))' - (\log p_{X,r}(\alpha))'| \\ &= \left| \frac{d}{d\alpha} \log \left(\frac{Z_{\Phi,X}(\alpha)}{p_{X,r}(\alpha)} \right) \right| \\ &= \left| \frac{d}{d\alpha} \log \left(1 + \frac{R_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right) \right|. \end{aligned} \quad (\text{B.19})$$

Using Lemma B.2 together with our assumptions on $|\alpha|$, we observe that

$$\begin{aligned} |R_{X,r}(\alpha)| &\leq (CK^2)^{r+1}|\alpha|^{r+1} \left(\frac{n^{r+1}}{(r+1)!} + 1 \right) \left(e^{CK^2|\alpha|n} + \frac{CK^2|\alpha|}{1 - CK^2|\alpha|} \right) \\ &\leq (CK^2)^{r+1}|\alpha|^{r+1} \left(\frac{n^{r+1}}{(r+1)!} + 1 \right). \end{aligned} \quad (\text{B.20})$$

In particular, by sharpening the constant C in our assumption on $|\alpha|$ if necessary, we can ensure that this quantity is less than $\frac{1}{4}$. In this case, we have

$$\left| \frac{R_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| \leq \frac{1}{2},$$

so that

$$\begin{aligned} \left| \frac{d}{d\alpha} \log \left(1 + \frac{R_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right) \right| &= \left| \log' \left(1 + \frac{R_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right) \right| \left| \left(\frac{R_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right)' \right| \\ &\leq 2 \left| \left(\frac{R_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right)' \right| \\ &\leq 2 \left(\left| \frac{R'_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| + \left| \frac{R_{X,r}(\alpha)p'_{X,r}(\alpha)}{p_{X,r}(\alpha)^2} \right| \right). \end{aligned} \quad (\text{B.21})$$

By our bounds on these functions (B.13), (B.14), and (B.20), we have

$$\left| \frac{R_{X,r}(\alpha) p'_{X,r}(\alpha)}{p_{X,r}(\alpha)^2} \right| \leq (CK^2)^{r+1} |\alpha|^{r+1} \left(\frac{n^{r+1}}{(r+1)!} + 1 \right). \quad (\text{B.22})$$

Furthermore, by using the moment bounds (D.5) as before, one can show that

$$|R'_{X,r}(\alpha)| \leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right).$$

so that the first term is also bounded according to

$$\left| \frac{R'_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| \leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right). \quad (\text{B.23})$$

As such, combining (B.19) and (B.21) tells us that

$$\begin{aligned} \left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r}(\alpha)} \right| &\leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right) + (CK^2)^{r+1} |\alpha|^{r+1} \left(\frac{n^{r+1}}{(r+1)!} + 1 \right) \\ &\leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right). \end{aligned} \quad (\text{B.24})$$

We can now use this estimate together with (B.18) to continue (B.16), writing

$$\begin{aligned} \left| \frac{Z'_{\Phi,X}(\alpha)}{Z_{\Phi,X}(\alpha)} - \frac{p'_{X,r}(\alpha)}{p_{X,r-1}(\alpha)} \right| &\leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right) + \frac{(CK^2)^r (n^r + r!) |\alpha|^r}{r!} \\ &\leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right). \end{aligned} \quad (\text{B.25})$$

Notice that same methods also give us

$$\left| \frac{p'_{g,r}(\alpha)}{p_{g,r-1}(\alpha)} - \frac{Z'_{\Phi,g}(\alpha)}{Z_{\Phi,g}(\alpha)} \right| \leq (CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right). \quad (\text{B.26})$$

We may therefore finally substitute these last two bounds, together with (B.15), into (B.12). This yields

$$|(\log Z_{\Phi,X})'(\alpha) - (\log Z_{\Phi,g})'(\alpha)| \geq \frac{2\Delta |\alpha|^{r-1}}{3(r-1)!} - C(CK^2)^r |\alpha|^r \left(\frac{n^{r+1}}{r!} + r + 1 \right). \quad (\text{B.27})$$

We now claim that with our assumptions on $|\alpha|$, the first term dominates the second. This is a simple calculation, thereby completing the proof of (B.9). To prove (B.10), we repeat the entire argument, but using the relevant estimates for $Z_{\Psi,X}$ instead of those for $Z_{\Phi,X}$. \square

Applying the previous lemma in the setting of our NGCA model, we get the following result.

Theorem B.4 (Robustness for non-gaussian eigenvalues). *Let X be a subgaussian random vector satisfying the NGCA model (1.1), and with subgaussian norm bounded above by $K \geq 1$. Let $\lambda_1(\Phi_{\tilde{X},\alpha}), \dots, \lambda_d(\Phi_{\tilde{X},\alpha})$ denote the eigenvalues of $\Phi_{\tilde{X},\alpha}$. Suppose the moments of $\|\tilde{X}\|_2^2$ and $\|g_d\|_2^2$ agree up to order $r-1$, but there is a number $\Delta > 0$ such that $|\mathbb{E}\|\tilde{X}\|_2^{2r} - \mathbb{E}\|g_d\|_2^{2r}| \geq \Delta$, then there is an absolute constant C such that for $|\alpha| \leq \Delta r / (CK^2)^r (d^{r+1} + (r+1)!)$, we have*

$$\left| \frac{1}{d} \sum_{i=1}^d \lambda_i(\Phi_{\tilde{X},\alpha}) - \frac{1}{2\alpha+1} \right| \geq \frac{\Delta}{2d(r-1)!} |\alpha|^{r-1}. \quad (\text{B.28})$$

Similarly, let $\lambda_1(\Psi_{\tilde{X},\alpha}), \dots, \lambda_d(\Psi_{\tilde{X},\alpha})$ denote the eigenvalues of $\Psi_{\tilde{X},\alpha}$, and suppose the moments of $\langle X, X' \rangle$ and $\langle g, g' \rangle$ agree up to order $r-1$ but $|\mathbb{E}\langle X, X' \rangle^r - \mathbb{E}\langle g, g' \rangle^r| \geq \Delta$. Then for $|\alpha| \leq \Delta r / (CK^2)^r (d^{r+1} + (r+1)!)$, we have

$$\left| \frac{1}{d} \sum_{i=1}^d \lambda_i(\Psi_{\tilde{X},\alpha}) - \frac{\alpha}{\alpha^2-1} \right| \geq \frac{\Delta}{2d(r-1)!} |\alpha|^{r-1}. \quad (\text{B.29})$$

Proof. This is simply a translation of the previous theorem with the help of Lemma 4.2, which tells us that the log derivatives of the partition functions are equal to the traces of $\Phi_{X,\alpha}$ and $\Psi_{X,\alpha}$, and that of Lemma C.3, which tells us what the gaussian eigenvalue is. \square

Proof of Theorem 4.4. Combine the previous Corollary with Theorem 2.3. \square

C Identities for Φ_g and Ψ_g

In this section, we let g_1 denote a standard gaussian random variable, and g_n , a standard gaussian random vector in \mathbb{R}^n . First, notice that independence gives $Z_{\Phi,g_n}(\alpha) = Z_{\Phi,g_1}(\alpha)^n$ and $Z_{\Psi,g_n}(\alpha) = Z_{\Psi,g_1}(\alpha)^n$.

Lemma C.1. *We have the identities $Z_{\Phi,g_n}(\alpha) = (2\alpha+1)^{-n/2}$ when $\alpha > -1/2$ and $Z_{\Psi,g_n}(\alpha) = (1-\alpha^2)^{-n/2}$ when $|\alpha| < 1$.*

Proof. By the remarks above, it suffices to prove the formula when $n = 1$. These are then simple exercises in calculus. Notice that

$$Z_{\Phi,g_1}(\alpha) = \mathbb{E}e^{-\alpha g_1^2} = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\alpha t^2} e^{-\frac{t^2}{2}} dt = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\frac{(2\alpha+1)t^2}{2}} dt.$$

Now substitute $u = \sqrt{2\alpha+1} \cdot t$ to arrive at the formula for Z_{Φ,g_1} . For the next formula, we use conditional expectations to write

$$Z_{\Psi,g_1}(\alpha) = \mathbb{E}e^{-\alpha g_1 g_1'} = \mathbb{E}[\mathbb{E}[e^{-\alpha g_1 g_1'} | g_1]]. \quad (\text{C.1})$$

The inner expectation can be computed as

$$\mathbb{E}[e^{-\alpha g_1 g_1'} | g_1] = \frac{1}{\sqrt{2\pi}} \int_{-\infty}^{\infty} e^{-\alpha g_1 t} e^{-\frac{t^2}{2}} dt = e^{\frac{(\alpha g_1)^2}{2}}.$$

Substituting this back into (C.1) and using the same technique as above gives us what we want. \square

Lemma C.2. We have the identities $-(\log Z_{\Phi, g_n})'(\alpha) = n(2\alpha + 1)^{-1}$ when $\alpha > -1/2$ and $-(\log Z_{\Psi, g_n})'(\alpha) = n\alpha(\alpha^2 - 1)^{-1}$ when $|\alpha| < 1$.

Lemma C.3. We have the identities $\Phi_{g_n, \alpha} = (2\alpha + 1)^{-1}I_n$ when $\alpha > -1/2$ and $\Psi_{g_n, \alpha} = \alpha(\alpha^2 - 1)^{-1}I_n$ when $|\alpha| < 1$. Here, I_n is the n -dimensional identity matrix.

Proof. By rotational symmetry, we know that both matrices are multiples of the identity. To compute these scalars, it hence suffices to find the trace of both matrices. But

$$\text{tr}(\Phi_{g_n, \alpha}) = \frac{\mathbb{E}e^{-\alpha\|g_n\|_2^2}\|g_n\|_2^2}{\mathbb{E}e^{-\alpha\|g_n\|_2^2}} = -(\log Z_{\Phi, g_n})'(\alpha).$$

Dividing by n and using the previous lemma gives us what we want. \square

D Concentration and moment bounds

Theorem D.1 (Concentration of norm for general subgaussian vectors). *Let X be a subgaussian random vector in \mathbb{R}^n , with $\|X\|_{\psi_2} \leq K$. There is a universal constant C such that for each positive integer $r > 0$, the moments of $\|X\|_2$ and $\langle X, X' \rangle$ satisfy*

$$(\mathbb{E}\|X\|_2^r)^{1/r} \leq CK(\sqrt{n} + \sqrt{r}) \quad (\text{D.1})$$

$$(\mathbb{E}|\langle X, X' \rangle|^r)^{1/2r} \leq CK(\sqrt{n} + \sqrt{r}). \quad (\text{D.2})$$

Proof. The second bound follows from the first, since by Cauchy-Schwarz,

$$(\mathbb{E}|\langle X, X' \rangle|^r)^{1/2r} \leq (\mathbb{E}\|X\|_2^r \mathbb{E}\|X'\|_2^r)^{1/2r} = (\mathbb{E}\|X\|_2^r)^{1/r}$$

To prove (D.1), pick a $\frac{1}{2}$ -net \mathcal{N} on S^{n-1} . A volumetric argument shows that one may pick \mathcal{N} to have size no more than 5^n [?]. We then have

$$\|X\| = \sup_{v \in S^{n-1}} \langle X, v \rangle \leq 2 \sup_{v \in \mathcal{N}} \langle X, v \rangle.$$

By definition, there is a universal constant c such that for any fixed unit vector $v \in S^{d-1}$, $\mathbb{P}(\langle X, v \rangle > t) \leq 2 \exp\left(-\frac{ct^2}{K^2}\right)$. Taking a union bound over the net thus gives

$$\mathbb{P}(\|X\| > 2t) \leq 2 \exp\left(n \log 5 - \frac{ct^2}{K^2}\right). \quad (\text{D.3})$$

Next, we integrate out the tail bound (D.3) to obtain bounds for the moments. Observe that if $\frac{ct^2}{2K^2} \geq n \log 5$, we have $n \log 5 - \frac{ct^2}{K^2} \leq -\frac{ct^2}{2K^2}$. This condition on t is equivalent to $t \geq CK\sqrt{n}$, so we have

$$\mathbb{P}(\|X\| > 2t) \leq \begin{cases} 1 & t < CK\sqrt{n} \\ 2 \exp\left(-\frac{ct^2}{K^2}\right) & t \geq CK\sqrt{n} \end{cases} \quad (\text{D.4})$$

For any positive integer r , we integrate this bound to get

$$\begin{aligned}\mathbb{E}\|X\|_2^r &= \int_0^\infty r t^{r-1} \mathbb{P}(\|X\| > t) dt \\ &\leq \int_0^{CK\sqrt{n}} r t^{r-1} dt + \int_{CK\sqrt{n}}^\infty 2r t^{r-1} \exp\left(-\frac{ct^2}{K^2}\right) dt \\ &\leq C^r K^r n^{r/2} + C^r K^r r \int_0^\infty t^{r/2-1} e^{-t} dt.\end{aligned}$$

The integral in the last line is the gamma function, so in short, we have shown that

$$\mathbb{E}\|X\|_2^r \leq C^r K^r (n^{r/2} + \Gamma(r/2 + 1)). \quad (\text{D.5})$$

Taking r -th roots of both sides and using Hölder, together with the fact that $\Gamma(x)^{1/x} \lesssim x$, gives (D.1). \square

Lemma D.2 (Covariance estimation for subgaussian random vectors). *Let X be a centered subgaussian random vector in \mathbb{R}^n with covariance matrix Σ and subgaussian norm satisfying $\|X\|_{\psi_2} \leq K$ for some $K \geq 1$. Let $\hat{\Sigma}_N = \frac{1}{N} \sum_{i=1}^N X_i X_i^T$ denote the sample covariance matrix from N independent samples. Then there is an absolute constant C such that for any $0 < \epsilon, \delta < 1$, we have $\mathbb{P}\left(\left\|\hat{\Sigma}_N - \Sigma\right\| > \epsilon\right) \leq \delta$ so long as $N \geq CK^2(n + \log(1/\delta))\epsilon^{-2}$.*

Proof. Refer to [21]. \square

Lemma D.3 (Moments of spherical marginals). *Let θ be uniformly distributed on the sphere S^{n-1} . Then for any unit vector $v \in S^{n-1}$ and any positive integer k , we have*

$$\mathbb{E}\langle \theta, v \rangle^{2k} = \frac{1 \cdot 3 \cdots (2k-1)}{n \cdot (n+2) \cdots (n+2k-2)} \quad (\text{D.6})$$

Proof. There are several ways to prove this identity. We shall prove this by computing gaussian integrals. Let g_1 and g_n denote standard gaussians in 1 dimension and n dimensions respectively. Then using the radial symmetry of g , we have

$$\mathbb{E}g_1^{2k} = \mathbb{E}\langle g_n, v \rangle^{2k} = \mathbb{E}\langle \|g_n\|_2 \theta, v \rangle^{2k} = \mathbb{E}\|g_n\|_2^{2k} \mathbb{E}\langle \theta, v \rangle^{2k}.$$

Rearranging gives

$$\mathbb{E}\langle \theta, v \rangle^{2k} = \frac{\mathbb{E}g_1^{2k}}{\mathbb{E}\|g_n\|_2^{2k}}.$$

We then compute

$$\mathbb{E}\|g_n\|_2^{2k} = \frac{\omega_n}{(2\pi n)^{n/2}} \int_0^\infty r^{2k} r^{n-1} e^{-r^2/2} dr, \quad (\text{D.7})$$

where ω_n is the volume of the sphere S^{n-1} . It is well known that

$$\omega_n = \frac{2\pi^{n/2}}{\Gamma(n/2)},$$

while we also have

$$\int_0^\infty r^{2k} r^{n-1} e^{-r^2/2} dr = 2^{n/2+k-1} \Gamma(n/2 + k).$$

Substituting these back into (D.7) gives

$$\mathbb{E} \|g_n\|_2^{2k} = 2^k \frac{\Gamma(n/2 + k)}{\Gamma(n/2)} = n \cdot (n+2) \cdots (n+2k-2). \quad (\text{D.8})$$

This yields the denominator in (D.6). A similar calculation for $\mathbb{E} g_1^{2k}$ yields the numerator. \square

Proof of Theorem 5.1. Let $Y = e^{-\alpha \|X\|_2^2} X$. Then Y is a subgaussian random vector with $\|Y\|_{\psi_2} \leq K$. Let Σ and $\hat{\Sigma}$ denote its covariance and empirical covariance matrices respectively. Then $\|\Sigma\| \leq 1$ and by Lemma D.2, we have $\|\hat{\Sigma} - \Sigma\| \leq \epsilon/2$ with probability at least $1 - \delta/2$. Next, observe that $\Phi_{X,\alpha} = Z_{\Phi,X}(\alpha)^{-1} \Sigma$ and $\hat{\Phi}_{X,\alpha} = \hat{Z}_{\Phi,X}(\alpha)^{-1} \hat{\Sigma}$, where $\hat{Z}_{\Phi,X}(\alpha) = \sum_{j=1}^N e^{-\alpha \|X_j\|_2^2} / N$. As such, we have

$$\|\hat{\Phi}_{X,\alpha} - \Phi_{X,\alpha}\| \leq |\hat{Z}_{\Phi,X}(\alpha)^{-1}| \|\hat{\Sigma} - \Sigma\| + |\hat{Z}_{\Phi,X}(\alpha)^{-1} - Z_{\Phi,X}(\alpha)^{-1}| \|\Sigma\|. \quad (\text{D.9})$$

Combining our lower bound on α with the power series formula for Z_Φ from Lemma B.1, we have $Z_{\Phi,X}(\alpha) \geq 1/2$. Furthermore, we may apply Hoeffding's inequality to see that $|\hat{Z}_{\Phi,X}(\alpha) - Z_{\Phi,X}(\alpha)| \leq \epsilon/2$ with probability at least $1 - \delta/2$. We can now combine all of this together to get the probability bound. \square

Proof of Theorem 5.2. First, define $\Sigma = \mathbb{E}[e^{-\alpha \langle X, X' \rangle} X(X')^T]$ and $\hat{\Sigma} = \sum_{i=1}^N e^{-\alpha \langle X_i, X'_i \rangle} (X_i(X'_i)^T + X'_i X_i^T) / 2N$, so that $\Psi_{X,\alpha} = Z_{\Psi,X}(\alpha)^{-1} \Sigma$ and $\hat{\Psi}_{X,\alpha} = \hat{Z}_{\Psi,X}(\alpha)^{-1} \hat{\Sigma}$. As in the previous theorem, we can write

$$\|\hat{\Psi}_{X,\alpha} - \Psi_{X,\alpha}\| \leq |\hat{Z}_{\Psi,X}(\alpha)^{-1}| \|\hat{\Sigma} - \Sigma\| + |\hat{Z}_{\Psi,X}(\alpha)^{-1} - Z_{\Psi,X}(\alpha)^{-1}| \|\Sigma\|. \quad (\text{D.10})$$

This time however, we cannot immediately invoke Lemma D.2 because we can no longer view Σ and $\hat{\Sigma}$ as the covariance and empirical covariance matrices of a random vector. Nonetheless, we can follow the same proof scheme with a few adjustments.

The basic idea is to use a net argument to transform the operator deviation bound into a scalar bound for random variables. Let \mathcal{N} be a $\frac{1}{4}$ -net on S^{n-1} . By a volumetric argument, we may pick \mathcal{N} to have size no more than 9^n [21]. For any n by n real symmetric matrix M , we then have

$$\|M\| = \sup_{v \in S^{n-1}} |\langle v, Mv \rangle| \leq 2 \sup_{v \in \mathcal{N}} |\langle v, Mv \rangle|. \quad (\text{D.11})$$

As such, by taking a union bound, we can hope to bound $\|\hat{\Sigma} - \Sigma\|$ by bounding $|\langle v, (\hat{\Sigma} - \Sigma)v \rangle|$ for a fixed unit vector $v \in S^{n-1}$. Let us do just this. We have

$$\langle v, \hat{\Sigma} v \rangle = \frac{1}{N} \sum_{i=1}^N e^{-\alpha \langle X_i, X'_i \rangle} \langle X_i, v \rangle \langle X'_i, v \rangle,$$

so that

$$\langle v, (\hat{\Sigma} - \Sigma)v \rangle = \frac{1}{N} \sum_{i=1}^N (Y_i - \mathbb{E} Y_i), \quad (\text{D.12})$$

where

$$Y_i = e^{-\alpha \langle X_i, X'_i \rangle} \langle X_i, v \rangle \langle X'_i, v \rangle. \quad (\text{D.13})$$

Observe that the Y_i 's are i.i.d. random variables. At this point in the proof of covariance estimation, one observes that the resulting random variables are subexponential, so one may apply Bernstein's inequality [21]. Unfortunately, our Y_i 's are not subexponential because of the $e^{-\alpha \langle X_i, X'_i \rangle}$ factor. The way we overcome this is to condition on the size of these factors being uniformly small. Indeed, by Lemma D.4 to come, we have $e^{-\alpha \langle X_i, X'_i \rangle} \leq e$ for all samples i with probability at least $1 - \delta$. We call this event A .

Next, define $\tilde{Y}_i := Y_i 1_A$. The Y_i 's are i.i.d random variables with subexponential norm bounded by eK^2 . We can then apply Bernstein and our assumption on the sample size N to get

$$\mathbb{P} \left(\left| \frac{1}{N} \sum_{i=1}^N (\tilde{Y}_i - \mathbb{E} \tilde{Y}_i) \right| > \epsilon \right) \leq e^{-N\epsilon^2 / CK^4} \leq \frac{\delta}{9^n}. \quad (\text{D.14})$$

Conditioning on the set A , we have $Y_i = \tilde{Y}_i$ for each i . We can also rewrite the bound on the right hand side using our assumption on N . Doing this gives us

$$\mathbb{P} \left(\left| \frac{1}{N} \sum_{i=1}^N (Y_i - \mathbb{E} \tilde{Y}_i) \right| > \epsilon \mid A \right) \leq \frac{\delta}{9^n}. \quad (\text{D.15})$$

We would like to replace $\mathbb{E} \tilde{Y}_i$ with $\mathbb{E} Y_i$, but the two quantities are not necessarily equal. Nonetheless, we can bound their difference as follows. We have

$$\mathbb{E} Y_i - \mathbb{E} \tilde{Y}_i = \mathbb{E} Y_i 1_{A^c} = \mathbb{E} [e^{-\alpha \langle X_i, X'_i \rangle} \langle X_i, v \rangle \langle X'_i, v \rangle 1_{A^c}]. \quad (\text{D.16})$$

We apply generalized Hölder to write

$$|\mathbb{E} [e^{-\alpha \langle X_i, X'_i \rangle} \langle X_i, v \rangle \langle X'_i, v \rangle 1_{A^c}]| \leq \left(\mathbb{E} e^{-4\alpha \langle X_i, X'_i \rangle} \right)^{1/4} \left(\mathbb{E} [\langle X_i, v \rangle^4 \langle X'_i, v \rangle^4] \right)^{1/4} \mathbb{P}(A^c)^{1/2}. \quad (\text{D.17})$$

We now use the moment bounds for subgaussian random variables and Lemma D.5 to bound the first two multiplicands on the right. This gives us

$$|\mathbb{E} [e^{-\alpha \langle X_i, X'_i \rangle} \langle X_i, v \rangle \langle X'_i, v \rangle 1_{A^c}]| \leq CK^2 \mathbb{P}(A^c)^{1/2}. \quad (\text{D.18})$$

Next, we use Lemma D.4 together with our assumption on $|\alpha|$, tightening the constant if necessary, to see that $\mathbb{P}(A^c) \leq \epsilon^2 / C^2 K^4$. We combine this together with the last few equations to obtain $|\mathbb{E} Y_i - \mathbb{E} \tilde{Y}_i| \leq \epsilon$, and combining this with (D.15), we obtain

$$\mathbb{P} \left(\left| \frac{1}{N} \sum_{i=1}^N (Y_i - \mathbb{E} Y_i) \right| > 2\epsilon \mid A \right) \leq \frac{\delta}{9^n}. \quad (\text{D.19})$$

Recall that Y_i 's were defined for a fixed $v \in \mathcal{N}$. We can take a union bound over all vectors in \mathcal{N} to get

$$\mathbb{P} \left(\sup_{v \in \mathcal{N}} |\langle v, (\hat{\Sigma} - \Sigma)v \rangle| > 2\epsilon \mid A \right) \leq \delta. \quad (\text{D.20})$$

Combining this with (D.11) then gives

$$\mathbb{P}\left(\|\hat{\Sigma} - \Sigma\| > 4\epsilon \mid A\right) \leq \delta. \quad (\text{D.21})$$

Let us continue to bound the other terms in (D.10) conditioned on the set A . Notice that on this set, $\hat{Z}_{\Psi,X}(\alpha)$ is an average of terms that are each bounded in absolute value by e . Using Hoeffding's inequality together with a similar argument as above to bound $|\mathbb{E}\hat{Z}_{\Psi,X}(\alpha)1_A - Z_{\Psi,X}(\alpha)|$, one may show that

$$\mathbb{P}\left(|\hat{Z}_{\Psi,X}(\alpha) - Z_{\Psi,X}(\alpha)| > \epsilon/2 \mid A\right) \leq \delta. \quad (\text{D.22})$$

We may also use the power series formula for $Z_{\Psi,X}$ from Lemma B.1 together with our bound on $|\alpha|$ to show that $Z_{\Psi,X}(\alpha) \geq \frac{1}{2}$.

It remains to bound $\|\Sigma\|$. To do this, we let v again be an arbitrary unit vector, and use Cauchy-Schwarz to compute

$$|\mathbb{E}[e^{-\alpha\langle X_i, X'_i \rangle} \langle X_i, v \rangle \langle X'_i, v \rangle]| \leq \left(\mathbb{E}e^{-2\alpha\langle X_i, X'_i \rangle}\right)^{1/2} \left(\mathbb{E}[\langle X_i, v \rangle^2 \langle X'_i, v \rangle^2]\right)^{1/2}. \quad (\text{D.23})$$

We have already seen that moment bounds and Lemma D.5 imply that this is bounded by an absolute constant C . In fact, we can take $C = 3$.

Putting everything together, we see that on the set A , we can continue writing (D.10) as

$$\begin{aligned} \|\hat{\Psi}_{X,\alpha} - \Psi_{X,\alpha}\| &\leq |\hat{Z}_{\Psi,X}(\alpha)^{-1}| \|\hat{\Sigma} - \Sigma\| + |\hat{Z}_{\Psi,X}(\alpha)^{-1} - Z_{\Psi,X}(\alpha)^{-1}| \|\Sigma\| \\ &\leq C\epsilon. \end{aligned}$$

Using our bound for $\mathbb{P}(A)$, we can therefore uncondition to get

$$\mathbb{P}\left(\|\hat{\Psi}_{X,\alpha} - \Psi_{X,\alpha}\| > C\epsilon\right) \leq \delta + \mathbb{P}(A) \leq 2\delta. \quad (\text{D.24})$$

Finally, note that we can massage the constants so that the multiplying constants in front of ϵ and δ disappear. \square

Lemma D.4. *For any $0 < \delta < 1$ and $N \in \mathbb{N}$, if $|\alpha| \leq \left(CK^2\sqrt{\log(N/\delta)}(\sqrt{n} + \sqrt{\log(N/\delta)})\right)^{-1}$, then*

$$\mathbb{P}\left(\sup_{1 \leq i \leq N} e^{-\alpha\langle X_i, X'_i \rangle} > e\right) \leq \delta. \quad (\text{D.25})$$

Proof. Without loss of generality, assume that $\alpha > 0$. Using the union bound, it suffices to prove that

$$\mathbb{P}(\langle X, X' \rangle < -1/\alpha) = \mathbb{P}\left(e^{-\alpha\langle X, X' \rangle} > e\right) \leq \frac{\delta}{N}. \quad (\text{D.26})$$

To compute this, we first condition on X' and use the subgaussian tail of X to get

$$\mathbb{P}(\langle X, X' \rangle < -1/\alpha \mid X') \leq \exp\left(-\frac{1}{CK^2\alpha^2\|X'\|_2^2}\right),$$

and integrating out X' , then gives

$$\mathbb{P}(\langle X, X' \rangle < -1/\alpha) \leq \mathbb{E}e^{-(CK^2\alpha^2\|X'\|_2^2)^{-1}}. \quad (\text{D.27})$$

To compute this expectation, let A be the event that $\|X'\|_2 \leq CK(\sqrt{n} + \sqrt{\log(N/\delta)})$. Then by equation (D.4) in Theorem D.1, we have $\mathbb{P}(A^c) \leq \delta/N$. As such, we can break up the expectation into the portion over A and the the portion over A^c to obtain

$$\begin{aligned} \mathbb{E}e^{-(CK^2\alpha^2\|X'\|_2^2)^{-1}} &= \mathbb{E}[e^{-(CK^2\alpha^2\|X'\|_2^2)^{-1}} \mid A]\mathbb{P}(A) + \mathbb{E}[e^{-(CK^2\alpha^2\|X'\|_2^2)^{-1}} \mid A^c]\mathbb{P}(A^c) \\ &\leq \mathbb{E}[e^{-(CK^2\alpha^2\|X'\|_2^2)^{-1}} \mid A] + \mathbb{P}(A^c) \\ &\leq \exp\left(-\frac{1}{CK^4\alpha^2(n + \log(N/\delta))}\right) + \frac{\delta}{N}. \end{aligned} \quad (\text{D.28})$$

As such, we just need the first term to be less than δ/N , which corresponds to the requirement that

$$\frac{1}{CK^4\alpha^2(n + \log(N/\delta))} \geq \log(N/\delta).$$

This is simply a rearrangement of our assumption on $|\alpha|$. \square

Lemma D.5 (Better bound for Z_Ψ). *There is an absolute constant C such that if $|\alpha| \leq 1/CK^2\sqrt{n}$, then $Z_{\Psi,X}(\alpha) \leq 3$.*

Proof. The idea of the proof is similar to that of the previous lemma. We first condition on X' and use the subgaussian nature of X to bound its Laplace transform, thereby obtaining

$$\mathbb{E}[e^{-\alpha\langle X, X' \rangle} \mid X'] \leq e^{CK^2\alpha^2\|X'\|_2^2}.$$

Integrating out X' gives

$$\begin{aligned} Z_{\Psi,X}(\alpha) &\leq \mathbb{E}e^{CK^2\alpha^2\|X'\|_2^2} \\ &= \int_0^\infty \mathbb{P}\left(e^{CK^2\alpha^2\|X'\|_2^2} > t\right) dt \\ &\leq e + \int_e^\infty \mathbb{P}\left(e^{CK^2\alpha^2\|X'\|_2^2} > t\right) dt \end{aligned} \quad (\text{D.29})$$

Next, we use our assumption on $|\alpha|$ to write

$$\begin{aligned} \mathbb{P}\left(e^{CK^2\alpha^2\|X'\|_2^2} > t\right) &= \mathbb{P}\left(\|X'\|_2 > \frac{\sqrt{\log t}}{CK|\alpha|}\right) \\ &\leq \mathbb{P}\left(\|X'\|_2 > \sqrt{\log t}CK\sqrt{n}\right). \end{aligned} \quad (\text{D.30})$$

For $t > e$, we have $\sqrt{\log t} > 1$, so we may apply (D.4) to get

$$\mathbb{P}\left(\|X'\|_2 > \sqrt{\log t}CK\sqrt{n}\right) \leq e^{-\log t Cn} = t^{-Cn}. \quad (\text{D.31})$$

Plugging this into (D.29) gives

$$Z_{\Psi,X}(\alpha) \leq e + \frac{e^{-Cn}}{Cn} \leq 3 \quad (\text{D.32})$$

if we choose C to be large enough. \square

E Eigenvector perturbation theory

If two n by n matrices are close in spectral norm, one can use minimax identities to show that their eigenvalues are also close. It is less trivial to show that their eigenvectors are also close, which is the case in the presence of an “eigengap”. This is addressed by the well-known Davis-Kahan theorems [7].

Definition E.1. Let E and \hat{E} be two subspaces of \mathbb{R}^n of dimension d . Let V and \hat{V} be n by d matrices with orthonormal columns forming a basis for E and \hat{E} respectively. Let $\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_d$ be the singular values of $V^T \hat{V}$. We define the *principal angles* of E and \hat{E} to be $\theta_i(E, \hat{E}) = \arccos \sigma_i$ for $1 \leq i \leq d$.

Lemma E.2. Let E , \hat{E} , U and \hat{U} be as in the previous definition. We have

$$\|\hat{V}\hat{V}^T - VV^T\|_F^2 = 2 \sum_{i=1}^d \sin^2 \theta_i(E, \hat{E}). \quad (\text{E.1})$$

In particular, the quantity depends only on E and \hat{E} and not the choice of bases.

Proof. We expand

$$\|\hat{V}\hat{V}^T - VV^T\|_F^2 = \|\hat{V}\hat{V}^T\|_F^2 + \|VV^T\|_F^2 - 2\langle VV^T, \hat{V}\hat{V}^T \rangle. \quad (\text{E.2})$$

Observe that

$$\|VV^T\|_F^2 = \text{tr}(VV^T VV^T) = \text{tr}(V^T V V^T V) = \text{tr}(I_d) = d. \quad (\text{E.3})$$

Similarly, we have

$$\|\hat{V}\hat{V}^T\|_F^2 = d. \quad (\text{E.4})$$

Next, we compute

$$\langle VV^T, \hat{V}\hat{V}^T \rangle = \text{tr}(VV^T \hat{V}\hat{V}^T) = \text{tr}(\hat{V}^T V V^T \hat{V}) = \|\hat{V}^T V\|_F^2. \quad (\text{E.5})$$

Next, we use the fact that the squared Frobenius norm of a matrix is the sum of squares of its singular values to write

$$\|\hat{V}^T V\|_F^2 = \sum_{i=1}^d \sigma_i^2 = \sum_{i=1}^d \cos^2 \theta_i(E, \hat{E}). \quad (\text{E.6})$$

We may then combine these identities to write

$$\|\hat{V}\hat{V}^T - VV^T\|_F^2 = 2 \sum_{i=1}^d (1 - \cos^2 \theta_i(E, \hat{E})) = 2 \sum_{i=1}^d \sin^2 \theta_i(E, \hat{E}). \quad (\text{E.7})$$

as was to be shown. \square

Using the previous lemma, it is easy to see that the distance between subspaces is preserved under taking orthogonal complements.

Lemma E.3. *Let F and F' be subspaces of \mathbb{R}^n of dimensions m , and let F^\perp and F'^\perp denote their orthogonal complements. We have $d(F, F') = d(F^\perp, F'^\perp)$.*

We can now use these observations to state Theorem 2 from [23] in a convenient form.

Theorem E.4 (Davis-Kahan theorem). *Let Σ and $\hat{\Sigma}$ be two n by n symmetric real matrices, with eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$ and $\hat{\lambda}_1 \geq \dots \geq \hat{\lambda}_n$. Fix $1 \leq r \leq s \leq n$, and assume that $\min\{\lambda_r - \lambda_{r+1}, \lambda_s - \lambda_{s+1}\} > 0$, where we define $\lambda_0 = \infty$ and $\lambda_{n+1} = -\infty$. Let $d = r + n - s$, and let $V = (v_1, v_2, \dots, v_r, v_{s+1}, \dots, v_n)$ and $\hat{V} = (\hat{v}_1, \hat{v}_2, \dots, \hat{v}_r, \hat{v}_{s+1}, \dots, \hat{v}_n)$ be n by d matrices whose columns are orthonormal eigenvectors to $\lambda_1, \lambda_2, \dots, \lambda_r, \lambda_{s+1}, \dots, \lambda_n$ and $\hat{\lambda}_1, \hat{\lambda}_2, \dots, \hat{\lambda}_r, \hat{\lambda}_{s+1}, \dots, \hat{\lambda}_n$ respectively. Then*

$$\|\hat{V}\hat{V}^T - VV^T\|_F \leq \frac{2\sqrt{2d}\|\hat{\Sigma} - \Sigma\|}{\min\{\lambda_r - \lambda_{r+1}, \lambda_s - \lambda_{s+1}\}}. \quad (\text{E.8})$$